

Domain Time II Documentation

These pages contain the documentation for the current version of Domain Time II (v 5.2). The information is current as of 1 Aug 2021. Note: the online version of this documentation is always current. Please refer to our website for up-to-date information.

NOTE: For easier navigation, turn on Bookmarks ([View -> Show/Hide -> Navigation Panes -> Bookmarks](#))

Installing Domain Time II

- System Requirements

- Planning

 - Recommended Configurations

 - Public Time Servers

 - Regulatory Compliance

 - FINRA

 - CAT NMS PLAN

 - 21 CFR Part 11

 - EU MiFID II

- Upgrades

 - 4.x to 5.x Considerations

- Setup

 - Command-line Options

- Network Rollout

 - Active Directory Policies

 - DHCP Server Options

 - Auto-Manage Windows Firewall

Configuring the Domain Time II Components

- Domain Time II Server

- Domain Time II Client for Windows

- Domain Time II Client for Linux (DTLinux)

- Domain Time II Manager

 - Domain Time II Monitor Service

 - Domain Time II Update Server

 - Domain Time II Other Management Tools

- Domain Time II Audit Server

- Windows Time Agent

This page explains the recommended versions and basic system requirements for using Domain Time II.

Recommended versions

This table shows the recommended version for use on the indicated operating systems.

Note: Older versions may be available for compatibility on the indicated platforms, however support and documentation are provided only for the current release version (v5.2). Many of the older versions will run on other operating systems than the ones shown here, however, the indicated version provides the best performance for the operating system listed.

OS	Recommended Version	32-bit	64-bit	Server	Client	Manager	Audit Server	SDK
Workstation: XP, Vista, Win7, Win8.x, Win10, Win11 Server: 2003 & R2, 2008 & R2, 2012 & R2, 2016, 2019, 2022	Version 5.2							
Linux (64-bit Intel, Kernel 2.6.30 or later)	Version 5.2							
NT4 (Intel), 2000	Version 4.1							
NT4 (Alpha)	Version 3.2 *							
NT3.51	Version 3.2 *							
NT3.51	Version 2.1 *							
Win9x	Version 4.1							
Windows for Workgroups	Version 1.1 *							
Linux/Solaris (SPARC/Intel)	Version 2.5							

= Supported version.

Older versions are provided for backwards-compatibility only. Fixes, updates, and technical support are available only for the supported version.

* If you need this older version, please write to [Microsemi Tech Support](#).

Version 5.2 Basic Requirements

These are the system requirements for the current version of Domain Time II (except where noted). System requirements for older versions are [here](#).

- Windows XP, 2003 (&R2), Vista, 2008 (&R2), Win7, Win8.x, 2012 (&R2), Win10, 2016, 2019
[32-bit \(x86\) and 64-bit \(x64\) versions available](#). ARM is not supported. [TCP/IP Port Usage](#)
- Linux x86_64 distro running in 64-bit mode. Kernel version 2.6.30 (minimum), version 3.0 or later (preferred), with systemd init service controller (not the older inet.d "SysV" type)
[64-bit \(x64\) Intel-architecture](#). ARM is not supported. [TCP/IP Port Usage](#)
- Hardware: Typically, any Intel-architecture machine capable of running Windows (32 or 64-bit) or Linux (64-bit) may run Domain Time. Machines running Server require sufficient memory for packet buffering and maintaining the number of TCP/UDP connections used (varies by the number of client machines contacted). See [Choose the right machine\(s\)](#) for more info. Machines running Manager/Audit Server also require sufficient disk space for holding collected audit records (see the [Application Notes](#) below).

Specific Application Notes:

- [Domain Time II Client for Windows](#) (includes DT-Set command-line client)

- Domain Time II Client may be used on virtual machines. Please see [How to configure Domain Time on a virtual machine](#).
- DTTray System Tray applet not accessible when running on Windows Server Core.
- [Nano Server](#) may require MS SNMP trap support (wsnmp32.dll) be manually copied to the /system32 folder before Domain Time installation.
- Installed disk usage (not including log files): ~ 6.0 Meg

► **Domain Time II Client for Linux (DTlinux)**

- DTLinux may be used on virtual machines. Please see [How to configure Domain Time on a virtual machine](#).
- Installed disk usage (not including log files): ~ 1.0 Meg

► **Domain Time II Server**

- Domain Time II Server may be used *with caution* on a virtual machine. Please see [How to configure Domain Time on a virtual machine](#).
- DTTray System Tray applet not accessible when running on Windows Server Core.
- [Nano Server](#) may require MS SNMP trap support (wsnmp32.dll) be manually copied to the /system32 folder before Domain Time installation.
- Installed disk usage (not including log files): ~ 6.4 Meg

► **Domain Time II Management Tools**

- The Management Tools consist of Domain Time II Manager plus a number of special-use utilities and services.
- Any tools that calculate comparative time variances (such as Domain Time II Manager, Domain Time II Monitor Service, DTCheck utility, etc.) provide less accurate results when executed from a virtual OS. They should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Some functions require file access through administrative file shares and Remote Registry access using Windows Networking.
- Manager makes numerous UDP and TCP connections to remote machines. You must have sufficient system resources (memory and processing bandwidth) to handle the number of network connections, which increase proportionally to the number of machines being managed.
- Manager and many of the Management Tools will not run on Nano Server.
- Installed disk usage (not including database or log files): ~ 14.5 Meg

► **Domain Time II Manager**

- Manager requires that Domain Time II Server be installed first. Server and Manager must be installed on the same machine and be the same version.

► **Domain Time II Monitor Service**

- Monitor can be installed and run from the Manager machine or installed independently on any other physical machine.
- Monitor will provide less accurate results when executed from a virtual OS. It should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Installed disk usage (included with Manager)

► **Domain Time II Update Server**

- Update Server must be installed on a machine running Domain Time II Manager v5.1 or later.
- Installed disk usage (included with Manager)

► **Domain Time II Audit Server**

- Audit Server requires both Domain Time II Server **and** Manager. All three components must installed on the same physical machine and be the same version.
- Audit Server acts as a plug-in to Domain Time Manager and shares its view of the network.
- Audit Server is not supported on Windows 2003/2003R2/XP or earlier versions.
- Audit Server will not run on Nano Server.
- Domain Time II Audit Server will provide less accurate results when executed from a virtual OS. It should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Audit Server makes numerous UDP and TCP connections to remote machines. You must have sufficient system resources (memory and processing bandwidth) to handle the number of network connections, which increase proportionally to the number of machines being audited. If you use the new PTP Monitor, these requirements can significantly increase due to the "chatty" nature of the additional PTP protocol and management message traffic.

- The Audit Server machine must have sufficient disk space to hold all audit records and sync logs collected. You may use the [Audit Disk Space Estimator](#) to calculate disk usage for Audit Records. Sync log sizes can be estimated based on ~20 bytes of log space per synchronization.
- Some functions require file access through administrative file shares using Windows Networking.
- Collected Sync Log files, Daily Audit Summaries, and Audit records cannot be viewed from Audit Server running on Windows Server Core.
- Installed disk usage (not including log files): ~ 3.6 Meg

► **Domain Time II Windows Time Agent**

The Windows Time Agent (WTA) is a special-use Control Panel applet for configuring and monitoring the Windows Time Service (W32Time).

- WTA will act as a snap-in to Domain Time II Server or Client, or will run as a stand-alone utility.
- Will run on Windows 2000 but with limited functionality.
- Windows Time Agent will not run on Nano Server.
- As of v5.1, WTA is no longer included by default in Server and Client installations, but it remains available in the distribution setup files

► **Domain Time LMCheck Utility**

LMCheck uses Windows Networking LanMan protocols to give a rough variance report of the local subnet.

- NetBIOS (Windows Networking Browse List) must be enabled on each machine being sampled.
- Included with the Domain Time II Management Tools. It is also available as a stand-alone freeware utility.
- LMCheck will not run on Nano Server.

Domain Time II Software distributed by [Microsemi, Inc.](#)

Documentation copyright © 1995-2021 Greyware Automation Products, Inc.

All Rights Reserved

All Trademarks mentioned are the properties of their respective owners.

- Domain Time II Client may be used on virtual machines. Please see [How to configure Domain Time on a virtual machine](#).
- DTTray System Tray applet not accessible when running on Windows Server Core.
- [Nano Server](#) may require MS SNMP trap support (wsnmp32.dll) be manually copied to the /system32 folder before Domain Time installation.
- Installed disk usage (not including log files): ~ 6.0 Meg

► **Domain Time II Client for Linux (DTlinux)**

- DTLinux may be used on virtual machines. Please see [How to configure Domain Time on a virtual machine](#).
- Installed disk usage (not including log files): ~ 1.0 Meg

► **Domain Time II Server**

- Domain Time II Server may be used *with caution* on a virtual machine. Please see [How to configure Domain Time on a virtual machine](#).
- DTTray System Tray applet not accessible when running on Windows Server Core.
- [Nano Server](#) may require MS SNMP trap support (wsnmp32.dll) be manually copied to the /system32 folder before Domain Time installation.
- Installed disk usage (not including log files): ~ 6.4 Meg

► **Domain Time II Management Tools**

- The Management Tools consist of Domain Time II Manager plus a number of special-use utilities and services.
- Any tools that calculate comparative time variances (such as Domain Time II Manager, Domain Time II Monitor Service, DTCheck utility, etc.) provide less accurate results when executed from a virtual OS. They should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Some functions require file access through administrative file shares and Remote Registry access using Windows Networking.
- Manager makes numerous UDP and TCP connections to remote machines. You must have sufficient system resources (memory and processing bandwidth) to handle the number of network connections, which increase proportionally to the number of machines being managed.
- Manager and many of the Management Tools will not run on Nano Server.
- Installed disk usage (not including database or log files): ~ 14.5 Meg

► **Domain Time II Manager**

- Manager requires that Domain Time II Server be installed first. Server and Manager must be installed on the same machine and be the same version.

► **Domain Time II Monitor Service**

- Monitor can be installed and run from the Manager machine or installed independently on any other physical machine.
- Monitor will provide less accurate results when executed from a virtual OS. It should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Installed disk usage (included with Manager)

► **Domain Time II Update Server**

- Update Server must be installed on a machine running Domain Time II Manager v5.1 or later.
- Installed disk usage (included with Manager)

► **Domain Time II Audit Server**

- Audit Server requires both Domain Time II Server **and** Manager. All three components must installed on the same physical machine and be the same version.
- Audit Server acts as a plug-in to Domain Time Manager and shares its view of the network.
- Audit Server is not supported on Windows 2003/2003R2/XP or earlier versions.
- Audit Server will not run on Nano Server.
- Domain Time II Audit Server will provide less accurate results when executed from a virtual OS. It should be run on a physical machine, if possible. Please see [How to configure Domain Time on a virtual machine](#).
- Audit Server makes numerous UDP and TCP connections to remote machines. You must have sufficient system resources (memory and processing bandwidth) to handle the number of network connections, which increase proportionally to the number of machines being audited. If you use the new PTP Monitor, these requirements can significantly increase due to the "chatty" nature of the additional PTP protocol and management message traffic.

- The Audit Server machine must have sufficient disk space to hold all audit records and sync logs collected. You may use the [Audit Disk Space Estimator](#) to calculate disk usage for Audit Records. Sync log sizes can be estimated based on ~20 bytes of log space per synchronization.
- Some functions require file access through administrative file shares using Windows Networking.
- Collected Sync Log files, Daily Audit Summaries, and Audit records cannot be viewed from Audit Server running on Windows Server Core.
- Installed disk usage (not including log files): ~ 3.6 Meg

► **Domain Time II Windows Time Agent**

The Windows Time Agent (WTA) is a special-use Control Panel applet for configuring and monitoring the Windows Time Service (W32Time).

- WTA will act as a snap-in to Domain Time II Server or Client, or will run as a stand-alone utility.
- Will run on Windows 2000 but with limited functionality.
- Windows Time Agent will not run on Nano Server.
- As of v5.1, WTA is no longer included by default in Server and Client installations, but it remains available in the distribution setup files

► **Domain Time LMCheck Utility**

LMCheck uses Windows Networking LanMan protocols to give a rough variance report of the local subnet.

- NetBIOS (Windows Networking Browse List) must be enabled on each machine being sampled.
- Included with the Domain Time II Management Tools. It is also available as a stand-alone freeware utility.
- LMCheck will not run on Nano Server.

This page describes how to choose time sources, select time server hardware, and how to prepare your network for using Domain Time.

■ Decide on your time source(s)

Choosing good time sources for your network is the first implementation decision you need to make.

IMPORTANT: It is essential that your time servers have sufficient performance, hardware, and OS stability to serve time reliably. The quality of the time sync on your network can only be as good as the accuracy of the time servers themselves.

Time sources should be located as close physically and network-topologically to the machines that use them as possible. A symmetrical, low-latency network connection between all machines will provide the most accurate time.

Your network will need a top-level (trusted) source of time. This can be obtained from GPS or CDMA receivers, cesium or other directly attached time servers, known good public Internet time servers, etc. On networks with no access to other time sources, you may decide to use a Domain Time Server as your trusted time source. If so, the internal system clock on the Domain Time Server will be the trusted time source.

■ If you will be using the NTP and/or DT2 protocols

If your Domain Time Server(s) are connecting to the top-level time source(s) over a network, you will want to use multiple time sources to provide redundancy, increase the accuracy of your time, and to prevent wild time from being served should any of your time sources have an error. The best accuracy and redundancy is achieved by using at least three or more reliable time sources.

Ideally, Domain Time Servers should be set to obtain at least three time samples from each time source during each time check. See the [About Time Samples](#) sidebar for detailed information.

For example, an excellent minimum configuration for your top-level time sources would be to have at least two GPS time servers located on a local LAN with at least one additional stable public server included as a sanity-check.

If you will be obtaining time from public time servers, please refer to the list of [public time servers](#) and abide by the published rules for each time source.

■ If you will be using PTP (IEEE 1588-2008/1588-2019)

PTP provides the best accuracy when connecting to a hardware-based Grandmaster clock on the same subnet.

You should have at least one other machine capable of becoming Grandmaster online for redundancy. PTP using the Default or Enterprise [profiles](#) provides for a master election among available machines should the current Grandmaster be offline. PTP using the Telecom [profile](#) uses a configured list of possible masters. Domain Time Server can be configured to be one of these backup master clocks for the Default or Enterprise profile (see [How to configure Domain Time Server as a PTP Master](#)). Domain Time Server cannot be a Telecom master. Domain Time Client cannot become a PTP master of any flavor.

All Domain Time Servers or Clients running PTP should also be configured to have at least one fallback NTP/DT2 time source (see [Configuring Domain Time II for PTP](#)).

■ Choose the right machine(s)

Review the [Software Requirements](#) for Domain Time II.

Due to how the system clock on operating systems are maintained, some systems are unsuitable for keeping accurate time. The guidelines below apply to both time servers and clients, however they are of particular concern to any machine you want to use as a time server.

A good candidate machine for accurate timekeeping will have sufficient processor power, memory, and network hardware to be able to service the operating system and applications without hitting bottlenecks under load that cause delays in servicing interrupts, packets, and threads in a smooth and timely fashion. A heavily-used machine will typically have more clock-drift problems than a lightly-used system, so be sure that your machines are not experiencing bursty periods of excessively-high load or other performance problems.

Some system motherboard designs, BIOS and firmware issues, multi-processor implementations, system/video/network drivers, or other system components can cause problems with servicing the system clock correctly and may require updates from the manufacturer. Be sure to check with your vendor(s) to be sure you are up-to-date with all necessary patches.

Most modern operating systems and motherboards have integrated power-saving features. Unfortunately, many of these have serious detrimental effects on system timekeeping. In general, you will want to disable all power-saving features on all of your time servers, and also on any clients where precise timing is required.

In general, the best processors/chipsets for time synchronization are Intel's Core i7 line (or later) or Xeon E7 line (or later). Earlier chips are not as stable or as precise as the newer models. The newer processors also have an invariant timestamp counter, which allows Domain Time II to measure the passage of time accurately regardless of SpeedStep or other power-saving mechanisms. Issuing `DTCHECK /cpuid` from the command-line will show you whether or not your processor supports an invariant TSC.

Win8/2012 or newer versions are preferred and are more predictable than Vista, Windows 7, or Windows 2008 for high-accuracy timing. The older XP/Server 2003 platform is also more stable than the problematic Vista/Win7/2008 versions.

Virtual Machines

In addition to the problem with heavily-loaded systems mentioned above, virtual environments (VMWare, Hyper-V, etc.) often have significant issues in servicing the clock in a timely manner, making them less than ideal for highly-accurate time synchronization. Domain Time will help you achieve the best synchronization possible on virtual systems, but you should be aware of the limitations. You can only determine if a virtual system will perform to your expectations by testing in your environment under your normal workloads.

- In general, Domain Time Server should be run from a physical machine, if possible. Also, any tools that calculate comparative time variances (such as Domain Time II Audit Server, Domain Time II Monitor Service, the Domain Time II Manager variance report, DTCheck utility, etc.) give less accurate results when executed from a virtual guest. These should be run on physical machines, if possible.
- Domain Time Clients may be run on an OS in a virtual machine guest, although you should be aware that regardless of the time service configuration, the clock will still have inherent inaccuracies. Any time-critical system should run directly on physical hardware.

See [this article](#) from our knowledgebase for more information on use with virtualization systems.

■ Prepare your network to pass the necessary traffic

Your network routers, switches, and firewalls must be able to pass the proper traffic to allow Domain Time to function correctly. Here are some basic guidelines:

- Domain Time II uses the DT2 (Domain Time II) protocol to communicate not only time sync data, but control messages and data streams between Servers, Clients, Management Tools, and Audit Server.

IMPORTANT: You should always configure your internal network to pass both port 9909 UDP **AND** port 9909 TCP traffic bi-directionally between all subnets, ***even if you will be using a different protocol to sync the time.***

- If you will be obtaining time from an external time source (such as from a public time server) through a firewall using the DT2 protocol, you may use either port 9909 UDP or 9909 TCP. UDP has lower overhead and latency than TCP so it tends to be slightly more accurate, however, some firewall administrators prefer to allow only TCP connections. DT2 also has a special "DT2 over HTTP" protocol available to allow synchronization with Domain Time II Servers over HTTP (default port 80), which can allow synchronization through most existing firewalls.
- You will need to configure your firewalls/switches to pass any other time protocols you want to use (i.e. port 123 UDP for NTP, ports 319 & 320 UDP for PTP, etc.). See the protocol table below.
- Domain Time uses standard IP networking calls, made via the WinSock stack on Windows and the standard TCP/IP stack in Linux. Traffic therefore conforms to IP protocol standards, including use of ephemeral source ports for originating traffic directed at remote target ports. You should be sure your firewall(s) permits traffic originating from ephemeral ports directed toward the defined listening ports in the protocol table below.
- Domain Time natively supports both IPv4 and IPv6. You may pass traffic over either version of TCP/IP.
- Domain Time components work best when able to transmit multicasts to discover machines on other subnets, so you will want to allow multicast traffic between your subnets, even if you will be using unicast protocols to synchronize the time. **Note:** A multicast-capable router must be present on each subnet and configured to pass the multicast traffic. Servers and Clients must be configured with sufficient TTL/Hop Count settings to cross the number intervening routers/switches. Some Domain Time components may also use broadcasts to local subnets and/or directed broadcasts to remote subnets for discovery purposes. See [Network Discovery](#).
- If using the PTP protocol, Domain Time will use multicasts, or a combination of multicasts and unicasts (if using the hybrid or

Enterprise [profiles](#)). Domain Time can also transmit DT2 and NTP time packets using multicast and broadcasts, if desired. See [Broadcasts and Multicasts](#) for more information.

- Your internal network should have correctly configured and functioning routing, DNS, Active Directory, WINS, and Windows Network browsing (if using NetBIOS).
- Some functions of Domain Time components (such as remote installation/upgrade/configuration) require Windows Networking file and remote registry access through administrative shares. Those programs or services must be run under a user account with sufficient administrative privileges to make such connections.
- ICMP traffic (esp. PING) should be permitted to all machines.

Note: As of Version 5.2.b.20150828, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

Domain Time components may use these network ports for various functions (default ports shown):

Protocol	Default Port/Address	Type
DT2	9909 UDP and 9909 TCP (Required for all Domain Time Components)	Time sync, auditing, and control messages
DT2 Multicast	IPv4: 239.192.99.9 IPv6: FF05::9909	Network discovery (optional broadcast time sync)
DT2 over HTTP	80 TCP	Time sync, stats webpage on Server, Version update checking
NTP/SNTP SNMP: RFC 1769 v3: RFC 1305 v4: RFC 2030	123 UDP	Time sync
NTP Multicast	IPv4: 224.0.1.1 IPv6: FF05::101	Network discovery (optional broadcast time sync)
PTP v2.0 (IEEE 1588-2008) PTP v2.1 (IEEE 1588-2019)	319 and 320 UDP	Time sync
PTP v2.0 (IEEE 1588-2008) Multicast PTP v2.1 (IEEE 1588-2019) Multicast	IPv4: 224.0.1.129 IPv4: 224.0.0.107 IPv6: FF05::181 IPv6: FF02::6B	Time sync
TIME/ITP (RFC 868)	37 UDP and/or 37 TCP	Time sync (Server only)
Daytime (RFC 867)	13 TCP	Time sync (Server only)
DT Alert Control	9910 TCP	Domain Time Real-time Alert Sharing/Alert Viewer Audit Server Standby-mode Replication
DT Status	9911 UDP and/or 9911 TCP	Domain Time Service Status Monitor

Recommended Configurations

Use these configuration examples to create an efficient and robust time distribution hierarchy for your network.

Choose the Time Distribution Model that Fits your Network

Find the example below that most closely matches your network. Then, follow the simple installation plan instructions indicated for your network model to quickly and successfully install Domain Time II.

If you are in an industry that has regulations regarding time synchronization, you'll also want to see the [Regulatory Compliance](#) pages.

Using NTP and/or DT2 protocols:

- ▶ [Single Machine Model](#)
- ▶ [Workgroup Model](#)
- ▶ [Single Domain Model](#)
- ▶ [Multi-Domain Model](#)
- ▶ [Multiple Networks without Masters/Slaves](#)

Using PTP (IEEE 1588-2008/2019):

- ▶ [Hardware Grandmaster](#)
- ▶ [Software Grandmaster](#)

Stand-Alone Single Machine Model

A machine that is not part of a domain.

Domain Time II running on any stand-alone machine should be manually configured to get its time from trusted sources.

Installation Plan:

(click the link to get detailed instructions for each component listed)

- Install Domain Time II [Server](#) or [Client](#).
- Configure it to get time from your chosen time source(s).




Workgroup Model

For networks without a Windows domain controller.

In a small workgroup without a Windows domain, one machine should run Domain Time Server. It will be configured to get its time from trusted sources and distribute it to clients on the network.

All other Windows machines on the network should run Domain Time II Client. Client can be set to either use the specific IP address or DNS name of the Server or automatically discover the time server. You may choose from the following options:



- Manually configure the Client to specify which Servers to use.
- Set the Client to [Discover sources automatically](#) using Broadcast/Multicast. See the [Discovery](#) page for more information.
- Set the Client to [Discover sources automatically](#) using DHCP. Client will use time servers listed in [DHCP Time Server options](#). Machines do not need to use DHCP for ip-address assignment to be able to get time server addresses from DHCP servers. See the [Discovery](#) page for more information.

Note: These client settings can be pre-configured and rolled-out to multiple machines using Domain Time Manager.

Any other time-capable machines and devices should be configured to get their time from the Server using whatever time protocol they use (such as NTP, TIME/ITP, etc.)

Installation Plan:

(click the link to get detailed instructions for each component listed)


- Install Domain Time II [Server](#).
- Configure the Server to get its time from your chosen time source(s).
- Install [Client](#) on all other Windows machines. Configure the Clients for either automatic discovery or manually select their time sources.
- Configure any third-party clients or devices on the network to get their time from the Domain Time II Server.

Single Domain Model

Networks with a single Windows domain (or single Active Directory Tree).

Domain Controllers should run Domain Time Server.

The machine holding the PDC-Emulator role (FSMO) automatically becomes the Master server and should be configured to obtain the time from trusted time sources. Domain Time Server installed on all other DCs automatically becomes a Slave to the Master. You may also set any other Domain Time Server on the domain to be a Slave. Slaves provide for important redundancy and efficient distribution of time. See the [Domain Role](#) page for more info.



Install Domain Time Client on all other Windows servers and workstations on the network. Client can be set to either use specific time sources or automatically discover time servers. Clients automatically acquire important redundancy and failover advantages when Masters and Slaves are present on the domain, regardless of which configuration options are selected.

You may choose from the following configuration options on Clients:

- Manually configure the Client to specify which Servers to use.
- Set the Client to [Discover sources automatically](#) using Broadcast/Multicast. See the [Discovery](#) page for more information.
- Set the Client to [Discover sources automatically](#) using DHCP. Client will use time servers listed in [DHCP Time Server options](#). Machines do not need to use DHCP for ip-address assignment to be able to get time server addresses from DHCP servers. See the [Discovery](#) page for more information.
- You may also use [Active Directory policies](#) to specify which Servers the Clients should use. Active Directory policies override any other settings you make on the Client.

Configure any other time-capable machines and devices to get their time from the nearest Domain Time Slave Server.

Note: Server and Client settings can be pre-configured and rolled-out to multiple machines using Domain Time Manager. See [Network Rollout](#) for details.

Installation Plan:

(click the link to get detailed instructions for each component listed)

- If you will be using [Active Directory policies](#) to specify Domain Time settings, Use your Group Policy Management Editor to install the domtime.adm policy file from the distribution files as a template into the Computer Configuration\Policies\Administrative Templates section of your desired Group Policy object(s). Then, configure the settings for each Domain Time policy item you want to apply to that object.
- If you will be [using DHCP](#) to specify time servers for your Clients to use, configure Option 004 of your DHCP servers to provide the IP address(es) of the desired Domain Time II Server(s) or Option 024 to specify NTP servers.
- Use [Setup](#) to install both Domain Time II [Server](#) and the [Management Tools](#) on any machine you want to use as your management workstation. If you will be using [Audit Server](#), install it on this machine also. (Each instance of Server, Manager and Audit Server requires a separate license)
- Use [Manager](#) to perform each of the following steps from your management workstation:
 - Install [Server](#) on the PDC/FSMO (It will assume the Master role). Configure the Master to get its time from your chosen trusted time source(s). Server averaging ("[Analyze all listed servers and choose the best...](#)") should be *enabled*.
 - Install Server on all other DCs (they will automatically assume the Slave role).
 - If you want to pre-configure your Client installation settings for network rollout:
 - Install [Client](#) on a test machine to prepare an installation template .reg file for Manager to use.
 - Connect to the Client's Control Panel applet to set up the Client exactly the way you want it to be configured.
 - Use the Client's [Import/Export](#) utility to export the Client settings to a .reg file. Copy the the .reg file to the Manager's [Program Files\Domain Time II](#) folder to be available as a template for installation.
 - Install Client on all other Windows machines. Select the template .reg file if you have created one to preset the settings, or connect to the Clients after installation to set them for either automatic discovery or manually select their time sources.
 - Configure any third-party clients or devices on the network to get their time from the nearest Domain Time II Server.
 - Use Manager to install the [Monitor Service](#) and [Update Server](#) to automatically monitor your network and keep it updated.

Multi-Domain Model

Networks with multiple Windows domains or Active Directory Forests with multiple trees.

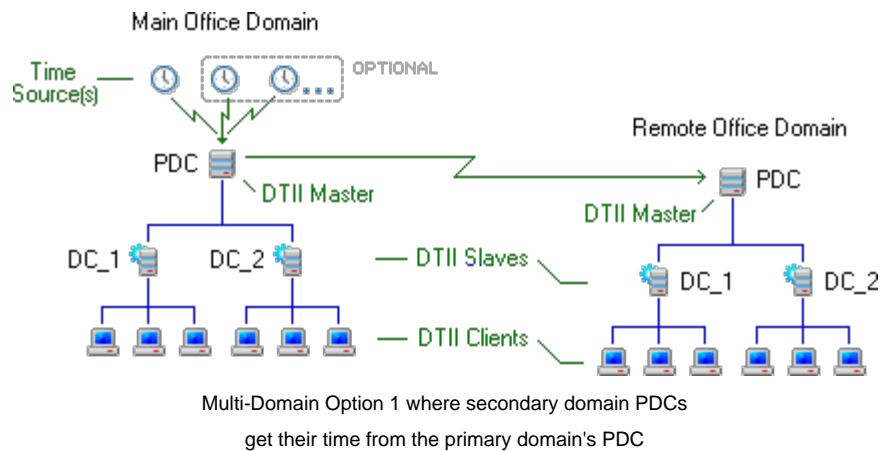
The Single Domain Model described above should be implemented on each individual domain (tree), except that the Master time server (the PDC-emulator) for each domain needs to be configured to get its time:

- from the Master Server on the primary domain, or
- from the same trusted time source(s) as the main domain's Master, or
- from its own local trusted time source(s), or
- Using a combination of the above methods (mesh configuration).

You can configure each domain's Master to get its own time in the various ways described below:

Option 1

In the first configuration option, the PDC for the master domain gets its time from its trusted source(s), while the PDCs for each of the resource domains are manually configured to use the master domain's PDC as their external time source.



The main advantages to this configuration are:

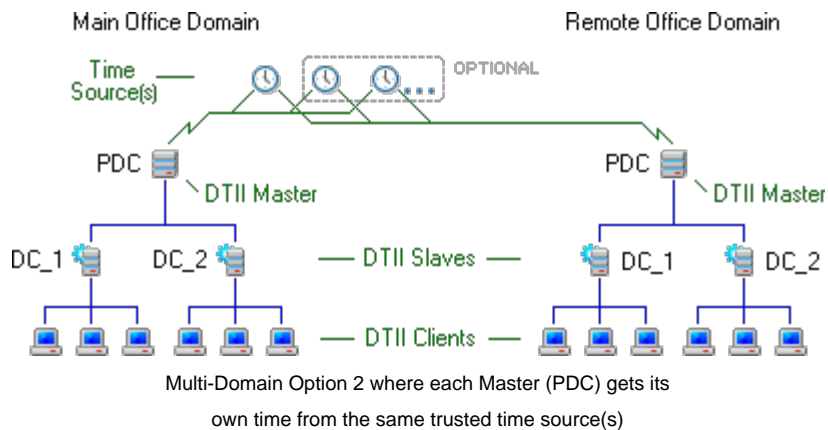
- The Master (PDC) running on each of the resource domains can automatically look up and use the Master of the main domain's PDC.
- The time hierarchy mirrors the Windows domain structure.
- Time in each domain will closely match the time in all other domains.

The main disadvantage to this configuration is:

- Using only the Master PDC of the main domain as a time source is a single point of failure for the resource domains.

Option 2

In the second option, the Master (PDC) of each domain gets its time from the same trusted time source(s).



The main advantage to this configuration is:

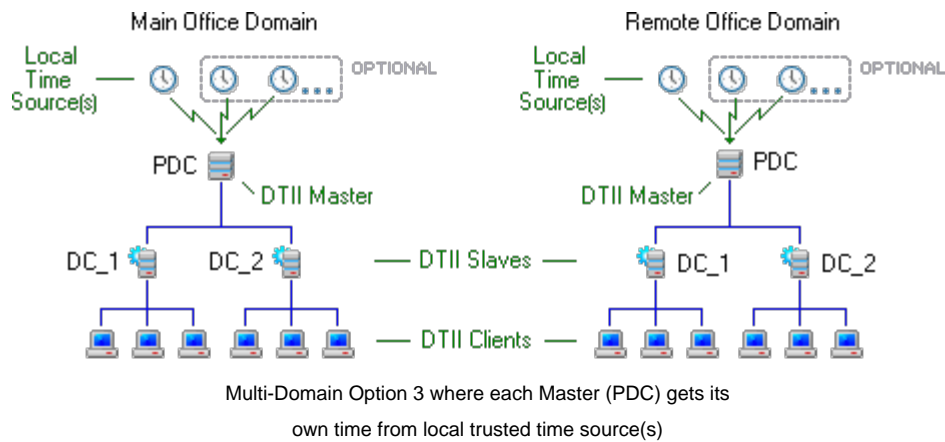
- Each domain has its own connection to the time source(s), If multiple sources are specified, there is no single point of failure.

The disadvantages to this configuration are:

- You must manually configure each Domain Time Server with the address of the time source(s).
- Each time check by each Server causes traffic to all time sources, which may be across WAN links.
- Time in each domain may differ slightly from each other (depending on which sources are local to the domain).

Option 3

The third option is similar to option 2, except the Master (PDC) of each domain gets its time from its own local trusted time sources.



The advantages to this configuration are:

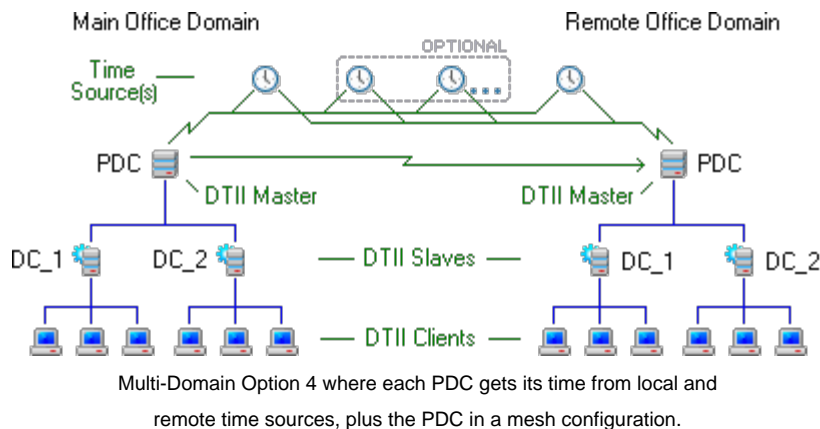
- Each domain has its own connection to the time source(s), If multiple sources are specified, there is no single point of failure.
- Accuracy is improved since local time sources have lower latency than remote ones.

The disadvantages to this configuration are:

- You must manually configure each Domain Time Server with the address of the time source.
- Time in each domain may vary somewhat from other domains since the time is not compared between sites.

Option 4

The mesh configuration shown below represents an excellent configuration for using Domain Time across multiple domains or for an Active Directory forest. Each PDC gets its time from both local and remote time sources, and also from other PDCs.



There a number of advantages to a mesh configuration:

- Accuracy is improved across your entire enterprise since Variances among the various Servers and time sources are compensated for automatically.
- Stable local time sources are automatically preferred when server averaging ("[Analyze all listed servers and choose the best...](#)") is *enabled*.
- The network is more robust. Domain Time Masters adjust automatically to changes in the availability of any time sources. If any source becomes unavailable, alternate sources are automatically used.
- The Master (PDC) running on each of domains can automatically look up and use the Master Server of any other domain.

The disadvantages to this configuration are:

- You must manually configure each Domain Time Server with the address of any non-Domain Time time source.

- Requires that each server be able to communicate with each other and each time source.
- Each time check by each Server causes traffic to all other Servers and time sources, which may be across WAN links.

Installation Plan:

(click the link to get detailed instructions for each component listed)


- Use [Setup](#) to install both Domain Time II [Server](#) and the [Management Tools](#) on any machine you want to use as your management workstation. Manager works best if you have trusts to all domains you want to manage. If not, you should also install the Server and Management Tools on a machine in each of the untrusted domains and perform installations to those domains from there. If you will be using [Audit Server](#), install it on this machine also. (Each instance of Server, Manager and Audit Server requires a separate license)
- You may use a single instance of [Audit Server](#) across multiple networks. Alternately, you may want to install additional Audit Servers on individual networks to spread the Audit workload, if you want to use different types of machines on multiple schedules, or to keep separate audit data for individual domains/companies. Audit Server also has a special [Standby Mode](#) for use in Disaster Recovery scenarios. (Each instance of Audit Server requires a separate license for Server, Manager, and Audit Server)
- Perform all of the tasks in the [Single Domain Model](#) Installation Plan above on each domain (tree), starting with the top-level domain.
- Use Manager to configure the Master time servers (PDCs) of each domain to obtain time from available local trusted time source(s) and from each other.

Multiple Networks without Masters/Slaves

Installing Domain Time in multiple locations without using Masters or Slaves.

When possible, you should install Domain Time using one of the Master/Slave configurations above. Masters and Slaves automatically provide important accuracy and redundancy benefits. However, it is possible to construct a robust time hierarchy across multiple physical locations without using Master and Slaves, if necessary.

Each physical location should be installed according to the instructions in the [Workgroup Model](#) above. Then, each Server should be set to get time using server averaging ("[Analyze all listed servers and choose the best...](#)" is *enabled*) from all available time sources. This creates a mesh configuration that harmonizes time among each of your locations, plus provides redundancy in case any time source becomes unavailable to the Servers. Alternately, you can turn off server averaging and have your Domain Time Servers get their time using a fallback list of sources, where each Server would get their time from their primary local time source, but would fall back to remote source(s) if the primary fails. Use this option if you have high or variable latency network connections between your locations.



To provide redundancy to your Clients, you will need to make manual changes to the time sources list. First, configure the Clients in each location to use their listed time sources as a fallback-list ("[Analyze all listed servers and choose the best...](#)" is *disabled*) so that each Client first contacts its local Domain Time Server for the time. Then, list any other Domain Time Servers (local or remote) and/or other available time sources (local sources first) so that Clients can fallback to those if no Domain Time Servers are reachable.

For example, Clients might have these servers listed in their [Obtain the time](#) time sources lists (Fallback-mode):

Main Office

Source 1:	Main Office Domain Time Server
Source 2:	

Remote Office

Source 1:	Remote Office Domain Time Server
Source 2:	

Main Office GPS Server	Remote Office GPS Server
Source 3: Remote Office Domain Time Server	Source 3: Main Office Domain Time Server
Source 4: Remote Office GPS Server	Source 4: Main Office GPS Server

Installation Plan:


(click the link to get detailed instructions for each component listed)

- Use [Setup](#) to install both Domain Time II [Server](#) and the [Management Tools](#) on any machine you want to use as your management workstation. Manager works best if you have trusts to all domains you want to manage. If not, you should also install the Management Tools on a machine in each of the untrusted domains and perform installations to those domains from there. If you will be using [Audit Server](#), install it on this machine also. (Each instance of Server, Manager and Audit Server requires a separate license)
- You may use a single instance of [Audit Server](#) across multiple networks. Alternately, you may want to install additional Audit Servers on individual networks to spread the Audit workload, if you want to use different types of machines on multiple schedules, or to keep separate audit data for individual domains/companies. Audit Server also has a special [Standby Mode](#) for use in Disaster Recovery scenarios. (Each instance of Audit Server requires a separate license for Server, Manager, and Audit Server)
- Perform all of the tasks in the [Workgroup Model](#) Installation Plan above in each physical location.
- Use Manager to configure the Servers in each location to obtain time from available trusted time source(s) and from each other.
- Turn off Server Averaging on the Clients, and configure the Time Sources list so that the local Domain Time Server is listed first, then add Domain Time Servers in other locations, and finally, all other time sources (local sources first).

PTP Using a Hardware Grandmaster Model

For synchronizing machines using PTP from a hardware Grandmaster.

There must be a hardware Grandmaster clock available, preferably on the same subnet as the other PTP devices. The device should provide IEEE 802.3 implementations of either the Default, Enterprise, or Telecom [PTP profile](#). Although the PTP protocols may be routed to other subnets, the additional latency and possible queuing or discard of UDP packets by intervening routers may make this problematic. Boundary or Transparent clocks are preferred for distributing PTP to subnets.



For redundancy, we recommend there be at least one additional machine (preferably another hardware clock) capable of becoming Grandmaster should the primary go offline. Domain Time Server can also be configured to be a backup clock capable of assuming the Grandmaster role if you are using the Default or Enterprise [PTP profile](#).

All other Windows machines on the network should run Domain Time II Client.

IMPORTANT: All Domain Time II machines should have at least one NTP or DT2 Server set as a fallback time source.

Note: These settings can be pre-configured and rolled-out to multiple machines using Domain Time Manager.

Installation Plan:

(click the link to get detailed instructions for each component listed)

- If using Domain Time Server as a backup PTP master (Default or Enterprise profile), install [Server](#) on at least one machine.
 - First, configure Server to become a PTP slave (see [Configuring Domain Time II for PTP](#)).
 - Then, configure the Server to become a PTP master (see [How to configure Domain Time Server as a PTP Master](#)). If using the Default or Enterprise profile, set this machine's Master priority so that it has lower precedence (a higher numeric value) than the Grandmaster's priority setting.
- Install [Client](#) on all other Windows machines. Configure the Clients to become PTP slaves (see [Configuring Domain Time II for PTP](#)).

If you will be using Domain Time Manager/Audit Server:

- Use [Setup](#) to install both Domain Time II [Server](#) and the [Management Tools](#) on any machine you want to use as your management workstation. Configure Server to become a PTP slave (see [Configuring Domain Time II for PTP](#)). You may also configure it to be a backup PTP master, as described above. If you will be using [Audit Server](#), install it on this machine also. Manager works best if you have trusts to all domains you want to manage. If not, you should also install the Management Tools on a machine in each of the untrusted domains and perform installations to those domains from there. If you will be using [Audit Server](#), install it on this machine also. (Each instance of Server, Manager and Audit Server requires a separate license)
- You may use a single instance of [Audit Server](#) across multiple networks. Alternately, you may want to install additional Audit Servers on individual networks to spread the Audit workload, if you want to use different types of machines on multiple schedules, or to keep separate audit data for individual domains/companies. Audit Server also has a special [Standby Mode](#) for use in Disaster Recovery scenarios. (Each instance of Audit Server requires a separate license for Server, Manager, and Audit Server)
- Use [Manager](#) to perform each of the following steps from your management workstation:
 - If you want to pre-configure your Client installation settings for network rollout:
 - Install [Client](#) on a test machine to prepare an installation template .reg file for Manager to use.
 - Connect to the Client's Control Panel applet to set up the Client exactly the way you want it to be configured.
 - Use the Client's [Import/Export](#) utility to export the Client settings to a .reg file. Copy the the .reg file to the Manager's **Program Files\Domain Time II** folder to be available as a template for installation.
 - Install Client on all other Windows machines. Select the template .reg file if you have created one to preset the settings, or connect to the Clients after installation to set them for either automatic discovery or manually select their time sources.

PTP Using a Software Grandmaster Model


For synchronizing machines using PTP from a Software Grandmaster.

There must be a machine configured to be a software Grandmaster clock available, preferably on the same subnet as the other PTP devices. The device should provide IEEE 802.3 implementations of either the Default, Enterprise, or Telecom [PTP profile](#). Although the PTP protocols may be routed to other subnets, the additional latency and possible queuing or discard of UDP packets by intervening routers may make this problematic. Boundary or Transparent clocks are preferred for distributing Default or Enterprise profile PTP to subnets.

For redundancy, we recommend there be at least one additional machine capable of becoming Grandmaster should the primary go offline. Domain Time Server can be configured to be a backup clock capable of assuming the Grandmaster role using the Default or Enterprise [PTP profile](#).

All other Windows machines on the network should run Domain Time II Client.

IMPORTANT: All Domain Time II machines should have at least one NTP or DT2 Server set as a fallback time source.



Note: These settings can be pre-configured and rolled-out to multiple machines using Domain Time Manager.

Installation Plan:

(click the link to get detailed instructions for each component listed)

- To use Domain Time Server as the Grandmaster clock, install [Server](#). Configure this machine to get its own time from at least one NTP or DT2 time source. If using the Default or Enterprise profile, set this machine's Master priority so that it has higher precedence (a lower numeric value) than any other PTP device on the network.
- If using Domain Time Server as a backup PTP master (Default or Enterprise profile), install [Server](#) on at least one machine.
 - First, configure backup Server to become a PTP slave (see [Configuring Domain Time II for PTP](#)).
 - Then, configure the Server to become a PTP master (see [How to configure Domain Time Server as a PTP Master](#)). Set this machine's Master priority so that it has lower precedence (a higher numeric value) than the Grandmaster's priority setting (if using the Default or Enterprise profiles).
- Install [Client](#) on all other Windows machines. Configure the Clients to become PTP slaves (see [Configuring Domain Time II for PTP](#)).

If you will be using Domain Time Manager/Audit Server:

- Use [Setup](#) to install both Domain Time II [Server](#) and the [Management Tools](#) on any machine you want to use as your management workstation. Configure Server to become a PTP slave (see [Configuring Domain Time II for PTP](#)). You may also configure it to be a backup PTP master, as described above. If you will be using [Audit Server](#), install it on this machine also. Manager works best if you have trusts to all domains you want to manage. If not, you should also install the Management Tools on a machine in each of the untrusted domains and perform installations to those domains from there. If you will be using [Audit Server](#), install it on this machine also. (Each instance of Server, Manager and Audit Server requires a separate license)
- You may use a single instance of [Audit Server](#) across multiple networks. Alternately, you may want to install additional Audit Servers on individual networks to spread the Audit workload, if you want to use different types of machines on multiple schedules, or to keep separate audit data for individual domains/companies. Audit Server also has a special [Standby Mode](#) for use in Disaster Recovery scenarios. (Each instance of Audit Server requires a separate license for Server, Manager, and Audit Server)
- Use [Manager](#) to perform each of the following steps from your management workstation:
 - If you want to pre-configure your Client installation settings for network rollout:
 - Install [Client](#) on a test machine to prepare an installation template .reg file for Manager to use.
 - Connect to the Client's Control Panel applet to set up the Client exactly the way you want it to be configured.
 - Use the Client's [Import/Export](#) utility to export the Client settings to a .reg file. Copy the the .reg file to the Manager's **Program Files\Domai n T i m e I I** folder to be available as a template for installation.
 - Install Client on all other Windows machines. Select the template .reg file if you have created one to preset the settings, or connect to the Clients after installation to set them for either automatic discovery or manually select their time sources.



If you don't have your own GNSS (GPS) or CDMA receiver, cesium clock, or other trusted time source on your local network, you will need to obtain the time over the Internet from a public time server. This page will help you locate an appropriate server.

IMPORTANT: Internet time servers are not always reliable; they can be overloaded, misconfigured, subject to network delays, or simply go away without notice. Be sure to enable the [Analyze time samples...](#) option and include at least three time sources in your configuration (setting each source to request 3 samples) to help compensate for these problems.

You should test each time server you choose for reliability and accuracy before committing to use it in production. The Domain Time II [Time Server Test](#) (included with Domain Time II Manager) and [NTPCheck](#) (included with Domain Time II Server and Client) utilities are ideal for this purpose.

Do not assume you can use the default time servers listed by Domain Time II Server. These servers are listed as a convenience, to help you get started. You *must* honor the requirements posted by each time service provider. Usually nothing more than asking permission is required. Some servers have regional restrictions; others only allow specific NTP strata; others may have other requirements. It is your responsibility to ensure compliance. Time service providers operate on the honor system, and the system will only continue functioning smoothly if everyone abides by the rules.

Domain Time Servers

Licensed Domain Time users may use them as primary and secondary time sources. You may, of course, also use these servers for testing purposes, but **please** do not use these servers regularly unless you are a Domain Time customer!

Note: these servers are provided for convenience only; we do not guarantee the availability and accuracy of these systems. In particular, these servers are not guaranteed to be traceable to NIST, so they are not suitable for meeting compliance standards that require such traceability.

■ tick.greyscale.com

Domain Time II Server

Both IPv4 and IPv6 available

NTP Stratum 2-4

Protocols: DT2 (UDP, TCP, DT2 over HTTP), NTP (UDP), TIME-ITP (UDP), Daytime (UDP).

Note: TIME-ITP (TCP) and Daytime (TCP) no longer offered as of 23 May 2017.

Access: Any registered customer may use this server as a time source for Domain Time Server. Please do not point individual clients at this server. Please do not check more than once per minute.

■ tock.greyscale.com

Domain Time II Server

Both IPv4 and IPv6 available

NTP Stratum 2-4

Protocols: DT2 (UDP, TCP, DT2 over HTTP), NTP (UDP), TIME-ITP (UDP), Daytime (UDP).

Note: TIME-ITP (TCP) and Daytime (TCP) no longer offered as of 23 May 2017.

Access: Any registered customer may use this server as a time source for Domain Time Server. Please do not point individual clients at this server. Please do not check more than once per minute.

Public NTP Servers

There are hundreds of NTP/SNTP servers available on the Internet. Be a good network citizen; if you have more than 100 clients on your network, your main server *may* qualify for connecting to a Stratum 1 server, otherwise you should pick a

Stratum 2 server.

- [NTP.Servers Public Time Servers List](#)
- [NIST Servers](#)
- [USNO Servers](#)

Using Domain Time for Regulatory Compliance

Many organizations and government regulatory agencies already have or are implementing regulations regarding the synchronization of time on computer systems. Further regulations exist requiring the establishment of an audit trail of time synchronization that can be used to verify the validity of electronic timestamps.

Domain Time II is uniquely suited to ensure compliance with such requirements. When properly implemented, Domain Time substantially exceeds all existing and proposed regulatory standards for time synchronization. Domain Time also provides the ability to automatically create the audit trail necessary to demonstrate this compliance.

These pages describe in detail how to configure Domain Time II to satisfy these regulations:

- ▶ [FINRA](#)
- ▶ [CAT NMS PLAN \(SEC\)](#)
- ▶ [US Federal Drug Administration 21 CFR Part 11](#)
- ▶ [European Union MiFID II](#)

FINRA Compliance

The Financial Industry Regulatory Authority (FINRA) regulates registered trading brokers and broker-dealer firms in the United States. The FINRA regulations include various rules and requirements for the synchronization of clocks, which are summarized below.

FINRA Rule 4590 Synchronization of Member Business Clocks specifies that:

- 50ms Tolerance against NIST time for NMS securities and OTC Equity Securities, 1 second Tolerance for all others.
- Clocks must be synchronized every business day before market open and re-synchronized, as necessary throughout the day.
- Clock synchronization procedures must be documented.
- Logs must be kept of synchronization events and any drift out of tolerance should be flagged.
- Logs need to be retained for the period indicated and in the format specified under SEC Rule 17a-4(f).

FINRA Rule 6820 Clock Synchronization (Consolidated Audit Trail) includes the same synchronization tolerances as Rule 4590, but adds that:

- Logs need to be retained for 5 years with the 2 most recent years easily accessible.
- Members need to certify to FINRA that they meet these requirements.
- Violations need to be reported according to SEC CAT NMS Plan rules

CAT NMS Plan Section 6.8(a)(i) & (ii) and other guidance adds the following to the above requirements:

- CAT Participants tolerance to NIST is 100 microseconds. Members remain at the regular tolerances as specified in Rules 4590/6820.
- Participants must self-report violations that exceed tolerance for at least one second.
- Members must self-report violations that exceed 2x the normal tolerance (i.e. $\geq 200\text{ms}$)
- Members must self-report any system that exceeds the tolerance 10 times in one rolling 24-hour period.
- Clock granularity should at minimum be recorded/reported in milliseconds and possibly finer increments as required by NMS.

Definition of Synchronization Tolerance

The regulations allow for the use of any time source for the synchronization, as long as all clocks stay synchronized within the specified tolerance of the NIST clock. The tolerance is defined as including all of the following:

- The difference between the NIST standard and a time provider's clock
- Transmission delay from the source (Latency)
- Amount of drift of the member's clock (Drift)


The total of all of the above for any clock must not exceed the specified tolerance.

How to Use Domain Time II to comply with the FINRA Requirements

Domain Time II meets or exceeds all of the specific FINRA requirements detailed above. Properly configured, Domain Time will allow you to easily comply with all of the computer clock synchronization requirements.

Domain Time II is designed specifically to provide both accurate time synchronization and a complete history of that synchronization. Each Domain Time II time sync component (Servers and Clients) have the ability to keep detailed logs and statistics of their own activity - and, critically, to **report that information automatically** to monitoring and auditing systems when requested.

This diagram shows the basic structure of the Domain Time II system, showing how time synchronization and audit data collection are handled. Your time distribution hierarchy may vary.



Configuring for compliance

There are two basic steps necessary to use Domain Time II to achieve compliance:

- Configure Domain Time II to provide accurate time synchronization to all clocks
- Configure Domain Time II to collect and maintain sync records in an audit trail
- **Configure Domain Time II to provide accurate time sync to all clocks**

Domain Time II, when installed according to the instructions found on the [Recommended Configurations](#) page of the Domain Time II documentation, will meet most of the NASD requirements for time synchronization. However, there are a few additional configuration considerations beyond the standard recommended installation instructions for FINRA compliance. Let's consider each of the requirements and what is required to ensure Domain Time II fulfills them.

- **FINRA Requirement:** 50ms (or 1 sec) Tolerance to NIST (or 100us if a CAT Participant)
Solution: Configure Domain Time to get its time from a local hardware time appliance

It is unlikely you will be able to consistently achieve 50ms sync to NIST without having a local GPS/GNSS or CDMA-derived time server appliance on your local network to act as your reference clock. Although you can get access to [NIST Echo Servers](#) over the Internet, the high load and variable latency of these servers make them unsuitable for higher accuracy. You may be able to use them if you only need to achieve the 1-second target, however, you may find them only sporadically available.

- **FINRA Requirement:** Regular synchronization
Solution: Configure Domain Time to synchronize on a Fixed Schedule

Domain Time II Clients and Servers are background services that remain continuously synchronized with their source. Be sure to set the synchronization period on the Timings property page of the Domain Time applet to **Fixed** of at least 1/minute. This may require some trial and error to achieve the correct rate for machines that have large amounts of drift, such as virtual systems. Synchronize more often if you machines drift outside the target. See the [Timings](#) page in the documentation for more info.

- **FINRA Requirement:** Synchronize every business day before Market Open
Solution: If Domain Time is set to a fixed schedule (see above) it will automatically synchronize before Market Open

For example, if Domain Time is set to a fixed schedule of 1/minute, the clock will always be synched no more than 1 minute before Market Open.

How to Configure Domain Time II to collect and maintain sync records in an audit trail

The information below is based on meeting FINRA regulatory requirements, but gives a good overview of how Domain Time II can assist in creating and maintaining an audit trail of time synchronization.

- **FINRA Requirement:** Documentation of clock synchronization procedures

Solution: Use Domain Time II documentation as necessary to write your procedures. At minimum, document which time sources you use and how often each component synchronizes with them. Be sure to indicate the path back to NIST time (i.e. your GNSS/GPS source is traceable to NIST).

Domain Time II is thoroughly documented, and the behavior of the Domain Time II system and each time component and how it synchronizes is detailed in the [online documentation](#). These documents can be used to provide any level of detail of the system operation for compiling your documented procedures.

- **FINRA Requirement:** Keep Logs of every time a clock is synchronized and the results of that synchronization

Solution: Use Domain Time II Audit Server to collect sync logs.

See the [Audit Server](#) documentation for details on configuring and using Audit Server.

Domain Time II Audit Server is capable of collecting a log of time sync activity from Domain Time II components into a central location for easy analysis and archiving. Information retrieved includes when a sync occurred, with whom the component synced, and amount the clock was corrected. Log retention is configurable to match archival schedules.

Audit Server also keeps an audit record which can be used to demonstrate on-demand that any particular machine was synchronized, with what source, and with what accuracy.

Domain Time II Server and Client also keep a local log that includes not only time sync events, but all other events activity and events by the component. These logs can be manually collected and archived to meet the log retention requirements, however doing so is typically much more complex than using Audit Server to do so, and results in significantly larger log files to be archived. In most cases, using Audit Server to collect sync logs is optimal.

Suggested Configuration Changes to Audit Server

Audit Server shares Domain Time Manager's view of the network. Adjust Manager's [discovery](#) settings to be sure you are able to see all the machines you need to audit. Be sure to Enable Auditing on your selected machines.

- **Set clock display granularity to your desired resolution:**

Domain Time keeps clock granularity internally to the limits of the operating system, however, you may display the time with any granularity you wish. This can be adjusted from the Manager menu. Choose *Options -> Appearance and Interface -> Format Options* and adjust the **Significant digits to show** item.

- **Discover machines for audit from all Domain Time II Servers:**

If you want to automatically audit all machines that synchronize with Domain Time II Server (this is a very robust choice), choose the [Audit Server -> Advanced -> Audit List Management](#) option from the Manager menu, enable the "Add machines that have synchronized with Domain Time II Server" option and enter the list of Domain Time II Servers you want to contact for their list.

- **Manually Enter Other Machines:**

Manually enter any machines not automatically discovered by the methods above. Enter machines to be added one at a time by right-clicking on the category where you want them to appear on Manager's Tree pane, or use Manager's [Batch Add](#) process for adding multiple machines.

- **Enable Central Log Collection:**

Use the [Audit Server -> Synchronization Logs -> Configure](#) menu item of Domain Time II Manager to collect Time Synchronization logs. Choose retention settings that correspond with your archival processes to ensure that all logs are transferred to archival storage before being deleted from the Audit Server.

- **Use the Conversions -> Daily Drift CSV file for analysis and compliance:** Audit Server can compile all Synchronization (Drift) Logs into a single [Daily Drift CSV file](#) which is ideal for complying with violation reporting or other tasks. This file can be provided to your Compliance Department for their use in compiling the necessary reporting forms and documents.

The Daily Drift CSV file can also be configured as an exception report, which only shows machines that have violated your alert thresholds. You may find this more useful than providing the entire drift data to compliance. To enable this function, select the **Only include error records - omit all records within defined tolerances** checkbox on the **CSV File Configuration** dialog.

- **FINRA Requirement:** The log should include notice of any time the clock drifts more than 50ms second from NIST time.

Solution: Domain Time II Audit Server has the capability to generate alerts when any monitored system's variance from a reference clock exceeds a threshold you set. Warning entries of these events are also included in the logs.

Reference Clock

Audit Server can compare the sampled time of any audited machine to a reference clock. The reference clock's time is used to calculate certain variances and alerts. By default, Audit Server shares the Reference Clock settings of Domain Time II Manager. Since FINRA specifies that variances be shown in relation to NIST, the [reference clock setting](#) on Manager must be changed to include a clock with as short a path to NIST time as possible (preferably a NIST server or a clock derived directly from it, such as a GPS time source).

Alert Thresholds

Audit Server has the ability to generate an alert if the time variance on any system exceeds a particular threshold. The FINRA-specified requirement is that the log for any machine drifts outside 50ms from NIST time should include a notice to that effect. Audit Server will automatically add a warning to the log when any machine exceeds the **Any machine time off by...** setting on the Audit Server [Alerts](#) dialog page.

Required Configuration Changes to Audit Server

Set the Reference Clock to NIST sources: Use Manager's [Options -> Network Options -> Reference Time...](#) menu selection to set the Reference Clock setting to use at least, preferably more of the official [NIST Servers](#) (note, you must have the NTP port 123 UDP open on your firewall to allow Manager/Audit Server to contact a NIST time server). You may also choose reliable local NIST-derived clocks, such as a GPS receiver.

Note that if your Audit Server is synchronizing via PTP to a GPS/GNSS synchronized Grandmaster, you may set the Reference Clock to "Use this machine's clock" as recommended in the configuration documentation. This still maintains traceability to the NIST source.

Set the Alert Threshold: On Audit Server's [Alerts](#) dialog page, make sure the **Any machine time off by** setting is set to 50ms or less.

- **FINRA Rule Requirement:** Logs must be maintained and preserved for the period of time and with the accessibility specified in SEC Rule 17a-4(b)

Solution: Use Domain Time II to collect audit logs and sync data and archive as necessary.

Rule 4590 specifies the retention period for this type of record is 3 years, the most recent 2 years of which must be in an easily accessible location.

Rule 6820 specifies the retention period for this type of record is 5 years, the most recent 2 years of which must be in an easily accessible location.

The Domain Time II Audit Server automatically collects detailed time synchronization data from the network into local disk storage. You may choose to keep the records locally or archive them into offline storage.

- **FINRA Rule 4590 Requirement:** Logs must be maintained in a format permitted under SEC Rule 17a-4(f)
Domain Time II does not directly address the specific provisions of this regulation (such as the use of non-erasable storage for electronic data records), however it does provide the data in an easily collected and stored manner that can be transferred to your required format.
- Violation self-reporting under CAT NMS. The [Daily Drift CSV file](#) can be used by your Compliance Department to determine if violations have occurred and provide the necessary documentation of events.

References

[FINRA Rule 4590. Synchronization of Member Business Clocks \(Financial and Operational Rules\)](#)

[FINRA Rule 6820. Clock Synchronization \(Consolidated Audit Trail Compliance Rule\)](#)

[SEC Rule 17 CFR 240 17a-4. Records to Be Preserved by Certain Exchange Members, Brokers and Dealers](#)

[CAT NMS PLAN Clock Synchronization Topics](#)

[CAT Alert 2020-02 - Standards for Self Reporting Deviations of Clock Synchronization Standards to FINRA CAT](#)

Disclaimer

This document is provided for informational and planning purposes only. The information used in compiling this document was obtained from publicly available sources and no representation is made as to the accuracy of the information, nor as to the accuracy of any reading or interpretation thereof. No warranty is made or implied regarding the usefulness or suitability of this information for a particular purpose. Further, Greyware Automation Products, Inc. is not liable for any damages, real or consequential, arising from use of this information.

Consolidated Audit Trail (SEC Reg NMS)

The United States Security and Exchange Commission (SEC) has established new rules governing financial market information, commonly referred to as Regulation or Reg NMS (National Market System). Reg NMS includes requirements for the synchronization of clocks and maintaining related audit trail information. These are covered in paragraph (d) of 17 CFR 242.613 (Rule 613):

- Clocks should be synchronized to NIST (National Institute of Standards and Technology) time.
- Synchronization standards should be reevaluated annually.
- Clock granularity should at minimum be recorded/reported in milliseconds and possibly finer increments as required by NMS.

FINRA has well-established procedures and policies regarding time synchronization that allow you to address the Reg NMS requirement for "reasonable policies and procedures". We recommend you use our instructions on [configuring for compliance with FINRA](#) to ensure the clocks on your various systems are synchronized and generating correct time-stamps and are collecting the necessary audit trail records to satisfy your data retention requirements.

References

[Rule 613 \(Consolidated Audit Trail\) Overview](#)

[17 CFR § 242.613 Consolidated Audit Trail regulations](#)

Disclaimer

This document is provided for informational and planning purposes only. The information used in compiling this document was obtained from publicly available sources and no representation is made as to the accuracy of the information, nor as to the accuracy of any reading or interpretation thereof. No warranty is made or implied regarding the usefulness or suitability of this information for a particular purpose. Further, Greyware Automation Products, Inc. is not liable for any damages, real or consequential, arising from use of this information.

The U.S. Food and Drug Administration (FDA) has adopted wide-ranging regulations governing activities in the life sciences industry. Among these are regulations regarding the use of electronic records and electronic signatures, contained in 21 CFR Part 11 of the Federal Register.

Implicit in these regulations is the expectation that systems used in producing electronic records or signatures must have the correct date and time so that time stamps produced are correct and verifiable. Extracts of the applicable sections of the Rule, current FDA Guidance to Industry memos, as well as relevant commentary from private industry that help define the time stamp requirements follow.

21 CFR Part 11 Final Rule, March 20, 1997 Section 11.10 Controls for closed systems specifies:

"(e) Use of secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records. Such audit trail documentation shall be retained for a period at least as long as that required to for the subject electronic records and shall be available for agency review and copying."

Draft Guidance for Industry on Part 11, Electronic Records, Electronic Signatures - Scope and Application:

Section III. C. 2. Audit Trail

"Even if there are no predicate rule requirements to document, for example, date, time, or sequence of events in a particular instance, it may nonetheless be important to have audit trails or other physical, logical, or procedural security measures to ensure the trustworthiness and reliability of the records."

Draft Guidance for Industry: Computerized Systems Used in Clinical Trials:

IV. Standard Operating Procedures, Data Entry, Section B. Electronic Signatures

"2. Personnel who create, modify, or delete electronic records should not be able to modify the audit trails."

"5. Audit trails should be created incrementally, in chronological order, and in a manner that does not allow new audit trail information to overwrite existing data in violation of §11.10(e)."

IV. Standard Operating Procedures, Data Entry, Section C. Date/Time Stamps

"1. Controls should be in place to ensure that the system's date and time are correct."

"2. The ability to change the date or time should be limited to authorized personnel and such personnel should be notified if a system date or time discrepancy is detected. Changes to date or time should be documented."

Summary of Requirements

Four basic requirements regarding time synchronization for FDA regulatory purposes can be derived from the above:

■ Time and Date on Systems Must Be Correct

Standard best practices to ensure the time and date on all computer clocks is correct is to verify they are reliably synchronized to a commonly used time standard, such as provided by the National Institute of Standards and Technology (NIST) or United States Naval Observatory (USNO) atomic clocks, or Stratum 1 time sources such as GPS receivers.

■ Ability to Change the System Clock Must Be Restricted

Users must not be allowed to change the time on systems at will.

■ Alerts to System Date/Time Discrepancies Must Be Generated

Administrators must be notified immediately of problems with the time on systems.

■ An Audit Trail to Demonstrate the Validity of Time Stamps Must Be Maintained


The accuracy of the date/time used to produce time stamps is integral to proving the trustworthiness and reliability of all electronic records and signatures. Non-modifiable records of the generation of time stamps (including the time synchronization process) are essential.

How to Use Domain Time II to comply with the FDA 21 CFR Part 11 Requirements

Domain Time II meets or exceeds all of the guidance and directives detailed above. Properly configured, Domain Time will allow you to easily comply with all current and proposed requirements to assure the validity of the date and time of time stamps and related audit requirements.

Domain Time II is designed specifically to provide both accurate time synchronization and a complete history of that synchronization. Each Domain Time II time sync component (Servers and Clients) have the ability to keep detailed logs and statistics of their own activity - and, critically, to **report that information automatically** to monitoring and auditing systems when requested.

This diagram shows the basic structure of the Domain Time II system, showing how time synchronization and audit data collection are handled.



Configuring for compliance

These are the steps necessary to use Domain Time II to assist in compliance with FDA 21 CFR Part 11:

- **Requirement:** Time and Date on Systems Must Be Correct

Solution: Configure Domain Time II to provide accurate time sync to all clocks

Domain Time II, when installed according to the instructions found on the [Recommended Configurations](#) page of the Domain Time II documentation, will meet the FDA 21 CFR Part 11 requirements for time synchronization. However, there are a few additional configuration considerations beyond the standard recommend installation instructions for FDA 21 CFR Part 11 compliance.

Required Configuration Change

Set the Domain Time II Master to use NIST, USNO, or other reliable Stratum 1 system as its time source:

Use the **Time Sources** tab page of the Domain Time II Server Control Panel Applet to include any of the following entries in the Sources fields (in any order):

- time.nist.gov
- ntp2.usno.navy.mil
- The IP address or DNS name of a Stratum 1 source, such as GPS receiver

- **Requirement:** Ability to Change the System Clock Must Be Restricted

Solution: Domain Time II Server and Clients have a function called **Clock Change Monitor** that prevents users without administrative rights from changing the time on their systems. This function is enabled by default. No additional

configuration is necessary.

- **Requirement:** Alerts to System Date/Time Discrepancies Must Be Generated

Solution: Use Domain Time II Audit Server (see below) or [Monitor Service](#) to provide alerts when time variance on systems exceeds the thresholds you set.

- **Requirement:** An Audit Trail to Demonstrate the Validity of Time Stamps Must Be Maintained

Solution: Use Domain Time II Audit Server to collect and preserve tamper-proof audit records

Domain Time II has exceptional built-in functionality to easily and automatically provide an audit trail of time synchronization. The following steps will provide unimpeachable records of the state of time synchronization. You will know exactly when your changed, why, when, and with whom they synchronized.

- **Set Reference Clock and System Alerts.**

Domain Time II Audit Server has the capability to generate alerts when any monitored system's variance from a reference clock exceeds a threshold you set. Warning entries of these events are also included in the logs.

Reference Clock

Audit Server can compare the sampled time of any audited machine to a reference clock. The reference clock's time is used to calculate certain variances and alerts. By default, Audit Server automatically locates the nearest Domain Time II Server to use as its reference clock. Since OATS specifies that variances be shown in relation to NIST, the reference clock setting must be changed.

Alert Thresholds

Audit Server has the ability to generate an alert if the time variance on any system exceeds a particular threshold. The OATS-specified requirement is that the log for any machine drifts outside 3 seconds from NIST time should include a notice to that effect. Audit Server will automatically add a warning to the log when any machine exceeds the **Any machine time off by...** setting on the Audit Server [Alerts & Logs](#) tab page.

Required Configuration Changes to Audit Server

Set the Reference Clock to NIST: Go to Audit Server's **Advanced** tab and change the Reference Clock setting to use time.nist.gov (note, you must have the NTP port 123 UDP open on your firewall to allow Monitor to contact the NIST time server).

Set the Alert Threshold: On Audit Server's **Alerts & Logs**, make sure the **Any machine time off by** setting is set to 3 seconds.

- **Set Audit Record Collection and Retention Options**

Use Domain Time II to collect audit data and archive as necessary.

The Domain Time II Audit Server automatically collects detailed time synchronization data from the network. Data can be kept in local storage for the required three year period, or archived off to permanent offline storage as necessary.

- **Document procedures**

Use Domain Time II documentation as necessary to write your procedures and as a baseline for your "black box" system validation testing.

Domain Time II is thoroughly documented, and the behavior of the Domain Time II system and each time component and how it synchronizes is detailed in the [Technical Information](#) and [Configuration](#) and sections of the online documentation. These documents can be used to provide any level of detail of the system operation for compiling your documented procedures.

In particular, the [Resource Impact Statement](#) contains a great deal of useful material as a starting point for "black box" testing of Domain Time II for use on your validated systems.

- **Optional: Keep Detailed Logs of every time a clock is synchronized and the results of that synchronization**
Use Domain Time II Audit Server to collect the synchronization logs of all machines.

See the [Audit Server](#) documentation for details on configuring and using Audit Server.

Domain Time II Audit Server is capable of collecting a log of time sync activity from Domain Time II components into a central location for easy analysis and archiving. Information retrieved includes when a sync occurred and with whom the component synced, and amount the clock was corrected. Log retention is configurable to match archival schedules.

Audit Server also keeps an audit record which can be used to demonstrate on-demand that any particular machine was synchronized, with what source, and with what accuracy.

Domain Time II Server and Client also keep a local log that includes not only time sync events, but all other events activity and events by the component. These logs can be manually collected and archived to meet the log retention requirements, however doing so is typically much more complex than using Audit Server to do so, and results in significantly larger log files to be archived. In most cases, using Audit Server to collect sync logs is optimal.

Required Configuration Changes to Audit Server

See [Audit Server](#) documentation for configuration details.

References

[21 CFR Part 11 Final Rule, March 20, 1997](#)

[Draft Guidance for Industry on Part 11, Electronic Records, Electronic Signatures - Scope and Application, February 2003](#)

[Guidance for Industry: General Principles of Software Validation;](#)

[Final Guidance for Industry and FDA Staff, January 11, 2002](#)

[Guidance for Industry: Computerized Systems Used in Clinical Trials, April 1999](#)

Disclaimer

This document is provided for informational and planning purposes only. The information used in compiling this document was obtained from publically available sources and no representation is made as to the accuracy of the information, nor as to the accuracy of any reading or interpretation thereof. No warranty is made or implied regarding the usefulness or suitability of this information for a particular purpose. Further, Greyware Automation Products, Inc. is not liable for any damages, real or consequential, arising from use of this information.

Markets in Financial Instruments Directive II (MiFID II)

The European Union has established new comprehensive rules for investment services, commonly referred to as MiFID II (or MiFID 2). The regulations include requirements for the synchronization of clocks and maintaining related audit trail information. In particular, Article 50 of MiFID II applies to trading venues and their members and participants and requires them to comply with accuracy requirements regarding the maximum divergence of their business clocks from UTC and to timestamp reportable events to a specific granularity.

Basic Regulatory Requirements

The MiFID II Draft regulatory technical standards on clock synchronization (Regulatory Technical Standards, RTS, 25) specifies these basic requirements:

- **UTC Reference Time**

Clocks must be synchronized to UTC (Coordinated Universal Time).

"Operators of trading venues and their members or participants shall synchronise the business clocks they use to record the date and time of any reportable event with the Coordinated Universal Time (UTC) issued and maintained by the timing centres listed in the latest Bureau International des Poids et Mesures Annual Report on Time Activities. Operators of trading venues and their members or participants may also synchronise the business clocks they use to record the date and time of any reportable event with UTC disseminated by a satellite system, provided that any offset from UTC is accounted for and removed from the timestamp."

- **Maximum divergence from UTC**

Clocks must be synchronized to UTC within a minimum accuracy target. The accuracy target differs by the type of trading system (computerized trading vs. high-frequency trading). Machines using high frequency algorithmic trading technique must synchronize to within ± 100 microseconds of UTC. All other computerized trading systems must be synchronized to within ± 1 millisecond of UTC. (RTS Annex Table 1)

- **Timestamp Granularity**

Timestamps must be recorded to specified precision for the type of trading system (computerized trading vs. high-frequency trading). Machines using High frequency algorithmic trading technique must have a timestamp precision of at least 1 microsecond. Note that does not mean clocks must be synchronized to microsecond accuracy, but that the number of decimal places used in measuring/recording the timestamps shows the data in microseconds. All other computerized trading systems must record timestamps with at least 1 millisecond precision. (RTS Annex Table 2)

- **UTC Traceability**

Timestamps must be documented as traceable to UTC.

"Operators of trading venues and their members or participants shall establish a system of traceability to UTC. They shall be able to demonstrate traceability to UTC by documenting the system design, functioning and specifications. They shall be able to identify the exact point at which a timestamp is applied and demonstrate that the point within the system where the timestamp is applied remains consistent. Reviews of the compliance with this Regulation of the traceability system shall be conducted at least once a year."

How to use Domain Time II to comply with MiFID II time sync regulations.

You can easily meet or exceed the MiFID II requirements by using Domain Time II Servers, Clients, Manager, and Audit Server components to synchronize, measure, and document your time synchronization.

- **UTC Reference Time**

All Domain Time components can be set to use UTC time sources, either for time synchronization or to use as reference time for accuracy measurements. How to do so specifically depends somewhat on your chosen [time distribution hierarchy](#). In general, it's best to have at least one satellite-synchronized hardware time appliance (GPS, GNSS, etc.) on premises to use as your top-level UTC source. This device may serve NTP and/or PTP protocols.

You may use the distributed domain hierarchy common with NTP or DT2 protocols (i.e. UTC hardware clock -> Domain Time Server(s) -> Domain Time Clients) or a flat-hierarchy such as used by PTP (all Clients/Servers connect directly to the UTC hardware clock). In either case, Domain Time components will log all transactions showing their time source, so you have direct traceability to UTC for time synchronization.

In addition, Domain Time Manager (and by extension, Audit Server) can be set to use UTC sources as the [reference clock](#) for all time delta measurements, alerting, and time auditing so you can use and create documentation of synchronization showing UTC traceability. Domain Time Manager/Audit Server can [monitor](#), [audit](#), and [retain](#) time synchronization records from Windows, Linux, PTP devices, and more.

■ Maximum divergence from UTC

As indicated, the regulations require different levels of time synchronization accuracy depending on your trading application.

High-frequency algorithmic trading systems

These systems must be synchronized within ± 100 microseconds of UTC. Although recent versions of Windows (2012/Win8/2016/Win10 and later) may be able to achieve this accuracy on well-provisioned modern hardware using only NTP, the most practical way to achieve this level of accuracy (and the only way on Vista, 2008, 2008r2, or Win7 systems) is to use PTP (IEEE1588-2008/2019). See the [PTP documentation](#) on how to configure Domain Time for PTP synchronization. Follow the instructions there carefully to achieve maximum accuracy.

Other computerized trading systems

These systems must be synchronized within ± 1 millisecond of UTC. This is easily achievable on any version of Windows using the NTP, DT2, or PTP protocols. Use these recommended settings:

- If using a distributed hierarchy (UTC clock -> Domain Time Server(s) -> Domain Time Clients), use at least 3 time sources on your top-level Servers (or, if only one UTC source is available, set to sample it 3 times)
- Use a fixed synchronization rate of at least 1 per minute on all your Servers and Clients

■ Timestamp Granularity

All Domain Time components are capable of hectonanosecond precision internally (that's one more decimal point of accuracy beyond the required 1 microsecond precision).

Domain Time Server and Client record all time transactions in hectonanoseconds. Domain Time Manager and Audit Server can also record time deltas, measurements, and audit data in up to hectonanosecond precision. The level of precision is configurable on Manager's [Options -> Appearance and Interface -> Format Options](#) menu page. Use the *Significant digits to show*: dropdown to select Microseconds.

Reading the Windows local clock in high precision from your application

Note that although both Domain Time and the operating system internally keep time to hectonanosecond precision, the normal Windows clock API `GetSystemTimeAsFileTime` only provides time to a best-case precision of 1 millisecond.

If your version of Windows is XP, 2003, Vista, 2008, or Win7, you will need to use a .dll from the Domain Time Software Development Kit (SDK) in order to read the system clock in the required microsecond precision for HFT systems. The SDK can be obtained from [Greyware Automation Products, Inc.](#)

For OS versions Win8, 2012, 2012r2, Win10, 2016 or later, you may use the built-in `GetSystemTimePreciseAsFileTime` API.

■ UTC Traceability

As mentioned above, you may configure Domain Time to obtain its time from UTC sources, and Domain Time Manager and Audit Server can also be configured to use UTC as the [reference clock](#) for all all measurement and auditing purposes.

This information is maintained in all [audit records and synchronization logs](#) collected by Audit Server. Using this information, you may easily demonstrate the synchronization status and UTC provenance of any machine at any historical point.

You may use this information as part of your documentation in demonstrating design and compliance.

References

[Final Report - Guidelines on transaction reporting, order record keeping and clock synchronisation under MiFID II](#)

[ESMA draft Technical Standards submitted to the European Commission on 28 September 2015 \(ESMA/2015/1464\)](#)

[Consultation Paper on Guidelines on transaction reporting, reference data, order record keeping & clock synchronisation \(ESMA/2015/1909\)](#)

Disclaimer

This document is provided for informational and planning purposes only. The information used in compiling this document was obtained from publically available sources and no representation is made as to the accuracy of the information, nor as to the accuracy of any reading or interpretation thereof. No warranty is made or implied regarding the usefulness or suitability of this information for a particular purpose. Further, Greyware Automation Products, Inc. is not liable for any damages, real or consequential, arising from use of this information.

Upgrading from previous versions

Information about upgrading to Domain Time II 5.2 from a previous version or from other time software.

Note: Please review the version 5.2 [System Requirements](#) and [Changelogs](#) carefully. They contain critical information you will need to successfully upgrade and use Domain Time II v5.2.

■ Upgrading from Domain Time 1.x (DT1)

If you are upgrading an existing Domain Time 1.x installation to 5.2, upgrade your PDC first, then any DCs, then finally any clients. This will ensure that all down-level machines continue working while the upgrade is in progress. Don't forget that you can install or upgrade over the network. See the [Rolling Out Domain Time II](#) page for instructions on performing remote upgrades.

■ Upgrading from Domain Time II, versions 2.1 or later

IMPORTANT: If you are upgrading from version 4.x, be sure to read the [4.x to 5.x Considerations](#) page

In general, Domain Time II Version 2.1 and later Servers and Clients are completely interoperable (with the exceptions noted below), see below), so you can upgrade any machine at any time without disturbing the program's operation.

For best results, however, you should upgrade in this order:

- Manager (and Audit Server, if installed)
- Master Server
- Slave Servers
- Any other Domain Time Server
- Clients

We recommend using the [Domain Time II Manager](#) to easily upgrade all your existing machines. Manager can investigate your network, automatically upgrade all your existing Domain Time machines, and provide a list of other machines that need to be upgraded manually. See the [Network Rollout](#) page for instructions on performing remote upgrades.

Upgrades are available using the following methods:

- Download the latest upgrade patch file from our website.
- Contact us to obtain the new distribution files on physical media.

Potential incompatibilities:

Items that may cause issues when interoperating with older versions of Domain Time II Versions:

- Servers prior to build 2.5.b.20030212 cannot provide the Domain Time II TCP protocol.
- Only version 3.1 and later Clients can understand DT2 Broadcasts that include timestamps ("Heartbeats (with Data)").
- As of v5.1, Domain Time uses both TCP and UDP port 9909 (DT2) for basic operations. Firewalls will need to pass both types of DT2 traffic, even if NTP (port 123 UDP) is actually being used to synchronize the time.
- Masters and Slaves from versions prior to v5.x will interoperate with newer Master and Slaves, but will not replicate security parameters or other advanced v5.x (or later) information.
- The Foreign Slave function of Masters and Slaves has been deprecated as of v5.1. Any older Servers set as Foreign Slaves will need to be re-configured to get their time explicitly from v5.1 or later Masters.
- v5.1 or later Clients have stricter requirements for assigning time servers using DHCP Time Server Options than do older versions, and may interpret existing DHCP settings differently. See the [Discovery Options](#) section of the

Client's **Obtain the Time** page for details.

- v5.1 or later Clients only use the DT2 and/or NTP time protocols to obtain the time. v5.1 or later Servers provide obsolete protocols such as TIME/ITP, Daytime, etc. for compatibility with other devices, but newer Domain Time Clients cannot use them.
- Proxy support for obtaining time using the DT2 over HTTP protocol through firewalls has been deprecated as of v5.1. Any existing proxy connections on older versions will need to be converted to direct-to-IP-port connections after upgrade.
- v5.x or later Clients are only able to match an older Server's timezone if the Server is configured to "Recommend timings to clients that ask for guidelines" (found on the [Client Timings](#) page). v5.1 or later Servers will also need to enable the "Allow client's to match this server's timezone" checkbox on the [Recommendations](#) property page.
- As of v5.1, Audit Server requires **both** Domain Time II Server and Domain Time II Manager be installed on the local machine before installation/upgrade.
- As of v5.1, Windows Time Agent (WTA) is no longer installed automatically along with Server or client. WTA can be installed manually from the distribution files, if needed.
- As of v5.1, Manager cannot install Windows Time Agent (WTA), although it can monitor or remove it. WTA can be installed manually from the distribution files, if needed.
- If you are using [SNMP Trap](#) functions, be sure to extract the latest MIB file from Domain Time Server or Client Control Panel applet after the upgrade and update your Network Management System(s).

■ Upgrading from other time sync software

You should not install two time sync programs on the same machine. The results of running two time services simultaneously are unpredictable, and usually not pretty. With the exception of Microsoft's Windows Time Service on Win2k or later (which Domain Time II handles automatically during installation), you must remove any other time management program before installing Domain Time. See the removal instructions provided by the manufacturer.

Microsoft's [Windows Time](#) service

Domain Time is much more accurate, reliable and easier to maintain than Windows Time, and also provides many additional important benefits without sacrificing any functions that Windows Time provides.

If you do a fresh install Domain Time Server v5.1 or later on a machine running Windows Time, Domain Time will automatically disable the Windows Time service (since Server provides all necessary time services). A fresh install of Domain Time Client will also disable Windows Time, except on Domain Controllers and Cluster Servers, where Windows Time should be set to run in **NoSync** mode for compatibility reasons.

If you are upgrading a machine running Domain Time, any existing Windows Time settings will not be changed. You should manually disable Windows Time on any Domain Time Servers after upgrade. See the [Co-existing with the Windows Time Service](#) page for more details.

4.x to 5.x Considerations

- Please see [Changelog](#) for detailed information on all changes and updates.
- Review this page for information on important differences from version 4.x and earlier.

Domain Time II Verison v5.x introduces significant performance improvements, new features, and other enhancements to the Domain Time II product line. This document does not cover all of the changes, but it does address important items to be aware of when installing or upgrading to version v5.x from v4.x.

See the online documentation for more complete information on these items.

Notable changes from version v4.x

- Version v5.x only runs on Windows XP and above. Win9x, NT, or Win2000 systems should continue using version v4.x.
- The three main types of v4.x Windows Client services (Full Client, Thin Client, and Ultra Thin Client) have been combined into a single 5.x Domain Time II Client. Existing v4.x Full, Thin, or Ultra Thin Clients will be converted to v5.x Domain Time Client during upgrade. Existing settings are preserved during the upgrade, so the v5.x Clients will continue to perform the same functions as the v4.x Clients did. EXCEPTION: The Thin client and Ultra Thin Clients may need to be configured via the Control Panel Applet if you weren't using the default settings on those Clients.

Note that v4.x Thin and Ultra Thin Clients did not have Control Panel Applets, but the v5.x Client does. You may delete the domtimec.cpl file from the \System32 folder after installation/upgrade if you do not want the applet to appear.

- The v4.x DOMTIME.INI template configuration file is no longer used to determine the installation defaults for Server and Client. Instead, a standard Windows registry file (dtserver.reg for Server or dtclient.reg for Client) holds all the default settings.

The new .reg file can be edited in any text editor. The Control Panel's advanced Import/Export page will read or write the file, making sure that only appropriate entries are loaded or saved. The Import/Export page also checks to make sure the .reg file is the correct version (v5.x) and type (server or client) for the machine.

- v5.x introduces the ability to configure many Client and Server options using Windows Group Policies. The distribution folder contains an administrative template file (domtime.adm) that you can use to populate objects in your Group Policy Object Editor. Load the domtime.adm template into the object as a "Computer Configuration" template. See the Microsoft documentation for details on using Group Policy templates.
- The Domain Time Windows Time Agent is no longer installed by default during installation of Domain Time Client or Server. The "Windows Time" button on the Advanced property page will only operate if the Windows Time Agent is already present (either from an upgrade from version v4.x or if manually installed from the distribution files).

Version v5.x allows you to disable the Windows Time service entirely in nearly all circumstances, so the additional configuration options the Agent provides for W32Time are not required.

- Masters and Slaves

The master server for any given domain may use the following options for obtaining the time:

- use its own clock (not recommended unless you have a third-party hardware device of some kind maintaining the clock)
- use a list of time sources
- use the PDC of another domain (by specifying the domain name on the time source setup dialog; the PDC is looked up dynamically)

Note that option (c) is not a v4.x-style foreign slave relationship; it merely allows the local PDC to easily obtain the time from another PDC in the forest. The v4.x "foreign slave" relationship where local PDCs would receive slave synchronization signals and replication settings from the foreign PDC does not exist in v5.x.

Option (c) can also be used by any other server or client, not just the domain master server. Again, using this option does not make the machine a slave; it merely lets you specify the domain without having to specify the PDC. When the PDC-emulator role shifts, machines will automatically start using the new PDC-emulator without additional configuration.

v4.x masters and slaves interoperate with v5.x masters and slaves, but will not replicate security parameters or other v5.x information.

- Clients using DHCP Discovery

Option 004 ("Time Servers") is used only for discovering DT2 servers. If a server is listed in option 004 that doesn't support DT2 UDP, it will be ignored.

DHCP Option 042 ("NTP Servers") is used to discover both NTP servers and DT2 servers. If a server is listed in option 042, it will be checked for NTP first. If NTP fails, it will be checked for DT2 UDP. If it does not provide time under either of these two protocols, it will be ignored.

- Clients set to match their server's time zone.

v5.x clients will only request timezone information if the server is configured to provide *both* recommended timings and timezone info. Clients may use either recommended times or timezone (or both or neither), but it won't ask for the server's timezone unless the server is also set to provide recommended timings. This is to cut down on the relatively-expensive timezone calculations needed by both the server and client when sharing timezones. It also reduces unnecessary network traffic.

Obtaining and Serving the time

- Domain Time components, except for test programs, only use DT2, PTP, and NTP protocols for getting the time. "DT2" includes the entire DT2 family: DT2 over UDP, DT2 over TCP, and DT2 over HTTP. All versions of NTP (1-4) are supported. Only version 2 (IEEE1588-2008/2019) of PTP is supported. Domain Time II Server *can* become a master PTP server, but normally operates as a slave, since software-based PTP lacks the accuracy of a proper PTP appliance. Domain Time II Client can only operate as a PTP slave.

Other time protocols (TIME/ITP, Daytime, etc.) are of insufficient resolution for good timekeeping, and are therefore not used for obtaining the time on v5.x. Server continues to provide them, primarily to service legacy devices.

- HTTP proxy servers are no longer supported for getting the time via Domain Time over HTTP. You must have a direct connection in order to use this protocol. You may specify a non-standard port number by appending a colon and port to the server's name (e.g. timeserver.mynet.com:91). You may also specify the port in the port number field when editing a time source.
- You are no longer limited to four time sources and you may make multiple requests (samples) of each time source, with a configurable delay between samples.
- Time sample analysis is much more sophisticated than in version v4.x. Sample averaging and analysis is automatic based on how many servers (and how many samples per server) you select. Averaging is enabled by default, but you may turn it off. You may have multiple samples per server with or without using averaging (the multiple samples will go through the same statistical analysis as if they were individual samples from multiple servers).
- Symmetric Key Authentication

Version v5.x supports symmetric authentication (MD5 hash of shared secrets). This type of authentication works between

Domain Time v5.x servers and clients, or between Domain Time and any properly-configured NTP v3.x or higher daemon (such as ntpd or xntp on UNIX/Linux machines, hardware GPS clocks, etc.).

Domain Time server supports serving time with symmetric authentication for client-server requests on NTP, DT2-UDP, DT2-TCP, and DT2-HTTP. Domain Time server also supports broadcasting (both NTP and DT2-UDP) with a shared key and MD5 hash. Clients configured with the same key validate the sending server by comparing the computed hash.

Slaves automatically replicate shared secrets from their Master. Masters, Independent Servers and Clients must be provided with a shared secrets list, either by manually entering them into the Control Panel applet, by being upgraded, importing a Domain Time v5.x configuration .reg file, or importing/exporting a standard ntp.keys file. Clients running on domain member machines may also receive their shared secrets from a Windows Group policy.

Interaction with Windows Time (W32Time)

Windows Time clients using NT5DS mode (the default) search the Active Directory hierarchy to find a server. They send a request for the time using their machine RID as the authentication key, and expect the returned timestamp to be authenticated by the server. Only a DC in the client machine's domain can provide this type of authentication.

Domain Time v4.x Servers provided for Windows Time clients by setting the W32Time service's client portion to "NoSync" mode and allowing the W32Time service's server portion to serve NTP directly. Although the quality of the timestamps provided by W32Time is significantly degraded, this approach allowed the DC running Domain Time to continue serving Windows Time clients. This workaround is no longer necessary.

Domain Time v5.x provides integrated Windows authentication natively for both NTP and DT2 protocols. This means that W32Time clients in NT5DS-mode can get their time directly from any Domain Time II Server running on a DC, exactly as if getting the time from the Windows Time Service on that DC.

Additionally, Domain Time v5.x clients can use the same Windows authentication model to obtain NTP time from DCs running either the Windows Time service or Domain Time.

Windows authentication only works on domain member machines. The machine on which the client is running must be joined to the domain (or the forest) from which it gets the time. Windows authentication is automatic; no configuration is necessary. NOTE: While the domain member getting the time may be any kind of machine, the domain member providing the time must be a DC. Only a DC can validate the request. Other servers will not know the shared secrets.

W32Time in NT5DS-mode has distinct disadvantages:

- The W32Time NTP Server is inaccurate, so even if the DC's clock itself is well-synchronized, the time being served may not be.
- Other ntp clients (such as ntpd or xntp) cannot synchronize with it.

Domain Time's NTP Server has none of these disadvantages. It can provide standard NTP (with or without NTP auth) at the same time it provides NT5DS-mode timestamps, and all at extreme accuracy.

It is therefore highly recommended you install Domain Time II v5.x Server on all DCs. You *can* install Domain Time v5.x Clients on a DC, but you will then need to enable W32Time in "NoSync" mode to provide NT5DS-mode.

Recommended settings:

For v5.x Server on a DC:

- Verify that the "NTP Server Enabled" checkbox is checked on the Domain Time II Server "Serve the Time" property page AND
- the "Windows Time mode" dropdown on the "Advanced" property page is set to "Disabled"

For v5.x Client on a DC:

- The "Windows Time mode" dropdown on the "Advanced" property page should be set to "NoSync"

■ Reliable Time Provider

DcDiag and other tools sometimes expect the Windows Time service to be running, even if it's not actually doing anything. This error may be safely ignored as long as your DC is advertising as a time source. You can check your domain using our supplied tool ntpcheck.exe. Use `ntpcheck -ad` to check the domain for all types of advertised time sources, and then test each source.

Starting with v5.x, Domain Time Server, when installed on a DC, sets the system flags to indicate the machine is serving time and is a reliable time source. The DsGetDcName() function will report Domain Time Server v5.x machines on DCs as both time servers and reliable time sources. Domain Time Server on a non-DC will not change the existing system flags.

You may override this behavior by editing the registry. In HKEY_LOCAL_MACHINE\Software\Greyware\Domain Time Server\Parameters, edit (or create) a REG_SZ (string) value called "Set Reliable Time Provider" and set its value to either "True" or "False" (the English words, without the quotation marks). If this value is present and set to True, Domain Time Server will set the two flags even if it is not running on a DC. This configuration has no meaning for Active Directory, since only DCs are examined for the flags. Other tools, however, may benefit from knowing that a reliable time source is present. If this value is present and set to False, then Domain Time Server will not change the flags.

■ Cluster Service

The Windows Cluster has a default startup dependency on W32Time. It does not require the time service for any other purpose. Thus, the simple recommendation for installing Domain Time on clusters is to set W32Time to NoSync mode, which allows the service to be running to satisfy the startup dependency, but allows Domain Time to set the clock.

However, you may replace the cluster's startup dependency if you want. After installing Domain Time Client or Server, you can edit the "DependOnService" value in the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clussvc key to replace "W32Time" with "Domain Time Client" (or "Domain Time Server"). The cluster service will then wait until Domain Time has started before starting. You can then set the "Windows Time" setting on the Domain Time applet to Disabled.

Network considerations

- Domain Time now uses both TCP and UDP port 9909 (DT2) for basic operations. Firewalls will need to pass both types of DT2 traffic, even if NTP (port 123 UDP) is actually being used to synchronize the time.
- Version v5.x can use either IPv4 or IPv6 (or both) for obtaining and serving the time. IPv6 requires operating system support, which is present by default in Vista or above, but must be specifically installed/enabled on XP. Domain Time will function in IPv4-only mode if IPv6 is not present. If both are present, you may choose which to use, or let the system figure it out (see the Network property page on the Server or Client Control Panel applet).
- Version v5.x does not use the MS Networking "browse list" for primary machine discovery. Functions that formerly only used the browse list now use various methods of automatic discovery and configuration, including broadcast, multicast, and Active Directory enumeration via LDAP. The Browse list remains available as an additional secondary discovery method on some components
- Broadcasts and Multicasts

Previous versions of Domain Time depended on directed broadcasts to discover or signal machines on remote subnets. Multicasting is now preferred for signals that are sent to other subnets. Broadcasts are still used on the local subnet, primarily for automatic client/server discovery or signaling of advisories and cascades (see below). Some Domain Time components (such as Monitor and Audit Server) may still allow directed broadcasts for backwards-compatibility, but this is discouraged.

The "Broadcasts and Multicasts" property page on the Server or Client Control Panel applet shows the addresses and hop-count/TTL used for broadcasts and multicasts. These values pertain to the following functions:

- Server uses them to send cascades and advisories
- Server uses them as the listen addresses for IPv4 and IPv6 multicast requests
- Server uses them to send broadcast/multicast time (DT2 heartbeats and NTP time)
- Tools that don't have their own settings (for example, dtcheck.exe) use them for discovery and testing. Clients use them to discover DT2 and NTP sources. Clients use them as the listen addresses for IPv4 and IPv6 multicast requests

If you are trying to control overall broadcasting and multicasting, it is better to enable or disable the particular functions that use the addresses (such as on the Serve the Time property page) rather than disabling them on the Broadcasts and Multicasts page. Enabling or disabling on the Broadcasts and Multicasts page can have unintended consequences -- you may be trying to keep clients from sending multicasts for discovery, but end up preventing servers from communicating with their peers.

Note: NTP time broadcasts/multicasts

In order for Domain Time to send or receive NTP time broadcasts or multicasts, the Domain Time service must control the NTP port (123 UDP). If running Domain Time II Server, the Windows Time service must be disabled (this is the recommended configuration anyway). If running Client, the W32Time service must either be disabled (preferred) or set to NoSync mode.

- Cascades and Advisories

Cascade signals are used to keep the hierarchy in sync when a server sets its clock. v5.x cascade signals may be unicast, broadcast, or multicast (any combination). Each server has its own settings for whether or not it sends cascades, and if so, what type. Server can send broadcast IPv4 only, multicast IPv6 only, multicast IPv4 only, or any combination. IPv4 broadcasts are sent to 255.255.255.255. This is not configurable. If you need cascades to cross routers, you must use IPv4 or IPv6 multicast instead.

Clock Control

Domain Time v5.x includes significant improvements and optimization of all timekeeping functions to maximize the accuracy and precision of clock synchronization and timestamps. The default timing settings are usually sufficient to obtain superior synchronization on most systems.

However, in order to provide for tuning to achieve maximum accuracy (and to deal with the occasional poor-performing clock), v5.x exposes or adds a large number of advanced clock control options. See the online documentation for details.

- Leap Seconds

The Windows family of operating systems does not support leap seconds natively. Leap seconds are simply unexpected one-second corrections as far as the operating system is concerned.

Version 4.x of Domain Time applied leap seconds at the first timecheck following the leap, discovering its time was off by one second, and performing a correction. Version 5.x applies leap seconds at 23:59:59 UTC on the last day of the month in which the leap is scheduled. You may disable the new behavior by unchecking the "Enable advance scheduling of leap second corrections" checkbox on the Advanced tab of the Control Panel applet, but we recommend you leave it enabled, since the leap second specification requires that all clocks everywhere in the world should change at exactly the same time (UTC), regardless of time zone or physical location.

Domain Time acquires pending leap second information from NTP, PTP, or other GPS-derived time sources. Once leap second information is acquired, Domain Time Server will advertise the leap second when it serves NTP, PTP, or any of the DT2 protocols. In order for Domain Time to acquire, schedule, and advertise a leap second, all of its queried sources must agree that a leap is pending. If the sources disagree, then the leap will be handled at the next timecheck after it

occurs, and a warning notice that the leap indicators are inconsistent will be placed in the log. GPS-derived time sources acquire knowledge of a pending leap second from satellites well in advance of the event, but each manufacturer is free to decide how long in advance to pass this information to their clients. Most start advertising a pending leap second between 1 and 24 hours prior to the event, although some will advertise it days early, and some (broken) servers will continue to advertise it for some time after the event has occurred. Since leap seconds are to be applied at 23:59:59 UTC on the last day of the month, it doesn't matter if the event is advertised early. However, Domain Time has protection built in to prevent another leap second being scheduled for the following month if a broken source continues to advertise after the event.

Pending leap information is queried with each timecheck, and maintained only while the Domain Time service is running. Restarting the Domain Time service will clear any pending leap second corrections. If the leap is still pending when the Domain Time service is restarted, it will be rescheduled for the appropriate time. If the leap occurs while the Domain Time service is stopped, the leap will be applied at the first timecheck after startup. Domain Time also acquires and remembers the current TAI-UTC offset (only available from PTP sources), and this information (if ever acquired) is automatically updated after the application of a leap second.

■ Clock Corrections vs. Alignments

Domain Time can correct the clock either by "stepping" (immediately changing the time) or "slewing" (changing the time slowly). Stepping and slewing only operate on variances of 1 millisecond or more.

Variances of less than 1 millisecond are "aligned," which is a process very similar to slewing. Aligning involves temporarily speeding up or slowing down the computer to make it match the time source more closely. Sub-millisecond alignments are NOT considered corrections, and will not show as corrections in Audit Server, Domain Time statistics, or the drift graph. Variances of less than 1 millisecond will be reported as zero milliseconds, except in the log file.

If your machine is stepped, the log file will say "Local clock stepped" (followed by details on which direction, by how much, and the protocol used to obtain the time).

If your machine is slewed, the log file will say "Local clock slewed" (followed by the same details as for stepping).

If your machine is aligned, the log file will say "Local clock aligned" (followed by the same details as for stepping or slewing).

Alignments happen automatically as long as slewing is enabled. The only important thing to remember about alignments is that they are not reported as clock corrections.

■ Never Step

By default, Domain Time will step corrections too large to slew (or if slewing in that direction is disabled), and will also step the very first correction after rebooting. In v4.x, you could change this behavior by enabling the "Never Step Clock" option. In v4.x, "Never Step" really meant "Do not step except on first boot or when triggered by an administrator," which was a bit confusing.

In v5.x, if Never Step Clock is enabled, Domain Time really will never step the clock. The slew limits and direction permissions are not overridden by triggers, the Control Panel applet, or reboot detection. As a result, if you have Never Step Clock enabled, you will probably have to set the clock manually after every boot to get the time within range to begin slewing.

To provide greater control of the stepping process, v5.x introduces the "Allow Stepping" setting. Allow Stepping is a bitmask of reasons to allow stepping. If your v4.x machine had Never Step specified in the registry, the value will be translated to an Allow Stepping value of zero when upgrading to v5.x. In all cases, stepping will only be applied if slewing is disabled or cannot correct the variance.

Wait for first synch

Some third-party time-sensitive applications or services are set to auto-start when the machine boots, but may need to

have the clock synchronized before providing services. Recall that the CMOS clock chip may be wildly inaccurate, and therefore the first synchronization after boot is normally treated specially, allowing jumps in time either backward or forward.

NOTE: Setting your service to have a dependency on Domain Time is not sufficient, because this will only make your service wait until Domain Time is running. Service startup dependencies don't have any way to check to see if Domain Time has finished synchronizing the clock after starting.

v5.x Servers and Clients export a Win32 named event your processes can monitor to determine when the clock is synchronized. If the event is unsignalled, Domain Time could not synchronize the clock (or has not synchronized it yet). If the event is signalled, Domain Time has successfully synchronized the clock at least once since the service started.

To monitor this event in your own application or service, use the Win32 API `OpenEvent` to obtain a handle to "Global\domtime-sync-status-synchronized" (case-sensitive), and then use any of the Win32 wait functions. For example,

```
DWORD WaitForSync()  
{  
    HANDLE hHandle = OpenEvent(STANDARD_RIGHTS_READ | SYNCHRONIZE,  
                               FALSE,  
                               "Global\\domtime-sync-status-synchronized");  
  
    if (hHandle == NULL) return GetLastError();  
    WaitForSingleObject(hHandle, INFINITE);  
    CloseHandle(hHandle);  
    return NO_ERROR;  
}
```

The code snippet above tries to open the named event. If unsuccessful, it will return the error code. Otherwise, it will wait for the event to become signalled. If the event is already signalled, the wait will complete immediately. As soon as the event becomes signalled, the snippet closes the handle and returns `NO_ERROR` to let you know that Domain Time has successfully synchronized the clock.

In your own code, you probably want to include more error checking, and allow for a timeout in case Domain Time isn't running or never manages to synchronize the clock.

Other Items

■ New command-line option on DTServer and DTClient

v5.x adds a command-line option "-reset" or "-re". This option is useful only when combined with the upgrade option "-upgrade" or "-u" -- if specified, the upgrade will read the initialization .reg file as with a fresh install (i.e., overwriting any existing values). If not specified, upgrade will leave any existing values intact, other than necessary housekeeping to convert values to the current version's format.

■ Service Status Monitor protocol

Version v4.x supported the service status monitor, but it was an undocumented feature used by only a few OEM customers. Version v5.x supports the same protocol unchanged, but exposes it on the Control Panel applet for easier configuration.

The service status monitor is a simple TCP/IP listener to which your own programs can connect to check the status of the Domain Time service. By default, it supports both UDP and TCP on port 9911.

For UDP, use `sendto` to send an empty (zero length) packet to the target port and then use `recvfrom` to get the reply. For TCP, use `connect` and then `recv` (you do not need to send any data). The service status monitor will reply to either

connection with a single text line (CRLF terminated) indicating the status and version.

- **Audit Server autodiscovery of Linux domtimed clients**

Older Linux clients do not send their serial number with a time request, so Domain Time Server does not record them when they get time, and Audit Server does not know of them by examining the ephemera data. Upgrade to the newest Linux client if you need Audit Server to discover your Linux machines automatically.

- **DTRCPL (DT Remote CPL program)**

The x64 version will run only on x64 systems, and can control either x64 or x86 remote systems. The x86 version will run only on x86 systems, but can control either x64 or x86 remote systems. If the architectures of the local system and the remote system match (both x86 or both x64), then DTRCPL will try to load the CPL installed in the remote system32 folder (in case it happens to be an older or newer version of the CPL). If the architectures do not match, or if the CPL was not found on the remote system, DTRCPL will look in the local system32 folder and the Manager folder (if Manager is installed) to find the appropriate CPL.

- **DTSLEW**

This program may produce less precise results on Vista/2008/Win7 than it does on NT4/XP/2003 or on newer operating systems. If the imprecision shows, it is only with very small corrections (milliseconds or seconds) over large periods of time (minutes or hours). This is a known limitation arising from Microsoft's virtualization of the timing scheme on those operating systems.

Domain Time II Manager

When performing a remote upgrade of a software component using Manager, the selected template .reg file's contents do not replace the existing registry contents on the target machine. Only settings present in the .reg file that do not exist in the registry will be added. Existing registry values will be unchanged.

During remote installation, or during Reset Configuration of a software component from Manager, the contents of the .reg file will be added to the target machine's registry, overwriting any matching values already present. Values present in the registry but absent from the reg file will not be deleted (unless the reg file contains delete instructions for values or keys).

Manager's Reset Configuration procedure doesn't touch the dtserver.reg or dtclient.reg installation defaults files on the target machine; it creates a dtupdate.reg file instead, which the service then imports and renames.

Setup

Domain Time II components can be installed on individual machines using the Setup program detailed below. If you've purchased a Domain Time II package that includes [Manager](#), the recommended procedure to easily install Domain Time on multiple machines is detailed on the [Network Rollout](#) page. You may also install some components using the [command-line](#) interface.

Several Domain Time II components for Windows (Server, Client, Management Tools, Windows Time Agent, etc.) can be obtained ala carte, in which case they each come with their own installation package. If you need to install a component separately, click a link below to jump directly to its own installation instructions page:

- ▶ [Server](#)
- ▶ [Client](#)
- ▶ [Manager](#) (Management Tools)
- ▶ [Audit Server](#)
- ▶ [Windows Time Agent](#)
- ▶ [LMCheck](#)
- ▶ [Software Development Kit \(SDK\)](#)

IMPORTANT: If you will be installing Domain Time II onto machines with AMD processors, we highly recommend you update your processor drivers (a.k.a. PowerNow!) to the current version for your operating system available from AMD's website to avoid known hardware timing issues. Please see this article from our knowledgebase for more info: [KB2007.817](#).

Distribution Files

If you downloaded your Domain Time II distribution files, they came compressed as either a .ZIP or a self-extracting .EXE file. Make a new, blank folder on the machine on which you will be installing and copy the distribution file(s) to it. **Note:** Do not use a temp folder that has other files in it - they can cause conflicts with the file extraction or installation.

Run the .EXE file (if that's what you downloaded) or use an unzip utility to extract the contents of the .ZIP file into this blank folder. Be sure to tell your unzip program to extract using the original folder names option so that files will be extracted into the correct folders.

If you received Domain Time II on physical media (such as a CD), the files will already be in uncompressed format and you can run the Setup program directly from the media.

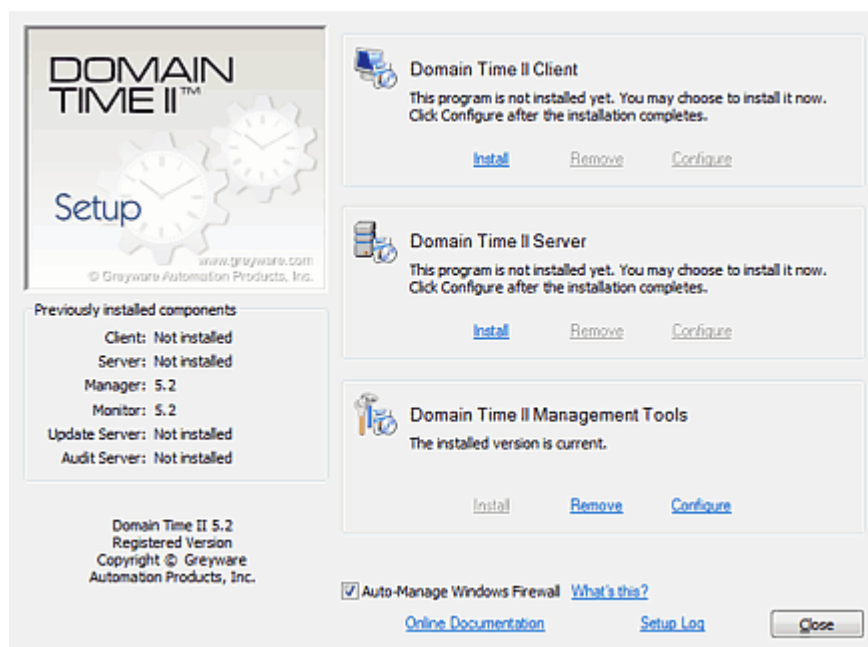
To 32-bit or not 32-bit?

Domain Time II programs come in either 32-bit or 64-bit versions. The 32-bit files are found in the **/i386** folder of the distribution; the 64-bit version are in the **/amd64** folder. The Setup program located in the root folder of the distribution will detect and install the correct version for the machine on which you are installing.

There are also setup programs located in both the **/i386** and **/AMD64** subfolders. Running those setup programs will run and install only on their own version (i.e. the **/AMD64/SETUP.EXE** program will only run on 64-bit systems). Although the 32-bit version of Domain Time programs may run on 64-bit systems, you should always install the 64-bit version on 64-bit systems (and vice versa).

Using the Setup Program

To start Setup, click on the **SETUP. EXE** program in root folder of the distribution files.



Installing Components

Setup will allow you to install any Domain Time component included in your distribution files.

If your distribution files include the Management Tools, install it first. You can then use Domain Time II Manager to install Domain Time II components to any other machines on your network (see the [Network Rollout](#) page for details).

Upgrading Components

Setup will also assist you in upgrading Domain Time II from previous versions. If you're running an earlier version of Domain Time, the setup program will offer to upgrade it for you.

If you have Management Tools installed, upgrade it first. Then use Manager to upgrade other machines on your network (see the [Network Rollout](#) page for details).

Removing Components

Domain Time II components can also be removed using the Setup program. Run the Setup program, and if it finds Domain Time software is installed, you'll have the option to remove it.

Notes:

- As of Version 5.2.b.20150828, Setup includes the **Auto-Manage Windows Firewall** checkbox. When checked, components you install will have their automatic management of the Windows Firewall features enabled to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.
- Most Domain Time II components can be installed, upgraded, or removed from the command-line as well. See the [Command-line Setup Options](#) page for details.
- To avoid required reboots or failed upgrades/removals, be sure all Domain Time programs, utilities, and control panels are closed before running Setup.
- Setup.exe cannot remove the Management Tools if it is run from the **\Program Files\Domain Time II** directory. Copy the initial distribution files to a blank folder and run Setup from there.
- If you use cloned OS images to install machines, please read [this article](#) from our knowledgebase about configuring Domain Time II properly for your clone image.

Command-line Options

Most Domain Time II components can be installed or removed from the command-line, including on remote machines. If you prefer to write your own batch files, or want to install/remove individual components manually, you can use the command-line parameters.

Select Program

Pick the program you're interested in from the drop-down list, then click the "Select" button to see the command-line parameters for that program:

Domain Time II Client

Client Command-line Options:

`dtclient.exe [directive] [options] [\\targetmachine]`

If no command-line parameters are provided, the program will display a dialog box that lets you choose whether to install, upgrade, or remove the program.

Command line parameters can be provided with either a leading dash or leading forward slash. Slashes are shown below. Parameters can usually be abbreviated to the first letter of the command. For example, you can use `/install`, `/i`, `-install`, or `-i` to specify the install directive.

Directives:

Specify one of these directives to install, remove, upgrade, or display version information:

`dtclient.exe /version` - displays version information
`dtclient.exe /install` - installs the program
`dtclient.exe /remove` - removes the program
`dtclient.exe /upgrade` - installs the program, upgrading if necessary

Options:

The following options may be added to any directive:

Example: `dtclient.exe -install -nocpl -vq`

Installs the program without showing the control panel applet or any confirmation messages.

`-noack`

Suppresses confirmation questions and successful completion notices (but not error messages).

`-nocpl`

Does not start the control panel applet after installation.

`-quiet` or `-q`

Suppresses confirmation questions, successful completion notices, and the progress bar. `-quiet` implies `-noack`. Like `-noack`, `-quiet` will not suppress error messages.

`-veryquiet` or `-vq`

Suppresses all messages during installation or removal.

`-reset` or `-re`

Reloads the DTCLIENT.REG default settings template file (overwriting any existing values). This option may only be used with the `/upgrade` directive.

Settings Templates:

The Domain Time Client and Server installation routines use [template files](#) to configure most of the default settings of the component. These files are simply modified registry settings files (.reg) in REGEDIT4 format. You may use the included default settings .reg files or, as of version 5.2.b.20180303, you may specify custom files of your own.

The .reg files containing the manufacturer's default settings for Server and Client are located in the installation folders (i.e. `\i386` or `\amd64`) for each component (**DTSERVER.REG** for Server, **DTCLIENT.REG** for Client). If you decide to modify these files to change the overall defaults, be sure to make a backup copy so you can revert if necessary.

To specify a custom template, copy your custom .reg file to the installation folders (both /i386 and /amd64) and add this directive to your command line:

-template=TemplateFilename.reg

IMPORTANT: A number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. You should not simply export the **HKLM\SOFTWARE\Greyware\Domain Time Client** key and apply it to other systems. Files exported using Windows Registry Editor in registry formats other than REGEDIT4 may not even import correctly. Rather, use the [Import/Export](#) function to export a .reg template. Alternately, you may use the PrepClone feature of the DTCheck utility to prepare a machine before exporting the key (see this [KB article](#)).

Suppress Shortcuts

You may opt to suppress creation of the Domain Time shortcuts in the All-Users startup menu. This is controlled through a registry setting. Add the following REG_SZ (String) key if it does not exist. The value may be *True* or *False*.

HKLM\Software\Greyware\Domain Time Client\Parameters\SuppressShortcuts

If you are editing the installation template, locate the Parameters section and add a new entry if it doesn't already exist:

"SuppressShortcuts"="True"

This will prevent the shortcuts from ever being created. If you want to remove the shortcuts from a previous installation, edit the registry and add (or edit) the REG_SZ (String) value named SuppressShortcuts and set its value to *True*. When the service restarts, it will remove the shortcuts.

Target Machine:

You may optionally specify a remote target machine (machine where the install, removal, or upgrade is to take place).

Targets are specified by the Windows Networking computer name, also known as the NetBIOS name. You may also use the target machine's IP address or fully-qualified DNS name instead of the NetBIOS name, as long as your WINS and DNS subsystems are functioning correctly.

You may combine any directive and option with the target machine. Below are some examples:

dtclient.exe -install -noack \\fred - installs the service, showing the progress bar but not any confirmation questions or success messages, to the machine named \\fred

dtclient.exe -remove \\barney - removes the service from the machine \\barney and displays both a progress bar while working and a confirmation message when the removal has completed successfully.

dtclient.exe -upgrade -quiet \\172.16.240.1 - installs the program, upgrading if necessary, to the machine with IP address 172.16.240.1, without confirmations or progress bar.

Note: For remote installation or removal using the command-line to work, both the machine you are working on and the target machine must be running on the same hardware platform (32-bit or 64-bit), and you must be logged on under an account that has administrative privileges on the target machine.

Command-line Options

Most Domain Time II components can be installed or removed from the command-line, including on remote machines. If you prefer to write your own batch files, or want to install/remove individual components manually, you can use the command-line parameters.

Select Program

Pick the program you're interested in from the drop-down list, then click the "Select" button to see the command-line parameters for that program:

Domain Time II Server

Server Command-line Options:

`dtserver.exe [directive] [options] [\\targetmachine]`

If no command-line parameters are provided, the program will display a dialog box that lets you choose whether to install, upgrade, or remove the program.

Command line parameters can be provided with either a leading dash or leading forward slash. Slashes are shown below. Parameters can usually be abbreviated to the first letter of the command. For example, you can use `/install`, `/i`, `-install`, or `-i` to specify the install directive.

Directives:

Specify one of these directives to install, remove, upgrade, or display version information:

`dtserver.exe /version` - displays version information
`dtserver.exe /install` - installs the program
`dtserver.exe /remove` - removes the program
`dtserver.exe /upgrade` - installs the program, upgrading if necessary

Options:

The following options may be added to any directive:

Example: `dtserver.exe -install -nocpl -vq`

Installs the program without showing the control panel applet or any confirmation messages.

`-noack`

Suppresses confirmation questions and successful completion notices (but not error messages).

`-nocpl`

Does not start the control panel applet after installation.

`-quiet` or `-q`

Suppresses confirmation questions, successful completion notices, and the progress bar. `-quiet` implies `-noack`. Like `-noack`, `-quiet` will not suppress error messages.

`-veryquiet` or `-vq`

Suppresses all messages during installation or removal.

`-reset` or `-re`

Reloads the DTSERVER.REG default settings template file (overwriting any existing values). This option may only be used with the `/upgrade` directive.

Settings Templates:

The Domain Time Client and Server installation routines use [template files](#) to configure most of the default settings of the component. These files are simply modified registry settings files (.reg) in REGEDIT4 format. You may use the included default settings .reg files or, as of version 5.2.b.20180303, you may specify custom files of your own.

The .reg files containing the manufacturer's default settings for Server and Client are located in the installation folders (i.e. `\i386` or `\amd64`) for each component (**DTSERVER.REG** for Server, **DTCLIENT.REG** for Client). If you decide to modify these files to change the overall defaults, be sure to make a backup copy so you can revert if necessary.

To specify a custom template, copy your custom .reg file to the installation folders (both /i386 and /amd64) and add this directive to your command line:

-template=TemplateFilename.reg

IMPORTANT: A number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. You should not simply export the **HKLM\SOFTWARE\Greyware\Domain Time Server** key and apply it to other systems. Files exported using Windows Registry Editor in registry formats other than REGEDIT4 may not even import correctly. Rather, use the [Import/Export](#) function to export a .reg template. Alternately, you may use the PrepClone feature of the DTCHECK utility to prepare a machine before exporting the key (see this [KB article](#)).

Suppress Shortcuts

You may opt to suppress creation of the Domain Time shortcuts in the All-Users startup menu. This is controlled through a registry setting. Add the following REG_SZ (String) key if it does not exist. The value may be *True* or *False*.

HKLM\Software\Greyware\Domain Time Server\Parameters\SuppressShortcuts

If you are editing the installation template, locate the Parameters section and add a new entry if it doesn't already exist:

"SuppressShortcuts"="True"

This will prevent the shortcuts from ever being created. If you want to remove the shortcuts from a previous installation, edit the registry and add (or edit) the REG_SZ (String) value named SuppressShortcuts and set its value to *True*. When the service restarts, it will remove the shortcuts.

Target Machine:

You may optionally specify a remote target machine (machine where the install, removal, or upgrade is to take place).

Targets are specified by the Windows Networking computer name, also known as the NetBIOS name. You may also use the target machine's IP address or fully-qualified DNS name instead of the NetBIOS name, as long as your WINS and DNS subsystems are functioning correctly.

You may combine any directive and option with the target machine. Below are some examples:

dtserver.exe -install -noack \\fred - installs the service, showing the progress bar but not any confirmation questions or success messages, to the machine named \\fred

dtserver.exe -remove \\barney - removes the service from the machine \\barney and displays both a progress bar while working and a confirmation message when the removal has completed successfully.

dtserver.exe -upgrade -quiet \\172.16.240.1 - installs the program, upgrading if necessary, to the machine with IP address 172.16.240.1, without confirmations or progress bar.

Note: For remote installation or removal using the command-line to work, both the machine you are working on and the target machine must be running on the same hardware platform (32-bit or 64-bit), and you must be logged on under an account that has administrative privileges on the target machine.

Command-line Options

Most Domain Time II components can be installed or removed from the command-line, including on remote machines. If you prefer to write your own batch files, or want to install/remove individual components manually, you can use the command-line parameters.

Select Program

Pick the program you're interested in from the drop-down list, then click the "Select" button to see the command-line parameters for that program:

Domain Time II Client

Monitor Service Command-line Options:

`dtmonitor.exe [directive] [options] [\\targetmachine]`

If no command-line parameters are provided, the program will display a dialog box that lets you choose whether to install, upgrade, or remove the program.

Command line parameters can be provided with either a leading dash or leading forward slash. Slashes are shown below. Parameters can usually be abbreviated to the first letter of the command. For example, you can use `/install`, `/i`, `-install`, or `-i` to specify the install directive.

Directives:

Specify one of these directives to install, remove, upgrade, or display version information:

`dtmonitor.exe /version` - displays version information
`dtmonitor.exe /install` - installs the program
`dtmonitor.exe /remove` - removes the program
`dtmonitor.exe /upgrade` - installs the program, upgrading if necessary

Options:

The following options may be added to any directive:

Example: `dtmonitor.exe -install -nocpl -vq`
Installs the program without showing the control panel applet or any confirmation messages.

-noack

Suppresses confirmation questions and successful completion notices (but not error messages).

-nocpl

Does not start the control panel applet after installation.

-quiet or -q

Suppresses confirmation questions, successful completion notices, and the progress bar. -quiet implies -noack. Like -noack, -quiet will not suppress error messages.

-veryquiet or -vq

Suppresses all messages during installation or removal.

Target Machine:

You may optionally specify a remote target machine (machine where the install, removal, or upgrade is to take place).

Targets are specified by the Windows Networking computer name, also known as the NetBIOS name. You may also

use the target machine's IP address or fully-qualified DNS name instead of the NetBIOS name, as long as your WINS and DNS subsystems are functioning correctly.

You may combine any directive and option with the target machine. Below are some examples:

`dtmonitor.exe -install -noack \\fred` - installs the service, showing the progress bar but not any confirmation questions or success messages, to the machine named \\fred

`dtmonitor.exe -remove \\barney` - removes the service from the machine \\barney and displays both a progress bar while working and a confirmation message when the removal has completed successfully.

`dtmonitor.exe -upgrade -quiet \\172.16.240.1` - installs the program, upgrading if necessary, to the machine with IP address 172.16.240.1, without confirmations or progress bar.

Note: For remote installation or removal using the command-line to work, both the machine you are working on and the target machine must be running on the same hardware platform (32-bit or 64-bit), and you must be logged on under an account that has administrative privileges on the target machine.

Command-line Options

Most Domain Time II components can be installed or removed from the command-line, including on remote machines. If you prefer to write your own batch files, or want to install/remove individual components manually, you can use the command-line parameters.

Select Program

Pick the program you're interested in from the drop-down list, then click the "Select" button to see the command-line parameters for that program:

Domain Time II Update Server

Update Server Command-line Options:

dtupdate.exe [*directive*] [*options*]

If no command-line parameters are provided, the program will display a dialog box that lets you choose whether to install, upgrade, or remove the program.

Command line parameters can be provided with either a leading dash or leading forward slash. Slashes are shown below. Parameters can usually be abbreviated to the first letter of the command. For example, you can use **/install**, **/i**, **-install**, or **-i** to specify the install directive.

Directives:

Specify one of these directives to install, remove, upgrade, or display version information:

dtupdate.exe /version - displays version information
dtupdate.exe /install - installs the program
dtupdate.exe /remove - removes the program
dtupdate.exe /upgrade - installs the program, upgrading if necessary

Options:

The following options may be added to any directive:

Example: **dtupdate.exe -install -nocpl -vq**
Installs the program without showing the control panel applet or any confirmation messages.

-noack
Suppresses confirmation questions and successful completion notices (but not error messages).

-nocpl
Does not start the control panel applet after installation.

-quiet or **-q**
Suppresses confirmation questions, successful completion notices, and the progress bar. -quiet implies -noack. Like -noack, -quiet will not suppress error messages.

-veryquiet or **-vq**
Suppresses all messages during installation or removal.

Command-line Options

Most Domain Time II components can be installed or removed from the command-line, including on remote machines. If you prefer to write your own batch files, or want to install/remove individual components manually, you can use the command-line parameters.

Select Program

Pick the program you're interested in from the drop-down list, then click the "Select" button to see the command-line parameters for that program:

Domain Time II Client

Audit Server Command-line Options:

`dtaudit.exe [directive] [options] [\\targetmachine]`

If no command-line parameters are provided, the program will display a dialog box that lets you choose whether to install, upgrade, or remove the program.

Command line parameters can be provided with either a leading dash or leading forward slash. Slashes are shown below. Parameters can usually be abbreviated to the first letter of the command. For example, you can use `/install`, `/i`, `-install`, or `-i` to specify the install directive.

Directives:

Specify one of these directives to install, remove, upgrade, or display version information:

`dtaudit.exe /version` - displays version information

`dtaudit.exe /install` - installs the program

`dtaudit.exe /remove` - removes the program

`dtaudit.exe /upgrade` - installs the program, upgrading if necessary

Options:

The following options may be added to any directive:

Example: `dtaudit.exe -install -nocpl -vq`

Installs the program without showing the control panel applet or any confirmation messages.

`-noack`

Suppresses confirmation questions and successful completion notices (but not error messages).

`-nocpl`

Does not start the control panel applet after installation.

`-quiet` or `-q`

Suppresses confirmation questions, successful completion notices, and the progress bar. `-quiet` implies `-noack`. Like `-noack`, `-quiet` will not suppress error messages.

`-veryquiet` or `-vq`

Suppresses all messages during installation or removal.

Target Machine:

You may optionally specify a remote target machine (machine where the install, removal, or upgrade is to take place).

Targets are specified by the Windows Networking computer name, also known as the NetBIOS name. You may also

use the target machine's IP address or fully-qualified DNS name instead of the NetBIOS name, as long as your WINS and DNS subsystems are functioning correctly.

You may combine any directive and option with the target machine. Below are some examples:

dtaudit.exe -install -noack \\fred - installs the service, showing the progress bar but not any confirmation questions or success messages, to the machine named \\fred

dtaudit.exe -remove \\barney - removes the service from the machine \\barney and displays both a progress bar while working and a confirmation message when the removal has completed successfully.

dtaudit.exe -upgrade -quiet \\172.16.240.1 - installs the program, upgrading if necessary, to the machine with IP address 172.16.240.1, without confirmations or progress bar.

Note: For remote installation or removal using the command-line to work, both the machine you are working on and the target machine must be running on the same hardware platform (32-bit or 64-bit), and you must be logged on under an account that has administrative privileges on the target machine.


Network Rollout

You can remotely install, upgrade, configure, or remove Domain Time II components to other machines on your network using Domain Time II [Manager](#). There's no need to visit individual machines to install or configure them!

You may also easily integrate Domain Time into your existing package distribution systems or cloned-install routines. Domain Time Server and Client have a built-in export utility to create standard Windows Registry files that allow you to preset your systems with all of your desired configuration options, while omitting any settings specific to individual machines.

- For setup instructions for use in installer packages, see the [Command-line Options](#) page.
- For cloned-machine installation, see this [KB article](#).

Follow these steps to rollout Domain Time to your entire network:



- Carefully review the [System Requirements](#) and [Planning](#) pages to be sure you understand the basic requirements for using this version of Domain Time.
- Decide on a [Recommended Configuration](#) and adapt it to your environment to create a detailed installation plan. You should follow your organization's standard quality-control procedures for testing and validation of new software (i.e. using separate test, staging, and production networks).
- Verify your environment meets the minimum requirements for performing remote operations using Domain Time components. In order to be able to install, upgrade, or configure remote machines:
 - Your network must be a correctly-configured Windows network, i.e. configured with working name resolution (DNS, WINS, NetBIOS, etc.), correct and functioning Active Directory (if used), working inter-domain trusts, etc.
 - Your network must pass both UDP and TCP network traffic sent to destination port 9909. Switches and firewalls must pass this traffic bi-directionally, since traffic will originate either from Manager or the remote machines. Your network must pass this traffic, regardless of what time protocols are used to actually synchronize the time.

Note: As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

- The remote machine must respond to PING requests from the connecting machine.
- The connecting Domain Time program, utility, or service must be run using credentials with sufficient privileges to connect to and write files to the administrative shares on the remote machine using Microsoft Networking (Domain Admin if the target is a domain member, Local Machine Administrator if the target is in a workgroup).
- The Remote Registry Service must be running on the remote systems and its registry keys must be accessible to the connecting program.
- All files from the original distribution for each type of product you want to install (Server, Client, etc.) must be extracted and present on your connecting machine. Setup copies these to the proper locations in the `\Program Files\Domain Time II` folder for you automatically when you install the Management Tools.
- Use the [Setup](#) program to install [Manager](#) (Management Tools) on the machine that you want to use as your Management Workstation.

- Since many functions of Domain Time II Manager depend on accurate time calculations, it should always be run on a physical (not virtual) machine.
- If you will be using [Audit Server](#), you should install both Manager and Domain Time II [Server](#) on the same machine you intend to use for Audit Server.
- If desired, create template .reg files containing your desired pre-set Server and/or Client settings to use during installation or upgrades.

Template .reg files are very handy for multiple-machine installations or ensuring that repeated single-machine installations will be identical. They are also ideal for use with software installation packaging tools (like creating .MSI packages) or creating cloned images.

You do not have to create template files for simple rollouts or single-machine installations if you would prefer not to. Simply use Setup or Manager to install the Domain Time Server or Client with the default settings and change the configuration as needed afterwards.

Here's how to create custom template .reg files:

- Use Manager to install Domain Time II Server and/or Domain Time II Client on machines you want to use as exemplar systems for creating your template(s). Exemplar machines should be comparable to systems you will actually use in production, if possible.
- Use Manager to connect remotely to the Domain Time Control Panel applet on each exemplar machine. Configure and test them with the exact settings you want your remotely-installed systems to use. Be sure to verify all of the settings carefully.
- Use the exemplar Server and/or Client's **Save Settings to File** utility (found on the [Import/Export](#) property page of the Control Panel applet) to save the template .reg file into the **C:\Program Files\Domain Time II\Templates\[Server][Client]** folder of the Domain Time II Manager machine. Template .reg files located in those folders will automatically be made available for use when installing or upgrading using Manager.

Note: If you launch the Control Panel applet of your exemplar machine using Manager, the export will automatically offer to save the .reg file in the proper directory on the Manager machine. Otherwise, you will have to manually save or copy the file to the Manager machine.

The .reg files containing the manufacturer's default settings for Server and Client are located in the installation folders (i.e. **\i386** or **\amd64**) for each component (**DTSERVER.REG** for Server, **DTCLIENT.REG** for Client). You should not change these under most circumstances (create new custom template .reg files instead), however, if you must change the overall defaults, be sure to make a backup copy so you can revert if necessary.

IMPORTANT: Although .reg files created using this utility are saved in standard Windows registry file format, it is **not** equivalent to exporting the registry keys using Windows' RegEdit program. A number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. Those settings are automatically excluded when you export using this utility, so you should always use this utility to create a Domain Time .reg file.

- If you will be using [Active Directory policies](#) and/or [DHCP Server Options](#) to globally set Domain Time settings on your network, configure them on your domains, then test the results of those settings on your exemplar machines and a representative sample of other systems on remote subnets before rolling out Domain Time on a broad scale.
- Use Manager to complete the detailed Installation Plan for your network, rolling out first to Domain Controllers, other Servers, and then Clients. See Manager's [instructions pages](#) on how to perform batch installation or upgrades on all of the various machines on your network.

How to use Active Directory Group Policies to specify Domain Time II settings.

As of version 5.1, certain settings of Domain Time Servers and Clients can be enforced using Windows Group Policies. This allows an administrator to fully integrate Domain Time into the existing Active Directory structure.

Policies give enormous flexibility and a very fine degree of control over which settings machines should use, even when managing a large number of systems. Using policies, you can, for example, specify which time servers the workstations in each of your remote offices will use to obtain their time, while simultaneously specifying different settings for all of your domain controllers.

Domain Time group policies work the same way as other Group Policies, so your Windows administrators will already be familiar with how to use them.

Installing the Domain Time Administrative Template file

In order to use Group Policies to configure Domain Time, you must install the **DOMTIME.ADM** administrative template file into your desired Group Policy object.

The **DOMTIME.ADM** template file is located in the Domain Time distribution files (and also in the **\Program Files\Domain Time II** folder, if you have installed Domain Time II **Manager**).


The method for installing Administrative Templates varies slightly among different versions of the operating system.

Note: A full discussion of Group Policies is outside the scope of this document. Always follow Microsoft's recommendations for installing and using custom Administrative Templates with Group Policies.

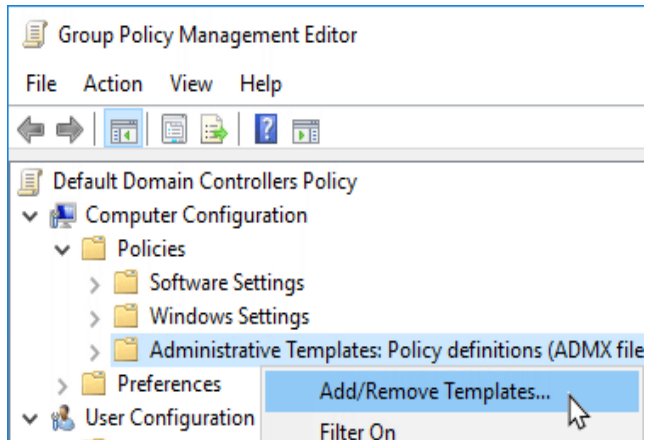
You can use the *Group Policy Management Console* (GPMC) to load the template. GPMC is included with newer versions of Windows, but it is also [available for download](#) for XP/2003 systems. If your system does not have GPMC, you may be able to use the Group Policy Editor instead.

Here's an example of the template installation process on Windows Server 2008:

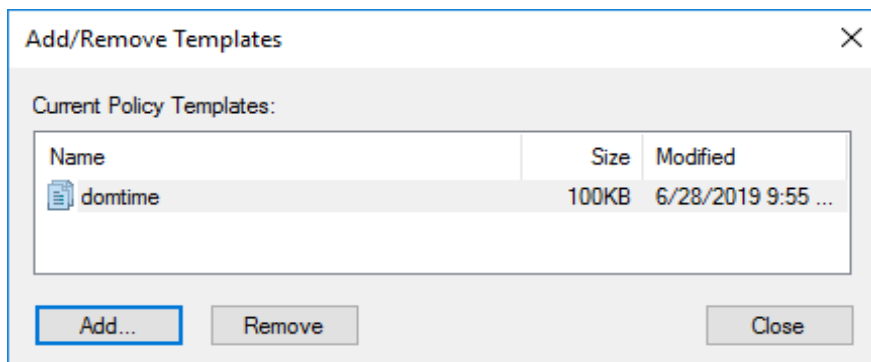
- Copy the **DOMTIME.ADM** file into the **C:\Windows\Inf** folder on a Domain Controller (the DC holding the PDC-Emulator role is preferred).
- Open the *Group Policy Management Console*, expand the tree to the **Group Policy Objects** section and locate (or create) the Group Policy object you want to use to apply Domain Time settings. Right-click the object and select **Edit** from the context menu to bring up the Policy Editor.



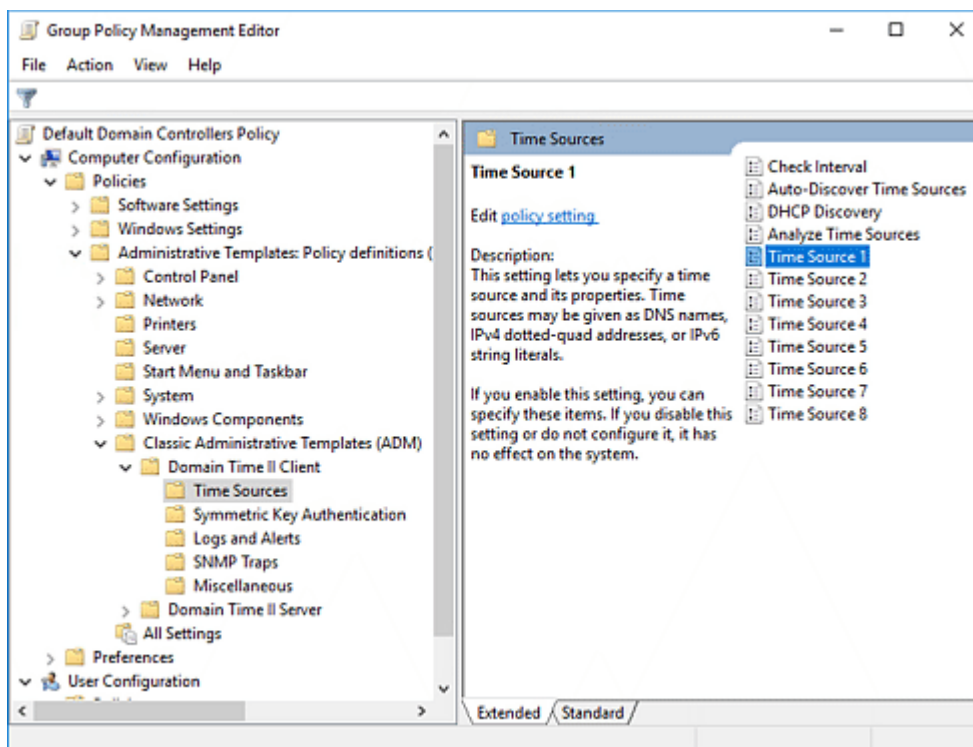
- Expand the tree to find the **Computer Configuration\Policies** section, highlight **Administrative Templates** and right-click it to select **Add/Remove Templates...** from the context menu.




- Click **Add** to browse for the **DOMTIME.ADM** file. Select it and close the **Add/Remove Templates** dialog box.



- You'll now see the Domain Time categories listed in your Administrative Templates (on newer operating systems, they may be listed in the **Classic Administrative Templates (ADM) section**). Select the Domain Time policy you want to configure. You may configure settings for Domain Time Client, Server, or both.



Notes:

- The Domain Time policy templates contain a subset of commonly-used settings on Client or Server. If you need to preset settings that are not included in the template, you should use the Import/Export function of Client and/or Server to create an installation template .reg file and apply it during [network rollout/upgrade](#) of your machines.
- Group Policy settings will override any settings made on the local machine.
- Once a policy has been applied to the local machine, Domain Time will start using it the next time it restarts or synchronizes its time.
- The  **Group Policy applied** indicator will appear in the lower-left corner of the Domain Time Control Panel applet when there are settings on the page that are being overridden by a Group Policy. The settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change to the setting.


DHCP Server Options

How to use DHCP Servers to specify time server addresses to Domain Time Client


Domain Time II Client using automatic discovery mode can be configured to check for the presence of a DHCP server on the local subnet. If a DHCP server is found, the DHCP options 004 and/or 042 can be examined to provide the IP address(es) of time servers for the Client to use.

Note: It is *not* necessary to set the TCP/IP protocol stack to get its IP address from a DHCP server in order for the Domain Time II Client to get a time server address. The Client uses its own independent inquiry of the DHCP server to discover the time server options. Therefore DHCP discovery of time servers can be used on a machine with either a static or a DHCP-assigned IP address.


An automatic Domain Time client with DHCP enabled will broadcast to locate a DHCP Server. Note that DHCP broadcasts usually do not cross routers.



If DHCP options 004 or 042 are configured, the DHCP server will respond with the IP address of the time server.



The client then uses the IP address provided by the DHCP server to contact and sync with the designated time server, even across a router.



The DHCP Time Server Options and Discovery Order

These are the DHCP options that can be set to provide time server addresses. If enabled, they are examined in the order listed:

- **Option 004** ("Time Servers") can contain a list of servers supplying the DT2 protocol. If a server is listed in option 004 that doesn't support DT2 UDP, it will be ignored.

Note: That this is a change in behavior from Domain Time v4.1, where Option 004 was used to specify servers providing the TIME/ITP protocol. TIME/ITP service has been deprecated in v5.1 and later Clients.

- **Option 042** ("NTP Servers") can contain a list of servers supplying either the NTP or DT2 protocol. If a server is listed in option 042, it will be checked for NTP first. If NTP fails, it will be checked for DT2 UDP. If it does not provide time under either of these two protocols, it will be ignored.

Note: That this is a change in behavior from Domain Time v4.1, where Option 042 was evaluated for the DT2 protocol before NTP. The new behavior more accurately reflects the purpose of option 042 (to specify NTP servers) and also allows administrators to set Clients to prefer NTP. You should now use option 004 if you want to prefer the DT2 protocol.

Use the [Client Discovery Options](#) section of the **Obtain the Time** property page on the Client Control Panel Applet to select which DHCP options are enabled.

Auto-Manage Windows Firewall Settings

How to use Domain Time's auto-management of Windows Firewall settings.

► [Be sure to read the Prepare your network to pass the necessary traffic section of the Planning page.](#)

As of Version 5.2.b.20150828, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. In prior versions, the administrator either needed to use the DTCheck utility to open the necessary ports, or needed to manage the firewall manually.

Only the built-in Windows Firewall is supported. If you are using a third-party firewall, you must open the needed ports yourself and should not use Domain Time's automatic firewall management.

Note: The Windows Firewall can be controlled by Group Policies. Changes made by Domain Time to a firewall managed by Group Policies will either fail, or only last until the next time Group Policies are applied, depending on the operating system and your settings.

Necessary Ports

Various services and time protocols require certain ports to be open to operate correctly.

- **Server** needs to have its listening ports opened to allow clients (either Domain Time or other clients) to obtain the time, and for monitoring and auditing purposes. Server ports include 9909 UDP and 9909 TCP for the DT2 protocol, 123 UDP for the NTP protocol, a TCP port (normally 80) for DT2-HTTP, 13 UDP and 13 TCP for the Daytime protocol, and 37 UDP and 37 TCP for the TIME/ITP protocol.
- **Client** needs to have ports 9909 UDP and 9909 TCP opened for monitoring and auditing purposes, and 123 UDP open for NTP broadcast reception or ntpq-style querying.
- Both **Server** and **Client** need ports 319 UDP and 320 UDP opened if they are deriving time using IEEE-1588 Precision Time Protocol (PTP) as a time source. Additionally, if the Service Status Monitor is being used, Server and Client need ports 9911 UDP and 9911 TCP opened for remote access to Service Status Monitor.
- **Audit Server** needs to have port 9910 TCP opened for Real-Time Alert sharing (the DTAAlert program), and for communication between an active Audit Server and a stand-by Audit Server.

Only three ports used by Domain Time are user-configurable; the others are governed by RFCs or convention, and cannot be changed. The configurable ports are:

- *Real-Time Alert Sharing*, used by Audit Server, normally 9910 (TCP only)
- *Service Status Monitor*, used by Client and Server, normally 9911 (UDP, TCP, or both)
- *DT2-HTTP*, used by Server to serve DT2-HTTP, normally 80 (TCP only)

If automatic management of the Windows Firewall is enabled, and you change any of these ports (or are already using non-standard ports), then Domain Time will make sure the firewall rule matches the port you are using. The change in the firewall happens when you make the change in Domain Time. If you are not letting Domain Time manage the Windows Firewall, you will have to adjust the firewall rules yourself. Note that this dynamic management will only create rules if they don't already exist, or update the port numbers if the rules do exist. Dynamic management will not change any other aspects of the rules.

Firewall Rules and Profiles

The Windows Firewall service must be running (but need not be enabled) in order for Domain Time to check or change rules. If automatic management of the Windows Firewall is enabled, Domain Time will check at service startup to ensure the needed rules are present. If the rules already exist, Domain Time will not change them (other than to correct port numbers if

wrong).

When a Domain Time service creates a firewall rule on XP or 2003, it will enable it for both the Standard and Domain profiles, regardless of which one is currently in use.

When a Domain Time service creates a firewall rule on Vista or above, it will enable it for the Private profile, and, if the machine is a domain member, for the Domain profile as well.

If a rule already exists, Domain Time services do not change that rule's profile, enabled/disable status, or other rule settings, such as interface restrictions. This allows you to let Domain Time create the rules for you, then fine-tune them yourself without further interference from Domain Time. For example, you could disable the rules for the ports you don't need or want shared, or change the profile to which the rule applies, or add IP range restrictions, etc.

Note: Domain Time will never change ICMP echo (ping) firewall settings. Ping is required for some operations that use TCP, or for Manager to install, upgrade, remove, or control remote machines. Ping is also required for Client or Server to send Real-Time Alerts (Status Reports) to Audit Server when using TCP. If your firewall blocks ICMP echo, you must allow it if you desire these functions.

Domain Time services record all firewall change activity in a text file called DTFirewall.log in the System32 folder.

Firewall Rule Names

These are the names used by Domain Time II for its firewall rules. If you have used DTCheck from an earlier version to open the firewall, the old names will be replaced by the names below:

Rule Name	Component	Default Port/Address
Domain Time II DT2-UDP	Client or Server	always 9909 UDP
Domain Time II DT2-TCP	Client or Server	always 9909 TCP
Domain Time II PTPv2-Event	Client or Server	always 319 UDP
Domain Time II PTPv2-General	Client or Server	always 320 UDP
Domain Time II NTP	Client or Server	always 123 UDP
Domain Time II Daytime-UDP	Server only	always 13 UDP
Domain Time II Daytime-TCP	Server only	always 13 TCP
Domain Time II TIME/ITP-UDP	Server only	always 37 UDP
Domain Time II TIME/ITP-TCP	Server only	always 37 TCP
Domain Time II Status Monitor-UDP	Client or Server	defaults to 9911 UDP
Domain Time II Status Monitor-TCP	Client or Server	defaults to 9911 TCP
Domain Time II Real-Time Alert Sharing	Audit Server only	defaults to 9910 TCP
Domain Time II DT2-HTTP	Server only	defaults to 80 TCP

Firewall Settings on Specific Domain Time II Components

The Setup utility

You may use the Setup utility to enable or disable the Auto-Manage Windows Firewall function on components when you install or upgrade them.

The **Auto-Manage Windows Firewall** checkbox defaults to checked. This is a tri-state checkbox (checked, unchecked, indeterminate). Click the box multiple times to cycle through the options.

When installing or upgrading Client or Server:

- If checked, Client or Server will automatically manage the Windows Firewall, regardless of prior settings.
- If unchecked, Client or Server will not automatically manage the Windows Firewall, regardless of prior settings.
- If indeterminate, Client or Server's firewall management will not be changed. Note that on new installs, Client and Server default to not managing the Windows Firewall.

When installing or upgrading the Management Tools:

- If checked, Manager's default setting for installs, upgrades, or reset configurations will be to "Force Auto-Manage Windows Firewall" on its *Remote Computer Operation* dialog. In addition, if Audit Server is installed, the "Auto-Manage Windows Firewall" setting for Alert Sharing will be enabled.
- If unchecked, the default *Remote Computer Operation* setting in Manager will not be set to "Force Auto-manage Windows Firewall". The Audit Server firewall management setting for Alert Sharing will also be disabled.
- If indeterminate, Manager's default setting will be unchanged. If you are upgrading from a previous version of Manager that does not have firewall management settings, then Manager (and Audit Server, if installed) will default to checked.

Domain Time II Client and Server

Client and Server both have the ability to auto-manage the Windows Firewall for their respective necessary ports.

The **Auto-Manage Windows Firewall** checkbox appears on the *Network\Security* property page of the applet. This is a binary checkbox (checked or unchecked). The default is unchecked, unless you have used a customized template that specifies otherwise, or have pushed the installation from Manager and specified that you want firewall management enabled.

- If checked, Client or Server will automatically manage the Windows Firewall.
- If unchecked, Client or Server will not touch the Windows Firewall.

Domain Time II Manager

Domain Time II Manager has the ability to enable or disable the Auto-Manage Windows Firewall function on Clients or Servers it installs, upgrades, or resets configuration on remotely.

The **Force Auto-Manage Windows Firewall** checkbox appears on the *Remote Computer Operation* dialog that is used for installing, upgrading, or resetting the configuration on other machines. This is a tri-state checkbox (checked, unchecked, indeterminate). Click the box multiple times to cycle through the options.

When using Manager to install, upgrade, or reset configuration on Client or Server:

- If checked, Client or Server will automatically manage the Windows Firewall, regardless of prior settings and template entries.
- If unchecked, Client or Server will not automatically manage the Windows Firewall, regardless of prior settings and

template entries.

- If indeterminate, Client or Server's firewall management settings will be affected as follows:
 - On Installs: Client and Server default to not managing the Windows Firewall unless you are using a custom template that enables it.
 - During Upgrades: Manager will either preserve the existing Client or Server settings or override them with the settings from the selected template (as configured on the [Choose Templates](#) dialog).
 - Reset configuration: Client or Server will either be reset to their default configuration (not auto-managed) or use the settings specified in the selected template (as chosen on the *Remote Computer Operation* dialog).

NOTE: Custom templates you create by exporting settings from Client or Server will contain that machine's setting for management of the Windows Firewall. The template setting is called "Auto-Manage Firewall" and can either be **True** or **False** (string). If present in the selected install/upgrade/reset template, this setting will be applied unless you tell Manager to override the setting. The default templates shipped with Domain Time do not contain the "Auto-Manage Firewall" setting; only templates you create yourself might have it.

Domain Time II [Audit Server](#)

Audit Server, if installed, has the ability to auto-manage firewall settings for its Alert Sharing and Standby-mode replication features.

The **Auto-Manage Windows Firewall** checkbox for Alert Sharing is located on the *Advanced Real-Time Alert Configuration* dialog (found via the Manager menu - choose Audit Server -> Alerts -> Configure, then click the Advanced button). This is a binary checkbox (checked or unchecked), and controls only whether or not Audit Server should automatically manage the Windows Firewall for the Real-Time Alert sharing port (default 9910 TCP).

The [DTCheck](#) Utility

DTCheck is a command-line utility program that ships with Domain Time. By default, it is installed in the System32 folder when you install Client or Server, and is in the Manager folder when you install the management tools.

DTCheck has two commands that affect the Windows Firewall:

■ **dtcheck -firewall:close**

Deletes all Domain Time II firewall rules. There are no optional parameters.

■ **dtcheck -firewall:open [optional parameter]**

Creates rules that if they don't exist, enables existing rules that aren't enabled, and sets the profile(s) used.

Optional parameters for -firewall:open are **-public** **-private** **-domain** and **-standard**. (Only -domain and -standard are allowed for XP/2003; only -public -private and -domain are allowed for newer operating systems.) The default, if you specify no parameters is -domain and -standard for XP/2003, and -private and -domain (if the machine is a domain member) for newer operating systems.

Note that unlike the Domain Time services described above, DTCheck's -firewall:open forces the rules to be present, enabled, and set to the profile(s) you choose. DTCheck will also force the port numbers to match the protocols in use, including the user-configurable port numbers mentioned above.

DTCheck's -firewall:open detects whether you have Client or Server installed, and opens only the ports needed for Client or Server. In addition, DTCheck detects if Audit Server is installed, and opens the Real-Time Alert Sharing port if needed.

Unlike firewall management by services, DTCheck writes its results to the command window instead of to the System32\DTFirewall.log file.



Information and configuration instructions for using the Precision Time Protocol (IEEE 1588-v2) with Domain Time II.

- ▶ [Introduction](#)
- ▶ [Optimum Environment for PTP](#)
- ▶ [PTP Profiles](#)
- ▶ [Configure Domain Time for Windows as a PTP Slave](#)
- ▶ [Configure Domain Time Client for Linux \(DTLinux\) as a PTP Slave](#)
- ▶ [Configure Domain Time Server as a PTP Master](#)
- ▶ [PTP Monitor/Audit Server](#)
- ▶ [PTPCheck](#)

Introduction

PTP, like other time protocols, is used to determine the current time offset and discipline the local clock. PTP typically achieves a much higher degree of accuracy than other types of time synchronization. In most cases, the accuracy is several orders of magnitude better than NTP or DT2 alone.

- Domain Time II supports Precision Time Protocol version 2.0 (IEEE 1588-2008) and version 2.1 (IEEE 1588-2019). PTPv1 is not supported.
- Domain Time Servers and Clients can be configured to run in slave-only PTP mode using the Default (multicast or hybrid), Enterprise, or (as of v5.2.b.20180606) the Telecom PTP profile (see [PTP Profiles](#) below).
- Domain Time II Server can assume the master clock role using either the Default or Enterprise (Hybrid) [PTP Profiles](#). As of v5.2.b.20180801, Domain Time Server can also act as a Telecom master. This option should only be enabled if you have no hardware-based PTP grandmaster, or if you need a software-based backup to your grandmaster. PTP is designed to work with hardware-based clocks. Software-based clocks cannot offer the same degree of precision or accuracy as hardware-based clocks.
- Using the Default or Enterprise profiles, Domain Time attempts to automatically detect and synchronizes to the best master on the local segment. When using the Telecom profile, you must manually select the Telecom profile and configure a list of acceptable masters.
- PTP has two delay-measurement mechanisms: End-to-End and Peer-to-Peer. Most master clocks can support either, but only one at a time. By default, Domain Time auto-detects the delay mechanism and uses whichever one the master supports.
- Domain Time is compatible with masters providing both one-step Syncs and two-step Syncs (Sync + Followup). When running as a Master, Domain Time Server defaults to acting as a one-step clock, but you may change it to two-step.
- PTP is a “chatty” protocol. The master clock sends announce, sync, and other messages on a regular basis, and slaves send queries and responses of their own. The default sync frequency is one per second, but most master clocks allow this number to be adjusted, and administrators may choose a higher number.

PTP Precision and Timestamping

A PTP network of hardware clocks on a well-regulated segment can usually synchronize to within a few tens or hundreds of nanoseconds. Software implementations can theoretically achieve the same results, but are subject to limitations of the host operating system. In particular, Windows and Linux are not real-time operating systems, and the connecting networks are often congested or have unpredictable or asymmetric performance characteristics. PTP implementations in these environments must estimate delays caused by servicing interrupts, stack traversal, latencies introduced by expiring time slice quanta, and vagaries in the network itself.

By default, when an application tells the operating system to place a packet on the wire, it can know only when the packet left the application, not when the packet finished traversing the operating system layers and was actually placed on the wire. Likewise, incoming packets may suffer delays in the network driver or the operating system before being delivered to the application. This degree of uncertainty produces “jitter,” or small variations in the measured time that are artifacts of the environment rather than the protocol. In addition, any procedure can be preempted, leading to erroneous measurements of elapsed time.

Some operating systems, with assistance from the network driver, can provide a hardware timestamp indicating when a particular packet

Theoretically, the higher the rate of messages, the more accurately the slave can measure the master's information, but this is primarily true only for hardware appliances. Domain Time is heavily optimized to perform best at one sync packet per second, so increasing the sync frequency can actually lead to worse performance.

- A standard-compliant PTP node using the Default Profile operates using only multicast messages. This means that not only does every node see the master's messages, but every node sees every request made by every slave, and every reply. Domain Time II tries to reduce this traffic by sending unicast delay requests (Hybrid Mode). The Enterprise Profile also uses unicast delay requests. If the master clock supports multicast with unicast replies, Domain Time II will continue using only unicast for delay measurement with that clock. If the master does not support unicast for delay requests, Domain Time II will fall back to using multicast. Domain Time can also operate as a Telecom profile slave (as of v5.2.b.20180606) or master (as of v5.2.b.20180801). Telecom masters and slaves communicate solely over unicast.
- When PTP is enabled on Domain Time, it runs continuously in the background, exchanging messages with the master and collecting statistics. You can see this activity using the logs and graphs from the IEEE 1588 *Status* and *Graph* links on the [Obtain the Time](#) property page of the Windows control panel applet, or the [Graphs and Statistics](#) menu of the DTLinux control panel app in Domain Time II Manager.
- Domain Time does not support the experimental PTP authentication mechanisms described in Annex K of 1588-2008. Domain Time supports PTP 1588-2019 v2.1 authentication as of version 5.2.b.20190331. See [KB article 2019.331](#) for more information.

Optimum Environment for PTP

The hardware, software, and network environment is critical for PTP synchronization at the sub-millisecond level. This section will help customers hoping to achieve sub-millisecond synchronization across your network of machines:

■ CPU

The best processors for time synchronization are Intel's Core i7 line (or later) or Xeon E7 line (or later). Earlier chips are not as stable or as precise as the newer models. The newer processors also have an invariant timestamp counter, which allows Domain Time II to measure the passage of time accurately regardless of SpeedStep or other power-saving mechanisms. On Windows machines, issuing `dt check -cpuid` from the command-line will show you whether or not your processor supports an invariant TSC. On Linux machines, this information is included in the startup banner of the text log.

■ Operating System

Win8-Win10, or Server 2012-2019 are preferred and are more predictable than Vista, Windows 7, Windows 2008, or Windows 2008r2 for high-accuracy timing. The older XP/Server 2003 platform is also more stable than the problematic Vista/Win7/2008/2008r2 versions.

■ Domain Time Version

PTP is a loose and evolving standard. You should always use the most current version of Domain Time Server or Client to

left or entered the system. Using this information, the application can determine how much extra time was spent handling the packet between the application and the wire. The network socket options `SO_TIMESTAMPING` (hardware), `SO_TIMESTAMPNS` (software), and or `SO_TIMESTAMP` (low-resolution software) are used to enable this functionality on Linux. If the operating system and network driver support either hardware or software timestamps, the measurements of PTP can be significantly more precise than without such support. DTLinux will automatically detect and use the best type of timestamp support. `SO_TIMESTAMPING` requires a NIC capable of providing hardware timestamps.

Current Windows operating systems do not support `SO_TIMESTAMPING`, `SO_TIMESTAMPNS`, or `SO_TIMESTAMP`. However, as of Server 2019 and recent builds of Win10, they do offer NDIS software timestamping measurement of internal network stack delays. Domain Time can use this function, see [KB2019.708](#) for more information.

The theoretical limit of precision for Domain Time on Windows using any protocol is the hectonanosecond. The limit of precision on Linux is the nanosecond. A hectonanosecond is 100 nanoseconds, or 1/10th of a microsecond. The real-world behavior of Domain Time's implementation of PTP achieves results as good or better than any other supported protocol.

be sure you are getting optimal performance and compatibility. Check the [Changelog](#) to see details of the current version.

■ Network Switches

The lower the latency in your network switches or hubs, the better. Switches typically use store-and-forward technology, which means that packets can be delayed even under ideal circumstances. During times of high traffic, the delay time can quickly soar to unacceptable levels. Even a small delay will affect timing, even though other applications may not even notice it. If you have managed switches, examine the statistics to determine if packets are being delayed. You may also set up a test environment and compare the synchronization accuracy in your test environment vs. your production environment. If the performance suffers in your production environment with identical equipment, suspect underpowered switches.

Machines using multicast profiles of PTP require IGMP support from either network routers or switches to join multicast groups. Be sure a device acting as an IGMP Querier (often provided with IGMP Snooping features) is present on each of your network segments.

Some switches and routers have PTP-awareness (such those capable of operating as boundary clocks) which can intercept or interfere with PTP communications. Be sure you have updated all of these to the latest firmware. Many implementations of these products have had significant bugs, such as preventing PTP management messages from being propagated correctly.

If your managed switch supports QoS, you should set it so that PTP traffic (to/from UDP ports 319 and 320) have the highest possibly priority.

■ Other Applications

Most user-mode applications do not affect timing significantly. However, programs that change the system clock resolution, or that overtax the hardware, can degrade performance. If possible, avoid Java-based applications, Flash applications, or any program that exercises the CPU excessively.

■ Miscellaneous Hardware and Drivers

Drivers for NICs, video, and disk subsystems can cause unpredictable delays in servicing interrupts. Windows timing relies on the regularity and predictability of interrupts. If sub-millisecond timing is critical for your environment, use only the best and fastest hardware. Some Linux distros benefit from changing the NOHZ setting (see [Linux Kernel Documentation](#) for more information).

Some switches and routers have PTP-awareness (such those capable of operating as boundary clocks) which can intercept or interfere with PTP communications. Be sure you have updated all of these to the latest firmware. Many implementations of these products have had significant bugs, such as preventing PTP management messages from being propagated correctly.

■ Virtualization

Any machine running either Hyper-V or VMWare will not perform as well as a stand-alone machine, although modern versions of Hyper-V (as of Server 2012) offer dramatically improved timing performance. This applies to both the host operating system and any guests, using any time synchronization protocol. On older hardware or systems with high resource use/low memory, the extra overhead of handling the continual PTP network activity may actually result in more interrupt delays among guests, resulting in additional clock drift.

■ Use a hardware clock as your master, if possible

PTP works best with hardware-based grandmaster clocks. Although you may design a network using only software components, your accuracy will be somewhat diminished. We recommend careful planning to ensure you have hardware clocks available, and that their capacity is commensurate with your expected load. Your grandmaster manufacturer can help you plan your network.



Domain Time's PTP support was developed using the Microsemi SyncServer S600

[\[Click for larger size\]](#)

PTP Profiles

The Precision Time Protocol IEEE1588-2008 (PTP) specification includes definition of several operating profiles. A profile is simply a selection of settings and attributes optimized for different environments or purposes. The idea behind profiles is that, as long as all nodes use the same profile, devices can interoperate with a minimum of user configuration.

The default profiles

IEEE 1588-2008 Annex J specifies two main default profiles for use with PTP. The profiles are identical except for the method of delay measurement specified. Both of the default profiles use the IEEE1588-2008 defined "Best Master Clock" (BMC) algorithm to determine which node on the network is master. It is not unusual for a network to consist of nodes using either of the default protocols.

Although the default profiles specify values for certain parameters, administrators have the ability to change some of them. Many administrator-selectable values are distributed automatically, but several important ones are not. Domain Time automatically queries the master using management messages to discover important variables when possible, minimizing manual configuration. Other PTP devices or implementations may need to be manually configured for these items.

- The **Delay Request-Response Default PTP Profile** (identifier code 00-1B-19-00-01-00) uses End-to-End (E2E) delay measurement. Slaves send periodic delay requests via multicast to the entire network. A master overhearing the request sends a multicast delay response packet with the slave request's sequence number and portIdentity. All nodes on the network hear each request and each response, even though delay requests are only answered by masters, and a particular delay response is only meaningful to the slave listening for that particular packet. The delay request frequency is normally controlled by the master when using E2E. The required interval is sent as part of the master's regular packets. Domain Time II allows you to override this automatic request frequency and specify a fixed interval.
- The **Peer-to-Peer Default PTP Profile** (identifier code 00-1B-19-00-02-00) uses Peer-to-Peer (P2P) for delay measurements. The operation is very similar to E2E, but is primarily used by Boundary Clocks, or between multiple appliance-type grandmasters, rather than by individual nodes. The P2P request packet is larger, and the response may be either one-step or two-step. Two-step responses are used by masters who can measure the exact departure time of the first reply, but cannot dynamically rewrite the contents of the first reply to include the correct time. The second reply contains a more precise timestamp. P2P uses the master's announce frequency times as an implementation-dependent multiplier to derive the delay request frequency. Domain Time II allows you to override this automatic request frequency and specify a fixed interval. P2P suffers the same excessive bandwidth issue as E2E, in that all requests and responses are sent to the entire network via multicast.

Bandwidth reduction

Since using multicast for all PTP traffic is an inefficient use of network resources, methods have been introduced to reduce the amount of multicast traffic, primarily by using directed unicast packets for the delay measurement portion of the PTP conversation between masters and slaves.

■ Hybrid Mode

Hybrid Mode is not defined by IEEE1588-2008, but it is in common use. It does not have an identifier code of its own; rather it operates using either E2E or P2P with one important difference: Delay requests are sent directly to the master using unicast. Masters that support hybrid mode will send the reply by unicast directly back to the requesting slave. This allows messages that need to be seen by the entire network (a master's Announces and Syncs) to continue being multicast, but limits delay measurement exchanges to only the nodes that need to see them. Other nodes do not need to process and discard inapplicable delay measurement requests or replies. Domain Time, PTPd, and most other implementations of IEEE1588-2008, support hybrid mode. Most grandmaster appliances also support it. Appliances that do not support hybrid mode may respond to unicast requests with multicast responses, or may not respond at all (the behavior is manufacturer-defined). Domain Time automatically detects both which profile to use and whether or not hybrid mode is supported by the master.

■ Unicast Negotiation

Unicast negotiation is described in Table 34 of the IEEE1588-2008 standard (REQUEST_UNICAST_TRANSMISSION). Negotiated unicast works somewhat like DHCP, in that leases are requested, granted or denied, have expiry times, and must be renewed. This is primarily useful in networks that do not allow multicast, such as older telecommunication

networks. Domain Time supports unicast negotiation as a slave (as of v5.2.b.20180606) or master (as of v5.2.b.20180801) using the [Telecom Profile](#).

PTP Domains

IEEE 1588-2008 v2.0 specifies a "domain" to be a number in the range of 0-127 (default 0), that logically separates multiple timing networks on the same wire. Version 2.1 allows domain numbers to extend to 239. Each node is required to have an operating domain number, and to ignore messages originating from or addressed to a different domain number.

■ Dynamic Domains

Two equal grandmasters in the same domain will not both send Announce and Sync messages. Using the BMC algorithm, they will decide which is primary and which is secondary. The secondary grandmaster will enter passive mode, waiting for the primary to go offline or suffer a clock degradation, at which point the roles will switch. This mechanism works as expected as long as the grandmasters use the same domain number.

Multiple grandmasters visible on the same wire may all act as primaries as long as they each use a different domain number. In this way, networks can be logically separated even if not physically segmented.

Slaves, too, have an operating domain, and IEEE1588-2008 requires them to exercise the BMC algorithm only among masters advertising in the same domain. The Enterprise Profile ([see below](#)) provides for slaves to follow masters in more than one domain, although the mechanism and selection criteria are implementation-dependent. As of version 5.2.b.20161215, Domain Time supports the dynamic domain concept, whether or not the Enterprise Profile is specified. When dynamic domain is enabled, Domain Time slaves will choose the best master from among all visible masters in domains 0 through 127 (or through 239 if using v2.1), and change their operating domain to match the domain of the selected master.

■ The "All Domains" Domain Number

Domain numbers are conveyed in an unsigned 8-bit value. Ordinary v2.0 clocks are restricted to domains 0-127. Values from 128-239 are available for version 2.1 use. Other values are reserved.

IEEE1588-2008 supports special values for clockIdentities and portNumbers that act as wildcards for management messages. A clockIdentity of all 1's is addressed to "all clocks," and a portNumber of all 1's is addressed to "all ports." Therefore a portIdentity of all 1's means "all ports of all clocks" but is limited to nodes using the same domain.

Oddly, IEEE1588-2008 does not include a wildcard "all domains" value. An 8-bit value of all 1's is 255 (0xFF), and is reserved by the standard, so is the perfect candidate for a domain wildcard value. Domain Time will respond to management messages addressed to domain 255. The reply will contain the node's current operating domain. In theory, a management station could send a single multicast addressed to all clocks, all ports, all domains, and, by collecting the responses, discover the identities of each node, including all of the logical domains in use.

Because the "all domains" domain wildcard is not in PTPv2.0 or v2.1, Domain Time does not place messages addressed to the wildcard domain on the wire, but will respond correctly to management messages addressed this way.

Other Profiles

Various additional profiles have been introduced to address specialized needs or industry requirements.

- The **Enterprise Profile** (identifier code 00-00-5E-00-01-00, see [draft 09](#) et sequelae) is an emerging standard that is already widely deployed. It is a specialized version of the default End-to-End Profile, and is interoperable with the default End-to-End Profile. It provides potentially tighter synchronization, allows for slaves to follow masters in multiple domains, and eliminates the need to query the masters for network variables.

Domain Time, in auto-detect mode, will use the values provided by the master, regardless of whether the master is using the End-to-End Default Profile or the Enterprise Profile. As of version 5.2.b.20170101, you may force Domain Time to use the Enterprise Profile.

When using the Enterprise Profile as a slave, Domain Time will not query the master; it always uses the values specified by the profile, and always uses hybrid mode. When using the Enterprise Profile as a master, Domain Time uses the values specified by the profile, will respond to multicast queries with multicast responses, and will respond to unicast

queries with unicast responses.

Note that although the Enterprise Profile's basic goal is largely designed around codification of hybrid mode, it also sets several values intended to be fixed. As with the default profiles, manufacturers can vary some of these values. In particular, a master may send Sync messages more often than once per second, and it may provide, when queried, a different delay measurement interval for slaves to use.

Key Points

- This profile is a codification of hybrid mode, with fixed values for packet intervals
 - Uses a fixed announce multiplier of 3 through the network
 - Requires masters to use a fixed announce interval of 1
 - Requires masters to use a fixed sync interval of 1
 - Requires masters to use multicast for Announce and Sync messages
 - Requires masters to respond to unicast delay requests with unicast delay responses
 - Slaves use a fixed delay measurement interval of 1
 - Slaves are forbidden from attempting unicast negotiation
 - Slaves are required to use E2E instead of P2P for delay measurement
 - Slaves are allowed to choose from masters in multiple domains
 - Slaves are encouraged to use unicast for delay requests
 - Nodes are not forbidden from using multicast for management messages
 - Masters are not forbidden from responding to multicast messages via multicast
- The **Telecom Profile** (identifier code 00-19-A7-00-01-00) is not one of the default profiles and is not compatible with them. Each slave node must be configured with the IP address and domain number of the master(s) it is allowed to use. The slave then obtains separate leases for Announces, Syncs, and E2E delay measurements by exchanging unicast messages with the master(s). The Telecom Profile uses a custom BMC algorithm, as defined by ITU-T G.8265.1, using the master's clockClass, the master's priority 2 value, and finally the local priority. Local priority is based on the order in which you provision the list of masters. The Telecom Profile requires that a given node is either always a slave or always a master; this is why Domain Time, when configured to use the Telecom Profile as a slave, can never assume the master role. Domain Time supports unicast negotiation as a slave (as of v5.2.b.20180606) or master (as of v5.2.b.20180801).
- Since you must provision a Telecom Profile slave with the master's IP and domain number, the Dynamic Domain checkbox will be grayed-out. Domain Time can follow up to sixteen Telecom masters in any combination of domains.
- The **802.1AS Profile** (identifier code 00-80-C2-00-01-00) is aimed primarily at the audio-video industry. This is a specialized version of the default Peer-to-Peer Profile using multicast for all messages, and different requirements for certain types of message contents. Its primary goal is to provide reliable timestamps across multiple nodes used for streaming source material. It is not, generally, interoperable with the default Peer-to-Peer Profile, and there are several mutually-incompatible versions in common use. Domain Time can usually slave successfully to an 802.1AS master, but compatibility must be determined on a trial-and-error basis due to the conflicting definitions.
- The **SMPTE ST-2059-2 Professional Broadcast Environment Profile** (identifier code 68-97-E8-00-01-00) for video synchronization. The profile is similar to the Default Profile, but requires different scheduling of Announce and Sync messages, as well as using Domain 127. Domain Time should be able to synchronize as a slave using this profile.



Please review the [Planning](#) and [Recommended Configurations](#) pages before proceeding.

In order to use Domain Time for Windows as a PTP Slave to an existing Master, you must make the following configuration choices on the [Obtain the Time](#) property page on the Domain Time applet:

- Choose the **Set this machine's clock by querying a list of servers** radio-button.
- On versions prior to 5.2.b.20150828, you **must** uncheck the **Analyze time samples from all servers and choose the best** checkbox to prevent skewing of the time by the additional sample analysis from sources other than PTP. However, on current versions, you may leave this box checked, since non-PTP samples are now automatically excluded. This allows you to have better performance when falling back to DT2 or NTP sources without degrading the PTP performance in normal operation.
- **IMPORTANT:** Specify at least one reliable non-PTP server (either NTP or DT2) in the **List of backup time sources to use** field to act as fallback server(s). This way, when PTP hasn't synchronized yet or when it fails because the master clock goes offline, Domain Time will have another source for time. This also aids in quick recovery should an unexpected large excursion occur (PTP takes a longer time to correct large variances).

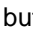
Note: PTP does not begin synchronizing immediately when starting. It must listen to the network to detect the master clock, and must then calibrate to it before usable time samples are available to the Domain Time II clock management algorithms. This means that for a period after startup, or if PTP sync is lost, Domain Time II will not be able to use PTP as a time source. Having a backup NTP or DT2 time source ensures the system has a good source of time in the interim. Also, PTP can take a fairly long time to correct large variances, so having a fallback server avoids having a long training period for PTP to bring the clock into conformance.

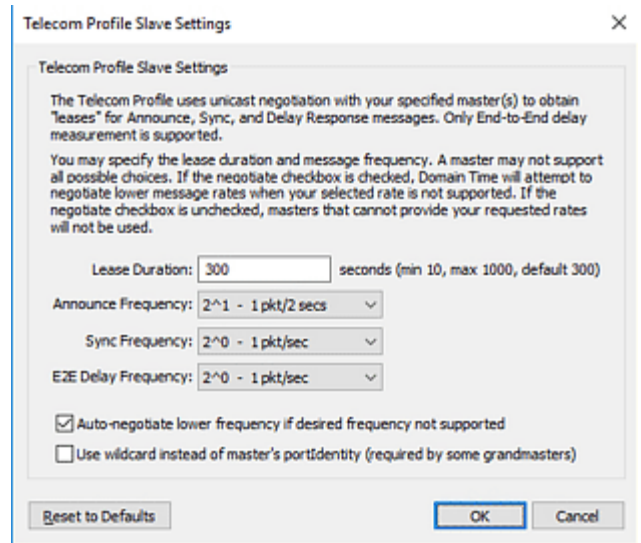
On networks where you don't have any alternative sources of time other than the PTP Grandmaster, you may opt to check the **Accept First PTP Timestamp** checkbox to accelerate convergence from large clock disparities. There are important consequences to this, however. Please see the full description of this option [below](#).

- If your PTP Default or Enterprise profile Master is not on the same subnet as the Client, be sure the Multicast IPv4 TTL (Hop count for IPv6) settings on the Client's [Broadcasts and Multicasts](#) property page are set high enough to allow multicasts to cross intervening switches and routers. Typically you would set this value at least one higher than the number of router hops involved. Note your Master must also have its router hops values configured to allow its multicasts to cross any boundaries. When using the Telecom profile, hop count settings are irrelevant.
- On the **Logs & Status** property page, enable a lazy write delay of at least 5 seconds. This helps prevent disruptions to accuracy from the logging process.
- On the [Timings](#) property page, set both the **Check Interval when able to get and correct the time** and **Check Interval when getting the time fails** settings to use the same fixed period; we HIGHLY recommend a Fixed schedule of every 1 minute. Domain Time's PTP implementation is optimized for this setting. **Do not set these values lower than every 30 seconds when using PTP.** Although these settings do not control PTP synchronization rates, they do affect sample collection and control how often Domain Time synchronizes if falling back to DT2 or NTP, and also how often it records sync status in the text logs and performs other alerting functions. Too short of a period may not allow PTP enough time to collect valid samples.
- Check the **Enable IEEE 1588 Precision Time Protocol (PTP)** checkbox.
- If you are going to be synchronizing using the Telecom profile, click the "Options" link and make the following

additional selections on the PTP Options dialog (otherwise, leave the options at the default settings):

- Choose the "End-to-End Telecom Profile" from the **PTP Profile** dropdown list.
- Select the **Only select best master clock from among those listed below** radio button in the Best Master Clock Options section. Enter the IP address of the Telecom master server(s) you want to use - one per line. Domain Time will subscribe to announce messages from each listed server, and decide which to use based on their advertised quality. If all servers are equal, Domain Time will attempt to use the servers in the order listed.

- If your telecom server is non-standard, you may click the  button to set Telecom-specific settings. Do not change these values unless your server requires them.



The Telecom Profile Slave Options dialog

[Click for larger size]

- Click the **Apply** button to save your changes. Domain Time will then start listening for Master announcements, and will automatically calibrate and start synchronizing with the Master. You may check the activity of the PTP conversation by clicking the **Status** and **Graph** links.
- OPTIONAL: As of v5.2.b.20190701, Domain Time Client and Server support Windows NDIS software timestamping, which allows measurement of network stack delays. Software timestamping is only available on Server 2019 (or newer) and recent updates of Win10. You may want to experiment with this setting to see if it improves your accuracy. See the [Use software timestamping](#) section on the Advanced property page for more details.

Troubleshooting

The default PTP settings should be sufficient for Domain Time to synchronize with most Master clocks. However, you may encounter some network issues or non-compliant Master servers that prevent synchronization. Here are some steps you can try:

- Verify that multicasts are allowed on your network. See the [Planning](#) documentation for network requirements.
- Ensure that there is an IGMP (if using IPv4) or MLD (IPv6)-enabled router or switch attached to the same subnet as your PTP Master and Slaves. IGMP/MLD is used to join multicast groups and provides the most robust multicast behavior. Unlike with unicast protocols like NTP, a direct cable connection to a clock with only a crossover cable may not work. If you do try this, be sure all of your machines are booted and have working IP addresses before starting the Domain Time service. Direct-connection is not recommended.
- Enable network communication through firewalls. You may quickly open all relevant time protocol ports in the Windows Firewall using the following command as Administrator from a command-prompt (Version 5.2 and later):

```
dtcheck -firewall:open -all
```

NOTE: Certain versions of Windows block these ports by default regardless of whether the Windows Firewall service is running. The Windows Firewall service must be started to run the above command. If the service is disabled, you must re-enable it temporarily to run the command, then re-disable the service if you wish.

- Ensure that the following checkboxes are checked on the [Network](#) property page:

Initiate rebind and resync if IP address changes

Enumerate multicast interfaces during IPv4/IPv6 bind

Reply to multicasts using incoming interface if possible

- To easily see which PTP masters are visible on your network, Click the *Status* link on the **Obtain the Time** property page and then click the *PTP Masters* link at the bottom of the stats display. Available master servers will appear in this utility. Select the master you want to use, and click the [Add](#) button. This will often be able to automatically configure Domain Time with the correct settings to use that master.

You may also issue the following command from a command-prompt:

```
dtcheck -ptpmasters
```

- Be sure you have not excluded the IP address of your PTP Grandmaster using the **Best Master Clock Options** settings.
- If your machine is multihomed, you may need to restrict Domain Time to listen only on the proper subnet using the **Listen only on these addresses:** list box on the [Network](#) property page.
- Try relaxing strict IEEE 1588 requirements. Click the **Options** link to display the **IEEE Precision Time Protocol Options** page. Then, click the [Advanced](#) button to display the **IEEE 1588 Precision Time Protocol Advanced Settings** dialog page.

First, uncheck the **Require strict IEEE 1588-2008 standard compliance (recommended)** checkbox. Test to see if this configuration now allows synchronization.

Next, uncheck the **Require ptpTimescale (IEEE 15878-2008 section 8.2.4.8)** checkbox. Test to see if this configuration now allows synchronization. Take care with this option. If unchecked, your Master clock may be using a different timescale than expected, resulting in the time being off. Check with your clock manufacturer to be sure the Master clock is using the correct timescale. The PTP timescale is the preferred choice.

- Specify a PTP profile if one isn't being autodetected. On the **IEEE Precision Time Protocol Options** page, you may use the **PTP Profile** dropdown to specify the profile used by your Master clock.
- Specify a defined **Delay Transport** option instead of "Auto-detect." Disable Unicast support if necessary.
- Ensure the **Clock Identity** value is unique on your network. The default value is derived from the MAC address of your primary network adapter during installation. You must change this value to a different guaranteed-unique value if the PTP master to which you want to sync is on the same NIC and/or has the same Clock Identity. Otherwise, do not change this value.
- As mentioned above, newer switches and routers may have PTP-awareness (such as those capable of acting as boundary clocks). Many of these have had firmware bugs that actually interfere with proper PTP operation. If you are using such a device, try swapping in a standard non-PTP capable switch to see if you can communicate properly. If so, contact your hardware vendor to obtain a fix/firmware update.
- If you are using PTP v2.1 authentication, temporarily disable authentication on the PTP Options dialog. If things then start working, you need to verify that your authentication settings match your Master. Check the [Symmetric Keys](#) property page to be sure your SHA256 hash(es) match your Master, and that you are either using an SPP of 0 (wildcard) or the exact SPP being used by the Master.
- Use the [PTPCheck](#) utility to verify your PTP devices are providing management messages. Copy the program to various machines on either side of switches and routers to verify the PTP traffic is being passed by your network equipment without interference.

Configurable Options

While PTP is designed to operate on the local segment without any configuration by the administrator, Domain Time II nevertheless exposes a few options for customization.

■ PTP Domain

Networks of clocks using PTP are separated into logical “domains,” or groups of clocks. Nodes will only respond to messages sent to their domain. The domain is specified by a number, and the default domain is 0 (zero). Any number between 0 and 127 inclusive is a valid domain number. If you change this setting on your master clock, you must change each Domain Time II machine to match.

As of version 5.2.b.20170101, you may use the Dynamic Domains feature by checking the **Dynamic (allow any domain when slave)** checkbox.

■ PTP Profile

The default operation of Domain Time II is to detect the delay mechanism used by the master clock. The options are:

- Auto-detect
- End-To-End Default Profile (00-1B-19-00-01-00)
- End-To-End Enterprise Profile (00-00-5E-00-01-00)
- End-To-End Telecom Profile (00-19-A7-00-01-00) - added as of v5.2.b.20180606. This option cannot be auto-discovered.
- Peer-to-Peer Default Profile (00-1B-19-00-02-00)
- Disable link delay measurement

■ Delay Transport

The default operation of Domain Time II is to test the master clock for unicast delay support and use unicast if possible. The options are:

- Auto-detect
- Unicast (hybrid mode)
- Multicast (standard mode)

Note: If **Auto-detect** is selected for **PTP Profile** and/or **Delay Transport**, Domain Time will send a limited number of discovery Delay Requests. If no response is received, it will stop sending requests until the next service restart (to conserve bandwidth). If you need Domain Time to continue to send requests even if it never receives a reply, choose something other than **Auto-detect** for these items.

■ IPv6 Scope

The default operation of Domain Time II's support for PTP over IPv6 is to listen and send on the site-local scope. You may override this to specify link-local scope instead.

■ Statistical options

Domain Time applies statistical analysis to acquired PTP sync and delay samples. In most cases, the default settings are optimal, however, if both your PTP master and slaves are not on a stable, local subnet with low-latency you may find some improvement by adjusting these values. Pay close attention to the results of your changes, as they can have adverse consequences on your accuracy, and even your ability to synchronize at all.

Use smoothing for meanPathDelay: Disable only if you have significant variable latency between your master and slaves.

Cap latency at hectos: Ignores any samples with latency larger than this value. Do **NOT** use unless you have significant variable latency between your master and slaves. Enabling this setting can result in a complete inability to synchronize.

Use smoothing for delta calculations: Disable only if you have significant variable latency or timestamps between your master and slaves.

Enable Trend Filter: On most machines, enabling this filter provides optimal synchronization, however, some machines may do slightly better with it disabled. Uncheck this box only if your machine is extremely stable and normally

has an average delta of only a few microseconds.

Coalesce PTP samples separately: Does not combine samples collected during the sync period into a single logged value, but logs them individually in Trace mode.

Accept First PTP Timestamp after samples: Immediately steps or slews the clock to match the first observed PTP timestamp(s). The **samples** entry indicates the number of timestamps that must be received before the clock is corrected. This option was only available via a registry setting in versions prior to 5.2.b.20200930. Older versions always stepped their correction.

Note: You should enable this option ONLY in situations where you have no other time sources available. **Accept First PTP Timestamp** may STEP the clock, either forward or backward, if the time exceeds slew boundaries.

Accept First PTP Timestamp is designed for isolated networks, where only a PTP grandmaster is available (i.e, no NTP or DT2 sources), and where the machine's startup time of day may be wildly off. It will trigger at the start of the Domain Time service, or when a resume-from-standby or clock-change signal is recognized. If no NTP or DT2 time sources are configured and working, then **Accept First PTP Timestamp** WILL slew or step the clock.

Crosscheck with other sources if delta exceeds ms: If the delta (variance) between the slave and the master exceeds this threshold, Domain time will fall back to the NTP or DT2 time sources listed on the [Obtain the Time](#) property page. There should always be at least one valid time source listed there. PTP is quite slow at correcting large variances both during startup and in the event of a large excursion. NTP and/or DT2 will recover much faster. Once the delta is back within tolerance, PTP will resume synchronization. You should not disable this option under most circumstances.

The **Delay/Pdelay Request Interval** and **Announce Receipt Timeout Multiplier** values should be set to Automatic in almost all cases.

■ PTP v2.1. Authentication Options

As of v5.2.b.20190331, Domain Time supports PTP v2.1 authentication. These options allow you decide whether to use authentication and whether or not to reject unsigned PTP messages. Please see the [Domain Time PTP v2.1. Authentication](#) knowledgebase article for a detailed discussion of how Domain Time implements the v2.1 specification.

Notes:

Authentication requires the use of at least one SHA256 symmetric key shared between both the Master and Slave(s). Symmetric keys are configured on the [Symmetric Keys](#) property page.

Message authentication can introduce a significant amount of overhead, decreasing PTP accuracy. You should enable authentication only if your environment requires the additional security and enable only the minimum number of options. Typically, signing announces is the least invasive option.

Take extreme care when enabling any of the message rejection options, particularly if you are using more than one potential Master clock. Rejecting messages will prevent Domain Time from becoming a Slave if any of your Master clocks do not sign their messages with the same options in a compatible fashion.

PTP v2.1. Authentication TLV processing enabled: Turn Authentication on or off.

Reject unsigned Announces: If checked, Domain Time will reject any unsigned Announce messages.

Reject unsigned Syncs/Follow-Ups: If checked, Domain Time will reject any unsigned Sync or Sync Follow-Up messages.

Reject unsigned Delay Responses: If checked, Domain Time will reject any unsigned Delay Response messages.

Select best master by quality: If selected, Domain Time will use any master selected by the Best Master Clock (BMC) process, regardless of whether its announces are signed.

Prefer signed Announces: If selected, Domain Time will give Masters that sign their announce messages higher priority in the Best Master Clock (BMC) process.

Require signed Announces: If selected, Domain Time will ONLY use Masters that sign their announce messages. Use this option with care.

(None) **E2E Delay Request KeyId**

(None) **PTP Delay request KeyId**

Use these dropdown lists to select the Symmetric Key ID being used by the Master to sign either End-to-End or Peer-to-Peer delay requests. Only KeyIds you've configured on the [Symmetric Keys](#) property page will appear in the list. In most cases, you will not need to configure this setting. Do so only if your Master requires it.

■ Best Master Clock Options

Domain Time II is a software-based implementation of PTP. The default operation according to the standard is for Domain Time II to discover and synchronize with the best master clock available on the local segment. If that clock goes offline and another clock assumes the master role, or if another clock claims to be more authoritative, Domain Time II will begin synchronizing to the new master automatically. The standard behavior is sometimes undesirable. A user could bring up a new master clock on the network that, either by intention or mistake, assumes the master role without providing better time.

Domain Time II optionally allows you to override the automatic selection of master clock by specifying the IP address(es) or CIDR mask(s) of the master(s) to which you are willing to synchronize. When you override the automatic selection of master clock, Domain Time II only pays attention to the IP addresses you specify; messages from other IP addresses are silently discarded.

If you are using the Telecom profile, you must enter the IP addresses of Telecom master(s) in this box. In Telecom mode, servers will be used in the order listed. Default and Enterprise profile PTP nodes automatically choose a master according to the algorithm specified in the standard.

Note: In Default or Enterprise profile mode, specifying a list of machines for synchronization may result in no server being available because the other nodes may elect a master that you have chosen to ignore. Use this option with caution.

As of v5.2.b.20190701, you may enter comments in the specified masters list. Comments are defined as text following a hashtag or semicolon. (If the hashtag or semicolon is the first character, the entire line is considered a comment.) For example, you may use this syntax:

```
; These are our masters:
172.16.13.3    # main server
192.168.33.21 # backup server
```

For how to configure other available options, please see the IEEE 1588-2008/2019 specifications for explanations and recommended settings. Also, contact your vendor for device-specific recommendations/requirements for compatibility with their PTP products.



Please review the [Planning](#) and [Recommended Configurations](#) pages before proceeding.

In order to use Domain Time for Linux (DTLinux) as a PTP Slave to an existing Master, you must make the following configuration choices to the [dtlinux.conf](#) file:

- In the [PTP Settings](#) section of the .conf file, set the **ptp:enabled** setting to **True**.
- **IMPORTANT:** Specify at least one reliable non-PTP server (either NTP or DT2) in the [NTP and DT2 Time Sources](#) section to act as fallback server(s). This way, when PTP hasn't synchronized yet or when it fails because the master clock goes offline, Domain Time II will have another source for time. This also aids in quick recovery should an unexpected large excursion occur (PTP takes a longer time to correct large variances).

Note: PTP does not begin synchronizing immediately when starting. It must listen to the network to detect the master clock, and must then calibrate to it before usable time samples are available to the Domain Time II clock management algorithms. This means that for a period after startup, or if PTP sync is lost, Domain Time II will not be able to use PTP as a time source. Having a backup NTP or DT2 time source ensures the system has a good source of time in the interim. Also, PTP can take a fairly long time to correct large variances, so having a fallback server avoids having a long training period for PTP to bring the clock into conformance.

- In the [Loop Variables](#) section, set the **loop:checkInterval** to 60 seconds. Domain Time's PTP implementation is optimized for this setting. **Do not set these values lower than every 30 seconds when using PTP.** Although this setting does not control the PTP synchronization rate, it does affect sample collection and controls how often Domain Time synchronizes if falling back to DT2 or NTP, and also how often it records sync status in the text logs. Too short of a period may not allow PTP enough time to collect valid samples.
- If your PTP Default or Enterprise profile Master is not on the same subnet as the Client, be sure the **network:multicastTTL** value in the Network Settings section of the .conf file is set high enough to allow multicasts to cross intervening switches and routers. Typically you would set this value at least one higher than the number of router hops involved. Note your Master must also have its router hops values configured to allow its multicasts to cross any boundaries. When using the Telecom profile, this setting is irrelevant.
- If you are going to be synchronizing using the Telecom profile, click the "Options" link and make the following additional selections:

- Set the **ptp:profile** value to **Telecom**
- Use the **telecomMaster** setting to specify the IP address(es) and domain of the Telecom master server(s) you want to use - one per line. Only one server will be used at a time. Domain Time will attempt to use the servers in the order listed, i.e.

```
telecomMaster=192.168.1.3,0 ; master 192.168.1.3, domain 0
telecomMaster=192.168.1.6,2 ; master 192.168.1.6, domain 2
```

- If your telecom server is non-standard, you may need to adjust other telecom options to match your Master. Do not change these values unless your server requires them.
- From a Linux command line, issue **sudo systemctl reload dtlinux.service** to have your changes recognized. If you used [DTLinux control panel](#) on Domain Time II Manager to edit the .conf file, then click either **Apply** or **Save & Exit**; changes will take effect immediately without any need to use **systemctl** (the same feature applies to any .conf changes you make using the DTLinux control panel). Domain Time will then start listening for Master announcements, and will automatically calibrate and start synchronizing with a suitable Default or Enterprise profile Master. You may

verify this by examining the `/var/log/ntpdate/ntpdate.log` file. You can also check the activity of the PTP conversation from the **Graphs & Statistics** menu of the [DTLinux control panel](#) on Domain Time Manager.

Troubleshooting

The default PTP settings should be sufficient for DTLinux to synchronize with most Master clocks. However, you may encounter some network issues or non-compliant Master servers that prevent synchronization. Here are some steps you can try:

- Verify that multicasts are allowed on your network. See the [Planning](#) documentation for network requirements.
- Be sure network communication for the PTP ports 319/UDP & 320/UDP is allowed bidirectionally through all switches and firewalls.
- For multicast profiles, ensure that there is an IGMP (if using IPv4) or MLD (IPv6)-enabled router or switch attached to the same subnet as your PTP Master and Slaves. IGMP/MLD is used to join multicast groups and provides the most robust multicast behavior. Unlike with unicast protocols like NTP, a direct cable connection to a clock with only a crossover cable may not work. If you do try this, be sure all of your machines are booted and have working IP addresses before starting the Domain Time service.

Some IGMP routers need to have multicast group memberships refreshed occasionally or machines can be dropped from the group. If you encounter this behavior, you can use the **network:igmpRefresh** value to resend group membership info on a schedule.

- To easily see which PTP masters are visible on your network, choose **Open PTP Statistics & Masters** from the Graphs and Statistics menu of the [DTLinux control panel](#) on Domain Time Manager. Click the *PTP Masters* link at the bottom of the stats display.

You may also issue the following command from a command-prompt:

```
dtcheck -ptpmasters
```

- If you specified authorized Masters using the **ptp:multicastMaster** settings, be sure you haven't excluded your master in the list.
- If your machine is multihomed, you may need to restrict DTLinux to use a specific network interface using the **network:adapterName** setting.
- Use the **ptp:profile** value to specify the PTP profile being used by your Master if one isn't being autodetected.
- Use the **ptp:delayTransport** value to specify the Delay Transport instead of "Auto."
- Ensure the PTP ClockID/port value is unique on your network. The default value is derived from the MAC address of your primary network adapter during installation. You must change this value to a different guaranteed-unique value if the PTP master to which you want to sync is on the same NIC and/or another machine has the same Clock Identity. The easiest way to do this is to increment the **ptp:portNumber** value. You can also use the following command to reset the ClockID to a random value:

```
sudo dtlinux -resetClockId
```

- Switches and routers may have PTP-awareness (such as those capable of acting as boundary clocks). Many of these have had firmware bugs that actually interfere with proper PTP operation. If you are using such a device, try swapping in a standard non-PTP capable switch to see if you can communicate properly. If so, contact your hardware vendor to obtain a fix/firmware update.
- If you are using PTP v2.1 authentication, temporarily disable authentication. If things then start working, you need to verify that your authentication settings match your Master. Check your `dtlinux.keys` file to be sure your SHA256

hash(es) match your Master, and that you are either using an SPP of 0 (wildcard) or the exact SPP being used by the Master.

- Use the [PTPCheck](#) utility to verify your PTP devices are providing management messages. Copy the program to various machines on either side of switches and routers to verify the PTP traffic is being passed by your network equipment without interference.



Please review the [Planning](#) and [Recommended Configurations](#) pages before proceeding.

In order to use Domain Time as a PTP Master, you must make the following configuration choices on the [Obtain the Time](#) property page on the Domain Time applet:

- Choose the **Set this machine's clock by querying a list of servers** radio-button.
- You may configure Domain Time Server as either the primary Grandmaster clock (using the Default, Enterprise, or Telecom profile), or as a backup master (using the Default or Enterprise profile).

If this machine will be the primary Grandmaster on the network:

- Configure Domain Time Server to get its time from at least one (preferably three or more) trusted external time sources using either the DT2 or NTP time protocols. Note: You cannot configure Domain Time Server to obtain the time via PTP from a different master and also simultaneously serve the time using PTP (i.e. act as a boundary clock).

By default, Domain Time Server will not act as a PTP Master unless it is correctly setting its own time from external sources. If you are on a closed network with no access to external time sources (or you have an internal clock card that sets the internal system clock time), you may set Domain Time Server to use the internal Windows clock as its source. To do this:

- Uncheck the **Refuse to serve time until this machine's clock has been set** checkbox.
- Uncheck or remove all listed time sources from the **Time Sources** list box

Note: On versions, prior to 5.2.b.20200930, per the IEEE 1588 2008 RFC's the Priority 1, Priority 2, and clock accuracy are automatically set to minimum values (regardless of how you have them configured), which may prevent the server from winning the Best Master Clock election. To workaround this limitation on older versions, uncheck all time sources in the list box, then **Add** a new NTP time source using **local host** as the host name. That way, Domain Time Server will look at itself to get the time when it starts. The PTP Server will then use the local clock as its source, but serve time using PTP with the Priority 1, Priority 2, and Clock Accuracy settings as you've configured them.

On version 5.2.b.20200930 and later, PTP configuration settings are always honored, so we recommend you upgrade if possible.

See the [Obtain the Time](#) property page for more info on configuring time sources.

- Check the **Analyze time samples and choose the best, or average equally good samples** checkbox. This ensures Domain Time Server will select the most accurate time from among your listed sources.
- On the [Timings](#) property page, set both the **Check Interval when able to get and correct the time** and **Check Interval when getting the fails** settings to use the same fixed period, i.e. every 1 minute.
- On the [Domain Role](#) property page, uncheck both the **Enable cascade signals** and **Enable Advisory signals** checkboxes. If you have other Domain Time Servers, you may consider disabling these settings on those machines as well. These signals help Domain Time machines running NTP or DT2 protocols to converge more quickly across your domain, but the extra sync triggers can cause problems for Domain Time PTP slaves by shortening their sample collection period unexpectedly.
- If this machine will be a backup master using the Default or Enterprise profiles, also configure the machine to be a [PTP slave](#) as described above. Telecom profile masters cannot be configured as backup slaves.

Be sure the **Enable IEEE 1588 Precision Time Protocol (PTP)** checkbox is checked, then click the **Options** link to display the **IEEE Precision Time Protocol Options** page. Review the listed options for PTP Domain, PTP Profile, Delay Transport, and IPv6 Scope to be sure they're set as you'd like. In most cases, the default settings are correct. Note that domain numbers from 0-127 are assumed to be PTP v2.0 domains unless overridden on the Master Options dialog. Domains above that range are always assumed to be v2.1.

- Then, click the [Options](#) button to display the **Master Options** dialog page. Check the **Allow this machine to become a PTP server** checkbox. Domain Time will then participate in the PTP Best Master Clock election process.
- If the **Send PTP v2.1 messages regardless of Domain number** checkbox is checked, Domain Time will act as a v2.1 Master regardless of the domain number selected on the PTP Options page. Otherwise, it will serve v2.0 if the domain is between 0-127 and v2.1 if it is above that range.
- Decide if you want to be a **Two-Step Master**. As of v5.2.b.20190331, Domain Time can act as either a One-Step (Sync messages only) or Two-Step (Sync and Sync Followup messages) Master clock. In some environments, Two-Step operation can help slightly with accuracy, at a cost of doubling Sync traffic. In most cases, acting as a One-Step master is preferred.
- Select the delay transport you want Domain Time Server to use:
 - Unicast Telecom subscriptions** if you want Server to act as a Telecom master server.
 - Multicast using IPv4** if you want Server to act as a Default or Enterprise profile server over IPv4. (default)
 - Multicast using IPv6** if you want Server to act as a Default or Enterprise profile server over IPv6.
- Set multicast hop counts if using the Default or Enterprise profiles. If your PTP Master is not on the same subnet as its slaves, be sure the Multicast IPv4 TTL (Hop count for IPv6) settings on the Server's [Broadcasts and Multicasts](#) property page are set high enough to allow multicasts to cross intervening switches and routers. Typically you would set this value at least one higher than the number of router hops involved. Note your Slaves must also have their router hops values configured to allow their multicasts to cross any boundaries.
- If you want Domain Time to sign messages with PTP v2.1 authentication, you need to select the KeyId for each type of message you want to authenticate. You can only choose KeyIds for SHA256 hashes that have been enabled on the [Symmetric Keys](#) property page. Signing creates extra overhead that can affect accuracy, so choose to use only those messages that you require. In most cases, signing Announce messages provides a sufficient level of security at minimum impact. See the [PTP v2.1 Authentication](#) knowledgebase article for more information.
- Verify the Priority 1 and 2 values domain Time will when participating in the PTP Best Master Clock election process.
 - If this machine is to be the Grandmaster, set the Priority 1 and 2 values so that it has higher precedence (a lower numeric value) than any other PTP device on the network.
 - If this machine is not to be the Grandmaster, set the Priority 1 and 2 values so that it has lower precedence (a higher numeric value) than the Grandmaster priority settings.
- Set the TAI-UTC offset to reflect the current number of UTC leap seconds. The simplest way to get the correct current global count of UTC leap seconds is to issue this command from an elevated command prompt which imports the data from IETF servers:

```
dtcheck -leapfile -y
```

Otherwise, you can manually enter the TAI-UTC offset into the registry by entering the number of leapseconds into this registry key (create it if not present):

```
Node: HKLM\SOFTWARE\Greyware\Domain Time Server\Time Sources\PTPv2 (IEEE 1588)
Type: Reg_DWORD
Key Name: TAI-UTC Offset Discovered (seconds)
```

- Restart the Domain Time Server service. You may check the activity of the PTP conversations by clicking the **Status** and **Graph** links.

Domain Time II Server

Version 5.2

Domain Time II Server is a Windows system service that can be configured to obtain time from various time sources (such as GPS clocks and Internet time servers) and match the system clock to them with extreme accuracy and precision. Server then provides the correct time reliably and securely to other systems on the network.

IMPORTANT: If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[Installation Instructions](#)

[System Requirements](#)

Configuring the Server

Once [installed](#), the **Domain Time II Server** service will start automatically when the system boots. All settings used by the service are read from the [Windows registry](#).

There are several ways to configure the Server's settings:

- Use the Domain Time II Server Control Panel applet on the Server machine itself.
- Remotely from another machine:
 - running Domain Time II (see *Connect to..* below).
 - using [Domain Time II Manager](#).
 - using the [RemoteCPL](#) tool (part of the Management Tools).
- [Import settings](#) from a saved configuration file.
- Use [Active Directory policies](#).

The following instructions describe the settings found on the Domain Time II Server Control Panel applet. Follow the links above for instructions on the other configuration methods. The applet can be configured whether the Domain Time II Server service is running or not.

Launch the applet

There are several ways to launch the Domain Time applet:

- From the System Tray (Notification Area) Icon: Double-click the Domain Time icon to launch the Domain Time II Server applet. You may also right-click the tray icon to launch the applet, as well as many other installed Domain Time II components and utilities from the context menu.
- From the Windows Control Panel: Click the **Domain Time Server** icon (it may be located in the **Clock, Language, and Region** section).

Note: On systems with User Account Control (UAC) enabled, you may need to *Shift+Right Click* the icon and choose **Run As...** or **Run As Administrator** from the context menu to launch the Control Panel applet.


- From the the command-prompt: Launch the applet by typing `domtimes. cpl` in the Windows *Start --> Run* dialog or at a command prompt (the file itself is located in the \System32 folder).

The Control Panel Applet


The Domain Time applet has two panels. On the left is the navigation tree, which lets you pick a configuration property page to view by left-clicking an item in the tree. Right-clicking in the navigation tree will bring up a context menu with shortcuts to various functions, such as *Connect to another computer...*, the text and drift logs, online help, etc. On the right-hand side of the applet is the currently-selected configuration page.

Click the **Domain Time Server (local)** item on the navigation tree to display information about the installed service, including version information, Serial Number, stats, and Start/Stop control.

The *Service Status* display gives you a quick overview of the state of the Domain Time service. This section will be blank if the service is not started, and may take a few moments to display after a service restart. Click the Refresh button to update the display.



Use the *Service Control* section to stop and restart the Domain Time service. Most changes you make using the Control Panel applet are dynamic and should not require you to restart the service.




Connect to another machine running Domain Time

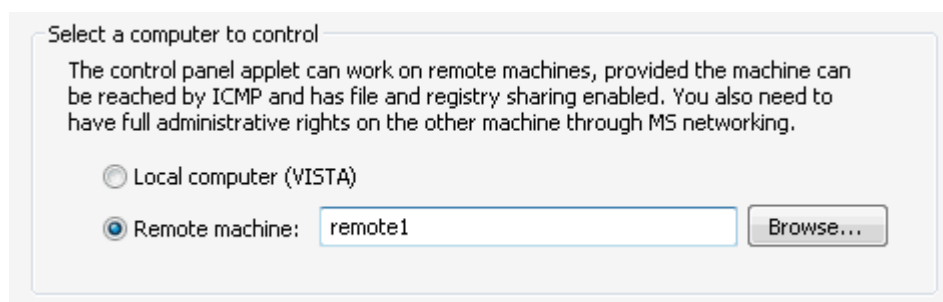
You can also use the Domain Time Control Panel applet to connect to and configure other machines running Domain Time version 5.1 or later. This is particularly useful for quick configuration changes to a few machines, or for configuring Domain Time on Windows Server Core systems. If you need to configure many machines, you will want to use [Domain Time Manager](#) and/or use [Active Directory policies](#) instead.

You must be able to log on to the remote machine with an account that has administrative rights to the remote system. Your machine must also have sufficient network connectivity to authenticate with the remote system using Microsoft networking (see the [Planning](#) page for complete network information).

Right-click any item on the navigation tree and select *Connect to another computer...* from the context menu. If necessary, you'll be prompted to enter an administrative account and password to the remote machine.



When connected, the **Domain Time Server (local)** item in the navigation tree will change to display the name of the remote computer (you'll also see the name in the title bar and the background color of the menu tree will change). Now, any changes you make to the applet will be made directly on the remote system. All of the functions of the Control Panel applet, such as the log viewers, stats display, etc. also behave as if you were directly using the applet on the remote machine.



To disconnect from the remote system, you can either right-click -> *Connect to another computer...* again and choose **Local computer**, or simply close the applet.



Please read the [Installation Topics](#) pages before installing.

In most cases, you'll want to use [Domain Time Manager](#) to install and configure Domain Time Servers and Clients remotely across your network from your management workstation. If that's what you'd like to do, you can skip these installation instructions and read the [Network Rollout](#) documentation instead.

It's also possible to use an existing installation of Server install to another machine remotely using the command-line. This option is for advanced users. See the [Command-line Options](#) page for details.

If you use cloned OS images to install machines, please read [this article](#) from our knowledgebase about configuring Domain Time properly on your clone image.

NOTES:

- Windows Nano Server has special installation requirements. See the [Nano Server FAQ](#) for details.
- See the [Planning](#) page for information on picking a suitable machine for serving time.
- If you will be installing Domain Time Server on a virtual machine, see [this article](#) from our knowledgebase for more information on proper use with virtualization systems.
- When installing Domain Time Server, the software will automatically disable the Windows Time service. On Domain Controllers, Domain Time supplies the necessary authenticated NTP time packets for domain members running Windows Time NT5DS-mode as well as standard NTP for other clients. The software sets announce flags so that the machine is seen as a reliable time provider to Windows Time clients (and utilities such as DCDIAG). However, these announce flags may not always be visible until the machine is rebooted, so you should plan to reboot the DCs after Domain Time Server is installed.
- If you are installing Domain Time Server onto a Windows Cluster, there are considerations regarding Cluster Service startup dependency on the Windows Time Service. Please see the **NoSync** section of the [Advanced -> Windows Time Mode](#) settings for more information.
- Check your routers and firewalls to be sure the ports for the time protocols you'll be using are open. Port 9909 TCP & UDP should always be open bi-directionally between Domain Time Servers and Clients. Port 123 UDP should be allowed if you will be using the NTP protocol for time synchronization. Ports 319 UDP & 320 UDP should be open bi-directionally for PTP use.

Domain Time version 5.2 and later includes a handy utility for adding the correct ports to the internal Windows firewall. Issue the following command from a command-prompt elevated with administrator privileges:

```
dtcheck /firewall:open
```

Hint: Run this command on any machine running Domain Time Server, even if the Windows Firewall is disabled.

- If you will be installing Domain Time onto machines with AMD processors, we highly recommend you update your processor drivers (a.k.a. PowerNow!) to the current version for your operating system available from AMD's website to avoid known hardware timing issues. Please see this article from our knowledgebase for more info: [KB2007.817](#).

Installation/Upgrade

- To install or upgrade Domain Time Server directly to a single machine from the distribution setup files:
 - Run the Setup program from the CD to install the program. (If you have an older version of Server installed, Setup will give you an upgrade option. Your original configuration settings will be preserved during the upgrade). See [this page](#) for details on using the Setup utility.
 - Start the Domain Time Server applet from the icon in the Windows Control Panel to configure it.

Note: On systems with User Account Control (UAC) enabled, you may need to *Shift+Right Click* and choose **Run As...** from the context menu to launch the Control Panel applet. On Windows Server Core, type in `domtimes.cpl` on the command line)

- Use the [Obtain the Time](#) property page to set Server to get the time:
 - from local GPS Network Time Servers, Domain Time Servers, or from other reliable network time sources, or
 - from an internal clock card (or if you don't have an external time source), choose the **Do not set this machine's time** checkbox.
- Use the [Serve the Time](#) property page to choose which protocols Server should provide to clients.
- If the machine is a Domain Controller, reboot to ensure the announce flags are advertising correctly (see NOTES: section above).
- Test your installation
 - Click the **Sync** button on the Control Panel Applet.
 - Click **View Log** button to see the service activity log. You should see messages indicating that the Domain Time service set its time correctly from the time source(s) you selected and is now serving time to clients.


Removal

- Use [Domain Time Manager](#) to remove the program remotely.
- Or, use the **Programs and Features** (Add/Remove Programs) utility from the Windows Control Panel to remove the program.
- You may also use the original [Setup program](#) to remove the program. Run Setup and choose the Remove option.
- The program can also be uninstalled from the command-line. See the [Command-line Options](#) page for details.



Use this page to configure where Domain Time will get the time to set the local system clock.

Note: If Domain Time Server on this machine is set as a Slave server, the options described below will be unavailable since Slaves inherit these settings from the Master server. You'll see a Slave Time Source statistics page instead. See [Domain Roles](#) for more information on Master, Slave, and Independent Server roles.

If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

If the machine is either a Master or Independent Server, you have three basic methods of obtaining the time:

External Time Sources

- Set this machine's time by querying a list of servers (recommended)
- Set this machine's time from broadcast or multicast sources (deprecated)
- Do not set this machine's time

The first two selections are methods Domain Time can use to acquire the time from external network source(s):

Set this machine's time by querying a list of servers (recommended)

This selection instructs Domain Time to make outgoing unicast time requests to the servers you list on this page. Domain Time will query this list on the schedule you set on the [Timings](#) property page.

See the [Time Sources \(Unicast\)](#) section below for details on configuring Domain Time for this method.

As of version 5.2, The IEEE 1588-2008 (PTP) protocol options will also become available when this method is selected. The protocol is enabled using the checkbox on this page (see the Additional Options section below), but is configured on its own dialog screens. Click the IEEE 1588-2008 (PTP) [Options](#) link to configure PTP. See the [IEEE 1588-2008 \(PTP\)](#) documentation page for details.

Set this machine's time from broadcast or multicast sources (deprecated)

This selection sets Domain Time to listen for incoming broadcast or multicast DT2/NTP time packets that are being transmitted from the sources you list on this page. Domain Time will set the local clock whenever it receives a time packet from the listed source(s). This selection is marked (deprecated) because very few administrators choose to use broadcast/multicast DT2/NTP for distributing the time, but the option is still supported.

See the [Time Sources \(Broadcast/Multicast\)](#) section below for details on configuring Domain Time for this method.

The third basic method disables Domain Time's time-setting functions on this machine:

Do not set this machine's time

If you select this option, Domain Time will not attempt to obtain the time from an external time server. Choose this option if you do not have access to an external time source, or if you have another time product installed that is synchronizing the local clock (such as if you have a GPS unit directly connected to the machine with its own driver software). You may still use Domain Time Server to serve the time when this option is selected, but note that the time being served can only be as

accurate as the local clock.

Additional Options

The following options may be available depending on which of the three basic methods of obtaining the time you've selected (see above):

Analyze time samples from all servers and choose the best
Refuse to serve time until this machine's clock has been set

Analyze time samples from all servers and choose the best

This controls whether Domain Time applies advanced analysis algorithms to all of the collected time samples.

When this box is checked, Domain Time contacts all of the listed servers to collect a group of time samples (if you're querying servers) or waits until it has collected the specified number of incoming time packets (if you're using broadcast/multicast sources). It then performs statistical analysis on the collected samples to determine the reliability and uses the most reliable samples to derive the correct time.

See the "About Time Samples" sidebar on the right side of this page for more information and rule-of-thumb suggestions on acquiring time samples.

If you are collecting multiple samples from unicast or broadcast sources using the NTP or DT2 protocols, checking this box will almost always improve your machine's accuracy and reliability.

Note: If you are using the IEEE 1588-2008 (PTP) protocol to synchronize your time, including other time sources in the time calculations can cause inaccuracies at very high levels of precision. Therefore, as of version 5.2.b.20150828, Domain Time automatically excludes all other sources from time calculations when using PTP, falling back to them only if PTP fails, so you may leave this box checked. However, on versions prior to 5.2.b.20150828, you must manually uncheck this checkbox to prevent skewing of the time from additional non-PTP sources.

If this box is unchecked, no comparative analysis among samples is performed. In addition, the list of time servers to query becomes a **fallback-only list**. In other words, the Server will only contact the first listed time server. This server will always be used unless it becomes unavailable, at which point the next listed server will be used. If that server is unavailable, the next server in the list will be tried, etc. When the first listed server becomes available again, the Server will revert to using it exclusively.

When analysis is enabled and more than one time source is used in a time calculation, the logs (when set to the default "Information" detail level) and other display fields without room for multiple entries will show the source for the time as "Averaged Time", otherwise the IP address of the single time source used will be displayed.

Refuse to serve time until this machine's clock has been set

This ensures that Domain Time Server will not serve incorrect time to clients if it

About Time Samples

When obtaining time from external time sources, Domain Time uses Time Samples to determine the correct time.

A time sample is collected either by (a) sending a unicast time request to a time server and receiving a unicast reply, or by (b) accepting an incoming time packet sent from a broadcasting or multicasting server.

By default, Domain Time analyzes the collected time samples using sophisticated statistical methods to reject bad samples and derive the correct time. It then sets the local clock to the correct time using the greatest accuracy possible.

Any single time sample from any time source may not reflect the correct time, either because of network delays, operating system load, or many other transient causes. Therefore, it's usually a good idea to collect more than one sample. If querying a list of servers, you may specify multiple time servers and also set the number of samples to request from each source. If accepting incoming broad/multicast packets, you can specify the number of samples that must be received from the source before making a correction.

In general, time will be more accurate the more samples you collect; however, there is a point of diminishing returns. Each sample takes a fixed amount of time to collect. If the overall time taken to collect the samples is too long, the clock may drift significantly in the interim so that any additional accuracy you obtain from the larger

hasn't yet set its own time or if all of its time sources have become unavailable.

This is the default setting. Note that this function is automatically turned off if you have selected the **Do not set this machine's time** method above.

Time Sources (Unicast)

If you have selected the **Set this machine's time by querying a list of servers** method of obtaining time, Domain Time will query the machines you list (and enable) on this page for the current time.

☐ Enable IEEE 1588 Precision Time Protocol (PTPv2) [Options](#) [Status](#) [Graph](#)

List of time sources to use. Unchecked sources will be ignored. [Import/Export](#)

Server Name or IP	Protocol	Auth	Reps	Delay	Comment
<input checked="" type="checkbox"/> time-a.nist.gov	NTP	None	1	n/a	
<input checked="" type="checkbox"/> time-b.nist.gov	NTP	None	1	n/a	
<input checked="" type="checkbox"/> nist1.symmetricom.com	NTP	None	1	n/a	
<input checked="" type="checkbox"/> tick.greyware.com	DT2-UDP	None	3	512	
<input type="checkbox"/> tock.greyware.com	DT2-HTTP	None	1	n/a	

[Local Time Sources](#)

Unicast Time Source List [\[Click for larger size\]](#)

The machine list can consist of servers that use the NTP, DT2 (UDP or TCP), or DT2-HTTP protocols. As of version 5.2, you may also select the IEEE 1588-2008 (PTP) Precision Time Protocol as a time source. See the [IEEE 1588-2008 \(PTP\)](#) page for details on using the Precision Time Protocol.

You may add machines to the list manually or by scanning for them on your network automatically.

- To easily identify available time servers on your network, click the [Local Time Servers](#) link at the bottom of the list box. This brings up the **Time Sources Search** dialog, where you can scan your network for time servers and then add your choice(s) to the Time Sources list automatically.

Local Time Sources Search

☒ Use IPv4 broadcast
☒ Use IPv4 multicast
☒ Use IPv6 multicast

5 local sources found

Search Results

Server IP Address	Protocol	Discovery Method
<input checked="" type="checkbox"/> 192.168.10.2	DT2-UDP	discovered indie
<input type="checkbox"/> 192.168.10.2	NTP	discovered ntp server
<input type="checkbox"/> 192.168.10.3	NTP	discovered ntp server
<input checked="" type="checkbox"/> 192.168.198.1	DT2-UDP	discovered indie
<input checked="" type="checkbox"/> fe80::a5cd:49e5...	NTP	discovered ntp server

Check the box(es) for the server(s) you want to add to your server list.

Search for Time Sources Automatically [\[Click for larger size\]](#)

number of samples will be offset by the additional drift.

A good rule of thumb for querying servers is to configure at least three or more sources, which provides additional sanity checks and fallback in case any one server is unavailable. Then, specify an odd number of samples from each server; three samples each is a good choice. An odd-number of samples makes the calculations necessary to reject a bad ticker more likely to be accurate. You can then use trial-and-error to determine if adding more samples increases your accuracy.

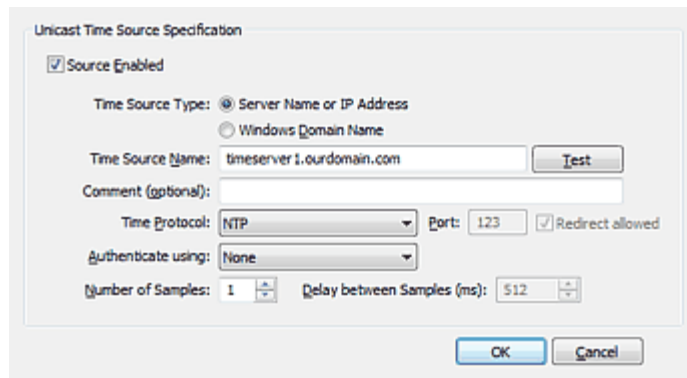
If taking multiple samples from any time server, take care to request a reasonable number of samples and set a delay between the samples to avoid being flagged as making a Denial-of-Service attack.

If you're using broadcasting/multicasting, you can require that multiple samples be collected before setting the time. However, multiple samples may or may not increase accuracy, depending on a number of factors. Consider this option only if the broadcast/multicast time pulse is occurring rapidly enough to collect your required number of samples before the clock drifts outside your target tolerance.

As of v5.2.b.20190701, Domain Time Client and Server support Windows NDIS software timestamping, which allows measurement of network stack delays. Software timestamping is only available on Server 2019 (or newer) and recent updates of Win10. You may want to experiment with this setting to see if it improves your accuracy. See [KB2019.708](#) for more information.

To manually add a time server to your list of time sources, click the

button. This brings up the **Add Time Source**

The image shows a Windows dialog box titled "Unicast Time Source Specification". It has a checkbox "Source Enabled" which is checked. Below it are two radio buttons for "Time Source Type": "Server Name or IP Address" (selected) and "Windows Domain Name". The "Time Source Name" field contains "timeserver1.ourdomain.com" and has a "Test" button next to it. There is an empty "Comment (optional):" field. The "Time Protocol" is set to "NTP" in a dropdown menu, with a "Port" field set to "123" and a checked "Redirect allowed" checkbox. The "Authenticate using:" dropdown is set to "None". At the bottom, "Number of Samples" is set to "1" and "Delay between Samples (ms)" is set to "512". "OK" and "Cancel" buttons are at the bottom right.

[Add Unicast Time Source](#) [\[Click for larger size\]](#)

If you will be using time servers over the Internet, please click the [Public Time Servers](#) link to find reliable servers.

Use the **Time Source Type**: radio buttons to indicate whether you want to contact a server directly using its machine name or IP address, or to automatically find and use the domain controller holding the PDC Emulator role on the specified Windows domain.

If you enter a machine name in the **Time Source Name** field, it must be resolvable to an IP address using DNS, WINS, Active Directory, from the HOSTS file, etc. If entering the IP address, you may use either the IPv4 or IPv6 address of the server.

You may use the **Comment** field to annotate this entry, if you want.

Note: As of v5.2.b.20190701, NetBIOS and DNS names are checked first for IPv4, and only use IPv6 if the IPv4 lookup fails. (If you use an IP literal, Domain Time will use the protocol family associated with what you entered, and the information in this section does not apply.)

To force a NetBIOS name or DNS name to use either IPv4 or IPv6, enter either the text "IPv4" or "IPv6" anywhere in the comment field. For example, if your source is specified as `ntp.mydomain.com` without specifying either IPv4 or IPv6 in the Comment field, Domain Time will first try to resolve the name using IPv4. If that lookup fails, Domain Time will try to resolve the name using IPv6. If, however, you put either "IPv4" or "IPv6" in the comment line, Domain Time will look up `ntp.mydomain.com`'s IP address using only the IP family you specify.

Use the **Time Protocol**: drop-down list to indicate which time protocol to use when contacting this server. You can use DT2-UDP, DT2-TCP, DT2-HTTP, or NTP.

The **Port**: field displays the IP port used by the selected protocol. This is a display-only field for all protocols except the DT2-HTTP protocol. The DT2-HTTP protocol port may be changed to match the DT2-HTTP listen port set on the [Serve the Time](#) page of the target Domain Time II Server. The default value for this is port 80. The **Redirect allowed** checkbox specifies whether the DT2-HTTP time requests will honor HTTP 301 and 302 redirects. Only one level of HTTP redirection is permitted.

The **Authenticate using**: drop-down list selects which authentication key to use when exchanging packets with this server. A key will show up in the list if it has been configured on the [Symmetric Keys](#) property page of the Control Panel applet.

Domain Time supports MD5 symmetric-key authentication compatible with NTP version 3 and later (AutoKey is not supported), and as of v5.2.b.20170922, SHA1 authentication as well. Windows Authentication compatible with Windows Time NT5DS-mode timestamps is also supported. Either authentication method can be used over any supported time protocol (NTP, DT2-UDP, etc.) See the [Symmetric Keys](#) page for details on using authentication.

Hint: When possible, be sure all of your time systems are working correctly before enabling authentication. Authentication requires a correct setup on both ends of the connection, and changes at either end can cause a previously-working connection to fail. Disabling authentication temporarily should always be one of the first steps when troubleshooting a connection issue.

Number of Samples: sets how many individual requests Domain Time will make of this server during each time check.

CAUTION: Take extreme care with this setting. Many time servers have Denial-of-Service (DOS) protection to prevent abuse. Issuing too many time requests in a row to one server over a short period of time can cause your machine to be locked out or even be permanently blacklisted.

Use the **Delay between samples (ms)** setting to space out your sample requests over a reasonable length of time. You may want to contact the administrator of any time server you will be using to find out what the acceptable retry period is on that server. Another option is to use fewer samples per server and simply check against more servers if you need to increase your sample count.

Click the button to be sure the server you've selected is reachable using the protocol specified.

Note: The Control Panel applet you're using for the test is running in the foreground security context of the currently logged-in user, but, in normal operation, the Domain Time service will use the context of the background service account under which it runs (by default, LocalSystem). There are some circumstances where the foreground test will succeed in contacting a source but the Domain Time service will fail, or vice versa. If this occurs, check your firewall and security settings to allow the Domain Time service the necessary network access to send/receive time protocols.

Import/Export

You may easily save or restore the Time Sources list settings by clicking the [Import/Export](#) link. You can use this function to quickly update just the list of time sources used without affecting any other settings.

This function does not preserve IEEE 1588-2008 (PTP) settings, it only saves/restores machines in the time sources list.

If you have multiple machines to update or need to configure other settings, you should use the full Import/Export features found on the [Advanced -> Import/Export](#) property page or use Domain Time II Manager's [Templates](#) feature.

Time Sources (Broadcast/Multicast)

If you have selected the **Set this machine's time from broadcast or multicast sources** method of obtaining time, Domain Time will listen for broadcast or multicast DT2/NTP packets from the listed time sources and extract time data from them.

In addition to the options described above, you'll see the following settings when you select **Set this machine's time from broadcast or multicast sources**:

Only accept signed packets
Log rejected packets Samples required for sync:
Only accept from well-known source port

Only accept signed packets

About xcasting

Using broadcast or multicast (sometimes referred to in this document as "xcast") time packets to obtain time has distinct advantages and disadvantages.

One advantage of using xcast to obtain time is that there is often lower processing overhead than when you're sending unicast time queries to a server. The unicast method must send queries to various servers, receive the responses, compensate for latency on each

If checked, only packets using authentication will be accepted. See the [Symmetric Keys](#) page for more information on packet authentication.

Log rejected packets

When checked, rejected packets will be noted in the log.

Samples required for sync:

This sets how many time packets with time data must be received before a correction occurs.

This is also the number of samples used for analysis if the **Analyze time samples and choose the best...** checkbox (discussed above) is checked.

Be careful not to specify a number of samples that would result in long period before the clock is corrected, since the clock may drift significantly before all the samples have been collected.

Only accept from well-known source port

If checked, only packets originating from port 123 UDP (if using the NTP protocol) or port 9909 UDP (if using the DT2 protocol).

Use this setting with caution, since the default behavior of many servers is to send outgoing traffic from a random source port.

Broadcast/Multicast Time Source List

Shows the currently configured time sources. Domain Time will only listen for time packets from sources listed (and enabled) here.

Selected broadcast/multicast time sources. Samples from unchecked servers will be ignored.

Server IP Address	Protocol	Estimated Delay
<input checked="" type="checkbox"/> 192.168.10.4	DT2-UDP	1

[Add](#) [Delete](#) [Edit](#) [Local Broadcast Sources](#)

[Broadcast/Multicast Time Source List](#) [\[Click for larger size\]](#)

You may add machines to the list manually or listen for broadcasting servers on your network.

- To easily identify available broadcast time servers on your network, click the [Local Broadcast Sources](#) link at the bottom of the list box. This brings up the **Time Sources** dialog, where you can listen for broadcast sources on your network and then add your choice(s) to the Time Sources list automatically.


sample, and then analyze the samples to determine the correct time. The xcast process is comparatively simple. The listening machine merely accepts the time presented in the packet as valid, subtracts out the estimated latency of the connection (see Estimated delay below) and sets the clock.

This can be useful in tightly-controlled networks where network propagation delays are known and unchanging. Under those circumstances, it is sometimes possible to achieve higher accuracy on clients by using a very rapid xcast pulse rather than by having the clients each make many individual requests of the server.

However, there are many disadvantages to xcast time under normal network operations. The most significant of these is that network conditions are rarely static, and the latency between time server and client can vary substantially in a short period of time. This can severely affect the accuracy of the incoming time stamp, causing jumps in time. In addition, broadcast time can be very susceptible to interference from rogue broadcast servers on the network, packet-spoofing (although signed packets can help avoid this), and other disruptions which can adversely affect reliability.

In most cases, modern time request/reply protocols with sophisticated round-trip latency detection such as NTP and DT2 are the better choice. However, broadcast time is still used by some legacy equipment, so it may be the only time synchronization option available for those devices.

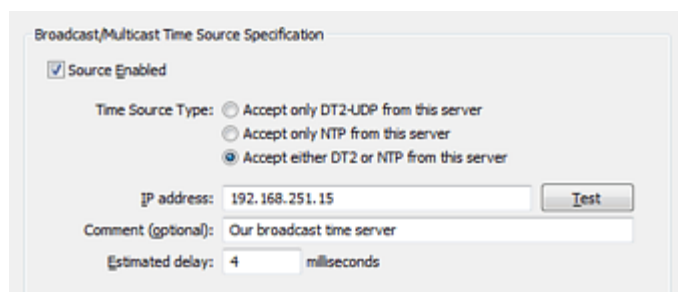
Domain Time allows you to configure to receive broadcasts and/or multicast packets using either the NTP or DT2-UDP protocols. There are efficiencies in the DT2-UDP protocol



that result in slightly-higher accuracy than NTP overall; otherwise, the packets function very similarly.

[Discover Broadcast Time Sources Automatically](#) [\[Click for larger size\]](#)

- To manually add a broadcast time server to your list of time sources, Click the button. This brings up the **Add Time Source** dialog.



[Add Broadcast/Multicast Time Source](#) [\[Click for larger size\]](#)

Use the **Time Source Type**: radio buttons to indicate the type of time packet to listen for from this server. You can accept either DT2-UDP, NTP, or both.

IMPORTANT: Only one service may own a particular port. If you will be accepting NTP broadcast packets with Domain Time, you will need to disable any other service that may be using the NTP port (such as the Windows Time service).

You may use either the IPv4 or IPv6 address of the broadcast server in the **IP Address**: field.

You may use the **Comment** field to annotate this entry, if you want.

Estimated delay is the expected amount of latency in milliseconds a time packet will encounter between the transmitting server and this machine. Domain Time will adjust the time contained in the timestamp by subtracting this value to improve accuracy. The closer this value is to the actual latency on your network connection, the more accurate your time synchronization will be. You may enter this value yourself, or click the button to calculate it for you.

You may need to adjust this value if the overall propagation delay changes on your network.

Import/Export


You may easily save or restore the Time Sources list settings by clicking the [Import/Export](#) link. You can use this function to quickly update just the list of time sources used without affecting any other settings. If you have multiple machines to update or need to configure other settings, you should use the full Import/Export features found on the [Advanced -> Import/Export](#) property page or use Domain Time II Manager's [Templates](#) feature.



Settings on this page control how often time checks are performed when querying servers for the time (if using normal NTP or DT2 protocols) or how often samples are coalesced for statistics/alerting if using PTP.

Note: If Domain Time Server on this machine is set as a Slave server, this page will not appear since Slaves inherit their timing settings from the Master server. See [Domain Roles](#) for more information on Master, Slave, and Independent Server roles.

Timings for broadcast NTP or DT2 are set on the [Serve the Time](#) property page. The settings on this page do not apply if you are using the Broadcast/Multicast method of obtaining network time.

If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

There are two scheduling options for determining how often Domain Time checks the time/reports statistics:

Check Interval when able to get and correct the time

Variable - check as often as needed to maintain approximately _____ milliseconds sync with server

Fixed - check once every _____ DaysHoursMinsSecs

Variable - check as often as needed to maintain approximately milliseconds sync with server

When this option is selected, Domain Time will automatically adjust how often it synchronizes with time sources to attempt to keep the clock within the threshold limit you set.

The wait period between time checks is called the *window size*. You can see the window size Domain Time is currently using by examining the [Text Log](#). Domain Time will adjust the window size based on how accurate previous time checks have been. If previous time corrections have been small the window size will be increased, and vice versa. The range of adjustment for the window size is between 15 seconds to 2 hours. On most machines, it will average between 10 - 30 minutes.

The **Variable** scheduling method is intended for use on machines with relatively constant clock drift and moderate accuracy requirements (where the acceptable tolerance for clock drift is more than ~25ms). This method is a good for general-purpose use, primarily on Clients, since it strikes a good balance between maintaining the target accuracy while minimizing network traffic.

Variable is not a good selection if the machine is under heavy and/or variable

Finding the "Sweet Spot"

The system clock on every machine runs at a different rate because of differences in operating system, applications, hardware, and environment. Even when well-managed by a time program such as Domain Time, the clock will always eventually drift either slower or faster than the actual time,

Since the clock can't be made to run at a perfect rate, it is necessary to correct it when it drifts. In general, the more often you can correct the clock, the more accurate it will be. The corollary to this is that the worse the clock drifts, the more often it must be corrected.

The goal of time correction is to synchronize often enough to keep the system clock within your accuracy target but not so often as to generate excessive network traffic or system overhead. We refer to this ideal rate as the synchronization sweet spot.

In many cases, Domain Time does a very good job of automatically

load that causes the clock to drift by significant amounts on an irregular basis. Since Domain Time may select a large window size if the clock on the machine has been well-behaved, anything that causes sudden clock drift during the Window period between checks can cause the clock to drift outside the specified threshold before the next correction. If this describes your machine, you should use the **Fixed** schedule instead.

Fixed - check once every

DaysHoursMinsSecs

If this option is selected, Domain Time will synchronize regularly on the schedule you specify.

This method is a good choice when you want to discipline the clock to stay within very tight synchronization tolerances. It's also the best choice for machines with highly variable load, poor timekeeping hardware, or any other issue that causes significant clock drift.

You should check the time often enough to keep your clock within your desired accuracy.

Since having highly-accurate time at all times is usually more critical on Servers, you will likely want to check often using a fixed schedule on Servers.

CAUTION: Take care with this setting. Many time servers have Denial-of-Service (DOS) protection to prevent abuse. Issuing too many time requests in a row to one server over a short period of time can cause your machine to be locked out or even be permanently blacklisted.

This problem can be exacerbated if you have opted to collect multiple samples per time source. See [About Time Samples](#).

See the "Finding the Sweet Spot" sidebar on the right for more info on picking the correct sync rate.

discovering the sweet spot (using the Variable scheduling method). It also, by default, uses an overall clock-rate adjustment, to train the clock to run more accurately over time.

However, the more accurate you need the clock to be (or the worse the clock itself is), the more difficult it is for these algorithms to make correct decisions to compensate correctly for drift.

In those cases, you will need to manually set a Fixed synchronization rate, using trial-and-error to find the sweet spot. You may want to start with a reasonable rate such as 5 minutes, and then if that's not sufficient, try every 3 minutes, then 1 minute, etc. On machines with highly-variable drift, you may also need to disable Domain Time's long-term clock adjustment function (see the Correction Reduction section of the [Clock Control](#) page).

Even with severe correction, some systems simply cannot be disciplined enough for every purpose. For example, virtual machines are often too inherently poor at timekeeping to be used for time-critical systems and you must change to physical hardware to achieve your desired accuracy. However, in most cases, you will usually be able to find the sweet spot for your systems by adjusting the synchronization rate appropriately.

Check Interval when getting the time fails

If Domain Time cannot obtain the time, it should try again every: DaysHoursMinsSecs

You may set a different rate for Domain Time to use if it cannot contact any time source.

If Domain Time cannot obtain the time, it should try again every:

DaysHoursMinsSecs


sets how often Domain Time will retry its time sources if it is unable to successfully obtain the time. You will probably want to check more often than during normal operation (unless you're already using a frequent synchronization schedule) to reacquire the correct time quickly when your time source(s) become available.

The same caution about synchronizing too often against your time sources (discussed above) applies.



Settings on this page control what corrections to the system clock are acceptable.

Note: If Domain Time Server on this machine is set as a Slave server, some of these items will be unavailable since Slaves inherit their timing settings from the Master server. See [Domain Roles](#) for more information on Master, Slave, and Independent Server roles.

If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings may adversely affect your clock accuracy or even prevent clock corrections entirely.

IMPORTANT: Settings on this page do not apply (or have minimal effect) if you are synchronizing using [IEEE 1588 \(PTP\)](#).

AS OF v5.2.b.20170922, the Minimum Correction setting is obsolete and has been removed from the applet

The documentation in this section is for reference purposes on older versions only. If you still need to change this value, please see the [Server Settings](#) registry key.

Minimum Correction

Variances smaller than milliseconds will not cause a correction (unless overridden below)

This setting controls the minimum amount the clock must be off from the time source(s) before it is corrected during a time check.

If you are using a variable synchronization schedule (see the [Timings](#) page), you will probably not want to incur the extra overhead of correcting the clock if the variance found during a time check is smaller than your selected synchronization target.

For example, if you have instructed Domain Time to aim for a target variance of 25 ms, you do not need to make a clock correction if the detected clock variance during a time check is only 10 ms. Making a correction in this case would only result in extra processing overhead and clock slewing without much affecting your overall accuracy. So, in general, if you are using variable targeting, you will want to set this value to be the same value or less than your target variance.

However, if you are using a fixed sync schedule, you will want to be sure the clock is corrected to the maximum accuracy on every synchronization. In that case, this value should be set to 1 millisecond. This also enables Domain Time's high-precision sub-millisecond alignment function, so that any variances detected that are less than 1 millisecond will have the clock aligned to match, giving you an added order of magnitude of possible accuracy.

This setting can be overridden under certain circumstances so that the clock can be forced to be corrected. See the **Limit Override** section below for details.

Maximum Correction

No maximum correction

Variances larger than _____ will not cause a correction (unless overridden)
HoursMinsSecs

This setting controls the maximum variance that should be corrected during a time check.

This setting provides a vital sanity check to prevent wild time changes in the event your time source(s) provide a rogue time value (such as sometimes occurs when bounds limits are exceeded or error conditions occur in time clocks or the network).

For example, assume you have restricted this value to not allow corrections for variances larger than 2 hours. If a time source suddenly goes crazy and provides a time/date from 1980, the rogue time correction will be rejected.

The default setting for this value is fairly generous, so you may want to restrict this more in your environment. Do be careful to not restrict this value too tightly. If you have clocks on the network that drift significantly under normal circumstances without restarting (such some laptops do when resuming from sleep modes), setting this value too low may prevent them from ever correcting the clock until they are rebooted.

This setting can (and usually must) be overridden under certain circumstances so that the clock can be forced to be corrected. See the **Minimum/Maximum Limit Override** section below for details.

This setting also interacts directly with the clock slewing settings (which control whether corrections are made by slewing or stepping). See the [Clock Control](#) page for details.

Limit Override

Override the minimum and maximum:

For sync signals, at startup, during training, and when triggered by Clock-Change monitor

For sync signals, at startup, during training, but NOT when triggered by Clock-Change monitor

Only on first correction after machine startup (within _____ seconds of boot)

Use these settings to control when Domain Time will override the correction limits to force a time correction.

The default selection is usually the best option since there are a number of situations where you typically want the time to always be corrected regardless of how far off it may be from the correct time such as:

- when the time service is started
- when triggered to force a correction
- when the clock is being trained (see the [Correction Reduction](#) section of the Clock Control page)
- when Doman Time's Clock-Change Monitor detects that a user or application has unexpectedly attempted to change the time

However, your particular needs may require the ability to restrict corrections even further. If so, you will want to select one of the other listed options. Do be **sure** you fully understand the effects of this selection if you change it from the default.

This setting interacts directly with the clock slewing settings (which control whether corrections are made by slewing or

stepping). You may use the override settings in combination with slewing/stepping limits to ensure that corrections are made only under the precise conditions you desire. See the [Clock Control](#) page for more details.

Advanced Sample Validation

Discard time samples that exceed milliseconds of historical average variance
Discard time samples whose latency exceeds milliseconds, regardless of history

These settings set boundaries on the maximum variance and/or latency permitted in individual time samples.

Discard time samples that exceed milliseconds of historical average variance

Checking this box enables an additional check which may help protect against significantly delayed or skewed time samples. See [About Time Samples](#) on the **Obtain the Time** page for details on how time samples are used.

This setting is turned off by default to minimize overhead, since the default expectation is that your network and time sources will generally be well-behaved. However, if you experience unusual spikes in the time from otherwise reliable sources, you may want to enable this setting to help screen out the errant samples that would otherwise skew your time calculations.

When enabled, Domain Time will keep a historical record of time samples from your selected time sources. It will then reject any time samples that exceed the historical average of the time source by the threshold value you select here.

For example, assume you have set this threshold to 500ms. The historical average variance of time server tick.mydomain.com to date has been +50ms but suddenly a time sample is received with a variance of -475ms. This new sample varies from the historical variance by more than the 500ms threshold value you've set, so the sample will be rejected.

CAUTION: As with other settings on this page, you should only enable this function if you fully understand the ramifications. If you set this value incorrectly, you may end up rejecting so many samples Domain Time will not be able to correctly identify the correct time, or may even be unable to set the clock entirely. Use the Trace or Debug log detail of the [Text Log](#) to see if samples are being rejected and how the accepted samples are being analyzed.


Discard time samples whose latency exceeds milliseconds, regardless of history

This setting is similar to the function described above, except that it rejects individual samples based on a latency limit you specify. This feature was introduced in version 5.2.b.20110309.

CAUTION: As with other settings on this page, you should only enable this function if you fully understand the ramifications. If you set this value incorrectly, you may end up rejecting so many samples Domain Time will not be able to correctly identify the correct time, or may even be unable to set the clock entirely. Use the Trace or Debug log detail of the [Text Log](#) to see if samples are being rejected and how the accepted samples are being analyzed.



Use this page to set which time protocols are served by Domain Time II Server.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Domain Time II on port 9909/udp and tcp

DT2 Server Enabled

Send IPv4 broadcasts

Send IPv4 multicasts

Send IPv6 multicasts

Bcast version: 3

Bcast interval:

HoursMinsSecs

This section controls whether Domain Time II Server will provide time using the Domain Time II (DT2) protocol.

The DT2 protocol is a high-efficiency protocol optimized for time synchronization that also allows for bi-directional communication between Domain Time components using Domain Time control messages that enable advanced management and monitoring functions. This setting only controls whether Domain Time will respond to DT2 time requests and/or transmit DT2 time broadcasts or multicasts.

These settings do **not** affect DT2 control messages. Even if the protocol is turned off on this page, Domain Time II Server will still send/receive DT2 control messages on port 9909 UDP and TCP. See the [Advanced... Command Restrictions](#) section of the **Security Settings** page for information on controlling command messages.

When **DT2 Server Enabled** is checked, Domain Time Server will reply to either UDP or TCP time requests. UDP is an excellent choice for time packets, since it is connectionless and does not incur the overhead necessary to establish and tear down a full TCP session. TCP is often used for sending larger amounts of data in control messages. Your network will need to pass **both** port 9909 UDP and TCP packets in order to work correctly with Domain Time II version 5.x and above. See [Installation Planning](#) for more information.

Should I choose DT2 or NTP? DT2 UDP sync packets tend to be smaller and therefore somewhat more efficient than NTP time packets. In normal networking applications, either protocol is a good choice. However, if you are attempting to achieve very high-accuracy, the difference in overhead between DT2 and NTP packets can become significant and you will want to use DT2 UDP whenever possible.

DT2 Broadcasts/Multicasts

The normal operation of Domain Time Server is to respond to unicast NTP time requests from time clients. This is a "pull" method of obtaining the time, where clients request the time on their own schedule. Domain Time Server can also be set to transmit NTP on a regular schedule. This is a "push" method of distributing time where clients listen for time packets from the server.

The Broadcast and Multicast addresses used by Domain Time are configured on the [Broadcasts and Multicasts](#) page.

The Send IPv4 broadcasts
 Send IPv4 multicasts
 Send IPv6 multicasts

checkboxes on this page set the method used to transmit DT2 packets when using the "push" method of distributing time.

- **IPv4 Broadcasts** are sent to the configured broadcast address (usually the local network). Broadcasts do not typically pass routers, so you will likely use this option only if your clients are on the same local subnet as the server. Broadcasts are deprecated in favor of multicast starting with version 5.x.
- **IPv4/IPv6 Multicasts** are usually able to pass routers/switches without additional configuration and are therefore usually a better choice for transmitting heartbeat packets. Domain Time handles all of the details of joining the correct multicast group for the protocol for you. You may select IPv4 and/or IPv6 multicasts.

Be sure the TTL (IPv4) or Hop Count (IPv6) settings on the [Broadcasts and Multicasts](#) page are set correctly so that multicasts reach your desired subnet(s). Multicasts are supported as of version 5.x.

Bcast Version: selects whether a Broadcast pulse contains a timestamp (version 4) or not (version 3).

If you select version 3, the Client will be triggered to synchronize its time when the pulse is received. It will then make a unicast request for the time from whatever Domain Time Server it is configured to use. Version 3 heartbeats are compatible with Domain Time version 3.1 and above.

If you select Bcast version 4, the Server will include the current time in the broadcast time packet. Clients will set their time to the included timestamp and do not make any requests of a server. Version 4 heartbeats are compatible with Domain Time version 4.x and above.

Multicasts always include the current time in the timestamp. This setting has no effect on Multicasts.

NTP on port 123/udp

NTP Server Enabled

Send IPv4 broadcasts

Send IPv4 multicasts

Send IPv6 multicasts

Bcast version: 1

Bcast interval:

 HoursMins Secs

This section controls whether Domain Time II Server will provide time using the NTP protocol, either by unicast requests from clients or by broadcast and/or multicast.

There are special considerations for this setting when you are running on a Windows Domain Controller. See the **What About Windows Time?** sidebar for details.

When **NTP Server Enabled** is checked, Domain Time Server will reply NTP UDP time requests. Your network will need to pass port 123 UDP packets in order to communicate with devices that use NTP.

Although port 123 UDP is sufficient to pass time traffic, you will also want to pass DT2 control messages if you want to use many of Domain Time's advanced features. See the Network section of the [Planning](#) page for more information.

What about Windows Time?

In most cases, Domain Time II will completely replace Windows Time, so the Windows Time service can and should be disabled in most cases*.

In previous Domain Time versions, it was necessary to set the Windows Time service to run in a special mode called NoSync on Domain Controllers (DCs) to allow Windows Time to provide the proprietary NT5DS-mode authentication timestamp required by domain members running Windows Time.

As of version 5.x, this is no longer required since Domain Time II can provide the correct NT5DS-mode timestamp directly to all domain members.

Installing Domain Time on all your systems is preferred for many reasons, but if this is not practical

When enabled, Server will also be able to respond to a subset of requests from the standard UNIX/Linux **ntp** and **ntpd** query utilities. See the [ntpd Compatibility](#) page for details.

NTP Stratum

By default, Domain Time selects an NTP stratum number based on strata reported by its time sources (including PTP, which uses "stepsAway" to correspond, roughly, with NTP strata). Since Domain Time can use multiple sources with multiple protocols, there may not be an ultimate single stratum from which time was received. In these cases, Domain Time takes the highest-level stratum reported by all used sources, and adds one to derive its own stratum number. So, for example, if Domain Time obtains its time from a PTP grandmaster directly, it will report itself as stratum 2. If it receives its time from three NTP sources, two of which report stratum 1, and one of which reports stratum 2, Domain Time will report itself as stratum 3. If all three NTP sources reported stratum 1, Domain Time would report stratum 2, and so forth.

If there are no upstream timesources (such as when you select the "Do not set this machine's time" radio button on the [Obtain the Time](#) page), Domain Time will assign itself stratum 15, which is the lowest valid stratum number. However, some NTP clients will refuse to synchronize with lower stratum levels. In this case, you will need to use the **HKLM\SOFTWARE\Greyware\Domain Time Server\NTP Server Stratum** registry key to assign a higher stratum, i.e. a registry value of 1 assigns stratum 1.

Versions prior to 5.2.b.20150516, automatic assignment applied a default Stratum depending on the type of Server Role selected:

- Master Server: Stratum 2
- Slave Server: Stratum 3
- Independent Server: Stratum 2

NTP Broadcasts/Multicasts

The normal operation of Domain Time Server is to respond to unicast time requests from time clients. This is a "pull" method of obtaining the time, where clients request the time on their own schedule. Domain Time Server can also be set to transmit time on a regular schedule (also referred to as "heartbeats"). Heartbeats are a "push" method of distributing time where clients do not obtain the time until signaled by the server.

The Broadcast and Multicast addresses used by Domain Time are configured on the [Broadcasts and Multicasts](#) page.

The Send IPv4 broadcasts
 Send IPv4 multicasts
 Send IPv6 multicasts

checkboxes on this page set the method used to transmit NTP packets when using the "push" method of distributing time.

- **IPv4 Broadcasts** are sent to the configured broadcast address (usually the local network). Broadcasts do not typically pass routers, so you will likely use this option only if your clients are on the same local subnet as the server.
- **IPv4/IPv6 Multicasts** are usually able to pass routers/switches without additional configuration and are therefore usually a better choice than broadcasts. Domain Time handles all of the details of joining the correct multicast group for the protocol for you. You may select IPv4 and/or IPv6 multicasts.

Be sure the TTL (IPv4) or Hop Count (IPv6) settings on the [Broadcasts and Multicasts](#) page are set correctly so that

for all your machines, having Domain Time provide NTP with Windows authentication directly to Windows Time clients is still a great advantage in accuracy and reliability over W32time.

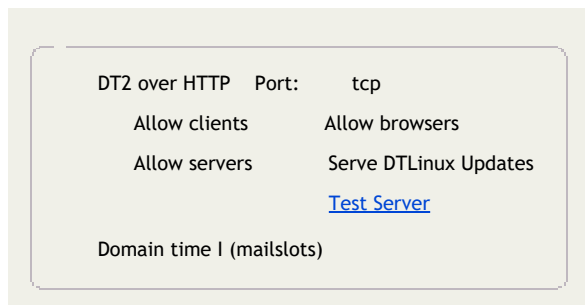
Also, unlike with w32time, you can serve NTP with Windows Authentication from any domain member machine, not just a DC.

* Windows Time will still need to be set to NoSync mode on older (Server 2003 or earlier) Windows Cluster Servers, since the Cluster Service has a dependency on the w32time service in order to start on those versions. w32time is not required after the Cluster Service has started. See the [Advanced](#) page for more information.

When Windows Time is disabled on a DC, you may see spurious messages in the log or in DCDIAG about Windows Time not being synchronized. You can disregard these faux warnings, since Domain Time is indeed accurately synchronizing the domain.

multicasts reach your desired subnet(s). Multicasts are supported as of version 5.x.

Bcast Version: selects what NTP version the time packet reports for compatibility with various NTP implementations. The default value is 3, which should be correct for most situations. Don't select version 1 or 2 if you are using authentication.



The screenshot shows a configuration window with the following settings:

DT2 over HTTP	Port:	tcp
Allow clients	Allow browsers	
Allow servers	Serve DTLinux Updates	
Test Server		
Domain time I (mailslots)		

Select whether to enable the DT2 over HTTP and/or the Domain Time I (mailslots) protocols.

■ DT2 over HTTP

This is a special time protocol that is communicated using standard HTTP packets. The typical port used is TCP port 80.

(Note: some older versions of Domain Time shipped with this setting set to port 90 as the default)

Since this protocol uses the same standard HTTP port used by web browsers, it can pass through most firewalls transparently. This can easily solve firewall connectivity issues for obtaining time.

However, we recommend you only use this protocol for time synchronization if your firewall administrator is unable/unwilling to allow the regular Domain Time II protocol (UDP port 9909) or other native time protocols to pass the firewall. Due to the nature of HTTP and the variability of proxy configurations and load, the DT2 over HTTP protocol is subject to unpredictable latencies that can affect the accuracy of the time being sent. Also, many of the advanced features of Domain Time II such as remote control, monitoring, and upgrade/installation are not available using this protocol.

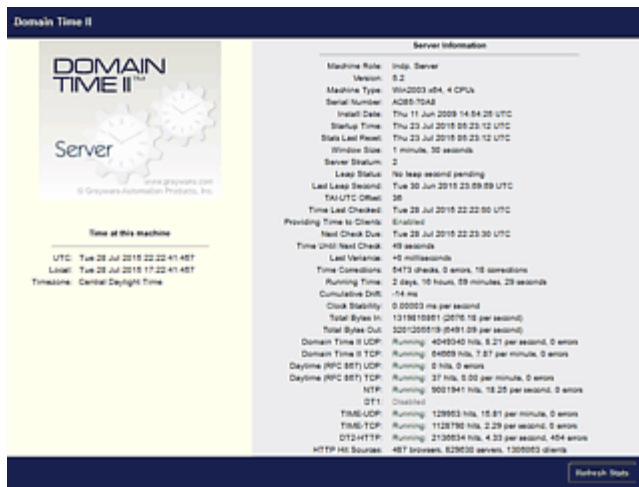
CAUTION: Only one service can "own" the HTTP port 80 on a machine at a time. Domain Time II Server starts very quickly at boot-up. If the Domain Time II over HTTP protocol is enabled, it will very likely attach to port 80 before any other service can start (such as standard web servers like Microsoft's IIS). This will prevent the web server from responding correctly.

If you are running Microsoft's IIS or another web server on your machine, you should change the Domain Time II over HTTP IP port to an unused number to avoid these conflicts.

If you enable the DT2 over HTTP protocol, Domain Time Server can provide a compact DT2 over HTTP time protocol packet (if you have either the **Allow clients** and/or **Allow servers** checkbox checked).

Web Page Stats Display

Domain Time can also provide a special human-readable web page if you have the **Allow browsers** checkbox checked.



Web Page Stats Display [Click for larger size]

The web page provides a great deal of useful information about the status of the Domain Time II Server, and is a great way to monitor the operation of your servers from any machine with a browser. To view the webpage, enter the name or IP address of your Domain Time II server (plus ":" and the port number, if needed) into your web browser address bar, i.e. <http://tick.greyscale.com:80>

You can customize the appearance of the web page using standard HTML code. See the **Customizing the appearance of Domain Time over HTTP** registry entry on the [Server Registry Options](#) page for details.

DTLinux Updates

The **Serve DTLinux Updates** checkbox allows Domain Time Server to act as an alternate data source for DTLinux updates (Domain Time Manager must also be installed on this machine). The [Test Server](#) link will display the available DTLinux distribution files. Please see the [UpdatingDTLinux.html](#) file located in the DTLinux distribution files or in the `/opt/domtime` folder of your DTLinux system for full instructions.

■ Domain Time I (mailslots)

Domain Time I is a proprietary protocol used by version 1.x of Domain Time.

It uses LAN Manager's Mailslots function to obtain the time from a Domain Time Server. You will only need this protocol if you have any Domain Time version 1.x clients in use.

Precision Time Protocol

Allow IEEE 1588 PTP Master mode [Options](#)

You may enable PTP Master mode operation by checking this box. Click the *Options* link to bring up the [PTP configuration](#) pages.


NIST Daytime on port 13/udp or tcp

NIST Daytime UDP Enabled
NIST Daytime TCP Enabled

NIST Daytime is an older, human-readable time protocol.

Domain Time II Server will serve time via the Daytime protocol, but no Domain Time II component uses Daytime for time

synchronization purposes. Domain Time II Server's Daytime string format defaults to the NIST standard (see below), but may be customized to other formats using a [registry setting](#).



TIME/ITP is another older protocol used by some legacy systems.

Domain Time II Server will serve time via the TIME/ITP protocol for compatibility with Domain Time II version 2x - 4x, but as of version 5.x, Domain Time II does not use TIME/ITP for obtaining the time.



Domain Time is designed to easily set up and use a *Time Hierarchy* to distribute the time to machines on your network.


The recommended Domain Time hierarchy follows the Windows domain hierarchy, where the domain controller machine holding the FSMO PDC (Primary Domain Controller) Emulator role acts as the Master Time Server for the domain and all the other domain controllers act as Slave Time Servers who automatically get their time directly from the Master. Clients on the network can auto-discover and get their time from their nearest Slave Server. Domain Time components automatically set themselves up correctly for this hierarchy when installed.

See the [The Domain Time Cascading Hierarchy](#) section below for more details on the Domain Time hierarchy.

Of course, you may choose to set up your own time distribution hierarchy instead of using the suggested Domain Time II Master/Slave method. You do this by using Domain Time II Independent Servers which can be installed on any machine, whether it is part of the domain or not.

Independent Servers do not have the automatic settings replication, client recommendations, auto-promotion, and failover benefits of the Master/Slave relationship, but your clients will respond to Independent Server advisory signals (see below) so that they can auto-locate the server and you can still achieve excellent timekeeping and propagation of time changes to all of your clients.

The [Installation Planning](#) page describes the best practices for configuring your time hierarchy.



This section shows the Domain Role assigned to this machine.

■ Master Time Server

Domain Time II Server installed on the machine with the PDC-Emulator role is the Master Server. This setting cannot be changed and no other machine can be set to be the Master.

If the PDC Emulator role is moved to or seized by another domain controller running Domain Time II Server as a Slave, that machine will automatically become the Master Server and it will take over performing exactly as the previous Master did, using the same settings. Slaves will automatically begin synchronizing with the new machine.

The cascading time hierarchy (see below) will assure that all other machines match this machine's time as accurately as possible and that any corrections propagate nearly instantaneously to every machine.

The Master Server also has the ability to set the recommended timing settings for other Domain Time machines in the time hierarchy so that you can centrally control the accuracy of your time network from one location. See the [Recommendations](#) page for more info.


Time Sources

The Master should be set to get its time from multiple stable, trusted time sources, such as GPS Clocks or other known good network time sources. However, it can also get its time by auto-locating the PDC of another domain by entering the name of the domain in the [Time Sources](#) list. This is useful if you have multiple domains and you want to easily set up a time structure among your domains.



These settings control which configuration recommendations this Domain Time II Server makes to Slaves and/or Clients that synchronize with it.

Note: The recommendations available depend upon the domain role of the server. Options not available for this server role will be greyed-out. See [Domain Roles](#) for more information on Master, Slave, and Independent Server roles.

If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Client Recommendations

Allow clients to match this server's timezone
Recommend timings and correction limits to
clients that ask for guidelines

This section sets which recommendations are available to Domain Time II Clients that synchronize with this machine (or its Slaves) using the DT2 protocol.

Note: The corresponding setting for the selected recommendation must also be enabled on the Client in order for the recommendation to be used. Think of these settings as permission for the Server to answer the settings question if a Client asks. You must also configure the Client to ask the Server for the setting before it can know the answer.

Allow clients to match this server's timezone

When this setting is checked, Clients that ask for this recommendation using the DT2 protocol will automatically set their local machine's timezone to match the timezone of the server. This will override any settings the user may make to their own Time & Date Control Panel.

CAUTION: Be **sure** you understand the effects of this setting before you enable it. For example, assume you have a laptop user with this recommendation enabled in California who usually synchronizes with a Slave in Los Angeles. Each time it synchronizes, Domain Time ensures the laptop timezone matches the Slave which is set to the Pacific timezone.

If the user travels to the New York office, the laptop would change to the Eastern timezone when it begins synchronizing with the Slave in the New York office. However, if the local Slave server becomes unavailable, it's possible that the laptop would fall back to synchronizing with the Master server in Chicago, and then the timezone on the machine would suddenly change to Central, which may not be your intent.

This is the only recommendation available from an Independent Server.

Recommend timings and correction limits to clients that ask for guidelines

This checkbox enables the Master (and its Slaves) to pass along timings and correction limits to Clients that synchronize

using the DT2 protocol. This option is only available on the Domain Time II Master Server.

When these recommendations are active, you can set the overall accuracy of your entire domain from your Master Server by adjusting the recommended settings on the [Client Timings](#) and [Client Correction Limits](#) pages.

Slave Recommendations

Slaves should use master's security settings

This checkbox selects whether the Master's Security Settings are replicated to its Slaves.

Slaves always use the timing settings set for them on the Master's [Slave Timings](#) page. Slaves also automatically replicate most of the Master's other settings to be able to take over should the Master become unavailable. These two functions are always enabled on a Slave.

However, this checkbox allows you to set whether each of your Slaves will use their own individual Security Settings or if they will always use the Master's Security Settings. For example, it may be useful to have individual security if some of your Slaves will be in locations where you want to limit which machines can synchronize with them, but you'd like to have other Slaves accept connections from everyone.

See the [Security Settings](#) page for details on these settings.

Note this is not a "foreign slave" relationship; the remote PDC will not send slave synchronization signals, and the local PDC will not replicate the remote PDC's settings. The older "Foreign Slave" function in versions prior to v5.1 is no longer used.

Alternately, you can set the Master to use its own internal clock as the time source, but this is not recommended unless you have a clock card or other external time clock attached directly to the machine that is correcting the local clock. You can use the uncorrected local clock as a last resort, but be aware that the time will almost certainly drift unpredictably across your domain as your Master's clock drifts.

■ **Slave Time Server**

Any domain member machine other than the PDC-Emulator may be set to be a Slave Time Server. Domain Time II Server installed on a domain controller will automatically become a Slave. Slaves will automatically find and synchronize their time with the Master.

Domain Time II Server in Slave mode performs two important functions in the time hierarchy:

- Slaves allow you to place time servers as close as possible to your time clients (such as on other subnets or remote locations) to allow clients to auto-discover their servers, and maximize accuracy while minimizing the network traffic load.
- Slaves provide robust automatic-failover protection for the network. If the Master becomes unavailable, the Slaves will automatically start obtaining the time using the Master's settings. Also, as described above, if any Slave machine is promoted to the PDC-Emulator role, it will automatically become the new Master machine.

■ **Independent Time Server**

Any machine running Domain Time II Server (other than the PDC-Emulator) may be set to be an Independent Time Server. Independent Servers do not participate in the Master/Slave time hierarchy and are therefore responsible for obtaining their own time from time source(s) you configure.

It is usually better to use Masters and Slaves if you have a domain structure. However, you may use Independent Servers on a domain if you prefer, and you may manually construct your time distribution structure using Independents without using Master and Slaves.

However, if you deploy Independent Servers on a subnet where a Master or Slaves are also visible to clients, you should be sure the independents get their time directly from the Master or the Slaves to prevent sending conflicting time and/or cascade messages to clients.

Cascades and Advisories

These settings control whether Domain Time II Server sends cascade and/or advisory messages to other Domain Time components.

Domain Time II Cascades and Advisories

Enable Cascade signals

Enable Advisory signals

IMPORTANT: These signals greatly enhance the overall synchronization of your time network. Disable only if necessary.

Cascade messages are used to propagate time corrections quickly down the hierarchy without having to wait for each component to synchronize on its own. This causes your clocks to converge on the correct time across your network in seconds instead of minutes/hours it takes using other non-signaled methods such as standard NTP or Windows Time. Cascade messages are considered mandatory and are always acted upon by the receiving component (unless explicitly disabled).

Advisory messages are used to help components determine the structure of the time hierarchy, such as by helping clients auto-locate available time servers. Components may act on or ignore advisory signals depending on their current configuration.

Cascade and Advisory signals may be unicast, broadcast, or multicast (any combination).

Each server has its own settings for whether or not it sends cascades, and if so, what type. A server can send broadcast IPv4 only, multicast IPv6 only, multicast IPv4 only, or any combination. IPv4 broadcasts are sent to 255.255.255.255. This is not configurable. If you need your cascades to cross routers, you must use IPv4 or IPv6 multicast instead.

See the hierarchy description below to see which type of cascade/advisory is used.

The Domain Time Cascading Hierarchy

Domain Time II is designed to use a cascading time hierarchy to distribute the time. The Domain Time hierarchy is more robust than the inbound time partner structure of Windows Time and much simpler than manually configuring NTP peering and strata. In the hierarchy, each server is responsible for matching its time with the server above it, and providing the time to servers or clients below it.

■ Level 1 (Master)

Domain Time II Server installed on the domain controller with the Primary Domain Controller (PDC Emulator) role becomes the *Master* time server for the domain. When the Master's time is corrected to match its time source(s), the Master directs a Level 1 unicast cascade signal to each known Slave. These are the only unicast cascade signals, and they cannot be disabled.

The Master expects an acknowledgement to the Level 1 cascade from the Slave. If a Slave fails to acknowledge (perhaps because it is currently offline), this is noted in the Master's log file.

After signaling each known slave, the master broadcasts/multicasts a Level 1 cascade signal to the network. It will be an advisory if at least one slave was contacted successfully, else mandatory (the assumption being that there are no slaves to relay the signal to waiting clients). A master may be configured to skip sending this Level 1 signal to the network.

■ Level 2 (Slave)

Any other domain controller, member server, or member workstation can be configured as a *Slave* (this is the default for domain controllers). Slaves automatically discover and synchronize with the Master Server. Slaves synchronize time with the Master using the DT2-TCP and protocol, so any intervening firewalls, routers, and/or switches must pass port 9909 TCP (note, it is always a good idea to pass both 9909 TCP **and** UDP traffic).

When a Slave receives a Level 1 cascade signal from the Master, it immediately synchronizes its clock and acknowledges the signal.

After resynchronizing with the Master, each Slave will broadcast/multicast a Level 2 signal to the network. If the resync was due to a Master's Level 1 trigger, the packet will be advisory, else mandatory. The Slave uses the Master's Level 1 sequence number, so any client that happens to hear from multiple Slaves, or from Slave(s) and the Master itself will not resynchronize multiple times. Slaves may be configured to skip sending Level 2 signals to the network.

■ Level 3 (Independent Server)

Any machine running Domain Time II Server (except the domain controller with the PDC Emulator role, which must be a Master) can be configured as an *Independent Server*.

When an Independent Server corrects its clock, it broadcasts/multicasts a Level 3 advisory. Independent servers may be configured to skip sending Level 3 signals to the network.

An Independent Server does not actively participate in the cascading hierarchy with Masters and Slaves. Independent Servers acknowledge Level 1 cascade signals from the Master, but do not act upon them. Independent Servers ignore Level 2 cascade signals from Slaves.

■ Level 4 (Client)

A client both listens for cascade signals and sets its own time independently based on its timing settings.

If a client is in automatic mode, it uses discovery broadcasts at startup to determine if any Level 2 (Slave) machines exist on the local subnet. If so, the client enters normal operating mode, and will synchronize its clock upon receipt of a Level 2 cascade signal.

If no Level 2 machine is found (perhaps because all Slaves are currently offline, or the client is not connected to the network), the client enters pessimistic mode. In pessimistic mode, the client listens for all cascade and advisory signals and responds by synchronizing its time with whatever machine sent the cascade signal and taking the following actions:

If the cascade signal comes from a Master Server, the client assumes that the network has only a Master and clients. From that point on, the client will ignore Level 3 cascade signals (Independent Servers). This is called Master-only mode. Master-only mode converts to normal mode upon receipt of the first cascade signal from a Slave.

- If the cascade signal comes from a Slave Server, the client assumes that Slaves are now present, and from that point on ignores both Level 1 (Master) and Level 3 (independent server) signals. This is normal mode.
- If the cascade signal comes from an Independent Server, the client will sync with the Independent Server, and assume the network has neither Master nor Slaves. The client continues in pessimistic mode until a Master or Slave signal is seen.

Note that this procedure allows the time hierarchy to automatically collapse levels so that clients respond only to the next-highest level at any time. If a Slave comes online after the client is started, the client will note this fact and move from Master-only mode to normal mode immediately. If, while in normal mode, no Slave can provide the time, the client will automatically move into Master-only or pessimistic mode, as needed.

This hands-free configuration allows you to have any mix of Master, Slaves, and independent servers on your network, any of which may be working or not at any given time, and still use Domain Time II's hierarchy to (a) limit network traffic, and (b) ensure quick, uniform updating of all levels when the highest-level time source is updated.



Settings on this page control how often Slaves synchronize their time with the Master.

Note: This property page will only appear on the Master Server. Slaves will always use these settings. See [Domain Role](#) for more information on the Master, Slave, and Independent Server roles.

Check Interval when able to get and correct the time

Variable - check as often as needed to maintain approximately milliseconds sync with server

Fixed - check once every
DaysHoursMinsSecs

Variable - check as often as needed to maintain approximately milliseconds sync with server

When this option is selected, Domain Time will automatically adjust how often it synchronizes with time sources to attempt to keep the clock within the threshold limit you set.

The wait period between time checks is called the *window size*. You can see the window size Domain Time is currently using by examining the [Text Log](#). Domain Time will adjust the window size based on how accurate previous time checks have been. If previous time corrections have been small the window size will be increased, and vice versa. The range of adjustment for the window size is between 15 seconds to 2 hours. On most machines, it will average between 10 - 30 minutes.

The **Variable** scheduling method is intended for use on machines with relatively constant clock drift and moderate accuracy requirements (where the acceptable tolerance for clock drift is more than ~25ms). This method strikes a good balance between maintaining the target accuracy while minimizing network traffic.

Variable is not a good selection if the machine is under heavy and/or variable load that causes the clock to drift by significant amounts on an irregular basis. Since Domain Time may select a large window size if the clock on the machine has been well-behaved, anything that causes sudden clock drift during the Window period between checks can cause the clock to drift outside the specified threshold before the next correction. If this describes your machine, you should use the **Fixed** schedule instead.

Since having highly-accurate time at all times is usually more critical on Servers, you will probably want to choose a fixed schedule on your Servers.

Fixed - check once every

DaysHoursMinsSecs

If this option is selected, Domain Time will synchronize regularly on the schedule you specify.

This method is a good choice when you want to discipline the clock to stay within very tight synchronization tolerances. It's also the best choice for machines with highly variable load, poor timekeeping hardware, or any other issue that causes significant clock drift.

You should check the time often enough to keep your Slaves within your desired accuracy.

Check Interval when getting the time fails

If Domain Time cannot
obtain the time, it should
try again every: DaysHoursMinsSecs

If Domain Time cannot obtain the time, it should try again every:

DaysHoursMinsSecs

sets how often Domain Time will retry its time sources if it is unable to successfully obtain the time. You will probably want to check more often than during normal operation (unless you're already using a frequent synchronization schedule) to reacquire the correct time quickly when your time source(s) become available.

The same caution about synchronizing too often against your time sources (discussed above) applies.



These synchronization settings are recommended to Clients by the Master and Slaves:

Note: This property page will only appear on the Master Server. See [Domain Role](#) for more information on the Master, Slave, and Independent Server roles.

These settings will only take effect on clients if:

- The **Recommend timings and correction limits to clients that ask for guidelines** checkbox is checked on the Master Server's [Recommendations](#) page.
- The **Accept server's recommended settings (if provided)** checkbox is checked on the Client's [Timings](#) page.

Check Interval when able to get and correct the time

Variable - check as often as needed to maintain approximately milliseconds sync with server

Fixed - check once every
DaysHoursMinsSecs

Variable - check as often as needed to maintain approximately milliseconds sync with server

When this option is selected, Domain Time will automatically adjust how often it synchronizes with time sources to attempt to keep the clock within the threshold limit you set.

The wait period between time checks is called the *window size*. You can see the window size Domain Time is currently using by examining the [Text Log](#). Domain Time will adjust the window size based on how accurate previous time checks have been. If previous time corrections have been small the window size will be increased, and vice versa. The range of adjustment for the window size is between 15 seconds to 2 hours. On most machines, it will average between 10 - 30 minutes.

The **Variable** scheduling method is intended for use on machines with relatively constant clock drift and moderate accuracy requirements (where the acceptable tolerance for clock drift is more than ~25ms). This method strikes a good balance between maintaining the target accuracy while minimizing network traffic.

Variable is not a good selection if the machine is under heavy and/or variable load that causes the clock to drift by significant amounts on an irregular basis. Since Domain Time may select a large window size if the clock on the machine has been well-behaved, anything that causes sudden clock drift during the Window period between checks can cause the clock to drift outside the specified threshold before the next correction. If this describes your machine, you should use the **Fixed** schedule instead.

Fixed - check once every

DaysHoursMinsSecs

If this option is selected, Domain Time will synchronize regularly on the schedule you specify.

This method is a good choice when you want to discipline the clock to stay within very tight synchronization tolerances. It's also the best choice for machines with highly variable load, poor timekeeping hardware, or any other issue that causes significant clock drift.

You should check the time often enough to keep your Clients within your desired accuracy.

CAUTION: Take care with this setting. Many time servers have Denial-of-Service (DOS) protection to prevent abuse. Issuing too many time requests in a row to one server over a short period of time can cause your machine to be locked out or even be permanently blacklisted.

This problem can be exacerbated if you have opted to collect multiple samples per time source. See [About Time Samples](#).

[Check Interval when getting the time fails](#)

If Domain Time cannot
obtain the time, it should
try again every: DaysHoursMinsSecs

If Domain Time cannot obtain the time, it should try again every:

DaysHoursMinsSecs

sets how often Domain Time will retry its time sources if it is unable to successfully obtain the time. You will probably want to check more often than during normal operation (unless you're already using a frequent synchronization schedule) to reacquire the correct time quickly when your time source(s) become available.

The same caution about synchronizing too often against your time sources (discussed above) applies.



These correction limits are recommended to Clients by the Master and Slaves:

Note: This property page will only appear on the Master Server. See [Domain Role](#) for more information on the Master, Slave, and Independent Server roles.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings may adversely affect your clock accuracy or even prevent clock corrections entirely.

These settings will only take effect on clients if:

- The **Recommend timings and correction limits to clients that ask for guidelines** checkbox is checked on the Master Server's [Recommendations](#) page.
- The **Accept server's recommended settings (if provided)** checkbox is checked on the Client's [Timings](#) page.

AS OF v5.2.b.20170922, the Minimum Correction setting is obsolete and has been removed from the applet

The documentation in this section is for reference purposes on older versions only. If you still need to change this value, please see the [Client Settings](#) registry key on the Clients themselves.

Minimum Correction

Variances smaller than milliseconds will not cause a correction (unless overridden below)

This setting controls the minimum amount the clock must be off from the time source(s) before it is corrected during a time check.

If you are using a variable synchronization schedule (see the [Timings](#) page), you will probably not want to incur the extra overhead of correcting the clock if the variance found during a time check is smaller than your selected synchronization target.

For example, if you have instructed Domain Time to aim for a target variance of 25 ms, you do not need to make a clock correction if the detected clock variance during a time check is only 10 ms. Making a correction in this case would only result in extra processing overhead and clock slewing without much affecting your overall accuracy. So, in general, if you are using variable targeting, you will want to set this value to be the same value or less than your target variance.

However, if you are using a fixed sync schedule, you will want to be sure the clock is corrected to the maximum accuracy on every synchronization. In that case, this value should be set to 1 millisecond. This also enables Domain Time's high-precision sub-millisecond alignment function, so that any variances detected that are less than 1 millisecond will have the clock aligned to match, giving you an added order of magnitude of possible accuracy.

Maximum Correction

No maximum correction

Variances larger than will not cause a correction (unless overridden)
HoursMinsSecs

This setting controls the maximum variance that should be corrected during a time check.

This setting provides a vital sanity check to prevent wild time changes in the event your time source(s) provide a rogue time value (such as sometimes occurs when bounds limits are exceeded or error conditions occur in time clocks or the network).

For example, assume you have restricted this value to not allow corrections for variances larger than 2 hours. If a time source suddenly goes crazy and provides a time/date from 1980, the rogue time correction will be rejected.

The default setting for this value is fairly generous, so you may want to restrict this more in your environment. Do be careful to not restrict this value too tightly. If you have clocks on the network that drift significantly under normal circumstances without restarting (such some laptops do when resuming from sleep modes), setting this value too low may prevent them from ever correcting the clock until they are rebooted.

This setting also interacts directly with the clock slewing settings (which control whether corrections are made by slewing or stepping). See the [Clock Control](#) page for details.



Settings on this page control how Domain Time listens and sends traffic to the network.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change.

These settings are unique to each machine, and therefore cannot be saved or imported using the [Import/Export](#) utilities.

Network Listen

Address Family:	Both IPv4 and IPv6	Join DT2/NTP IPv4 multicast groups
		Join DT2/NTP IPv6 multicast groups
		Initiate rebind and resync if IP address changes
		Enumerate multicast interfaces during IPv4/IPv6 bind
Listen on all IP addresses available		Reply to multicasts using incoming interface if possible
Listen only on these addresses:		
Enter one IP address, hostname, or mask per line.		
You may use any combination of IPv4 and IPv6 addresses. The addresses you enter must exist and be permanently assigned to this machine		
You may also use CIDR notation to specify any/all addresses matching the mask you supply.		

Use these settings to tell Domain Time which IP addresses to use when listening for incoming network traffic and whether to join multicast groups.

IMPORTANT: Keep in mind that there can only be one network service listening on any one network port on any one IP address. Domain Time will attempt to control all the enabled time protocol ports on the selected IP address(es) when the service starts.

You may select which protocols are enabled on the [Serve the Time](#) property page.

DT2: 9909 UDP and TCP (cannot be disabled)

DT2 over HTTP: 80 TCP (configurable on the [Serve the Time](#) property page)

DT2 Status Monitor: 9911 UDP or TCP (configurable on the [Status Reports](#) property page)

NTP: 123 UDP

NIST Daytime: 13 UDP and TCP

TIME/ITP: 37 UDP and TCP

IPv6 requires operating system support, which is present by default in Vista or above, but must be specifically installed/enabled on XP. Domain Time will function in IPv4-only mode if IPv6 is not present. If both are present, you may choose which to use, or let the system figure it out.

Domain Time assumes your TCP/IP and Windows networks are configured properly, i.e. name resolution is functioning, rules are in place to permit traffic through switches, routers, and firewalls, any Active Directory/Domain structure is functioning correctly, etc.

Join DT2/NTP IPv4 multicast groups

Join DT2/NTP IPv6 multicast groups

These checkboxes control whether Domain Time will join multicast groups to listen for either NTP or DT2 protocol multicast traffic.

Multicasts may include not only time sync information, but client discovery packets from Manager and Audit Server. You should not disable this function unless you have a compelling reason to do so. Note this setting only applies to joining groups. You can control the multicast address and whether multicasts are sent on the [Broadcasts and Multicasts](#) property page.

Initiate rebind and resync if IP address changes

If this checkbox is checked, Domain Time will rebind interfaces if it notices that an IP address has been changed/added while the service is running. Use this option if your machine is likely to have its list of IP addresses change during operation.

Enumerate multicast interfaces during IPv4/IPv6 bind

Instructs Domain Time to include all addresses capable of multicast when it binds to the network interfaces. Check this box if the machine is multihomed.

Reply to multicasts using incoming interface if possible

Use this option to reduce unnecessary multicast traffic. When checked, and if the machine is multihomed and listening using a specified list of IP addresses (see below), Domain Time will attempt to reply to multicasts (such as PTP messages) over the same interface on which they were received. If unchecked, Domain Time will attempt to reply on all known interfaces. This feature has no effect if the **Listen on all IP addresses available** option is selected. Check this box only if the machine is multihomed and you have addresses listed in the **Listen only on these addresses:** box.

Listen on all IP addresses available

Listen only on these addresses:

By default, Domain Time attempts to listen for network traffic on all available interfaces. However, you can restrict this to specific IP addresses and/or address ranges if necessary.

IMPORTANT: Be sure to include the hostname **localhost** in your list of IP addresses to ensure that foreground Domain Time processes (such as the Domain Time applet) can continue to talk to the Domain Time service correctly.

As of v5.2.b.20190701, you may enter comments in the specified listen-only list. Comments are defined as text following a hashtag or semicolon. (If the hashtag or semicolon is the first character, the entire line is considered a comment.) For example, you may use this syntax:

```
; These are our subnets:
172.16.13.0/24 # main network
192.168.33.0/24 # internal network
```

Comments in the list (other than commenting out an entire line) are not backward-compatible with previous versions of Domain Time, so don't use them in templates until all of your machines have been upgraded.

Network Send

DT2:

NTP:

TIME/ITP:

Use blank (or zero) to mean the system should choose a random source port

Use this section to specify a fixed source port for time protocol traffic.

These settings should be left blank unless you have a specific requirement to send traffic from a specific source port. If you do need to set this port, you must not select the same port that is used for listening.

Note: Take care in assigning this port to avoid conflicts with any other port that may be used by any other service. In particular, you should not assign a port number in the ranges Windows will use for ephemeral source port assignment.



These are the broadcast and multicast addresses used by Domain Time Server.

CAUTION: The default settings on this page are correct for most situations. Only make changes if you are sure you need them and you fully understand the effects of the change.

Broadcasts and Multicasts are used for a variety of purposes by the various Domain Time components. For example:

- Server uses them to send cascades and advisories
- Server uses them as the listen addresses for IPv4 and IPv6 multicast requests
- Server uses them to send broadcast/multicast time (using the DT2, NTP and PTP protocols)
- Tools that don't have their own settings (for example, dtcheck.exe) use them for discovery and testing
- Clients can use them to discover DT2, NTP, and PTP time sources
- Clients use them as the listen addresses for IPv4 and IPv6 multicast requests

The settings on this page control whether Domain Time will enable multicasts and broadcasts and what addresses to use for these functions. Note, Domain Time joins multicast groups by default even if these settings are disabled. Also, multicasts will always be used if you enable PTP. Multicast group join settings are set on the [Network](#) property page.

IPv4 Broadcasts

Enabled	Broadcast address:
---------	--------------------

This section configures the Domain Time IPv4 Broadcast address.

By default, Domain Time Server will use broadcasts only on subnets local to this machine. If you want to reach machines on remote subnets, you should enable multicast (see below).

IMPORTANT: Disabling all broadcasts on this page may prevent your time hierarchy from operating correctly. If you are looking to disable a particular type of broadcast (such as NTP broadcasts) you are better off disabling that specific broadcast type on its property page instead. See the [Serve the Time](#) and [Domain Role](#) property pages for details.

IPv4 Multicasts

Enabled	TTL:
	DT2 multicast address:
	NTP multicast address:

This section configures the Domain Time IPv4 Multicast address.

The **TTL** (Time-to-live) entry sets how many router hops a multicast should make when propagating through your network. Choose a value that allows your multicasts to reach all of your subnets. This is usually the only setting you should change in this section.

Changing the multicast address listed here is usually an error.

IMPORTANT: Disabling all IPv4 Multicasts on this page may prevent your time hierarchy from operating correctly. If you are looking to disable a particular type of multicasts (such as NTP multicasts) you are better off disabling that specific multicast type on its property page instead. See the [Serve the Time](#) and [Domain Role](#) property pages for details.

IPv6 Multicasts

Enabled

Hop Count:

DT2 multicast address:

NTP multicast address:

This section configures the Domain Time IPv6 Multicast address.

The **Hop Count** entry sets how many router hops a multicast should make when propagating through your network. Choose a value that allows your multicasts to reach all of your subnets. This is usually the only setting you should change in this section.

Changing the multicast address listed here is usually an error.

IMPORTANT: Disabling all IPv6 Multicasts on this page may prevent your time hierarchy from operating correctly. If you are looking to disable a particular type of multicast (such as NTP multicasts) you are better off disabling that specific multicast type on its property page instead. See the [Serve the Time](#) and [Domain Role](#) property pages for details.



Settings on this page control the Domain Time Security settings.

Denial of Service (Flooding) Protection

DoS Protection Enabled

If any one machine sends more than requests in a -second period, ban
for: seconds

Auto-extend ban if abuse continues while IP is banned

Domain Time II has automatic protection against Denial-of-Service (DoS) disruption caused by intentional or accidental flooding of the network.

Any system that exceeds the DoS traffic thresholds you specify here has its access automatically blocked for a period of time.

Use the **Auto-extend ban if abuse continues while IP is banned** option if you have persistent bad actors whose bans expire, only to be re-blocked. You can also block them by IP address (see below).

Note: Even legitimate traffic can be blocked if it occurs too frequently. Take care that time sync requests from any individual machine or any tools that send repeated inquiries/commands to this machine do not exceed your DoS threshold.

Access Permissions

No restrictions IP ranges

Permit only listed range(s)

Deny any in listed range(s)

First IP in range

Last IP in range

Allow Domain Time II Manager to change the time zone on this machine

Auto-Manage Windows Firewall

Your time service can potentially be degraded by responding to audit inquiries, sync triggers, and/or time requests from clients or servers on other network subnets over which you have little control. For example, this can happen if your Domain Time Server is accessible from a public network and many other users discover and start to use your server as a time source.

To prevent this kind of problem, you may specify whether Domain Time should accept or reject time protocol traffic from certain IP addresses. You can specify whether to **Permit** or **Deny** traffic from multiple ranges of addresses. This allows you to easily restrict your incoming traffic to only the intended machines.

If you wish to permit or deny a single IP address, enter it as both the First and Last IP address in the range.

Allow Domain Time II Manager to change the time zone on this machine

When checked, you may change the timezone on this machine remotely from Manager.

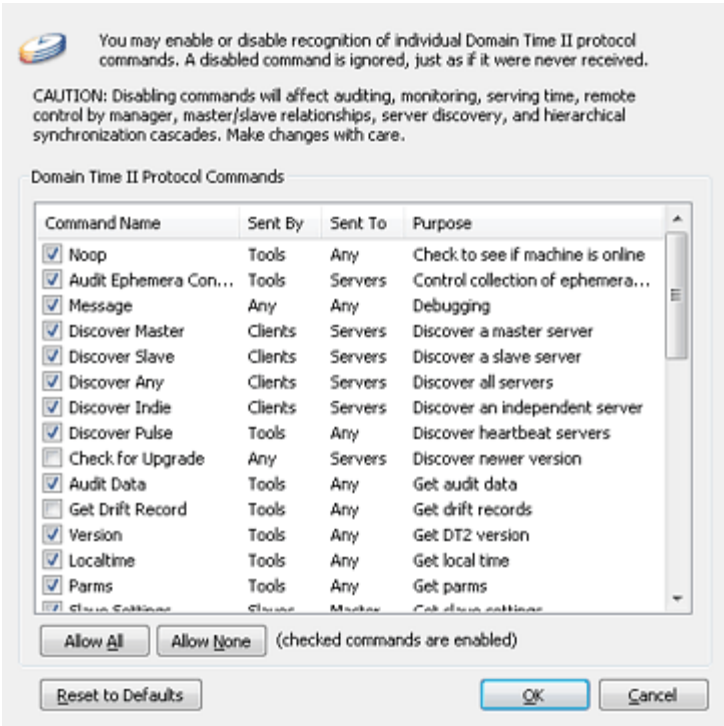
Auto-Manage Windows Firewall

As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to

the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

Command Restrictions

When you click on the [Command Restrictions](#) button you'll be presented with the *Command Restrictions* dialog window. You can use these settings to restrict what kind of Domain Time II control and sync messages your server listens for on the network.



Command Restrictions Dialog [\[Click for larger size\]](#)

The default protocol restriction settings assure both maximum functionality and a high degree of security; in most cases you will have no need to adjust them from the defaults. Domain Time II components communicate with each other primarily through directed communication, and are therefore highly resistant to spoofing and other malign interference.


The Domain Time II protocol command restriction capability is intended for use by system administrators in environments where an extra level of security is required, such as running a Server on the open Internet. Using the restrictions list, you can determine exactly what Domain Time II protocol command messages the service is allowed to listen for. Think of the command restriction list as an application-level "firewall" allowing in only the desired Domain Time II commands and blocking any others. Keep in mind that the restriction list only affects incoming DTII protocol commands - outgoing commands are not affected.

Warning:

Disabling protocol commands can have unintended consequences on the operation of your entire time distribution network, including the prevention of cascade triggers and sync notifications, which may result in inaccurate clocks. Problems resulting from disabled protocol messages can be quite hard to troubleshoot later, particularly by the next system administrator after you. Make adjustments only if you understand and require them, and be sure you document the changes so you can maintain the consistency and smooth operation of your time network.



This page configures the symmetric keys used by authenticating network protocols such as NTP and DT2.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

As of version 5.1, Domain Time supports two methods of network packet authentication: Windows Authentication and Symmetric Key Authentication.

Windows Authentication

Windows Authentication refers to the proprietary authentication method Microsoft uses to validate time packets between domain member machines and domain controllers (DCs). As of v5.1, Domain Time fully supports integrated Windows Authentication for both serving and obtaining the time within a domain.

Important considerations (read the sidebar for more details):

- Windows Authentication only works within a domain. That is, all machines exchanging time using Windows Authentication must be joined to the same domain (or forest).
- Windows Authentication only works between domain members and domain controllers.

While the domain member getting the time may be any kind of machine, the machine providing authenticated time must be a DC. Only a DC can validate the request. Other machines will not know the necessary shared secrets.
- Windows Authentication is automatic, requiring no additional configuration on servers or clients.
- Domain Time's Windows Authentication works with NTP, DT2-UDP, and DT2-TCP protocols between Domain Time Servers and Clients. W32time only authenticates using NTP.
- Windows Time (W32time) clients using "NT5DS" mode (the default domain member setting) can get authenticated NTP time natively from Domain Time Server running on a DC.

NT5DS-mode W32time clients may also get authenticated time from DCs running Domain Time Client, however, W32time must be set to NoSync mode on the DC so that it provides the authenticate NTP timestamp.

Although Domain Time Client will keep the DC's local clock highly accurate, using the W32time service to provide the authenticated time will result in reduced accuracy to the clients.

- Domain Time Clients can obtain authenticated time from Windows DCs running only the W32time service. However, this is not recommended, since

Interaction with Windows Time (W32time)

Windows Time clients using NT5DS mode (the default) search the Active Directory hierarchy to find a server. They send a request for the time using their machine RID as the authentication key, and expect the returned timestamp to be authenticated by the server. Only a DC in the client machine's domain can provide this type of authentication.

Domain Time v4.x Servers provided for Windows Time clients by setting the W32time service's client portion to "NoSync" mode and allowing the W32time service's server portion to serve NTP directly. Although the timekeeping ability of W32time is poor, this approach allowed the DC running Domain Time to continue serving Windows Time clients. This workaround is no longer necessary.

As of v5.x, Domain Time provides integrated Windows authentication natively for both NTP and DT2 protocols. This means that W32time clients in NT5DS-mode can get their time directly from any Domain Time II Server running on a DC exactly as if getting the time from the Windows Time Service on that DC.

- the Windows Time NTP service does a poor job of keeping the DC's local clock accurate, and
- the W32time NTP server itself adds additional inaccuracy to the network time being served.

Additionally, Domain Time v5.x clients can obtain authenticated time from DCs running either the Windows Time service or Domain Time.

W32time in NT5DS-mode has distinct disadvantages:

- The W32time NTP Server is inaccurate, so even if the DC's clock itself is well-synchronized, the time being served may not be. W32time clients receiving the time add to the problem, since they are unpredictable as well.
- Other ntp clients (such as xntp) cannot synchronize with it.

Domain Time's NTP Server has none of these disadvantages. It can provide standard NTP (with or without NTP auth) at the same time it provides NT5DS-mode timestamps, and all at extreme accuracy.

It is therefore highly recommended you install Domain Time II v5.x Server on all DCs. You can install Domain Time v5.x Clients on a DC, but you will then need to enable W32time in "NoSync" mode to provide NT5DS-mode time if you have clients that need it.

Recommended settings:

Using Domain Time on your DCs is *highly* recommended.

For v5.x Server on a DC:

- Verify that the *NTP Server Enabled* checkbox is checked on the Domain Time II Server [Serve the Time](#) property page AND
- the *Windows Time mode*: dropdown on the Server's [Advanced](#) property page is set to **Disabled**.

For v5.x Client on a DC:

- the *Windows Time mode*: dropdown on the [Advanced](#) property page is set to **NoSync**.

Other considerations

■ Cluster Service

The Windows Cluster has a default startup dependency on W32time. It does not require the time service for any other purpose. Thus, the simple recommendation for installing Domain Time on clusters is to set the *Windows Time mode*: dropdown on the [Advanced](#) property page to **NoSync**, which allows the service to be running to satisfy the startup dependency, but allows Domain Time to set the cluster's clock.

However, you may replace the cluster's startup dependency if you want. After installing Domain Time Client (or Server) on the cluster, use RegEdit to navigate to the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clussvc

Change the **DependOnService** value (omitting the quotation marks) from "W32time" to "Domain Time Client" (or "Domain Time Server" if that's what's installed).

The cluster service will then wait until Domain Time has started before starting the cluster. You can then set the *Windows Time mode*: dropdown on the [Advanced](#) property page to **Disabled**.

■ Reliable Time Provider

DcDiag and other tools sometimes expect the Windows Time service to be running on DCs, even if it's not actually doing anything. These tools often depend upon the DC being flagged as a reliable time provider.

Starting with v5.x, Domain Time Server, when installed on a DC, sets the system flags to indicate the machine is serving time and is a reliable time source. The **DsGetDcName()** function will report Domain Time Server v5.x machines on DCs as both time servers and reliable time sources. Domain Time Server on a non-DC will not change the existing system flags.

You may override this behavior by editing the registry. In

HKEY_LOCAL_MACHINE\Software\Greyware\Domain Time Server\Parameters

edit (or create) a **REG_SZ (string)** value called "Set Reliable Time Provider" and set its value to either "True" or "False" (the English words, without the quotation marks). If this value is present and set to True, Domain Time Server will set the two flags even if it is not running on a DC. This configuration has no meaning for Active Directory, since only DCs are examined for the flags. Other tools, however, may benefit from knowing that a reliable time source is present. If this value is present and set to False, then Domain Time Server will not change the flags.

Symmetric Key Authentication

v5.x Domain Time Clients and Servers support Symmetric Key Authentication (hash of shared secrets). Domain Time supports MD5, SHA1 (as of version 5.2.b.20170922), SHA256 (as of version 5.2.b.20190331), and SHA512 (as of version 5.2.b.20190701). Older versions support MD5-only.

Domain Time Server and Client support symmetric authentication (using SHA1 and/or MD5) of client-server requests using NTP (version 3 and later; AutoKey is not supported), DT2-UDP, DT2-TCP, and DT2-HTTP protocols. Domain Time also supports broadcasting (both NTP and DT2-UDP) with a shared key and hash. SHA256 hashes are only valid for use with PTP v2.1. Clients configured with the same key validate packets from the sending server by comparing the computed hash.

Note: Although v5.2.b.20190701 added support for SHA512, this option is reserved for future use. You should not create SHA512 keys. If an SHA512 key exists in the keyring of an older version of Domain Time, it will be (mis)interpreted as a very long MD5 key.

SHA1 keys are always exactly forty hex characters long. MD5 keys are ASCII text; different implementations of the NTP daemon have allowed different maximum key lengths. In general, an MD5 key should be composed only from 7-bit ASCII-printable text, excluding space, tab, and the # character. MD5 keys should be at least 8 characters long, and should not exceed 20 characters. Some versions of NTP daemons allow lengths of 32, while others have a maximum of 8 or 16. You will need to choose MD5 keys that are interoperable with all of your various devices and daemons. SHA256 keys must be exactly 64 hex characters long. SHA512 keys are a 128-byte hex string, corresponding to a 64-byte key.

The Keyring

Symmetric Keys are kept in a list containing the Key number and the Key secret (password). This list is also known as a *keyring*.

The keyring may contain a combination of *trusted* and *untrusted* keys.

A trusted key means the key is available to be selected by the component, but the trusted key is not active until its key number is selected when configuring a unicast time source in the [time sources list](#) (or by using the **Broadcast/multicast key** section of this page for broadcasts/multicasts). Untrusted keys are ignored.

Checked items are "trusted" or active keys; only trusted keys will be used:

Key	Type	Password
<input checked="" type="checkbox"/> 1	MD5	DomainTimeII
<input checked="" type="checkbox"/> 2	MD5	TTnts200
<input checked="" type="checkbox"/> 3	SHA1	97d870fe734e05bd449d476b9fbeb3b332234003
<input checked="" type="checkbox"/> 9909	MD5	greyware

There are various ways to configure the keyring on Domain Time II components:

- Master and Independent Servers
 - manually using the Control Panel applet (on this page)
 - importing the keys from a properly-formatted keyring (ntp.keys) file

- importing a previously-exported Domain Time .reg file.
- Shared secrets are automatically replicated between Domain Time Masters and Slaves. No configuration of the Slaves is necessary.
- Clients
 - manually using the Control Panel applet
 - importing the keys from a properly-formatted keyring (ntp.keys) file
 - importing a previously-exported Domain Time .reg file.
 - using Windows Group Policies (See the [Active Directory](#) page.)


PTP v2.1 Security Parameter Pointer (SPP):

In addition to shared SHA256 hashes, PTP v2.1 requires that Masters and Slaves use the same SPP value to be able to validate v2.1 Authentication TLVs. The SPP stored in the keyring may either be zero (which acts like a wildcard) or must match what the grandmaster sends. If there is a potential for your Slaves to discover more than one Master (such as with a fallback server), we recommend you use the wildcard setting (0) to avoid synchronization failure if each server has a different SPP.

Import/Export the Keyring (keys file)

Click the **Import/Export** link in the *Symmetric Keys* section, which brings up a dialog where you can import or export a keyring (*.keys) text file to share among your devices.

This function is very useful if you are sharing a keyring file with other systems running a daemon such as dtlinux (dtlinux.keys), chronyd (i.e chrony.keys), or ntpd (i.e. ntp.keys).



Secrets Import/Export Dialog [\[Click for larger size\]](#)

Hints:

If you have multiple time sources, each may have its own set of symmetric keys. Be sure to import all the keys from all time sources into Domain Time.

If you are signing outgoing PTP v2.1 delay requests, all of your grandmasters should be configured to accept the same KeyId for incoming delay requests.

When possible, be sure all of your time systems are working correctly before enabling authentication. Authentication requires a correct setup on both ends of the connection, and changes at either end can cause a previously-working connection to fail. Disabling authentication temporarily should always be one of the first steps when troubleshooting a connection issue.

As of v5.2.b.20190701, you may use Manager to push out the symmetric keyring to multiple machines at once. See the [Reset Keyring](#) command.

Broadcast/Multicast Key


Broadcast/Multicast Key: 1

This dropdown selects the trusted key to be used when signing Broadcast or Multicast time packets. Note this refers specifically to the "heartbeat" type of time packets sent to the network on a fixed schedule, as configured on the [Serve the Time](#) page.

As with normal Symmetric Key authentication, Clients receiving the broadcast/multicast must also be using the same authentication key to decode the packet.



The Logs and Status page contains the settings for the Domain Time service text log.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Logs are kept in the `%SystemRoot%\System32\` folder. There are at least four main log files collected when the service is running:


- **domtimes.log**

This is the currently active text log file.

If log archiving is enabled (see below), additional archived log files will be created using a `domt i mes. YYYYMMDD. l og` naming scheme (i.e. `domtimes.20090928.log`).

- **domtimes.log.startup**

A detailed text log of the service startup process. Only data from the latest startup is included.

To view these two text logs, click the  button on the applet, which launches the Domain Time Log Viewer.

- **drift.dt**

A binary file containing information on each time check/correction made using the NTP and DT2 protocols, or the aggregate of corrections made by PTP (if enabled) during the check interval configured on the [Timings](#) page. Drift logs can also collected remotely by Domain Time [Audit Server](#).

To view this log, click the  button on the applet, which launches the Domain Time Drift Log Viewer.

- **driftptp.dt**

IF PTP is enabled, a binary file containing information on each PTP Sync. PTP Drift logs can also collected remotely by Domain Time [Audit Server](#).

To view this log, click the [Graph](#) link on the [Obtain the Time](#) property page, which launches the Domain Time Drift Log Viewer.

Text Log

Log Level: Information

Max Size: KB (use zero to mean unlimited size)

Enable lazy write delay of up to seconds (range 1-600)

Include info-level timeset success messages in warning and error-level logs

Include client accesses (requests for time and control messages)

Enable UDP packet tracing Enable TCP packet tracing

Enable Time Change Event Monitor

Enable NTP4 peerstats Enable NTP4 loopstats

NTP Stats Folder:

This section selects the properties of the **domtimes.log** service text log.

The **Log Level** drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

This switch will only disable the **domtimes.log** file. Other system logs, such as **domtimes.startup.log** and **drift.dt** cannot be disabled.

- **Errors**

Only messages marked as Errors will be logged

- **Warnings**

Logs will include Errors and Warnings

- **Information**

Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.

- **Trace**

Includes all of the above, plus detailed information on time setting and time sample analysis.

- **Debug**

Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

When you select **Debug** level, the button becomes enabled. Clicking this brings up a dialog where you can select exactly which type of debug messages to include in the logs. This allows you to limit the size of the logs while troubleshooting a particular issue. Please re-enable all messages if submitting a log for analysis.

Max size: sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to troubleshoot any issues that may arise.

Enable lazy write delay of up to seconds (range 1-600)

This value specifies the maximum amount of time to wait before flushing data to the log.

Logging large amounts of information on an underpowered or busy machine can generate a great deal of overhead, which can adversely affect system performance and diminish timing accuracy. When enabled, Domain Time will try to buffer log data in memory until the delay period is reached instead of attempting to write all events to the disk in real time. This may provide some "breathing space" for the disk system to process outstanding writes that may accumulate from constant log activity. If the buffer fills before the limit is reached, it will be flushed to disk even if the full wait period has not expired.

Include client accesses (requests for time and control messages)

When checked, all client requests made of this machine will be logged. This includes queries for time or other information such as statistics, auditing, or status requests.

CAUTION:

Enabling this option can generate a large amount of logging data and system overhead if you have a lot of

clients synchronizing with this server. Examine the log after the server has been running for a while to see if this option generates an onerous amount of data.

Enable UDP packet tracing

Enable TCP packet tracing

These checkboxes cause Domain Time to log additional useful details about the time packets being used by Domain Time. The output is similar to the output from a packet analyzer, showing you the actual packet contents and payloads.

CAUTION:

Enabling this option can generate a large amount of logging data and system overhead. You should enable these only when actively troubleshooting network issues and for short periods.

Enable Time Change Event Monitor

Tells Domain Time to use Windows auditing to help identify the user or process responsible for changing the system clock.

When checked, Domain Time will attempt to enable the Windows audit category "System" success auditing, and then watch the security event log for events pertaining to date/time changes made by programs other than Domain Time. If such an event is detected, Domain Time will parse the security event log entry and issue a warning in its own log. The warning will show the user and process that changed the time, and by how much (if the information is available). These warning messages are informational only, and should be enabled only to help track down environments where another user or process is interfering with the system clock.

If the audit policy for the machine is controlled by a group policy, then Domain Time's change to the audit policy will succeed, but only until the next group policy refresh is applied. If you are using this feature in a domain, either undefine the group policy setting (Local Policies/Audit Policy/Audit system events) or set it to enabled for success events.

Note: This option is complementary to, but independent of the Clock Change Monitor function, which resets the system time if unauthorized changes to the system clock are detected. See the description of Clock Change Monitor on the [Advanced](#) page.

CAUTION:

This option causes additional system overhead and uses additional memory and resources. You should enable this option only if you are experiencing rogue time changes on your system and are having difficulty identifying the cause. You should not run with this option enabled in normal operation.

Enable NTP4 peerstats

Enable NTP4 loopstats

As of version 5.2.b.20170101, these checkboxes enable creation of ntpd-style peerstats and loopstats statistic files. See the [ntpd Compatibility](#) page for details.

When enabled, the path where the files are collected will be displayed in the *NTP Stats Folder:* field. Note: NTP must be enabled on the [Serve the Time](#) property page in order to collect these stats.

[Text Log Archiving](#)

Log Roll: Daily at Midnight

Delete old logs

Keep up to old logs

Domain Time can automatically archive the text log on a daily, weekly, or monthly schedule.


When the log is archived, all existing log events in the **domtimes.log** file will be written to an archive file named

domtimes.YYYYMMDD.log (i.e. domtimes.20090928.log) and the current log file will then be cleared to accept new data.

You can choose how many archived log files to keep on the machine. When the indicated limit is reached, the oldest log file will be deleted.



This page specifies whether Domain Time service activity will be echoed to the Windows Event logs.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Some levels of logging can create a significant amount of data. The Windows Event logs can be difficult to read, or the Event Log process may even have problems recording all the data when large amounts of log activity are generated.

You should consider using only the **Error** level when using the Event Logs unless you generate a very small amount of logging data overall. In general, Text or Syslog logging is a better choice for keeping more detail.

Event Viewer

Log Level: Information

The **Log Level** drop-down chooses what type of entries to include in the Event logs. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**
Domain Time will not log events to the Windows Event Logs.
- **Errors**
Only messages marked as Errors will be logged
- **Warnings**
Logs will include Errors and Warnings
- **Information**
Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.
- **Trace**
Includes all of the above, plus detailed information on time setting and time sample analysis.
- **Debug**
Includes all available information provided by the service.

Warning:

The amount of data generated by Debug logging can easily overwhelm the Event Log system. Use Text or Syslog logs for debugging instead.

Event IDs


Event ID	Category	Meaning of Event ID
1000	Success	Generic success (examine text for details)
1001	Warning	Generic warning (examine text for details)
1002	Error	Generic error (examine text for details)
2001	Success	Time set successfully
3000	Warning	Unable to set the time from any source
3001	Warning	Protocol error while trying to obtain the time
3009	Warning	Clock-Change Monitor is disabled
4009	Error	Clock-Change Monitor trigger detected

The listed Event ID codes can be used to filter for Domain Time events in the Event Viewer. The Event Source field on the Event records will be **Domain Time Server**.

If you're considering using the Event Viewer for live system monitoring purposes, you may want to investigate the [SNMP Traps](#) function or [Service Status Monitor](#) to be more efficient.



You may choose to have Domain Time send service activity events to Syslog servers.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Syslog

Log Level: Information

Server(s):

You may list up to eight IPv4/IPv6 target addresses. Separate with spaces

☐ RFC 3264 format

☐ RFC 5424 format (UDP only)

☐ Include timeQuality structured data element

☐ Send each PTP data point to syslog (trace or debug level only)

The **Server:** field should contain the DNS Name or IP address of the Syslog Server(s).

As of v5.2.20170922, you may list up to eight targets on the Syslog Server line. Older versions only support a single target. Separate targets with a space.

Important: If you will be assigning this value using [Templates](#) or [Active Directory Policies](#) to any Server or Client version older than v5.2.b.20170922, the first entry of a multiple target list **MUST** be an IPv4 address, otherwise the older version will read the field incorrectly. For maximum backward compatibility with older versions of Domain Time, avoid using a DNS name if you list more than one target. If all of your machines have been upgraded, then you may use either DNS names or IPv4 addresses for your targets.

The **Log Level** drop-down chooses what type of entries to include in the Syslog logs. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**
Domain Time will not log events to the Syslog server.
- **Errors**
Only messages marked as Errors will be logged
- **Warnings**
Logs will include Errors and Warnings
- **Information**
Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.

- **Trace**

Includes all of the above, plus detailed information on time setting and time sample analysis.

- **Debug**

Includes all available information provided by the service.

RFC 3264 format

RFC 5424 format (UDP only)

Include timeQuality structured data element

Added as of version 5.2.b.20171113. Use these radio buttons to select the RFC format that matches your syslog server. If the **Include timeQuality structured data element** checkbox is checked, the output will include timeQuality information, for example: `[timeQuality tzKnown="1" is Synced="1"]`

Send each PTP data point to syslog (trace or debug level only):


Added as of version 5.2.b.20180101. If checked, and if the syslog level is set to either trace or debug, then each PTP data point will be sent to syslog.

The format is `PTP sample offset ±0.0000000, mpd 0.0000000, source ipaddress` where 0.0000000 is that sample's delta and the current meanPathDelay. Syslog log collectors may parse for trace-level messages beginning with "PTP sample offset" to categorize these messages.

Caution: Enabling this output can create a large number of syslog messages. Enable only if you actually require this level of detail. You may want to enable lazy writes if you find the logging process is affecting your clock accuracy. See the [Logs and Status](#) page to enable lazy writes.



Domain Time II Server can send notifications of its status to Network Management Systems and other SNMP-capable monitoring devices using SNMP Traps.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Important: SNMP support depends upon wsnmp32.dll being present in the OS. All Windows versions except the initial release of Windows Server 2016 Nano Server have this installed by default. To install Domain Time on Nano Server, you must install the SNMP trap support .dll before installing Domain Time. See the [Nano Server FAQ](#) for more information.

SNMP Control

Enabled

Community:

Server:

Use this section to enable SNMP Traps. Enter the SNMP community name and password used by your Network Management System (NMS), as well as its DNS name or IP address. Your community name and password must match the one in use by the receiving system.

As of v5.2.b.20190331, you may also specify a port number to which SNMP traps are sent in the Server field. Examples: **snmp. mydomai n. com: 1214** or **[2002: 410: 1: 1: 2a0: 69ff: fe01: b0f4] : 2444**. If the port number is not specified, the default SNMP trap port of 162 will be used.

Best Practices for SNMP include using a unique community name and hard-to-guess password on production systems. The default community *public* should only be used for initial testing. Although Domain Time only sends outgoing trap information and is therefore not susceptible to SNMP remote control vulnerabilities, you should still be mindful of SNMP security for the benefit of your other SNMP devices.

SNMP Traps

Sync Successful	If variance exceeds	milliseconds
Sync Failed	Ignore variance alert for first sync	
Service Startups		
Service Shutdowns		

The settings in this section select which SNMP traps are sent by Domain Time.

SNMP v2 traps are generated whenever the selected event occurs. Keep in mind that SNMP Traps are sent via UDP, and are therefore not guaranteed to be delivered by the network.

Although useful for raising performance alarms or other monitoring functions, you should not depend upon the SNMP trap data for critical logging of time synchronization events, particularly if your logging is necessary for regulatory compliance. Use a product designed for more robust data collection, such as Domain Time [Audit Server](#) instead.

The **If variance exceeds** **milliseconds** setting lets you set a threshold value so that you can be alerted when the variance of any timecheck exceeds this value.

Warning:

Domain Time may generate a large number of timechecks (depending on your [Timings](#) settings) so it is very easy to swamp your monitoring system with alerts if you set this value too low. Also, it is normal even on the best-behaved networks for occasional timechecks to reflect a spurious large value due to latency or other network conditions.

It is therefore very easy to generate a large number of false alarms using this trap.

If you enable this trap, you should choose a threshold value that truly reflects a critical amount of clock drift to avoid unnecessary alarms for normal transient variances. If your monitoring software has the capability of further restricting alarms only after a certain number or percentage of traps have been raised, you can add extra protection against false alarms.

You may find the [Audit Server Notifications](#) feature to be a better way of immediately alerting you to poorly performing clocks than this trap.

Since the first timecheck on any system after the time service starts is likely to be quite large (due to the clock not having been set yet), the **Ignore variance alert for first sync** setting is highly recommended.


The Domain Time MIB File

Domain Time comes with a MIB file that you can use to compile on your SNMP monitoring system so that your traps are interpreted correctly. The MIB text file is generated when you click the [Generate MIB File](#) button on the Control Panel applet so you don't need to worry about locating it in some obscure installation folder or having online access.

Note: The MIB file generated here matches the version of Domain Time that's currently installed. Be sure to remember to update your SNMP Network Management Station(s) with the latest version of the MIB after any upgrades to Domain Time.



This page contains options for additional features to monitor the Domain Time service in real-time.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Audit Server Real-Time Alerts

Enabled

Primary Server:

Backup Server:

send to the backup only if primary is down

send to both primary and backup servers

Send reports using TCP (recommended)

Send reports using UDP

Always audit this machine

Never audit this machine

Do not change audited status

Domain Time Server and Client are designed to be centrally monitored and have their synchronization status recorded by Domain Time [Audit Server](#). This is usually done by Audit Server in a scheduled, background data collection process.

However, Domain Time II Server and Client can send an immediate notification packet to Audit Server containing feedback on the success or failure of each time check to provide real-time monitoring and alerting. Domain Time Manager on the Audit Server machine will display the collected data on its Real-Time Alerts panel. See the [Audit Server Alerts](#) pages for more information on setting up Real-Time Alerts.

If you want this function, check the **Enabled** checkbox and enter the DNS name or IP address of your Primary Audit Server.

As of version 5.2, you may specify addresses for both a Primary and Secondary server for redundancy purposes. Use the radio buttons to select whether to

send to the backup only if primary is down, or to

send to both primary and backup servers.

This function is designed to work well with the Audit Server "Hot Spare" standby mode functionality introduced in version 5.2.

Select whether to use TCP or UDP for the notifications. In general, TCP is more reliable than UDP since delivery of TCP traffic is given priority, whereas UDP can be delayed or dropped entirely by busy network hardware. However, TCP does require more resources on the network stack of the Audit Server to handle the mechanics of building and tearing down TCP connections. If you have many machines sending lots of real-time updates using TCP, it is possible to exceed your operating system's ability to handle the number of open network connections. If you run into those limitations, you may want to consider changing your alert notices to UDP.

CAUTION:

If you are using notifications on many machines that are set to synchronize frequently, the feature can generate a significant amount of network traffic toward your Audit Server. You will probably want to enable instant notifications only on critical systems where an immediate alert of time sync errors is desired.

Otherwise, you should use Audit Server's normal scheduled sync log collection instead, which is much more efficient. Audit Server's sync log collection can run in a background process to gather all sync records from audited systems with much less impact on overhead than active notifications. Audit Server can be programmed to provide alerts based on that data as well.

Audit controls

As of version 5.2.b.20110224, you have the option to have the Domain Time service control its own inclusion in or exclusion from the Audit Server Audit list.


Always audit this machine

Never audit this machine

Do not change audited status

These options override any existing settings on the Audit Server itself. You may use these functions to ensure that machines are (or are not) audited whether or not they are discovered by Audit Server.

Domain Time Service Status Monitor



The *Service Status Monitor* section controls whether Domain Time will provide a simple text response about its current status when asked by an application.

This monitor is provided to allow third-party applications a simple way to monitor the activity of the Domain Time service. The Status Monitor will respond to TCP or UDP requests on the specified port with a simple text string showing the current activity of the service.


Sample responses from the Status Monitor:

ACK Adjusting (Indp. Server 5. 2. b. 20130403R)

ACK Set 22 seconds ago (Indp. Server 5. 2. b. 20130403R)



This page gives you access to various advanced Domain Time II settings.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings may adversely affect your clock accuracy or even prevent clock corrections entirely.

Miscellaneous Options

- Enable Test Mode (if checked, time on this machine will NOT be corrected)
- Enable Clock-Change Monitor
- Force setting of CMOS clock at service shutdown
- Use software timestamping (if available and enabled)
- Enable advance scheduling of leap second corrections
- Signal resync if VM guest resumes from paused or saved state
- Truncate drift status records to milliseconds

Miscellaneous options include:

Enable Test Mode

Checking this box causes Domain Time to operate in all ways as it would in normal operation, except for actually setting the time or changing the machine's slew rate. This allows you to test or troubleshoot the server's ability to obtain and serve the time, but without actually changing the server's time.

Results of all operations are logged normally, so you can use the log in test mode to track down any communication or other synchronization issues. Note: Be sure to disable this option when you're through testing!

Enable Clock-Change Monitor

Domain Time's *Clock-Change Monitor* notifies Domain Time if another user or process attempts to change the time on this system.

When an unauthorized clock change is detected, Domain Time immediately re-synchronizes the time with its time source(s) and makes a warning entry in the logs. This prevents inadvertent or malicious tampering with the system clock.

This setting should always be enabled unless you are doing testing that requires you to change the system clock manually, either from the Windows Date/Time applet or from some other application.

Force setting of CMOS clock at service shutdown

Controls whether Domain Time should perform an API system call to write the current system time to the CMOS Real-Time Clock (RTC) on the motherboard each time the service stops.

On modern operating systems, the CMOS RTC clock is primarily used to provide something approaching the current

date/time to the operating system while booting until the operating system can take over timekeeping. The CMOS clock is subject to all manner of inaccuracies, and is therefore not used for timekeeping while the OS is running, nor is it updated often.

The CMOS clock can therefore go for long periods without having its time corrected, resulting in huge drift. By default, Domain Time will update the CMOS with the current time either when doing so doesn't cause a disruption to the operating system time (during stepped corrections) or just before shutdown so that the CMOS has its best chance to be accurate during the time the system is not running.

When this box is **unchecked** (the default), Domain Time writes the current time to the CMOS RTC clock:

- when making a stepped time correction.
- if the Domain Time service is running and it receives a system shutdown command from the operating system.

If the box is **checked**, then Domain Time will also write to the CMOS clock any time the Domain Time service is stopped, whether or not the stoppage is due to a system shutdown.

Although at first read this may seem desirable, there is a downside to writing to the CMOS clock if the machine isn't already being hard-set (stepped) or in the process of shutting down. The API used by the operating system to write to the system clock also immediately steps the system time to the same time as the RTC *but only at the resolution of the CMOS clock*. Since the RTC resolution varies on different machine, writing to the CMOS will cause the system clock to jump either forward or backward to the nearest increment of the RTC, which can mean an unpredictable jump of 1ms or more in the system time.

As a result, you should leave this switch turned off unless you need to force a CMOS update by manually stopping the Domain Time service.

Use software timestamping (if available and enabled)

As of v5.2.b.20190701, Domain Time can use the Microsoft NDIS software timestamping API to help compensate for network stack delays. Microsoft only offers this function on its most recent versions of the operating system (i.e. Server 2019 and updated versions of Win10). This option will be greyed-out if it is not supported on this OS.

Note this checkbox only controls Domain Time's ability to use software timestamping if present. Timestamping must also be enabled at the operating system level. You can do this by using Microsoft's own PowerShell scripts (see KB article linked below), or you may use the [DTCheck](#) utility. To use DTCheck, open an elevated command prompt and issue the desired command:

<code>dtcheck -swTimestamps</code>	displays current settings
<code>dtcheck -swTimestamps: Enable</code>	attempts to enable software timestamping
<code>dtcheck -swTimestamps: Disable</code>	attempts to disable software timestamping
<code>dtcheck -stats2</code>	shows statistics for NDIS timestamping
<code>dtcheck -interfaces</code>	shows which interfaces have software timestamping enabled

You'll need to restart the NICs (or reboot the machine) after enabling or disabling timestamping.

You can experiment with this setting to see if it improves your accuracy or not. The stack delay on most machines is minimal so the overhead of measuring it may outweigh the benefits, however, you may see improvement on very busy machines.

See [KB2019.708](#) for more information.

Enable advance scheduling of leap second corrections

Controls how Domain Time handles upcoming UTC leap second corrections.

NTP and PTP packets can contain a flag to indicate an upcoming UTC leap second. When this checkbox is enabled, Domain Time will apply leap seconds at 23:59:59 UTC on the last day of the month in which the leap occurs (typically June or December). If unchecked, leap seconds will be applied at the first timecheck following the leap.

Domain Time acquires pending leap second information only from NTP or PTP time sources. All queried NTP or PTP sources must agree that a leap is pending in order for Domain Time to schedule the leap. If the sources disagree, then the leap will be handled at the next timecheck after it occurs, and a warning notice that the leap indicators are inconsistent will be placed in the log.

Pending leap information is queried with each timecheck (NTP sources only), and maintained only while the Domain Time service is running. Restarting the Domain Time service will clear any pending leap second corrections. If the leap is still pending when the Domain Time service is restarted, it will be rescheduled for the appropriate time. If the leap occurs while the Domain Time service is stopped, the leap will be applied at the first timecheck after startup.

[Read more](#) about leap seconds.

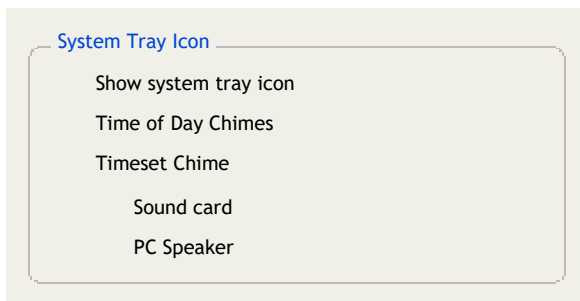
Signal resync if VM guest resumes from paused or saved state

Allows correction of clocks on virtual machines after resuming from pause/suspension.

Virtual machines are often paused/suspended, causing the clock to be incorrect when resumed. Domain Time may be able to sense a resumption by examining the Time Stamp Counter (TSC) and resync the clock. This setting is only useful on virtual machine guests. This option will be greyed-out if the machine is not a virtual guest.

Truncate drift status records to milliseconds

Delta values in the drift logs will be reported to the nearest millisecond if this checkbox is checked.



The Domain Time II System Tray applet (DTTRAY.EXE) is a foreground application that can load whenever a user logs into the system. When loaded, it will display as an icon in the System Tray.

The applet provides a number of very useful functions, including audio alerts and chiming, statistics, drift graphs, and a quick way to launch the various features of Domain Time installed on the machine.

The settings in this section determine whether or not to load the applet and which audio features are enabled.

Show system tray icon

This checkbox controls whether the [Domain Time System Tray](#) applet is loaded during login. If the icon is present in the System Tray, you can right-click it to choose from many additional features.

Note: The applet will unload if the Domain Time service is stopped. On XP and Server 2003, the applet will reload automatically when the service restarts. However, beginning with Windows Vista, Microsoft disabled the ability for background services to launch foreground programs, so on those systems you will need to either log out and back in or relaunch the applet manually. You can restart the applet manually by entering `dttray.exe` into the *Start -> Run* program field or at a command prompt.

Time of Day Chimes

The Time of Day Chimes feature plays sound files at specific times of the day, such as every 15, 30 45 minutes and on the hour.

This option will be unavailable if the **Show System Tray icon** checkbox on the Advanced tab is unchecked.

There must be a logged-in user and the Domain Time II System Tray icon must be present in the Windows System Tray for the chimes to play. You must also have installed at least one free Domain Time II [Chime Pack](#) for this feature to work.


See the documentation for the [System Tray Applet](#) for complete instructions and to download free Domain Time II Chime Packs.

Timeset Chime

Plays a sound whenever the Server successfully sets its time from its time source. If checked, the sound will play whether or not there is a logged-in user.

Sound card plays through the sound card if available

PC Speaker plays through the PC speaker



The settings in this section configure the Windows Time Service to co-exist with Domain Time.

Windows Time mode: Disabled

This drop-down box lets you determine how the Windows Time Service should behave on this machine. When the Domain Time II service starts, it will force the Windows Time service into this mode. The available options are:

■ Disabled

The service startup setting for Windows Time Service is set to Disabled. The Windows Time Service will not be allowed to run. This is the preferred setting for all machines except machines running Windows Cluster Service (see the *NoSync* description below).

Note: Some utilities on Windows Server will report spurious errors in the logs such as "Time not set for xxxxx seconds" when Windows Time is disabled on domain controllers. DCDIAG may also note the W32Time service is not running, but these warnings can be safely ignored.

■ NoSync

This mode makes sure the W32Time Server Provider portion of the Windows Time Service is running, but the W32Time Client Time Provider is disabled. In this mode, Domain Time II actually obtains the correct time and manages the local system clock; Windows Time merely answers NTP requests.

Note: On Domain Time versions prior to v5.1, *NoSync* was necessary either when Domain Time II Server was installed on a Windows Domain Controller to enable NT5DS mode to function properly, or on Cluster Servers to satisfy their startup dependency. However, as of version 5.1, Domain Time Server can supply NT5DS-mode clients with the necessary authenticated timestamp, so *NoSync* is no longer required on DC's. *Disabled* is the preferred Windows Time mode setting on all machines, except when used on Clusters (see below).

In *NoSync* mode, W32Time will attempt to bind the NTP port 123 UDP. Therefore, this mode will conflict with Domain Time II Server's NTP functions. If you enable *NoSync*, **all** of the NTP protocol boxes must be unchecked on the Domain Time Server [Serve the Time](#) page.

Although machines running in *NoSync* mode will provide NTP to NT5DS-mode machines, the accuracy of the timestamps provided will be constrained by the native inaccuracy of the Windows Time service. Also, non-Windows systems may have difficulty synchronizing with the machine, since W32Time is not compatible with many NTP daemons. If possible, we recommend you use the native Domain Time II Server NTP functions instead.

Cluster Service

Some versions of the Windows Cluster Service (i.e. Win2003 and earlier) have a default startup dependency on the w32Time (Windows Time) service. Cluster Server does not appear to require the time service for any other purpose. Thus, the simplest recommendation for installing Domain Time on clusters that have the W32Time startup dependency is to set the *Windows Time mode*: dropdown to **NoSync**, which allows the W32Time service to be running to satisfy the dependency, but allows Domain Time to set the cluster's clock.

However, given the limitations of the W32Time NTP service and the fact that only one service may serve NTP on a machine; if you will be serving NTP from the Cluster, you may want to remove or replace the cluster's startup dependency on W32Time so that you can disable Windows Time and use Domain Time II Server's native NTP services instead.

CAUTION: The following registry change is provided for your information. We're not aware of any issues with removing the dependency, but you should defer to Microsoft's guidance. Be sure to test any changes thoroughly on non-production servers before implementing on production systems.

To remove the W32Time startup dependency (if present):

After installing Domain Time on the cluster, use RegEdit to navigate to the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clussvc

The *DependOnService* value lists all services on which the Cluster Service startup depends. If the **w32Time** entry is present in the list, change it to **Domain Time Server** and save your changes. The cluster service will then wait until Domain Time has started before starting the cluster.

If the **w32Time** entry is not present in the list, there is no startup dependency on Windows Time in your version of Cluster Server and you do not need to make any changes to this registry value.

Once you have verified there is no startup dependency on W32Time on all nodes of the cluster, you can then set the Domain Time *Windows Time mode*: dropdown list to **Disabled** and restart the Domain Time service.

■ Not Touched

The existing configuration of the Windows Time Service is not changed. In this mode, Windows Time will operate however it is currently configured.

This mode may severely conflict with Domain Time unless:

- The "Do not set this machine's time" radio button is selected on the [Obtain the Time](#) page.
- The "Refuse to serve time until this machine's time has been set" checkbox has been unchecked on the [Obtain the Time](#) page.
- All of the NTP protocol boxes must be unchecked on the [Serve the Time](#) page.

With these settings, Domain Time will not set the time or manage the clock. Domain Time II Server will only answer time sync queries on enabled protocols other than NTP.

This option is not recommended.

■ NT5DS

The Windows Service is set to run and it obtains the time from the Active Directory hierarchy in NT5DS sync mode.

This mode will severely conflict with Domain Time unless:

- The "Do not set this machine's time" radio button is selected on the [Obtain the Time](#) page.
- The "Refuse to serve time until this machine's time has been set" checkbox has been unchecked on the [Obtain the](#)

[Time](#) page.

- All of the NTP protocol boxes must be unchecked on the [Serve the Time](#) page.

With these settings, Windows Time obtains the time, manages the local clock, and serves NTP. Domain Time will merely answer time sync queries on all other enabled protocols.

This option is not recommended.

■ AllSync

The Windows Service is set to run and it attempts to obtain the time from the Active Directory hierarchy in NT5DS sync mode and/or using NTP Client mode.

This mode will severely conflict with Domain Time unless:

- The "Do not set this machine's time" radio button is selected on the [Obtain the Time](#) page.
- The "Refuse to serve time until this machine's time has been set" checkbox has been unchecked on the [Obtain the Time](#) page.
- All of the NTP protocol boxes must be unchecked on the [Serve the Time](#) page.

With these settings, Windows Time obtains the time, manages the local clock, and serves NTP. Domain Time will merely answer time sync queries on all other enabled protocols.

This option is not recommended.

■ NTP

The Windows Service is set to run and it attempts to obtain the time using Windows Time's NTP Client mode.

This mode will severely conflict with Domain Time unless:

- The "Do not set this machine's time" radio button is selected on the [Obtain the Time](#) page.
- The "Refuse to serve time until this machine's time has been set" checkbox has been unchecked on the [Obtain the Time](#) page.
- All of the NTP protocol boxes must be unchecked on the [Serve the Time](#) page.

With these settings, Windows Time obtains the time, manages the local clock, and serves NTP. Domain Time will merely answer time sync queries on all other enabled protocols.

This option is not recommended.

Disable Agent

This checkbox disables the Domain Time II [Windows Time Agent](#).

Note: In version 4.1, the Domain Time II Windows Time Agent was installed by default. In version 5.1 and newer, Domain Time is able to replace Windows Time entirely, so the agent is not installed and the option defaults to disabled. If you have upgraded from 4.1, the agent is still present but not required. You may use this option to disable it.

This option has no effect if Agent is not installed.

If you would like to use the Agent, you may install it from the distribution files, or by using Manager, or [download the software](#) from the website, if desired. You must close and re-open the Server Control Panel applet after installing the Agent.

If Agent is installed, the [Agent button](#) launches the Domain Time II Windows Time Agent to allow you to view and configure the settings for the Windows Time service. Depending on the settings above, various parts of the Windows Time Agent applet may be disabled. See the full [Windows Time Agent documentation](#) for more details.



Clock Control

Domain Time gives you extensive control over how corrections are applied to the system clock.

These advanced settings are provided to address special clock-correction requirements, poorly-behaving system clocks, and for fine-tuning for extreme clock accuracy. In most cases, you will not need to make changes here.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings **WILL** adversely affect your clock accuracy or even prevent clock corrections entirely.

Clock Corrections vs. Alignments

Domain Time can correct the clock either by "stepping" (immediately changing the time) or "slewing" (changing the time slowly). Stepping and slewing only operate on variances of 1 millisecond or more.

If slewing is enabled, variances of less than 1 millisecond are "aligned," which are very small slewed clock adjustments. Sub-millisecond alignments are NOT considered corrections, and will not show as corrections in some displays and reports, such as drift records, Audit Server reports, etc. Variances of less than 1 millisecond will be reported as zero milliseconds, except in the log files, drift graphs, or Manager's displays.

If your machine is stepped, the log file will say "Local clock stepped" (followed by details on which direction, by how much, and the protocol used to obtain the time).

If your machine is slewed, the log file will say "Local clock slewed" (followed by the same details as for stepping).

If your machine is aligned, the log file will say "Local clock aligned" (followed by the same details as for stepping or slewing).

Alignments happen automatically as long as slewing is enabled. The only important thing to remember about alignments is that they are not reported as clock corrections.

About Slewing and Stepping

There are two methods of correcting the system clock: Slewing or Stepping. Slewing means adjusting the system's overall clock rate so that the system either speeds up or slows down until it matches the wall clock. Stepping means an instantaneous jump to the new time, either forward or backward.

Slewing gives all processes a linear progression of ticks as time passes, and time is guaranteed never to go backward ("backwards" time corrections are actually made by causing the system clock to go forward at a slower pace than normal until the actual time catches up).

Slewing is critical to time-sensitive applications like databases, logging facilities, or auditing where a backwards jump in time would be highly disruptive. Slewing can also yield more precise corrections than stepping, and can correct variances of less than one millisecond. Slewing is therefore the correction method of choice.

Clock Corrections

- Slew the clock if possible, otherwise step it
- Only slew the clock, never step it
- Only step the clock, never slew it

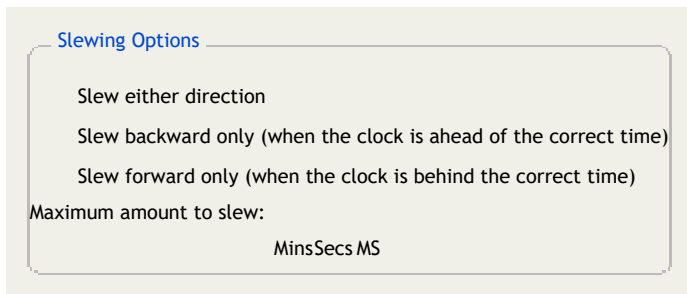
The **Clock Corrections** dialog selects whether corrections are made by slewing, stepping, or a combination of both.

Slew the clock if possible, otherwise step it

When this radio button is selected, Domain Time will step corrections too large to slew (or if slewing in that direction is disabled on the *Slewing Options* dialog page, see below), and will also step the very first correction after rebooting. This is

the default option and highly recommended.

The [Slewing Options](#) button brings up a dialog that lets you specify the direction of slew and the maximum amount of slewing that is permitted.



By default, Domain Time will slew both forward and backward to correct the clock, provided the correction being applied is within the set limit for slewing. The default limit is 30,000 milliseconds (30 seconds), but you may change this to anywhere in the range of 1 through 3,600,000 (1 millisecond through 1 hour). You may disable forward slewing, backward slewing, or both. Slewing large corrections can take an extended amount of time, so be careful if you modify this setting. If slews take a long time to complete, the clock can continue to drift significantly during the slew, making it very difficult to achieve accurate corrections.

Only slew the clock, never step it

In v4.x, you could change the default stepping behavior by modifying the "Never Step Clock" registry option. However, in v4.x, "Never Step" really meant "Do not step except on first boot or when triggered by an administrator," which was a bit ambiguous.

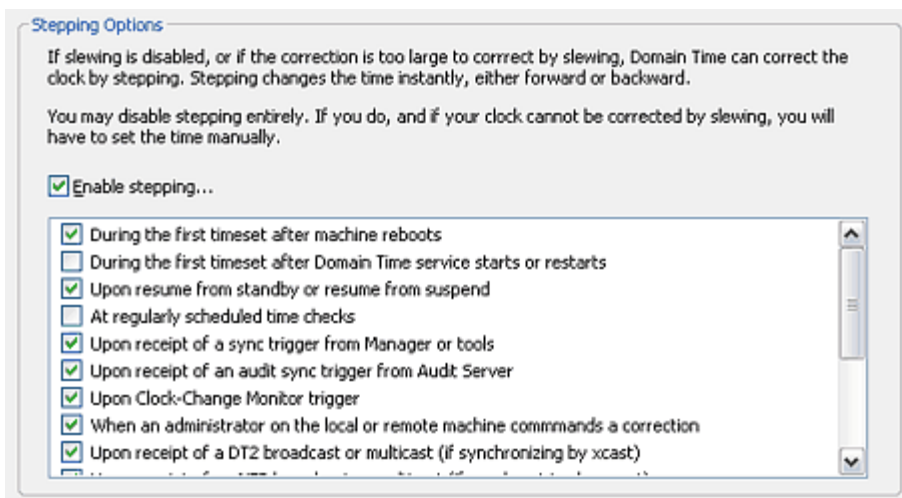
In v5.x, this setting is now available on the Control Panel applet. If it is enabled, Domain Time really will never step the clock. The slew limits and slew direction settings are not overridden by sync triggers, the control panel applet, or reboot detection. As a result, if you have this option selected, you will probably have to set the clock manually after every boot to get the time within the slew limit range to begin correcting the clock.

IMPORTANT: If this setting is enabled and the clock variance *is ever* outside the slew limit, the clock will never be corrected; a log entry indicating this condition will be made instead. Therefore, use care when choosing this option.

Only step the clock, never slew it

To provide greater control of the stepping process, v5.x introduces the [Stepping Options](#) dialog page which allows you to select the conditions under which stepping is allowed.

dialog page which allows you to

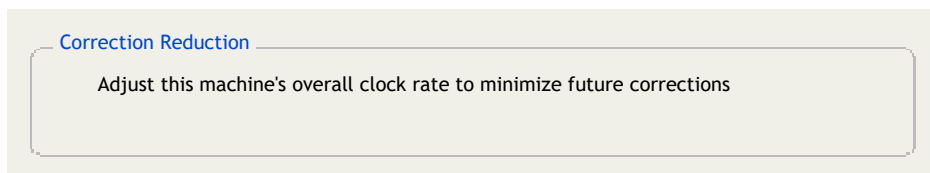


Stepping Options Dialog [\[Click for larger size\]](#)

The settings on this dialog correspond to the new "Allow Stepping" registry setting. "Allow Stepping" is a bitmask of the selected options. If your v4.x machine had "Never Step" specified in the registry, the value will be translated to an "Allow Stepping" value of zero when upgrading to v5.x. In all cases, stepping will only be applied if slewing is disabled or if the variance is outside the slewing limit (see above).

Correction Reduction and Advanced Clock Control

In addition to correcting the time during a time check, Domain Time can use highly-sophisticated clock control methods to ensure the clock on your machine runs more accurately, even between time corrections. This section shows whether these processes are enabled, as well as some statistics about the current clock management parameters.




This setting determines whether Domain Time will manage the rate of the system clock between time corrections. In nearly all cases, you will want to leave this checkbox checked.

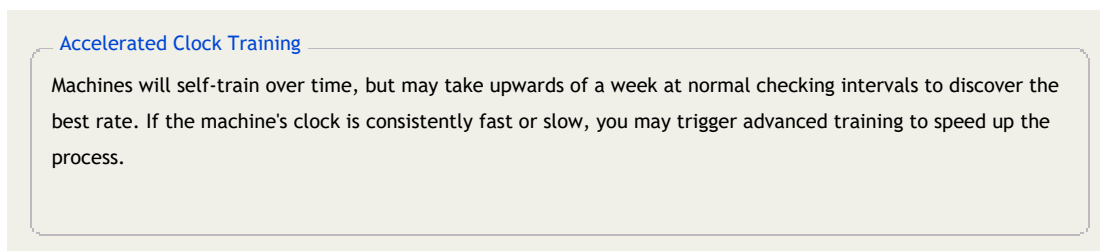
IMPORTANT: When unchecked, Windows will operate as if there is no time service controlling the clock between time corrections and the system clock will therefore run at whatever rate was last adjusted to (see the description of **Phase Adjustment** below). This may cause the clock to drift severely between corrections. More critically, because Windows thinks no process is managing the clock, it will periodically (usually ~ once an hour), hard-set the system clock to match the CMOS hardware clock on the motherboard, causing jumps in the system time. The CMOS clock is notoriously unreliable and thus the resulting time jumps may be very large.

The stability of the system clock on Windows is a result of complex and difficult-to-predict interactions of hardware and software timers, operating system processes, and running applications on each machine. Domain Time's Correction Reduction functions are extremely good at taking all these items into account; however, some machine's clocks are difficult to discipline automatically, and the default automatic algorithms can sometimes appear to make things worse.

If your machine gets progressively less accurate over time, please carefully investigate the options on the Advanced Clock Control dialog (see below) before turning off Correction Reduction entirely. It may be that resetting the timings to defaults, locking the clock rate, adjusting the Interphase settings, and/or changing the Timer Resolution settings will improve the accuracy dramatically on an errant system. Uncheck **Correction Reduction** only if you have a compelling reason to do so.

Advanced Clock Control

The  button will bring up the **Advanced Clock Control** dialog.



Machines running Domain Time II automatically train themselves over time to better match the speed of their sources. If a particular machine is consistently slow, for example, it will gradually speed up in order to reduce the frequency and magnitude of corrections. If it is consistently fast, it will gradually slow down. See the "Phase Adjustment" section below for a description of how this is done.

Self-training may take several days to several weeks to reach equilibrium across all the machines on your network since each machine has to collect data from its own time synchronizations, determine what type of adjustment to attempt, and then re-sample to determine if the adjustment was correct. In most cases, the self-training process will hone in on the best clock rate for each machine, and very few training changes will occur thereafter.

However, it may sometimes be desirable to take a shortcut to speed up the training process. The button initiates a special accelerated clock training sequence that causes an extended series of rapid clock synchronizations that allows the system to estimate what the correct clock rate is in a relatively short period of time.

Accelerated clock training is not as accurate as the automatic self-training over time, however it can get the machine close to the correct clock rate so that final self-training can fine-tune the rate more quickly. On machines with highly variable load where self-training cannot reliably determine a correct rate, accelerated training can be used to determine a decent clock rate value to use as a starting point to determine a locked rate (see below).

Recommendations:

- Use accelerated training on your master time server when you first deploy Domain Time II on your network. When the master is fully trained, then you can use accelerated training on each slave.
- Then, use accelerated training on your Independent Servers.
- Allow other machines to automatically train their clocks without accelerated training if possible.
- Do not use accelerated training on clients before their servers have finished either self-training or accelerated training, since this will result in inaccurate training of the client.

Multimedia Timer Resolution

Set multimedia timer to the maximum resolution for this machine (default checked)

Set OS timer to the maximum resolution for this machine (default unchecked)

In most cases, having these boxes checked will increase the accuracy of the system clock. However, they are global settings for all applications on the machine, and other applications may raise or lower the resolution unexpectedly. Also, changes to the OS timer in particular may have unintended consequences to other applications or system functions, so the default for that setting is unchecked.

If you find that your system's clock accuracy changes when certain applications are run (such as multimedia applications, Java applets, etc.) you may want to try disabling these settings, reset the clock rate settings on this page to the defaults (click the button, and then use the Accelerated Clock Training function (described above) to re-train the software while your application is active.

Phase Adjustment

Default phase adjustment: **156001**

Phase adjustment: (integral rate) Phase Locked

Interphase adjustment: (+-ms/minute) Interphase Locked Continuously variable (PTP)

Interphase period: 10 seconds Change rate limiter enabled: %

Interphase significance threshold: (average delta in hectonanoseconds)

Interphase reliability threshold: (minimum ms/minute change allowed)

These settings display (and set) the current clock rate settings on the machine.

In most cases, you will not need to adjust these settings as Domain Time will usually do an excellent job of finding the optimum tuning values. However, you may be able to achieve higher accuracy/clock stability by adjusting these settings.

Default Phase Adjustment: displays the original clock tick rate calculated by the OS for the hardware of this machine. Windows expects that this rate will equal one second per second. In reality, it almost never does. The actual number of ticks necessary to run at exactly one second per second on this machine is usually some value greater or less than the default.

Phase adjustment:

Domain Time attempts to automatically determine the optimum tick rate setting and will adjust this value in small increments over time. The current tick rate setting in use is displayed in the **Phase adjustment:** field.

Checking the **Phase Locked** box will lock the clock at the current rate. This should be used with great care as it prevents any further automatic clock training. Use this setting if automatic clock training results in incorrect phase adjustments. For example, auto-training on some systems may result in the phase adjustment continually incrementing, causing the clock to run ever faster (or slower). Locking the rate will prevent the rate from incrementing.

Locking the phase rate can also be useful if auto-training results in the clock consistently running slightly slower or faster than the correct time. You can change the phase rate and lock it so that corrections are slightly behind the actual time (so that corrections are always speeding up the clock slightly) Then, use the Interphase settings described below to fine-tune the corrections the rest of the way.

IMPORTANT: On versions of Windows other than Vista, 2008, 2008/R2, and Win7, the tick rate can reliably be set to the actual resolution shown (i.e. changing the tick rate from 156250 to 156251 will change the clock rate by an exact number of microseconds per second). However, with Vista, 2008, 2008/R2, and Win7, Microsoft made underlying changes to how Windows handles the clock timers, resulting in a significant reduction in the granularity of this setting. In effect, on those operating systems a change of at least 16 ticks must occur before a change in clock rate actually happens on those systems.

In the example above, changing 156250 to 156251 will not have any effect. You must change the rate to 156266 ($156250 + 16$) before the clock rate actually changes. The next rate change will occur at 156282 ($156250 + 32$), etc.

In most cases, this decrease in granularity will result in a significant loss of accuracy unless you allow Domain Time to compensate for it properly using the **Interphase adjustment** setting below.

As of Windows 2012/Windows 8, Microsoft appears to have addressed the granularity issue so that single digit changes to the tick value again respond as expected.

The **Interphase adjustment:** settings allows for fine-tuning the clock rate when the correct time falls somewhere between the resolution of the tick rates set in the **Phase Adjustment** setting above. As described above, this is much more likely to be necessary on newer versions of Windows, since the granularity of phase adjustment is larger than on older versions. This value is the automatic clock correction value applied every minute (at the rate selected by the **Interphase period** setting below) to provide highly precise clock adjustment (and compensate for underlying timer errors/corrections on newer versions of Windows).

These corrections are usually made automatically, but if your system is unable to achieve the desired level of accuracy, you can manually set this value and lock it using the **Interphase Locked** checkbox. Use this value for extreme fine tuning after getting the clock as accurate as possible first if you are using NTP or DT2 protocols to synchronize time.

If you are using the IEEE 1588-2008 (PTP) protocol to synchronize time, it is usually better to allow Domain Time to continuously adjust the Interphase settings automatically to achieve significantly more accurate and smoother clock corrections. This function is enabled by checking the **Continuously variable (PTP)** checkbox. If you are using PTP and this function is enabled, the **Interphase adjustment/Interphase period** settings are ignored, even if locked. Note that if PTP sync is lost, these settings resume effect if the machine falls back to using another time

protocol like NTP or DT2.

The **Interphase period: 10 seconds** selection determines how often per minute Domain Time applies a fraction of the Interphase adjustment when attempting to exactly match this clock's phase to a time source.

For example, a setting of 15 seconds will apply 1/4 of the Interphase setting every 15 seconds. A setting of 20 will apply 1/3 of the Interphase value every 20 seconds, etc.

Due to the complex nature of the underlying timers, there is usually a "sweet spot" for this setting that results in maximum accuracy, but it can only be discovered by trial and error. It is probably not necessary to adjust this value unless you have first fine-tuned your accuracy as much as possible with all of the other settings.

Change rate limiter enabled: % controls how quickly the software makes changes to the Interphase settings based on the collected time samples. Setting this value too high may result in overly-large Interphase adjustments; too low of a setting may result in unnecessarily-small adjustments. Either condition can affect how long it takes to converge on accurate timings for the machine.

Interphase significance threshold:

Interphase reliability threshold:

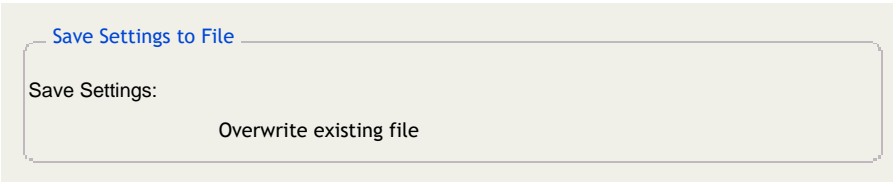
These settings fine-tune the criteria the automatic Interphase adjustment algorithm uses to calculate the Interphase rate. You should not change these values without advice from Technical Support.

Import/Export

You may import or export Domain Time II configuration settings using the utilities on this page.

Starting with version 5.1, Domain Time Server can import/export its settings using a standard Windows Registry .reg file. This allows you to easily make backup copies of the settings, create a custom .reg file to use as a template for configuring other machines, or to use in creating a custom installation package or script.

The Import/Export function automatically excludes any settings that are machine-specific so that the .reg file may safely be imported on any other copy of Domain Time Server (of the same version) without causing disruption.



This section allows you to save the current configuration to a .reg file.

The .reg file will be created using the currently configured options of the Server. You should review each of the settings of the Control Panel applet to be sure that they are correct before exporting the file.

If you expect to be importing the .reg file on machines that need differing configurations (such as for machines in different cities that use different time sources), you should configure the applet for each configuration, and then export a separate .reg file for it.

You may use exported .reg files as template files for installing/upgrading multiple machines using [Domain Time Manager](#). Template files need to be located in the **C:\Program Files\Domain Time II\Templates\[Server][Client]** folder of the Domain Time II Manager machine. Template .reg files located in those folders will automatically be made available for use when installing or upgrading using Manager.

Note, if you launched the Control Panel applet using Manager, the export utility will automatically offer to save the .reg file in the proper directory on the Manager machine. Otherwise, you will have to manually save or copy the file to the Templates folder on the Manager machine.

Exported templates may be edited manually using any text editor. Templates may contain all exported settings, or only the specific settings you want to change.

About Settings (.reg) Files

Domain Time Servers and Clients are background system services that obtain all of their running settings from the Windows Registry.

Domain Time components get their initial registry settings by importing default template files during installation or upgrade. As of version 5.1, the templates are standard Windows Registry Editor .reg files in either Unicode or ANSI text format. Previous versions used a proprietary .ini template file (domtime.ini), which has been deprecated.

The default template file for Server is **dtserver.reg**; Client uses **dtclient.reg**. These files are included in the original distribution files used during Setup or in the source files used for remote installation by Domain Time II [Manager](#).

During installation/upgrade, the appropriate template file is copied from the distribution files to the **Windows/System32** folder of the target machine.

Once installed, Domain Time will not use the default .reg template file again, unless the user clicks a Reset to Defaults button on the Control Panel Applet, or if a Reset Configuration is commanded remotely from Domain Time II [Manager](#).

IMPORTANT: Although .reg files created using this utility are saved in standard Windows registry file format, it is **not** equivalent to exporting the registry keys using Windows' RegEdit program. A number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. Those settings are automatically excluded when you export using this utility, so you should always use this utility to create a Domain Time .reg file.

Load Settings from File

Installation defaults (the settings used when Domain Time was first installed on this machine)

Choose File:

Do not prompt for confirmation

Use this section to import the default settings or a custom .reg file.

Installation defaults (the settings used when Domain Time was first installed on this machine)

This selection will reload the settings file that was used when the product was installed.

The file is named **dtserver.reg** and is located in the **/System32** folder on a running system. This is the default file used during installation. See the [Rollout](#) page of the installation instructions for more information on using this file as a setup template.

Choose File:

Use this to import an existing .reg file.

CAUTION: It may be necessary to restart the Domain Time Server service after importing the settings file.

Although the .reg file is saved in standard Windows Registry file format and you can install it by clicking on the .reg file in Windows Explorer, it is usually better practice to import the file using this utility since it does additional validation checking on the values and attempts to exclude items not appropriate to this version or machine.



This property page contains Domain Time Support information and utilities.

Problem Report

Your company:	Include main log file and drift graph
Your name:	Include most recent startup log file
Your phone:	Include registry settings

The **Problem Report** utility can compile a problem report to send to Domain Time technical support including important diagnostic information and log files which will greatly assist in troubleshooting any problems you may experience. You can either email it directly from the program or save the file to forward it manually later.

You have the option of including various items and logs in the report in a compressed (zipped) file. In most cases, you should include all items to provide the most information possible.

The utility will use the currently-selected default MIME email program on the machine where this utility is run to send the mail and compressed file attachment when you click the button.

If you don't want to (or can't) send email directly from the machine in question, remotely connect to the problem machine from a machine that does have email capability. You can connect remotely using Domain Time Manager, the Remote CPL utility, or from the Control Panel applet of another Domain Time Server or Client. Once you have the remote machine's Control Panel applet displayed, click the button. The utility will send the email from the email client on the local machine, but will automatically include the diagnostic and log files taken from the remotely-connected machine.

If you'd prefer to save the report file to disk and forward it to Tech Support manually, click the button.

Domain Time Server has the option of displaying a handy application (DTTray) in the Windows System Tray.

Note: The System Tray is a graphical function of Windows Explorer, so it is not available on Windows Server Core. This means the DTTray program is inaccessible on that version of the operating system.

DTTray gives you quick access to common tasks (such as manually triggering the machine to sync with its time source) as well as a quick way to launch Domain Time applications.

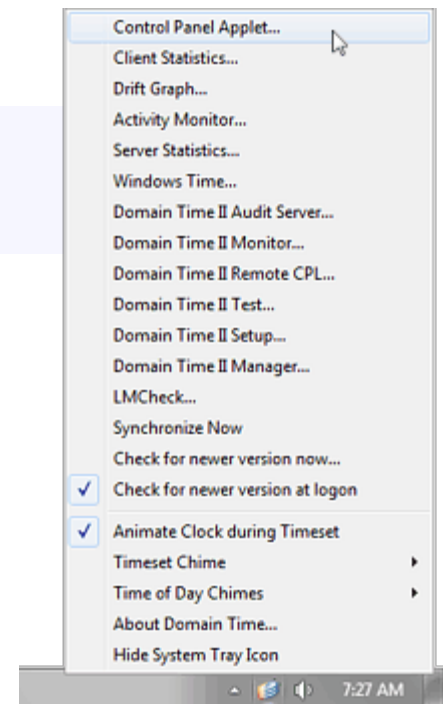
It also gives you visual and audible alerts to your current time sync status, traffic monitoring displays, drift analysis graphs, etc.

Starting the Applet

The DTTray applet loads into the System Tray automatically when you log in.

Double-clicking the system tray icon will pull up the Domain Time Server Control Panel applet. Right-clicking the icon will present you the full DTTray context menu. Only Domain Time applications currently installed will appear on the context menu.

Items installed or removed while the tray applet is loaded may not appear on the DTTray context menu until the user logs off or the system is restarted.



System Tray Applet [\[Click for larger size\]](#)

Hiding the System Tray Applet

You can hide the system tray applet by clicking the **Hide System Tray Icon** item on the DTTray context menu itself.

You can show or hide the system tray applet by setting the **"Show system tray icon"** checkbox on the [Advanced](#) property page of the Control Panel applet.

You can also set the value of the **HKLM\Software\Greyware\Domain Time Server\Parameters\SystemTrayIcon** registry key if you prefer.

Multiple Instances on Terminal Services/Remote Desktop

By default, the system tray applet will only appear in the first logged-on instance of the console. If you want the system tray to appear in all terminal sessions, you can set the value of the following registry key to **True** :

Location: **HKLM\Software\Greyware\Domain Time System Tray Icon\Parameters**

Key: (create it if it doesn't exist) **Allow Multiple Instances**

Type: **REG_SZ**

Note: Each running instance of the icon holds a file lock on the tray icon executable. If you enable this feature, it will be necessary for all users (local and remote) to log off before performing an upgrade of Domain Time on this machine.

System Tray Applet Command Functions

Use the context menu to trigger a sync and check for updates.

Triggering a Time Sync

You can trigger the time service to synchronize with its time source by right-clicking system tray icon and choosing **Synchronize Now** from the context menu.

Alert Functions

The DTTray Applet has a number of features that provide visual or audible feedback on the status of your clock synchronization.

Animate Clock During Timeset

While Domain Time is synchronizing the clock, the system tray icon will show a running clock icon. When the clock is successfully synchronized the icon will change to the standard Domain Time icon. You can turn off this feature by unchecking the option from the right-click context menu.

The "Time Not Synchronized" Alert (The Flashing Clock)

This alert flashes a clock in the system tray to indicate the time is not synchronized.


If Domain Time is unable to synchronize its time with a time source, it will alert you to the problem by changing the Domain Time icon in the system tray to a flashing clock icon. Once you have resolved the cause (usually due to a network issue preventing Domain Time from contacting its time source) and re-synchronized, the icon will return to the normal Domain Time icon.

Timeset Chime

The System Tray Applet can indicate a successful time synchronization with an audible signal.

Timeset Chimes are off by default. You can enable them by pulling up the context menu and selecting the sound device you wish to use to play the chimes. The Timeset Chimes will play whether or not there is a logged-in user.


Choose **Chime off** if you do not want the signal to play.



Timeset Chimes [\[Click for larger size\]](#)

Time of Day Chimes

The Time of Day Chimes feature is a special function that plays sound files at particular times of the day (on-the-hour, and at 15, 30, 45 minutes past the hour) to emulate a chiming clock. You may download selected free chime packs from our site or create your own.



Time of Day Chimes [\[Click for larger size\]](#)

The Time of Day Chimes are off by default. You can enable them by pulling up the context menu and selecting the chime pack you want to use. (You must have downloaded and installed at least one chime pack before this feature will be available.)

In order to play the Time of Day Chimes, you must meet these requirements :

- Your system must be configured with a sound card, drivers, and other hardware (such as speakers or headphones) necessary to play .WAV files.
- Time of Day Chimes are played by the DTTray Applet, so you must be logged in and have the System Tray Applet installed and visible in the system tray if you want the Time of Day chimes to play.

You must have downloaded and installed at least one chime pack (see below).

Choose **None** if you do not want the chimes to play.

How to Install Chime Packs

Chimes are standard .WAV sound files played using the Windows Media subsystem. You can choose from free chime packs that we've prepared for you to download or you may create your own.

Download	Hear a Sample
Grandfather Clock	Listen
Westminster Classic	Listen
Westminster Glass	Listen
Westminster Deep	Listen
Cuckoo Clock	Listen

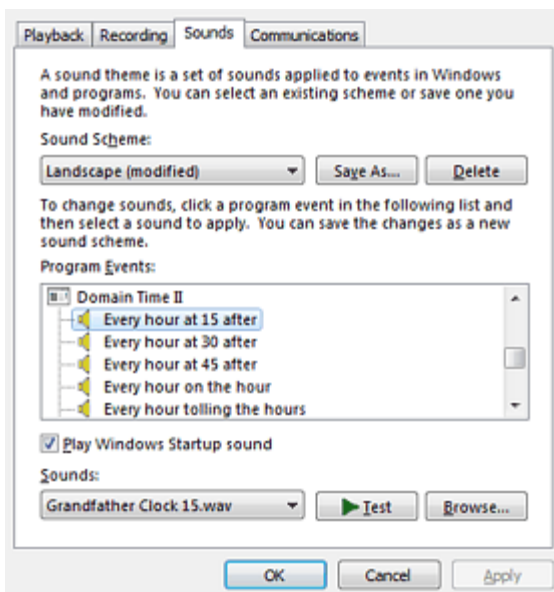
To install a chime pack, download the zip file and unzip the contents into your Media folder (usually **C: \Windows\Media**). There will be one text file (.txt) and one or more sound files (.wav) contained in the zip. The text file contains a description of the chime pack, and instructions for which sound file goes with which event. These instructions are used by the DTTay Applet when you chose a chime pack.

Important: The .txt file from the chime pack must be copied into the Media folder along with the .wav files. It must be present for the chime function to work.

After unzipping the files into your Media folder, right-click the Domain Time II system tray applet. Your newly-installed chime pack should show up by name under Time of Day chimes.

Configure and Customize your Chimes

Chimes are fully configurable using the Windows **Sounds** Control Panel applet (called *Sounds and Audio Devices* on older versions of Windows).



Time of Day Chimes [\[Click for larger size\]](#)

To customize your chimepack, launch the **Sounds** (Sounds and Audio Devices) applet, then click the **Sounds** tab. Scroll down through the list until you see the Domain Time II sound scheme. Listed under Domain Time II, you'll see entries for "Every hour at 15 after," "Every hour at 30 after," and so on. You may associate any .WAV files you like with the various sound events.

The "Every hour tolling the hours" sound is played after any other chimes on the hour, and is played a number of times corresponding to the hour (using a 12-hour scheme). At one o'clock, it will play once, at two o'clock it will play twice, and so forth.

You can make your own chime pack by collecting the .WAV files you want to use and creating a text file for them. Download one of the chime packs from the list above and look at how the text file specifies the sounds. To have your own chimepack show up so you can select it from the System Tray Applet context menu, simply create a new text file with a .txt extension. The file should be in the same format as the sample chimepack, listing which .WAV plays for

which event. Use a full or relative path before the filename if your .WAV files aren't located in the Media folder. You'll notice from the sample .txt files that you can create a very elaborate set of chimes played at various times and events. Feel free to experiment!

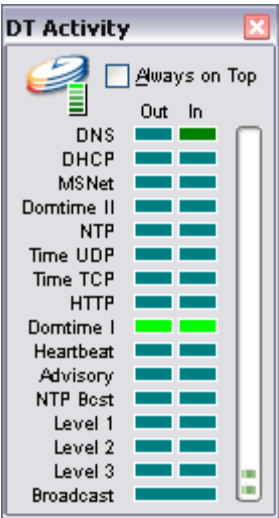
Activity Monitor

This window displays all Domain Time-related network traffic sent or received by this machine.

The Activity Monitor provides a visual indicator of when the various types of protocol and control messages are sent or received. Besides being a great deal of fun to watch, it is a useful diagnostic tool you can use to verify that your server or client is receiving and responding to time sync requests and cascade signals.

For example, click the Synchronize Now option from the DTTray context menu to watch how your system performs a sync request.

Click the **Always on top** box if you wish the Activity Monitor to always be visible on your screen.



Statistical Functions

View detailed statistics and view drift graphs

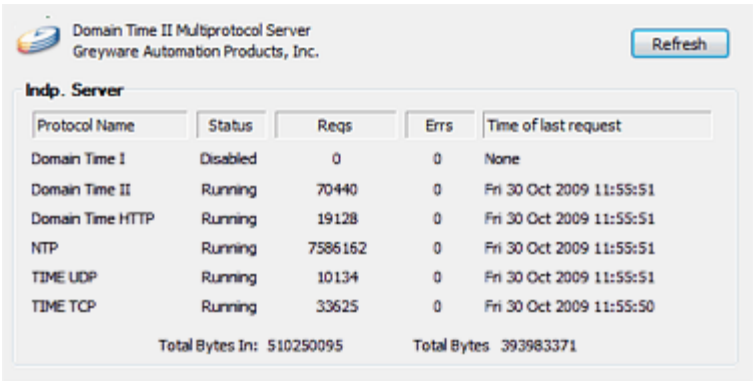
Viewing System Statistics

You may view the current time statistics on this machine by double-clicking the system tray icon, or by right-clicking the icon and choosing **Client (and/or Server) Statistics** from the menu.

Note that Domain Time II Server acts as both a time client and a time server, so there are two statistical displays available on Server.




Client Statistics (on Server) Display [\[Click for larger size\]](#)



Server Statistics [\[Click for larger size\]](#)

Drift Graph

In addition to detailed summary statistics, DTTray can show you a graphical representation of the accuracy of the clock on your machine sampled at the time it synchronized with its time source. You may scroll through the entire drift data to see how your clock has been performing over time.



The Drift Graph [Click for larger size]

You can also see the actual sync data in text format by clicking the [Raw Data](#) button. This data can also be collected and analyzed centrally using [Audit Server](#).



Domain Time Server keeps its settings in the Windows Registry. Most of the service options are best set using the Domain Time Server Control Panel applet. However, some advanced options can only be set by changing the registry. This page explains many of these special registry entries used by Domain Time Server.

CAUTION:

Modifying Registry entries requires basic familiarity with the Windows Registry and its operations. Incorrect changes to the Registry can result in unpredictable, perhaps non-repairable, damage. We cannot be responsible for registry problems.

The Domain Time II Server settings are located in these keys (click the names to jump to details):

HKEY_LOCAL_MACHINE

Software

Greyware

Domain Time Server

[Enabled Protocols](#)

[HTML](#)

[Keyring](#)

[Logs and Alerts](#)

[Parameters](#)

[Time Sources](#)

Enabled Protocols

The Domain Time II Server Enabled Protocols settings are located in this key:

HKEY_LOCAL_MACHINE

Software

Greyware

Domain Time Server

[Enabled Protocols](#)

The values listed in the **Enabled Protocols** registry key represent the protocol types Domain Time will listen for. They correspond to checkboxes on the [Serve the Time](#) and [Status Reports](#) property pages of the Control Panel applet. You should not make manual changes to this key or its subkeys.

HTML

The Domain Time II Server HTML settings are located in this key:

HKEY_LOCAL_MACHINE

Software

Greyware

Domain Time Server

HTML

If you enable the Domain Time over HTTP protocol to be served by the Domain Time II Server, Domain Time will optionally provide both a human-readable stats web page (when the server is visited by a browser) and/or a compact time data packet (when the server is queried by a Domain Time Server or Client).

You may add a custom header and footer to the human-readable web page. For example, everything on tick.greyscale.com and tock.greyscale.com outside the rectangular frame is either a custom header or footer.

The custom header and footer are optional. Add them to the registry as follows (you will have to create the HTML registry key and Header and Footer values manually):

Value Name: Header

Value Type: REG_MULTI_SZ

Default Data: (blank)

Notes: Any HTML you want to appear above the time display

This key and value do not exist unless you create them. Whatever HTML you place in the Header and Footer values will appear on the web page along with the standard logo and status report. If you include a <body...> statement in the Header value, Domain Time will use your <body...> tag instead of its own.

Value Name: Footer

Value Type: REG_MULTI_SZ

Default Data: (blank)

Notes: Any HTML you want to appear below the time display

This key and value do not exist unless you create them.

Value Name: Display Network Info

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: Controls whether or not network information such as IP addresses and the server name are displayed on the web page.

This key and value do not exist unless you create them.

Value Name: robots.txt

Value Type: REG_MULTI_SZ

Default Data: (see Notes below)

Notes: This sets the contents of the robots.txt file provided by the HTML server.
Use this value to control whether search engines should index your time server's web page.
Default contents:

```
User-agent: *  
Disallow /  
Disallow *
```

Value Name: style.css

Value Type: REG_MULTI_SZ

Default Data: (see Notes below)

Notes: This value contains the CSS code used to style the HTML status web page.

Default contents:

```
body {background:#e0ddd0; margin:15px 0px}  
table.main {background-color:#ffe; border-collapse:collapse; border:solid #000 0px;  
width:95%; margin:0px; margin-left:auto; margin-right:auto;}  
table.main td {padding:0px 1%; line-height:130%;}  
.maintop {font: bold small sans-serif; text-align:left; color:#fff; background-color:#18204f;  
height:40px;}  
.mainleft {font:x-small sans-serif; text-align:center; vertical-align:top; background-  
color:#ffe;}  
.mainright {font:x-small sans-serif; text-align:center; color:black; vertical-align:middle;  
background-color:#ececce;}  
.mainbottom {font:bold x-small sans-serif; text-align:right; background-color:#18204f;  
height:40px;}  
table.left {background-color:#ffe; width:100%;}  
table.right {background-color:#ececce; width:100%;}  
.l {text-align:right;}  
.r {text-align:left;}  
.refresh {display:block; float:right; margin:0 5px 0 0; background-color:#18204f;  
cursor:pointer; padding:5px 1%; border-top:1px solid #999; border-left:1px solid #888;  
border-right:1px solid #444; border-bottom:1px solid #666; line-height:120%; text-  
decoration:none; color:#ccc;}  
a:hover {background-color:#456; color:#ffe;}  
hr {border:0; color:#000; background-color:#999; height:1px; width:70%; text-align:center;  
vertical-align:top;}
```

Keyring

The Domain Time II Server Keyring settings are located in this key:

HKEY_LOCAL_MACHINE
Software

Greyware
Domain Time Server
Keyring

The values listed in the **Keyring** registry key contain various items related to authentication. They correspond to settings on the [Symmetric Keys](#) property page of the Control Panel applet. You should not make manual changes to this key or its subkeys.

Logs and Alerts

The Domain Time II Server Logs and Alerts settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time Server
Logs and Alerts

The values listed in the **Logs and Alerts** registry key contain various items related to logging and alerting functions. They correspond to settings on the [Logs](#), [Windows Event Viewer](#), [Syslog](#), [SNMP](#), and [Status Reports](#) property pages of the Control Panel applet. You should not make manual changes to this key or its subkeys.

Parameters

The Domain Time II Server Parameter settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time Server
Parameters

The values listed in the **Parameters** registry key control a wide variety of Domain Time functions. In most cases, they are auto-generated or correspond to settings on the property pages of the Control Panel applet. In general, you will not need to make manual changes to this key or its subkeys.

However, some values require additional explanation or control functions not exposed on the Control Panel. Those items are listed here.

Value Name: Accept First PTP Timestamp

Value Type: REG_SZ

Default Data: *False*

Options: *True or False*

Notes: If set to *True* and no other time sources are configured, then the clock will be stepped or slewed if within slewing limits to match the first PTP timestamp(s) received (the number of samples required are configured using the **Accept Firest PTP Sample Count** registry entry described below). This initial correction will bring the clock into close enough sync for normal PTP operations to govern the clock. Note, on versions prior to 5.2.b.20200930, the clock adjustment was always stepped.

IMPORTANT: Changing this setting to *True* is discouraged in networks with fallback NTP/DT2 timesources, since a restart of the service may step the clock, including possibly stepping the clock backwards. This option should only be used in closed environments where PTP is the only possible source of time and the initial startup delta takes an excessively long time to correct (i.e. if the motherboard CMOS clock is wrong).

Value Name: Accept First PTP Sample Count

Value Type: REG_DWORD

Default Data: 3

Range: 1-15

Notes: Introduced in v5.2.b.20200930. Applies only if **Accept First PTP Timestamp** is enabled. Specifies the number of timestamps that must be received before the clock is adjusted.

Value Name: Allow Browser-based HTTP requests

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: Controls whether or not this Server will allow browsers to contact the server's built-in web server to view the running server statistics.

The **Domain Time II over HTTP** protocol must also be enabled on the [Serve the Time](#) property page of the Control Panel applet.

Value Name: Allow Client-based HTTP requests

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: Controls whether or not this Server will allow Domain Time II Clients to obtain the time via the Domain Time over HTTP protocol.

The **Domain Time II over HTTP** protocol must also be enabled on the [Serve the Time](#) property page of the Control Panel applet.

Note: This setting has a more limited scope than in versions prior to 5.1; it only applies to Client access. In earlier versions it also controlled Server and browser access. See the **Allow Server-based HTTP requests** and **Allow Browser-based HTTP requests** values for enabling these items individually in v5.1 and above.

Value Name: Allow Server-based HTTP requests

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: Controls whether or not this Server will allow Domain Time II Servers to obtain the time via the Domain Time over HTTP protocol.

The **Domain Time II over HTTP** protocol must also be enabled on the [Serve the Time](#) property page of the Control Panel applet.

Value Name: Allow Remote Timezone Change

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: Enables Domain Time II Manager to change the timezone on this machine.

Value Name: Allow Stepping

Value Type: REG_DWORD

Default Data: Varies

Notes: New as of v5.1, this value is a hex bitmap representing the settings made on **Stepping Options** dialog of the [Clock Control](#) property page. Do not edit this value.

These values will be overridden if the **Never Step Clock** setting (see below) is enabled.

Value Name: Client Settings

Value Type: REG_BINARY

Default Data: Varies

Notes: This value is a hex bitmap of various settings recommended by a Domain Time Master Server to Domain Time Clients. See the [Recommendations](#) page for an explanation of how they operate.

Note: On Domain Time Client, this key controls miscellaneous [settings for the Client](#) itself. See the [Server Settings](#) key for the miscellaneous settings for Domain Time Server.

Value Name: Clock Adjustment Bucket Size

Value Type: REG_DWORD

Default Data: 7

Range: 3-32

Notes: The bucket size is the number of time samples collected before a particular clock adjustment rate is evaluated. The specified value is used except during accelerated clock training, where a fixed value of 5 is employed.

You should not change this number unless instructed by techsupport.

Important: This is a machine-specific setting and should not be included in installation templates or copied to other machines via mass registry imports.

Value Name: Clock Adjustment Statistical Method

Value Type: REG_SZ

Default Data: Automatic

Notes: Sets the type of statistical analysis Domain Time performs on collected time samples from a time source when deciding whether and how much to adjust the clock rate to compensate for drift. Domain Time then uses the calculated clock performance to evaluate and remember each integral clock adjustment rate it tries. Changing this value may improve or degrade timing accuracy (or have no effect).

Changes to this setting take effect immediately after the next group of collected samples is ready for analysis. You do not need to restart the service. You should clear your clock history using the command **dtcheck -resettimings** before changing this value. Allowed values are:

- **Automatic** - On Vista/2008/Win7/2008r2 machines, Automatic will use the median value from each group of samples. On all other versions of Windows, it will use the arithmetic mean (average) of each group of samples.
- **Average** - The arithmetic mean of values
- **Median** - The median number in the array of values
- **Toss** - average of values excluding the highest high and lowest low
- **RMS** - the quadratic mean (signed root-mean-square) of the array of values
- **Disabled** - no statistical analysis is retained for future comparison

Important: This is a machine-specific setting and should not be included in installation templates or copied to other machines via mass registry imports.

Value Name: Clock Change Monitor

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: If enabled, Domain Time monitors changes to the system clock made by other programs (including the foreground user changing the time or date with the Control Panel applet or the command-line TIME and DATE commands). When the Clock Change Monitor is enabled on a Server and the clock changes unexpectedly, the Server will immediately resynchronize with its time source(s).

You may turn the Clock Change Monitor off if your setup requires having machines with different times (usually only in labs or testing environments). If Clock Change Monitor is disabled and you change a machine's time, it will stay changed until the next cascade signal or regular sync interval. Changes take effect immediately, and may be made by editing the registry or remotely from Domain Time II Manager.

Value Name: Clock Change Sensitivity

Value Type: REG_DWORD

Default Data: 0

Range: 0-255

Notes: This value represents the number of seconds the system clock must differ from the expected value in order for **Clock Change Monitor** to decide an unauthorized change has been made to the system clock.

If not present or set to zero, Domain Time will use a value of 2 seconds.

Increase this value only if **Clock Change monitor** is triggering on normal clock drift (unlikely). Decrease this value only if **Clock Change Monitor** is not flagging known clock change events by another user or process.

Value Name: Current Version

Value Type: REG_SZ

Default Data: Varies

Notes: This value is set by the system for informational purposes. Changing it has no effect.

Value Name: Critical Timing Processor Limit

Value Type: REG_SZ

Default Data: Depends on processor type (see below)

Options: *True or False*

Notes: This value is set to *False* during installation on machines with processors that have an Invariant TSC or if they are a Hyper-V guest; otherwise it is set to *True*. When *True*, Domain Time uses the last-processor-but-one for time-critical events, and any available processor for all other work. If set to *False*, Domain Time does not prefer one processor over another for

any task.

Modern CPUs (ones with Invariant TSC) generally have better timing performance with this value set to False. You can check to see if you have an Invariant TSC by running the command-line DTCheck program:

```
dtcheck -cpuid
```

Value Name: Daytime Format

Value Type: REG_SZ

Default Data: *NIST*

Notes: Specifies the format that Server will use for the Daytime (RFC-868) protocol. The available options are NIST, NISTLF, or a custom format.

- If set to NISTLF, the output consists of an LF, followed by the normal NIST format string, followed by a space and then a terminating LF.
- If set to plain NIST, there is no leading LF, and the terminator is a CRLF (no extra space).
- You may also use a *ddd MMM yyyy* format string . Any legal combination of specifiers in Microsoft's GetDateFormat API are acceptable.

Value Name: Dependent Services

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more services you'd like Domain Time to start.

Notes: Requires version 5.2.b.20150516 or later. Any services listed here will be started by Domain Time after the first successful timecheck, as long as the services are set to manual start. This is an alternative to using the built-in service database's dependencies. If you use the built-in functions, dependent services will wait for Domain Time to start, but won't know to wait until the first synchronization has completed.

You may list services by their display names (e.g. "Disk Defragmenter") or by their internal service names (e.g. "defragsvc"). List services one per line, without quotation marks. Domain Time will only attempt to start services that are listed, not yet running, and set to manual startup.

Important: If Domain Time cannot set the clock for some reason (invalid sources, firewall settings, etc.), then services you have set to manual start will not be started.

Value Name: DT2 Bias in Milliseconds

Value Type: REG_SZ

Default Data: +0

Range: -3600000 to +3600000

Notes: Corresponds to offset from the correct time, in milliseconds, the server will use when serving the time to clients or other servers using the Domain Time II or Domain Time over HTTP protocols. Useful chiefly for situations where you need the network to lead or trail the server by a set amount. This setting does not affect the server's own time, or the time it serves using protocols other than DT2.

This is a REG_SZ value, not a binary or DWORD value. Express the offset using a plus sign or a minus sign, followed by the number of milliseconds you want. **Note:** Variance reports are *not* affected by this setting. Do not set up multiple servers with different offsets! Changes to this setting take effect upon restart. The domtimes.log will indicate a warning message if this value is set to anything other than +0.

Value Name: Ephemerides

Value Type: REG_DWORD

Default Data: N/A

Notes: This value is used by the system. Do not edit.

Value Name: ICMP TTL (hop limit)

Value Type: REG_DWORD

Default Data: 32 (decimal)

Range: 1 to 255 (decimal)

Notes: This value controls the number of router hops that are allowed in an ICMP echo ("ping") request. Domain Time pings machines first to help eliminate long waits for machines that are unreachable. You should only need to adjust this value if you have an LAN/WAN configuration requiring more than the default 32 hops.

Value Name: Machine Statistics

Value Type: REG_BINARY

Default Data: N/A

Notes: This binary value contains the statistics, as of the last update, that can be viewed from DTCheck, the Domain Time II Manager, or the system tray icon. Do not edit.

Value Name: Max Slew Correction (milliseconds)

Value Type: REG_DWORD

Default Data: 30000 (decimal)

Range: 1 to 36000000 (decimal)

Notes: This value specifies the upper limit, in milliseconds, of variance that Domain Time will attempt

to correct by slewing instead of stepping the clock. This setting affects both forward and backward clock adjustments.

The older registry entry controlling this function, **Max Slew Correction (seconds)**, has been deprecated.

If the correction to be made is larger than this setting but less than the allowed MaxDisparity setting (Correction Limit), Domain Time II will step the correction (unless **Never Step Clock** is enabled, at which point no correction is made and a note to this effect will be entered in the Domain Time logs). See the **Never Step Clock** and **Override Max Disparity** registry settings for more info.

Value Name: Min Success Interval (seconds)

Value Type: REG_DWORD

Default Data: 5

Notes: Sets the minimum period allowed between timechecks. Do not change this value.

Value Name: Never Step Clock

Value Type: REG_SZ

Default Data: *False*

Options: *True or False*

Notes: When enabled, causes Domain Time to make clock corrections only by slewing. This prevents the clock from being stepped to make corrections such as those normally done during startup or from Clock Change Monitor, manual sync triggers, etc.

CAUTION: Enable this option with care. Use of this option may prevent Domain Time from successfully being able to synchronize with a time source if the time correction is too large to accomplish using slewing. See the **Max Slew Correction (milliseconds)** registry setting for more info.

IMPORTANT: Unlike with versions prior to v5.1, the behavior of this setting is **NOT** modified by the **Override Max Disparity** registry setting. If **Never Step Clock** is enabled, the clock will never be stepped, regardless of any other settings.

As of v5.1,, Domain Time uses the **Allow Stepping** setting (see above) to provide greater control of the stepping process. If your machine running an older version of Domain Time had **Never Step Clock** specified in the registry, the value will be translated to an **Allow Stepping** value of zero when upgrading to v5.x or later. See the **Stepping Options** dialog of the [Clock Control](#) property page to set the options.

In most cases, it is better to set the **Stepping Options** with the behavior you want than to enable **Never Step Clock**.

Value Name: NTP Client Version

Value Type: REG_DWORD

Default Data: 4 (was 3 on versions prior to 5.2.b.20150516)

Range: 1 to 7

Notes: Controls the reported NTP version. Any value from 1 to 7 is legal, although using anything but 3 or 4 is not recommended.

Value Name: NTP Server RefID

Value Type: REG_SZ

Default Data: *Dynamic*

Options: *Dynamic or Static*

Notes: Determines how the NTP RefID field (used to indicate the reference time source) is populated in NTP packets. *Dynamic* is the appropriate setting under most circumstances, which lets Domain Time decide what is a correct response. Change this value only if the automatically-selected RefID causes problems for NTP clients.

Note: This value is related to and may be overwritten in some circumstances based on the NTP Server Stratum value described below. When the NTP Server Stratum is set to 0 (Automatic) and NTP Server RefID is set to Dynamic, Domain Time will typically populate the RefID field with the IP address of its last known time source if the time was derived from a single machine, or 0.0.0.0 if multiple machines were included in the time analysis. LOCL may sometimes appear under certain configurations for compatibility with Active Directory.

If the NTP Server RefID value is set to Static, the RefID field will typically contain word GREY (or the dotted-quad numeric equivalent). Setting the NTP Server Stratum to 1 will force the NTP Server RefID to be Static. You may use either Dynamic or Static with any other value of NTP Server Stratum (2 - 15).

Note that when you set the Domain Time Server to "Do not set this machine's time" on the Obtain the Time property page, there is no external time source, so the RefID field will contain zeros. Some older NTP clients interpret this as the machine not being a valid time source and therefore refuse to synchronize with the Server. In this case, set the NTP Server Stratum value to 1 and restart the service.

Value Name: NTP Server Stratum

Value Type: REG_DWORD

Default Data: 0

Options: 0 through 15 (decimal)

Notes: Specifies the NTP Stratum of this Server. A value of 0 means automatic assignment.

As of 5.2.b.20150516, when set to automatic assignment Domain Time selects a stratum

number based on on strata reported by its time sources (including PTP, which uses "stepsAway" to correspond, roughly, with NTP strata).

Since Domain Time can use multiple sources with multiple protocols, there may not be an ultimate single stratum from which time was received. In these cases, Domain Time takes the highest-level stratum reported by all used sources, and adds one to derive its own stratum number. So, for example, if Domain Time obtains its time from a PTP grandmaster directly, it will report itself as stratum 2. If it receives its time from three NTP sources, two of which report stratum 1, and one of which reports stratum 2, Domain Time will report itself as stratum 3. If all three NTP sources reported stratum 1, Domain Time would report stratum 2, and so forth.

Versions prior to 5.2.b.20150516, automatic assignment applied a default Stratum depending on the type of Server Role selected:

- Master Server: Stratum 2
- Slave Server: Stratum 3
- Independent Server: Stratum 2

You may specify a different value if it better fits your time distribution configuration.

Value Name: Override Max Disparity

Value Type: REG_DWORD

Default Data: *Not present* (same as zero)

Options: 0, 1, 2, 3, or 4

Notes: Controls how Domain Time decides when to override the **Correction Limits** set in the Control Panel applets for Server, Slave, or Client Timings as explained below. This allows for setting the clock under certain conditions that would otherwise prevent a correction.

IMPORTANT: As of v5.1, none of these settings modify the **Never Step Clock** setting (see above). Note that this is a change in behavior from older versions. Enabling **Never Step Clock** effectively limits corrections to the **Maximum Slew Correction (milliseconds)** value, even if a larger correction would otherwise be permitted by **Override Max Disparity**.

- 0 or not present (Auto)
Domain Time will override the disparity settings during startup, on Clock Change Monitor event detection, receiving sync triggers/cascades from management components, or from Control Panel applet (CPL) signals.
- 1 (Always)
Domain Time will **always** override the disparity settings. This is the same as not having disparity settings at all. Always honors **Never Step Clock** setting.
- 2 (Never)
Domain Time will **never** override the disparity settings. Always honors **Never Step Clock**. This option may prevent your machine from syncing until you manually set the time to within the set Min/Max disparity range. If the machine is a Domain Time Server, it will normally refuse to serve the time until its own time has been set, so selecting a value of 2 may impact your entire network.

- 3 (Startup only)
Domain Time will override the disparity settings only until the first time after startup that it has set its own time correctly. Thereafter, it behaves as if you had set the option to 2.
- 4 (Limit CCM)
Clock Change Monitor signals do not override the disparity settings. Startup, management, or CPL signals **will** override the disparity settings.

Changes to this value take effect immediately. You do not have to stop and restart the service or reboot the machine.

Value Name: Override Sanity Checks

Value Type: REG_DWORD

Default Data: *False*

Options: *True or False*

Notes: To prevent accepting obviously-wild time corrections, Domain Time will (by default) refuse to set the time outside of a defined range of acceptable correction. Backwards-correction is limited to the build date of the software - 1 year. Forward-correction is limited to 11:59:59 on 12/31/2036 due to NTP and UNIX Year 2038 date calculation issues.

However, Windows itself will allow setting the local clock outside of this sanity-checked range. Set this value to *True* to permit Domain Time to set the clock to any time/date the operating system will allow.

CAUTION: Change this value only if you have a clear requirement to do so.

Value Name: Send Port Generic

Value Type: REG_DWORD

Default Data: 0

Notes: Domain Time uses several sockets for generic outgoing messages. By default, the port used is an ephemeral port assigned by the system. This is the proper behavior for client-server systems; only the server should have a fixed listening port, and clients should use ephemeral ports. However, in rare cases, other applications have high-number ephemeral ports hard-coded as their communications ports. If Domain Time happens to start first, and happens to obtain those particular ports, the hard-coded applications may fail.

Set this value to the beginning port number (n) of a range you want Domain Time to use for its generic outgoing sockets. Domain Time will attempt to use (n) through (n + 50) to bind its generic outgoing sockets. If none of the ports (n) through (n + 50) are available, Domain Time will revert to letting the system choose an ephemeral port.

IMPORTANT: Be very careful not to specify any well-known ports or IANA-registered ports for your range, and only set this value if you have a specific problem that you know will be solved by changing the ephemeral ports Domain Time uses.

Value Name: Server Answer IP

Value Type: REG_MULTISZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses

Notes: This value corresponds to the "Listen only on these addresses" list on the Network tab of the Control Panel applet. If this value is not present or is blank, Domain Time will answer on all IP addresses bound to all interfaces present on the machine. Otherwise, Domain Time will only bind to the IP addresses you provide. You may provide IPv4 or IPv6 addresses, and may also use NetBIOS or DNS names. The addresses/names you provide must exist and be permanently assigned to the machine. This setting is useful chiefly in situations where the machine is multihomed and you want Domain Time restricted to particular interface(s). This setting affects all listening ports for Domain Time Server, unless individual protocols are overridden (see below). You must restart the service (or reboot the machine) for changes to take effect.

Note: Because this value is highly machine-specific, it is not included in template imports or exports. You must set it individually on each machine.

As of version 5.2.b.20130221, you may also use CIDR notation to specify ranges of addresses. For example, 192.168.10.0/24 would bind to any address between 192.168.10.1 and 192.168.10.254, as long as one or more of those addresses was assigned to the machine. This is useful for machines using DHCP: you may restrict Domain Time to a particular network without knowing what IP the machine will have.

Value Name: Server Answer IP Override DT2

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with DT2/udp and DT2/tcp protocols

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by DT2 traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Server Answer IP Override NTP

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with NTP

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by NTP traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Server Answer IP Override PTP

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with PTP

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by PTP traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Server Settings

Value Type: REG_BINARY

Default Data: Varies

Notes: This value is a hex bitmap of various settings used by Domain Time Server, such as [Timings](#), [Corrections](#), and other miscellaneous settings.

Note: On Domain Time Client, these functions are located in the [Client Settings](#) key.

In general, you should not edit these settings manually. Use the Control Panel applet to configure your settings instead.

However, as of v5.2.b.20170922, the applet setting for Minimum Correction (MinDisparity) has been removed. The setting for this value defaults to 0x1, but if you upgraded over a previous version with a higher setting, you may edit the binary key to change it. The Minimum Correction setting is a DWORD, starting at offset 14, stored in little-endian order, as shown below.

The default should be 01 00 00 00, see this example:

Value data:

0000	D1	7A	C4	59	01	01	00	01	NzAY....
0008	00	00	20	1C	00	00	01	00
0010	00	00	19	00	00	00	58	02X.
0018	00	00	2C	01	00	00	01	00
0020	00	00	78	00	00	00	78	00	..x...x.
0028	00	00	07	00	00	00	00	00
0030	00	00	00	00	00	00	00	00
0038	00	00	00	00	00	00	00	00
0040	00	00	00	00	00	00	00	00
0048	00	00	00	00	00	00	00	00
0050	00	00	00	00	00	00	00	00
0058	00	00	00	00	00	00	00	00
0060	00	00	00	00	00	00	00	00
0068	00	00	00	00	00	00	00	00

Value Name: Service Installed

Value Type: REG_SZ

Default Data: N/A

Options: *True or False*

Notes: Used internally. Do not edit.

Value Name: Service Log Filename

Value Type: REG_SZ

Default Data: [not present]

Notes: Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename **domtimes.log** in the **%SystemRoot%\System32** folder. The complete path and filename must be specified (i.e. **C: \Windows\System32\domtimes.log**) and the drive specified must be a local drive.

Value Name: Service Running

Value Type: REG_SZ

Default Data: N/A

Options: *True or False*

Notes: Used internally. Do not edit.

Value Name: Set Processor Affinity

Value Type: REG_DWORD

Default Data: 0

Options: 00-FF (hex)

Notes: **Note:** This value has been deprecated in version 5.x and later; see the *Critical Timing Processor Limit* value instead.

If not present or set to zero, Domain Time will not attempt to restrict time-sensitive operations to any particular processor in a multi-processor system. In some systems, the majority of hardware interrupt handling occurs on only one processor (typically processor 0), so it may provide increased accuracy if Domain Time uses only other processors during time-sensitive operations. This value is a hex bitmap representing the processors in the system, with bit 0 representing the first processor, bit 1 representing the second processor, and so forth.

Value Name: Test Mode

Value Type: REG_SZ

Default Data: *False*

Options: *True or False*

Notes: Corresponds to the Test Mode checkbox on the [Advanced](#) property page of the Control Panel applet. If enabled (*True*), Domain Time will go through all the motions of obtaining the time and calculating variances, but will not actually set the clock. If disabled (*False*, the default value), Domain Time will set the clock after obtaining the time from its time source(s). Changes to this value only take effect after restarting the service.

Value Name: TIME/ITP Offset (seconds)

Value Type: REG_DWORD

Default Data: 2208988800 (decimal)

Notes: Used internally by the system. Do not change this value unless instructed to do so by tech support.

Value Name: Time Sample PreFilters

Value Type: REG_SZ

Default Data: HighLow

Options: Allowed options are HighLow, Latency, Delta, and Stratum. Prefilters are applied in the order listed; separate filter names with a comma or semi-colon.

Notes: Requires version 5.2.b.20150828 or later. This value controls the prefilters used to discard samples before applying statistical analysis. Prefilters only operate when there are five or more samples available for analysis, and are chiefly useful when the number of samples is very large, or the sources are unstable. It is best to leave this value at the default, which eliminates only egregious spikes. Statistical analysis of the entire group of samples usually performs better than prefiltering more samples out of the mix.

For example, HighLow,Latency would apply first the Highlow filter, then if at least five samples remain, the latency Filter. Stratum,Latency,Delta would first apply the Stratum filter;

then if at least five samples remain, the Latency filter; then, if at least five samples remain, the Delta filter. Changes to the list of prefilters are recognized only when parameters are reloaded (server stop/restart, machine reboot, a CPL-initiated sync, or a DTCheck /reload).

Prefilter operations are:

- **HighLow** (default) - Rejects the most extreme samples, based on absolute magnitude delta (max of 2 samples rejected)
- **Latency** - Rejects highest latency samples (max 1/3 of samples rejected)
- **Delta** - Rejects highest magnitude delta (max 1/3 of samples rejected)
- **Stratum** - Rejects all but the lowest-stratum samples present. Be very careful with this filter. Example 1: If your selection of samples includes one sample from a stratum 1 server, and ten more from a mix of stratum 2 and stratum 3 servers, then all but the single stratum 1 sample would be rejected. Example 2: If your lowest-stratum samples are a mix of stratum 2 servers, then all the stratum 2 samples would survive, but all your samples from strata 3 and up would be rejected. It is probably better to use the "NTP Client Max Stratum" value introduced in version 5.2.b.20110224 to control the highest stratum acceptable for NTP sources. The Stratum filter introduced here applies to all sources that report a stratum, including NTP, DT2, and PTP (the PTP "stepsAway" value is used to mimic NTP strata, as documented in the release notes for 5.2.b.20150516). Samples that do not report a stratum are not eliminated by this filter.

Value Name: Wait for Network Startup

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: Present in version 5.2.b.20151102 or later. If set to True, Domain Time will wait up to 30 seconds after boot for an IPv4 address to be assigned to the machine. At boot time, some network adapter drivers report ready before assigning IP addresses to an interface, even if the IPs are pre-configured as fixed addresses. DHCP-obtained addresses can take several seconds longer. The wait period helps ensure that Domain Time's initial enumeration of adapters and IPs is correct before protocol listeners or timechecks are started.

Change this value only if instructed by Technical Support

Time Sources

The Domain Time II Server Time Sources settings are located in this key:

HKEY_LOCAL_MACHINE

Software

Greyware

Domain Time Server

Time Sources

Broadcast

PTPv2 (IEEE 1588)

The values listed in the **Time Sources** registry key represent the time sources Domain Time uses to obtain the time. They correspond to settings on the [Obtain the Time](#) property page of the Control Panel applet or are otherwise automatically set. You should not change items in this section unless instructed by Tech Support or you are familiar with the specific function.

PTPv2 (IEEE 1588) key

Value Name: Current Master

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Current Offset (signed 64-bit)

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Current Offset Enabled

Value Type: REG_SZ

Default Data: False

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160922. When false, Domain Time will not update the current offset value in the registry (*Current Offset (signed 64-bit)* described above) or fire the offset-changed event. Note, this reverses the behavior introduced in version 5.2.b.20160415. To regain this behavior, set *Current Offset Enabled* to True, then trigger a sync or issue dtcheck -reload. See the SDK.DOC file included with the Software Development Kit (SDK) for details.

Value Name: Current PortState

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Duplicate Node Detection Enabled

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160415. Controls whether Domain Time will detect and prevent duplicate Clock Identities on the network.

Value Name: TAI-UTC Offset Discovered (seconds)

Value Type: REG_DWORD

Default Data: N/A

Options: N/A

Notes: Contains the current TAI-UTC offset (number of UTC leap seconds) discovered from the upstream Master or by importing a leapfile using the DTCheck utility. If this machine is acting as a stand-alone PTP Master, you may manually enter the number of leap seconds (create the key if it doesn't exist). The service must be stopped/restarted for changes to this value to take effect.

Value Name: TAI-UTC Offset Locked

Value Type: REG_SZ

Default Data: False

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160922. If changed to True, DT will not adjust its discovered TAI-UTC offset to match a new master advertising a different offset. The service must be stopped/restarted for this change to take effect. You should use this setting only if you have a broken PTP master advertising an incorrect TAI-UTC offset.

Domain Time II Client

Version 5.2

Domain Time II Client is a Windows system service that can be configured to obtain time from various time sources (such as GPS clocks and Internet time servers) and match the system clock to them with extreme accuracy and precision.

IMPORTANT: If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[Installation Instructions](#)

[System Requirements](#)

Configuring the Client

Once [installed](#), the **Domain Time II Client** service will start automatically when the system boots. All settings used by the service are read from the [Windows registry](#).

There are several ways to configure the Client's settings:

- Use the Domain Time II Client Control Panel applet on the Client machine itself.
- Remotely from another machine:
 - running Domain Time II (see *Connect to..* below).
 - using [Domain Time II Manager](#).
 - using the [RemoteCPL](#) tool (part of the Management Tools).
- [Import settings](#) from a saved configuration file.
- Use [Active Directory policies](#).

The following instructions describe the settings found on the Domain Time II Client Control Panel applet. Follow the links above for instructions on the other configuration methods. The applet can be configured whether the Domain Time II Client service is running or not.

Launch the applet

There are several ways to launch the Domain Time applet:

- From the System Tray (Notification Area) Icon: Double-click the Domain Time icon to launch the Domain Time II Client applet. You may also right-click the tray icon to launch the applet, as well as many other installed Domain Time II components and utilities from the context menu.
- From the Windows Control Panel: Click the **Domain Time Client** icon (it may be located in the **Clock, Language, and Region** section).

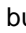
Note: On systems with User Account Control (UAC) enabled, you may need to *Shift+Right Click* the icon and choose **Run As...** or **Run As Administrator** from the context menu to launch the Control Panel applet.

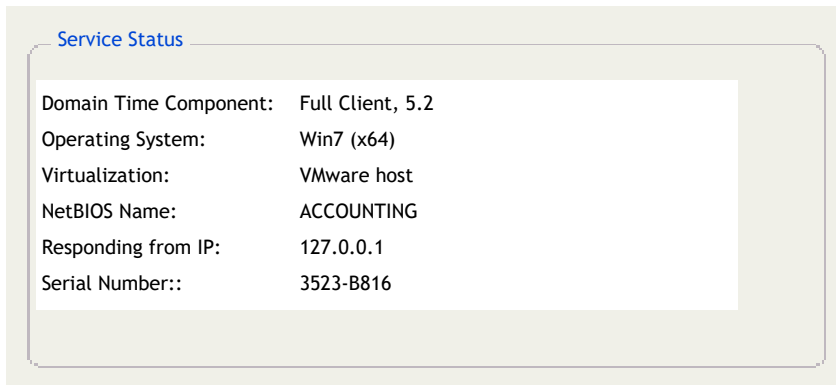
- From the the command-prompt: Launch the applet by typing **domtimec.cpl** in the Windows *Start --> Run* dialog or at a command prompt (the file itself is located in the \System32 folder).

The Control Panel Applet


The Domain Time applet has two panels. On the left is the navigation tree, which lets you pick a configuration property page to view by left-clicking an item in the tree. Right-clicking in the navigation tree will bring up a context menu with shortcuts to various functions, such as *Connect to another computer...*, the text and drift logs, online help, etc. On the right-hand side of the applet is the currently-selected configuration page.

Click the **Domain Time Client (local)** item on the navigation tree to display information about the installed service, including version information, Serial Number, stats, and Start/Stop control.

The *Service Status* display gives you a quick overview of the state of the Domain Time service. This section will be blank if the service is not started, and may take a few moments to display after a service restart. Click the  button to update the display.



Use the *Service Control* section to stop and restart the Domain Time service. Most changes you make using the Control Panel applet are dynamic and should not require you to restart the service.




Connect to another machine running Domain Time

You can also use the Domain Time Control Panel applet to connect to and configure other machines running Domain Time version 5.1 or later. This is particularly useful for quick configuration changes to a few machines, or for configuring Domain Time on Windows Server Core systems. If you need to configure many machines, you will want to use [Domain Time Manager](#) and/or use [Active Directory policies](#) instead.

You must be able to log on to the remote machine with an account that has administrative rights to the remote system. Your machine must also have sufficient network connectivity to authenticate with the remote system using Microsoft networking (see the [Planning](#) page for complete network information).

Right-click any item on the navigation tree and select *Connect to another computer...* from the context menu. If necessary, you'll be prompted to enter an administrative account and password to the remote machine.



When connected, the **Domain Time Client (local)** item in the navigation tree will change to display the name of the remote computer (you'll also see the name in the title bar and the background color of the menu tree will change). Now, any changes you make to the applet will be made directly on the remote system. All of the functions of the Control Panel applet, such as the log viewers, stats display, etc. also behave as if you were directly using the applet on the remote machine.

Select a computer to control

The control panel applet can work on remote machines, provided the machine can be reached by ICMP and has file and registry sharing enabled. You also need to have full administrative rights on the other machine through MS networking.

☐ Local computer (VISTA)

☒ Remote machine:

To disconnect from the remote system, you can either right-click -> *Connect to another computer...* again and choose **Local computer**, or simply close the applet.



Please read the [Installation Topics](#) pages before installing.

In most cases, you'll want to use [Domain Time Manager](#) to install and configure Domain Time Servers and Clients remotely across your network from your management workstation. If that's what you'd like to do, you can skip these installation instructions and read the [Network Rollout](#) documentation instead.

It's also possible to use an existing installation of Server install to another machine remotely using the command-line. This option is for advanced users. See the [Command-line Options](#) page for details.

If you use cloned OS images to install machines, please read [this article](#) from our knowledgebase about configuring Domain Time properly on your clone image.

NOTES:

- Windows Nano Server has special installation requirements. See the [Nano Server FAQ](#) for details.
- If you will be installing Domain Time Client on a virtual machine, see [this article](#) from our knowledgebase for more information on proper use with virtualization systems.
- If you are installing Domain Time Client onto a Windows Cluster, there are considerations regarding Cluster Service startup dependency on the Windows Time Service. Please see the **NoSync** section of the [Advanced -> Windows Time Mode](#) settings for more information.
- Check your routers and firewalls to be sure the ports for the time protocols you'll be using are open. Port 9909 TCP & UDP should always be open bi-directionally between Domain Time Servers and Clients. Port 123 UDP should be allowed if you will be using the NTP protocol for time synchronization. Ports 319 UDP & 320 UDP should be open bi-directionally for PTP use.

Domain Time version 5.2 and later includes a handy utility for adding the correct ports to the internal Windows firewall. Issue the following command from a command-prompt elevated with administrator privileges:

```
dtcheck /firewall:open
```

- If you will be installing Domain Time onto machines with AMD processors, we highly recommend you update your processor drivers (a.k.a. PowerNow!) to the current version for your operating system available from AMD's website to avoid known hardware timing issues. Please see [this article](#) from our knowledgebase for more info: [KB2007.817](#).

Installation/Upgrade

- To install or upgrade Domain Time Client directly to a single machine from the distribution setup files:
 - Run the Setup program from the CD to install the program. (If you have an older version of Client installed, Setup will give you an upgrade option. Your original configuration settings will be preserved during the upgrade). See [this page](#) for details on using the Setup utility.
 - Start the Domain Time Client applet from the icon in the Windows Control Panel to configure it.

Note: On systems with User Account Control (UAC) enabled, you may need to *Shift+Right Click* and choose **Run As...** from the context menu to launch the Control Panel applet. On Windows Server Core, type in `domt i mec. cpl` on the command line)


- Use the [Obtain the Time](#) property page to set Client to get the time:
 - from local GPS Network Time Servers, Domain Time Servers, or from other reliable network time sources, or
 - by auto-discovering a server (as configured by the Discovery Options on the [Obtain the Time](#) property page).
 - Using [Active Directory policies](#).
- Test your installation
 - Click the **Sync** button on the Control Panel Applet.
 - Click **View Log** button to see the service activity log. You should see messages indicating that the Domain Time service set its time correctly from the time source(s) you selected.

Removal

- Use [Domain Time Manager](#) to remove the program remotely.
- Or, use the **Programs and Features** (Add/Remove Programs) utility from the Windows Control Panel to remove the program.
- You may also use the original [Setup program](#) to remove the program. Run Setup and choose the Remove option.
- The program can also be uninstalled from the command-line. See the [Command-line Options](#) page for details.



Use this page to configure where Domain Time will get the time to set the local system clock.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Domain Time Client has four basic methods of obtaining the time; three of these are selected using the Control Panel applet as shown below. The fourth method is to assign time sources using Active Directory Group Policies. See the [Active Directory](#) page for more information on using Group Policies.

External Time Sources

- Set this machine's time by querying a list of servers (recommended)
- Set this machine's time from broadcast or multicast sources (deprecated)
- Discover sources automatically

Set this machine's time by querying a list of servers (recommended)

This selection instructs Domain Time to make outgoing unicast time requests to the servers you list on this page. Domain Time will query this list on the schedule you set on the [Timings](#) property page.

See the [Time Sources \(Unicast\)](#) section below for details on configuring Domain Time for this method.

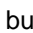
As of version 5.2, The IEEE 1588-2008 (PTP) protocol options will also become available when this method is selected. The protocol is enabled using the checkbox on this page (see the Additional Options section below), but is configured on its own dialog screens. Click the IEEE 1588-2008 (PTP) [Options](#) link to configure PTP. See the [IEEE 1588-2008 \(PTP\)](#) documentation page for details.

Set this machine's time from broadcast or multicast sources (deprecated)

This selection sets Domain Time to listen for incoming broadcast or multicast DT2/NTP time packets that are being transmitted from the sources you list on this page. Domain Time will set the local clock whenever it receives a time packet from the listed source(s). This selection is marked (deprecated) because very few administrators choose to use broadcast/multicast DT2/NTP for distributing the time, but the option is still supported.

See the [Time Sources \(Broadcast/Multicast\)](#) section below for details on configuring Domain Time for this method.

Discover sources automatically, instructs Domain Time Client to attempt to auto-discover a time server to use.

Client uses a reliable and sophisticated process to discover available servers to use. You may customize the discovery process by clicking the  button to pull up the **Discovery Options** dialog box:

Discovery Options

- Use last-known-good servers (recommended)
- Check DHCP option 004 for DT2 servers
- Check DHCP option 042 for NTP servers
- Ignore DT2 masters, slaves, or independent servers observed via cascades
- Broadcast/multicast for DT2 slave servers
- Broadcast/multicast for a DT2 master server

Broadcast/multicast for DT2 independent servers
Broadcast/multicast for NTP servers
Use servers in configured list of time sources
Use Windows domain hierarchy for DT2 or NTP servers
Require authentication from domain controllers

Client will try the selected discovery methods in the order listed and use the servers found according to the specified criteria.

If the **Use last-known-good servers (recommended)** option Client will remember which server(s) provided valid time. The last-known-good cache expires when the machine is rebooted or the service restarted. If "Use last-known-good servers" is unchecked, Client will perform the entire discovery process again at each timecheck.

If the **Use servers in configured list of time sources** is checked, Client will attempt to use the same list of Time Sources displayed when the **Set this machine's time by querying a list of servers** radio button is selected on the *Obtain the Time* property page.

If all discovered servers becomes unavailable, Client will automatically re-start the discovery process to find another server to use.

As of v5.2.b.20110831, you have the option of enabling/disabling the Windows Authentication function when communicating with domain controllers discovered using the Windows domain hierarchy. This option was automatically selected on previous versions. See the [Symmetric Keys](#) page for more information.

DHCP Options

You may use DHCP to assign time sources to Clients. If enabled in the **Discovery Options** above, the Client will do a DHCP discovery broadcast to find a local DHCP Server. If the IP address of a time source is defined in DHCP Option 004 or 042, the Client will use the specified source.

Note: You may assign a time server using DHCP Options whether or not the machine on which Client is running is using DHCP to assign an IP address. Client's DHCP discovery broadcast to determine the value of Options 004 and/or 042 is completely separate from the IP address assignment used by the network stack.

- **Option 004** ("Time Servers") is used only for discovering DT2 servers. If a server is listed in option 004 that doesn't support DT2 UDP, it will be ignored.
- **Option 042** ("NTP Servers") is used to discover both NTP servers and DT2 servers. If a server is listed in option 042, it will be checked for NTP first. If NTP fails, it will be checked for DT2 UDP. If it does not provide time under either of these two protocols, it will be ignored.

Additional Options

The following options may be available depending on which of the three basic methods of obtaining the time you've selected (see above):

Analyze time samples from all servers and choose the best
If all listed servers fail, try to discover sources automatically
Match server's timezone if available (DT2 protocol only)

About Time Samples

When obtaining time from external time sources, Domain Time uses Time Samples to determine the correct time.

A time sample is collected either by

Analyze time samples from all servers and choose the best

This controls whether Domain Time applies advanced analysis algorithms to all of the collected time samples.

When this box is checked, Domain Time contacts all of the listed servers to collect a group of time samples (if you're querying servers) or waits until it has collected the specified number of incoming time packets (if you're using broadcast/multicast sources). It then performs statistical analysis on the collected samples to determine the reliability and uses the most reliable samples to derive the correct time.

See the "About Time Samples" sidebar on the right side of this page for more information and rule-of-thumb suggestions on acquiring time samples.

If you are collecting multiple samples from unicast or broadcast sources using the NTP or DT2 protocols, checking this box will almost always improve your machine's accuracy and reliability.

Note: If you are using the IEEE 1588-2008 (PTP) protocol to synchronize your time, including other time sources in the time calculations can cause inaccuracies at very high levels of precision. Therefore, as of version 5.2.b.20150828, Domain Time automatically excludes all other sources from time calculations when using PTP, falling back to them only if PTP fails, so you may leave this box checked. However, on versions prior to 5.2.b.20150828, you must manually uncheck this checkbox to prevent skewing of the time from additional non-PTP sources.

If this box is unchecked, no comparative analysis among samples is performed. In addition, the list of time servers to query becomes a **fallback-only list**. In other words, the Server will only contact the first listed time server. This server will always be used unless it becomes unavailable, at which point the next listed server will be used. If that server is unavailable, the next server in the list will be tried, etc. When the first listed server becomes available again, the Server will revert to using it exclusively.

When analysis is enabled and more than one time source is used in a time calculation, the logs (when set to the default "Information" detail level) and other display fields without room for multiple entries will show the source for the time as "Averaged Time", otherwise the IP address of the single time source used will be displayed.

If all listed servers fail, try to discover sources automatically

This selection causes the Client to use the *Discover Sources Automatically* process (described above) to try to automatically find an available server if it cannot communicate with your specified time sources.

Do not enable this option if you always want your Client to attempt to use only the specified sources under all circumstances.

Match server's timezone if available (DT2 protocol only)

When selected, Client will change the local machine's Windows timezone settings to match the timezone setting of the Domain Time Server it contacts.

Note that this is a global change to the operating system which will affect all programs that display local time (the same way that manually changing the timezone using Windows' Date & Time Applet does).

(a) sending a unicast time request to a time server and receiving a unicast reply, or by (b) accepting an incoming time packet sent from a broadcasting or multicasting server.

By default, Domain Time analyzes the collected time samples using sophisticated statistical methods to reject bad samples and derive the correct time. It then sets the local clock to the correct time using the greatest accuracy possible.

Any single time sample from any time source may not reflect the correct time, either because of network delays, operating system load, or many other transient causes. Therefore, it's usually a good idea to collect more than one sample. If querying a list of servers, you may specify multiple time servers and also set the number of samples to request from each source. If accepting incoming broad/multicast packets, you can specify the number of samples that must be received from the source before making a correction.

In general, time will be more accurate the more samples you collect; however, there is a point of diminishing returns. Each sample takes a fixed amount of time to collect. If the overall time taken to collect the samples is too long, the clock may drift significantly in the interim so that any additional accuracy you obtain from the larger number of samples will be offset by the additional drift.

A good rule of thumb for querying servers is to configure at least three or more sources, which provides additional sanity checks and fallback in case any one server is unavailable. Then, specify an odd number of samples from each server; three samples each is a good choice. An odd-number of samples makes the calculations necessary to reject a bad ticker more likely to be

In order for this feature to work, the Domain Time Server you are contacting must be set to recommend the Time Zone to Clients (see the *Allow clients to match this server's timezone* setting on the Server's [Recommendations](#) property page) and the Client must be using the DT2 protocol to synchronize its time with the Server.

Time Sources (Unicast)

If you have selected the **Set this machine's time by querying a list of servers** method of obtaining time, Domain Time will query the machines you list (and enable) on this page for the current time.

Server Name or IP	Protocol	Auth	Reps	Delay	Comment
<input checked="" type="checkbox"/> time-a.nist.gov	NTP	None	1	n/a	
<input checked="" type="checkbox"/> time-b.nist.gov	NTP	None	1	n/a	
<input checked="" type="checkbox"/> nist1.symmetricon.com	NTP	None	1	n/a	
<input checked="" type="checkbox"/> tick.greyware.com	DT2-UDP	None	3	512	
<input type="checkbox"/> tock.greyware.com	DT2-HTTP	None	1	n/a	

Unicast Time Source List [\[Click for larger size\]](#)

The machine list can consist of servers that use the NTP, DT2 (UDP or TCP), or DT2-HTTP protocols. As of version 5.2, you may also select the IEEE 1588-2008 (PTP) Precision Time Protocol as a time source. See the [IEEE 1588-2008 \(PTP\)](#) page for details on using the Precision Time Protocol.

You may add machines to the list manually or by scanning for them on your network automatically.

- To easily identify available time servers on your network, click the [Local Time Servers](#) link at the bottom of the list box. This brings up the **Time Sources Search** dialog, where you can scan your network for time servers and then add your choice(s) to the Time Sources list automatically.

Server IP Address	Protocol	Discovery Method
<input checked="" type="checkbox"/> 192.168.10.2	DT2-UDP	discovered indie
<input type="checkbox"/> 192.168.10.2	NTP	discovered ntp server
<input type="checkbox"/> 192.168.10.3	NTP	discovered ntp server
<input checked="" type="checkbox"/> 192.168.198.1	DT2-UDP	discovered indie
<input checked="" type="checkbox"/> fe80::a5cd:49e5...	NTP	discovered ntp server

Search for Time Sources Automatically [\[Click for larger size\]](#)

accurate. You can then use trial-and-error to determine if adding more samples increases your accuracy.

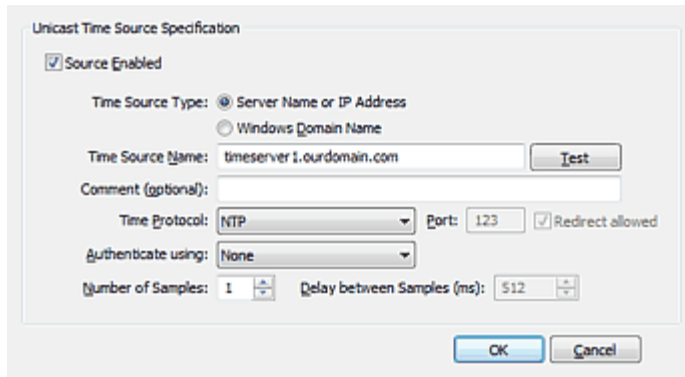
If taking multiple samples from any time server, take care to request a reasonable number of samples and set a delay between the samples to avoid being flagged as making a Denial-of-Service attack.

If you're using broadcasting/multicasting, you can require that multiple samples be collected before setting the time. However, multiple samples may or may not increase accuracy, depending on a number of factors. Consider this option only if the broadcast/multicast time pulse is occurring rapidly enough to collect your required number of samples before the clock drifts outside your target tolerance.

As of v5.2.b.20190701, Domain Time Client and Server support Windows NDIS software timestamping, which allows measurement of network stack delays. Software timestamping is only available on Server 2019 (or newer) and recent updates of Win10. You may want to experiment with this setting to see if it improves your accuracy. See [KB2019.708](#) for more information.

- To manually add a time server to your list of time sources, click the dialog.

button. This brings up the **Add Time Source**

The image shows a Windows dialog box titled "Unicast Time Source Specification". It has a checkbox labeled "Source Enabled" which is checked. Below it are two radio buttons for "Time Source Type": "Server Name or IP Address" (selected) and "Windows Domain Name". The "Time Source Name" field contains "timeserver1.ourdomain.com" and has a "Test" button next to it. There is an empty "Comment (optional):" field. The "Time Protocol" is set to "NTP" in a dropdown menu, with a "Port" field set to "123" and a checked "Redirect allowed" checkbox. The "Authenticate using:" dropdown is set to "None". At the bottom, "Number of Samples" is set to "1" and "Delay between Samples (ms)" is set to "512". "OK" and "Cancel" buttons are at the bottom right.

[Add Unicast Time Source](#) [\[Click for larger size\]](#)

If you will be using time servers over the Internet, please click the [Public Time Servers](#) link to find reliable servers.

Use the **Time Source Type**: radio buttons to indicate whether you want to contact a server directly using its machine name or IP address, or to automatically find and use the domain controller holding the PDC Emulator role on the specified Windows domain.

If you enter a machine name in the **Time Source Name** field, it must be resolvable to an IP address using DNS, WINS, Active Directory, from the HOSTS file, etc. If entering the IP address, you may use either the IPv4 or IPv6 address of the server.

You may use the **Comment** field to annotate this entry, if you want.

Note: As of v5.2.b.20190701, NetBIOS and DNS names are checked first for IPv4, and only use IPv6 if the IPv4 lookup fails. (If you use an IP literal, Domain Time will use the protocol family associated with what you entered, and the information in this section does not apply.)

To force a NetBIOS name or DNS name to use either IPv4 or IPv6, enter either the text "IPv4" or "IPv6" anywhere in the comment field. For example, if your source is specified as `ntp.mydomain.com` without specifying either IPv4 or IPv6 in the Comment field, Domain Time will first try to resolve the name using IPv4. If that lookup fails, Domain Time will try to resolve the name using IPv6. If, however, you put either "IPv4" or "IPv6" in the comment line, Domain Time will look up `ntp.mydomain.com`'s IP address using only the IP family you specify.

Use the **Time Protocol**: drop-down list to indicate which time protocol to use when contacting this server. You can use DT2-UDP, DT2-TCP, DT2-HTTP, or NTP.

The **Port**: field displays the IP port used by the selected protocol. This is a display-only field for all protocols except the DT2-HTTP protocol. The DT2-HTTP protocol port may be changed to match the DT2-HTTP listen port set on the [Serve the Time](#) page of the target Domain Time II Server. The default value for this is port 80. The **Redirect allowed** checkbox specifies whether the DT2-HTTP time requests will honor HTTP 301 and 302 redirects. Only one level of HTTP redirection is permitted.

The **Authenticate using**: drop-down list selects which authentication key to use when exchanging packets with this server. A key will show up in the list if it has been configured on the [Symmetric Keys](#) property page of the Control Panel applet.

Domain Time supports MD5 symmetric-key authentication compatible with NTP version 3 and later (AutoKey is not supported), and as of v5.2.b.20170922, SHA1 authentication as well. Windows Authentication compatible with Windows Time NT5DS-mode timestamps is also supported. Either authentication method can be used over any supported time


protocol (NTP, DT2-UDP, etc.) See the [Symmetric Keys](#) page for details on using authentication.

Hint: When possible, be sure all of your time systems are working correctly before enabling authentication. Authentication requires a correct setup on both ends of the connection, and changes at either end can cause a previously-working connection to fail. Disabling authentication temporarily should always be one of the first steps when troubleshooting a connection issue.

Number of Samples: sets how many individual requests Domain Time will make of this server during each time check.

CAUTION: Take extreme care with this setting. Many time servers have Denial-of-Service (DOS) protection to prevent abuse. Issuing too many time requests in a row to one server over a short period of time can cause your machine to be locked out or even be permanently blacklisted.

Use the **Delay between samples (ms)** setting to space out your sample requests over a reasonable length of time. You may want to contact the administrator of any time server you will be using to find out what the acceptable retry period is on that server. Another option is to use fewer samples per server and simply check against more servers if you need to increase your sample count.

Click the  button to be sure the server you've selected is reachable using the protocol specified.

Note: The Control Panel applet you're using for the test is running in the foreground security context of the currently logged-in user, but, in normal operation, the Domain Time service will use the context of the background service account under which it runs (by default, LocalSystem). There are some circumstances where the foreground test will succeed in contacting a source but the Domain Time service will fail, or vice versa. If this occurs, check your firewall and security settings to allow the Domain Time service the necessary network access to send/receive time protocols.

Import/Export

You may easily save or restore the Time Sources list settings by clicking the [Import/Export](#) link. You can use this function to quickly update just the list of time sources used without affecting any other settings.

This function does not preserve IEEE 1588-2008 (PTP) settings, it only saves/restores machines in the time sources list.

If you have multiple machines to update or need to configure other settings, you should use the full Import/Export features found on the [Advanced -> Import/Export](#) property page or use Domain Time II Manager's [Templates](#) feature.

Time Sources (Broadcast/Multicast)

If you have selected the **Set this machine's time from broadcast or multicast sources** method of obtaining time, Domain Time will listen for broadcast or multicast DT2/NTP packets from the listed time sources and extract time data from them.

In addition to the options described above, you'll see the following settings when you select **Set this machine's time from broadcast or multicast sources**:

Only accept signed packets

Log rejected packets Samples required for sync:

Only accept from well-known source port

About xcasting

Using broadcast or multicast (sometimes referred to in this document as "xcast") time packets to obtain time has distinct advantages and disadvantages.

One advantage of using xcast to obtain time is that there is often lower processing overhead than when you're sending unicast time queries to a server. The unicast method must send queries to various servers, receive the responses,

Only accept signed packets

If checked, only packets using authentication will be accepted. See the [Symmetric Keys](#) page for more information on packet authentication.

Log rejected packets

When checked, rejected packets will be noted in the log.

Samples required for sync:

This sets how many time packets with time data must be received before a correction occurs.

This is also the number of samples used for analysis if the **Analyze time samples and choose the best...** checkbox (discussed above) is checked.

Be careful not to specify a number of samples that would result in long period before the clock is corrected, since the clock may drift significantly before all the samples have been collected.

Only accept from well-known source port

If checked, only packets originating from port 123 UDP (if using the NTP protocol) or port 9909 UDP (if using the DT2 protocol).

Use this setting with caution, since the default behavior of many servers is to send outgoing traffic from a random source port.

compensate for latency on each sample, and then analyze the samples to determine the correct time. The xcast process is comparatively simple. The listening machine merely accepts the time presented in the packet as valid, subtracts out the estimated latency of the connection (see Estimated delay below) and sets the clock.

This can be useful in tightly-controlled networks where network propagation delays are known and unchanging. Under those circumstances, it is sometimes possible to achieve higher accuracy on clients by using a very rapid xcast pulse rather than by having the clients each make many individual requests of the server.

However, there are many disadvantages to xcast time under normal network operations. The most significant of these is that network conditions are rarely static, and the latency between time server and client can vary substantially in a short period of time. This can severely affect the accuracy of the incoming time stamp, causing jumps in time. In addition, broadcast time can be very susceptible to interference from rogue broadcast servers on the network, packet-spoofing (although signed packets can help avoid this), and other disruptions which can adversely affect reliability.

In most cases, modern time request/reply protocols with sophisticated round-trip latency detection such as NTP and DT2 are the better choice. However, broadcast time is still used by some legacy equipment, so it may be the only time synchronization option available for those devices.

Domain Time allows you to configure to receive broadcasts and/or multicast packets using either the

Broadcast/Multicast Time Source List

Shows the currently configured time sources. Domain Time will only listen for time packets from sources listed (and enabled) here.

Selected broadcast/multicast time sources. Samples from unchecked servers will be ignored.


Server IP Address	Protocol	Estimated Delay
<input checked="" type="checkbox"/> 192.168.10.4	DT2-UDP	1

[Add](#) [Delete](#) [Edit](#) [Local Broadcast Sources](#)

[Broadcast/Multicast Time Source List](#) [\[Click for larger size\]](#)


You may add machines to the list manually or listen for broadcasting servers on your network.

- To easily identify available broadcast time servers on your network, click the [Local Broadcast Sources](#) link at the bottom of the list box. This brings up the **Time Sources** dialog, where you can listen for broadcast sources on your network and then add your choice(s) to the Time Sources list automatically.



[Discover Broadcast Time Sources Automatically](#) [\[Click for larger size\]](#)

- To manually add a broadcast time server to your list of time sources, Click the [Add Time Source](#) button. This brings up the **Add Time Source** dialog.



[Add Broadcast/Multicast Time Source](#) [\[Click for larger size\]](#)

Use the **Time Source Type**: radio buttons to indicate the type of time packet to listen for from this server. You can accept either DT2-UDP, NTP, or both.

IMPORTANT: Only one service may own a particular port. If you will be accepting NTP broadcast packets with Domain Time, you will need to disable any other service that may be using the NTP port (such as the Windows Time service).

You may use either the IPv4 or IPv6 address of the broadcast server in the **IP Address**: field.

You may use the **Comment** field to annotate this entry, if you want.

Estimated delay is the expected amount of latency in milliseconds a time packet will encounter between the transmitting server and this machine. Domain Time will adjust the time contained in the timestamp by subtracting this value to improve accuracy. The closer this value is to the actual latency on your network connection, the more accurate your time synchronization will be. You may enter this value yourself, or click the [Calculate](#) button to calculate it for you.

You may need to adjust this value if the overall propagation delay changes on your network.

Import/Export


You may easily save or restore the Time Sources list settings by clicking the [Import/Export](#) link. You can use this function to quickly update just the list of time sources used without affecting any other settings. If you have multiple machines to update or need to configure other settings, you should use the full Import/Export features found on the [Advanced -> Import/Export](#) property page or use Domain Time II Manager's [Templates](#) feature.

NTP or DT2-UDP protocols. There are efficiencies in the DT2-UDP protocol that result in slightly-higher accuracy than NTP overall; otherwise, the packets function very similarly.



Settings on this page control how often time checks are performed when querying servers for the time (if using normal NTP or DT2 protocols) or how often samples are coalesced for statistics/alerting if using PTP.

Note: The settings on this page do not apply if you are using the Broadcast/Multicast method of obtaining network time.

If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

There are two scheduling options for determining how often Domain Time checks the time/reports statistics:

Check Interval when able to get and correct the time

Variable - check as often as needed to maintain approximately milliseconds sync with server

Fixed - check once every

DaysHoursMinsSecs

Variable - check as often as needed to maintain approximately milliseconds sync with server

When this option is selected, Domain Time will automatically adjust how often it synchronizes with time sources to attempt to keep the clock within the threshold limit you set.

The wait period between time checks is called the *window size*. You can see the window size Domain Time is currently using by examining the [Text Log](#). Domain Time will adjust the window size based on how accurate previous time checks have been. If previous time corrections have been small the window size will be increased, and vice versa. The range of adjustment for the window size is between 15 seconds to 2 hours. On most machines, it will average between 10 - 30 minutes.

The **Variable** scheduling method is intended for use on machines with relatively constant clock drift and moderate accuracy requirements (where the acceptable tolerance for clock drift is more than ~25ms). This method is a good for general-purpose use, primarily on Clients, since it strikes a good balance between maintaining the target accuracy while minimizing network traffic.

Variable is not a good selection if the machine is under heavy and/or variable load that causes the clock to drift by significant amounts on an irregular basis. Since Domain Time may select a large window size if the clock on the machine has been well-behaved, anything that causes sudden clock drift during the Window period between checks can cause the clock to drift outside the

Finding the "Sweet Spot"

The system clock on every machine runs at a different rate because of differences in operating system, applications, hardware, and environment. Even when well-managed by a time program such as Domain Time, the clock will always eventually drift either slower or faster than the actual time,

Since the clock can't be made to run at a perfect rate, it is necessary to correct it when it drifts. In general, the more often you can correct the clock, the more accurate it will be. The corollary to this is that the worse the clock drifts, the more often it must be corrected.

The goal of time correction is to synchronize often enough to keep the system clock within your accuracy target but not so often as to generate excessive network traffic or system overhead. We refer to this ideal rate as the synchronization sweet spot.

In many cases, Domain Time does a very good job of automatically discovering the sweet spot (using the Variable scheduling method). It also, by default, uses an overall clock-rate adjustment, to train the clock to run

specified threshold before the next correction. If this describes your machine, you should use the **Fixed** schedule instead.

Fixed - check once every

DaysHoursMinsSecs

If this option is selected, Domain Time will synchronize regularly on the schedule you specify.

This method is a good choice when you want to discipline the clock to stay within very tight synchronization tolerances. It's also the best choice for machines with highly variable load, poor timekeeping hardware, or any other issue that causes significant clock drift.

You should check the time often enough to keep your clock within your desired accuracy.

Since having highly-accurate time at all times is usually more critical on Servers, you will likely want to check often using a fixed schedule on Servers.

CAUTION: Take care with this setting. Many time servers have Denial-of-Service (DOS) protection to prevent abuse. Issuing too many time requests in a row to one server over a short period of time can cause your machine to be locked out or even be permanently blacklisted.

This problem can be exacerbated if you have opted to collect multiple samples per time source. See [About Time Samples](#).

See the "Finding the Sweet Spot" sidebar on the right for more info on picking the correct sync rate.

Check Interval when getting the time fails

If Domain Time cannot obtain the time, it should try again every: DaysHoursMinsSecs

You may set a different rate for Domain Time to use if it cannot contact any time source.

If Domain Time cannot obtain the time, it should try again every:

DaysHoursMinsSecs

sets how often Domain Time will retry its time sources if it is unable to successfully obtain the time. You will probably want to check more often than during normal operation (unless you're already using a frequent synchronization schedule) to reacquire the correct time quickly when your time source(s) become available.

The same caution about synchronizing too often against your time sources (discussed above) applies.

Accept server's recommended settings (if provided)

more accurately over time.

However, the more accurate you need the clock to be (or the worse the clock itself is), the more difficult it is for these algorithms to make correct decisions to compensate correctly for drift.

In those cases, you will need to manually set a Fixed synchronization rate, using trial-and-error to find the sweet spot. You may want to start with a reasonable rate such as 5 minutes, and then if that's not sufficient, try every 3 minutes, then 1 minute, etc. On machines with highly-variable drift, you may also need to disable Domain Time's long-term clock adjustment function (see the [Correction Reduction](#) section of the [Clock Control](#) page).

Even with severe correction, some systems simply cannot be disciplined enough for every purpose. For example, virtual machines are often too inherently poor at timekeeping to be used for time-critical systems and you must change to physical hardware to achieve your desired accuracy. However, in most cases, you will usually be able to find the sweet spot for your systems by adjusting the synchronization rate appropriately.

Domain Time Client can inherit its timing settings from the Domain Time Master Server.

When the **Accept server's recommended settings (if provided)** box is checked, Client will adopt the timing settings specified on the Domain Time Master Server's [Client Timings Recommendations](#) page.

Client will obtain these settings by synchronizing either directly with the Master or with any Slave server (Slaves automatically replicate the Master's settings).


Having this option enabled on your Clients allows you to set the overall accuracy of your time hierarchy from one central location. Note that any manual settings you make on this page will be overwritten when this box is checked the next time the Client synchronizes. Remember to uncheck this box if you want to make manual adjustments the timings.

Note: These settings will only take effect on clients if:

- The *Recommend timings and correction limits to clients that ask for guidelines* checkbox is checked on the Master Server's [Recommendations](#) page.
- The *Accept server's recommended settings (if provided)* checkbox is checked on the Client's [Timings](#) page.



Settings on this page control what corrections to the system clock are acceptable.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings may adversely affect your clock accuracy or even prevent clock corrections entirely.

IMPORTANT: Settings on this page do not apply (or have minimal effect) if you are synchronizing using [IEEE 1588 \(PTP\)](#).

AS OF v5.2.b.20170922, the Minimum Correction setting is obsolete and has been removed from the applet

The documentation in this section is for reference purposes on older versions only. If you still need to change this value, please see the [Client Settings](#) registry key.

Minimum Correction

Variances smaller than milliseconds will not cause a correction (unless overridden below)

This setting controls the minimum amount the clock must be off from the time source(s) before it is corrected during a time check.

If you are using a variable synchronization schedule (see the [Timings](#) page), you will probably not want to incur the extra overhead of correcting the clock if the variance found during a time check is smaller than your selected synchronization target.

For example, if you have instructed Domain Time to aim for a target variance of 25 ms, you do not need to make a clock correction if the detected clock variance during a time check is only 10 ms. Making a correction in this case would only result in extra processing overhead and clock slewing without much affecting your overall accuracy. So, in general, if you are using variable targeting, you will want to set this value to be the same value or less than your target variance.

However, if you are using a fixed sync schedule, you will want to be sure the clock is corrected to the maximum accuracy on every synchronization. In that case, this value should be set to 1 millisecond. This also enables Domain Time's high-precision sub-millisecond alignment function, so that any variances detected that are less than 1 millisecond will have the clock aligned to match, giving you an added order of magnitude of possible accuracy.

This setting can be overridden under certain circumstances so that the clock can be forced to be corrected. See the **Limit Override** section below for details.

Maximum Correction

No maximum correction

Variances larger than _____ will not cause a correction (unless overridden)
HoursMinsSecs

This setting controls the maximum variance that should be corrected during a time check.

This setting provides a vital sanity check to prevent wild time changes in the event your time source(s) provide a rogue time value (such as sometimes occurs when bounds limits are exceeded or error conditions occur in time clocks or the network).

For example, assume you have restricted this value to not allow corrections for variances larger than 2 hours. If a time source suddenly goes crazy and provides a time/date from 1980, the rogue time correction will be rejected.

The default setting for this value is fairly generous, so you may want to restrict this more in your environment. Do be careful to not restrict this value too tightly. If you have clocks on the network that drift significantly under normal circumstances without restarting (such some laptops do when resuming from sleep modes), setting this value too low may prevent them from ever correcting the clock until they are rebooted.

This setting can (and usually must) be overridden under certain circumstances so that the clock can be forced to be corrected. See the **Minimum/Maximum Limit Override** section below for details.

This setting also interacts directly with the clock slewing settings (which control whether corrections are made by slewing or stepping). See the [Clock Control](#) page for details.

Limit Override

Override the minimum and maximum:

For sync signals, at startup, during training, and when triggered by Clock-Change monitor

For sync signals, at startup, during training, but NOT when triggered by Clock-Change monitor

Only on first correction after machine startup (within _____ seconds of boot)

Use these settings to control when Domain Time will override the correction limits to force a time correction.

The default selection is usually the best option since there are a number of situations where you typically want the time to always be corrected regardless of how far off it may be from the correct time such as:

- when the time service is started
- when triggered to force a correction
- when the clock is being trained (see the [Correction Reduction](#) section of the Clock Control page)
- when Doman Time's Clock-Change Monitor detects that a user or application has unexpectedly attempted to change the time

However, your particular needs may require the ability to restrict corrections even further. If so, you will want to select one of the other listed options. Do be **sure** you fully understand the effects of this selection if you change it from the default.

This setting interacts directly with the clock slewing settings (which control whether corrections are made by slewing or stepping). You may use the override settings in combination with slewing/stepping limits to ensure that corrections are made only under the precise conditions you desire. See the [Clock Control](#) page for more details.

Advanced Sample Validation

Discard time samples that exceed milliseconds of historical average variance
Discard time samples whose latency exceeds milliseconds, regardless of history

These settings set boundaries on the maximum variance and/or latency permitted in individual time samples.

Discard time samples that exceed milliseconds of historical average variance

Checking this box enables an additional check which may help protect against significantly delayed or skewed time samples. See [About Time Samples](#) on the **Obtain the Time** page for details on how time samples are used.

This setting is turned off by default to minimize overhead, since the default expectation is that your network and time sources will generally be well-behaved. However, if you experience unusual spikes in the time from otherwise reliable sources, you may want to enable this setting to help screen out the errant samples that would otherwise skew your time calculations.

When enabled, Domain Time will keep a historical record of time samples from your selected time sources. It will then reject any time samples that exceed the historical average of the time source by the threshold value you select here.

For example, assume you have set this threshold to 500ms. The historical average variance of time server tick.mydomain.com to date has been +50ms but suddenly a time sample is received with a variance of -475ms. This new sample varies from the historical variance by more than the 500ms threshold value you've set, so the sample will be rejected.

CAUTION: As with other settings on this page, you should only enable this function if you fully understand the ramifications. If you set this value incorrectly, you may end up rejecting so many samples Domain Time will not be able to correctly identify the correct time, or may even be unable to set the clock entirely. Use the Trace or Debug log detail of the [Text Log](#) to see if samples are being rejected and how the accepted samples are being analyzed.

Discard time samples whose latency exceeds milliseconds, regardless of history

This setting is similar to the function described above, except that it rejects individual samples based on a latency limit you specify. This feature was introduced in version 5.2.b.20110309.

CAUTION: As with other settings on this page, you should only enable this function if you fully understand the ramifications. If you set this value incorrectly, you may end up rejecting so many samples Domain Time will not be able to correctly identify the correct time, or may even be unable to set the clock entirely. Use the Trace or Debug log detail of the [Text Log](#) to see if samples are being rejected and how the accepted samples are being analyzed.



Settings on this page control how Domain Time listens and sends traffic to the network.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change.

These settings are unique to each machine, and therefore cannot be saved or imported using the [Import/Export](#) utilities.

Network Listen

Address Family:	Both IPv4 and IPv6	<input type="checkbox"/> Join DT2/NTP IPv4 multicast groups
		<input type="checkbox"/> Join DT2/NTP IPv6 multicast groups
		<input type="checkbox"/> Initiate rebind and resync if IP address changes
		<input type="checkbox"/> Enumerate multicast interfaces during IPv4/IPv6 bind
<input type="checkbox"/> Listen on all IP addresses available		<input type="checkbox"/> Reply to multicasts using incoming interface if possible
<input type="checkbox"/> Listen only on these addresses:		
Enter one IP address, hostname, or mask per line.		
You may use any combination of IPv4 and IPv6 addresses. The addresses you enter must exist and be permanently assigned to this machine		
You may also use CIDR notation to specify any/all addresses matching the mask you supply.		

Use these settings to tell Domain Time which IP addresses to use when listening for incoming network traffic and whether to join multicast groups.

IMPORTANT: Keep in mind that there can only be one network service listening on any one network port on any one IP address. The DT2 protocol (both 9909 UDP and 9909 TCP) will be reserved by the Domain Time service and cannot be disabled.

IPv6 requires operating system support, which is present by default in Vista or above, but must be specifically installed/enabled on XP. Domain Time will function in IPv4-only mode if IPv6 is not present. If both are present, you may choose which to use, or let the system figure it out.

Domain Time assumes your TCP/IP and Windows networks are configured properly, i.e. name resolution is functioning, rules are in place to permit traffic through switches, routers, and firewalls, any Active Directory/Domain structure is functioning correctly, etc.

Join DT2/NTP IPv4 multicast groups

Join DT2/NTP IPv6 multicast groups

These checkboxes control whether Domain Time will join multicast groups to listen for either NTP or DT2 protocol multicast traffic.

Multicasts may include not only time sync information, but client discovery packets from Manager and Audit Server. You should not disable this function unless you have a compelling reason to do so. Note this setting only applies to joining groups. You can control the multicast address and whether multicasts are sent on the [Broadcasts and Multicasts](#) property page.

Initiate rebind and resync if IP address changes

If this checkbox is checked, Domain Time will rebind interfaces if it notices that an IP address has been changed/added while the service is running. Use this option if your machine is likely to have its list of IP addresses change during operation.

Enumerate multicast interfaces during IPv4/IPv6 bind

Instructs Domain Time to include all addresses capable of multicast when it binds to the network interfaces. Check this box if the machine is multihomed.

Reply to multicasts using incoming interface if possible

Use this option to reduce unnecessary multicast traffic. When checked, and if the machine is multihomed and listening using a specified list of IP addresses (see below), Domain Time will attempt to reply to multicasts (such as PTP messages) over the same interface on which they were received. If unchecked, Domain Time will attempt to reply on all known interfaces. This feature has no effect if the **Listen on all IP addresses available** option is selected. Check this box only if the machine is multihomed and you have addresses listed in the **Listen only on these addresses:** box.

Listen on all IP addresses available

Listen only on these addresses:

By default, Domain Time attempts to listen for network traffic on all available interfaces. However, you can restrict this to specific IP addresses and/or address ranges if necessary.

IMPORTANT: Be sure to include the hostname **localhost** in your list of IP addresses to ensure that foreground Domain Time processes (such as the Domain Time applet) can continue to talk to the Domain Time service correctly.

As of v5.2.b.20190701, you may enter comments in the specified listen-only list. Comments are defined as text following a hashtag or semicolon. (If the hashtag or semicolon is the first character, the entire line is considered a comment.) For example, you may use this syntax:

```
; These are our subnets:
172.16.13.0/24 # main network
192.168.33.0/24 # internal network
```

Comments in the list (other than commenting out an entire line) are not backward-compatible with previous versions of Domain Time, so don't use them in templates until all of your machines have been upgraded.

Network Send

DT2:

NTP:

TIME/ITP:

Use blank (or zero) to mean the system should choose a random source port

Use this section to specify a fixed source port for time protocol traffic.

These settings should be left blank unless you have a specific requirement to send traffic from a specific source port. If you do need to set this port, you must not select the same port that is used for listening.

Note: Take care in assigning this port to avoid conflicts with any other port that may be used by any other service. In particular, you should not assign a port number in the ranges Windows will use for ephemeral source port assignment.



These are the broadcast and multicast addresses used by Domain Time Client.

CAUTION: The default settings on this page are correct for most situations. Only make changes if you are sure you need them and you fully understand the effects of the change.

Broadcasts and Multicasts are used for a variety of purposes by the various Domain Time components. For example:

- Server uses them to send cascades and advisories
- Server uses them as the listen addresses for IPv4 and IPv6 multicast requests
- Server uses them to send broadcast/multicast time (using the DT2, NTP and PTP protocols)
- Tools that don't have their own settings (for example, dtcheck.exe) use them for discovery and testing
- Clients can use them to discover DT2, NTP, and PTP time sources
- Clients use them as the listen addresses for IPv4 and IPv6 multicast requests

The settings on this page control whether Domain Time will enable multicasts and broadcasts and what addresses to use for these functions. Note, Domain Time joins multicast groups by default even if these settings are disabled. Also, multicasts will always be used if you enable PTP. Multicast group join settings are set on the [Network](#) property page.

IPv4 Broadcasts

Enabled Broadcast address:

This section configures the Domain Time IPv4 Broadcast address.

By default, Domain Time Client will use broadcasts only on subnets local to this machine. If you want to reach machines on remote subnets, you should enable multicast (see below).

IPv4 Multicasts

Enabled TTL:

DT2 multicast address:

NTP multicast address:

This section configures the Domain Time IPv4 Multicast address.

The **TTL** (Time-to-live) entry sets how many router hops a multicast should make when propagating through your network. Choose a value that allows your multicasts to reach all of your subnets. This is usually the only setting you should change in this section.

Changing the multicast address listed here is usually an error.

IPv6 Multicasts

Enabled Hop Count:

DT2 multicast address:

NTP multicast address:

This section configures the Domain Time IPv6 Multicast address.

The **Hop Count** entry sets how many router hops a multicast should make when propagating through your network. Choose a value that allows your multicasts to reach all of your subnets. This is usually the only setting you should change in this section.

Changing the multicast address listed here is usually an error.



Settings on this page control the Domain Time Security settings.

Denial of Service (Flooding) Protection

DoS Protection Enabled

If any one machine sends more than requests in a -second period, ban
for: seconds

Auto-extend ban if abuse continues while IP is banned

Domain Time II has automatic protection against Denial-of-Service (DoS) disruption caused by intentional or accidental flooding of the network.

Any system that exceeds the DoS traffic thresholds you specify here has its access automatically blocked for a period of time.

Use the **Auto-extend ban if abuse continues while IP is banned** option if you have persistent bad actors whose bans expire, only to be re-blocked. You can also block them by IP address (see below).

Note: Even legitimate traffic can be blocked if it occurs too frequently. Take care that any tools that send repeated inquiries/commands to this machine do not exceed your DoS threshold.

Access Permissions

No restrictions IP ranges

Permit only listed range(s)

Deny any in listed range(s)

First IP in range

Last IP in range

Allow Domain Time II Manager to change the time zone on this machine

Auto-Manage Windows Firewall

Your Client's performance can potentially be degraded by responding to audit inquiries, sync triggers, or other traffic from machines on other network subnets over which you have little control.

To prevent this kind of problem, you may specify whether Domain Time should accept or reject time protocol traffic from certain IP addresses. You can specify whether to **Permit** or **Deny** traffic from multiple ranges of addresses. This allows you to easily restrict your incoming traffic to only the intended machines.

If you wish to permit or deny a single IP address, enter it as both the First and Last IP address in the range.


Allow Domain Time II Manager to change the time zone on this machine

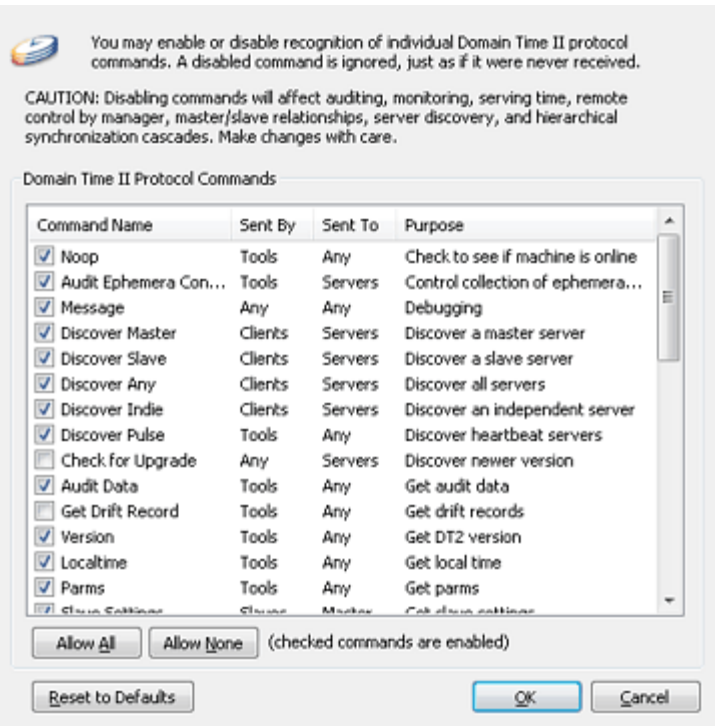
When checked, you may change the timezone on this machine remotely from Manager.

Auto-Manage Windows Firewall

As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

Command Restrictions

When you click on the  button you'll be presented with the *Command Restrictions* dialog window. You can use these settings to restrict what kind of Domain Time II control and sync messages your server listens for on the network.



Command Restrictions Dialog [\[Click for larger size\]](#)

The default protocol restriction settings assure both maximum functionality and a high degree of security; in most cases you will have no need to adjust them from the defaults. Domain Time II components communicate with each other primarily through directed communication, and are therefore highly resistant to spoofing and other malign interference.


The Domain Time II protocol command restriction capability is intended for use by system administrators in environments where an extra level of security is required, such as running a Server on the open Internet. Using the restrictions list, you can determine exactly what Domain Time II protocol command messages the service is allowed to listen for. Think of the command restriction list as an application-level "firewall" allowing in only the desired Domain Time II commands and blocking any others. Keep in mind that the restriction list only affects incoming DTII protocol commands - outgoing commands are not affected.

Warning:

Disabling protocol commands can have unintended consequences on the operation of your entire time distribution network, including the prevention of cascade triggers and sync notifications, which may result in inaccurate clocks. Problems resulting from disabled protocol messages can be quite hard to troubleshoot later, particularly by the next system administrator after you. Make adjustments only if you understand and require them, and be sure you document the changes so you can maintain the consistency and smooth operation of your time network.



This page configures the symmetric keys used by authenticating network protocols such as NTP and DT2.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

As of version 5.1, Domain Time supports two methods of network packet authentication: Windows Authentication and Symmetric Key Authentication.

Windows Authentication

Windows Authentication refers to the proprietary authentication method Microsoft uses to validate time packets between domain member machines and domain controllers (DCs). As of v5.1, Domain Time fully supports integrated Windows Authentication for both serving and obtaining the time within a domain.

Important considerations (read the sidebar for more details):

- Windows Authentication only works within a domain. That is, all machines exchanging time using Windows Authentication must be joined to the same domain (or forest).
- Windows Authentication only works between domain members and domain controllers.

While the domain member getting the time may be any kind of machine, the machine providing authenticated time must be a DC. Only a DC can validate the request. Other machines will not know the necessary shared secrets.
- Windows Authentication is automatic, requiring no additional configuration on servers or clients.
- Domain Time's Windows Authentication works with NTP, DT2-UDP, and DT2-TCP protocols between Domain Time Servers and Clients. W32time only authenticates using NTP.
- Windows Time (W32time) clients using "NT5DS" mode (the default domain member setting) can get authenticated NTP time natively from Domain Time Server running on a DC.

NT5DS-mode W32time clients may also get authenticated time from DCs running Domain Time Client, however, W32time must be set to NoSync mode on the DC so that it provides the authenticate NTP timestamp.

Although Domain Time Client will keep the DC's local clock highly accurate, using the W32time service to provide the authenticated time will result in reduced accuracy to the clients.

- Domain Time Clients can obtain authenticated time from Windows DCs running only the W32time service. However, this is not recommended, since

Interaction with Windows Time (W32time)

Windows Time clients using NT5DS mode (the default) search the Active Directory hierarchy to find a server. They send a request for the time using their machine RID as the authentication key, and expect the returned timestamp to be authenticated by the server. Only a DC in the client machine's domain can provide this type of authentication.

Domain Time v4.x Servers provided for Windows Time clients by setting the W32time service's client portion to "NoSync" mode and allowing the W32time service's server portion to serve NTP directly. Although the timekeeping ability of W32time is poor, this approach allowed the DC running Domain Time to continue serving Windows Time clients. This workaround is no longer necessary.

As of v5.x, Domain Time provides integrated Windows authentication natively for both NTP and DT2 protocols. This means that W32time clients in NT5DS-mode can get their time directly from any Domain Time II Server running on a DC exactly as if getting the time from the Windows Time Service on that DC.

- the Windows Time NTP service does a poor job of keeping the DC's local clock accurate, and
- the W32time NTP server itself adds additional inaccuracy to the network time being served.

Recommended settings:

Using Domain Time on your DCs is *highly* recommended.

For v5.x Server on a DC:

- Verify that the *NTP Server Enabled* checkbox is checked on the Domain Time II Server [Serve the Time](#) property page AND
- the *Windows Time mode*: dropdown on the Server's [Advanced](#) property page is set to **Disabled**.

For v5.x Client on a DC:

- the *Windows Time mode*: dropdown on the [Advanced](#) property page is set to **NoSync**.

Other considerations

■ Cluster Service

The Windows Cluster has a default startup dependency on W32time. It does not require the time service for any other purpose. Thus, the simple recommendation for installing Domain Time on clusters is to set the *Windows Time mode*: dropdown on the [Advanced](#) property page to **NoSync**, which allows the service to be running to satisfy the startup dependency, but allows Domain Time to set the cluster's clock.

However, you may replace the cluster's startup dependency if you want. After installing Domain Time Client (or Server) on the cluster, use RegEdit to navigate to the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clussvc

Change the **DependOnService** value (omitting the quotation marks) from "W32time" to "Domain Time Client" (or "Domain Time Server" if that's what's installed).

The cluster service will then wait until Domain Time has started before starting the cluster. You can then set the *Windows Time mode*: dropdown on the [Advanced](#) property page to **Disabled**.

■ Reliable Time Provider

DcDiag and other tools sometimes expect the Windows Time service to be running on DCs, even if it's not actually doing anything. These tools often depend upon the DC being flagged as a reliable time provider.

Starting with v5.x, Domain Time Server, when installed on a DC, sets the system flags to indicate the machine is serving time and is a reliable time source. The **DsGetDcName()** function will report Domain Time Server v5.x machines on DCs as both time servers and reliable time sources. Domain Time Server on a non-DC will not change the existing system flags.

You may override this behavior by editing the registry. In

Additionally, Domain Time v5.x clients can obtain authenticated time from DCs running either the Windows Time service or Domain Time.

W32time in NT5DS-mode has distinct disadvantages:

- The W32time NTP Server is inaccurate, so even if the DC's clock itself is well-synchronized, the time being served may not be. W32time clients receiving the time add to the problem, since they are unpredictable as well.
- Other ntp clients (such as xntp) cannot synchronize with it.

Domain Time's NTP Server has none of these disadvantages. It can provide standard NTP (with or without NTP auth) at the same time it provides NT5DS-mode timestamps, and all at extreme accuracy.

It is therefore highly recommended you install Domain Time II v5.x Server on all DCs. You can install Domain Time v5.x Clients on a DC, but you will then need to enable W32time in "NoSync" mode to provide NT5DS-mode time if you have clients that need it.

HKEY_LOCAL_MACHINE\Software\Greyware\Domain Time Server\Parameters

edit (or create) a **REG_SZ (string)** value called "Set Reliable Time Provider" and set its value to either "True" or "False" (the English words, without the quotation marks). If this value is present and set to True, Domain Time Server will set the two flags even if it is not running on a DC. This configuration has no meaning for Active Directory, since only DCs are examined for the flags. Other tools, however, may benefit from knowing that a reliable time source is present. If this value is present and set to False, then Domain Time Server will not change the flags.

Symmetric Key Authentication

v5.x Domain Time Clients and Servers support Symmetric Key Authentication (hash of shared secrets). Domain Time supports MD5, SHA1 (as of version 5.2.b.20170922), SHA256 (as of version 5.2.b.20190331), and SHA512 (as of version 5.2.b.20190701). Older versions support MD5-only.

Domain Time Server and Client support symmetric authentication (using SHA1 and/or MD5) of client-server requests using NTP (version 3 and later; AutoKey is not supported), DT2-UDP, DT2-TCP, and DT2-HTTP protocols. Domain Time also supports broadcasting (both NTP and DT2-UDP) with a shared key and hash. SHA256 hashes are only valid for use with PTP v2.1. Clients configured with the same key validate packets from the sending server by comparing the computed hash.

Note: Although v5.2.b.20190701 added support for SHA512, this option is reserved for future use. You should not create SHA512 keys. If an SHA512 key exists in the keyring of an older version of Domain Time, it will be (mis)interpreted as a very long MD5 key.

SHA1 keys are always exactly forty hex characters long. MD5 keys are ASCII text; different implementations of the NTP daemon have allowed different maximum key lengths. In general, an MD5 key should be composed only from 7-bit ASCII-printable text, excluding space, tab, and the # character. MD5 keys should be at least 8 characters long, and should not exceed 20 characters. Some versions of NTP daemons allow lengths of 32, while others have a maximum of 8 or 16. You will need to choose MD5 keys that are interoperable with all of your various devices and daemons. SHA256 keys must be exactly 64 hex characters long. SHA512 keys are a 128-byte hex string, corresponding to a 64-byte key.

The Keyring

Symmetric Keys are kept in a list containing the Key number and the Key secret (password). This list is also known as a *keyring*.

The keyring may contain a combination of *trusted* and *untrusted* keys.

A trusted key means the key is available to be selected by the component, but the trusted key is not active until its key number is selected when configuring a unicast time source in the [time sources list](#) (or by using the **Broadcast/multicast key** section of this page for broadcasts/multicasts). Untrusted keys are ignored.

Checked items are "trusted" or active keys; only trusted keys will be used:

Key	Type	Password
<input checked="" type="checkbox"/> 1	MD5	DomainTimeII
<input checked="" type="checkbox"/> 2	MD5	TTnts200
<input checked="" type="checkbox"/> 3	SHA1	97d870fe734e05bd449d476b9fbeb3b332234003
<input checked="" type="checkbox"/> 9909	MD5	greyware

There are various ways to configure the keyring on Domain Time II components:

- Master and Independent Servers
 - manually using the Control Panel applet (on this page)
 - importing the keys from a properly-formatted keyring (ntp.keys) file

- importing a previously-exported Domain Time .reg file.
- Shared secrets are automatically replicated between Domain Time Masters and Slaves. No configuration of the Slaves is necessary.
- Clients
 - manually using the Control Panel applet
 - importing the keys from a properly-formatted keyring (ntp.keys) file
 - importing a previously-exported Domain Time .reg file.
 - using Windows Group Policies (See the [Active Directory](#) page.)

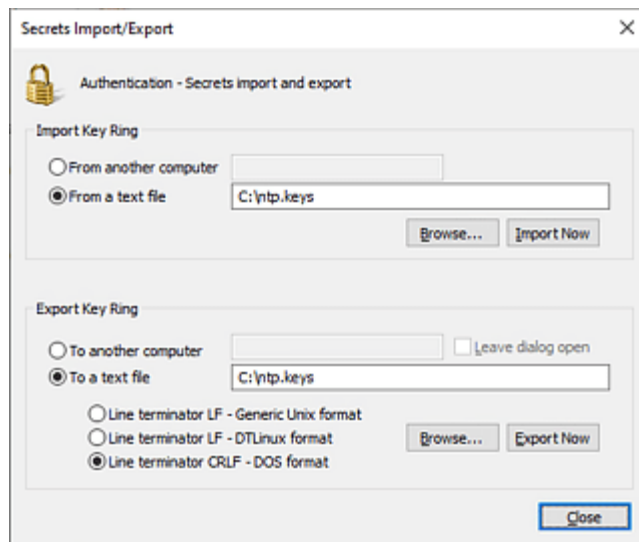
PTP v2.1 Security Parameter Pointer (SPP):

In addition to shared SHA256 hashes, PTP v2.1 requires that Masters and Slaves use the same SPP value to be able to validate v2.1 Authentication TLVs. The SPP stored in the keyring may either be zero (which acts like a wildcard) or must match what the grandmaster sends. If there is a potential for your Slaves to discover more than one Master (such as with a fallback server), we recommend you use the wildcard setting (0) to avoid synchronization failure if each server has a different SPP.

Import/Export the Keyring (keys file)

Click the **Import/Export** link in the *Symmetric Keys* section, which brings up a dialog where you can import or export a keyring (*.keys) text file to share among your devices.

This function is very useful if you are sharing a keyring file with other systems running a daemon such as dtlinux (dtlinux.keys), chronyd (i.e chrony.keys), or ntpd (i.e. ntp.keys).



Secrets Import/Export Dialog [\[Click for larger size\]](#)

Hints:

If you have multiple time sources, each may have its own set of symmetric keys. Be sure to import all the keys from all time sources into Domain Time.

If you are signing outgoing PTP v2.1 delay requests, all of your grandmasters should be configured to accept the same KeyId for incoming delay requests.

When possible, be sure all of your time systems are working correctly before enabling authentication. Authentication requires a correct setup on both ends of the connection, and changes at either end can cause a previously-working connection to fail. Disabling authentication temporarily should always be one of the first steps when troubleshooting a connection issue.

As of v5.2.b.20190701, you may use Manager to push out the symmetric keyring to multiple machines at once. See the [Reset Keyring](#) command.

Broadcast/Multicast Key


Broadcast/Multicast Key: 1

This dropdown selects the trusted key to be used when signing Broadcast or Multicast time packets. Note this refers specifically to the "heartbeat" type of time packets sent to the network on a fixed schedule, as configured on the [Serve the Time](#) page.

As with normal Symmetric Key authentication, Clients receiving the broadcast/multicast must also be using the same authentication key to decode the packet.



The Logs and Status page contains the settings for the Domain Time service text log.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Logs are kept in the `%SystemRoot%\System32\` folder. There are at least four main log files collected when the service is running:


- **domtimec.log**

This is the currently active text log file.

If log archiving is enabled (see below), additional archived log files will be created using a `domt i mec. YYYYMMDD. l og` naming scheme (i.e. domtimec.20090928.log).

- **domtimec.log.startup**

A detailed text log of the service startup process. Only data from the latest startup is included.

To view these two text logs, click the  button on the applet, which launches the Domain Time Log Viewer.

- **drift.dt**

A binary file containing information on each time check/correction made using the NTP and DT2 protocols, or the aggregate of corrections made by PTP (if enabled) during the check interval configured on the [Timings](#) page. Drift logs can also be collected remotely by Domain Time [Audit Server](#).

To view this log, click the  button on the applet, which launches the Domain Time Drift Log Viewer.

- **driftptp.dt**

If PTP is enabled, a binary file containing information on each PTP Sync. PTP Drift logs can also be collected remotely by Domain Time [Audit Server](#).

To view this log, click the [Graph](#) link on the [Obtain the Time](#) property page, which launches the Domain Time Drift Log Viewer.

Text Log

Log Level: Information

Max Size: KB (use zero to mean unlimited size)

Enable lazy write delay of up to seconds (range 1-600)

Include info-level timeset success messages in warning and error-level logs

Include client accesses (requests for time and control messages)

Enable UDP packet tracing Enable TCP packet tracing

Enable Time Change Event Monitor

Enable NTP4 peerstats Enable NTP4 loopstats

NTP Stats Folder:

This section selects the properties of the **domtimec.log** service text log.

The **Log Level** drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

This switch will only disable the **domtimec.log** file. Other system logs, such as **domtimec.startup.log** and **drift.dt** cannot be disabled.

- **Errors**

Only messages marked as Errors will be logged

- **Warnings**

Logs will include Errors and Warnings

- **Information**

Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.

- **Trace**

Includes all of the above, plus detailed information on time setting and time sample analysis.

- **Debug**

Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

When you select **Debug** level, the button becomes enabled. Clicking this brings up a dialog where you can select exactly which type of debug messages to include in the logs. This allows you to limit the size of the logs while troubleshooting a particular issue. Please re-enable all messages if submitting a log for analysis.

Max size: sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to troubleshoot any issues that may arise.

Enable lazy write delay of up to seconds (range 1-600)

This value specifies the maximum amount of time to wait before flushing data to the log.

Logging large amounts of information on an underpowered or busy machine can generate a great deal of overhead, which can adversely affect system performance and diminish timing accuracy. When enabled, Domain Time will try to buffer log data in memory until the delay period is reached instead of attempting to write all events to the disk in real time. This may provide some "breathing space" for the disk system to process outstanding writes that may accumulate from constant log activity. If the buffer fills before the limit is reached, it will be flushed to disk even if the full wait period has not expired.

Include client accesses (requests for time and control messages)

When checked, all client requests made of this machine will be logged. This includes queries for time or other information such as statistics, auditing, or status requests.

CAUTION:

Enabling this option can generate a large amount of logging data and system overhead if you have a lot of

clients synchronizing with this server. Examine the log after the server has been running for a while to see if this option generates an onerous amount of data.

Enable UDP packet tracing

Enable TCP packet tracing

These checkboxes cause Domain Time to log additional useful details about the time packets being used by Domain Time. The output is similar to the output from a packet analyzer, showing you the actual packet contents and payloads.

CAUTION:

Enabling this option can generate a large amount of logging data and system overhead. You should enable these only when actively troubleshooting network issues and for short periods.

Enable Time Change Event Monitor

Tells Domain Time to use Windows auditing to help identify the user or process responsible for changing the system clock.

When checked, Domain Time will attempt to enable the Windows audit category "System" success auditing, and then watch the security event log for events pertaining to date/time changes made by programs other than Domain Time. If such an event is detected, Domain Time will parse the security event log entry and issue a warning in its own log. The warning will show the user and process that changed the time, and by how much (if the information is available). These warning messages are informational only, and should be enabled only to help track down environments where another user or process is interfering with the system clock.

If the audit policy for the machine is controlled by a group policy, then Domain Time's change to the audit policy will succeed, but only until the next group policy refresh is applied. If you are using this feature in a domain, either undefine the group policy setting (Local Policies/Audit Policy/Audit system events) or set it to enabled for success events.

Note: This option is complementary to, but independent of the Clock Change Monitor function, which resets the system time if unauthorized changes to the system clock are detected. See the description of Clock Change Monitor on the [Advanced](#) page.

CAUTION:

This option causes additional system overhead and uses additional memory and resources. You should enable this option only if you are experiencing rogue time changes on your system and are having difficulty identifying the cause. You should not run with this option enabled in normal operation.

Enable NTP4 peerstats

Enable NTP4 loopstats

As of version 5.2.b.20170101, these checkboxes enable creation of ntpd-style peerstats and loopstats statistic files. See the [ntpd Compatibility](#) page for details.

When enabled, the path where the files are collected will be displayed in the *NTP Stats Folder:* field. Note: The NTP Listener must be enabled on the Client's [Advanced](#) property page in order to collect these stats.

[Text Log Archiving](#)

Log Roll: Daily at Midnight

Delete old logs

Keep up to old logs

Domain Time can automatically archive the text log on a daily, weekly, or monthly schedule.


When the log is archived, all existing log events in the **domtimec.log** file will be written to an archive file named

`domtimec. YYYYMMDD. log` (i.e. domtimec.20090928.log) and the current log file will then be cleared to accept new data.

You can choose how many archived log files to keep on the machine. When the indicated limit is reached, the oldest log file will be deleted.



This page specifies whether Domain Time service activity will be echoed to the Windows Event logs.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Some levels of logging can create a significant amount of data. The Windows Event logs can be difficult to read, or the Event Log process may even have problems recording all the data when large amounts of log activity are generated.

You should consider using only the **Error** level when using the Event Logs unless you generate a very small amount of logging data overall. In general, Text or Syslog logging is a better choice for keeping more detail.

Event Viewer

Log Level: Information

The **Log Level** drop-down chooses what type of entries to include in the Event logs. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**
Domain Time will not log events to the Windows Event Logs.
- **Errors**
Only messages marked as Errors will be logged
- **Warnings**
Logs will include Errors and Warnings
- **Information**
Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.
- **Trace**
Includes all of the above, plus detailed information on time setting and time sample analysis.
- **Debug**
Includes all available information provided by the service.

Warning:

The amount of data generated by Debug logging can easily overwhelm the Event Log system. Use Text or Syslog logs for debugging instead.

Event IDs


Event ID	Category	Meaning of Event ID
1000	Success	Generic success (examine text for details)
1001	Warning	Generic warning (examine text for details)
1002	Error	Generic error (examine text for details)
2001	Success	Time set successfully
3000	Warning	Unable to set the time from any source
3001	Warning	Protocol error while trying to obtain the time
3009	Warning	Clock-Change Monitor is disabled
4009	Error	Clock-Change Monitor trigger detected

The listed Event ID codes can be used to filter for Domain Time events in the Event Viewer. The Event Source field on the Event records will be **Domain Time Server**.

If you're considering using the Event Viewer for live system monitoring purposes, you may want to investigate the [SNMP Traps](#) function or [Service Status Monitor](#) to be more efficient.



You may choose to have Domain Time send service activity events to Syslog servers.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Syslog

Log Level: Information

Server(s):

You may list up to eight IPv4/IPv6 target addresses. Separate with spaces

☐ RFC 3264 format

☐ RFC 5424 format (UDP only)

☐ Include timeQuality structured data element

☐ Send each PTP data point to syslog (trace or debug level only)

The **Server:** field should contain the DNS Name or IP address of the Syslog Server(s).

As of v5.2.20170922, you may list up to eight targets on the Syslog Server line. Older versions only support a single target. Separate targets with a space.

Important: If you will be assigning this value using [Templates](#) or [Active Directory Policies](#) to any Server or Client version older than v5.2.b.20170922, the first entry of a multiple target list **MUST** be an IPv4 address, otherwise the older version will read the field incorrectly. For maximum backward compatibility with older versions of Domain Time, avoid using a DNS name if you list more than one target. If all of your machines have been upgraded, then you may use either DNS names or IPv4 addresses for your targets.

The **Log Level** drop-down chooses what type of entries to include in the Syslog logs. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**
Domain Time will not log events to the Syslog server.
- **Errors**
Only messages marked as Errors will be logged
- **Warnings**
Logs will include Errors and Warnings
- **Information**
Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.

- **Trace**

Includes all of the above, plus detailed information on time setting and time sample analysis.

- **Debug**

Includes all available information provided by the service.

RFC 3264 format

RFC 5424 format (UDP only)

Include timeQuality structured data element

Added as of version 5.2.b.20171113. Use these radio buttons to select the RFC format that matches your syslog server. If the **Include timeQuality structured data element** checkbox is checked, the output will include timeQuality information, for example: `[timeQuality tzKnown="1" is Synced="1"]`

Send each PTP data point to syslog (trace or debug level only):


Added as of version 5.2.b.20180101. If checked, and if the syslog level is set to either trace or debug, then each PTP data point will be sent to syslog.

The format is `PTP sample offset ±0.0000000, mpd 0.0000000, source ipaddress` where 0.0000000 is that sample's delta and the current meanPathDelay. Syslog log collectors may parse for trace-level messages beginning with "PTP sample offset" to categorize these messages.

Caution: Enabling this output can create a large number of syslog messages. Enable only if you actually require this level of detail. You may want to enable lazy writes if you find the logging process is affecting your clock accuracy. See the [Logs and Status](#) page to enable lazy writes.



Domain Time II Client can send notifications of its status to Network Management Systems and other SNMP-capable monitoring devices using SNMP Traps.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Important: SNMP support depends upon wsnmp32.dll being present in the OS. All Windows versions except the initial release of Windows Server 2016 Nano Server have this installed by default. To install Domain Time on Nano Server, you must install the SNMP trap support .dll before installing Domain Time. See the [Nano Server FAQ](#) for more information.

SNMP Control

Enabled

Community:

Server:

Use this section to enable SNMP Traps. Enter the SNMP community name and password used by your Network Management System (NMS), as well as its DNS name or IP address. Your community name and password must match the one in use by the receiving system.

As of v5.2.b.20190331, you may also specify a port number to which SNMP traps are sent in the Server field. Examples: **snmp.mydomain.com: 1214** or **[2002: 410: 1: 1: 2a0: 69ff: fe01: b0f4] : 2444**. If the port number is not specified, the default SNMP trap port of 162 will be used.

Best Practices for SNMP include using a unique community name and hard-to-guess password on production systems. The default community *public* should only be used for initial testing. Although Domain Time only sends outgoing trap information and is therefore not susceptible to SNMP remote control vulnerabilities, you should still be mindful of SNMP security for the benefit of your other SNMP devices.

SNMP Traps

Sync Successful	If variance exceeds	milliseconds
Sync Failed	Ignore variance alert for first sync	
Service Startups		
Service Shutdowns		

The settings in this section select which SNMP traps are sent by Domain Time.

SNMP v2 traps are generated whenever the selected event occurs. Keep in mind that SNMP Traps are sent via UDP, and are therefore not guaranteed to be delivered by the network.

Although useful for raising performance alarms or other monitoring functions, you should not depend upon the SNMP trap data for critical logging of time synchronization events, particularly if your logging is necessary for regulatory compliance. Use a product designed for more robust data collection, such as Domain Time [Audit Server](#) instead.

The **If variance exceeds** **milliseconds** setting lets you set a threshold value so that you can be alerted when the variance of any timecheck exceeds this value.

Warning:

Domain Time may generate a large number of timechecks (depending on your [Timings](#) settings) so it is very easy to swamp your monitoring system with alerts if you set this value too low. Also, it is normal even on the best-behaved networks for occasional timechecks to reflect a spurious large value due to latency or other network conditions.

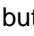
It is therefore very easy to generate a large number of false alarms using this trap.

If you enable this trap, you should choose a threshold value that truly reflects a critical amount of clock drift to avoid unnecessary alarms for normal transient variances. If your monitoring software has the capability of further restricting alarms only after a certain number or percentage of traps have been raised, you can add extra protection against false alarms.

You may find the [Audit Server Notifications](#) feature to be a better way of immediately alerting you to poorly performing clocks than this trap.

Since the first timecheck on any system after the time service starts is likely to be quite large (due to the clock not having been set yet), the **Ignore variance alert for first sync** setting is highly recommended.


The Domain Time MIB File

Domain Time comes with a MIB file that you can use to compile on your SNMP monitoring system so that your traps are interpreted correctly. The MIB text file is generated when you click the  button on the Control Panel applet so you don't need to worry about locating it in some obscure installation folder or having online access.

Note: The MIB file generated here matches the version of Domain Time that's currently installed. Be sure to remember to update your SNMP Network Management Station(s) with the latest version of the MIB after any upgrades to Domain Time.



This page contains options for additional features to monitor the Domain Time service in real-time.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

Audit Server Real-Time Alerts

Enabled

Primary Server:

Backup Server:

send to the backup only if primary is down

send to both primary and backup servers

Send reports using TCP (recommended)

Send reports using UDP

Always audit this machine

Never audit this machine

Do not change audited status

Domain Time Server and Client are designed to be centrally monitored and have their synchronization status recorded by Domain Time [Audit Server](#). This is usually done by Audit Server in a scheduled, background data collection process.

However, Domain Time II Server and Client can send an immediate notification packet to Audit Server containing feedback on the success or failure of each time check to provide real-time monitoring and alerting. Domain Time Manager on the Audit Server machine will display the collected data on its Real-Time Alerts panel. See the [Audit Server Alerts](#) pages for more information on setting up Real-Time Alerts.

If you want this function, check the **Enabled** checkbox and enter the DNS name or IP address of your Primary Audit Server.

As of version 5.2, you may specify addresses for both a Primary and Secondary server for redundancy purposes. Use the radio buttons to select whether to

send to the backup only if primary is down, or to

send to both primary and backup servers.

This function is designed to work well with the Audit Server "Hot Spare" standby mode functionality introduced in version 5.2.

Select whether to use TCP or UDP for the notifications. In general, TCP is more reliable than UDP since delivery of TCP traffic is given priority, whereas UDP can be delayed or dropped entirely by busy network hardware. However, TCP does require more resources on the network stack of the Audit Server to handle the mechanics of building and tearing down TCP connections. If you have many machines sending lots of real-time updates using TCP, it is possible to exceed your operating system's ability to handle the number of open network connections. If you run into those limitations, you may want to consider changing your alert notices to UDP.

CAUTION:

If you are using notifications on many machines that are set to synchronize frequently, the feature can generate a significant amount of network traffic toward your Audit Server. You will probably want to enable instant notifications only on critical systems where an immediate alert of time sync errors is desired.

Otherwise, you should use Audit Server's normal scheduled sync log collection instead, which is much more efficient. Audit Server's sync log collection can run in a background process to gather all sync records from audited systems with much less impact on overhead than active notifications. Audit Server can be programmed to provide alerts based on that data as well.

Audit controls

As of version 5.2.b.20110224, you have the option to have the Domain Time service control its own inclusion in or exclusion from the Audit Server Audit list.


Always audit this machine

Never audit this machine

Do not change audited status

These options override any existing settings on the Audit Server itself. You may use these functions to ensure that machines are (or are not) audited whether or not they are discovered by Audit Server.

Domain Time Service Status Monitor



The *Service Status Monitor* section controls whether Domain Time will provide a simple text response about its current status when asked by an application.

This monitor is provided to allow third-party applications a simple way to monitor the activity of the Domain Time service. The Status Monitor will respond to TCP or UDP requests on the specified port with a simple text string showing the current activity of the service.


Sample responses from the Status Monitor:

ACK Adjusting (Indp. Server 5. 2. b. 20130403R)

ACK Set 22 seconds ago (Indp. Server 5. 2. b. 20130403R)



This page gives you access to various advanced Domain Time II settings.

Note: If you see the  **Group Policy applied** indicator in the lower-left corner of the applet, there are settings on this page that are being overridden by an Active Directory Group Policy. Settings controlled by policy may be greyed-out or you may be otherwise prevented from making a change here. See the [Active Directory](#) page for more information on using Group Policies.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings may adversely affect your clock accuracy or even prevent clock corrections entirely.

Miscellaneous Options

- Enable Test Mode (if checked, time on this machine will NOT be corrected)
- Enable Clock-Change Monitor
- Force setting of CMOS clock at service shutdown
- Use software timestamping (if available and enabled)
- Enable advance scheduling of leap second corrections
- Signal resync if VM guest resumes from paused or saved state
- Truncate drift status records to milliseconds

Miscellaneous options include:

Enable Test Mode

Checking this box causes Domain Time to operate in all ways as it would in normal operation, except for actually setting the time or changing the machine's slew rate. This allows you to test or troubleshoot the server's ability to obtain and serve the time, but without actually changing the server's time.

Results of all operations are logged normally, so you can use the log in test mode to track down any communication or other synchronization issues. Note: Be sure to disable this option when you're through testing!

Enable Clock-Change Monitor

Domain Time's *Clock-Change Monitor* notifies Domain Time if another user or process attempts to change the time on this system.

When an unauthorized clock change is detected, Domain Time immediately re-synchronizes the time with its time source(s) and makes a warning entry in the logs. This prevents inadvertent or malicious tampering with the system clock.

This setting should always be enabled unless you are doing testing that requires you to change the system clock manually, either from the Windows Date/Time applet or from some other application.

Force setting of CMOS clock at service shutdown

Controls whether Domain Time should perform an API system call to write the current system time to the CMOS Real-Time Clock (RTC) on the motherboard each time the service stops.

On modern operating systems, the CMOS RTC clock is primarily used to provide something approaching the current

date/time to the operating system while booting until the operating system can take over timekeeping. The CMOS clock is subject to all manner of inaccuracies, and is therefore not used for timekeeping while the OS is running, nor is it updated often.

The CMOS clock can therefore go for long periods without having its time corrected, resulting in huge drift. By default, Domain Time will update the CMOS with the current time either when doing so doesn't cause a disruption to the operating system time (during stepped corrections) or just before shutdown so that the CMOS has its best chance to be accurate during the time the system is not running.

When this box is **unchecked** (the default), Domain Time writes the current time to the CMOS RTC clock:

- when making a stepped time correction.
- if the Domain Time service is running and it receives a system shutdown command from the operating system.

If the box is **checked**, then Domain Time will also write to the CMOS clock any time the Domain Time service is stopped, whether or not the stoppage is due to a system shutdown.

Although at first read this may seem desirable, there is a downside to writing to the CMOS clock if the machine isn't already being hard-set (stepped) or in the process of shutting down. The API used by the operating system to write to the system clock also immediately steps the system time to the same time as the RTC *but only at the resolution of the CMOS clock*. Since the RTC resolution varies on different machine, writing to the CMOS will cause the system clock to jump either forward or backward to the nearest increment of the RTC, which can mean an unpredictable jump of 1ms or more in the system time.

As a result, you should leave this switch turned off unless you need to force a CMOS update by manually stopping the Domain Time service.

Use software timestamping (if available and enabled)

As of v5.2.b.20190701, Domain Time can use the Microsoft NDIS software timestamping API to help compensate for network stack delays. Microsoft only offers this function on its most recent versions of the operating system (i.e. Server 2019 and updated versions of Win10). This option will be greyed-out if it is not supported on this OS.

Note this checkbox only controls Domain Time's ability to use software timestamping if present. Timestamping must also be enabled at the operating system level. You can do this by using Microsoft's own PowerShell scripts (see KB article linked below), or you may use the [DTCheck](#) utility. To use DTCheck, open an elevated command prompt and issue the desired command:

<code>dtcheck -swTimestamps</code>	displays current settings
<code>dtcheck -swTimestamps: Enable</code>	attempts to enable software timestamping
<code>dtcheck -swTimestamps: Disable</code>	attempts to disable software timestamping
<code>dtcheck -stats2</code>	shows statistics for NDIS timestamping
<code>dtcheck -interfaces</code>	shows which interfaces have software timestamping enabled

You'll need to restart the NICs (or reboot the machine) after enabling or disabling timestamping.

You can experiment with this setting to see if it improves your accuracy or not. The stack delay on most machines is minimal so the overhead of measuring it may outweigh the benefits, however, you may see improvement on very busy machines.

See [KB2019.708](#) for more information.

Enable advance scheduling of leap second corrections

Controls how Domain Time handles upcoming UTC leap second corrections.

NTP and PTP packets can contain a flag to indicate an upcoming UTC leap second. When this checkbox is enabled, Domain Time will apply leap seconds at 23:59:59 UTC on the last day of the month in which the leap occurs (typically June or December). If unchecked, leap seconds will be applied at the first timecheck following the leap.

Domain Time acquires pending leap second information only from NTP or PTP time sources. All queried NTP or PTP sources must agree that a leap is pending in order for Domain Time to schedule the leap. If the sources disagree, then the leap will be handled at the next timecheck after it occurs, and a warning notice that the leap indicators are inconsistent will be placed in the log.

Pending leap information is queried with each timecheck (NTP sources only), and maintained only while the Domain Time service is running. Restarting the Domain Time service will clear any pending leap second corrections. If the leap is still pending when the Domain Time service is restarted, it will be rescheduled for the appropriate time. If the leap occurs while the Domain Time service is stopped, the leap will be applied at the first timecheck after startup.

[Read more](#) about leap seconds.

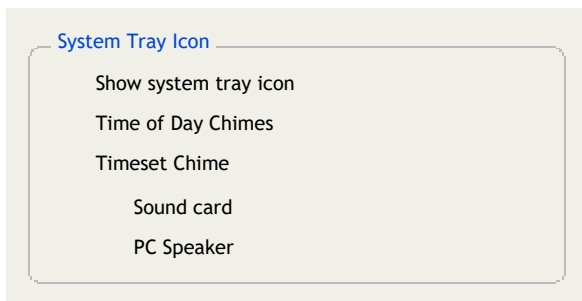
Signal resync if VM guest resumes from paused or saved state

Allows correction of clocks on virtual machines after resuming from pause/suspension.

Virtual machines are often paused/suspended, causing the clock to be incorrect when resumed. Domain Time may be able to sense a resumption by examining the Time Stamp Counter (TSC) and resync the clock. This setting is only useful on virtual machine guests. This option will be greyed-out if the machine is not a virtual guest.

Truncate drift status records to milliseconds

Delta values in the drift logs will be reported to the nearest millisecond if this checkbox is checked.



The Domain Time II System Tray applet (DTTRAY.EXE) is a foreground application that can load whenever a user logs into the system. When loaded, it will display as an icon in the System Tray.

The applet provides a number of very useful functions, including audio alerts and chiming, statistics, drift graphs, and a quick way to launch the various features of Domain Time installed on the machine.

The settings in this section determine whether or not to load the applet and which audio features are enabled.

Show system tray icon

This checkbox controls whether the [Domain Time System Tray](#) applet is loaded during login. If the icon is present in the System Tray, you can right-click it to choose from many additional features.

Note: The applet will unload if the Domain Time service is stopped. On XP and Server 2003, the applet will reload automatically when the service restarts. However, beginning with Windows Vista, Microsoft disabled the ability for background services to launch foreground programs, so on those systems you will need to either log out and back in or relaunch the applet manually. You can restart the applet manually by entering `dttray.exe` into the *Start -> Run* program field or at a command prompt.

Time of Day Chimes

The Time of Day Chimes feature plays sound files at specific times of the day, such as every 15, 30 45 minutes and on the hour.

This option will be unavailable if the **Show System Tray icon** checkbox on the Advanced tab is unchecked.

There must be a logged-in user and the Domain Time II System Tray icon must be present in the Windows System Tray for the chimes to play. You must also have installed at least one free Domain Time II [Chime Pack](#) for this feature to work.

See the documentation for the [System Tray Applet](#) for complete instructions and to download free Domain Time II Chime Packs.

Timeset Chime

Plays a sound whenever the Server successfully sets its time from its time source. If checked, the sound will play whether or not there is a logged-in user.

Sound card plays through the sound card if available

PC Speaker plays through the PC speaker

Client Options

Ignore incoming DT2 Cascade signals

Ignore incoming DT2 Advisory signals

The first two settings in this section control whether Domain Time II Client responds to cascade and/or advisory messages from other Domain Time components.

Cascade messages are used to propagate time corrections quickly down the hierarchy without having to wait for each component to synchronize on its own. This causes your clocks to converge on the correct time across your network in seconds instead of minutes/hours it takes using other non-signaled methods such as standard NTP or Windows Time. Cascade messages are considered mandatory and are always acted upon by the receiving component (unless explicitly disabled).

IMPORTANT: These signals greatly enhance the overall synchronization of your time network.
Disable only if necessary.

Advisory messages are used to help components determine the structure of the time hierarchy, such as by helping clients auto-locate available time servers. Components may act on or ignore advisory signals depending on their current configuration.

Cascade and Advisory signals may be unicast, broadcast, or multicast (any combination).

Each server has its own settings for whether or not it sends cascades, and if so, what type (see the Server's [Broadcasts and Multicasts](#) page for details). A server can send broadcast IPv4 only, multicast IPv6 only, multicast IPv4 only, or any combination. IPv4 broadcasts are sent to 255.255.255.255. This is not configurable. If you need your cascades to cross routers, you must use IPv4 or IPv6 multicast instead.

See the hierarchy description below to see which type of cascade/advisory is used.

The Domain Time Cascading Hierarchy

Domain Time II is designed to use a cascading time hierarchy to distribute the time. The Domain Time hierarchy is more robust than the inbound time partner structure of Windows Time and much simpler than manually configuring NTP peering and strata. In the hierarchy, each server is responsible for matching its time with the server above it, and providing the time to servers or clients below it.

■ Level 1 (Master)

Domain Time II Server installed on the domain controller with the Primary Domain Controller (PDC Emulator) role becomes the *Master* time server for the domain. When the Master's time is corrected to match its time source(s), the Master directs a Level 1 unicast cascade signal to each known Slave. These are the only unicast cascade signals, and they cannot be disabled.

The Master expects an acknowledgement to the Level 1 cascade from the Slave. If a Slave fails to acknowledge (perhaps because it is currently offline), this is noted in the Master's log file.

After signaling each known slave, the master broadcasts/multicasts a Level 1 cascade signal to the network. It will be an advisory if at least one slave was contacted successfully, else mandatory (the assumption being that there are no slaves to relay the signal to waiting clients). A master may be configured to skip sending this Level 1 signal to the network.

■ **Level 2 (Slave)**

Any other domain controller, member server, or member workstation can be configured as a *Slave* (this is the default for domain controllers). Slaves automatically discover and synchronize with the Master Server. Slaves synchronize time with the Master using the DT2-TCP and protocol, so any intervening firewalls, routers, and/or switches must pass port 9909 TCP (note, it is always a good idea to pass both 9909 TCP **and** UDP traffic).

When a Slave receives a Level 1 cascade signal from the Master, it immediately synchronizes its clock and acknowledges the signal.

After resynchronizing with the Master, each Slave will broadcast/multicast a Level 2 signal to the network. If the resync was due to a Master's Level 1 trigger, the packet will be advisory, else mandatory. The Slave uses the Master's Level 1 sequence number, so any client that happens to hear from multiple Slaves, or from Slave(s) and the Master itself will not resynchronize multiple times. Slaves may be configured to skip sending Level 2 signals to the network.

■ **Level 3 (Independent Server)**

Any machine running Domain Time II Server (except the domain controller with the PDC Emulator role, which must be a Master) can be configured as an *Independent Server*.

When an Independent Server corrects its clock, it broadcasts/multicasts a Level 3 advisory. Independent servers may be configured to skip sending Level 3 signals to the network.

An Independent Server does not actively participate in the cascading hierarchy with Masters and Slaves. Independent Servers acknowledge Level 1 cascade signals from the Master, but do not act upon them. Independent Servers ignore Level 2 cascade signals from Slaves.

■ **Level 4 (Client)**

A client both listens for cascade signals and sets its own time independently based on its timing settings.

If a client is in automatic mode, it uses discovery broadcasts at startup to determine if any Level 2 (Slave) machines exist on the local subnet. If so, the client enters normal operating mode, and will synchronize its clock upon receipt of a Level 2 cascade signal.

If no Level 2 machine is found (perhaps because all Slaves are currently offline, or the client is not connected to the network), the client enters pessimistic mode. In pessimistic mode, the client listens for all cascade and advisory signals and responds by synchronizing its time with whatever machine sent the cascade signal and taking the following actions:

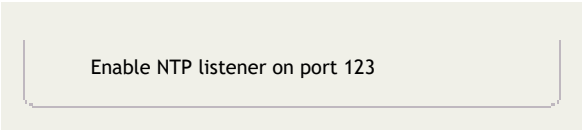
If the cascade signal comes from a Master Server, the client assumes that the network has only a Master and clients. From that point on, the client will ignore Level 3 cascade signals (Independent Servers). This is called Master-only mode. Master-only mode converts to normal mode upon receipt of the first cascade signal from a Slave.

- If the cascade signal comes from a Slave Server, the client assumes that Slaves are now present, and from that point on ignores both Level 1 (Master) and Level 3 (independent server) signals. This is normal mode.
- If the cascade signal comes from an Independent Server, the client will sync with the Independent Server, and assume the network has neither Master nor Slaves. The client continues in pessimistic mode until a Master or Slave signal is seen.

Note that this procedure allows the time hierarchy to automatically collapse levels so that clients respond only to the next-highest level at any time. If a Slave comes online after the client is started, the client will note this fact and move from Master-only mode to normal mode immediately. If, while in normal mode, no Slave can provide the time, the client will

automatically move into Master-only or pessimistic mode, as needed.

This hands-free configuration allows you to have any mix of Master, Slaves, and independent servers on your network, any of which may be working or not at any given time, and still use Domain Time II's hierarchy to (a) limit network traffic, and (b) ensure quick, uniform updating of all levels when the highest-level time source is updated.



Enable NTP listener on port 123

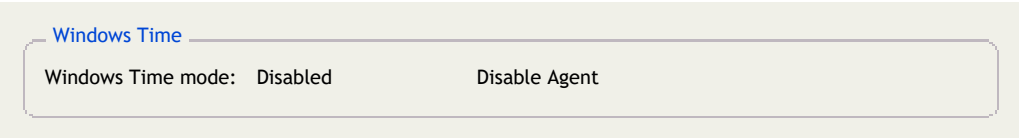
The **Enable NTP listener on port 123** checkbox controls whether Domain Time Client will listen for incoming NTP traffic (such as NTP broadcasts/multicast time packets or other requests from NTP status tools).

IMPORTANT: Only one program or service may own the NTP port on this machine. If this checkbox is checked, Domain Time will attempt to acquire ownership of the NTP UDP port 123 when the service starts. You ***must*** be sure any other time program that listens for incoming NTP requests (such as W32Time) is disabled before enabling this setting.

Note that this setting does not apply to the outgoing NTP time requests Domain Time makes of its time sources; it only applies to incoming NTP traffic.

When enabled, Client will be able to accept NTP broadcast/multicast time packets (assuming you have chosen that method on the [Obtain the Time](#) page). Client will also be able to respond to a subset of requests from the standard UNIX/Linux **ntp** query utility, i.e. **ntpq -np**.

Windows Time



Windows Time mode: Disabled Disable Agent

The settings in this section configure the Windows Time Service to co-exist with Domain Time.

Windows Time mode: Disabled

This drop-down box lets you determine how the Windows Time Service should behave on this machine. When the Domain Time II service starts, it will force the Windows Time service into this mode. The available options are:

- **Disabled**

The service startup setting for Windows Time Service is set to Disabled. The Windows Time Service will not be allowed to run. This is the preferred setting for all machines except domain controllers and machines running Windows Cluster Service (see the *NoSync* description below).

- **NoSync**

This mode makes sure the W32Time Server Provider portion of the Windows Time Service is running, but the W32Time Client Time Provider is disabled. In this mode, Domain Time II actually obtains the correct time and manages the local system clock; Windows Time merely answers NTP requests.

Note: This mode is necessary either when Domain Time II Client is installed on a Windows Domain Controller to enable NT5DS mode to function properly, or on versions of Cluster Server that have a startup dependency on W32Time (see below). This mode will be enabled automatically during installation of Client on a DC; you will need to manually enable it on Cluster Servers.

Although machines running in *NoSync* mode will provide NTP to NT5DS-mode machines, the accuracy

of the timestamps provided will be constrained by the native inaccuracy of the Windows Time service. Also, non-Windows systems may have difficulty synchronizing with the machine, since W32Time is not compatible with many NTP daemons. If possible, we recommend you use the [Domain Time II Server](#) on Domain Controllers instead, so that it can provide high-accuracy NTP to all clients.

Cluster Service

Some versions of the Windows Cluster Service (i.e. Win2003 and earlier) have a default startup dependency on the w32Time (Windows Time) service. Cluster Server does not appear to require the time service for any other purpose. Thus, the simplest recommendation for installing Domain Time on clusters that have the W32Time startup dependency is to set the *Windows Time mode*: dropdown to **NoSync**, which allows the W32Time service to be running to satisfy the dependency, but allows Domain Time to set the cluster's clock.

Although the *NoSync* setting is sufficient to allow the Cluster Service to start on machines running Domain Time Client, you may replace or remove the startup dependency if you want.

CAUTION: The following registry change is provided for your information. We're not aware of any issues with removing the dependency, but you should defer to Microsoft's guidance. Be sure to test any changes thoroughly on non-production servers before implementing on production systems.

To remove the W32Time startup dependency (if present):

After installing Domain Time on the cluster, use RegEdit to navigate to the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clussvc

The *DependOnService* value lists all services on which the Cluster Service startup depends. If the **w32Time** entry is present in the list, change it to **Domain Time Client** and save your changes. The cluster service will then wait until Domain Time has started before starting the cluster.

If the **w32Time** entry is not present in the list, there is no startup dependency on Windows Time in your version of Cluster Server and you do not need to make any changes to this registry value.

Once you have verified there is no startup dependency on W32Time on all nodes of the cluster, you can then set the Domain Time *Windows Time mode*: dropdown list to **Disabled** and restart the Domain Time service.

■ Not Touched

The existing configuration of the Windows Time Service is not changed. In this mode, Windows Time will operate however it is currently configured. With this setting, Domain Time will not set the time or manage the clock. The Client will only respond to audit and/or monitoring requests.

This option is not recommended.

■ NT5DS

The Windows Service is set to run and it obtains the time from the Active Directory hierarchy in NT5DS sync mode. With this setting, Domain Time will not set the time or manage the clock. The Client will only respond to audit and/or monitoring requests.

This option is not recommended.

■ AllSync

The Windows Service is set to run and it attempts to obtain the time from the Active Directory hierarchy in NT5DS sync mode and/or using NTP Client mode. With this setting, Domain Time will not set the time or manage the clock. The Client will only respond to audit and/or monitoring requests.

This option is not recommended.

■ NTP

The Windows Service is set to run and it attempts to obtain the time using Windows Time's NTP Client mode. With this setting, Domain Time will not set the time or manage the clock. The Client will only respond to audit and/or monitoring requests.

This option is not recommended.

Disable Agent

This checkbox disables the Domain Time II [Windows Time Agent](#).

Note: In version 4.1, the Domain Time II Windows Time Agent was installed by default. In version 5.1 and newer, Domain Time is able to replace Windows Time entirely, so the agent is not installed and the option defaults to disabled. If you have upgraded from 4.1, the agent is still present but not required. You may use this option to disable it.

This option has no effect if Agent is not installed.

If you would like to use the Agent, you may install it from the distribution files, or by using Manager, or [download the software](#) from the website, if desired. You must close and re-open the Server Control Panel applet after installing the Agent.

If Agent is installed, the [Agent button](#) launches the Domain Time II Windows Time Agent to allow you to view and configure the settings for the Windows Time service. Depending on the settings above, various parts of the Windows Time Agent applet may be disabled. See the full [Windows Time Agent documentation](#) for more details.



Clock Control

Domain Time gives you extensive control over how corrections are applied to the system clock.

These advanced settings are provided to address special clock-correction requirements, poorly-behaving system clocks, and for fine-tuning for extreme clock accuracy. In most cases, you will not need to make changes here.

CAUTION: The default settings on this page are usually correct for most applications. Only make changes if you are sure you need them and you fully understand the effects of the change. Incorrect settings **WILL** adversely affect your clock accuracy or even prevent clock corrections entirely.

Clock Corrections vs. Alignments

Domain Time can correct the clock either by "stepping" (immediately changing the time) or "slewing" (changing the time slowly). Stepping and slewing only operate on variances of 1 millisecond or more.

If slewing is enabled, variances of less than 1 millisecond are "aligned," which are very small slewed clock adjustments. Sub-millisecond alignments are NOT considered corrections, and will not show as corrections in some displays and reports, such as drift records, Audit Server reports, etc. Variances of less than 1 millisecond will be reported as zero milliseconds, except in the log files, drift graphs, or Manager's displays.

If your machine is stepped, the log file will say "Local clock stepped" (followed by details on which direction, by how much, and the protocol used to obtain the time).

If your machine is slewed, the log file will say "Local clock slewed" (followed by the same details as for stepping).

If your machine is aligned, the log file will say "Local clock aligned" (followed by the same details as for stepping or slewing).

Alignments happen automatically as long as slewing is enabled. The only important thing to remember about alignments is that they are not reported as clock corrections.

About Slewing and Stepping

There are two methods of correcting the system clock: Slewing or Stepping. Slewing means adjusting the system's overall clock rate so that the system either speeds up or slows down until it matches the wall clock. Stepping means an instantaneous jump to the new time, either forward or backward.

Slewing gives all processes a linear progression of ticks as time passes, and time is guaranteed never to go backward ("backwards" time corrections are actually made by causing the system clock to go forward at a slower pace than normal until the actual time catches up).

Slewing is critical to time-sensitive applications like databases, logging facilities, or auditing where a backwards jump in time would be highly disruptive. Slewing can also yield more precise corrections than stepping, and can correct variances of less than one millisecond. Slewing is therefore the correction method of choice.

Clock Corrections

- Slew the clock if possible, otherwise step it
- Only slew the clock, never step it
- Only step the clock, never slew it

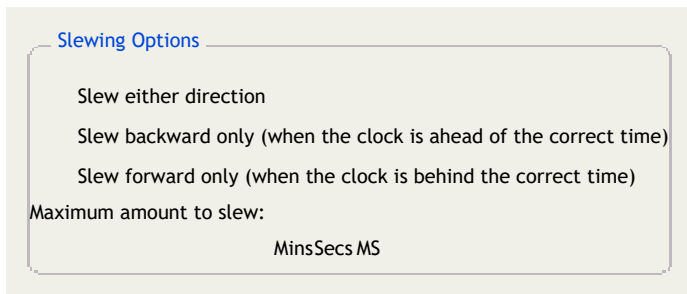
The **Clock Corrections** dialog selects whether corrections are made by slewing, stepping, or a combination of both.

Slew the clock if possible, otherwise step it

When this radio button is selected, Domain Time will step corrections too large to slew (or if slewing in that direction is disabled on the *Slewing Options* dialog page, see below), and will also step the very first correction after rebooting. This is

the default option and highly recommended.

The [Slewing Options](#) button brings up a dialog that lets you specify the direction of slew and the maximum amount of slewing that is permitted.



By default, Domain Time will slew both forward and backward to correct the clock, provided the correction being applied is within the set limit for slewing. The default limit is 30,000 milliseconds (30 seconds), but you may change this to anywhere in the range of 1 through 3,600,000 (1 millisecond through 1 hour). You may disable forward slewing, backward slewing, or both. Slewing large corrections can take an extended amount of time, so be careful if you modify this setting. If slews take a long time to complete, the clock can continue to drift significantly during the slew, making it very difficult to achieve accurate corrections.

Only slew the clock, never step it

In v4.x, you could change the default stepping behavior by modifying the "Never Step Clock" registry option. However, in v4.x, "Never Step" really meant "Do not step except on first boot or when triggered by an administrator," which was a bit ambiguous.

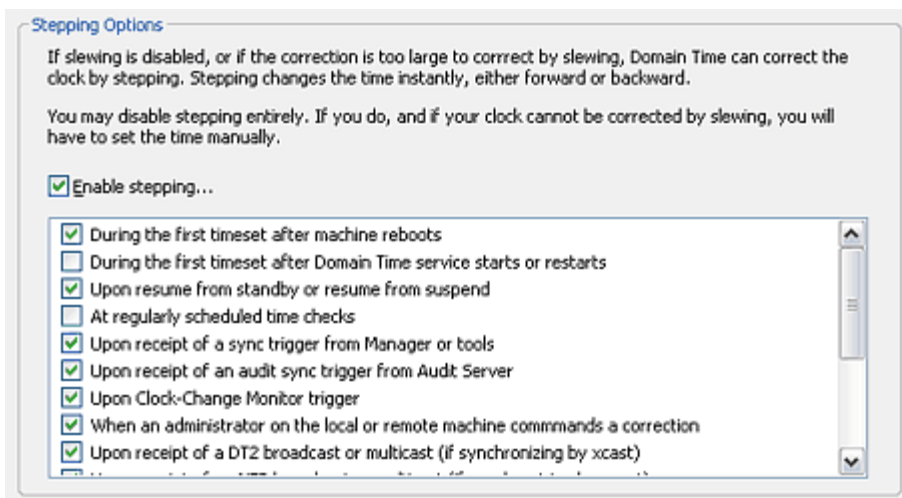
In v5.x, this setting is now available on the Control Panel applet. If it is enabled, Domain Time really will never step the clock. The slew limits and slew direction settings are not overridden by sync triggers, the control panel applet, or reboot detection. As a result, if you have this option selected, you will probably have to set the clock manually after every boot to get the time within the slew limit range to begin correcting the clock.

IMPORTANT: If this setting is enabled and the clock variance *is ever* outside the slew limit, the clock will never be corrected; a log entry indicating this condition will be made instead. Therefore, use care when choosing this option.

Only step the clock, never slew it

To provide greater control of the stepping process, v5.x introduces the [Stepping Options](#) dialog page which allows you to select the conditions under which stepping is allowed.

dialog page which allows you to

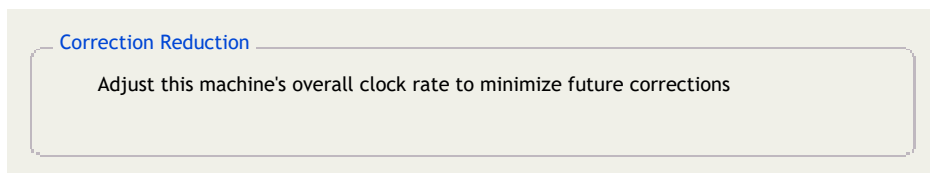


Stepping Options Dialog [\[Click for larger size\]](#)

The settings on this dialog correspond to the new "Allow Stepping" registry setting. "Allow Stepping" is a bitmask of the selected options. If your v4.x machine had "Never Step" specified in the registry, the value will be translated to an "Allow Stepping" value of zero when upgrading to v5.x. In all cases, stepping will only be applied if slewing is disabled or if the variance is outside the slewing limit (see above).

Correction Reduction and Advanced Clock Control

In addition to correcting the time during a time check, Domain Time can use highly-sophisticated clock control methods to ensure the clock on your machine runs more accurately, even between time corrections. This section shows whether these processes are enabled, as well as some statistics about the current clock management parameters.




This setting determines whether Domain Time will manage the rate of the system clock between time corrections. In nearly all cases, you will want to leave this checkbox checked.

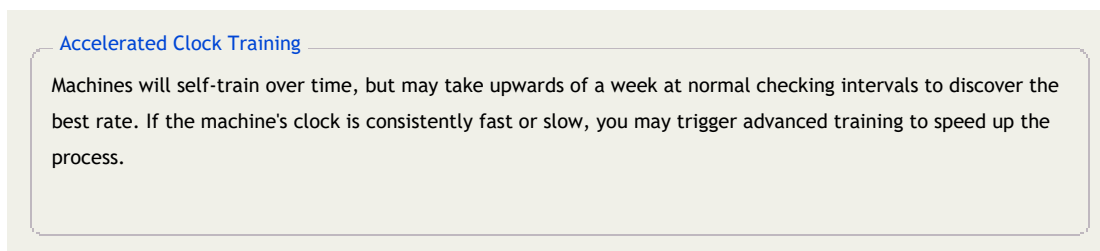
IMPORTANT: When unchecked, Windows will operate as if there is no time service controlling the clock between time corrections and the system clock will therefore run at whatever rate was last adjusted to (see the description of **Phase Adjustment** below). This may cause the clock to drift severely between corrections. More critically, because Windows thinks no process is managing the clock, it will periodically (usually ~ once an hour), hard-set the system clock to match the CMOS hardware clock on the motherboard, causing jumps in the system time. The CMOS clock is notoriously unreliable and thus the resulting time jumps may be very large.

The stability of the system clock on Windows is a result of complex and difficult-to-predict interactions of hardware and software timers, operating system processes, and running applications on each machine. Domain Time's Correction Reduction functions are extremely good at taking all these items into account; however, some machine's clocks are difficult to discipline automatically, and the default automatic algorithms can sometimes appear to make things worse.

If your machine gets progressively less accurate over time, please carefully investigate the options on the Advanced Clock Control dialog (see below) before turning off Correction Reduction entirely. It may be that resetting the timings to defaults, locking the clock rate, adjusting the Interphase settings, and/or changing the Timer Resolution settings will improve the accuracy dramatically on an errant system. Uncheck **Correction Reduction** only if you have a compelling reason to do so.

Advanced Clock Control

The  button will bring up the **Advanced Clock Control** dialog.



Machines running Domain Time II automatically train themselves over time to better match the speed of their sources. If a particular machine is consistently slow, for example, it will gradually speed up in order to reduce the frequency and magnitude of corrections. If it is consistently fast, it will gradually slow down. See the "Phase Adjustment" section below for a description of how this is done.

Self-training may take several days to several weeks to reach equilibrium across all the machines on your network since each machine has to collect data from its own time synchronizations, determine what type of adjustment to attempt, and then re-sample to determine if the adjustment was correct. In most cases, the self-training process will hone in on the best clock rate for each machine, and very few training changes will occur thereafter.

However, it may sometimes be desirable to take a shortcut to speed up the training process. The button initiates a special accelerated clock training sequence that causes an extended series of rapid clock synchronizations that allows the system to estimate what the correct clock rate is in a relatively short period of time.

Accelerated clock training is not as accurate as the automatic self-training over time, however it can get the machine close to the correct clock rate so that final self-training can fine-tune the rate more quickly. On machines with highly variable load where self-training cannot reliably determine a correct rate, accelerated training can be used to determine a decent clock rate value to use as a starting point to determine a locked rate (see below).

Recommendations:

- Use accelerated training on your master time server when you first deploy Domain Time II on your network. When the master is fully trained, then you can use accelerated training on each slave.
- Then, use accelerated training on your Independent Servers.
- Allow other machines to automatically train their clocks without accelerated training if possible.
- Do not use accelerated training on clients before their servers have finished either self-training or accelerated training, since this will result in inaccurate training of the client.

Multimedia Timer Resolution

Set multimedia timer to the maximum resolution for this machine (default checked)

Set OS timer to the maximum resolution for this machine (default unchecked)

In most cases, having these boxes checked will increase the accuracy of the system clock. However, they are global settings for all applications on the machine, and other applications may raise or lower the resolution unexpectedly. Also, changes to the OS timer in particular may have unintended consequences to other applications or system functions, so the default for that setting is unchecked.

If you find that your system's clock accuracy changes when certain applications are run (such as multimedia applications, Java applets, etc.) you may want to try disabling these settings, reset the clock rate settings on this page to the defaults (click the button, and then use the Accelerated Clock Training function (described above) to re-train the software while your application is active.

Phase Adjustment

Default phase adjustment: **156001**

Phase adjustment: (integral rate) Phase Locked

Interphase adjustment: (+-ms/minute) Interphase Locked Continuously variable (PTP)

Interphase period: 10 seconds Change rate limiter enabled: %

Interphase significance threshold: (average delta in hectonoseconds)

Interphase reliability threshold: (minimum ms/minute change allowed)

These settings display (and set) the current clock rate settings on the machine.

In most cases, you will not need to adjust these settings as Domain Time will usually do an excellent job of finding the optimum tuning values. However, you may be able to achieve higher accuracy/clock stability by adjusting these settings.

Default Phase Adjustment: displays the original clock tick rate calculated by the OS for the hardware of this machine. Windows expects that this rate will equal one second per second. In reality, it almost never does. The actual number of ticks necessary to run at exactly one second per second on this machine is usually some value greater or less than the default.

Phase adjustment:

Domain Time attempts to automatically determine the optimum tick rate setting and will adjust this value in small increments over time. The current tick rate setting in use is displayed in the **Phase adjustment:** field.

Checking the **Phase Locked** box will lock the clock at the current rate. This should be used with great care as it prevents any further automatic clock training. Use this setting if automatic clock training results in incorrect phase adjustments. For example, auto-training on some systems may result in the phase adjustment continually incrementing, causing the clock to run ever faster (or slower). Locking the rate will prevent the rate from incrementing.

Locking the phase rate can also be useful if auto-training results in the clock consistently running slightly slower or faster than the correct time. You can change the phase rate and lock it so that corrections are slightly behind the actual time (so that corrections are always speeding up the clock slightly) Then, use the Interphase settings described below to fine-tune the corrections the rest of the way.

IMPORTANT: On versions of Windows other than Vista, 2008, 2008/R2, and Win7, the tick rate can reliably be set to the actual resolution shown (i.e. changing the tick rate from 156250 to 156251 will change the clock rate by an exact number of microseconds per second). However, with Vista, 2008, 2008/R2, and Win7, Microsoft made underlying changes to how Windows handles the clock timers, resulting in a significant reduction in the granularity of this setting. In effect, on those operating systems a change of at least 16 ticks must occur before a change in clock rate actually happens on those systems.

In the example above, changing 156250 to 156251 will not have any effect. You must change the rate to 156266 ($156250 + 16$) before the clock rate actually changes. The next rate change will occur at 156282 ($156250 + 32$), etc.

In most cases, this decrease in granularity will result in a significant loss of accuracy unless you allow Domain Time to compensate for it properly using the **Interphase adjustment** setting below.

As of Windows 2012/Windows 8, Microsoft appears to have addressed the granularity issue so that single digit changes to the tick value again respond as expected.

The **Interphase adjustment:** settings allows for fine-tuning the clock rate when the correct time falls somewhere between the resolution of the tick rates set in the **Phase Adjustment** setting above. As described above, this is much more likely to be necessary on newer versions of Windows, since the granularity of phase adjustment is larger than on older versions. This value is the automatic clock correction value applied every minute (at the rate selected by the **Interphase period** setting below) to provide highly precise clock adjustment (and compensate for underlying timer errors/corrections on newer versions of Windows).

These corrections are usually made automatically, but if your system is unable to achieve the desired level of accuracy, you can manually set this value and lock it using the **Interphase Locked** checkbox. Use this value for extreme fine tuning after getting the clock as accurate as possible first if you are using NTP or DT2 protocols to synchronize time.

If you are using the IEEE 1588-2008 (PTP) protocol to synchronize time, it is usually better to allow Domain Time to continuously adjust the Interphase settings automatically to achieve significantly more accurate and smoother clock corrections. This function is enabled by checking the **Continuously variable (PTP)** checkbox. If you are using PTP and this function is enabled, the **Interphase adjustment/Interphase period** settings are ignored, even if locked. Note that if PTP sync is lost, these settings resume effect if the machine falls back to using another time

protocol like NTP or DT2.

The **Interphase period: 10 seconds** selection determines how often per minute Domain Time applies a fraction of the Interphase adjustment when attempting to exactly match this clock's phase to a time source.

For example, a setting of 15 seconds will apply 1/4 of the Interphase setting every 15 seconds. A setting of 20 will apply 1/3 of the Interphase value every 20 seconds, etc.

Due to the complex nature of the underlying timers, there is usually a "sweet spot" for this setting that results in maximum accuracy, but it can only be discovered by trial and error. It is probably not necessary to adjust this value unless you have first fine-tuned your accuracy as much as possible with all of the other settings.

Change rate limiter enabled: % controls how quickly the software makes changes to the Interphase settings based on the collected time samples. Setting this value too high may result in overly-large Interphase adjustments; too low of a setting may result in unnecessarily-small adjustments. Either condition can affect how long it takes to converge on accurate timings for the machine.

Interphase significance threshold:

Interphase reliability threshold:


These settings fine-tune the criteria the automatic Interphase adjustment algorithm uses to calculate the Interphase rate. You should not change these values without advice from Technical Support.

Import/Export

You may import or export Domain Time II configuration settings using the utilities on this page.

Starting with version 5.1, Domain Time Client can import/export its settings using a standard Windows Registry .reg file. This allows you to easily make backup copies of the settings, create a custom .reg file to use as a template for configuring other machines, or to use in creating a custom installation package or script.

The Import/Export function automatically excludes any settings that are machine-specific so that the .reg file may safely be imported on any other copy of Domain Time Client (of the same version) without causing disruption.



This section allows you to save the current configuration to a .reg file.

The .reg file will be created using the currently configured options of the Client. You should review each of the settings of the Control Panel applet to be sure that they are correct before exporting the file.

If you expect to be importing the .reg file on machines that need differing configurations (such as for machines in different cities that use different time sources), you should configure the applet for each configuration, and then export a separate .reg file for it.

You may use exported .reg files as template files for installing/upgrading multiple machines using [Domain Time Manager](#). Template files need to be located in the **C:\Program Files\Domain Time II\Templates\[Server][Client]** folder of the Domain Time II Manager machine. Template .reg files located in those folders will automatically be made available for use when installing or upgrading using Manager.

Note, if you launched the Control Panel applet using Manager, the export utility will automatically offer to save the .reg file in the proper directory on the Manager machine. Otherwise, you will have to manually save or copy the file to the Templates folder on the Manager machine.

Exported templates may be edited manually using any text editor. Templates may contain all exported settings, or only the specific settings you want to change.

IMPORTANT: Although .reg files created using this utility are saved in standard Windows registry file format, it is **not** equivalent to exporting the registry keys using Windows' RegEdit program. A number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. Those settings are automatically excluded when you export using this utility, so you should always use this utility to create a Domain Time .reg file.

About Settings (.reg) Files

Domain Time Servers and Clients are background system services that obtain all of their running settings from the Windows Registry.

Domain Time components get their initial registry settings by importing default template files during installation or upgrade. As of version 5.1, the templates are standard Windows Registry Editor .reg files in either Unicode or ANSI text format. Previous versions used a proprietary .ini template file (domtime.ini), which has been deprecated.

The default template file for Server is **dtserver.reg**; Client uses **dtclient.reg**. These files are included in the original distribution files used during Setup or in the source files used for remote installation by Domain Time II [Manager](#).

During installation/upgrade, the appropriate template file is copied from the distribution files to the **Windows/System32** folder of the target machine.

Once installed, Domain Time will not use the default .reg template file again, unless the user clicks a Reset to Defaults button on the Control Panel Applet, or if a Reset Configuration is commanded remotely from Domain Time II [Manager](#).

Load Settings from File

Installation defaults (the settings used when Domain Time was first installed on this machine)

Choose File:

Do not prompt for confirmation

Use this section to import the default settings or a custom .reg file.

Installation defaults (the settings used when Domain Time was first installed on this machine)

This selection will reload the settings file that was used when the product was installed.

The file is named **dtclient.reg** and is located in the **/System32** folder on a running system. This is the default file used during installation. See the [Rollout](#) page of the installation instructions for more information on using this file as a setup template.

Choose File:

Use this to import an existing .reg file.

CAUTION: It may be necessary to restart the Domain Time Client service after importing the settings file.

Although the .reg file is saved in standard Windows Registry file format and you can install it by clicking on the .reg file in Windows Explorer, it is usually better practice to import the file using this utility since it does additional validation checking on the values and attempts to exclude items not appropriate to this version or machine.



This property page contains Domain Time Support information and utilities.

Problem Report

Your company:	Include main log file and drift graph
Your name:	Include most recent startup log file
Your phone:	Include registry settings

The **Problem Report** utility can compile a problem report to send to Domain Time technical support including important diagnostic information and log files which will greatly assist in troubleshooting any problems you may experience. You can either email it directly from the program or save the file to forward it manually later.

You have the option of including various items and logs in the report in a compressed (zipped) file. In most cases, you should include all items to provide the most information possible.

The utility will use the currently-selected default MIME email program on the machine where this utility is run to send the mail and compressed file attachment when you click the button.

If you don't want to (or can't) send email directly from the machine in question, remotely connect to the problem machine from a machine that does have email capability. You can connect remotely using Domain Time Manager, the Remote CPL utility, or from the Control Panel applet of another Domain Time Server or Client. Once you have the remote machine's Control Panel applet displayed, click the button. The utility will send the email from the email client on the local machine, but will automatically include the diagnostic and log files taken from the remotely-connected machine.

If you'd prefer to save the report file to disk and forward it to Tech Support manually, click the button.

Domain Time Client has the option of displaying a handy application (DTTray) in the Windows System Tray.

Note: The System Tray is a graphical function of Windows Explorer, so it is not available on Windows Server Core. This means the DTTray program is inaccessible on that version of the operating system.

DTTray gives you quick access to common tasks (such as manually triggering the machine to sync with its time source) as well as a quick way to launch Domain Time applications.

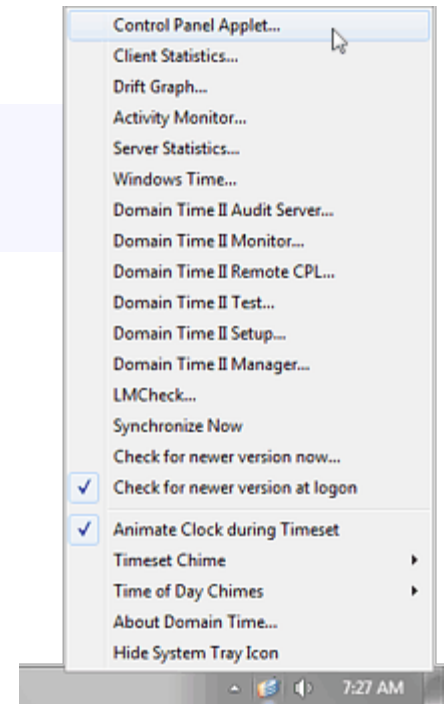
It also gives you visual and audible alerts to your current time sync status, traffic monitoring displays, drift analysis graphs, etc.

Starting the Applet

The DTTray applet loads into the System Tray automatically when you log in.

Double-clicking the system tray icon will pull up the Domain Time Client Control Panel applet. Right-clicking the icon will present you the full DTTray context menu. Only Domain Time applications currently installed will appear on the context menu.

Items installed or removed while the tray applet is loaded may not appear on the DTTray context menu until the user logs off or the system is restarted.



System Tray Applet [\[Click for larger size\]](#)

Hiding the System Tray Applet

You can hide the system tray applet by clicking the **Hide System Tray Icon** item on the DTTray context menu itself.

You can show or hide the system tray applet by setting the **"Show system tray icon"** checkbox on the [Advanced](#) property page of the Control Panel applet.

You can also set the value of the `HKLM\Software\Greyware\Domain Time Client\Parameters\SystemTrayIcon` registry key if you prefer.

Multiple Instances on Terminal Services/Remote Desktop

By default, the system tray applet will only appear in the first logged-on instance of the console. If you want the system tray to appear in all terminal sessions, you can set the value of the following registry key to **True** :

Location: `HKLM\Software\Greyware\Domain Time System Tray Icon\Parameters`

Key: (create it if it doesn't exist) `Allow Multiple Instances`

Type: `REG_SZ`

Note: Each running instance of the icon holds a file lock on the tray icon executable. If you enable this feature, it will be necessary for all users (local and remote) to log off before performing an upgrade of Domain Time on this machine.

System Tray Applet Command Functions

Use the context menu to trigger a sync and check for updates.

Triggering a Time Sync

You can trigger the time service to synchronize with its time source by right-clicking system tray icon and choosing **Synchronize Now** from the context menu.

Alert Functions

The DTTray Applet has a number of features that provide visual or audible feedback on the status of your clock synchronization.

Animate Clock During Timeset

While Domain Time is synchronizing the clock, the system tray icon will show a running clock icon. When the clock is successfully synchronized the icon will change to the standard Domain Time icon. You can turn off this feature by unchecking the option from the right-click context menu.

The "Time Not Synchronized" Alert (The Flashing Clock)

This alert flashes a clock in the system tray to indicate the time is not synchronized.


If Domain Time is unable to synchronize its time with a time source, it will alert you to the problem by changing the Domain Time icon in the system tray to a flashing clock icon. Once you have resolved the cause (usually due to a network issue preventing Domain Time from contacting its time source) and re-synchronized, the icon will return to the normal Domain Time icon.

Timeset Chime

The System Tray Applet can indicate a successful time synchronization with an audible signal.

Timeset Chimes are off by default. You can enable them by pulling up the context menu and selecting the sound device you wish to use to play the chimes. The Timeset Chimes will play whether or not there is a logged-in user.


Choose **Chime off** if you do not want the signal to play.



Timeset Chimes [\[Click for larger size\]](#)

Time of Day Chimes

The Time of Day Chimes feature is a special function that plays sound files at particular times of the day (on-the-hour, and at 15, 30, 45 minutes past the hour) to emulate a chiming clock. You may download selected free chime packs from our site or create your own.



Time of Day Chimes [\[Click for larger size\]](#)

The Time of Day Chimes are off by default. You can enable them by pulling up the context menu and selecting the chime pack you want to use. (You must have downloaded and installed at least one chime pack before this feature will be available.)

In order to play the Time of Day Chimes, you must meet these requirements :

- Your system must be configured with a sound card, drivers, and other hardware (such as speakers or headphones) necessary to play .WAV files.
- Time of Day Chimes are played by the DTTray Applet, so you must be logged in and have the System Tray Applet installed and visible in the system tray if you want the Time of Day chimes to play.

You must have downloaded and installed at least one chime pack (see below).

Choose **None** if you do not want the chimes to play.

How to Install Chime Packs

Chimes are standard .WAV sound files played using the Windows Media subsystem. You can choose from free chime packs that we've prepared for you to download or you may create your own.

Download	Hear a Sample
Grandfather Clock	Listen
Westminster Classic	Listen
Westminster Glass	Listen
Westminster Deep	Listen
Cuckoo Clock	Listen


To install a chime pack, download the zip file and unzip the contents into your Media folder (usually **C: \Windows\Media**). There will be one text file (.txt) and one or more sound files (.wav) contained in the zip. The text file contains a description of the chime pack, and instructions for which sound file goes with which event. These instructions are used by the DTTay Applet when you chose a chime pack.

Important: The .txt file from the chime pack must be copied into the Media folder along with the .wav files. It must be present for the chime function to work.

After unzipping the files into your Media folder, right-click the Domain Time II system tray applet. Your newly-installed chime pack should show up by name under Time of Day chimes.

Configure and Customize your Chimes

Chimes are fully configurable using the Windows **Sounds** Control Panel applet (called *Sounds and Audio Devices* on older versions of Windows).



Time of Day Chimes [\[Click for larger size\]](#)

To customize your chimepack, launch the **Sounds** (Sounds and Audio Devices) applet, then click the **Sounds** tab. Scroll down through the list until you see the Domain Time II sound scheme. Listed under Domain Time II, you'll see entries for "Every hour at 15 after," "Every hour at 30 after," and so on. You may associate any .WAV files you like with the various sound events.

The "Every hour tolling the hours" sound is played after any other chimes on the hour, and is played a number of times corresponding to the hour (using a 12-hour scheme). At one o'clock, it will play once, at two o'clock it will play twice, and so forth.

You can make your own chime pack by collecting the .WAV files you want to use and creating a text file for them. Download one of the chime packs from the list above and look at how the text file specifies the sounds. To have your own chimepack show up so you can select it from the System Tray Applet context menu, simply create a new text file with a .txt extension. The file should be in the same format as the sample chimepack, listing which .WAV plays for

which event. Use a full or relative path before the filename if your .WAV files aren't located in the Media folder. You'll notice from the sample .txt files that you can create a very elaborate set of chimes played at various times and events. Feel free to experiment!

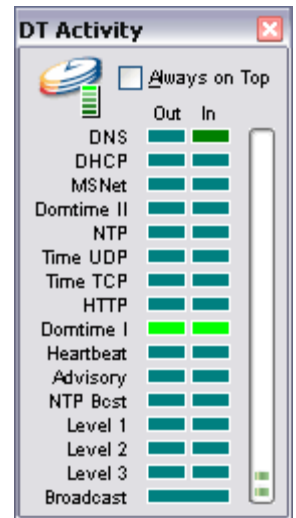
Activity Monitor

This window displays all Domain Time-related network traffic sent or received by this machine.

The Activity Monitor provides a visual indicator of when the various types of protocol and control messages are sent or received. Besides being a great deal of fun to watch, it is a useful diagnostic tool you can use to verify that your server or client is receiving and responding to time sync requests and cascade signals.

For example, click the Synchronize Now option from the DTTray context menu to watch how your system performs a sync request.

Click the **Always on top** box if you wish the Activity Monitor to always be visible on your screen.



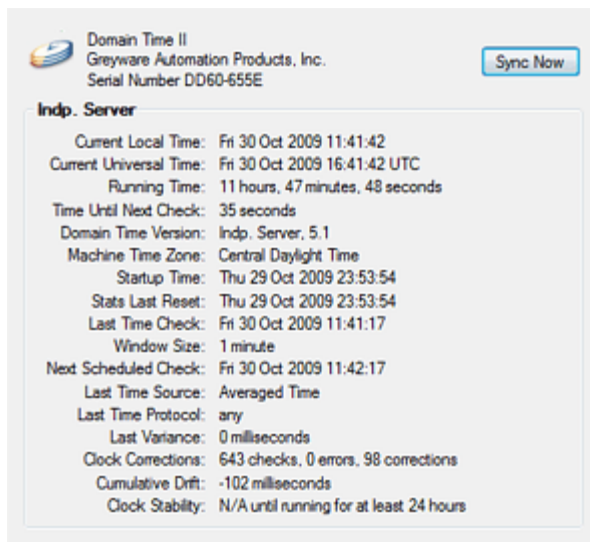
Statistical Functions

View detailed statistics and view drift graphs

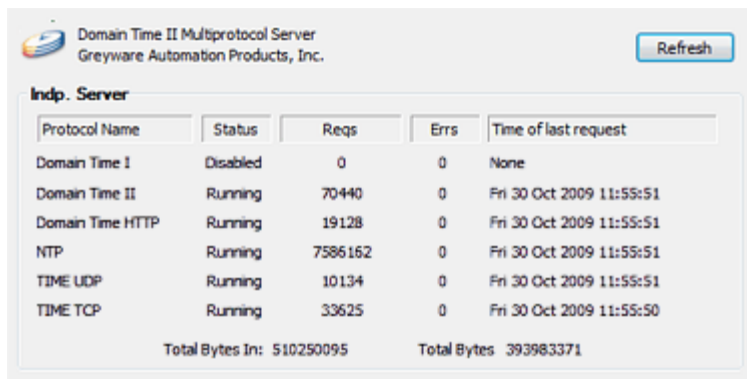
Viewing System Statistics

You may view the current time statistics on this machine by double-clicking the system tray icon, or by right-clicking the icon and choosing **Client (and/or Server) Statistics** from the menu.

Note that Domain Time II Server acts as both a time client and a time server, so there are two statistical displays available on Server.




[Client Statistics \(on Server\) Display](#) [\[Click for larger size\]](#)



[Server Statistics](#) [\[Click for larger size\]](#)

Drift Graph

In addition to detailed summary statistics, DTTray can show you a graphical representation of the accuracy of the clock on your machine sampled at the time it synchronized with its time source. You may scroll through the entire drift data to see how your clock has been performing over time.



The Drift Graph [Click for larger size]

You can also see the actual sync data in text format by clicking the [Raw Data](#) button. This data can also be collected and analyzed centrally using [Audit Server](#).



Domain Time Client keeps its settings in the Windows Registry. Most of the service options are best set using the Domain Time Client Control Panel applet. However, some advanced options can only be set by changing the registry. This page explains many of these special registry entries used by Domain Time Client.

CAUTION:

Modifying Registry entries requires basic familiarity with the Windows Registry and its operations. Incorrect changes to the Registry can result in unpredictable, perhaps non-repairable, damage. We cannot be responsible for registry problems.

The Domain Time II Client settings are located in these keys (click the names to jump to details):

[HKEY_LOCAL_MACHINE](#)

[Software](#)

[Greyware](#)

[Domain Time Client](#)

[Enabled Protocols](#)

[Keyring](#)

[Logs and Alerts](#)

[Parameters](#)

[Time Sources](#)

Enabled Protocols

The Domain Time II Client Enabled Protocols settings are located in this key:

[HKEY_LOCAL_MACHINE](#)

[Software](#)

[Greyware](#)

[Domain Time Client](#)

[Enabled Protocols](#)

The values listed in the **Enabled Protocols** registry key represent the protocol types Domain Time will listen for. They correspond to checkboxes on the [Status Reports](#) and [Advanced](#) property pages of the Control Panel applet. You should not make manual changes to this key or its subkeys.

Keyring

The Domain Time II Client Keyring settings are located in this key:

[HKEY_LOCAL_MACHINE](#)

[Software](#)

[Greyware](#)

[Domain Time Client](#)

[Keyring](#)

The values listed in the **Keyring** registry key contain various items related to authentication. They correspond to settings on the [Symmetric Keys](#) property page of the Control Panel applet. You should not make manual changes to this key or its subkeys.

Logs and Alerts

The Domain Time II Client Logs and Alerts settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time Client
Logs and Alerts

The values listed in the **Logs and Alerts** registry key contain various items related to logging and alerting functions. They correspond to settings on the [Logs](#), [Windows Event Viewer](#), [Syslog](#), [SNMP](#), and [Status Reports](#) property pages of the Control Panel applet. You should not make manual changes to this key or its subkeys.

Parameters

The Domain Time II Client Parameter settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time Client
Parameters

The values listed in the **Parameters** registry key control a wide variety of Domain Time functions. In most cases, they are auto-generated or correspond to settings on the property pages of the Control Panel applet. In general, you will not need to make manual changes to this key or its subkeys.

However, some values require additional explanation or control functions not exposed on the Control Panel. Those items are listed here.

Value Name:	Accept First PTP Timestamp
Value Type:	REG_SZ
Default Data:	<i>False</i>
Options:	<i>True or False</i>
Notes:	If set to <i>True</i> and no other time sources are configured, then the clock will be stepped or slewed if within slewing limits to match the first PTP timestamp(s) received (the number of samples required are configured using the Accept Firest PTP Sample Count registry entry described below). This initial correction will bring the clock into close enough sync for normal

PTP operations to govern the clock. Note, on versions prior to 5.2.b.20200930, the clock adjustment was always stepped.

IMPORTANT: Changing this setting to *True* is discouraged in networks with fallback NTP/DT2 timesources, since a restart of the service may step the clock, including possibly stepping the clock backwards. This option should only be used in closed environments where PTP is the only possible source of time and the initial startup delta takes an excessively long time to correct (i.e. if the motherboard CMOS clock is wrong).

Value Name: Accept First PTP Sample Count

Value Type: REG_DWORD

Default Data: 3

Range: 1-15

Notes: Introduced in v5.2.b.20200930. Applies only if **Accept First PTP Timestamp** is enabled. Specifies the number of timestamps that must be received before the clock is adjusted.

Value Name: Allow Remote Timezone Change

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: Enables Domain Time II Manager to change the timezone on this machine.

Value Name: Allow Stepping

Value Type: REG_DWORD

Default Data: Varies

Notes: New as of v5.1, this value is a hex bitmap representing the settings made on **Stepping Options** dialog of the [Clock Control](#) property page. Do not edit this value.

These values will be overridden if the **Never Step Clock** setting (see below) is enabled.

Value Name: Client Settings

Value Type: REG_BINARY

Default Data: Varies

Notes: This value is a hex bitmap of various settings used by Domain Time Client, such as [Timings](#), [Corrections](#), and other miscellaneous settings.

Note: On Domain Time Server, this key controls miscellaneous [settings recommended](#) by a Domain Time Master Server to its Clients. See the [Server Settings](#) key for the miscellaneous

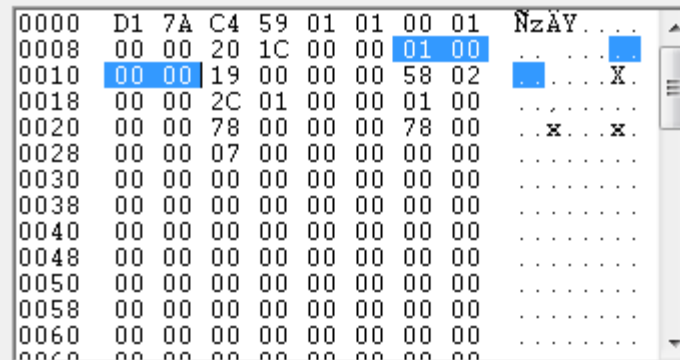
settings for Domain Time Server.

In general, you should not edit these settings manually. Use the Control Panel applet to configure your settings instead.

However, as of v5.2.b.20170922, the applet setting for Minimum Correction (MinDisparity) has been removed. The setting for this value defaults to 0x1, but if you upgraded over a previous version with a higher setting, you may edit the binary key to change it. The Minimum Correction setting is a DWORD, starting at offset 14, stored in little-endian order, as shown below.

The default should be **01 00 00 00**, see this example:

Value data:



0000	D1	7A	C4	59	01	01	00	01	NzAY...
0008	00	00	20	1C	00	00	01	00	...
0010	00	00	19	00	00	00	58	02	...X.
0018	00	00	2C	01	00	00	01	00	...
0020	00	00	78	00	00	00	78	00	...x...x.
0028	00	00	07	00	00	00	00	00	...
0030	00	00	00	00	00	00	00	00	...
0038	00	00	00	00	00	00	00	00	...
0040	00	00	00	00	00	00	00	00	...
0048	00	00	00	00	00	00	00	00	...
0050	00	00	00	00	00	00	00	00	...
0058	00	00	00	00	00	00	00	00	...
0060	00	00	00	00	00	00	00	00	...

Value Name: Clock Adjustment Bucket Size

Value Type: REG_DWORD

Default Data: 7

Range: 3-32

Notes: The bucket size is the number of time samples collected before a particular clock adjustment rate is evaluated. The specified value is used except during accelerated clock training, where a fixed value of 5 is employed.

You should not change this number unless instructed by techsupport.

Important: This is a machine-specific setting and should not be included in installation templates or copied to other machines via mass registry imports.

Value Name: Clock Adjustment Statistical Method

Value Type: REG_SZ

Default Data: Automatic

Notes: Sets the type of statistical analysis Domain Time performs on collected time samples from a time source when deciding whether and how much to adjust the clock rate to compensate for drift. Domain Time then uses the calculated clock performance to evaluate and remember each integral clock adjustment rate it tries. Changing this value may improve or degrade timing accuracy (or have no effect).

Changes to this setting take effect immediately after the next group of collected samples is ready for analysis. You do not need to restart the service. You should clear your clock history using the command **dtcheck -resettimings** before changing this value. Allowed values are:

- **Automatic** - On Vista/2008/Win7/2008r2 machines, Automatic will use the median value from each group of samples. On all other versions of Windows, it will use the arithmetic mean (average) of each group of samples.
- **Average** - The arithmetic mean of values
- **Median** - The median number in the array of values
- **Toss** - average of values excluding the highest high and lowest low
- **RMS** - the quadratic mean (signed root-mean-square) of the array of values
- **Disabled** - no statistical analysis is retained for future comparison

Important: This is a machine-specific setting and should not be included in installation templates or copied to other machines via mass registry imports.

Value Name: Clock Change Monitor

Value Type: REG_SZ

Default Data: *True*

Options: *True or False*

Notes: If enabled, Domain Time monitors changes to the system clock made by other programs (including the foreground user changing the time or date with the Control Panel applet or the command-line TIME and DATE commands). When the Clock Change Monitor is enabled on a Client and the clock changes unexpectedly, the Client will immediately resynchronize with its time source(s).

You may turn the Clock Change Monitor off if your setup requires having machines with different times (usually only in labs or testing environments). If Clock Change Monitor is disabled and you change a machine's time, it will stay changed until the next cascade signal or regular sync interval. Changes take effect immediately, and may be made by editing the registry or remotely from Domain Time II Manager.

Value Name: Clock Change Sensitivity

Value Type: REG_DWORD

Default Data: 0

Range: 0-255

Notes: This value represents the number of seconds the system clock must differ from the expected value in order for **Clock Change Monitor** to decide an unauthorized change has been made to the system clock.

If not present or set to zero, Domain Time will use a value of 2 seconds.

Increase this value only if **Clock Change monitor** is triggering on normal clock drift (unlikely). Decrease this value only if **Clock Change Monitor** is not flagging known clock

change events by another user or process.

Value Name: Current Version

Value Type: REG_SZ

Default Data: Varies

Notes: This value is set by the system for informational purposes. Changing it has no effect.

Value Name: Critical Timing Processor Limit

Value Type: REG_SZ

Default Data: Depends on processor type (see below)

Options: *True* or *False*

Notes: This value is set to *False* during installation on machines with processors that have an Invariant TSC or if they are a Hyper-V guest; otherwise it is set to *True*. When *True*, Domain Time uses the last-processor-but-one for time-critical events, and any available processor for all other work. If set to *False*, Domain Time does not prefer one processor over another for any task.

Modern CPUs (ones with Invariant TSC) generally have better timing performance with this value set to *False*. You can check to see if you have an Invariant TSC by running the command-line DTCheck program:

```
dtcheck -cpuid
```

Value Name: Dependent Services

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more services you'd like Domain Time to start.

Notes: Requires version 5.2.b.20150516 or later. Any services listed here will be started by Domain Time after the first successful timecheck, as long as the services are set to manual start. This is an alternative to using the built-in service database's dependencies. If you use the built-in functions, dependent services will wait for Domain Time to start, but won't know to wait until the first synchronization has completed.

You may list services by their display names (e.g. "Disk Defragmenter") or by their internal service names (e.g. "defragsvc"). List services one per line, without quotation marks. Domain Time will only attempt to start services that are listed, not yet running, and set to manual startup.

Important: If Domain Time cannot set the clock for some reason (invalid sources, firewall settings, etc.), then services you have set to manual start will not be started.

Value Name: Ephemerides

Value Type: REG_DWORD

Default Data: N/A

Notes: This value is used by the system. Do not edit.

Value Name: ICMP TTL (hop limit)

Value Type: REG_DWORD

Default Data: 32 (decimal)

Range: 1 to 255 (decimal)

Notes: This value controls the number of router hops that are allowed in an ICMP echo ("ping") request. Domain Time pings machines first to help eliminate long waits for machines that are unreachable. You should only need to adjust this value if you have an LAN/WAN configuration requiring more than the default 32 hops.

Value Name: Machine Statistics

Value Type: REG_BINARY

Default Data: N/A

Notes: This binary value contains the statistics, as of the last update, that can be viewed from DTCheck, the Domain Time II Manager, or the system tray icon. Do not edit.

Value Name: Max Slew Correction (milliseconds)

Value Type: REG_DWORD

Default Data: 30000 (decimal)

Range: 1 to 36000000 (decimal)

Notes: This value specifies the upper limit, in milliseconds, of variance that Domain Time will attempt to correct by slewing instead of stepping the clock. This setting affects both forward and backward clock adjustments.

The older registry entry controlling this function, **Max Slew Correction (seconds)**, has been deprecated.

If the correction to be made is larger than this setting but less than the allowed MaxDisparity setting (Correction Limit), Domain Time II will step the correction (unless **Never Step Clock** is enabled, at which point no correction is made and a note to this effect will be entered in the Domain Time logs). See the **Never Step Clock** and **Override Max Disparity** registry settings for more info.

Value Name: Min Success Interval (seconds)

Value Type: REG_DWORD

Default Data: 5

Notes: Sets the minimum period allowed between timechecks. Do not change this value.

Value Name: Never Step Clock

Value Type: REG_SZ

Default Data: *False*

Options: *True or False*

Notes: When enabled, causes Domain Time to make clock corrections only by slewing. This prevents the clock from being stepped to make corrections such as those normally done during startup or from Clock Change Monitor, manual sync triggers, etc.

CAUTION: Enable this option with care. Use of this option may prevent Domain Time from successfully being able to synchronize with a time source if the time correction is too large to accomplish using slewing. See the **Max Slew Correction (milliseconds)** registry setting for more info.

IMPORTANT: Unlike with versions prior to v5.1, the behavior of this setting is **NOT** modified by the **Override Max Disparity** registry setting. If **Never Step Clock** is enabled, the clock will never be stepped, regardless of any other settings.

As of v5.1,, Domain Time uses the **Allow Stepping** setting (see above) to provide greater control of the stepping process. If your machine running an older version of Domain Time had **Never Step Clock** specified in the registry, the value will be translated to an **Allow Stepping** value of zero when upgrading to v5.x or later. See the **Stepping Options** dialog of the [Clock Control](#) property page to set the options.

In most cases, it is better to set the **Stepping Options** with the behavior you want than to enable **Never Step Clock**.

Value Name: NTP Client Version

Value Type: REG_DWORD

Default Data: 4 (was 3 on versions prior to 5.2.b.20150516)

Range: 1 to 7

Notes: Controls the reported NTP version. Any value from 1 to 7 is legal, although using anything but 3 or 4 is not recommended.

Value Name: Override Max Disparity

Value Type: REG_DWORD

Default Data: *Not present* (same as zero)

Options: 0, 1, 2, 3, or 4

Notes: Controls how Domain Time decides when to override the **Correction Limits** set in the Control Panel applets for Server, Slave, or Client Timings as explained below. This allows for setting the clock under certain conditions that would otherwise prevent a correction.

IMPORTANT: As of v5.1, none of these settings modify the **Never Step Clock** setting (see above). Note that this is a change in behavior from older versions. Enabling **Never Step Clock** effectively limits corrections to the **Maximum Slew Correction (milliseconds)** value, even if a larger correction would otherwise be permitted by **Override Max Disparity**.

- 0 or not present (Auto)
Domain Time will override the disparity settings during startup, on Clock Change Monitor event detection, receiving sync triggers/cascades from management components, or from Control Panel applet (CPL) signals.
- 1 (Always)
Domain Time will **always** override the disparity settings. This is the same as not having disparity settings at all. Always honors **Never Step Clock** setting.
- 2 (Never)
Domain Time will **never** override the disparity settings. Always honors **Never Step Clock**. This option may prevent your machine from syncing until you manually set the time to within the set Min/Max disparity range. If the machine is a Domain Time Server, it will normally refuse to serve the time until its own time has been set, so selecting a value of 2 may impact your entire network.
- 3 (Startup only)
Domain Time will override the disparity settings only until the first time after startup that it has set its own time correctly. Thereafter, it behaves as if you had set the option to 2.
- 4 (Limit CCM)
Clock Change Monitor signals do not override the disparity settings. Startup, management, or CPL signals **will** override the disparity settings.

Changes to this value take effect immediately. You do not have to stop and restart the service or reboot the machine.

Value Name: Override Sanity Checks

Value Type: REG_DWORD

Default Data: *False*

Options: *True or False*

Notes: To prevent accepting obviously-wild time corrections, Domain Time will (by default) refuse to

set the time outside of a defined range of acceptable correction. Backwards-correction is limited to the build date of the software - 1 year. Forward-correction is limited to 11:59:59 on 12/31/2036 due to NTP and UNIX Year 2038 date calculation issues.

However, Windows itself will allow setting the local clock outside of this sanity-checked range. Set this value to *True* to permit Domain Time to set the clock to any time/date the operating system will allow.

CAUTION: Change this value only if you have a clear requirement to do so.

Value Name: Send Port Generic

Value Type: REG_DWORD

Default Data: 0

Notes: Domain Time uses several sockets for generic outgoing messages. By default, the port used is an ephemeral port assigned by the system. This is the proper behavior for client-server systems; only the server should have a fixed listening port, and clients should use ephemeral ports. However, in rare cases, other applications have high-number ephemeral ports hard-coded as their communications ports. If Domain Time happens to start first, and happens to obtain those particular ports, the hard-coded applications may fail.

Set this value to the beginning port number (n) of a range you want Domain Time to use for its generic outgoing sockets. Domain Time will attempt to use (n) through (n + 50) to bind its generic outgoing sockets. If none of the ports (n) through (n + 50) are available, Domain Time will revert to letting the system choose an ephemeral port.

IMPORTANT: Be very careful not to specify any well-known ports or IANA-registered ports for your range, and only set this value if you have a specific problem that you know will be solved by changing the ephemeral ports Domain Time uses.

Value Name: Server Answer IP

Value Type: REG_MULTISZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses

Notes: This value corresponds to the "Listen only on these addresses" list on the Network tab of the Control Panel applet. If this value is not present or is blank, Domain Time will answer on all IP addresses bound to all interfaces present on the machine. Otherwise, Domain Time will only bind to the IP addresses you provide. You may provide IPv4 or IPv6 addresses, and may also use NetBIOS or DNS names. The addresses/names you provide must exist and be permanently assigned to the machine. This setting is useful chiefly in situations where the machine is multihomed and you want Domain Time restricted to particular interface(s). This setting affects all listening ports for Domain Time Client, unless individual protocols are overridden (see below). You must restart the service (or reboot the machine) for changes to take effect.

Note: Because this value is highly machine-specific, it is not included in template imports or

exports. You must set it individually on each machine.

As of version 5.2.b.20130221, you may also use CIDR notation to specify ranges of addresses. For example, 192.168.10.0/24 would bind to any address between 192.168.10.1 and 192.168.10.254, as long as one or more of those addresses was assigned to the machine. This is useful for machines using DHCP: you may restrict Domain Time to a particular network without knowing what IP the machine will have.

Value Name: Server Answer IP Override DT2

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with DT2/udp and DT2/tcp protocols

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by DT2 traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Server Answer IP Override NTP

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with NTP

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by NTP traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Server Answer IP Override PTP

Value Type: REG_MULTI_SZ

Default Data: (blank)

Options: Blank, or a list of one or more IP addresses for use with PTP

Notes: Requires version 5.2.b.20130221 or later. If this value is not blank, Domain Time will use it to bind to the IP addresses you specify for use by PTP traffic. As with "Server Answer IP" above, you may use CIDR notation to specify networks without specifying individual IPs. Unlike "Server Answer IP," this value is included in template imports and exports. This value is not configurable using the Control Panel applet. You must restart the service (or boot the machine) for changes to take effect.

Value Name: Service Installed

Value Type: REG_SZ

Default Data: N/A

Options: *True or False*

Notes: Used internally. Do not edit.

Value Name: Service Log Filename

Value Type: REG_SZ

Default Data: [not present]

Notes: Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename **domtimec.log** in the **%SystemRoot%\System32** folder. The complete path and filename must be specified (i.e. **C:\Windows\System32\domtimec.log**) and the drive specified must be a local drive.

Value Name: Service Running

Value Type: REG_SZ

Default Data: N/A

Options: *True or False*

Notes: Used internally. Do not edit.

Value Name: Set Processor Affinity

Value Type: REG_DWORD

Default Data: 0

Options: 00-FF (hex)

Notes: **Note:** This value has been deprecated in version 5.x and later; see the *Critical Timing Processor Limit* value instead.

If not present or set to zero, Domain Time will not attempt to restrict time-sensitive operations to any particular processor in a multi-processor system. In some systems, the majority of hardware interrupt handling occurs on only one processor (typically processor 0), so it may provide increased accuracy if Domain Time uses only other processors during time-sensitive operations. This value is a hex bitmap representing the processors in the system, with bit 0 representing the first processor, bit 1 representing the second processor, and so forth.

Value Name: Test Mode

Value Type: REG_SZ

Default Data: *False*

Options: *True or False*

Notes: Corresponds to the Test Mode checkbox on the [Advanced](#) property page of the Control Panel applet. If enabled (*True*), Domain Time will go through all the motions of obtaining the time and calculating variances, but will not actually set the clock. If disabled (*False*, the default value), Domain Time will set the clock after obtaining the time from its time source(s). Changes to this value only take effect after restarting the service.

Value Name: TIME/ITP Offset (seconds)

Value Type: REG_DWORD

Default Data: 2208988800 (decimal)

Notes: Used internally by the system. Do not change this value unless instructed to do so by tech support.

Value Name: Time Sample PreFilters

Value Type: REG_SZ

Default Data: HighLow

Options: Allowed options are HighLow, Latency, Delta, and Stratum. Prefilters are applied in the order listed; separate filter names with a comma or semi-colon.

Notes: Requires version 5.2.b.20150828 or later. This value controls the prefilters used to discard samples before applying statistical analysis. Prefilters only operate when there are five or more samples available for analysis, and are chiefly useful when the number of samples is very large, or the sources are unstable. It is best to leave this value at the default, which eliminates only egregious spikes. Statistical analysis of the entire group of samples usually performs better than prefiltering more samples out of the mix.

For example, HighLow,Latency would apply first the Highlow filter, then if at least five samples remain, the latency Filter. Stratum,Latency,Delta would first apply the Stratum filter; then if at least five samples remain, the Latency filter; then, if at least five samples remain, the Delta filter. Changes to the list of prefilters are recognized only when parameters are reloaded (server stop/restart, machine reboot, a CPL-initiated sync, or a DTCheck /reload).

Prefilter operations are:

- [HighLow](#) (default) - Rejects the most extreme samples, based on absolute magnitude delta (max of 2 samples rejected)
- [Latency](#) - Rejects highest latency samples (max 1/3 of samples rejected)
- [Delta](#) - Rejects highest magnitude delta (max 1/3 of samples rejected)
- [Stratum](#) - Rejects all but the lowest-stratum samples present. Be very careful with this filter. Example 1: If your selection of samples includes one sample from a stratum 1 server, and ten more from a mix of stratum 2 and stratum 3 servers, then all but the single stratum

1 sample would be rejected. Example 2: If your lowest-stratum samples are a mix of stratum 2 servers, then all the stratum 2 samples would survive, but all your samples from strata 3 and up would be rejected. It is probably better to use the "NTP Client Max Stratum" value introduced in version 5.2.b.20110224 to control the highest stratum acceptable for NTP sources. The Stratum filter introduced here applies to all sources that report a stratum, including NTP, DT2, and PTP (the PTP "stepsAway" value is used to mimic NTP strata, as documented in the release notes for 5.2.b.20150516). Samples that do not report a stratum are not eliminated by this filter.

Value Name: Wait for Network Startup

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: Present in version 5.2.b.20151102 or later. If set to True, Domain Time will wait up to 30 seconds after boot for an IPv4 address to be assigned to the machine. At boot time, some network adapter drivers report ready before assigning IP addresses to an interface, even if the IPs are pre-configured as fixed addresses. DHCP-obtained addresses can take several seconds longer. The wait period helps ensure that Domain Time's initial enumeration of adapters and IPs is correct before protocol listeners or timechecks are started.

Change this value only if instructed by Technical Support

Time Sources

The Domain Time II Client Time Sources settings are located in this key:

```
HKEY_LOCAL_MACHINE
  Software
    Greyware
      Domain Time Client
        Time Sources
          Broadcast
            PTPv2 (IEEE 1588)
```

The values listed in the **Time Sources** registry key represent the time sources Domain Time uses to obtain the time. They correspond to settings on the [Obtain the Time](#) property page of the Control Panel applet or are otherwise automatically set. You should not change items in this section unless instructed by Tech Support or you are familiar with the specific function.

PTPv2 (IEEE 1588) key

Value Name: Current Master

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Current Offset (signed 64-bit)

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Current Offset Enabled

Value Type: REG_SZ

Default Data: False

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160922. When false, Domain Time will not update the current offset value in the registry (*Current Offset (signed 64-bit)* described above) or fire the offset-changed event. Note, this reverses the behavior introduced in version 5.2.b.20160415. To regain this behavior, set *Current Offset Enabled* to True, then trigger a sync or issue dtcheck -reload. See the SDK.DOC file included with the Software Development Kit (SDK) for details.

Value Name: Current PortState

Value Type: REG_SZ

Default Data: N/A

Notes: Introduced as of version 5.2.b.20160415. Read-only key for use with the Software Development Kit (SDK), purchased separately.

Value Name: Duplicate Node Detection Enabled

Value Type: REG_SZ

Default Data: True

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160415. Controls whether Domain Time will detect and prevent duplicate Clock Identities on the network.

Value Name: TAI-UTC Offset Discovered (seconds)

Value Type: REG_DWORD

Default Data: N/A

Options: N/A

Notes: Contains the current TAI-UTC offset (number of UTC leap seconds) discovered from the upstream Master or by importing a leapfile using the DTCheck utility. If this machine is acting as a stand-alone PTP Master, you may manually enter the number of leap seconds (create the key if it doesn't exist). The service must be stopped/restarted for changes to this value to take effect.

Value Name: TAI-UTC Offset Locked

Value Type: REG_SZ

Default Data: False

Options: *True or False*

Notes: Introduced as of version 5.2.b.20160922. If changed to True, DT will not adjust its discovered TAI-UTC offset to match a new master advertising a different offset. The service must be stopped/restarted for this change to take effect. You should use this setting only if you have a broken PTP master advertising an incorrect TAI-UTC offset.

Domain Time II Client for Linux (DTLinux)

Version 5.2

Domain Time II Client for Linux (DTLinux) is an easily-configured multi-protocol (PTP, NTP, and DT2) daemon that obtains time from various time sources (such as GPS/GNSS clocks, PTP Masters, Internet time servers, etc.) and matches the system clock to them with extreme accuracy and precision. Keeps full transaction logging and drift records that are easily audited by [Domain Time II Audit Server](#) or exported to other monitoring utilities.

As powerful as DTLinux is on its own, you can take full advantage of its capabilities by connecting to it from [Domain Time II Manager](#) which provides remote configuration, graphical drift displays, the ability to push out configuration templates to multiple DTLinux machines, and more. See the [Managing DTLinux Remotely](#) page for details.

[Installation Instructions](#)

[System Requirements](#)

Configuring and Controlling DTLinux

Once [installed](#), the DTLinux daemon will start automatically as a system service when the machine boots.

You may use systemctl to control the service, i.e.

<code>systemctl start dtlinux.service</code>	Start the service
<code>systemctl stop dtlinux.service</code>	Stop the service
<code>systemctl restart dtlinux.service</code>	Stop and restart the service
<code>systemctl reload dtlinux.service</code>	Reload the service configuration
<code>systemctl status dtlinux.service</code>	Show the service status

The `dtlinux.conf` file

All settings used by DTLinux are read from the `dtlinux.conf` configuration file located in the `/etc/opt/domtime/` folder. The [Configuration](#) page discusses the contents of this file in more detail.

You may edit the `dtlinux.conf` file directly on the machine itself, or you may use Domain Time II Manager (as of v5.2.b.20201116) to connect to the machine and configure it remotely. See the [Managing DTLinux Remotely](#) page for details.

Notes:

- If you manually make changes to the `/etc/opt/domtime/dtlinux.conf` file on the machine, you must tell dtlinux to reload the settings. You can do this using `systemctl reload dtlinux.service` or by sending a HUP signal, i.e. `kill -SIGHUP <pid>` If you configure DTLinux remotely using Manager, there is no need to reload or HUP the service after making changes.
- If you want to duplicate the settings from one DTLinux machine to another, you have two options:
 - Manually copy the `dtlinux.conf` and `dtlinux.keys` files.
 - Copy the `/etc/opt/domtime/dtlinux.conf` file from machine A to replace machine B's `/etc/opt/domtime/dtlinux.conf` file.
 - Copy the `/etc/opt/domtime/dtlinux.keys` file from machine A to replace machine B's `/etc/opt/domtime/dtlinux.keys` file.
 - Use `systemctl reload dtlinux.service` or send a HUP signal on machine B to load the new configuration.

Important: If you have specified a network adapter name in machine A's `dtlinux.conf` file

(network:adapterName in the [Network Settings](#) section of dtlinux.conf), you must either manually set this value on machine B to match the correct adapter name or leave the setting blank to let DTLinux auto-discover an interface. Use **dtcheck -adapters** on machine B to see the available interfaces.

- Use Manager to create a template and apply it one or more other machines (see the Templates section of [Managing DTLinux Remotely](#) for details):
 - Connect to machine A and use Templates -> Create New Template from the Control Panel menu to create and save a new template containing the settings you want to duplicate. We recommend you edit the template to be sure the network:adapterName (see [Network Settings](#) section of dtlinux.conf) is blank to let DTLinux auto-discover an interface when the template is applied to a new system.
 - Use Manager's [Reset Configuration](#) command to push out the settings to one or more other dtlinux machines.
 - Ensure the keyring on the Manager machine matches the keyring in use on DTLinux This is configured in Domain Time II Server on the Manager machine (see [Symmetric Keys](#)). Then, use Manager's [Reset Keyring](#) command to push out the keyring to one or more other dtlinux machines.

The dtlinux program

Although the **dtlinux** executable runs primarily as background system service, there are several commands you may issue from the command-line (You may execute it from any location; the program is mapped to your path):

dtlinux -help	Show the list of available commands
dtlinux -version	Display the software version
dtlinux -revision	Display the revision number
dtlinux -registration=xxxxxxxx-yyyy	Apply registration code
dtlinux -prepClone	Prepare system for cloning
dtlinux -resetSerial	Reset the serial number
dtlinux -resetClockId	Reset the PTP identity
dtlinux -problemReport	Create ClientSettings.zip in current folder Submit to tech support when opening a ticket
dtlinux --add-missing	Compares the dtlinux.conf file against the default conf and adds any missing or new settings

You may also view the man page for dtlinux by entering **man dtlinux** at the command-line.

The dtcheck utility

dtcheck is a handy multi-function utility that can act on the local machine and also remotely against other Domain Time components on the network. It can do useful things like trigger a synchronization, display synchronization statistics, show available network adapters, display active PTP Masters on the network, convert drift logs to either text or CSV-formatted files, backup/restore the configuration, and more. Enter this command to see the full list of options and syntax:

dtcheck -help

You may also view the man page for dtcheck by entering **man dtcheck** at the command-line.

Best Practices

Here are some suggestions for getting the best performance and security out of DTLinux:

- Make a backup copy of your dtlinux.config file before making changes so that you can revert if necessary. You may also select the *File -> Create Backup of Configuration File* option from the Control Panel menu when connected remotely from

Manager. You may also use the **dtcheck -backup** command from the command line.

- Make minimal changes to the configuration settings. The defaults are optimized for most applications, particularly the PTP settings. Only change settings that you completely understand.
- Configure dtlinux to use at least one DT2 or NTP timesource (three is even better), even if you will be synchronizing using PTP. PTP takes a significant amount of time to correct large time variances, particularly those that occur just after boot. Having active non-PTP timesources will bring the system into rough synchronization quickly at startup so that PTP can synchronize at high accuracy sooner. It also provides a robust fallback in case PTP sync is lost at some point. See the [NTP and DT2 Time Sources](#) section of the dtlinux.conf file.
- Use a hardware-timestamping capable NIC if possible. DTLinux will automatically take advantage of the additional accuracy available using this feature. If hardware-timestamping isn't available, less-accurate software-timestamping will be used.
- If using PTP, set the loop:checkInterval to 60 seconds. This provides sufficient time to collect enough PTP samples for statistical analysis and filtering. Values smaller than 60 may result in not having enough valid samples to synchronize. See the [Loop Variables](#) section of the dtlinux.conf file.
- Consider restricting remote configuration of DTLinux to individual machines or specific network subnets. By default, DTLinux will accept connections from Domain Time II Manager located on RFC 1918 private networks. You can restrict this (or deny it entirely) using the [Domain Time II Security](#) section of the dtlinux.conf file.
- The simplest way to obtain synchronization performance information from DTLinux is by using [Domain Time Audit Server](#). This allows you to collect synchronization information to a central repository, raise alerts, and easily meet regulatory requirements. However, you can output synchronization records directly from DTLinux using NTP-style loopstats and peerstats, or PTP data may be kept in a CSV file. These options are enabled in the **Loop Variables** section of the dtlinux.conf file and the output files are kept in the folders specified in the **Logs and Folders** section of dtlinux.conf.

You may also use the **dtcheck** program to convert Domain Time binary synchronization files (*.dt) to either text or CSV files. Issue the **dtcheck -help** command at the console for the correct syntax.



Please read the [README.txt](#) file before installing. (Note the latest version of this file will be in the distribution file you download)

NOTES:

- Disable any other time synchronization software to prevent conflicts, including chronyd, ntpd, or any PTP daemons.
- Check your routers and firewalls to be sure the ports for the time protocols you'll be using are open.

IEEE 1588-2008 or 1588-2019 (PTP) require inbound and outbound access to ports 319/udp and 320/udp. All packets sent will have a source port of either 319 or 320, and a target port of 319 or 320. Inbound packets will have a source port of 319 or 320, and a target port of 319 or 320. So, essentially, for PTP, you need to open 319/udp and 320/udp in both directions.

Other time protocols have a fixed target port, but requests are sent from ephemeral ports. Replies will have a source port matching the target port, but a target port matching the ephemeral port used for the request. You need to allow outgoing to ports 123/udp, 9909/udp, and 9909/tcp. You need to allow incoming replies. If this is a problem for your fire- wall, you may change sourcePortUDP above to any unused port on your system. If network:sourcePortUDP is zero, the operating system will choose an ephemeral port. If it is non-zero, it indicates the start of a range of five ports beginning with the port number you specify. For example, if you use network:sourcePortUDP=333, then the program will attempt to bind to port 333. If 333 is already in use, it will try 334, and so on, for up to five ports. If none of the ports in the range you specify are available, then the bind (and transaction) will fail. The range is required because multiple threads may be trying to sending unicast requests at the same time. Choose a starting port number where you are sure that the range (port through port + 4) is not being used by other protocols.

NOTE: network:sourcePortUDP only applies to UDP-based protocols. TCP will always use an ephemeral source port.

If you are using other Domain Time II products, such as Domain Time II Manager, Audit Server, or any of the command-line tools, you must open ports 9909/udp and 9909/tcp for incoming packets. Requests will come via multicast or unicast directed toward 9909/udp or 9909/tcp. Replies will be sent to the source port of the request.

If you want to use Manager's ssh feature, you will need configure ssh on your OS and allow incoming access from port 22/udp.

- If you will be installing DTLinux on a virtual machine, see [this article](#) from our knowledgebase for more information on proper use with virtualization systems.
- If you use cloned OS images to install machines, please read [this article](#) from our knowledgebase about configuring Domain Time properly on your clone image.

Distribution Types

There are three flavors of packaging available:

- **TGZ:** The most flexible option. Distro-independent, with included scripts to install, upgrade, and remove.
- **DEB:** For Debian-derived distros like Ubuntu.
- **RPM:** For distros using the RPM package manager, like Red Hat.

Caution about Graphical Installers:

Some distros offer a graphical RPM or DEB package installers. We recommend using the command line version of **rpm** or **apt** instead, because each graphical overlay behaves differently, and may not give complete installation, upgrade, or remove options. For example, some graphical installers "upgrade" by doing a remove followed by a re-install. This causes a loss of settings and registration info. The command line options for RPM (**rpm -U**) and DEB (**apt-get install**) will upgrade in place, preserving settings.

Installation

- To install DTLinux:
 - **Download** the latest version of the distribution file you prefer. If using TGZ, extract the files into a blank folder.
 - **cd** to the folder containing your installation file(s).
 - Run the installation command as root (or sudo):
 - **TGZ:** `./install.sh`
 - **DEB:** `apt-get install ./[filename]`
 - **RPM:** `rpm -U ./[filename]`
 - Configure the software either by:
 - Editing the `/etc/opt/domtime/dtlinux.conf` file locally from within the OS.
 - or
 - Editing the `dtlinux.conf` file remotely from Domain Time Manager. See the [Managing DTLinux Remotely](#) page.
 - Test your installation
 - Ensure the service is started correctly by issuing the `systemctl status dtlinux.service` command as root (or sudo).
 - Examine the `/var/log/domtime/dtlinux.log` file to see if the service is synchronizing correctly.
 - Use Domain Time II Manager to connect to DTLinux remotely, if desired. See the [Managing DTLinux Remotely](#) page.
 - Make a backup of your settings

After making config changes and/or applying your registration key, run `sudo dtcheck -backup` from the Linux command-line. This preserves your settings and license in case you mistakenly remove instead of upgrading. You may restore backed-up settings by running `sudo dtcheck -restore`. Use `dtcheck -help` for a more detailed explanation.

Upgrade

IMPORTANT: The latest instructions for updating DTLinx can be found in the [UpdatingDTLinux.html](#) file located in the DTLinux distribution files or in the `/opt/domtime` folder of your DTLinux system. Please review this file before proceeding.

- To upgrade DTLinux:

If the machine has access to the web (via port 80 TCP), dtlinux can download the update and upgrade for you, regardless of which installation package you used. (As an alternative, you may configure your Domain Time Manager as an alternate upstream data source so that upgrades do not require Internet access. See the [UpdatingDTLinux.html](#) file located in the DTLinux distribution files or in the `/opt/domtime` folder of your DTLinux system for instructions).

To see if there's a newer version available:

```
dtcheck -update
```

If so, run either of the following commands as root or sudo:

```
dtlinux-update
```

or

```
/opt/dontime/update.sh
```

To upgrade if the machine doesn't have Internet access, or if you prefer to download the software yourself:

- [Download](#) the latest version of the distribution file you prefer. If using TGZ, extract the files into a blank folder.
- `cd` to the folder containing your installation file(s).
- Run the upgrade command as root (or sudo):
 - **TGZ:** `./install.sh`
 - **DEB:** `apt-get install ./[filename]`
 - **RPM:** `rpm -U ./[filename]`

Your existing settings and registration info will be preserved during the upgrade.

Upgrading remotely using Domain Time Manager

As of version 5.2.b.2021025, you may push upgrades from Domain Time Manager, allowing you to easily update multiple machines at once. Please see the Upgrade instructions on the [Managing DTLinux Remotely](#) page for the requirements.

Removal

- Run the following command as root (or sudo) regardless of which package installation method you used:

```
/opt/dontime/remove.sh
```

Registration

This program is delivered as an evaluation version. You may use it for up to 30 days for testing. If you want to continue using it after that, you must obtain a registration code from your vendor. The registration code will look something like this: 692310601-94007843. The number of digits may vary.

To register your program, run the following command as root or sudo:

```
dtlinux -registration=692310601-94007843 *
```

* Use the real code provided by your vendor, not the example code shown above.

NOTE: You can also send the registration code from a Windows machine running Domain Time, using either the DTCheck command-line utility or Domain Time II Manager. If you use Manager, you may register multiple machines at the same time.

- From an elevated Windows command-prompt (requires version 5.2.b.20210103 or later):


```
dtcheck [name or ip address of the remote Linux machine] -registration=692310601-94007843 *
```

* Use the real code provided by your vendor, not the example code shown above.

or

- Using Domain Time II Manager (requires version 5.2.b.20210103 or later):

Highlight the machine(s) you want to register (this can be done from any Node List in Manager, i.e Domains & Workgroups, DT Nodes, etc.), right-click and choose **Register Eval Versions** from the context menu. Enter your registration code. Manager will apply the code to each selected system.



DTLinux - Register evaluation versions [\[Click for larger size\]](#)



DTLinux configuration is simple and straightforward. All configuration (with the exception of configuring the [dtlinux.keys](#) file for symmetric authentication) is done by editing the `dtlinux.conf` file. See a sample here: [dtlinux.conf.sample.txt](#).

Both the `dtlinux.conf` and `dtlinux.keys` files are heavily commented and are the primary documentation for DTLinux. You should always keep a copy of the original distribution files available for reference in case the comments in your running copies are inadvertently removed during editing. This online documentation page merely highlights a few of the topics for additional discussion.

The `dtlinux.conf` file

The file is divided into functional sections:

- File Locations
- File Format
- Other Time Software
- [NTP and DT2 Time Sources](#)
- [Loop Variables](#)
- Logs and Folders
- Network Settings
- [PTP Settings](#)
- [PTP 1588-2019 \(v2.1\) Security](#)
- [Domain Time II Real-Time Alerts](#)
- [Domain Time II Security](#)
- Firewall Advice
- Clock Stepping and Slewing
- Advanced Sample Filtering
- [Cloning](#)
- [License](#)

■ **NTP and DT2 Time Sources**

This section covers how to configure DTLinux to obtain time from NTP and/or DT2 time sources. You can manually specify the sources in the `dtlinux.conf` file and DTLinux can also obtain a list of sources from DHCP options.

Selecting the correct time sources are critical for accurate timing. The Internet time sources specified in the default `.conf` file are intended as examples only. Choose servers that are optimal for your environment. Stable time sources on a local subnet are best.

See the [Planning for effective time distribution](#) for help making the right choice.

■ **Loop Variables**

You can set the time check intervals using the parameters in the section. You also control whether to keep ntp-style loopstats and peerstats files.

A **loop:checkInterval** of 60 is recommended if you are using PTP to allow PTP time to collect enough valid samples to analyze statistically for best performance. Otherwise, if using NTP or DT2, set the value low enough to achieve the accuracy you require. Setting the value too low just increases overhead and network traffic.

Also, set a reasonable **loop:errorInterval**. The value should normally be 30 seconds or less. This affects the period between DTLinux detecting a loss of sync with time sources and when it retries a connection. A relatively short error interval is desirable to restore sync quickly when sources become available again.

The **loop:checkAll** setting determines whether all the configured NTP and DT2 time sources are included and analyzed in each time check or if the list is used for fallback, where the first server is used until it fails, at which point the next machine in the list is tried. You may set the log level to Trace (**log:logLevel** = Trace) if you want to see the details on which machines are used in each time synchronization.

■ **PTP Settings**

Use this section to enable/disable PTP and set its basic parameters.

See the [PTP Profiles](#) section of the main PTP page for information on which PTP Profiles DTLinux supports.

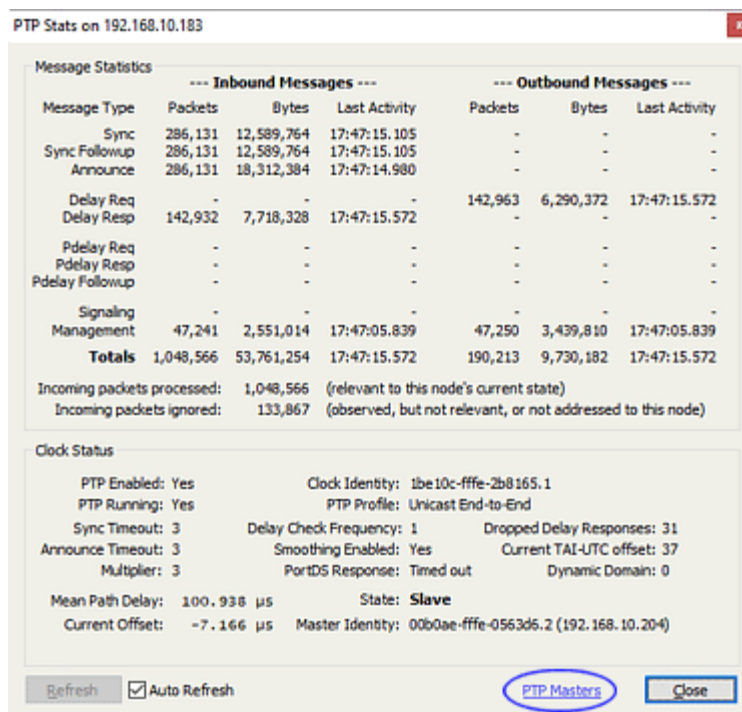
Be sure to use DTLinux's ability to view available PTP Masters when troubleshooting synchronization issues:

- **From the Terminal:**

`dtcheck -ptpmasters`

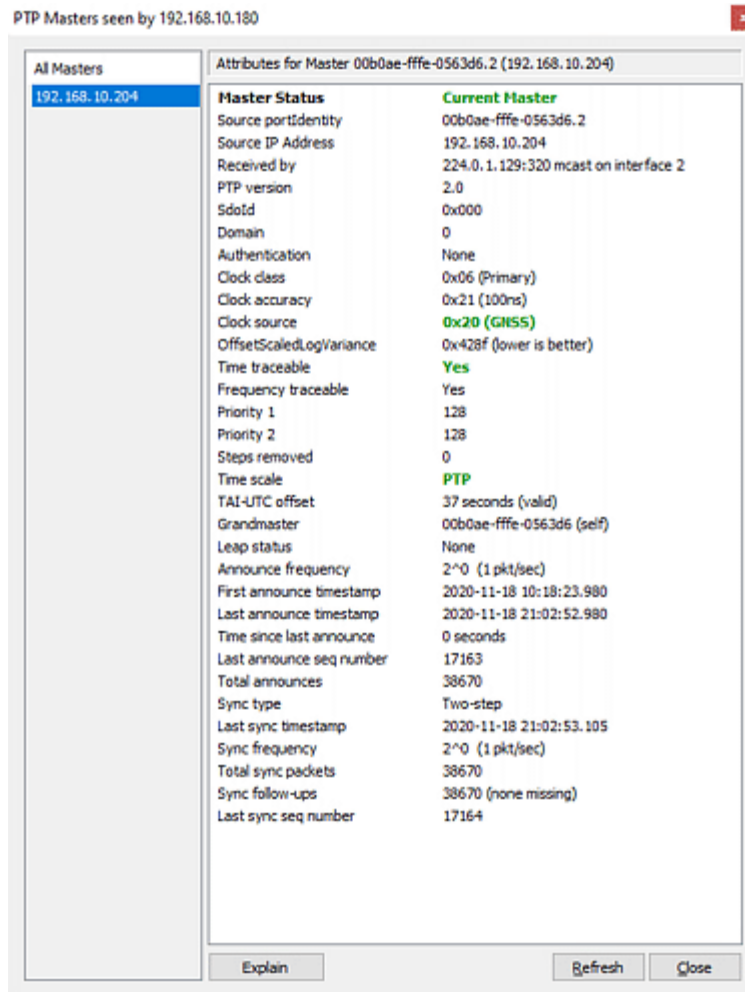
- **From Domain Time Manager:**

Choose *Graphs & Statistics* -> *Open PTP Statistics & Masters* from the Manager menu to display the PTP Statistics page.



Manager - DTLinux PTP Statistics [Click for larger size]

Then click the PTP Masters link on that page to see the PTP Masters list.



Manager - DTLinux PTP Masters [Click for larger size]

■ PTP 1588-2019 (v2.1) Security

DTLinux supports PTP v2.1 (IEEE 1588-2019) which features enhanced security features. You enable these features by setting the options in this section.

There are six true/false settings related to PTP v2.1 security. The default for all is false.

```
ptpSecurity: enabled = false ; enable v2.1 security?
ptpSecurity: preferSignedAnnounces = false ; prefer signed announces?
ptpSecurity: requireSignedAnnounces = false ; require signed announces?
ptpSecurity: requireSignedSyncs = false ; require signed syncs?
ptpSecurity: requireSignedDelayResp = false ; require signed delay responses?
ptpSecurity: signOutgoingDelay = false ; sign outgoing delay requests?
```

If `ptpSecurity: enabled` is `false`, the remaining settings are ignored. No PTP v2.1 security TLVs are processed.

If `ptpSecurity: enabled` is `true`, then the following options obtain:

- If `ptpSecurity: preferSignedAnnounces` is `true`, then Best-Master-Clock (BMC) algorithm is altered to give priority to masters that sign their Announces. This configuration allows a mix of v2.0 and v2.1, so that the normal BMC applies if no master signs Announces.
- If `ptpSecurity: requireSignedAnnounces` is `true`, then *only* masters that provide signed announces will be considered.
- If `ptpSecurity: requireSignedSyncs` is `true`, then *if* the selected master is v2.1 and sending signed Announces, signed Syncs/Follow-Ups are also required. If your master signs Announces but not Syncs, DTLinux will

not be able to follow it.

- If **ptpSecurity: requireSignedDelayResp** is **true**, then *if* the selected master is v2.1 and sending signed Announces, signed delay responses are also required. If your master signs Announces but not delay responses, then DTLinux will not be able to calculate the meanPathDelay, and your synchronization quality will suffer.
- If **ptpSecurity: signOutgoingDelay** is **true**, and the selected master is v2.1 and sending signed Announces, DTLinux will sign delay requests according to the key numbers you have selected. Some v2.1 masters send unsigned delay responses if the delay request is unsigned. Some always sign delay responses. Others will ignore unsigned delay requests. Check with your appliance manufacturer's documentation to decide whether or not to sign delay requests and require signed delay responses.

You should normally leave **ptpSecurity: requireSignedDelayResp** and **ptpSecurity: signOutgoingDelay** set to false, unless required by your grandmaster appliance. There is no real security benefit to signing delay requests or requiring signed delay responses.

The **/etc/opt/domtime/dtlinux.keys** file contains the symmetric keys and also the PTP v2.1 key numbers to use for signing packets.

MD5 and SHA1 keys are used with NTP and DT2. PTP v2.1 keys must be either SHA256 or SHA512. In order to verify signed incoming packets, you must have SHA256 or SHA512 keys corresponding to whatever the master sends for Announces, Syncs, Follow-Ups, and delay responses.

In order to sign a delay request, you must have an SHA256 or SHA512 key known to the master, and you must select the key number to use. For example, if SHA256 keys 24 and 26 exist in your keyring, you set the values in the dtlinux.keys file as follows:

```
ptpDelayReq    24  # End-to-End delay request key
ptpPDelayReq   26  # Peer-to-Peer delay request key
```

You should not use SHA512 keys. DTLinux supports them in the keyring and will use them if the master sends an SHA512 key, but most masters only support SHA256. In particular, the only defined algorithm for PTP v2.1 is HMAC-SHA256-128. This will change when NTS (RFC 8915) is adapted for use with PTP and appliance manufacturers accommodate it. In the meantime, use SHA256.

The **/etc/opt/domtime/dtlinux.keys** file also contains a Security Parameter Pointer (SPP) value, range 0-255. For example:

```
ptpSPP 0 # Security Parameter Pointer
```

If the ptpSPP value is zero, then DTLinux will accept any SPP value sent by the master. When sending delay requests, DTLinux will use the SPP of the master.

If the ptpSPP value is non-zero, it must match the SPP value sent by the master. If the values don't match, then neither incoming packets nor outgoing delay requests will verify correctly.

■ Domain Time II Real-Time Alerts

If you are using Domain Time II Audit Server, we suggest you enable Real-Time Alerts in this section, even if you haven't yet configured any Real-Time Alerts in Audit Server. This will cause the DTLinux machine to display in the Real-Time Alerts page of Manager, giving you up-to-date information on synchronization status and accuracy.

■ Domain Time II Security

Settings in this section control remote access to DTLinux from [Domain Time II Manager](#).

The **dt2Security:allow** entries specify IP addresses or CIDR masks that control which IPs are allowed to connect to DTLinux. The default entries include the RFC 1918 private network blocks. You should remove any entries that don't correspond to the networks where you have Domain Time II Manager or Audit Server installed. To prevent all remote access, change the first dt2Security:allow entry to 127.0.0.1 and comment out (or delete) the remainder.

Assuming you have granted IP access to Manager, the **dt2Security:managerReadOnly** entry controls whether or not Manager is allowed to make any changes. If you set this value to true, then Manager will still be able to view settings, and Audit Server will still be able to audit the machine, but Manager will be prevented from making any changes to the configuration. If you set this value to true, then you must manually edit the dtlinux.conf file on the remote machine to change it back to false. By design, there is no remote method available for changing this value from true to false.

The **dt2Security:managerRestart** entry controls whether or not Manager is allowed to restart the DTLinux service remotely. It has no effect if dt2Security:managerReadOnly is true.

The **dt2Security:managerUpgrade** entry controls whether or not Manager is allowed to upgrade the DTLinux service and associated file remotely. It has no effect if dt2Security:managerReadOnly is true.

Please read /opt/domtime/UpdatingDTLinux.html for information on using the **dtcheck** command-line utility to limit key functions to only specific Managers.

■ Cloning

If you use cloned OS images to install machines, please read [this article](#) from our knowledgebase about configuring Domain Time properly on your clone image.

■ License: Commercial Proprietary (registration required)

This section describes the evaluation period and how to register the software. The section will be removed when the software is registered.

The dtlinux.keys file

This file contains the authentication keys used for the DT2, NTP, and/or PTP v2.1 protocols. It's also referred to as your keyring. It's located in the `/etc/opt/domtime/` folder.

The keyring may contain a combination of trusted and untrusted keys. A trusted key means the key is available to be selected by the component, but trusted keys for DT2 and NTP are not active until their key number is specified when configuring a DT2 or NTP time source in the [time sources](#) list of the dtlinux.conf file (i.e. **timesource = 192.168.1.3 protocol NTP key 5**). Trusted keys for PTP v2.1 aren't active unless [PTP Security](#) has been enabled. Untrusted keys are ignored.

Here are values from a sample keyring, with MD5 keys available for use by DT2 or NTP, and SHA256 keys available for PTP v2.1:

Key #	Type	Secret
1	MD5	DomainTimeII
2	MD5	TTnts200
3	SHA256	bf14d67e2ddc8e6683ef574961ff698f61cdd11e9d9c167272e61df0844f4a71
4	SHA256	48d38f75e6d91d2ae5c0f72b788187440e5f5000d4618dbe7b0515073b338211
5	MD5	greyware

The **Trustedkey** line in the file specifies which keys in the keyring are trusted, i.e.:

```
Trustedkey 1 2 3 4 9909
```

The file also contains additional settings required for PTP v2.1 authentication.

ptpSPP sets the Security Parameter Pointer (SPP). PTP v2.1 requires that Masters and Slaves use the same SPP value to be able to authenticate. The SPP stored in the keyring may either be zero (which acts like a wildcard) or must match what the grandmaster sends. If there is a potential for your Slaves to discover more than one Master (such as with a fallback server), we recommend you use the wildcard setting (0) to avoid synchronization failure if each server has a different SPP.

These entries specify the key number of the secret that Masters use for signing outgoing packet types. They are included

here for compatibility when importing the .keys file into Domain Time Server. These parameters are ignored by Domain Time Client and DTLinux :

ptpAnnounce [key #]
ptpSync [key #]
ptpDelayResp [key #]
ptpPDelayResp [key #]

These entries specify the key number of the secret used for signing packet types sent by the Slave:

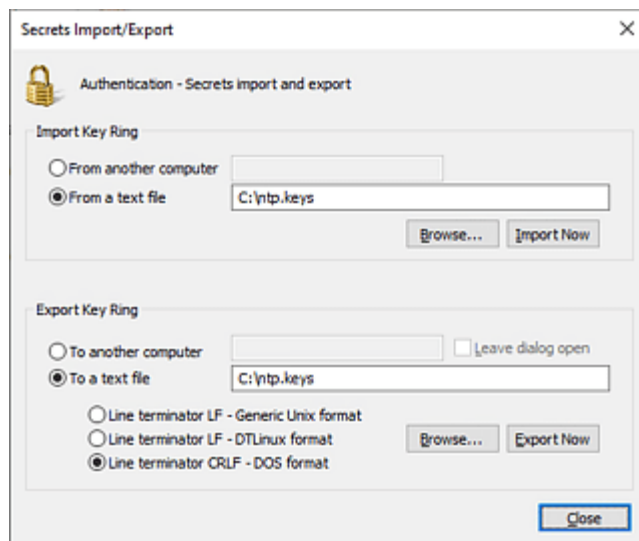
ptpDelayReq [key #]
ptpPDelayReq [key #]

Sharing the keyring file.

For symmetric authentication to work, the keyring must be shared among all devices that wish to use it. The dtlinux.keys file uses a format compatible with most time daemons (i.e. ntpd's ntp.keys, chrony's chrony.keys, etc.). You can usually simply copy the `/etc/opt/dontime/dtlinux.keys` file to your target system (rename it if necessary).

You can also copy the `/etc/opt/dontime/dtlinux.keys` file from one DTLinux machine to another.

You may also share the dtlinux.keys file with Domain Time Servers and Clients on Windows (and vice versa). Use the [Import/Export](#) link on the Symmetric Keys property page of the Server or Client's applet to import or export the .keys file.



Secrets Import/Export Dialog [\[Click for larger size\]](#)

If you are using Domain Time II Manager, you can use the [Reset Keyring](#) function to push out the keyring to all of your Windows Servers and Clients and DTLinux machines at once. The [Reset Keyring](#) function uses the keyring of the Domain Time Server on which Manager is installed. So, to easily share a DTLinux machine's keyring among all of your other Domain Time systems, you'd [import](#) the keyring file into Manager's Domain Time Server and then select the machines you want to update and use the [Reset Keyring](#) command from the right-click context menu.



DTLinux is designed to easily integrate with other Domain Time products, such as Domain Time II [Manager](#) and [Audit Server](#). When you connect to DTLinux from Manager, you gain the ability to remotely configure the system, graphically display its performance, create configuration templates to apply to other DTLinux systems, register multiple systems at once, perform remote upgrades and more.

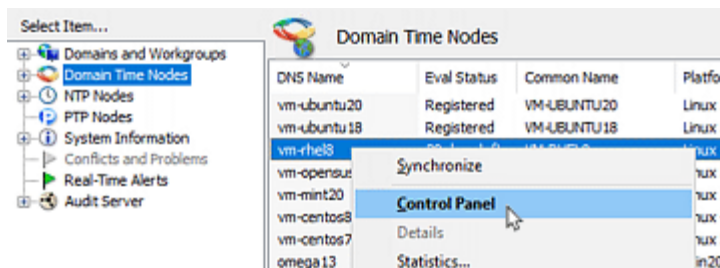
Requirements

In order to connect, you will need to be running Domain Time II Manager v5.2.b.20201116 or later.

Connection is made over ports 9909/udp and 9909/tcp, so these will need to be permitted through your network and on your DTLinux systems. Also, you must connect from an IP address that is permitted in the [Domain Time II Security](#) section of its dtlinux.conf file.

Additional permissions are required to perform remote upgrades and to remotely restart the DTLinux service. These are also configured in the [Domain Time II Security](#) section of its dtlinux.conf file.

Connect to DTLinux




If you have configured Manager's [Network Discovery](#) correctly, DTLinux machines should show up automatically in the **Domain Time Nodes** section of Manager. To connect, simply double-click the machine's name or right-click and choose **Control Panel** from the context menu. Note that you can also connect from any screen in Manager where the machine is displayed, such as the **Domains and Workgroups** section or the **Real-Time Alerts** page.

Manager - Connect from the Domain Time Nodes display [Click for larger size]

Adding machines manually

If your network does not permit broadcast or multicast Network Discovery, you can add DTLinux machines to Manager manually. To do this, right-click the **Domains and Workgroups** icon in the left-hand column and choose **Add Computer** from the context-menu. You may also press the INS key. Enter the DNS name or IP address of the DTLinux system. It will be added to the appropriate section under **Domains and Workgroups**. You can connect to it and audit it from there.



Manager - Manually Add a Machine to Manager [Click for larger size]

Note: DTLinux must already be installed on the remote machine in order to add it to Manager. For security reasons, remote install and removal are not supported for DTLinux. Use your organization's package deployment system to install DTLinux, or install it manually on each machine.

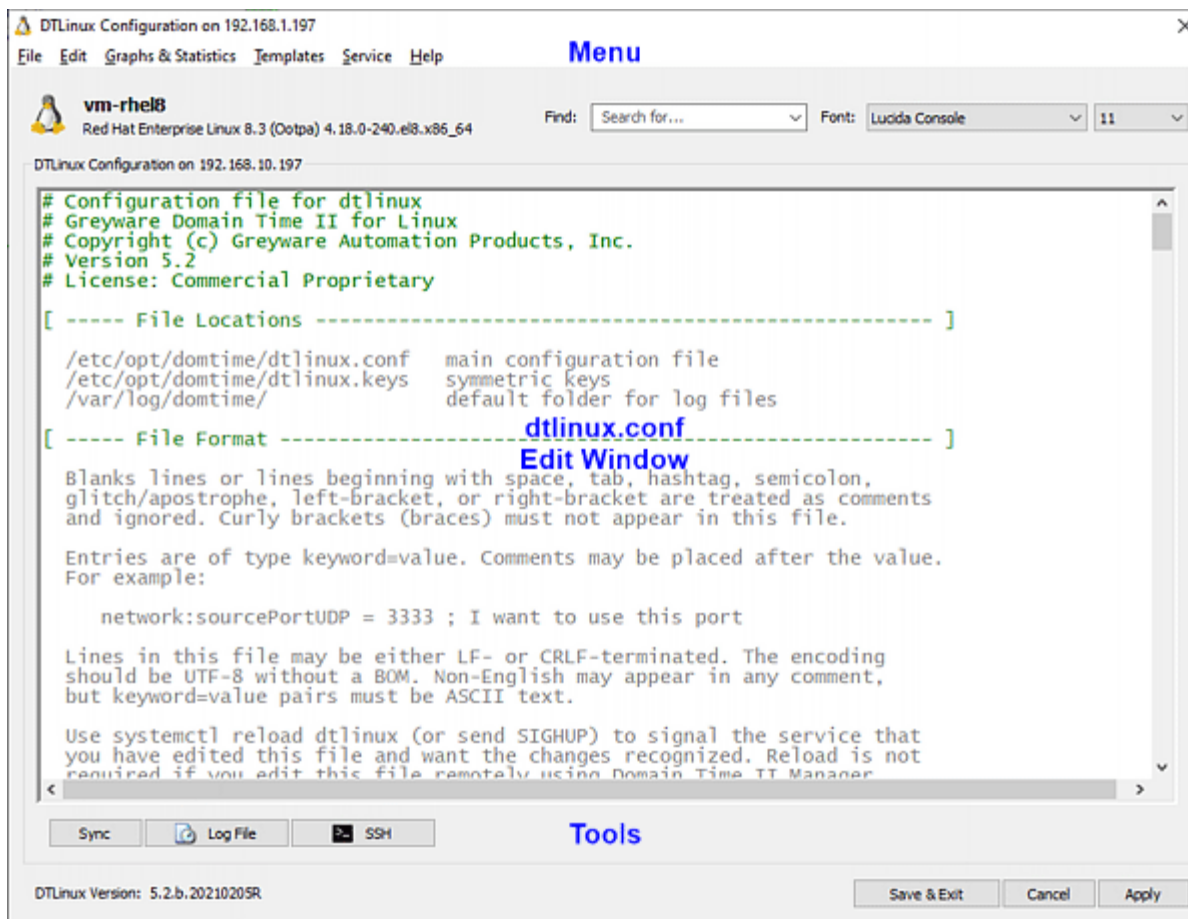
You can also add multiple machines to Manager at once. See Manager's [Batch Add](#) instructions.

The Control Panel

When you connect to DTLinux, you'll see the Control Panel. You can perform most management tasks from this screen.

Editing the dtlinux.conf file

The currently-loaded dtlinux.conf file is displayed in the editing window of the Control Panel. You can make any changes you like to the file. When you click the **Apply** button, the changed file is pushed up to DTLinux and it takes effect immediately. The **Save & Exit** button also pushes the file to DTLinux and exits the Control Panel. The **Close** button exits the Control Panel without making any changes.



Manager - The DTLinux Control Panel [\[Click for larger size\]](#)

Tool Buttons

The three buttons on the bottom-left of the Control Panel provide quick access to common tasks.

- The **Sync** button triggers DTLinux to do an immediate synchronization with its time sources. You'll receive a confirmation of the trigger.
- The **Log File** button displays the log viewer applet showing the active service log (/var/log/domtime/dtlinux.log). This contains DTLinux activity from this machine. Log Levels (Info, Trace, Debug, etc.) for this log are controlled in the [Logs and Folders](#) section of the dtlinux.conf file. You may also use the **File:** dropdown on the log viewer to view cached versions of other DTLinux machines you have viewed previously.
- The **SSH** button will allow you to connect to the machine via ssh. You must have an ssh client installed on your Manager machine. ssh must also be configured on the Linux system.

Menu Commands

The Control Panel menu contains a wide variety of useful features.

File

- **Backup Configuration**
Saves a copy of the dtlinux.conf file to a local file.
- **Restore Configuration**
Restores the backup copy created using the previous command to the editing window. You must then **Apply** the changes if you want the configuration pushed up to the running DTLinux system.
- **Service Log File**
Displays the DTLinux Service Log. This is the same as clicking the **Log File** tool button described above.

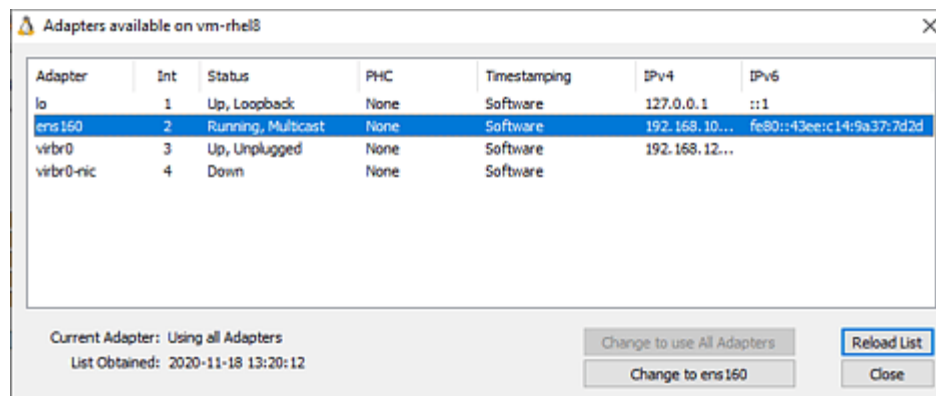
Edit

■ Find

Takes you to the Find search box to quickly locate text in the .conf file.

■ Network Adapters

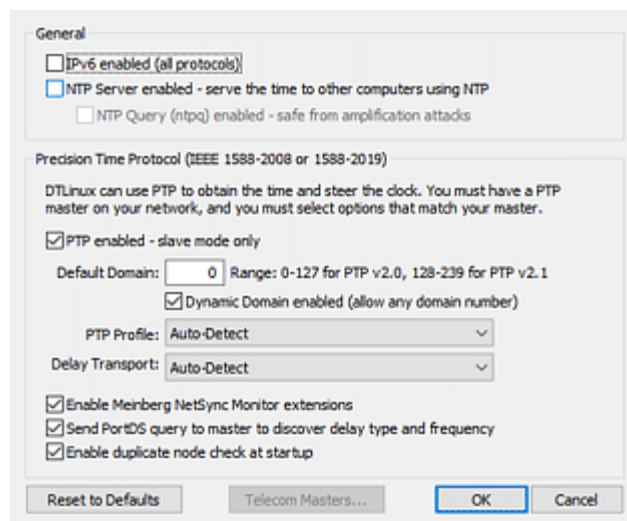
This command displays the list of network adapters on the DTLinux system. You can use this tool to restrict DTLinux to using a single adapter or tell it to use all adapters (the default). This corresponds to the **network:adapterName** entry in the [Network Settings](#) section of the dtlinux.conf file. Use caution when limiting DTLinux to a single adapter. You may lose connectivity from Manager if you select the wrong interface.



Manager - The List of Network Adapters [\[Click for larger size\]](#)

■ Protocol Settings

This dialog allows you to easily configure many of the time and network protocols used by DTLinux. You can enable IPv6, whether DTLinux will act as an NTP server, whether to use PTP to synchronize, set PTP Profile options, and more.



Manager - Protocol Settings dialog [\[Click for larger size\]](#)

■ Time Sources

You may edit the list of time sources DTLinux will contact using this dialog. Checked entries are active.

vm-rhel8 time sources and loop control

Loop Check Interval: Range 15-28800. How often to check the time and update statistics.

Error Retry Interval: Range 15-28800. How long to wait after failing to obtain the time.

☒ Analyze all timesources and choose the best, or average equally good samples (recommended)

Server Name or IP	Family	Protocol	Samples	Delay	Key	Comment
<input checked="" type="checkbox"/> 192.168.10.203	IPv4	NTP	3	256	None	Microsemi ...
<input checked="" type="checkbox"/> tick.greyscale.com	IPv4	NTP	1	n/a	None	Greyscale ...
<input checked="" type="checkbox"/> tock.greyscale.com	IPv4	DT2	1	n/a	None	Greyscale ...

Manager - Time Sources dialog [Click for larger size]

■ Symmetric Keys

Use this dialog to configure your symmetric authentication keys, and also set your PTP v2.1 authentication options. You may also use the **Import** and **Export** functions to create or import .keys files to share symmetric keys between machines.

vm-rhel8 Symmetric Keys

KeyId	Type	Password
<input checked="" type="checkbox"/> 1	MD5	DomainTimeII
<input checked="" type="checkbox"/> 2	MD5	TTnts200
<input checked="" type="checkbox"/> 3	SHA256	bf14d67e2ddc8e6683ef574961ff698f61cdd11e9d9c167272e6...
<input checked="" type="checkbox"/> 4	SHA256	48d38f75e6d91d2ae5c0f72b788187440e5f5000d4618dbe7b0...
<input checked="" type="checkbox"/> 9909	MD5	greyscale

PTP v2.1 Security Settings

☐ Enabled

Security Parameter Pointer (SPP): Range 0-255; use 0 to match any incoming SPP

☒ Select best master by quality ☐ Sign Outgoing Delay Requests

☐ Prefer signed Announces (alternate BMC) End-to-End Delay Request Key

☐ Require signed Announces (alternate BMC) Peer-to-Peer Delay Request Key

☐ Require signed Syncs/Follow-Ups ☐ Require signed Delay Responses

Manager - Symmetric Keys dialog [Click for larger size]

■ Log File Settings

Use this dialog to configure the log level and retention settings for DTLinux Service Log File. You can also set file path locations for the Service Log and Synchronization records files, and enable loopstats, peerstats, or PTP stats files. Settings on this dialog correspond to the [Logs and Folders](#) and [Loop Variables](#) sections of the dtlinux.conf file.

Text Log Level:

☐ Echo text log to syslogd

Log Retention: Range 0-365; use zero to keep only today's log

Log retention applies to the text log file, loopstats, peerstats, and ptpstats. Older files are named <filename>.<yyyymmdd>.log, and kept in the Log File Path folder.

Log File Path: Location for text log files

Drift File Path: Location for binary drift files

☐ NTP-style Loopstats enabled

☐ NTP-style Peerstats enabled

☐ CSV-style PTP stats enabled Caution: can grow very large

Manager - Log Settings dialog [Click for larger size]

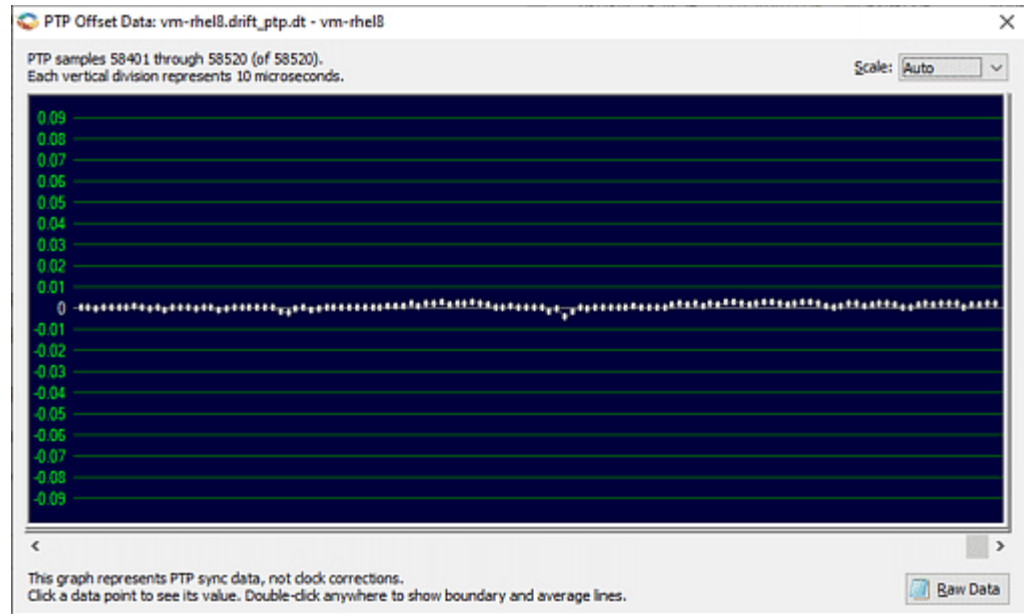
Graphs & Statistics

■ Open Drift Graph

This displays the drift graph showing time corrections made using DT2 or NTP, or summary information for the loop period if using PTP. The data is collected at the rate selected by the **loop:checkInterval** entry in the [Loop Variables](#) section of the dtlinux.conf file.

■ Open PTP Drift Graph

This displays the drift graph showing individual PTP sync data. The data is collected at the rate sync packets are received from the PTP Master. Click a dot to see its data, or you can click the Raw Data button to see detailed statistics and individual data records.



[Manager - A Sample Drift Graph](#) [\[Click for larger size\]](#)

■ Auto-Refresh Drift Graphs

Toggles whether the drift graphs are updated automatically while viewing them.

About Drift Logs

- Click a dot in a drift graph to see its data in the lower-left-hand corner of the graph, or you can click the Raw Data button to see detailed statistics and individual data records.
- Drift graph files are binary and have a .dt extension and are kept on the DTLinux machine in the `/var/logs/dontime/` folder. When you view the graph using Manager's Control Panel, a copy is made in the Manager's `\Program Files\Domain Time II\DTLinux\Drift\` folder.
- Drift files are limited in size and older data scrolls off as new information is acquired. You may, however, use Domain Time II Audit Server to collect logs centrally and maintain them for historical and audit purposes. See the [Synchronization \(Drift\) Logs](#) section of the Audit Server documentation.
- Drift files can be converted into .txt and .csv formats using either Linux or Windows:

■ DTLinux

```
dtcheck [filename].dt -txt
dtcheck [filename].dt -csv
```

■ Windows

```
dt drift - convert [filename].dt (converts to .txt)
dt drift - convert -csv [filename].dt (converts to .csv) You may add the -noheader switch to omit header info.
See the DTRIFT.EXE discussion for full details.
```

In addition, Audit Server has the ability to automatically convert drift files. See the [Conversions](#) section of the Audit Server documentation.

- **Statistics** This command displays computer information, synchronization status, and the current operating statistics for DTLinux.

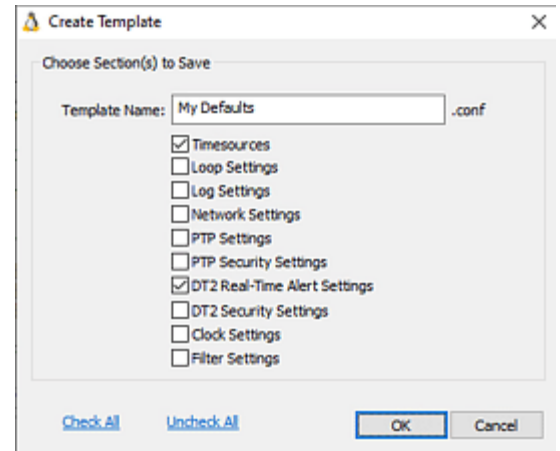
■ PTP Statistics & Masters

This command displays full statistics for the PTP protocol on DTLinux. In addition the **PTP Masters** link on that page will display all PTP Masters visible to the DTLinux machine. This is extremely useful in diagnosing PTP issues.

■ Reset Statistics

This resets all statistics and deletes all information in the drift logs. Use with care.

Templates



Manager - Create a New Template [Click for larger size]

> Templates are snippets of a configuration file (for example, the PTP settings section) that you can save to use again, either on the same DTLinux node or another one. Templates make it easy to update just a portion of a configuration file without searching through the file and manually typing or pasting. Once you have saved a template, you can access it through Manager's *Options* -> [Manage Templates](#) menu, and you may apply the template to one or more machines using Manager's [Reset Configuration](#) function.

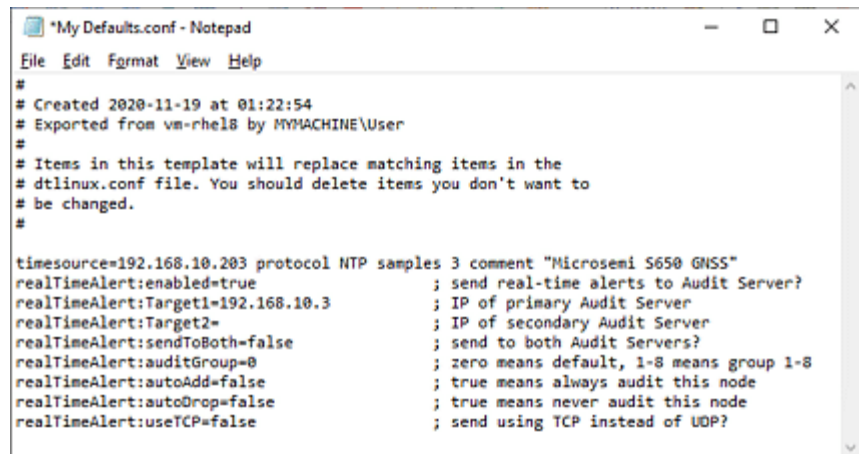
■ Create New Template

This option will allow you to decide which sections of the current dtlinux.conf file to include in your new template. Templates are saved in the `\Program Files\Domain Time II\Templates\Client\` folder on your Manager machine.

Once created, you'll have the option to edit the file.

■ View/Edit Templates

This lets you browse the Client templates folder to choose a template to edit. The selected file will open in Notepad.



Manager - Editing a DTLinux template [Click for larger size]

Service

■ View Service Status

Displays the current DTLinux service status (equivalent to running the `systemctl status dtlinux.service` command at the console).

■ Service Unit File

Shows the Service Unit file used by the Linux systemd init service controller.

■ Persistent Storage File

Displays the contents of the Persistent Storage file, which contains machine-specific information such as the PTP ClockID, clock timings, TAI-UTC offset, serial numbers, etc.

■ Startup Log

Shows the Debug-level log of the most recent startup of the DTLinux service.

■ Update Log

Displays the update history of this system.

■ Reload Service Parameters

Reloads the service parameters (equivalent to running the `systemctl reload dtlinux.service` command at the console).

■ Restart Service

Restarts the DTLinux service (equivalent to running the `systemctl restart dtlinux.service` command at the console). Note, the `dt2Security:managerRestart` permission must be set in the [Domain Time II Security](#) section of the `dtlinux.conf` file to permit this.

Help

■ Create Problem Report

Creates a compressed file named `[hostname]-Settings.zip` containing useful troubleshooting data, logs, and configuration info. Provide this file to Technical Support if you need to open a Trouble Ticket.

Other Manager Commands

In addition to the commands on the Control Panel, there are several important commands available from Manager's various displays. These are found on the context menu when you select a machine from the **Domains and Workgroups**, **Domain Time Nodes**, or **Real-Time Alerts** lists. Pick your machine, and right-click to display the context menu.

■ Synchronize

This command triggers the selected machine to synchronize with its configuration time sources. You'll receive a confirmation of the trigger.

■ Statistics

This command displays computer information, synchronization status, and the current operating statistics for DTLinux.

■ Auditing

Only available if you have Audit Server installed. Lets you choose to audit this machine, and select to which Audit Group it belongs. See the [Alerts and Audit Groups](#) page in the Audit Server documentation.

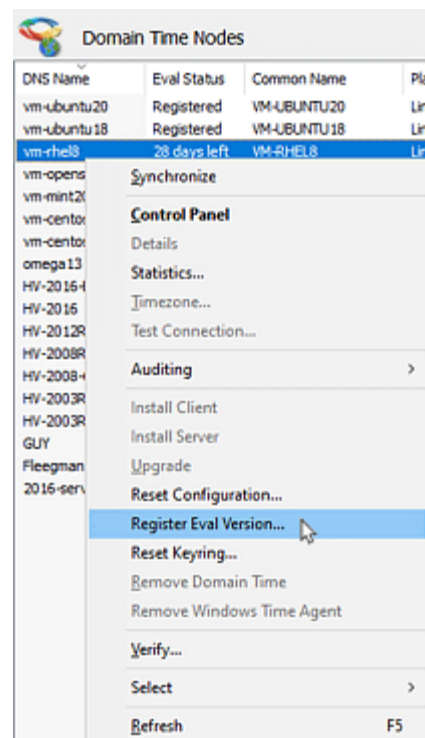
■ Upgrade

As of v5.2.b.20210205, you may use Domain Time Manager to push an upgrade to remote DTLinux machines. This allows you to quickly upgrade one system or many machines at once. The option will be available if

- the version of your Manager is newer than the version installed on the remote system and
- the Domain Time Manager has been configured as an upstream update source (see the [UpdatingDTLinux.html](#) file located in the DTLinux distribution files or in the `/opt/domtime` folder of your DTLinux system) and
- the `dt2Security:managerUpgrade` permission has been granted in the [Domain Time II Security](#) section of the `dtlinux.conf` file
- the `dt2Security:managerReadOnly` setting is set to `False` in the [Domain Time II Security](#) section of the `dtlinux.conf` file

■ Reset Configuration

Allows you to apply templates to one or more machines. See the [Reset Configuration](#) page in the Manager documentation.

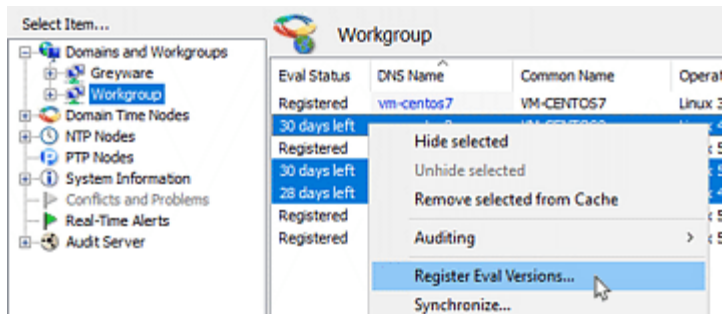


The Context Menu [Click for larger size]

■ Register Eval Version

This very useful feature allows you to apply your registration key to one or machines at once.

Simply select the machines you want to register from any Manager list (**Domains and Workgroups**, **Domain Time Nodes**, or **Real-Time Alerts**), right-click and choose **Register Eval Version** from the context menu.



Manager - Selecting Machines To Register [Click for larger size]

You'll be prompted to enter your registration key. Manager will then proceed to register your selected machines.

■ Reset Keyring

Use this command to push out your master symmetric keyring to one or more machines. Symmetric keys are used for authentication of NTP and DT2 packets. Machines must each have the same keys defined in order to authenticate traffic between them. These user-defined keys are kept in a keyring file (/etc/opt/domtime/dtlinux.keys) on DTLinux. They are kept in a protected registry key on Windows.

Domain Time Manager uses the keyring defined in the Domain Time Server installed on the Manager machine as the master keyring. You should configure Domain Time Server with the keys you want to use among your machines. You can then use **Reset Keyring** to push out the keys to your other machines.

See the [Symmetric Key Authentication](#) section of the Domain Time Server documentation for a full discussion of symmetric authentication and how to configure the keyring on Server.

■ Verify

Verify contacts the selected machine to verify connectivity and updates the Domain Time Manager database with the latest information about the machine. Use this command if you are having difficulty contacting the machine or after registering (to show the change).

Additional commands available on the Real-Time Alerts context menu if you have Audit Server installed:

■ History

Audit Server keeps a history file of Real-Time Alerts received from each reporting machine. This command displays the history file so you can review past alerts.

■ Configure Real-Time Alerts

This is a convenient way to set the Real-Time Alert configuration on one or more machines. Fill out the dialog and Manager will push out the settings to DTLinux. These settings may also be found in the [Domain Time II Real-Time Alerts](#) section of the dtlinux.conf file.

Change Real-Time Alert settings on VM-RHEL8

Audit Server Real-Time Alerts

☒ Enabled

Primary Server: 192.168.10.3

Backup Server: IP address, DNS name, or NetBIOS name

☒ Send to backup only if primary is down ☐ Send reports using TCP

☐ Send to both primary and backup servers ☒ Send reports using UDP (recommended)

NOTE: The values displayed are the current settings on VM-RHEL8.

OK Cancel

Manager - Change Real-Time Alert Settings [\[Click for larger size\]](#)

Domain Time II Manager

Version 5.2

Domain Time II Manager is an advanced and highly-capable network management and monitoring tool. You can remotely install/upgrade/remove Domain Time II components, and see the overall time synchronization across your network from a central Management Workstation.

Manager also comes with a comprehensive suite of other management tools for advanced diagnostics, testing, ongoing monitoring/alerting, and various time-related tasks.

IMPORTANT: If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[System Requirements](#)

[Installation Instructions](#)

[Configuration Instructions](#)

Also, see these pages for information on how to use the Domain Time Manager program:

- [Using the Manager Interface](#)
- [How to manage Domain Time remotely](#)

Installation

Use the [Setup](#) program to install both Domain Time II [Server](#) and the Manager/Management Tools on the machine that you want to use as your Management Workstation.

- Since many functions of Domain Time II Manager depend on accurate time calculations, it should be run on a physical (not virtual) machine, if possible.
- Manager uses Domain Time II [Server](#) for many of its network operations; Manager will not run using Domain Time Client.
- If you will be using [Audit Server](#), you must install both Manager and Domain Time II [Server](#) on the same machine you intend to use for Audit Server.

To launch Manager, click the **Domain Time Manager** icon in the *Start -> All Programs -> Domain Time II* program folder.

You may also launch the Domain Time II Manager program (and many other installed Domain Time II components) by right-clicking on the Domain Time icon in the System Tray to bring up the context menu.

Network Requirements

Verify that your environment meets the minimum requirements for performing remote operations using Domain Time components. In order to be able to install, upgrade, or configure remote machines:

- Your network must be a correctly-configured Windows network, i.e. configured with working name resolution (DNS, WINS, NetBIOS, etc.), correct and functioning Active Directory (if used), working inter-domain trusts, etc.
- Your network must pass both UDP and TCP network traffic sent to destination port 9909. Switches and firewalls must pass this traffic bi-directionally, since traffic will originate either from Manager or the remote machines. Your network must pass this traffic, regardless of what time protocols are used to actually synchronize the time.

Note: As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

- The remote machine must respond to PING requests from the connecting machine.

- The connecting Domain Time program, utility, or service must be run using credentials with sufficient privileges to connect to and write files to the administrative shares on the remote machine using Microsoft Networking (Domain Admin if the target is a domain member, Local Machine Administrator if the target is in a workgroup).
- The Remote Registry Service must be running on the remote systems and its registry keys must be accessible to the connecting program.
- All files from the original distribution for each type of product you want to install (Server, Client, etc.) must be extracted and present on your connecting machine. Setup copies these to the proper locations in the **\Program Files\Domain Time II** folder for you automatically when you install the Management Tools.

Configuration

Manager is configured primarily using the various selections found on its *Options* menu. You have extensive control over how Manager looks and operates.

Note: The *Options* menu items related specifically to discovering machines are discussed on the [Discovery](#) page. The [Using Templates](#) page covers installation Templates in detail.

Configure Reference Time

Before using Domain Time Manager, you should make sure you have decided on what time source(s) to use to act as Reference Time.

Reference Time is configured by selecting *Options -> Network Options -> Reference Time Sources...* from the menu.

Important: Stable reference time is critical to obtaining trustworthy variance data from your network. Choose sources that are known to be reliable and available over low-latency connections.

Reference Time Sources

Select how the service should determine the reference time. The reference time is the time against which other machines are measured.

Reference Clock Type: Specify a list of servers

☒ Analyze time samples and choose the best, or average equally good samples (recommended)

Server Name or IP	Protocol	Auth	Reps	Delay
<input checked="" type="checkbox"/> nist1.symmetricom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp1.symmetricom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp2.symmetricom.com	NTP	None	3	512

Buttons: Add, Delete, Edit, Move Up, Move Down, OK, Cancel

Reference Time Source Selection [\[Click for larger size\]](#)

The **Reference Clock Type:** Use this machine's clock list gives you multiple options for obtaining reference time:

■ Use this machine's clock

The local machine's clock is used as the reference. Use this setting if you have the Domain Time Server running on this machine set to synchronize using PTP. Otherwise, only use this setting if the local time on your machine is being well-corrected by a reliable process, either by Domain Time or another source, such as an internal GPS clock card.

■ Use this machine's sources

When selected, Manager will use the same time sources used by the Domain Time Server or Client installed on the local machine. This is an excellent option if you have already configured the local Server or Client to obtain time from reliable sources using the NTP or DT2 protocols.

- **Specify a list of servers**

Use this option to specify the exact machines you want to use for your reference time.

- **Discover DT2 server(s)**

- **Discover NTP server(s)**

- **Discover any available server(s)**

The auto-discovery options allows Manager to locate available servers of the selected type on the network. Discovery will use all discovered servers if the *Analyze all listed servers and choose the best...* checkbox is checked, otherwise it will use the first discovered server.

Note: To avoid the possibility of inadvertently using a free-running local clock, the discovery process will not use the local machine, even if the local machine is a time server.

Analyze time samples and choose the best, or average equally good samples (recommended)

This controls whether Manager applies advanced analysis algorithms to the collected time samples.

When this box is checked, Manager contacts all of the listed servers to collect a group of time samples. It then performs statistical analysis on the collected samples to determine the reliability and uses the most reliable samples to derives the correct time.

See the [About Time Samples](#) sidebar for more information and rule-of-thumb suggestions on acquiring time samples.

If you are collecting multiple samples, checking this box will almost always improve your reference time's accuracy and reliability.


If this box is unchecked, no comparative analysis among samples is performed. In addition, the list of time servers to query becomes a fallback-only list. In other words, Manager will only contact first listed time server. This server will always be used unless it is unavailable, at which point the next listed server will be used. If that server is unavailable, the next server in the list will be tried, etc. When the first listed server becomes available again, the Server will revert to using it exclusively.

Appearance and Interface

These items on the *Options -> Appearance and Interface* menu control how items are displayed in Manager:

- **Format Options**


These settings control how times and dates are displayed in the various fields and reports of Manager, and Audit Server (if installed). You can also control the precision of digits displayed in variance and timings calculation fields. Be sure you have selected sufficient significant digits to meet the timestamp granularity of any regulatory requirements you may have.



Format Options Dialog [Click for larger size]

■ Interface Options

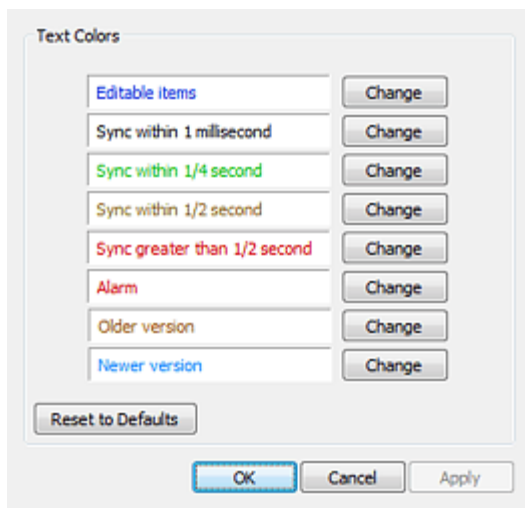
These settings control how you interact with Manager to add, delete, or change items in the program. If you have [Audit Server](#) installed, it also controls how many audits are displayed in Details Pane when you select the *Audit Server -> Audit Results* item from the Tree pane.



Interface Options Dialog [Click for larger size]

■ Interface Colors

Use these settings to set the display color of various items in Manager.



Interface Colors Dialog [\[Click for larger size\]](#)

Manager Log Settings

Manager keeps a log to allow you to easily see its activity and troubleshoot any problems you may encounter.

The Manager Log is named **dtman.log** and is kept in the **%SystemRoot%\System32** folder.

To view these logs, click the [button](#), which launches the Domain Time Log Viewer.

Log Level

This drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**
This switch will disable the **dtman.log** file.
- **Errors**
Only messages marked as Errors will be logged.
- **Warnings**
Logs will include both Errors and Warnings.
- **Information**
Includes Errors, Warnings, and Information messages.
- **Trace**
Includes all of the above, plus additional detailed trace messages.
- **Debug**
Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

Max size



This sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to help you determine the timeframe and scope of any issues you may encounter.

Using the Manager Interface

The Domain Time II Manager window is divided into two main Panes:

- The left-hand Pane  (Tree) contains a selection tree showing the various machine and report lists.
- The right-hand Pane  (Details) shows detailed information on the item selected in the Tree Pane.

Tree Pane

The Tree Pane consists of three machine lists (four if Audit Server is installed), and several additional informational sections:

Domains and Workgroups

This is a list of known machines on the network, grouped by domain. This list is useful for identifying machines on which to install Domain Time for the first time.

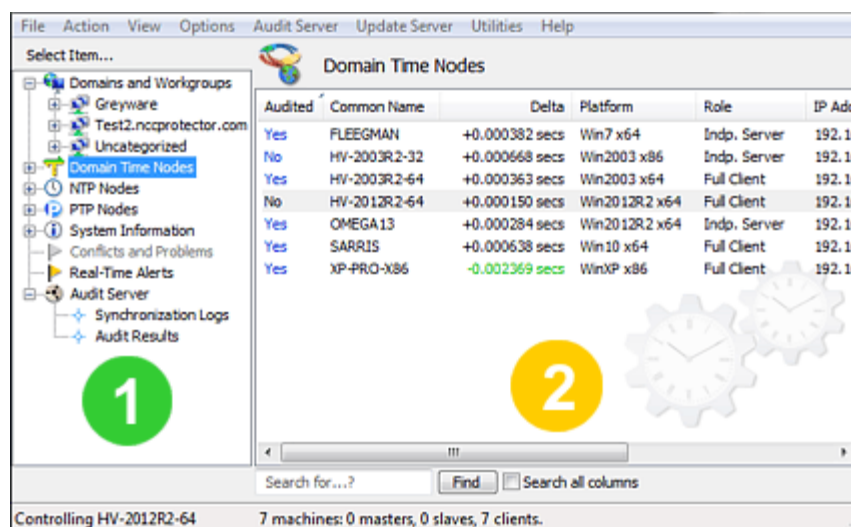
The list is derived primarily by enumeration from Active Directory (if available). You can also supplement this list by enumerating the Windows Networking Browse List. This list will also show any machines running Domain Time discovered by multicast/broadcast. See the [Discovery](#) page for details on how machines are enumerated and discovered.

Important: This list is semi-persistent.

- Machines enumerated from Active Directory or the Browse List are cached and will persist as long as they continue to be present in those lists. Manager can be set to automatically remove or tombstone machines that are removed from Active Directory (see [Computer List Enumeration](#) for details).
- Machines not automatically enumerated or discovered can be manually added to individual domains this list. To manually add a machine, highlight the desired domain in the Tree and right-click to select *Add Computer...* from the context menu. Manually-added machines will persist until manually removed.
- Machines may be manually removed from the list by expanding the tree list in the left-hand column (or in the right-hand column in recent versions), then right-clicking the the machine name and choosing *Remove from cache* from the context menu. However, if the machine is rediscovered/re-enumerated by any discovery method, it will reappear in the list.
- Large numbers of machines may be added or removed from the cache using batch files (see the [Batch Add](#) section below)
- Any machine marked as *Audited* (if [Audit Server](#) is installed) will persist, even if the machine is later removed from Active Directory or it disappears from the Browse list.
- To rediscover/re-enumerate machines in a particular category, highlight the category in the Tree and choose *Action -> Refresh* from the menu or right-click and choose *Refresh* from the context menu. To rediscover/re-enumerate all machines, highlight the top-level **Domains and Workgroups** item and choose *Refresh*.

Unknown, Uncategorized, and Hidden machines


Machines will show up in this list as unknown (greyed-out) if it cannot be determined that Domain Time is present on the machine (such as if the Domain Time service is not installed, not started, not responding to multicast/broadcast discovery,



Domain Time II Manager - Panes [\[Click for larger size\]](#)

etc.). If the machine is later determined to be running Domain Time, the details for this machine will be updated and the icon will change to show the running Domain Time machine type.

Manager makes a best-guess at categorizing machines into their proper domain. If the domain cannot be identified, machines will be listed as **Uncategorized**. A machine can be moved manually to any domain category by highlighting it in the Tree, then double-click to edit the *Domain* field on its Details Pane.



You may choose to hide machines on the list by right-clicking the computer name in the Tree and choosing *Hide Computer* from the context menu (or by double-clicking the *Hidden* editable field on the machine's Details pane).

You can toggle whether to temporarily display all hidden machines by selecting the *Show Hidden Computers* item on the *View* menu. Hidden machines displayed this way will be shown with a ~~struckthrough~~ font in the Tree list.


If you want to reset the hidden flag on all systems to the default setting, choose *View -> Unhide all Hidden Computers* from the menu.

Domain Time Nodes

This list shows all machines discovered to be up and running Domain Time when the list was last refreshed. This list is useful for showing current variances, statistics, and other Domain Time-specific information.

You cannot manually add machines to this list. The machines that appear in this list are discovered either by multicast or broadcast scans. To configure [Network Discovery](#), right click the Domain Time Nodes category and choose *Configure* from the context menu.

Important: This list is not persistent. It will be updated each time the program is started (if desired, see the [Scan Options](#) page) or when the list is manually refreshed. To refresh the list, highlight the **Domain Time Nodes** category in the Tree and choose *Action -> Refresh* from the menu or right-click and choose *Refresh* from the context menu.



If a new Domain Time II machine is found by multicast/broadcast, it will appear in this list and also be added automatically to the **Domains and Workgroups** list above. If a record already exists for the machine there, it will be updated with current information. Note, although you may mark a machine as *Audited* on the Details pane for a machine in this list (assuming you have Audit Server installed), the record for the machine will only persist in the **Domains and Workgroups** list.

As of v5.2.b.20190701, you may also manually scan IPv4 subnets for machines not discovered by multicast or broadcast. See [Manually Scan Ip4 Subnet](#) for details.

It's quite possible to get multiple entries on this list from a single machine, since there may be more than one IP address and more than one Domain Time Component installed on a machine. Each of these may send a reply to the discovery packets, resulting in duplicated entries. You can toggle whether to show all responses or hide the duplicated items by clicking the *Show Duplicate DT2 Responses* item on the *View* menu.

NTP Nodes

This list shows all machines that respond to standard NTP time requests that have been discovered by multicast/broadcast or that were manually added to Manager. This includes machines running ntpd or chrony on Linux, or other NTP daemons on other platforms. Even Domain Time Servers that have NTP enabled will appear on this list.

This list is useful for displaying the current variance of your NTP machines, and to select which NTP machines to audit (if you have Audit Server installed).

Important: The list is always persistent. No machines will be removed unless you manually delete them.

Machines newly discovered by multicast/broadcast will be added each time the program is started (if desired, see the [Scan Options](#) page) or when the list is manually refreshed. To trigger a new discovery scan, highlight the category in the

Tree and choose *Action -> Refresh* from the menu or right-click and choose *Refresh* from the context menu.

To configure [Network Discovery](#), right click the NTP Nodes category and choose *Configure* from the context menu.

As of v5.2.b.20190701, you may also manually scan IPv4 subnets for machines not discovered by multicast or broadcast. See [Manually Scan Ip4 Subnet](#) for details.

To manually add single NTP machines to the list, highlight the **NTP Nodes** category in the Tree pane, then right-click to select *Add NTP Node...* from the context menu. Enter the DNS Name or IP address of the NTP Node on the **Add Machine** dialog. Large numbers of machines may be added or removed from the cache using batch files (see the [Batch Add](#) section below)

PTP Nodes

As of version 5.2.b.20170101, this category will appear when Audit Server is installed. If PTP Monitor is enabled (it is off by default), this list shows all PTP nodes discovered by Audit Server's powerful PTP Monitor process. Using PTP Monitor, you may discover, audit, and track the availability of virtually any PTP device on your network, regardless of platform or manufacturer. Read more about [PTP Monitor](#).

This list is useful for displaying the current node identifying information, time deltas, up state, and to select which nodes to audit. Right-click the column titles to customize which information is displayed in the list.

Newly discovered nodes will be added each time Manager is started, when the list is manually refreshed., or when an audit runs. Discovery sweeps may also be set to run on a regular basis if desired, see the [PTP Monitor Configuration](#) page.

Important: This list is persistent. Machines will not be removed unless you manually delete them, or if you have enabled auto-removal of non-responding nodes on the [Configuration](#) page.

System Information

Click this category in the Tree to show important information about the system running Domain Time II Manager (this is also the same information shown when you select *Help -> About Domain Time II Manager...* from the menu).

There are two other items contained in the **System Information** category:

✦ **Reference Time** shows the current status of the reference clocks selected on the *Options -> Reference Time* page. Click [here](#) for more information on configuring Manager's Reference Time.

✦ **License Report** shows the license status of all known Domain Time components.

The remaining time on any evaluation versions will be displayed, otherwise machines will be shown as registered, along with their current version information and install dates.


Summary information displaying the totals of each type of license appears on the status bar on the bottom of the Manager window.

You may export this report using the *File -> Export List...* command on the menu.


Conflicts and Problems

Items will appear in this category if a problem is detected by Manager.

For example, a conflict will be flagged if Domain Time machines are detected running with the same Domain Time serial number. Domain Time serial numbers must be unique.

In most cases, you can resolve the conflict by highlighting the conflict item  in the Tree and double-clicking the [Repair](#) link on the Details Pane. This will bring up the **Reset Duplicated Serial Numbers** dialog which lets you select which

machines to renumber.



Domain Time II Manager - Reset Duplicated Serial Numbers dialog [\[Click for larger size\]](#)

If you are unable to correct the conflict using this utility, you will need to manually correct the serial number on the machines:

- On Windows machines, run the following command as administrator from the command-prompt on one of the affected machines:

```
dtcheck /resetserial
```

- On Linux and UNIX machines running domtimed, the serial number is derived from the HOSTID variable. Each flavor of 'nix sets this value in its own way; you will need to research how to change the HOSTID value for your particular system.

Real-Time Alerts

This category shows you any Real-Time Alerts reported from Servers and Clients. Real-Time Alerts are configured and collected by [Audit Server](#).

Audit Server

If you have installed [Audit Server](#), this category will appear and show you summary information about the service.

There are two other items contained in the **Audit Server** category:

✦ **Synchronization logs** shows you all machines for which you have collected Synchronization (Drift) Logs. You can double-click on the machine's name in the Details Pane to display its Drift Graph.

You can also see this information from the *Audit Server -> Synchronization Logs* menu item. You can also Right-click to choose *Open Containing Folder* from the context menu if you want to browse the actual log files using Windows Explorer.

✦ **Audit Results** shows a list of the most recent audit runs. Double-click the name of the audit in the Details Pane to display the complete Audit Record.

The number of results shown is configurable using the *Options -> Appearance and Interface -> Interface Options* menu item. You can also Right-click to choose *Open Containing Folder* from the context menu if you want to browse all of the audit results files using Windows Explorer.

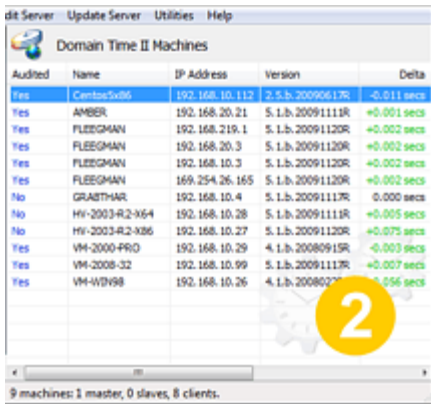
Details Pane

The Details Pane will change to display information on the item selected on the Tree Pane.

If the selected Tree item is a single machine or topic, individual statistics and information about that machine will be shown.

Information in the displayed Details Pane can be updated by clicking *Action -> Refresh* from the menu or by right- clicking and choosing *Refresh* from the context menu.

If the Tree item is a category that contains multiple items, you'll see either summary



Audited	Name	IP Address	Version	Delta
Yes	Centos506	192.168.10.112	2.5.3.20090617R	-0.011 secs
Yes	AMBER	192.168.20.21	5.1.3.20091111R	+0.001 secs
Yes	FLEEGMAN	192.168.219.1	5.1.3.20091120R	+0.002 secs
Yes	FLEEGMAN	192.168.20.3	5.1.3.20091120R	+0.002 secs
Yes	FLEEGMAN	192.168.10.3	5.1.3.20091120R	+0.002 secs
Yes	FLEEGMAN	169.254.26.165	5.1.3.20091120R	+0.002 secs
No	GRABTHAR	192.168.10.4	5.1.3.20091117R	0.000 secs
No	HV-2003-R2-X64	192.168.10.28	5.1.3.20091111R	+0.005 secs
No	HV-2003-R2-X86	192.168.10.27	5.1.3.20091120R	+0.075 secs
Yes	VM-2000-FRO	192.168.10.29	4.1.3.20080915R	-0.003 secs
Yes	VM-2008-32	192.168.10.99	5.1.3.20091117R	+0.007 secs
Yes	VM-WIN98	192.168.10.26	4.1.3.20080227R	-0.556 secs

Manager - Details Pane [\[Click for larger size\]](#)

information or a table display showing the items.

You can also control how Manager displays items, fields, and lists using the [Options -> Appearance and Interface](#) menu item.

Layout


Table displays are highly customizable. You can control which columns are displayed, and in what order. You can also sort the table list by any column, making it very easy to see all the machines running a particular version, which machine has the worst variance, etc.

To change which columns are displayed, right-click on any column header. You can pick from a list of all available columns.

To re-arrange the order of displayed columns, grab the column header and drag it to the desired position. Manager will remember your selections.

To sort a column, just click on its column header. Click again to change the sort order.

You may reset any view to the default settings by selecting the *View -> Reset Column Layout* from the menu. You may reset individual views or reset all views.




[\[Click for larger size\]](#)

Editable Items

The Details Pane will often contain fields that can be edited for selected machines, such as whether a machine is *Audited* (if *Audit Server* is installed), *Hidden*, etc. You may also change a machine's domain name, DNS name, location (if an NTP Node), etc. You can edit these fields by double-clicking them (this default behavior may be modified to be done via a single-click on the *Options -> User Interface Settings* menu page).

Here's an example of changing the timezone of a machine remotely by editing its Details record:



Domain Time II Manager - Timezone Change [\[Click for larger size\]](#)

Editable items are shown in the color selected for them on the *Options -> Appearance and Interface -> Interface Colors* page.

Export List

Most displays in the Details pane can be exported to a file using *File -> Export List...* from the main menu.

For example, you could export the current list of NTP Nodes to a comma-delimited text file (.CSV extension) for later import into a database, or export the License Report to an XML file for use in Excel, etc.

Exports can include only the visible columns in the Details pane, which allows you to easily select which data to include, or you may choose to include hidden columns as well (there is a checkbox for this on the *Save As* dialog).

You may export to CSV (comma-separated value), XML, or HTML formatted files.

License Report

Eval Status	Name	Type	Version	Install Date
Registered	FLEEGMAN	Monitor	5.1.b.20090802R	Fri 9 Oct 2009 22:16:11
Registered	FLEEGMAN	Audit Server	5.1.b.20090802R	Fri 9 Oct 2009 22:16:11
Registered	FLEEGMAN	Manager	5.1.b.20090802R	Fri 9 Oct 2009 22:16:11
Registered	VM-WIN98	Client	4.1.b.20080229R	Sat 1 Mar 2008 00:01:45
Registered	Centos5x86	Client	2.5.b.20090617R	Thu 18 Jun 2009 02:00:06
Registered	VM-VISTA-N-32	Client	5.1.b.20090723R	Thu 8 Oct 2009 03:50:43
Registered	OMEGA13	Client	4.1.b.20080916R	Tue 21 Apr 2009 13:22:38
Registered	HV-2008-R2-X64	Client	5.1.b.20090802R	Sat 26 Sep 2009 01:01:13
Registered	XP-PRO-X86	Client	4.1.b.20080916R	Wed 16 Jul 2008 10:37:56
Registered	VM-2003-R2-32	Server	5.1.b.20090727R	Thu 22 Oct 2009 14:26:10
Registered	VM-2000-PRO	Server	4.1.b.20080915R	Fri 29 Feb 2008 23:55:54
Registered	FLEEGMAN	Server	5.1.b.20090802R	Fri 9 Oct 2009 22:16:11
Registered	VM-2008-64	Server	5.1.b.20090720R	Wed 15 Apr 2009 17:28:29

Sample HTML Export of the License Report

Other Items

Backup/Restore Database

As of version 5.2, Manager can backup and restore its database using the *File -> Backup Database* and *File -> Restore Database* commands from the main menu.

You may choose to save the backup files in any file location you want, however, any files you want to restore must be located in the **\Program Files\Domain Time II\Backups** folder on your system drive.

Database backups are not intended as comprehensive disaster recovery files. Database backups only include information related to individual machines, such as whether particular machines are audited, machine type, etc. The backups **do not** include information that is kept in the registry, such as Manager's own settings, or Audit Server schedules, alert settings, email configuration, etc.

Backups are primarily for making working copies of the database, such as when making large-scale changes or transporting audit lists to another Manager/Audit Server.

Backup/Restore Audit List and Audit Group settings

As of version 5.2.b.20180606, Manager can backup and restore Audit list and group configuration settings using the *File -> Audit List and Audit Group settings* and *File -> Restore Audit List and Audit Group settings* commands from the main menu.

You may choose to save the backup files in any file location you want, however, any files you want to restore must be located in the **\Program Files\Domain Time II\AuditList** folder on your system drive.

Audit list and group settings backups are not intended as comprehensive disaster recovery files. They only include the limited set of current audit lists and group alerting settings.

Backups are primarily for making working copies of the audit settings, such as when making large-scale changes or transporting audit lists to another Manager/Audit Server.

Batch Add

Manager also allows you to add/remove machines in a batch from a text file. This is done by entering the following command from the command-line:

```
dtman import [full path]\filename.txt
```

The %ERRORLEVEL% result is zero if the file was parsed without errors or warnings. If %ERRORLEVEL% is non-zero, you should examine the Manager logs (*Options -> View Log File* from the Manager menu, or open `\Program Files\Domain Time II\Text Logs\Manager.log` with a text editor) to see the errors or warning messages.

The text file format is very simple:

SYNTAX: [add|drop|delete] [dt2|ntp|ptp|all] [ipaddress|name] (ipaddress|name[:domain] for ptp) [AuditGroup]

add all [AuditGroup]	- sets any existing unaudited machine to audited
add ntp <i>identifier</i> [AuditGroup]	- adds NTP machine to the database and sets it to audited
add dt2 <i>identifier</i> [AuditGroup]	- adds DT2 machine to the database and sets it to audited
add ptp <i>identifier[:domain]</i> [AuditGroup]	- adds a PTP machine to the database and sets it to audited
drop all	- sets any existing audited machine to unaudited
drop ntp <i>ipaddress</i>	- sets existing NTP machine to unaudited
drop dt2 <i>netbiosname</i>	- sets existing DT2 machine to unaudited
drop ptp <i>ipaddress</i>	- sets existing PTP machine to unaudited
del all	- removes all records from the database
del ntp <i>ipaddress</i>	- removes existing NTP machine from the database
del dt2 <i>netbiosname</i>	- removes existing DT2 machine from the database
del ptp <i>ipaddress</i>	- removes existing PTP machine from the database

AuditGroup is an optional parameter for ADD. If present, it must come after the *identifier*, and be a single-digit number 1-8, corresponding to the audit group to which you want the node added. As of version 5.2.20200405, you may specify 0 (zero) to mean you want the node unaudited. If you do not specify an AuditGroup, the default group number will be used.

Adding an NTP or DT2 machine (or setting an existing machine to audited) can use any kind of *identifier* that resolves to an IP address. You may use a plain NetBIOS name, an IP address, or a DNS name.

Batch add for PTP nodes (included as of v5.2.b.20180606) is limited to unicast-only slaves. You may use either an IP address or a DNS name as the *identifier*, optionally followed by a colon and a PTP domain number. For example, ADD PTP 192.168.1.200:3 would add the PTP slave at 192.168.1.200, using domain 3. If you leave the colon and domain number off, Manager will attempt to discover the domain. If you specify a domain number that isn't one of the ones being monitored by PTP Monitor, the node won't be added. If the node already exists in the database but with a different domain number, the domain number will be updated to the one you specify. If multiple PTP nodes reside at a single IP address, Manager will add all discovered nodes. Note: This also applies to the drop and del commands, so use them carefully.

Dropping or deleting machines, however, requires that you use the exact name by which the machine is known to the database. For NTP and PTP machines, this is the IP address (either IPv4 or IPv6). For DT2 machines, this is the NetBIOS name.

CAUTION:

Use the **all** options with great care, particularly the **drop all** command. Batch imports take effect immediately and cannot be undone.

Comment lines are ignored. A comment line is any line whose first non-whitespace character is a pound sign (#), a glitch ('), a slash (/), or a semicolon (;). Blank lines are also ignored.

The text file should be CRLF terminated. If you are importing from a *nix operating system, be sure to transform the line terminators to full CRLFs.

Example import file:

```
#  
# this is an example  
#  
  
add      ntp      tick.mydomain.com  
add      dt2      PDC01  
add      dt2      pdc02.south.brooklyn.ny.company.com  
drop     dt2      pdc03  
add      ntp      ntp2.ourpool.org  
add      ptp      192.168.1.200:3  
delete   dt2      pdc04  
delete   ntp      tock.mydomain.com  
  
#  
# end of list  
#
```

How to Manage Domain Time Remotely

You can use Domain Time II Manager to easily accomplish these common management tasks:

Click a link to jump to the discussion about each task:

Remote management

- ▶ [Install/Upgrade/Remove](#) Domain Time remotely
- ▶ [Reset](#) the configuration on one or more machines
- ▶ [Reset](#) the symmetric keyring on one or more machines
- ▶ [Remotely control](#) and configure individual machines
- ▶ [Set the timezone](#) on one or more machines
- ▶ [Get machine stats](#) and trigger synchronization
- ▶ [Troubleshoot](#) problems and resolve conflicts

Generate reports

- ▶ [Report the variance](#) of machines on the network
- ▶ [Report the license status](#) of Domain Time


Audit Server functions

- ▶ [Select machines to audit](#) with Audit Server
- ▶ [Real-time Alerts](#)

Install/Upgrade/Remove Domain Time remotely

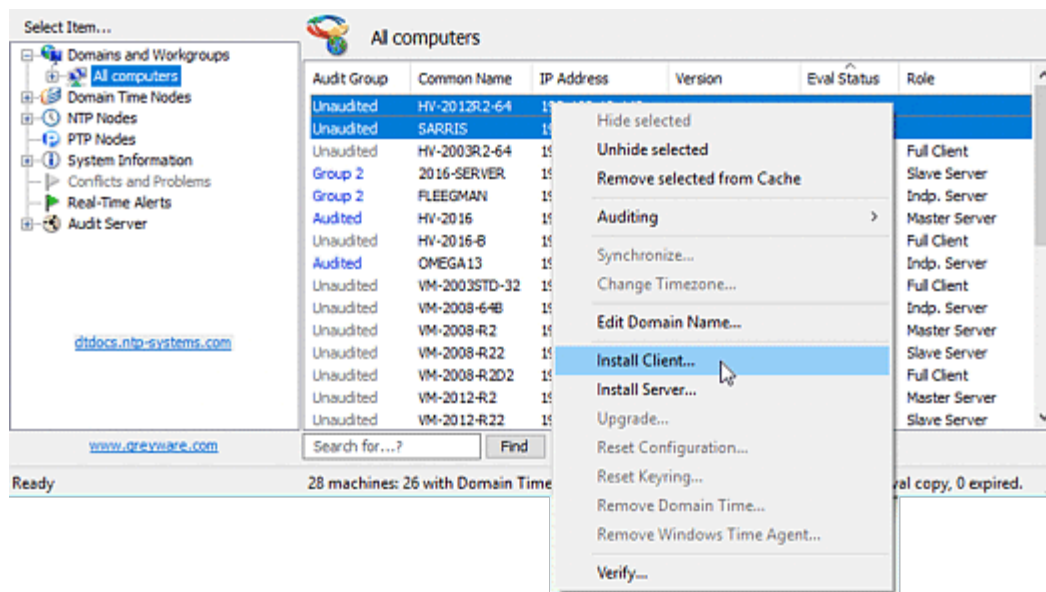
Be sure you meet the [minimum requirements for remote operations](#).

- On a single machine:
Highlight the target machine's name from either the list on the Tree pane or from a list on the Details pane, right-click and choose the action you want from the context-menu.
- On multiple machines:
 - Click on a domain category in the **Domains and Workgroups** list on the Tree pane to display all of its machines in the Details pane.
 - Sort the list by any column you'd like to help identify the machines you want to work with.
 - You can also refine which machines to choose by right-clicking in the Details Pane and choosing *Select* from the context menu to highlight different types of machines (Servers, Client, Installed machines, Uninstalled machines, etc.).
 - Select the machines you'd like as targets (using Shift+Click, Ctrl+Click, etc.) and then right-click to select the action you want to perform from the context-menu.



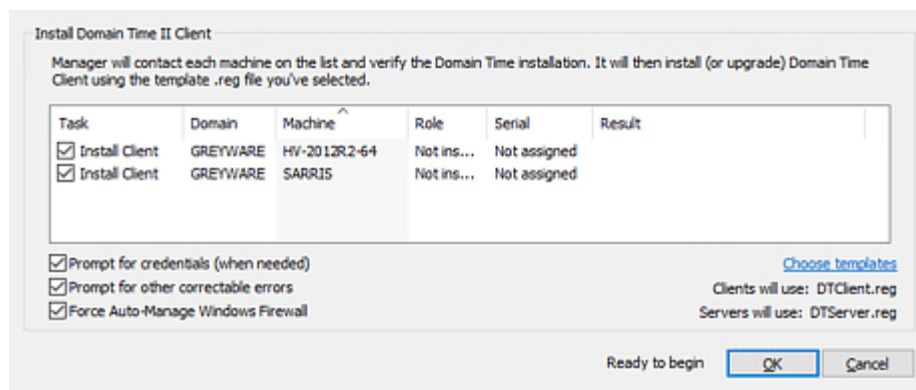
Installing Server on a single machine

[\[Click for larger size\]](#)



Select multiple machines to install [Click for larger size]

Before proceeding, you have the option to verify your selection(s) on the Remote Computer Operation dialog.



Remote Computer Operation dialog [Click for larger size]

Installations and upgrades will use Manager's default configuration Templates for Server or Client unless you select alternates by clicking the [Choose templates](#) link. [Read about using Templates](#)

Prompt for credentials (when needed): Manager uses the Windows credentials store to securely keep the usernames/passwords used to connect to machines. When this box is checked, Manager will prompt for any credentials it does not already know. When unchecked, Manager will move on to the next system without prompting if the connection fails.

Prompt for other correctable errors: When checked, Manager will display a prompt if it cannot connect to a machine, allowing you to attempt to correct the error (such as by supplying an IP address instead of a DNS name). When unchecked, Manager will merely move on to the next system without prompting when it encounters an error.

Force Auto-Manage Windows Firewall

As of Version 5.2.b.20150828, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. This checkbox is tri-state (click on the box several times to cycle through the options). See [Auto-Manage Windows Firewall Settings](#) for a detailed explanation of the effect of each option.

When you click **OK**, Manager will proceed, showing you the status of the operation on each machine. Any errors will be logged in Manager's log file. Be sure to review it by pressing F9 or choosing [View Manager Log File...](#) from the Options menu.

If Manager encounters an error processing a machine, use the **Connection Troubleshooter** (highlight the machine in the

Tree or Details pane and click *Test Connection...* from the right-click context menu) to help identify the cause of the problem. See the [Troubleshoot](#) task below for more information.

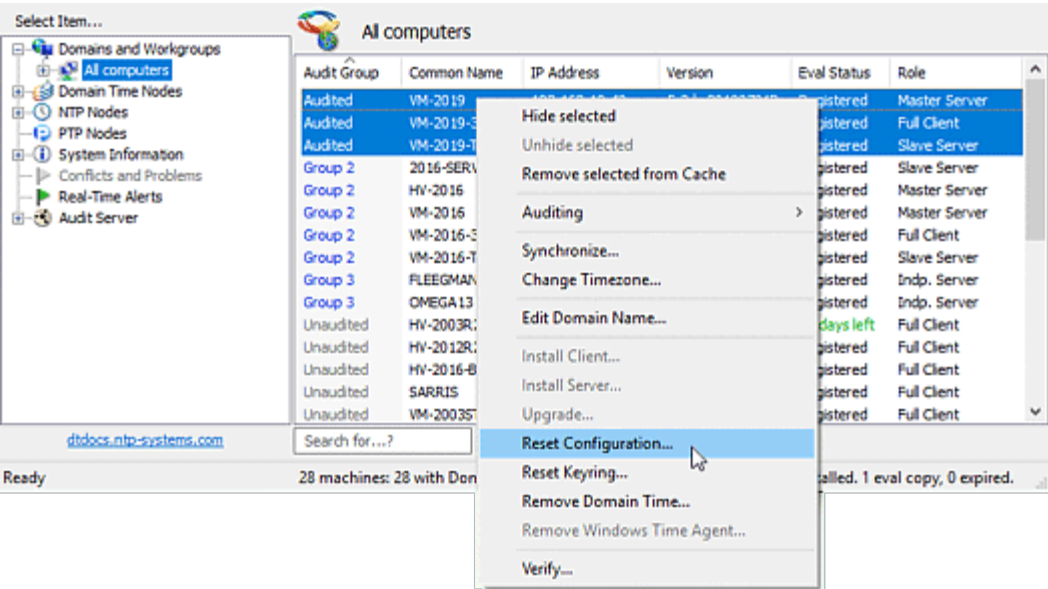
Reset the configuration on one or more machines

Be sure you meet the [minimum requirements for remote operations](#).

After Server or Client is installed on a machine, Manager can cause the component to reset itself to the default settings it was installed with, or Manager can push out new settings for the component to use. You can do this for any number of installed machines.

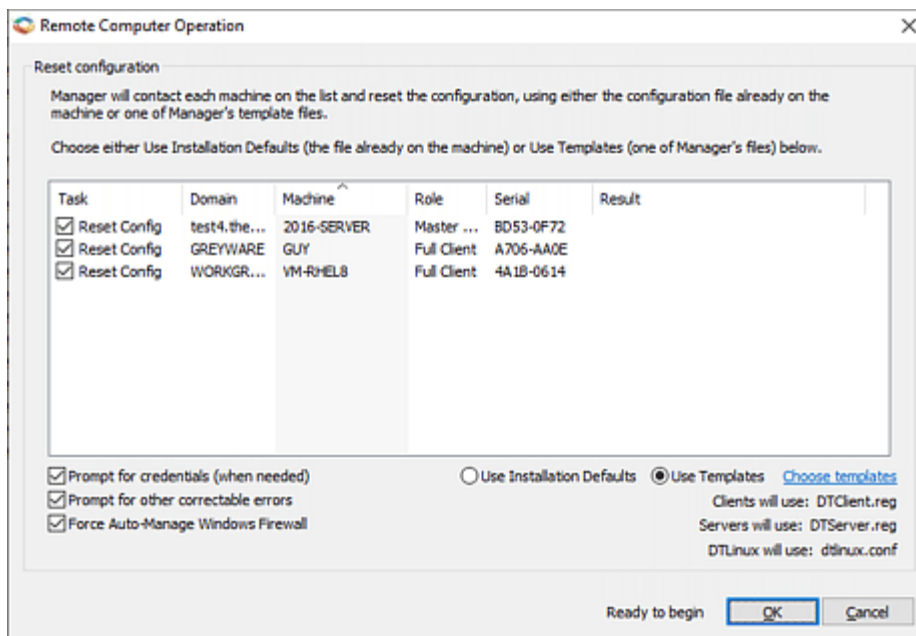
This function allows you to easily make large-scale changes to groups of machines if necessary. For example, you could choose to change the specified time sources for all the machines in a particular city, or adjust the sync timings of all your servers, etc.

The process is very similar to the [Install](#) procedure described above, except that you use the *Reset Configuration* command from the right-click context menu once you have selected the target machine(s).



Reset Configuration on selected machines [Click for larger size]

The Remote Computer Operation dialog lets you confirm your selected machines, and also decide what method to use for the reset:



Remote Computer Operation dialog [Click for larger size]

Use Installation Defaults: Causes the machine to revert to the exact configuration it used when it was first installed.

Use Templates: Templates can set all or just a portion of the machine's configuration parameters. Click the [Choose templates](#) link to specify a template to use. [Read about using Templates](#)

Prompt for credentials (when needed): Manager uses the Windows credentials store to securely keep the usernames/passwords used to connect to machines. When this box is checked, Manager will prompt for any credentials it does not already know. When unchecked, Manager will move on to the next system without prompting if the connection fails.

Prompt for other correctable errors: When checked, Manager will display a prompt if it cannot connect to a machine, allowing you to attempt to correct the error (such as by supplying an IP address instead of a DNS name). When unchecked, Manager will merely move on to the next system without prompting when it encounters an error.

Force Auto-Manage Windows Firewall

As of Version 5.2.b.20150828, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. This checkbox is tri-state (click on the box several times to cycle through the options). See [Auto-Manage Windows Firewall Settings](#) for a detailed explanation of the effect of each option.

When you click **OK**, Manager will proceed, showing you the status of the operation on each machine.

If Manager encounters an error processing a machine, use the **Connection Troubleshooter** (highlight the machine in the Tree or Details pane and click *Test Connection...* from the right-click context menu) to help identify the cause of the problem. See the [Troubleshoot](#) task below for more information.

Reset the symmetric keyring on one or more machines

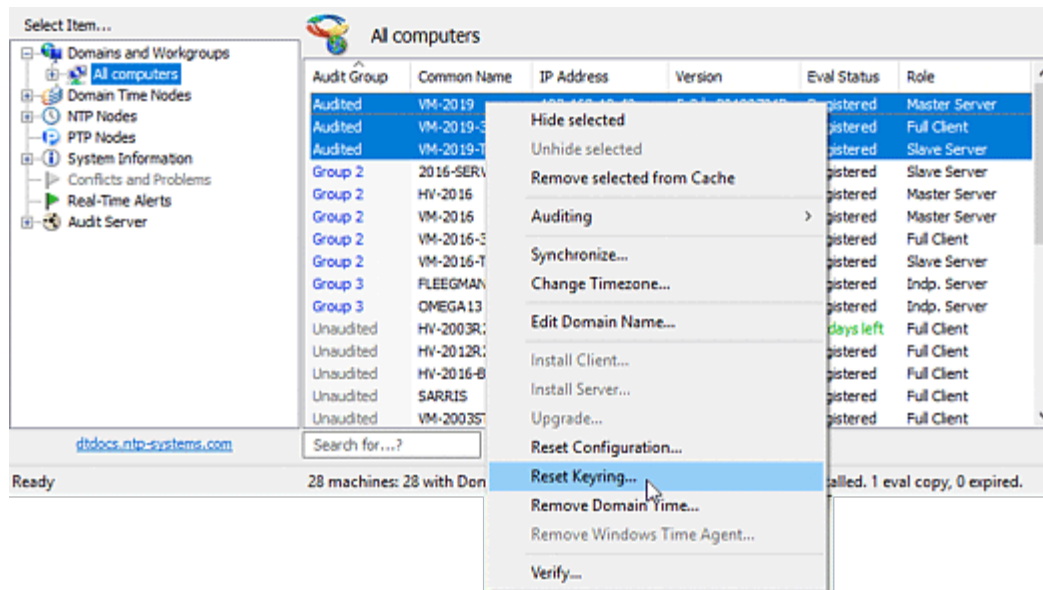
Be sure you meet the [minimum requirements for remote operations](#).

Manager can push out a common symmetric keyring to Domain Time Servers and Clients. You can do this for any number of installed machines. This allows you to easily ensure all your machines are using the same set of symmetric keys for authentication. Read more about [Symmetric Keys](#).

The process is very similar to the [Reset Configuration](#) procedure described above, except that you use the *Reset Keyring* command from the right-click context menu once you have selected the target machine(s). The right-click option will not be available for the Domain Time Server on the same machine as Manager, because it is, by definition, already up to date.

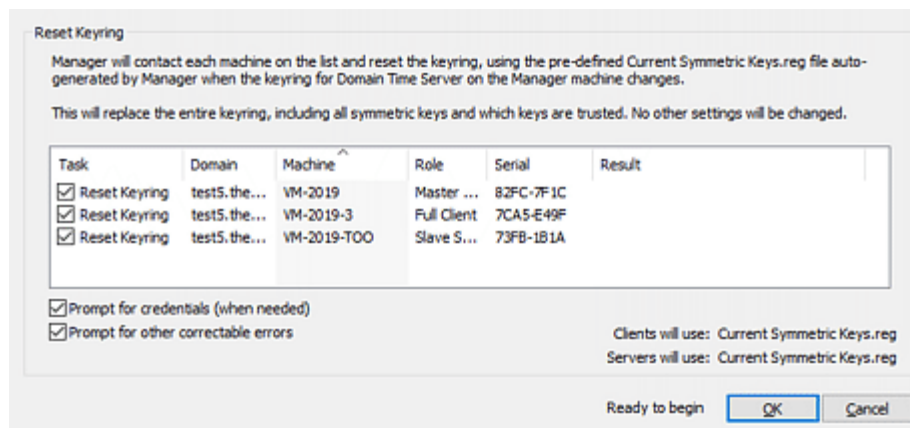
The keyring configured on the Domain Time Server on the Manager machine is used as the master keyring. Each time you modify the Server's keyring, registry files called **Current Symmetric Keys** are created/updated on Manager in the Client and Server template folders. Reset Keyring uses these files to reset the keyrings on your selected machines.

NTFS security restricts these files to Administrators and SYSTEM only. To change the information in these auto-generated templates, use the Control Panel applet for Domain Time Server on the Manager machine and edit the keyring. Do not attempt to edit the keys manually. (There are also "Compatible" versions of these two files for use with older Domain Time machines that don't understand the syntax for clearing a key's values before repopulating with information from the template. Manager will automatically select the Compatible version when required.)



Reset Keyring on selected machines [\[Click for larger size\]](#)

The Remote Computer Operation dialog lets you confirm your selected machines, and also decide what method to use for the reset:



Remote Computer Operation dialog - Reset Keyring [\[Click for larger size\]](#)

Prompt for credentials (when needed): Manager uses the Windows credentials store to securely keep the usernames/passwords used to connect to machines. When this box is checked, Manager will prompt for any credentials it does not already know. When unchecked, Manager will move on to the next system without prompting if the connection fails.

Prompt for other correctable errors: When checked, Manager will display a prompt if it cannot connect to a machine, allowing you to attempt to correct the error (such as by supplying an IP address instead of a DNS name). When unchecked, Manager will merely move on to the next system without prompting when it encounters an error.

When you click **OK**, Manager will proceed, showing you the status of the operation on each machine.


If Manager encounters an error processing a machine, use the **Connection Troubleshooter** (highlight the machine in the Tree or Details pane and click *Test Connection...* from the right-click context menu) to help identify the cause of the problem. See the [Troubleshoot](#) task below for more information.

Remotely control and configure individual machines

Be sure you meet the [minimum requirements for remote operations](#).

You can use Manager to connect to any installed Server or Client and bring up its Control Panel applet, which will allow you to remotely configure and control the machine. You can change settings, view the various logs and graphs, start and stop the service, etc. just as if you were sitting at the console of the remote system.


To connect, simply double-click on the machine's name in the Tree pane. You can also highlight the name in any Tree or Details pane list, and choose *Action -> Control Machine* from the main menu, press F7, or right-click and choose *Control Panel* from the context menu.



Connect to a remote machine

[\[Click for larger size\]](#)

Manager first uses the computer name to attempt the connection. If the name does not resolve to a running Domain Time system, you'll see the Connect error dialog where you can tell Manager to try a DNS name or IP address instead. You can also instruct Manager to try the DNS name or IP address instead of the machine name for all future connections.



Connection Error dialog [\[Click for larger size\]](#)


If Manager still cannot connect to the machine, use the **Connection Troubleshooter** (highlight the machine in the Tree or Details pane and click *Test Connection...* from the right-click context menu) to help identify the cause of the problem. See the [Troubleshoot](#) task below for more information.

Set the timezone on one or more machines

Be sure you meet the [minimum requirements for remote operations](#).

■ On a single machine:

Highlight the target machine's name from either the list on the Tree pane or from a list on the Details pane, right-click and choose the action you want from the context-menu.




Change timezone on a single machine

[\[Click for larger size\]](#)


■ On multiple machines:

- Click on category in the Tree pane to display a list of machines in the Details pane.
- Sort the list by any column you'd like to help identify the machines you want to work with.
- You can also refine which machines to choose by right-clicking in the Details Pane and choosing *Select* from the context menu to highlight different types of machines (Servers, Client, Installed machines, Uninstalled machines, etc.).
- Select the machines you'd like as targets (using Shift+Click, Ctrl+Click, etc.) and then right-click to select *Timezone...* from the context-menu.



Change the timezone on multiple machines [\[Click for larger size\]](#)

Choose the timezone you want from the **Change Timezone** dialog.



Change Timezone dialog [\[Click for larger size\]](#)

If Manager encounters an error setting the timezone, use the **Connection Troubleshooter** (highlight the machine in the Tree or Details pane and click *Test Connection...* from the right-click context menu) to help identify the cause of the problem. See the [Troubleshoot](#) task below for more information.

Get machine stats and trigger synchronization

Manager presents statistical information obtained from systems in a number of ways.

- You can see basic information on a machine by highlighting its name in any list the Tree pane.

The Details pane will show information derived from a variety of sources such as Active Directory, network discovery scans, and direct contact with Domain Time on the machine.

Some data may initially be missing or not available depending on what kind of contact has been made with the machine. You can sometimes fill in missing data by clicking *Action -> Refresh* from the main menu (or right-clicking and choosing *Refresh* from the context menu).

- You can also get statistical information from machines by customizing the columns displayed in lists shown on the Details pane.


Click a category name in the Tree pane to display a list of machines in that category. Choose *View -> Add/Remove Columns...* from the main menu (or right-click on any column label and choose which columns to display).

Update the list by clicking *Action -> Refresh* from the main menu (or right-clicking and choosing *Refresh* from the context menu).

- You can get extremely detailed statistics from any machine running Domain Time.

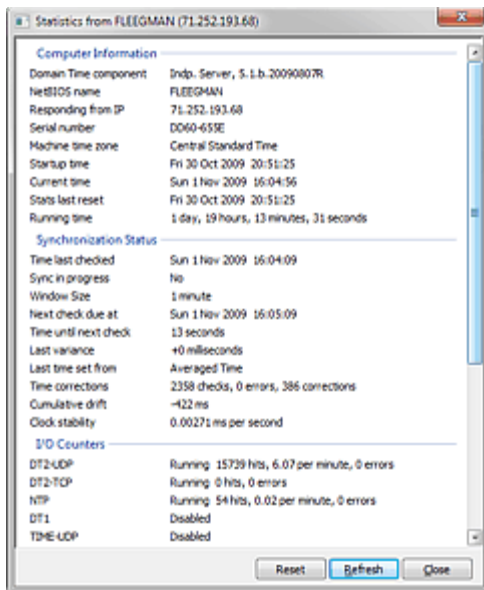
Status Information	
Alarm	None
Audited	Yes
Hidden	No
Role	Full Client
Version	5.1.b.20090802R
Timezone	Central Standard Time
Last Contact	Mon 2 Nov 2009 02:07:45
Eval Status	Registered
Serial	EEF3-C43F
Operating System/Network Information	
Domain	Test1
Domain derived from	Manual
Name	HV-2003-R2-x64
Distinguished Name	
Location	
Description	No description
DNS Name	HV-2003-R2-x64
IP Address	192.168.10.28
Operating System	Win2003 x64
OS Service Pack	
OS Version	5.2

Basic Stats [\[Click for larger size\]](#)



[\[Click for larger size\]](#)

Highlight the name of a machine running Domain Time in either a Tree or Details pane list and right-click to choose *Statistics* from the context menu. This information is provided by contacting the Domain Time II component directly, and it can therefore contain extremely detailed information about the machine and its operations. Versions of Domain Time 5.1 or newer provide additional statistics such as Network I/O Counters.



Detailed Domain Time statistics [Click for larger size]

Update the list by clicking the *Refresh* button on the Statistics dialog. You can also reset the statistics and counters on the remote machine by clicking the *Reset* button. Note that this will also clear the machine's drift graph data, so if you are using Audit Server to collect this information for historical purposes, be sure you have run a recent Audit before clearing the data. See Audit Server [Synchronization Logs](#) for more info.

- To see statistical information on the Domain Time II Manager itself, click the **System Information** category in the Tree pane. Click the **Reference Time** item underneath it to show current info on the Reference Clock.

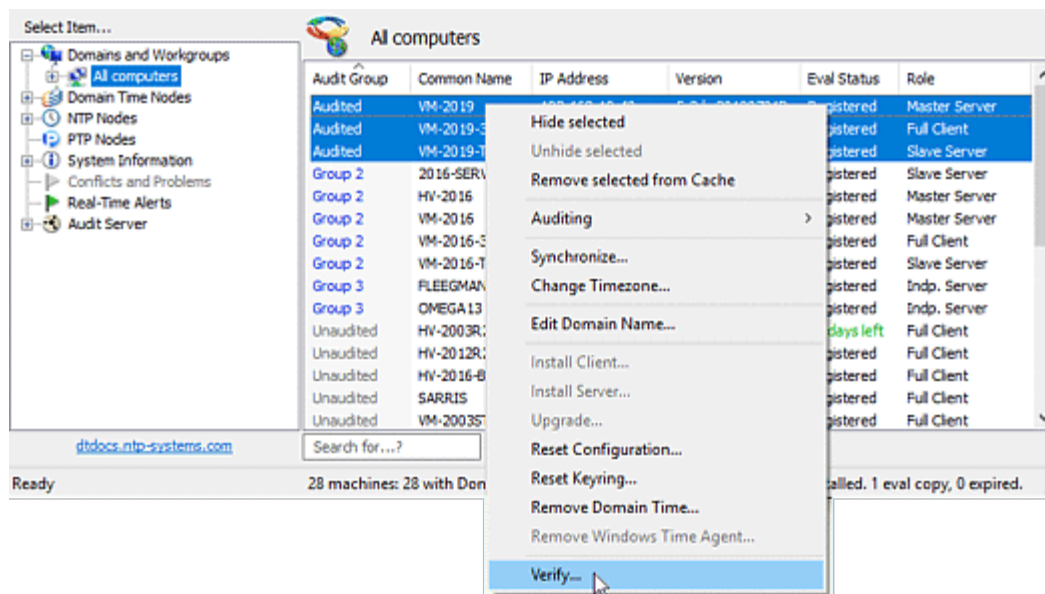
Troubleshoot problems and resolve conflicts

Manager gives you several tools to help troubleshoot connection issues and resolve serial number conflicts.

- The **Manager Log** can display extensive information on what Manager is doing. If you are having difficulties, the first step you'll want to take is to enable Manager's Debug logging using the *Options -> Manager Log File Settings* menu item.
- The **Verify** command tells Manager to attempt to identify whether Domain Time is installed on selected machines and obtain updated statistical information, if possible.

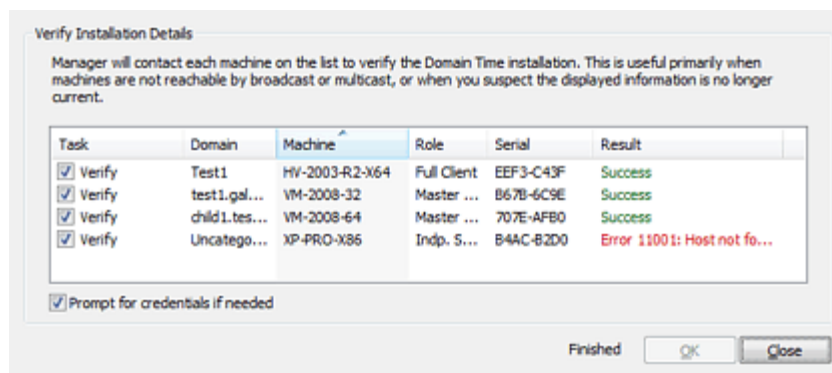
This is particularly useful if the normal [Network Discovery](#) processes aren't able to determine if Domain Time is installed, or if statistical information hasn't been filled in correctly.

- Click on a list category in the Tree pane to display all of its machines in the Details pane.
- Sort the list by any column you'd like to help identify the machines you want to work with.
- Select the machines you'd like as targets (using Shift+Click, Ctrl+Click, etc.) and then right-click to select *Verify...* from the context-menu.



Select multiple machines to verify [Click for larger size]

- Next, you'll be presented with the Network Operations dialog, which allows you to confirm which machines you want to verify:



Remote Computer Operation dialog [Click for larger size]

Prompt for credentials (when needed): Manager uses the Windows credentials store to securely keep the usernames/passwords used to connect to machines. When this box is checked, Manager will prompt for any credentials it does not already know. When unchecked, Manager will move on to the next system without prompting if the connection fails.

Prompt for other correctable errors: When checked, Manager will display a prompt if it cannot connect to a machine, allowing you to attempt to correct the error (such as by supplying an IP address instead of a DNS name). When unchecked, Manager will merely move on to the next system without prompting when it encounters an error.

When you click OK, Manager will proceed with the verification, showing you the status of the operation on each machine.

- The **Connection Troubleshooter** is an extremely helpful utility that will show you all of the steps Manager performs when connecting to a remote machine. This allows you to immediately pinpoint the exact problem preventing communication with any machine.

To launch the Troubleshooter, right-click a machine name in either the Tree or Details pane and select *Test Connection...* from the context menu. Alternately, you can select *Utilities -> Connection Troubleshooter* from the Main Menu.

Once the Troubleshooter is running, you can enter the DNS name or

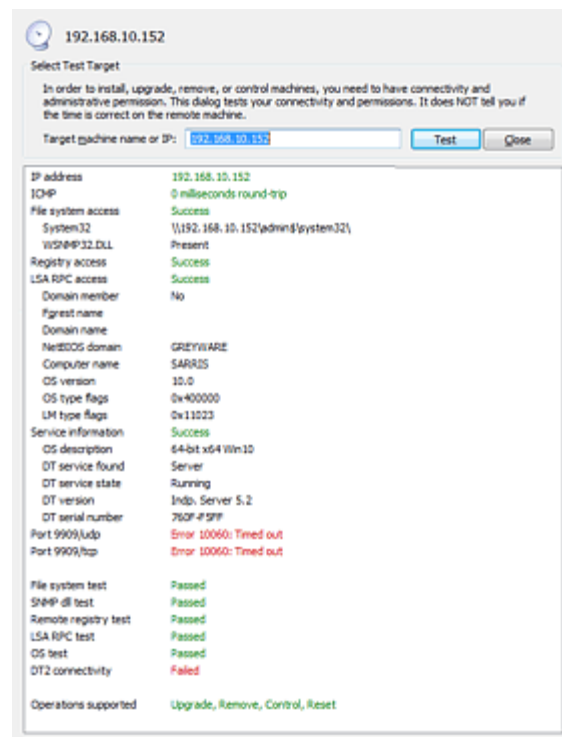
IP address of the system to test.

The Troubleshooter will perform each action required to connect to the remote machine and display the results of each test.

Critical errors that prevent particular Manager functions from operating will be shown in red, warnings that may represent a problem will be shown in yellow, and tests that pass will be shown in green.

Most connection issues are due to network configuration, name resolution, firewall, or permissions problems. Please refer to the [minimum requirements for remote operations](#) list for details on what types of access is required.

As of version 5.2.b.20170101, the troubleshooter checks for the presense of wsnmp32.dll as it must be present for Domain Time installation. Nano Server 2016 does not install this .dll by default; it must be manually installed before this test will pass. See the [Nano Server FAQ](#) for more information.




Connection Troubleshooter, showing a firewall problem
[\[Click for larger size\]](#)


■ The Conflicts and Problems Category

Items will appear in this category if a problem is detected by Manager.

For example, a conflict will be flagged if Domain Time machines are detected running with the same Domain Time serial number. Domain Time serial numbers must be unique.



Domain Time II Manager - Reset Duplicated Serial Numbers dialog [\[Click for larger size\]](#)

In most cases, you can resolve the conflict by highlighting the conflict item  in the Tree and double-clicking the [Repair](#) link on the Details Pane. This will bring up the **Reset Duplicated Serial Numbers** dialog which lets you select which machines to renumber. If renumbering the first machine listed doesn't work, try again and renumber the other(s). It doesn't matter which machine has which number, they just must be unique.

If you are repeatedly seeing machines appear with conflicts, you may be bringing up machines that were created by cloning. Please see this [KB article](#) on how to properly set up Domain Time on cloned images or manually reset the serial numbers if you cannot reset them using Manager.

Report the variance and current synchronization status of the network

You can create a custom report showing how much variance (delta) machines on the network currently have from Manager's Reference Time.

First, be sure you have set Manager to use correct and stable Reference Time. See the [Configure Reference Time](#) instructions.


Next, highlight the **Domain Time Nodes** category in the Tree pane to display a list of currently-running Domain Time machines. Choose which columns to include in your report by clicking *View -> Add/Remove Columns...* or right-clicking on any column header and choosing the columns. Be sure you include the **Delta** column, which shows the variance information. You can include any other details you want on your report, such as Latency, Last Timeset, Corrections, etc.

Finally, you can show the current variance by clicking *Action -> Refresh* or right-click and choose *Refresh* from the context menu to update the list.

If this list isn't showing you all Domain Time machines you expect, check your settings on the [Network Discovery](#) page.

The list in the Details pane can be exported to a file using *File -> Export* from the main menu. The exports will only include the currently-visible columns. You may export to CSV (comma-separated value), XML, or HTML formatted files.

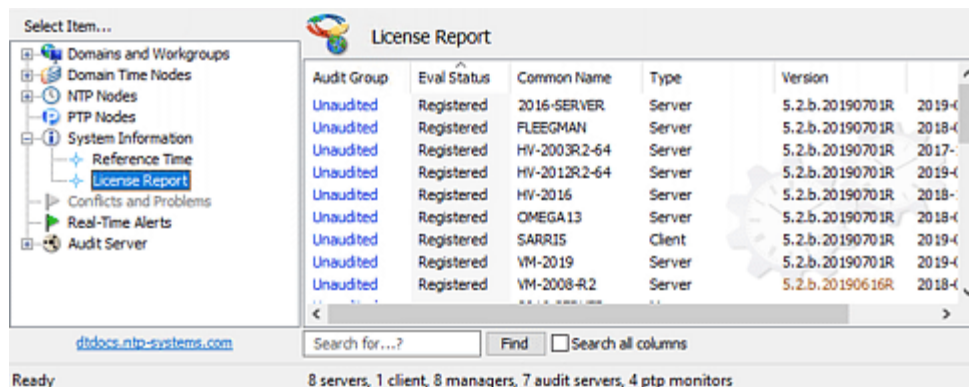
Also, see the [Real-Time Alerts](#) function for live monitoring of your time synchronization status.



Creating a Variance Report
[Click for larger size]

Report the license status of Domain Time

To show the license status and type of currently-running Domain Time machines, click the **License Report** item in the Tree pane (found under the **System Information** category).



The License Report [Click for larger size]

Choose which columns to include in your report by clicking *View -> Add/Remove Columns...* or right-clicking on any column header and choosing the columns.

If you have the **Eval Status** column visible, it will show whether a machine is running a registered copy of Domain Time, or show how many days remain on an evaluation version.

If this list isn't showing you all Domain Time machines you expect, check your settings on the [Network Discovery](#) page.


The list in the Details pane can be exported to a file using *File -> Export* from the main menu. The exports will only include the currently-visible columns. You may export to CSV (comma-separated value), XML, or HTML formatted files.

Select machines to audit with Audit Server

If you have [Audit Server](#) installed, you can use Manager to determine which machines will be included in audit runs.

You can select machines for Audit from the *Domains & Workgroups*, *Domain Time Nodes*, *NTP Nodes*, *Real-Time Alerts*, or *PTP Monitor* pane by exposing the **Audit** column and making your selection, or you can click on any individual machine and change its audited status by double-clicking on the Audited editable item shown on the Details pane. Clicking multiple times will toggle through all available Audit Groups.

To select using the **Audit** column, highlight any list category in the Tree pane to display a list of machines in the Details pane. Click *View -> Add/Remove Columns...* or right-click on any column header and then click to display the **Audit** column. You can change the audited state of any individual machine, or you can change multiple machines by selecting them, then right-click and choose *Auditing* from the context menu. You may choose from up to 8 custom audit groups. See [Audit Groups](#) for details on configuring Audit Groups.



Display Audit Column
[Click for larger size]

You may also set which machines should be audited from a batch file. See the [Batch Add](#) section of the **How to use the Manager Interface** page for details.

Note: The auditing of PTP Nodes is a separate function from other types Audit Server auditing. The "Audited" setting's column for PTP Monitor is independent of the "Audited" settings on the *Domains & Workgroups*, *NTP Nodes*, *Domain Time Nodes*, or *Real-Time Alerts* displays. Enabling/Disabling auditing on the PTP Monitor display will not change the audit settings on the other pages, and vice versa.

■ To audit Windows machines:

- Install Domain Time Client or Server on the machine.
- Choose to audit the machine either from the the *Domains & Workgroups* or *Domain Time Nodes* list.
- If the machine does not appear in either list, you manually add it to the *Domains & Workgroups* list by right-clicking and choosing **Add computer** from the context menu. You may also add multiple machines using the [Batch Add](#) function.

■ To audit Linux machines:

- If your machine is running an NTP daemon like *ntpd* or *chronyd*, set it to act as an NTP Server (responding to NTP time requests). If your Linux machine is running PTP, see the next section. You can test whether it will respond correctly by using the *ntpcheck* utility from any machine running Domain Time. Open an elevated command prompt and enter:

ntpcheck [name or ip address of your linux machine]

If you see a response like this, you're good to go:

```
C:\WINDOWS\system32>ntpcheck 192.168.1.203
Domain Time NTP Check 5.2
Copyright (c) Greyware Automation Products, Inc.
Hostname: MYHOST

Checking server 192.168.1.203 protocol ntp... okay

Timezone: UTC
```

Server	YYYY:MM:DD	HH:MM:SS.mss	Latency	Secs Delta
192.168.10.203	2019-08-26	17:51:04.145	0.007	+0.0003348

- In the *NTP Nodes* section of Domain Time Manager, right-click and choose "Add NTP Node". Enter the IP

address of your NTP Server. Once added, you can audit the machine just like you audit other machines. You may also add multiple machines using the [Batch Add](#) function.


- On the Audit Server -> Synchronization Logs -> Configure dialog, set your desired NTP drift collection schedule, either on the same schedule you've set for your other sync logs, or you can set a custom schedule. NTP data will then be written to the Synchronization Logs folder on that schedule.

- **To audit PTP devices:**

- Make sure your PTP device or daemon can respond to standard PTP management requests and is visible via both multicast and unicast from the Audit Server.
- Choose to audit the machine from the the *PTP Nodeslist*. Note, Windows machines using Domain Time's PTP should be audited from the the *Domains & Workgroups* or *Domain Time Nodes* list.
- On the Audit Server -> Synchronization Logs -> Configure dialog, set your desired PTP data collection methods.

Read more about how to use [PTP Monitor](#).

Real-time Alerts If you have [Audit Server](#) installed, Manager can show and reset the current status of the Real-Time Alerts notifications.

Click the  **Real-Time Alerts** category in the Tree pane to show the list of machines sending Real-Time Alert notifications to Audit Server. (see the [Audit Server Alerts](#) page for information on configuring Real-Time Alerts.)

The listed machines will show the current synchronization status reported by each machine, as well as the time of the last contact, number of errors since last reset, and last error encountered.

The **Status** column shows the currently-reported state of each machine. Machines that have not yet reported their status will be shown with no color; if the machine has reported with no errors, it will be shown in green.

If a machine has reported an error, the indicator will be red if the error condition still exists, or will be yellow if at least one red error has occurred since the last reset but the error condition has since resolved. The indicator will stay red or yellow until the alert has been cleared manually.

By default, the Real-Time Alerts list will auto-refresh while displayed. You can turn off auto-refresh by unchecking *View -> Auto-Refresh Real-Time Alerts* from the menu.

To clear an alert

Highlight the alert line(s) in the list, right-click and choose *Reset Alert Status* from the context-menu. This will also clear the error count and error message information for the machine. After an alert is cleared, the status indicator will be reset to the initial status and will not change until the next real-time notification is received from the machine.

Discovery

Manager uses multiple methods to discover machines on the network for remote control and monitoring purposes. You can control how Manager performs these discovery tasks using settings contained on three dialog pages from the *Options* item on the main menu.

Click a link to jump to the discussion about each page:

- ▶ [Computer List Enumeration...](#)
- ▶ [Network Discovery...](#)
- ▶ [Scan Options...](#)

Computer List Enumeration...

If you choose *Options -> Network Options -> Computer List Enumeration...* from Manager's menu, you'll see a dialog with the following settings:

Computer List Enumeration

Use Microsoft Networking to find domains and workgroups (can be slow and inaccurate)

Use Active directory to find domains (accurate, but requires Active Directory membership)

LDAP server:	When a computer is removed from Active Directory, Manager should:
Username:	Tombstone the record
Password:	Purge it from the cache

These options control how Manager populates the **Domains and Workgroups** list on the Tree pane of the display when discovering machines from domains, workgroups, and Active Directory.

Enumeration is a two-step process. When either of these options are checked, Manager will first use the selected method(s) to find any existing domains and/or workgroups on the network. Then, if the domain or workgroup is marked to be enumerated in the **Enumerate...** list (see below), Manager will attempt to discover the individual computers present in the selected domain/workgroups.

Use Microsoft Networking to find domains and workgroups... will cause Manager to examine the Windows Browse list for machines. This tends to be a slow process and can only discover machines that have NetBIOS enabled. In general you'll only want to use this method if there are machines on your network that aren't included in Active Directory (such as workgroup members).

In most cases, you will want to choose **Use Active directory to find domains**, which is usually much faster and (assuming your network is correctly configured and maintained) more accurate than by examining the Browse list.


By default, Manager will attempt to access the Active Directory LDAP server used by the domain of the Manager machine, using the credentials of the logged-on user account to connect to it. Manager will attempt to enumerate all listed domains using this server. If you prefer to use a different LDAP server for default enumeration, enter its DNS name or IP address and a user account/password with sufficient privileges to connect to the server.

You can set individual LDAP servers and logon permissions for each domain using the **Enumerate** list below. Select a domain and *Edit* the entry to use the LDAP server and user account you prefer.

Use the **Tombstone the record** and **Purge it from the cache** radio buttons to select how you want Manager to treat listed machines that are later removed from Active Directory.


The Enumerate List

Once domains and workgroups have been identified, they will appear in the **Enumerate computers in the following domains and workgroups** list. You can use the checkboxes to select which of these you want Manager to enumerate when the program starts, or when you manually refresh the **Domains and Workgroups** category in the Tree pane.



Manager - Select domains/workgroups to enumerate [Click for larger size]

You may also add any domain or workgroup that is not discovered automatically



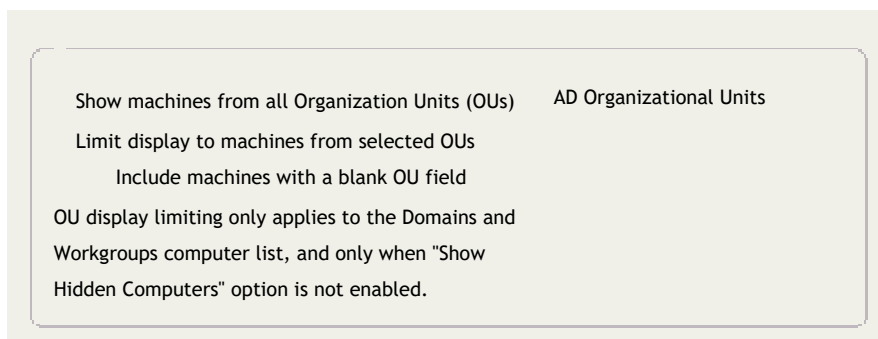
Manager - Add domains to the Enumerate list [Click for larger size]

Enumeration can take quite some time to complete, especially on large networks, so after machines have been enumerated the first time, you may want to only re-enumerate them manually when necessary.

Enumeration can be triggered manually for individual domains by highlighting the domain's name (category) in the Tree pane and right-clicking to choose *Refresh* from the context menu (or *Actions -> Refresh* from the main menu). Enumeration for all domains is initiated by highlighting the **Domains and Workgroups** category and choosing *Refresh*.

Limit display by Organizational Unit (OU)

You may further restrict the view of machines by selecting which OU's you'd prefer to display



Network Discovery...

Select *Options* -> *Network Options* -> *Network Discovery...* from Manager's menu and the Network Discovery dialog will appear:

The screenshot shows the 'Network Discovery' dialog box. It is divided into four main sections:

- IPv4 Multicast Discovery:** The 'Enabled' checkbox is checked. Two radio buttons are present: 'All local subnets (more thorough)' and 'Primary subnet only (faster)', with the latter selected. Below these are input fields for 'DT2 multicast address' (239 . 192 . 99 . 9), 'NTP multicast address' (224 . 0 . 1 . 1), and 'IPv4 TTL' (1).
- IPv6 Multicast Discovery:** The 'Enabled' checkbox is unchecked. It has similar radio buttons and input fields for 'DT2 multicast address' (FF05::9909), 'NTP multicast address' (FF05::101), and 'IPv6 Hops' (1).
- IPv4 Broadcast Discovery:** The 'Enabled' checkbox is checked. Three radio buttons are present: 'List of custom addresses', 'All local subnets (more thorough)', and 'Primary subnet only (faster)', with the middle one selected. There is a 'Custom broadcast addresses' list box containing '255.255.255.255' and an input field showing '255 . 255 . 255 . 255'.
- UDP and ICMP Timeouts:** 'Send' is set to 5 ms (range 1 to 5000) and 'Receive' is set to 500 ms (range 1 to 15000). There is a 'Presets' dropdown menu set to '[Choose]' and an 'Apply' button.

At the bottom, there is a 'Reset to Defaults' button on the left and 'OK' and 'Cancel' buttons on the right.

The Network Discovery Dialog [\[Click for larger size\]](#)

This dialog controls how Manager discovers currently running Domain Time components and/or NTP daemons on the local subnet and beyond.

When enabled, Manager broadcasts and/or multicasts a special discovery packet to the network. Any machine that hears the discovery packet will then respond back via a unicast UDP packet containing its synchronization status and other data. Broadcasts/Multicasts are sent to the networks indicated on this dialog, using the protocol types selected.

This method only discovers machines that are online and that respond to the discovery packet. Once machines are discovered, they are added to the Manager's cache list and are thereafter contacted using unicast for most operations.

IPv4/IPv6 Multicast Discovery

Domain Time can use multicast to discover Domain Time machines currently online over IPv4 and/or IPv6, both on the local subnet and on remote subnets. This method is preferred over the older IPv4 Broadcast Discovery method described below.

In order to use this method, your network must be enabled for multicast. You must have multicast-enabled routers/switches configured to allow multicasts sent to port 9909 UDP to reach all of your subnets. This may also require enabling PIM (Protocol-Independent Multicasts) on all relevant interfaces. Consult your network equipment vendor(s) for configuration details necessary for your environment.

The image shows two panels side-by-side, each representing the configuration for a specific protocol:

- IPv4 Multicast Discovery:** Contains an 'Enabled' checkbox, two radio buttons ('All local subnets (more thorough)' and 'Primary subnet only (faster)'), and three input fields labeled 'DT2 multicast address:', 'NTP multicast address:', and 'IPv4 TTL:'.
- IPv6 Multicast Discovery:** Contains an 'Enabled' checkbox, two radio buttons ('All local subnets (more thorough)' and 'Primary subnet only (faster)'), and three input fields labeled 'DT2 multicast address:', 'NTP multicast address:', and 'IPv6 Hops:'.

All local subnets (more thorough)
Primary subnet only (faster)

These options select whether the multicast discovery packet is sent out only over the interface the operating system has designated as the primary interface or through all network interfaces, or whether it is sent to specified remote subnets.

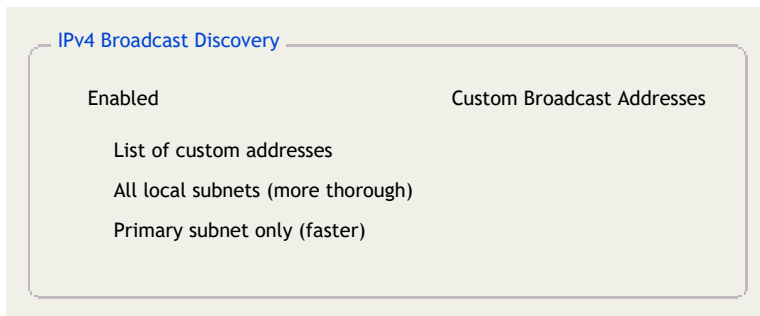
Sending through multiple interfaces also results in longer timeouts and processing overhead to send traffic and listen for responses on each interface. In general, you will want to send only through the primary interface unless you have machines that cannot be discovered otherwise. Also, If you have multi-homed machines visible to the Manager on multiple interfaces, sending discovery packets through all interfaces may result in duplicate responses from those machines.

The **DT2** and **NTP multicast address** boxes indicate the multicast address(es) Manager will send the discovery packet to. Only change these addresses if you have a compelling reason to do so. Changing these values is usually an error.

The **IPv4 TTL:** and **IPv6 Hops:** entries set how many router hops a multicast packet must cross when propagating through your network. Choose a value that allows multicasts to reach all of your subnets.

IPv4 Broadcast Discovery

By default, Manager only broadcasts to the local segment. This is an efficient way to discover machines local on local subnets, and you probably do not want to disable this function entirely. However, using IPv4 Broadcasts to discover machines on remote subnets is recommended only if multicast discovery is not possible or reliable on your network.



List of custom addresses

All local subnets (more thorough)

Primary subnet only (faster)

These options select whether the broadcast discovery packet is sent out only over the interface the operating system has designated as the primary interface, through all local network interfaces, or whether it is sent to specified subnets.

Sending through multiple interfaces also results in longer timeouts and processing overhead to send traffic and listen for responses on each interface. In general, you will want to send only through the primary interface unless you have machines that cannot be discovered otherwise. Also, If you have multi-homed machines visible to the Manager on multiple interfaces, sending discovery packets through all interfaces may result in duplicate responses from those machines.

The **255. 255. 255. 255** address is used to broadcast to the primary segment. Only change this address if you have a compelling reason to do so. Changing this value is usually an error.

Custom Broadcast Addresses

To discover machines on remote subnets, your routers/switches must be configured to allow broadcast traffic using port 9909 UDP to reach each subnet. You must also select the **List of custom addresses** radio button and specify the broadcast address for each subnet. Manager will then send directed broadcasts to the broadcast addresses specified.

This option is provided primarily for backwards-compatibility with versions prior to version 5.1. This method does not scale well and is unreliable due to firewall/router issues with passing broadcasts between subnets. We recommend you use multicasts to reach remote subnets instead, if possible (see above).

If you do choose to use this option, the **255. 255. 255. 255** entry is used to broadcast to the primary segment. You should leave this entry in the list unless you have a specific reason to remove it. Then, enter the broadcast address for any other remote subnets you want to scan via IPv4 broadcast. For example, you would enter **192. 168. 1. 255** as the broadcast address for the 192.168.1.x subnet.

Note about Broadcast Traffic across firewalls

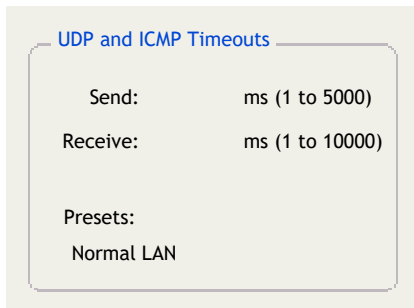
Broadcast traffic typically originates from an ephemeral source port sent to the broadcast address of designated subnets intended for destination port 9909. Domain Time II components that receive the broadcast then respond via unicast from source port 9909 back to the sending IP address's ephemeral port. This broadcast discovery process is substantially the same as the process used by DHCP, TFTP, and other standard broadcast discovery methods.

Stateful firewalls/routers will often require additional configuration to ensure broadcast discovery operates correctly, since unlike normal unicast UDP communication, the originating traffic is not sent to the same IP address from which the reply will come. Broadcast traffic is sent to the address xxx.xxx.xxx.255, but the unicast replies from that subnet may come from any (or all) addresses in the range xxx.xxx.xxx.1-254.

Normal stateful firewall rules typically only open the firewall for replies from the same IP address to which the originating traffic was sent, so even if unicast port 9909 UDP traffic is enabled and working, broadcast traffic may still fail. Therefore, most firewalls have special rules that can be applied to allow broadcasts to function correctly, (such as **ip helper-address**, **ip directed-broadcast** and **ip forward-port** functions on Cisco equipment, for example). Check with your firewall manufacturer for the correct broadcast address configuration instructions for your particular systems.

UDP & ICMP Timeouts

UDP packets are subject to numerous possible delays, particularly on busy networks or across slower links. This can cause some machines to not respond to scans reliably. You may be able to compensate for some of these issues by increasing the UDP/ICMP Send and Receive values.



UDP and ICMP Timeouts

Send:	ms (1 to 5000)
Receive:	ms (1 to 10000)
Presets:	Normal LAN


If you have configured broadcasts and multicasts correctly, and you know for certain that your routers, firewalls, and switches are configured correctly to pass the traffic, and yet you are not seeing all of your remote machines, you may need to increase your timeout settings.

The **Presets:** [Choose] drop-down list allows you to quickly pick recommended values for various network configurations. You can use these values as suggested starting points to adjust your timeouts.

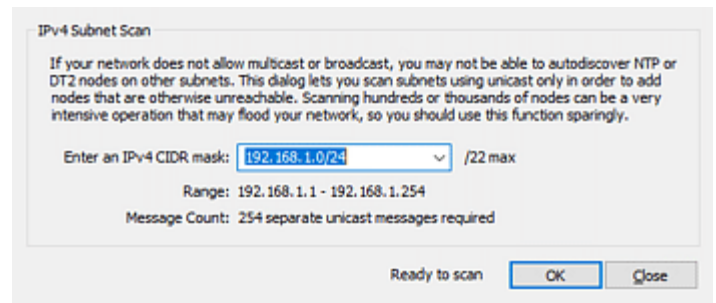
For troubleshooting, you may want to temporarily pick the longest timeout preset (Satellite Link) to eliminate timeouts as a connection issue. However, in production, set timeouts only long enough to receive responses from all of your systems. Setting too large of a timeout will result in scans taking an excessively long time to complete.

Manually Scan IP4 Subnet

As of v5.2.b.20190701, you may manually scan a selected IPv4 subnet to find machines that might not otherwise answer a multicast or broadcast discovery. You launch this by right-clicking on a blank part of either the Domain Time Nodes or NTP Nodes list (right-side) and choosing **Scan IPv4 Subnet** from the context menu.



Choose Scan IPv4 Subnet from the context menu [Click for larger size]

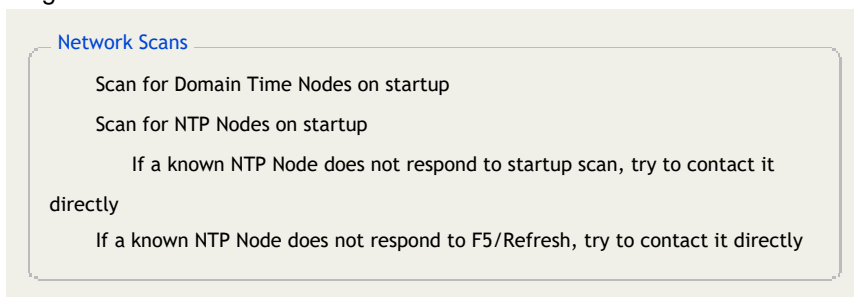


IPv4 Subnet Scan Dialog [Click for larger size]

Caution: Scanning subnets can generate a large amount of traffic, take a significant amount of time, and may cause issues with flooding or violate your network rules. Use with care. For this reason, the largest subnet scan allowed is a /22.

Scan Options...

If you choose *Options -> Network Options -> Scan Options...* from Manager's menu, you'll see a dialog with the following settings:



These items control how Manager handles discovery scanning.

If desired, Manager can transmit a special discovery packet to the network (using Multicast and/or IPv4 Broadcast) to locate existing Domain Time components or NTP daemons. Any machine that hears the discovery packet will then respond back to Manager via a unicast UDP packet containing its synchronization status and other data. Machines responding to these discovery queries will populate the **Domain Time Nodes** and **NTP Nodes** lists on the Manager Tree pane.

You can choose have Manager run the discovery scan each time it starts. Scans can take quite some time to complete, especially on large networks, so you may want to run them manually only when needed. Scans can be triggered manually by highlighting the **Domain Time Nodes** or **NTP Nodes** category in the Tree pane and right-clicking to choose *Refresh* from the context menu (or *Actions -> Refresh* from the main menu).

The **If a known NTP Node does not respond to _____, try to contact it directly** selections tell Manager to try a unicast time request to determine the status of an NTP machine if the original discovery attempt fails. This results in more reliable scans, at the cost of slightly more network traffic.

Using Templates

Domain Time II Manager uses configuration template files for remote installation, upgrade, and configuration resets. Template files contain Domain Time Server or Client registry settings in standard Windows Registry export file format (.reg).

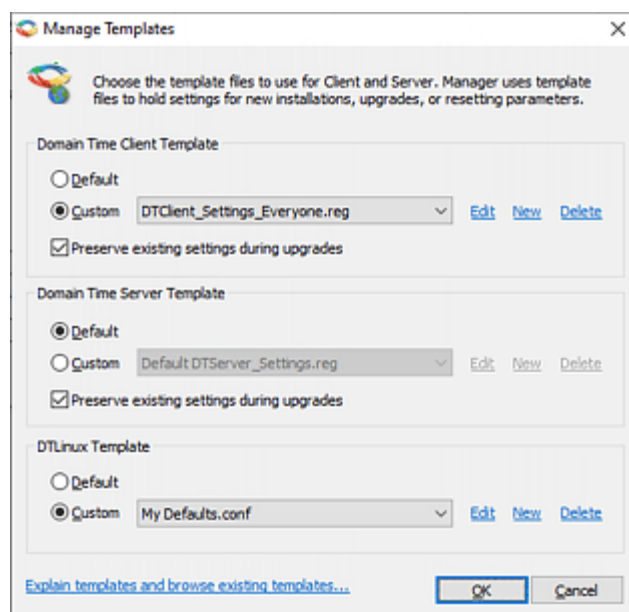
You have the option of using the default configuration templates or custom templates whenever you using Domain Time Manager to install, upgrade, or reset configurations.

Notes:

- ▶ The manufacturer's default configuration templates (dtserver.reg and dtclient.reg) are provided as part of the distribution files for each flavor of Domain Time Server and client. We recommend that you do not edit these files directly. Use Custom templates to make any configuration changes you require instead.
- ▶ The DOMTIME.INI template files used for this function in version 4.1 and earlier have been deprecated.

Managing Templates

Click *Options -> Manage Templates...* from Manager's menu to work with templates.



The Choose Templates dialog [Click for larger size]

The **Manage Templates** dialog allows you to select which template to use for installs, upgrades, or resets for both Client and Server. It will remember your choices so you can easily perform the same operation repeatedly without having to re-select your templates. You'll also be able to change these settings at any time while performing installations, upgrades, or resets.

Default tells Manager to use the manufacturer's default template included in the distribution files.

Custom tells Manager to use the selected custom template. Any custom templates you've created will be available in the drop-down list.

You can also create a new template from scratch (see below), edit an existing template, or delete custom templates from this dialog.

The **Preserve existing settings during upgrades** checkboxes indicate whether you want upgraded machines to

continue using the settings they had prior to the upgrade, or to overwrite their configuration with the settings contained in the selected template.

Creating Custom Templates

You can create a custom template either by creating one from scratch or by exporting one from a sample machine. It is usually easier to export settings from a sample machine unless you are familiar with all of the Domain Time registry settings.

Custom template files are located in the **C: \Program Files\Domain Time II\Templates\[Server][Client]** folder of the Domain Time II Manager machine. Template .reg files located in those folders will automatically be made available for use when installing or upgrading using Manager.

■ From a sample machine

Creating a custom template from a sample machine is a simple process. Follow these steps:

- Pick a sample machine to use as your configuration machine.
- [Install Server or Client](#) on the sample machine using the installation defaults.
- [Connect to](#) the sample machine's Domain Time Control Panel remotely using Manager.
- Configure the sample machine exactly as you want your target machines to be configured. Double-check all settings on all property pages.
- While still connected remotely, change to the *Advanced -> Import/Export* property page of the Control Panel applet.

Since you're connected remotely, the export utility will automatically offer to save the .reg file in the proper directory on the Manager machine. Note that it also suggests a filename that indicates the name of the sample machine and when the file was created. You may accept this name, or edit it if you prefer. Be sure to keep the .reg extension.

If you configure a machine while not connected remotely, you can save the export file locally, then manually copy the file to the Templates folder on the Manager machine.

IMPORTANT: Although .reg files created using this utility are saved in standard Windows registry REGEDIT4 file format, it is **not** equivalent to exporting the registry keys using Windows' Registry Editor program. Registry file formats other than REGEDIT4 may not import correctly. Also, a number of registry settings on a running Domain Time system are machine-specific, and are likely to cause problems if directly copied into another machine's registry. Those settings are automatically excluded when you export using this utility, so you should always use the **Import/Export** utility to create a Domain Time .reg file.

- Click the **Save** button to save the template.

Note: Exported templates may be edited manually using any text editor (or by clicking the [Edit](#) link on the **Choose Templates** dialog. You may allow the template to contain all exported settings, or edit it to contain only the specific settings you want to apply to the target machine(s).

Settings you omit from the template will either be set to the manufacturer's defaults if they do not already exist on the target (such as during initial installation), or left untouched if they do exist (during upgrades/resets).

■ From scratch

You can create a template file manually by creating a copy of the default settings template and editing it with any text editor.

Custom

[Edit](#) [New](#) [Delete](#)

Click *Options -> Manage Templates...* from Manager's menu to bring up the **Choose Templates** dialog (see above). Then, click the [New](#) link to create a copy of the manufacturer's default template (for either Client or Server).

You can then edit the registry settings as desired. You may allow the template to contain all existing settings, or trim it to contain only the specific settings you want to apply to the target machine(s).

Settings you omit from the template will either be set to the manufacturer's defaults if they do not already exist on the target (such as during initial installation), or left untouched if they do exist (during upgrades/resets).

Note: Some settings (such as Server or Client Timings) are kept in binary keys which are not easily edited manually. You will need to configure those items using the Control Panel applet and export them into a .reg file (see the Creating Custom Templates **From a sample machine** section above).

Domain Time II Monitor Service

Version 5.2

Domain Time II Monitor Service is a system service that periodically collects variance statistics from Domain Time Servers and Clients across your entire network. It also raises alerts if time variance on any running machine that exceeds the tolerances you specify.

Note: The Monitor Service is intended only for simple "snapshot" monitoring tasks. Use [Domain Time II Audit Server](#), to provide more capable monitoring and record-keeping. Audit Server adds auditing of NTP and Windows Time systems (using the Domain Time Windows Time Agent), more robust automatic discovery, notifications of non-responding machines, custom reports, SNMP and other real-time alerts, etc.

IMPORTANT: If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[Installation](#)

[System Requirements](#)

Additional Requirements


Verify that your environment meets the minimum requirements for performing remote monitoring using Domain Time Monitor.

- Since variance reporting and alerting depends on accurate time calculations, Monitor Service should always be run on a physical (not virtual) machine.
- To monitor machines on your local subnet, your local network must allow UDP broadcast traffic sent to destination port 9909.
- To monitor machines on remote subnets, your network must pass UDP multicast traffic sent to destination port 9909. Switches and firewalls must also pass UDP unicast traffic sent to destination port 9909 bi-directionally between subnets, since unicast traffic will originate either from Monitor or the remote machines.

Installation

The Monitor Service is installed and its applet launched from [Domain Time II Manager](#). Click *Utilities --> Monitor Service* on the Manager's menu. If Monitor Service is not installed, you'll be prompted to install it.

Once the Monitor service has been installed, you configure it using its Control Panel applet. You can launch the applet from Manager (see above), from the Control Panel, or you may also launch the Domain Time II Monitor Service program (and many other installed Domain Time II components) by right-clicking on the Domain Time icon in the System Tray to bring up the context menu.

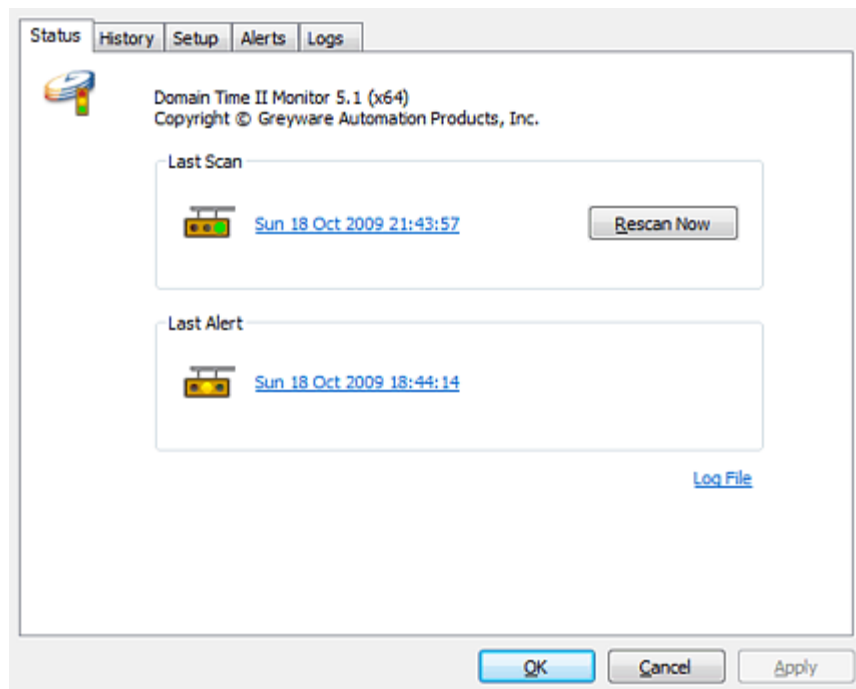


Note: On systems with User Account Control (UAC) enabled, you may need to *Shift+Right Click* and choose **Run As...** from the context menu to launch the Control Panel applet as Administrator. On Windows Server Core, type in `domt mem. cpl` on the command line.

The Control Panel Applet - Status tab

The first tab page of the applet shows the time/date and status of the last variance scan and last alert raised.

If this is the first time the applet has been launched, it will run a scan automatically. You can also trigger a Rescan from this tab page.

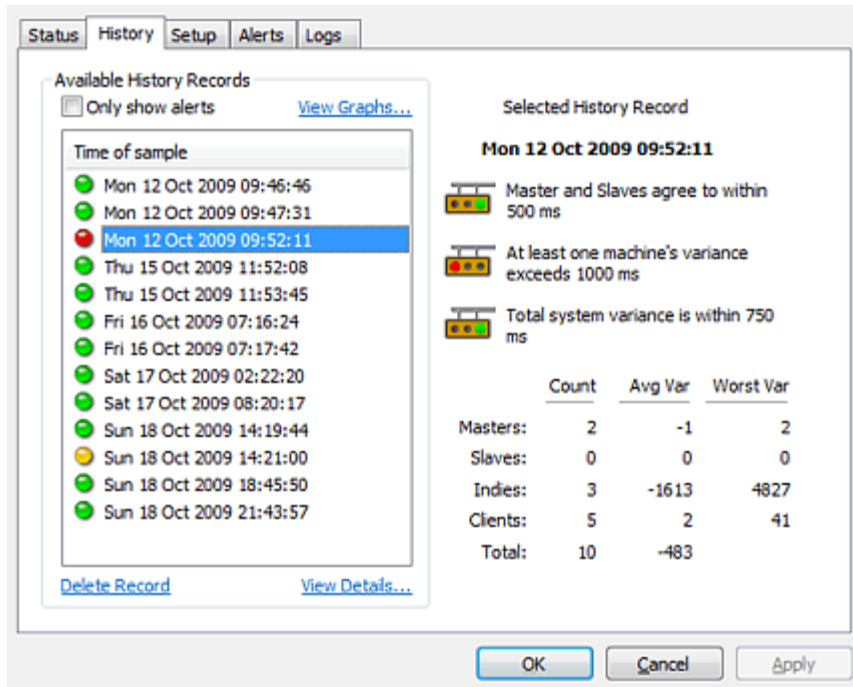


Domain Time II Monitor - Status tab [Click for larger size]

Click the link in each section to see the detail summary report from either the last scan or the most recent scan that resulted in an error. You can also view the service log by clicking the *Log File* link.

History

Use this page to view and manage the Monitor Service's scan history.



Domain Time II Monitor - History tab [Click for larger size]

Available History Records

Scan reports saved to disk are displayed in this window. Highlight a scan item in the list to see its *Selected History Record* summary on the right-hand panel. You may also double-click the list item (or click the *View Details...* link) to pull up that scan's detail report.

The status indicator lamp on each list item shows the highest severity level found during that scan. Note that this indicator may show a problem even if a specific alert threshold was not crossed. For example, you would see a severity indicator of yellow if the Monitor Service wasn't able to contact its reference clock, even if all of the scanned machines were within the selected range of variance.


The **Selected History Record** traffic lights on the right are a graphical representation of the alarm condition thresholds in place *when the selected scan was run*. If any machine exceeded the indicated threshold, the lights will indicate which alert was raised.

The average and worst variances for each type of machine contacted during the scan are also displayed.

Use the *Delete Record* link to remove the selected scan record.

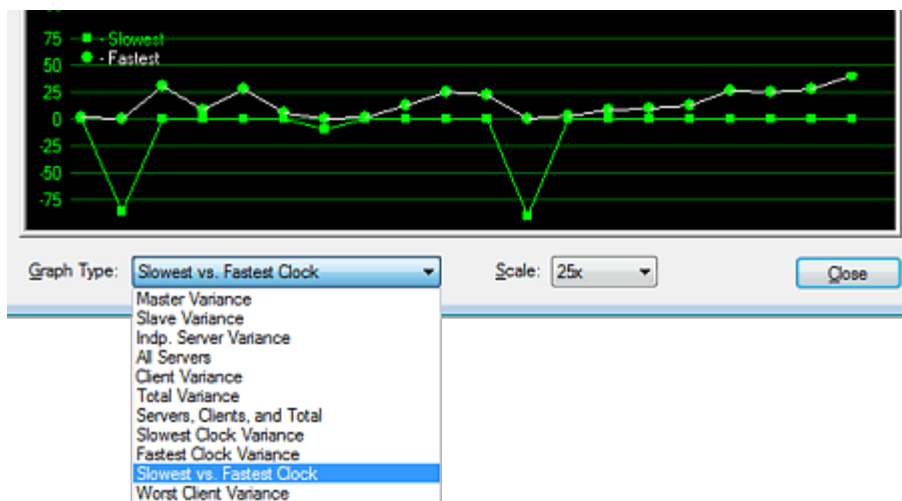
View Graphs...

Clicking this link will bring up a graphical display summarized from all of your collected data. You must have collected at least two scans to be able to generate the graphs.



Domain Time II Monitor - Sample Graph [\[Click for larger size\]](#)

You can select from a wide variety of graph types that let you see at a glance how well your network is being synchronized. You can also adjust the scale of the graph to your desired level of accuracy.



Domain Time II Monitor - Graph Types [\[Click for larger size\]](#)

Domain Time II Software distributed by [Microsemi, Inc.](#)

Documentation copyright © 1995-2021 Greyware Automation Products, Inc.

All Rights Reserved

All Trademarks mentioned are the properties of their respective owners.

Use the Setup page to configure important Domain Time II Monitor Service settings.

[Schedule](#)

Monitor should check the network

Every Five Minutes

Send email summaries

Every Day

Include all detail records

Include only error details

No detail records

The **Schedule** section controls which functions Monitor performs on a regular basis, as well as notification email configuration.

The **Monitor should check the network** Every Five Minutes drop-down list sets how often the service will run a network scan.* Choose a schedule that gives you the level of surveillance of your Domain Time machines you desire.

When a scan is run, Monitor sends a discovery packet to the network (see the **Discovery** section below). Any Domain Time Server or Client that hears the discovery packet responds to the Monitor machine with its current sync information.

Each time the scan is run, the results are analyzed to be sure all of the responding machines meet your alert criteria (see the [Alerts](#) tab page to set your alert definitions).

Note: Only machines that are running at the time of the scan will be included in the Monitor scan results. In addition, transient network issues may prevent UDP packet responses from being received reliably.

Domain Time [Audit Server](#) is a much more robust solution, and offers many additional monitoring and alerting functions than does Monitor Service. This includes persistent contact lists (status of both on- and off-line machines), monitoring of NTP daemons, real-time alarms, more complete historical record keeping, etc. You should only use Monitor for quick snapshots and "heads-up" alerting. Use Audit Server for more critical tasks.

* You may run a scan manually by clicking the **Rescan Now** button on the [Status](#) tab page.

Send email summaries Every Day

Include all detail records
Include only error details
No detail records

Use these controls to determine whether and how often to receive summary email reports of the Monitor Service scan activity. You can also select which level of detail you want included in the email. See the sample report below for an example of what's included in the summary section and in the optional detail records.

Domain Time II Monitor Summary Report

Report Date: Sun 18 Oct 2009 00:00:39

Scan machine: FLEEGMAN
Email Alerts: Enabled
Event Viewer: Enabled
First Scan: Sat 17 Oct 2009 02:22:20
Last Scan: Sat 17 Oct 2009 20:20:17
Period Covered: 17 hours, 57 minutes
Detail Records: All (see below)
Total Scans: 1
Scan Errors: 0
Tests Passed: 9 of 9
Tests Failed: 0 of 9
Average Delta: - 00 000 00:00:00.071

Detail Record 1 of 1

Scan Status: **Passed**
Scan Time: Sat 17 Oct 2009 07:22:20 UTC
Sat 17 Oct 2009 02:22:20 (local)
Reference Clock Type: This machine's sources (averaged)
Reference Clock Time: Sat Oct 17 2009 02:22:19
Reference Source: server tick.greyscale.com protocol DT2-UDP samples 3
Reference Offset: -0.2964544 seconds

Type and average delta yy ddd hh:mm:ss.mss

Clients: 3 - 00 000 00:00:00.119
Total: 3 - 00 000 00:00:00.119

Slowest Clock: Sat Oct 17 2009 07:22:19.972 UTC 192.168.10.2 (OMEGA13)
Fastest Clock: Sat Oct 17 2009 07:22:19.972 UTC 192.168.10.4 (Grabthar)

Worst Client: + 00 000 00:00:00.371 192.168.10.2 (OMEGA13)

Master/Slave Test: Limit: 500 ms Slave Var: + 00 000 00:00:00.000 **Test Passed**
Single Machine Test: Limit: 1000 ms Worst Var: + 00 000 00:00:00.371 **Test Passed**
Avg Variance Test: Limit: 750 ms Total Var: - 00 000 00:00:00.119 **Test Passed**

INDIVIDUAL VARIANCES (sorted by variance)

- 00 000 00:00:00.371 Full Client 192.168.10.2 (OMEGA13)
+ 00 000 00:00:00.000 Full Client 192.168.10.28 (HV-2003-R2-X64)
+ 00 000 00:00:00.013 Full Client 192.168.10.4 (Grabthar)

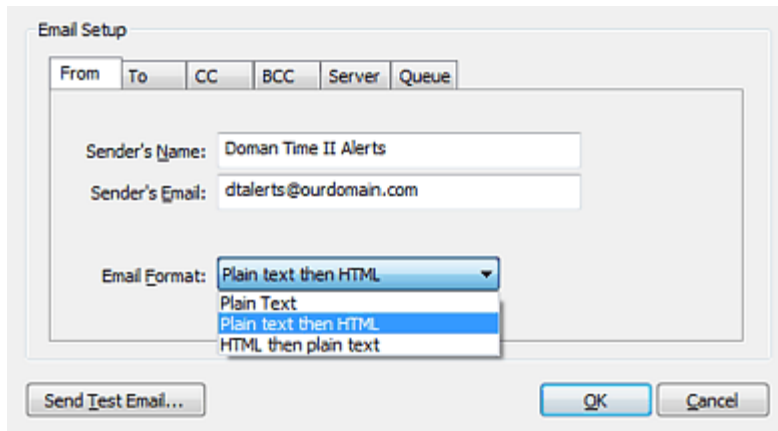
End of Summary Report

Sample Monitor Summary Report Email (reduced size)

Use _____ to configure SMTP servers, delivery addresses, and manage the notification email queue.

You must configure these email settings before Monitor Service can send notification emails.

Set the From address



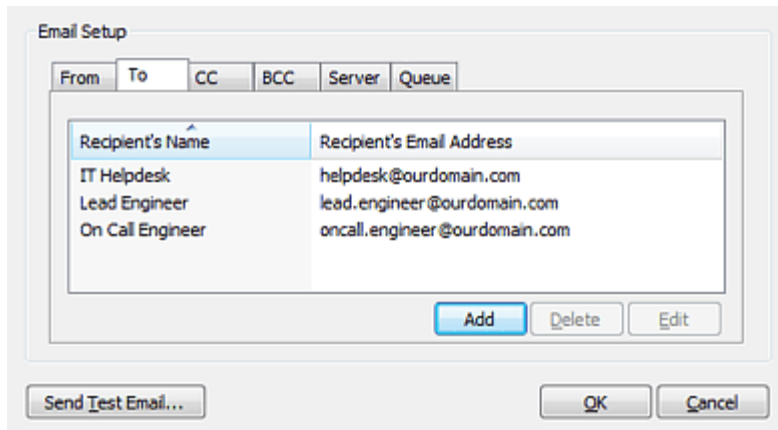
[Email Setup From and Format Selection](#) [Click for larger size]

Specify the From: email address that will appear on the notification emails. You can also specify the format and MIME part order of the emails:

- Plain Text
- Text part followed by HTML part
- HTML part followed by Text part

Choose the format that provides the best compatibility with your email system.

Set the TO/CC/BCC distribution lists



[Email Recipients List](#) [Click for larger size]

Use the **To**, **CC**, and **BCC** tabs to add the email addresses of your desired recipients.

Set the Outgoing Server

The 'Email Setup' dialog box has tabs for 'From', 'To', 'CC', 'BCC', 'Server', and 'Queue'. The 'Server' tab is selected. It contains the following fields:

- SMTP Server:** smtp.ourdomain.com
- Server Port:** 587 (with a dropdown menu set to 'STARTTLS if supported')
- Auth Username:** audit (optional)
- Auth Password:** [masked with dots] (optional)

At the bottom, there are three buttons: 'Send Test Email...', 'OK', and 'Cancel'.

Outgoing SMTP Server Settings [Click for larger size]

Enter the server address and account login information required for Monitor Service to send outgoing mail through your SMTP server.

Send Test Email

Once you have entered all of the above information, click the **Send Test Email** button to generate a test email.

If your test email encounters any errors, you'll receive a pop-up window showing the entire SMTP conversation so you can easily troubleshoot the problem:



The email could not be delivered. The system error code was error 8225:
A protocol error occurred,
and the last SMTP result code was 535.


A log of the SMTP conversation follows:

```
I: Looking up smtp.greypware.com
I: Connected to 71.252.193.51:25
R: 220-smtp.sff.net
R: 220-Greyware Mailman 1.5.b.20090107R
R: 220 ready
S: EHLO Fleegman
R: 250-smtp.sff.net hello Fleegman
(pool-173-74-57-68.dl1stb.fios.verizon.net [173.74.57.68])
R: 250-SIZE 15728640
R: 250-AUTH=LOGIN PLAIN
R: 250-AUTH LOGIN PLAIN
R: 250 HELP
S: AUTH PLAIN
R: 334 send authentication
S: SVRIZVxxwZGVzawBJVEhIbHBkZXNrAHNwYWFjZQ==
R: 535 <ITHelpdesk> mailbox does not exist
E: Protocol error: expected 235 but got 535
```

OK

Send Test Email, Showing SMTP Error [Click for larger size]

Check the email queue to troubleshoot delivery issues



Email Queue Settings and Email Logs [Click for larger size]

This page contains the settings for the email queue and email logs.

The **Queue Folder**: specifies the location of the folder where outgoing emails are queued. The **email.log** trace file is also kept in this folder.

Note: In most cases, you will not need to adjust this location. If you do decide to change the folder location, you must pick a location on a local disk (not a networked share) with sufficient permissions (Change) granted to the Monitor Service service account so that it is able to manage the queues.

Use the **SMTP Trace**: Disabled drop-down list to select the level of detail you want to keep in the **email.log** trace file. In general, you should only enable the trace file if you are troubleshooting an email delivery issue. Otherwise, your **email.log** file may grow to an unmanageable size over time.

Use the **Open** or **Browse** buttons to open the queue folder and locate the **email.log** file, which is a plain text file you can open in any editor, such as Notepad.

Advanced Configuration: Email-related registry settings

Depending on your email server configuration, you may also need to adjust these additional settings in the Windows registry.

Email registry settings are located in the **HKEY_CLASSES_ROOT\Gap\GWServiceSMTP** key.

TLSIgnoreCertErrors (REG_DWORD)

Introduced in v5.2.b.20140922 with default=0 (ignore no errors). As of v5.2.b.20160711, the default changed to 0x311 (accept certs that are self-signed, expired, or have the wrong CN)

If this value is zero, the server cert must pass all tests. If the value is non-zero, it is a bitmask specifying which particular types of errors may be ignored. See [Microsoft's documentation](#) for a list of certificate errors that may be ignored. Use a logical **OR** to combine multiple values.

- **0x00000080** - Ignore errors associated with certificate revocation
- **0x00000100** - Ignore errors associated with an unknown (or self-signed) certificate authority
- **0x00000200** - Ignore errors associated with wrong use of a certificate
- **0x00001000** - Ignore errors associated with an invalid/mismatched common name
- **0x00002000** - Ignore errors associated with an expired certificate

You may set the value to 0x10000000 in order to regain strict certificate checking, 0x0000FFFF to disable certificate checking altogether, or any combination of the above values.

TLSAcceptableProtocol s (REG_DWORD)

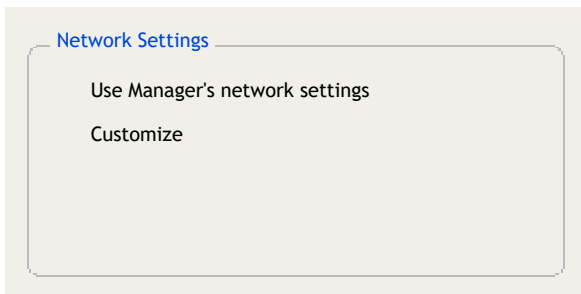
Introduced in v5.2.b.20160711. This is a bitmask of acceptable encryption protocols. The default value is 0x0AA0. Use a logical **OR** to combine multiple values.

- 0x00000002 - PTC1 (not recommended)
- 0x00000008 - SSL2 (not recommended)
- 0x00000020 - SSL3 (not recommended, but included in default for backward compatibility)
- 0x00000080 - TLS 1.0 (not recommended, but included in default for backward compatibility)
- 0x00000200 - TLS 1.1
- 0x00000800 - TLS 1.2
- 0xFFFFFFFF - any available protocol (not recommended)

FQDN (REG_SZ)

Introduced in v5.2.b.20160711. This value contains the name to use during SMTP envelope negotiations; specifically, it is the name presented as the HELO or EHLO name immediately after receiving the server's greeting.

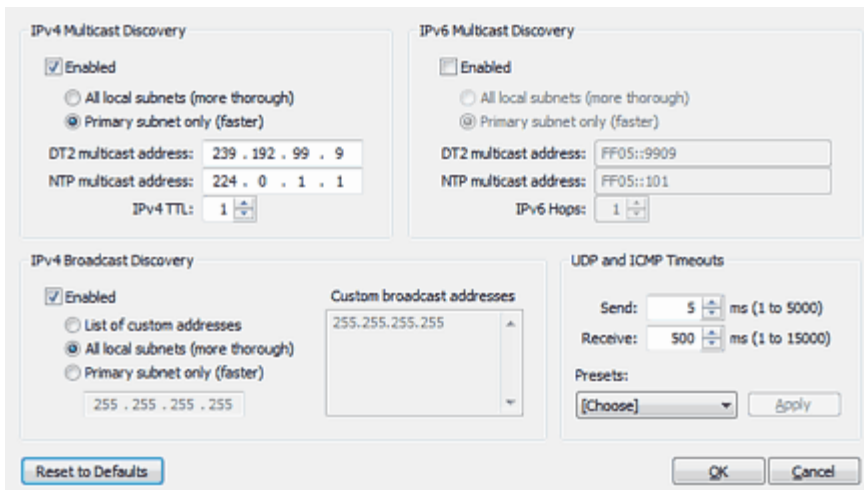
In previous versions, the name used was the sending machine's fully-qualified host name. However, workgroup members or machines just starting may only have a bald hostname available. This new value is set the first time an email is sent, and used thereafter for all subsequent emails. If a fully-qualified name is not discoverable, then Domain Time will use either a dotted-quad IP enclosed in brackets, or the computer name followed by .smtp.local. RFC 2821 section 4.1.1.1 requires one of these two forms. You may change the name if your particular email server requires an externally-verifiable DNS name to be presented.



The **Network Settings** section allows you to choose whether or not Monitor should use Manager's settings for [Network Discovery](#) and [Reference Clock](#) or to set these items independently.

Network Discovery

If you choose **Customize**, you can pull up Monitor's *Network Discovery* dialog by clicking the [button](#).



The Network Discovery Dialog [\[Click for larger size\]](#)

This dialog controls how Monitor Service discovers currently running Domain Time components on the local subnet and

beyond.

When enabled, Monitor Service broadcasts and/or multicasts a special discovery packet to the network. Any machine that hears the discovery packet will then respond back via a unicast UDP packet containing its synchronization status and other data. Broadcasts/Multicasts are sent to the networks indicated on this dialog, using the protocol types selected.

This method only discovers machines that are online and that respond to the discovery packet.

Note: Although NTP broadcast/multicast addresses are listed on this dialog, the Monitor Service only uses the DT2 protocol settings and only discovers Domain Time machines.

IPv4/IPv6 Multicast Discovery

Domain Time can use multicast to discover Domain Time machines currently online over IPv4 and/or IPv6, both on the local subnet and on remote subnets. This method is preferred over the older IPv4 Broadcast Discovery method described below.

In order to use this method, your network must be enabled for multicast. You must have multicast-enabled routers/switches configured to allow multicasts sent to port 9909 UDP to reach all of your subnets. This may also require enabling PIM (Protocol-Independent Multicasts) on all relevant interfaces. Consult your network equipment vendor(s) for configuration details necessary for your environment.

IPv4 Multicast Discovery	IPv6 Multicast Discovery
Enabled	Enabled
All local subnets (more thorough)	All local subnets (more thorough)
Primary subnet only (faster)	Primary subnet only (faster)
DT2 multicast address:	DT2 multicast address:
NTP multicast address:	NTP multicast address:
IPv4 TTL:	IPv6 Hops:

All local subnets (more thorough)
Primary subnet only (faster)

These options select whether the multicast discovery packet is sent out only over the interface the operating system has designated as the primary interface or through all network interfaces, or whether it is sent to specified remote subnets.

Sending through multiple interfaces also results in longer timeouts and processing overhead to send traffic and listen for responses on each interface. In general, you will want to send only through the primary interface unless you have machines that cannot be discovered otherwise. Also, If you have multi-homed machines visible to the Monitor Service on multiple interfaces, sending discovery packets through all interfaces may result in duplicate responses from those machines.

The **DT2** and **NTP multicast address** boxes indicate the multicast address(es) Monitor Service will send the discovery packet to. Only change these addresses if you have a compelling reason to do so. Changing these values is usually an error.

The **IPv4 TTL:** and **IPv6 Hops:** entries set how many router hops a multicast packet must cross when propagating through your network. Choose a value that allows multicasts to reach all of your subnets.

IPv4 Broadcast Discovery

By default, Monitor Service only broadcasts to the local segment. This is an efficient way to discover machines local on local subnets, and you probably do not want to disable this function entirely. However, using IPv4 Broadcasts to discover machines on remote subnets is recommended only if multicast discovery is not possible or reliable on your network.

IPv4 Broadcast Discovery

Enabled

Custom Broadcast Addresses

List of custom addresses

All local subnets (more thorough)

Primary subnet only (faster)

List of custom addresses

All local subnets (more thorough)

Primary subnet only (faster)

These options select whether the broadcast discovery packet is sent out only over the interface the operating system has designated as the primary interface, through all local network interfaces, or whether it is sent to specified subnets.

Sending through multiple interfaces also results in longer timeouts and processing overhead to send traffic and listen for responses on each interface. In general, you will want to send only through the primary interface unless you have machines that cannot be discovered otherwise. Also, If you have multi-homed machines visible to the Monitor Service on multiple interfaces, sending discovery packets through all interfaces may result in duplicate responses from those machines.

The **255. 255. 255. 255** address is used to broadcast to the primary segment. Only change this address if you have a compelling reason to do so. Changing this value is usually an error.

Custom Broadcast Addresses

To discover machines on remote subnets, your routers/switches must be configured to allow broadcast traffic using port 9909 UDP to reach each subnet. You must also select the **List of custom addresses** radio button and specify the broadcast address for each subnet. Monitor Service will then send directed broadcasts to the broadcast addresses specified.

This option is provided primarily for backwards-compatibility with versions prior to version 5.1. This method does not scale well and is unreliable due to firewall/router issues with passing broadcasts between subnets. We recommend you use multicasts to reach remote subnets instead, if possible (see above).

If you do choose to use this option, the **255. 255. 255. 255** entry is used to broadcast to the primary segment. You should leave this entry in the list unless you have a specific reason to remove it. Then, enter the broadcast address for any other remote subnets you want to scan via IPv4 broadcast. For example, you would enter **192. 168. 1. 255** as the broadcast address for the 192.168.1.x subnet.

Note about Broadcast Traffic across firewalls

Broadcast traffic typically originates from an ephemeral source port sent to the broadcast address of designated subnets intended for destination port 9909. Domain Time II components that receive the broadcast then respond via unicast from source port 9909 back to the sending IP address's ephemeral port. This broadcast discovery process is substantially the same as the process used by DHCP, TFTP, and other standard broadcast discovery methods.

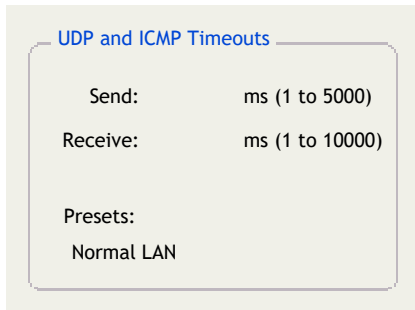
Stateful firewalls/routers will often require additional configuration to ensure broadcast discovery operates correctly, since unlike normal unicast UDP communication, the originating traffic is not sent to the same IP address from which the reply will come. Broadcast traffic is sent to the address xxx.xxx.xxx.255, but the unicast replies from that subnet may come from any (or all) addresses in the range xxx.xxx.xxx.1-254.

Normal stateful firewall rules typically only open the firewall for replies from the same IP address to which the originating traffic was sent, so even if unicast port 9909 UDP traffic is enabled and working, broadcast traffic may still fail. Therefore, most firewalls have special rules that can be applied to allow

broadcasts to function correctly, (such as **ip helper-address**, **ip directed-broadcast** and **ip forward-port** functions on Cisco equipment, for example). Check with your firewall manufacturer for the correct broadcast address configuration instructions for your particular systems.

UDP & ICMP Timeouts

UDP packets are subject to numerous possible delays, particularly on busy networks or across slower links. This can cause some machines to not respond to scans reliably. You may be able to compensate for some of these issues by increasing the UDP/ICMP Send and Receive values.



UDP and ICMP Timeouts

Send: ms (1 to 5000)

Receive: ms (1 to 10000)

Presets:

Normal LAN

If you have configured broadcasts and multicasts correctly, and you know for certain that your routers, firewalls, and switches are configured correctly to pass the traffic, and yet you are not seeing all of your remote machines, you may need to increase your timeout settings.

The **Presets:** [Choose] drop-down list allows you to quickly pick recommended values for various network configurations. You can use these values as suggested starting points to adjust your timeouts.

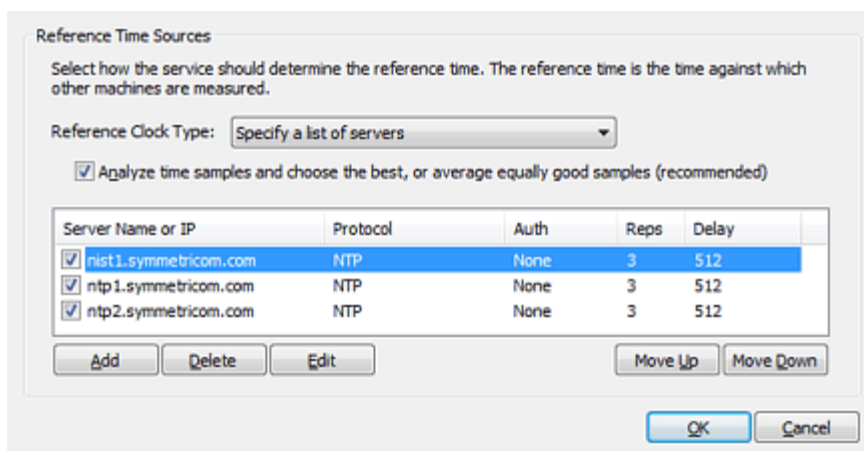
For troubleshooting, you may want to temporarily pick the longest timeout preset (Satellite Link) to eliminate timeouts as a connection issue. However, in production, set timeouts only long enough to receive responses from all of your systems. Setting too large of a timeout will result in scans taking an excessively long time to complete.

If, after adjusting the broadcast/multicast addresses and timeouts, you still encounter issues with seeing all of your remote systems, you should consider using [Audit Server](#) instead. Audit Server uses a persistent list of machines with more robust unicast communication to be sure remote machines are reachable.

Reference Clock

Specify the reference time that Monitor will use to compare against all sampled machines by clicking the button. All variances and alert thresholds will be calculated from this time.

Important: Stable reference time is critical to obtaining trustworthy variance data from your network. Choose sources that are known to be reliable and available over low-latency connections.



Reference Time Sources

Select how the service should determine the reference time. The reference time is the time against which other machines are measured.

Reference Clock Type: Specify a list of servers

☒ Analyze time samples and choose the best, or average equally good samples (recommended)

Server Name or IP	Protocol	Auth	Reps	Delay
<input checked="" type="checkbox"/> nist1.symmetriccom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp1.symmetriccom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp2.symmetriccom.com	NTP	None	3	512

Add Delete Edit Move Up Move Down OK Cancel

The **Reference Clock Type:** ☐ Use this machine's clock ☐ list gives you multiple options for obtaining reference time:

■ **Use this machine's clock**

The local machine's clock is used as the reference. Use this setting if you have the Domain Time Server running on this machine set to synchronize using PTP. Otherwise, only use this setting if the local time on your machine is being well-corrected by a reliable process, either by Domain Time or another source, such as an internal GPS clock card.

■ **Use this machine's sources**

When selected, Monitor Service will use the same time sources used by the Domain Time Server or Client installed on the local machine. This is an excellent option if you have already configured the local Server or Client to obtain time from reliable sources using the NTP or DT2 protocols.

■ **Specify a list of servers**

Use this option to specify the exact machines you want to use for your reference time.

■ **Discover DT2 server(s)**

■ **Discover NTP server(s)**

■ **Discover any available server(s)**

The auto-discovery options allows Monitor Service to locate available servers of the selected type on the network. Discovery will use all discovered servers if the *Analyze all listed servers and choose the best...* checkbox is checked, otherwise it will use the first discovered server.

Note: To avoid the possibility of inadvertently using a free-running local clock, the discovery process will not use the local machine, even if the local machine is a time server.

Analyze time samples and choose the best, or average equally good samples (recommended)

This controls whether Monitor Service applies advanced analysis algorithms to the collected time samples.

When this box is checked, Monitor Service contacts all of the listed servers to collect a group of time samples. It then performs statistical analysis on the collected samples to determine the reliability and uses the most reliable samples to derives the correct time.

See the [About Time Samples](#) sidebar for more information and rule-of-thumb suggestions on acquiring time samples.

If you are collecting multiple samples, checking this box will almost always improve your reference time's accuracy and reliability.

If this box is unchecked, no comparative analysis among samples is performed. In addition, the list of time servers to query becomes a fallback-only list. In other words, Monitor Service will only contact first listed time server. This server will always be used unless it is unavailable, at which point the next listed server will be used. If that server is unavailable, the next server in the list will be tried, etc. When the first listed server becomes available again, the Server will revert to using it exclusively.

Saved Histories

Sort by IP address

Sort by variance

Save up to histories

for up to days

The **Saved Histories** settings control how saved scan results are displayed on the [History](#) tab page, and how many scan

results to keep at any one time.

The amount of disk space used by your stored scan results varies significantly based on how many machines are scanned and how often scans are run. Monitor keeps a maximum of 999 scans, but you can reduce this amount if you have disk constraints.

Alerts

Settings on this page control Alerts.

[Alert Definitions](#)

If Master and Slaves disagree by	or more milliseconds
If any one machine is off by	or more milliseconds
If the average system variance is	or more milliseconds

The items in this section allow you to specify the type and amount of time variance that will be tolerated before Monitor generates an alert.

The default values are suggestions only. The acceptable range of variance depends on your individual network and sync requirements. Examine your scan results to get a good feel for the range of normal variance on your network, and then choose settings that will only flag problems that you consider serious enough for intervention.

If Master and Slaves disagree by or more milliseconds

The synchronization between the master and its slaves is critical to overall time accuracy on the network. It is a good idea to specify a closer tolerance for these systems than for others on the network.

If any one machine is off by or more milliseconds

This specifies the outside limit of variance for any single machine, regardless of its role (server or client).

If the average system variance is or more milliseconds

Raises an alert if the overall synchronization of your network exceeds your limits.

[Alert Actions](#)

Record details in Event Viewer Log	Summary event only
Send email alert notice	High Priority
Display Control Panel applet	
Play a sound	
Filename:	
Repeat every	seconds until acknowledged

Settings in this section define the actions that Monitor will take when an alert is triggered.

Record details in Event Viewer Log

Checking this box instructs Monitor to put alert information in the Windows Event logs. If the **Summary event only** checkbox is checked, only a single summary event will be entered in the Event logs. If **Summary event only** is unchecked, an individual event will be entered for each machine scanned.

Send email alert notice

When this box is checked, Monitor will send an error message similar to the sample below to the email recipients specified on the [Email Setup](#) pages. The **High Priority** checkbox will send the mail marked High Priority.

Domain Time II Monitor Alert

This is an automated alert from Domain Time II Monitor. The system time variance exceeds your minimum alert level. The attached text file contains the scan data. If your email client cannot open attachments, please use the Control Panel applet to examine the scan data.

- Scan Time: Sun 18 Oct 2009 18:44:14
- Scan Status: **Warning**

[Scan Summary.txt attachment]

Domain Time II Monitor 5.1.b.20090724R

Copyright © Greyware Automation Products, Inc.

Monitor Machine: FLEEGMAN

Data Collection: [Success]

Scan Time: Sun 18 Oct 2009 23:44:14 UTC

Sun 18 Oct 2009 18:44:14 (local)

Reference Clock Type: Discovered (NTP) (averaged)

Reference Clock Time: [ERROR]: could not obtain reference time; Time sources disabled

Reference Offset: 0.000000 seconds

Type and average delta yy ddd hh:mm:ss.mss

```
-----
Indies:          1  + 00 000 00:00:00.000
Clients:         3  + 00 000 00:00:00.001
Total:           4  + 00 000 00:00:00.000
```

Slowest Clock: Sun Oct 18 2009 23:44:14.973 UTC 192.168.10.2 (OMEGA13)

Fastest Clock: Sun Oct 18 2009 23:44:14.973 UTC 192.168.10.4 (Grabthar)

Worst Indie: + 00 000 00:00:00.000 N/A ()

Worst Client: + 00 000 00:00:00.003 192.168.10.4 (Grabthar)

Master/Slave Test: Limit: 500 ms Slave Var: + 00 000 00:00:00.000 [Passed]

Single Machine Test: Limit: 1000 ms Worst Var: + 00 000 00:00:00.003 [Passed]

Avg Variance Test: Limit: 750 ms Total Var: + 00 000 00:00:00.000 [Passed]

INDIVIDUAL VARIANCES (sorted by variance)

+ 00 000 00:00:00.000	Full Client	192.168.10.2 (OMEGA13)
+ 00 000 00:00:00.000	Indp. Server	192.168.10.3 (Fleegman)
+ 00 000 00:00:00.000	Full Client	192.168.10.28 (HV-2003-R2-X64)
+ 00 000 00:00:00.003	Full Client	192.168.10.4 (Grabthar)

Sample Monitor Alert Email (reduced size)

Display Control Panel applet

When this box is checked, Monitor will pop up the Control Panel applet on the console of the logged-on user.

Play a sound

When this box is checked, Monitor will play the selected sound file, repeating it as indicated.

Logs

Settings on this page control the Monitor Service's text logs and Windows Event logging.

[Text Log](#)

Log Level: Information

Max Size: KB (use zero to mean unlimited size)

Log Roll: Daily at Midnight

Delete old logs

Keep up to old logs

This section selects the properties of the **domtimem.log** service text log.

Text Logs are kept in the **%SystemRoot%\System32** folder. There are two main log files collected when the service is running:

- **domtimem.log**

This is the currently active service text log file.

If log archiving is enabled (see below), additional archived log files will be created using a **domtimem.YYYYMMDD.log** naming scheme (i.e. domtimem.20090928.log).

- **domtimem.startup.log**

A detailed text log of the service startup process. Only data from the latest startup is included.

To view these logs, click the button, which launches the Domain Time Log Viewer.

Log Level

This drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

This switch will only disable the **domtimem.log** file. The other system log, **domtimem.startup.log** cannot be disabled.

- **Errors**

Only messages marked as Errors will be logged

- **Warnings**

Logs will include Errors and Warnings

- **Information**

Includes Errors, Warnings, and Information messages.

- **Trace**

Includes all of the above, plus detailed Trace information .

- **Debug**

Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

Max size

This sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to help you determine the timeframe and scope of any issues you may encounter.

Log Roll

Domain Time can automatically archive the text log on a daily, weekly, or monthly schedule.

When the log is archived, all existing log events in the **domtimem.log** file will be written to an archive file named **domtimem.YYYYMMDD.log** (i.e. domtimem.20090928.log) and the current log file will then be cleared to accept new data.

You can choose how many archived log files to keep on the machine. When the indicated limit is reached, the oldest log file will be deleted.



This section specifies whether Monitor's service activity will be echoed to the Windows Event logs.

Note: Some levels of logging can create a significant amount of data. The Windows Event logs can be difficult to read, or the Event Log process may even have problems recording all the data when large amounts of log activity are generated.

You should consider using only the **Error** level when using the Event Logs unless you generate a very small amount of logging data overall. In general, Text logging is a better choice for keeping more detail.

The **Log Level** drop-down chooses what type of entries to include in the Event logs. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

Monitor will not log events to the Windows Event Logs.

- **Errors**

Only messages marked as Errors will be logged

- **Warnings**

Logs will include Errors and Warnings

- **Information**

Includes Errors, Warnings, and information on the activity of the time service, such as time sources contacted, amount of clock correction, etc.

- **Trace**

Includes all of the above, plus detailed information on time setting and time sample analysis.

- **Debug**

Includes all available information provided by the service.

Warning:

The amount of data generated by Debug logging can easily overwhelm the Event Log system. Use the Text log for debugging instead.

Click the button to launch the Windows Event Viewer.

Domain Time II Update Server

Version 5.2

Domain Time II Update Server is a system service that automatically installs and/or upgrades Domain Time Server and Client on machines on the network. Update Server can run unattended, updating systems as they are discovered, or it can run on demand.

Update Server is also an excellent way to update your machines to the latest version of Domain Time when it is updated. When you upgrade the Manager/Management Tools on your Management Workstation, Update Server will take note and can automatically update your entire network to the latest version without administrator intervention (if desired).

Update Server is included as part of the Domain Time II Management Tools. **IMPORTANT:** If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[Installation](#)

[System Requirements](#)

Additional Requirements

- Update Server requires that Domain Time Manager be installed on the same machine.
- The Manager/Update Server machine must be a domain member.
- The target machines must be members of the Update Server machine's domain, OR be members of a domain with a trust relationship to that domain OR (if in a workgroup or a stand-alone system) have an administrative account with the same username/password as the one used by the Update Server service.
- Your network must be a correctly-configured Windows network, i.e. configured with working name resolution (DNS, WINS, NetBIOS, etc.), correct and functioning Active Directory (if used), working inter-domain trusts, etc.
- Your network must pass both UDP and TCP network traffic sent to destination port 9909. Switches and firewalls must pass this traffic bi-directionally, since traffic will originate either from Manager or the remote machines. Your network must pass this traffic, regardless of what time protocols are used to actually synchronize the time.


Note: As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

- The remote machine must respond to PING requests from the connecting machine.
- The connecting Domain Time program, utility, or service must be run using credentials with sufficient privileges to connect to and write files to the administrative shares on the remote machine using Microsoft Networking (Domain Admin if the target is a domain member, Local Machine Administrator if the target is in a workgroup).
- The Remote Registry Service must be running on the remote systems and its registry keys must be accessible to the connecting program.
- All files from the original distribution for each type of product you want to install (Server, Client, etc.) must be extracted and present on your connecting machine. Setup copies these to the proper locations in the **\Program Files\Domain Time II** folder for you automatically when you install the Management Tools.

Installation

The Update Server is installed from [Domain Time II Manager](#). Click *Update Server --> Install Update Server service* on the Manager's menu to install it.

Update Server is fully integrated with Manager. Once the Update Server service has been installed, you configure it using the *Update Server* menu items.

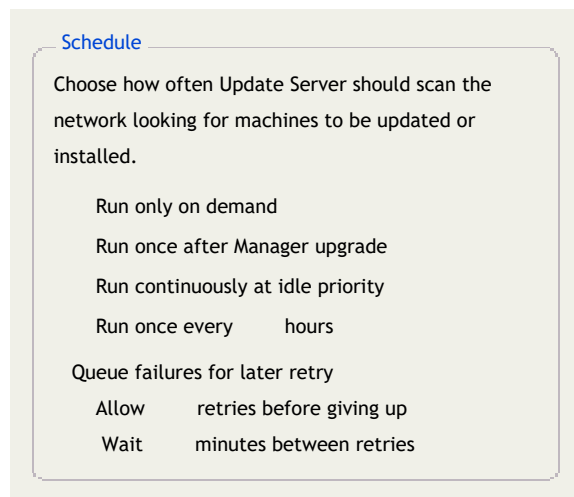


Configuration

Update Server is very simple to configure:

- Set an update schedule.
- Choose which components want installed on your systems (Update Actions).
- Select which domains you want Update Server to examine.
- Provide the access credentials for Update Server to access your target machines.

Schedule



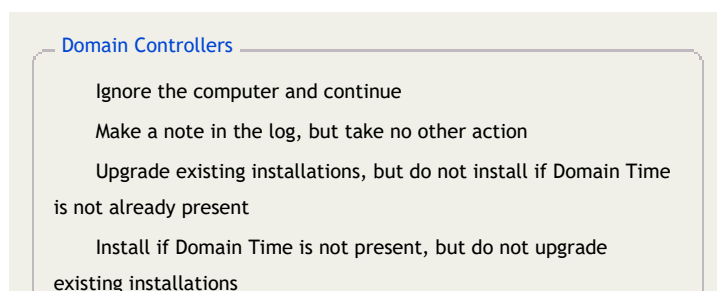
Click the *Update Server* -> *Schedule* item from the Manager menu to set the Update Server schedule. You have a great deal of control over when the updates are performed.

When Update Server performs an update, it examines the domains you've selected (see below) for machines that need installation or update, and then installs/upgrades Domain Time Client or Server with the template settings you've selected (see below).

Queue failures for later retry

You can tell Update Server how many attempts it should make to install or upgrade a machine if it encounters a problem during the process. You can also select how long to wait between installation/upgrade attempts. The retry period can increase the chance of a successful update, giving time for locked files to clear or for services to stop. Any failures to complete an install/upgrade will be logged in the Update Server logs.

Update Actions



Install or upgrade as needed

New Installations

Domain Time Server, using [Default Server Template]

Domain Time Client, using [Default Client Template]

Use the *Update Server -> Update Actions* items from the Manager menu to configure how you want Update Server to handle machines it discovers.

You can set separate options for **Domain Controllers** and for all other **Servers and Workstations** found. The options on the Update Actions dialog for each type is identical (the *Domain Controllers* dialog is shown above).

If you've selected to install Domain Time on newly discovered machines, you have the option of choosing whether to install Domain Time Server or Client, and which settings template to use for the installation. The templates presented in the drop-down lists are the ones currently configured in Domain Time Manager. See the Manager's [Using Templates](#) page for details on working with templates.

Select Domains

Domains and Workgroups

Choose which domains Update Server should examine.

NetBIOS Name	Domain Name	Data Source
<input checked="" type="checkbox"/> CHILD1	child1.test1.galaxyquest.com	Active Directory
<input type="checkbox"/> TEST1	test1.galaxyquest.com	Active Directory

OK Cancel

Select domains to examine. [\[Click for larger size\]](#)

Update Server shares Domain Time Manager's view of the network. The *Update Server -> Advanced -> Select Domains...* menu item allows you to pick which domains Update Server will process from the list of all Domains and Workgroups that Manager has discovered on your network. See Manager's [Interface](#) page for details on how the Domains and Workgroups list is generated.

Credentials

Credentials

Update Server needs credentials sufficient to install services and access remote file systems. In general, this should be a Domain Admin account.

Run the Update Server service using the LocalSystem account, but use the following credentials to access remote computers (recommended)

Domain:

Username:

Password:

Run the Update Server service using a Domain Admin account

Domain:

Username:

Password:

Update Server needs administrative rights to be able to install/upgrade machines remotely. The settings on the *Update Server* -> *Advanced* -> *Credentials...* dialog allow you to specify the account used by Update Server to perform installations and upgrades.

You have the choice of having the Update Server service itself run under the LocalSystem account and supply the administrative access credentials only when performing an update, or having the service running with the administrative privileges at all times. In general, the first option is preferred. In either case, account details are encrypted in the registry.

Update Server can install/upgrade on domains and workgroup members as long as the credentials supplied match an administrative account on the domain (or local machines in the workgroup). If you select a workgroup or domain to which Update Server does not have administrative access, the updates will fail and will be noted in the logs.

Update Server Log

Text Log

Log Level: Information

Max Size: KB (use zero to mean unlimited size)

Log Roll: Daily at Midnight

Delete old logs

Keep up to old logs

This section selects the properties of the **dtupdate.log** service text log.

Text Logs are kept in the `%SystemRoot%\System32\` folder. There are two main log files collected when the service is running:

- **dtupdate.log**

This is the currently active service text log file.

If log archiving is enabled (see below), additional archived log files will be created using a **dtupdate.YYYYMMDD.1log** naming scheme (i.e. dtupdate.20090928.log).

- **dtupdate.startup.log**

A detailed text log of the service startup process. Only data from the latest startup is included.

To view these logs, click the button, which launches the Domain Time Log Viewer.

Log Level

This drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

This switch will only disable the **dtupdate.log** file. The other system log, **dtupdate.startup.log** cannot be disabled.

- **Errors**

Only messages marked as Errors will be logged.

- **Warnings**

Logs will include both Errors and Warnings.

- **Information**

Includes Errors, Warnings, and Information messages.

- **Trace**

Includes all of the above, plus additional detailed trace messages.

- **Debug**

Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

Max size

This sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to help you determine the timeframe and scope of any issues you may encounter.

Log Roll

Domain Time can automatically archive the text log on a daily, weekly, or monthly schedule.

When the log is archived, all existing log events in the **dtupdate.log** file will be written to an archive file named **dtupdate.YYYYMMDD.log** (i.e. dtupdate.20090928.log) and the current log file will then be cleared to accept new data.

You can choose how many archived log files to keep on the machine. When the indicated limit is reached, the oldest log file will be deleted.

Other Management Tools

The Domain Time II Management tools include many useful diagnostic and utility programs. Many of these utilities are installed automatically when Manager is installed, and are located in the Domain Time II Program Folder (usually **C: \Program Files \Domain Time II**). Others are installed when Server or Client are installed and found in the **\System32** folder. Some are not installed by default, but only found in the original distribution file folders.

Click a link to jump to the description for the tool:

- ▶ [DTCheck](#) - Multi-purpose Utility
- ▶ [NTPCheck](#) - NTP Server Test
- ▶ [PTPCheck](#) - PTP Messages Test Utility
- ▶ [DTClean](#) - Complete Removal Tool
- ▶ [DTTest](#) - Time Server Test Utility
- ▶ [LMCheck](#) - Simple Variance Check
- ▶ [DTSync](#) - Sync Trigger
- ▶ [DTRCPL](#) - Remote CPL
- ▶ [DTSlew](#) - Manual Slew Utility
- ▶ [DTLockDN](#) - Security Lockdown Tool

DTCheck

This multi-purpose utility can check statistics, trigger Domain Time synchronizations, check clock accuracy, open firewall ports for Domain Time II use, generate high-accuracy variance reports, and more. This tool is installed in the **/System32** folder on Server, Client, and Manager.

```

Administrator: C:\Windows\system32\cmd.exe
C:\>dtcheck
Domain Time Check
Copyright (c) Greyware Automation Products, Inc.

Stats on Eppman
Machine Role: Indp. Server, 5.1
Startup Time: Thu 22 Oct 2009 19:36:04 UTC
Current Time: Mon 26 Oct 2009 14:10:13 UTC
Stats Last Reset: Thu 22 Oct 2009 19:36:04 UTC
Last Time Set: Mon 26 Oct 2009 14:10:01 UTC
Time Set From: Averaged Time using selected protocol
Time Corrections: 4994 checks, 0 errors, 2214 corrections

Sync in Progress: No
Window Size: 1 minute
Last Variance: 0 milliseconds
Running Time: 3 days, 18 hours, 34 minutes, 9 seconds
Cumulative Drift: 5526 ms
Clock Stability: 0.01695 ms per second
Time Zone: Central Daylight Time
Total Bytes In: 844048
Total Bytes Out: 7162687

Protocol Status Tot Reqs Tot Errs Time of Last Req
-----
DT2-UDP Running 700923 0 Mon 26 Oct 2009 14:10:06 UTC
NTP Running 59 0 Mon 26 Oct 2009 14:00:32 UTC
DTI Disabled 0 0 None since last startup
TIME-UDP Disabled 0 0 None since last startup
TIME-TCP Disabled 0 0 None since last startup
DT2-HTTP Running 0 0 None since last startup
  
```

Run **DTCheck /?** from a command prompt to see a list of all the available parameters and options.

You can examine the statistics ([sample](#)) of any Domain Time II server or client, force the synchronization of a particular machine (or of the entire time hierarchy), and generate a system-wide variance report ([sample](#)).

Note: DTCheck's variance reporting is much more accurate than LMCheck utility, since it uses higher accuracy protocols and sampling methods from installed Domain Time II components. Use this utility for variance reports on networks that have Domain Time Servers and Clients installed.

Domain Time II DTCheck utility [\[Click for larger size\]](#)

DTCheck can add Windows firewall exemptions for time protocols, view PTP Masters and display PTP traffic, act as an SNMP listener, and much more. See the Help (run with the **/?** switch) for details.

DTCheck can also be used to test your machine's clock for reliability. Run **DTCheck /test** to test your machine. You will need to reboot the OS after reliability testing, since parameters changed during the test can only be reset at boot-time.

NTPCheck

A utility for testing NTP/SNTP time servers. Use this utility if you need to save NTP server tests to a file, or want to run regular tests in a batch file. This tool is installed in the **/System32** folder on Server, Client, and Manager.

```

Administrator: C:\Windows\system32\cmd.exe

C:\>ntpcheck /?
Domain Time NTP Check
Copyright (c) Greyware Automation Products, Inc.

This tool retrieves the time using NTP/SNTP from the servers you specify
and displays the variance in milliseconds from the clock on this machine.

SYNTAX: ntpcheck [-raw] ["server1"] ["server2"] ["server3..."]

Server specification: For each "server" specify the DNS name or IP address
of an NTP server, followed optionally by a symmetric key to use, or the number
of samples. You only need the quotation marks if you are including extra parns.

EXAMPLES:
ntpcheck time.nist.gov                checks time.nist.gov
ntpcheck                               searches for servers
ntpcheck "ntp1.symmetricon.com samples 2" takes 2 samples
ntpcheck "time.windows.com key 1 samples 3" uses key 1, takes 3 samples

```

Domain Time II NTPCheck utility [Click for larger size]

NTPCheck provides clock test information similar to that of DTCheck, but uses the NTP/SNTP protocol to query servers instead of the Domain Time II protocol. It is useful for determining whether or not a particular server is reachable and operating, and for comparing the time reported by multiple servers.

NTPCheck is also useful for demonstrating the limits of NTP/SNTP accuracy. With the -raw option, you can see the results of other information derived from the NTP packets.

For example, here are two actual [sample reports](#) generated by querying time.nist.gov. The first query shows the standard NTPCheck response; the second query shows the results of the -raw option.

PTPCheck

A utility for viewing the available PTP protocol messages (in particular, management messages) visible from any PTP node. This tool is installed in the /System32 folder on Server, Client, and Manager (as of v5.2.b.20171101).

PortIdentity	Name	Dom	State	Delta	Err	Replies	IP
408d5c-ffe-e775e9.1	FLEEG...	0	Slave	+0.0000010	0	5 of 5	192.168.10.3
00155d-ffe-0a9f04.1	HV-20...	0	Slave	-0.0000012	0	5 of 5	192.168.10.148
00155d-ffe-0a9f02.1	HV-2016	0	Slave	+0.0000002	0	5 of 5	192.168.10.70
002618-ffe-2e8216.1	OMEG...	0	Master	+0.0003962	0	5 of 5	192.168.10.5
408d5c-ffe-e775e9.2	PTPM...	0	Passive	0.0000000	0	5 of 5	192.168.10.3
00155d-ffe-0a9f04.2	PTPM...	0	Passive	0.0000000	0	5 of 5	192.168.10.148
002618-ffe-2e8216.2	PTPM...	0	Passive	0.0000000	0	5 of 5	192.168.10.5

Replies: 7 nodes, 38 messages, 7706 bytes

Domain Time II PTPCheck utility [Click for larger size]

a scan

You can start PTPCheck either by running PTPCHECK.EXE directly or by launching the program by right-clicking and selecting to scan from the context menu in Manager's PTP Monitor section.

PTPCheck can scan for various message types across all available PTP domains using both multicast and unicast tests. You can (and should) restrict your scans to only domains and subnets you know exist, as well as choosing to scan using IPv4 and/or IPv6 traffic based on your network configuration. This speeds your test and cuts down on spurious traffic across your network.

Use the **PTP Domains**, **Muticast TTL**, and **Boundary Hops** settings to set the scope of your tests. Set the **PTP Domains** to the domain number(s) you know are in use. You may list them (comma-separated) or set a range if there are several. The **TTL** value specifies how many routers the packet can cross. The **boundary hops** setting is used by PTP-aware switches to

You may copy the PTPCheck.EXE to any Windows machine to run the tests, whether Domain Time is installed or not. This gives you the ability to see what messages are available across various routers, boundary clocks, and subnets.

Why do I need this tool?

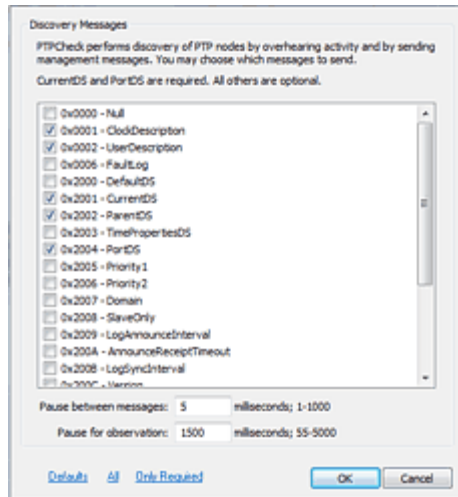
PTP management messages are optional in the PTP IEEE 1588-2008/2019 specifications and implemented differently by different vendors and services. The Domain Time [PTP Monitor](#) function of Domain Time Audit Server relies on management messages to collect variance and audit data, so it's important to know if (a) any particular node or device is able to respond to management queries, and (b) if those replies are visible throughout your network.

Configuring

Launch from PTP Monitor [Click for larger size]

limit traffic. You must set the TTL to a sufficiently high number to cross all routers (including VLANs). Keep in mind that the responding node's TTL must also be set high enough for the replies to make it back. Set the boundary hops to 1 more than the number of boundary clocks between the machine running PTPCheck and the target machine. (There is no corresponding boundary hops to set on the receiving machine).

Note, even if you set the hop counts properly, many boundary clocks and PTP-aware switches prevent the passage of management messages not only to other subnets, but even to other ports on the same switch. Use PTPCheck on multiple machines to see if your switches are actually passing management traffic. If your switch is not passing management messages, contact the manufacturer for a patch.



The Discovery Options Dialog [Click for larger size]

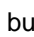
Configure the Discovery Options

Choose *Edit -> Discovery Options* from the PTPCheck menu. This displays all of the available message types PTPCheck can scan for. Note, this is not an exhaustive list of all available PTP messages. It is meant to focus on management messages using in monitoring, although several other standard message types may be selected. PTPCheck only sends GET requests and listens passively for Announces and Syncs; it cannot set variables on PTP nodes. Overheard Announces and Syncs are useful for discovering master nodes, and PTPCheck disassembles some of the information for you, but it cannot analyze packets it has not received. Use a protocol analyzer like Wireshark for detailed troubleshooting of PTP synchronization.

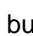
Click the **Defaults** link to reset the selections to the default message types used by **PTP Monitor**. This is the most useful selection for testing compatibility with Audit Server. You may, of course, add or subtract from the default selections to gather more information from your devices/nodes.


Each node/device takes a discrete amount of time to respond to requests, and there are often delays in its ability to respond to sequential requests of different types. You can set the amount of delay between sending requests and also how long PTPCheck waits to hear a response (the **Pause for Observation** value). The Pause for Observation time is most useful for overhearing Announces and Syncs. Many master clocks do not respond to management messages, so the only way to detect them is by listening to them talk to the network. You will need to experiment with your particular nodes to choose the optimal settings that give you consistent results.

Run a Scan

You can run a network-wide scan based on the configuration settings you made above by clicking the  button.

Scans send multicast messages to the wildcard portIdentity (AllClock.AllPorts). Responding machines will send back multicast packets directed to PTPCheck. The list of nodes is then filled with responders. PTPCheck cannot tell you anything about a node that does not respond at all to wildcard multicasts. A red indicator bulb means the node was detected, but did not respond to any management messages (this almost always means it is an overheard master that doesn't handle management messages). A yellow indicator bulb means the node was detected, but either didn't respond to all queries, or responded to one or more with an error response. A green indicator bulb means that the node correctly responding to all management messages sent by PTPCheck.

You can test a single IP or DNS name (either IPv4 or IPv6) by using the **Unicast Test** field. Click the  button to send your configured list of management messages to the selected IP. Note that PTP requires a targetPortIdentity, even when messages are unicast directly to a particular node. PTPCheck uses the wildcard portIdentity in the unicast test. Some interpretations of IEEE 1588-2008/2019 forbid nodes from responding to unicast requests if the targetPortIdentity is a wildcard, so devices that respond to multicast queries may or may not respond to unicast queries.




Use the Context Menu to run a test [Click for larger size]

Another convenient way to send multicast or unicast messages to an individual machine is to first run a Scan, which will display all visible nodes, then then right-click on any node in the list and choose to send either multicast or unicast requests directly to that node. (If the node is a master, you may also test the delay request-response cycle.) Unlike the Scan and Test

buttons, right-clicking a node sends a message with the targetPortIdentity set to the node's actual portIdentity instead of a wildcard. If you are testing unicast and the node has more than one IP, you may choose to which IP the requests should be sent.

View the Results

The node list will show the initial summary results of the scan. The number of messages sent and successful replies received is shown in the **Replies** column. The **Errs** column shows the number of messages to which a node replied, but with an error message instead of a management reply. For example, if you send six management queries, and a node responds to all of them, the Replies column will show "6 of 6" and the indicator bulb will be green, unless one or more replies were errors -- in that case, the Errs column will show how many of the replies were error messages, and the indicator bulb will be yellow.



Domain Time II PTPCheck utility [Click for larger size]

To view the detailed results of a scan or an individual test, double-click on the the node name in the node list. All messages types sent (as configured on the Discovery Options dialog above) will be listed. Failed types will be in red.

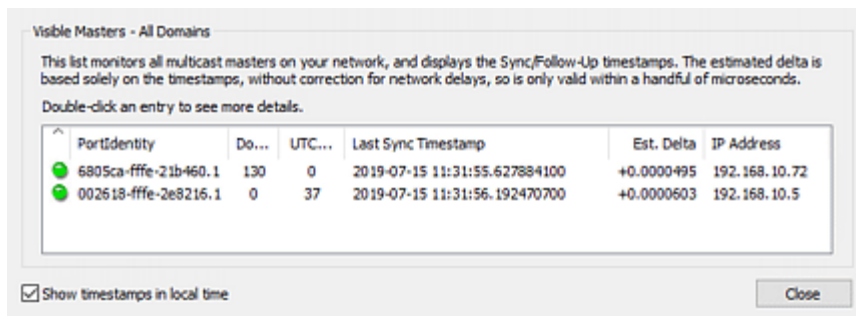
You can expand each green item to see the detailed contents of the replies received. Domain Time's [PTP Monitor](#) combines a number of these messages to provide the information displayed on its readout. This can help you determine if your device is providing correct and sufficient information. Recall that each manufacturer determines which messages to support and how, so you cannot expect every device to provide all messages.

As mentioned above, the Default settings on the Discovery Options page reflect the message types used by [PTP Monitor](#), so pay particular attention to those items.

You can save your test results to a text file by clicking the button. Here's a sample of [the output](#). Tech Support may request this file when troubleshooting PTP issues.

Master Monitor

Introduced in v5.2.b.20190701, Master Monitor shows all current masters and their sync timestamps, regardless of PTP domain they use. This is very useful to determine if there is more than one active master is on your network, and if they are serving roughly the correct time-of-day.



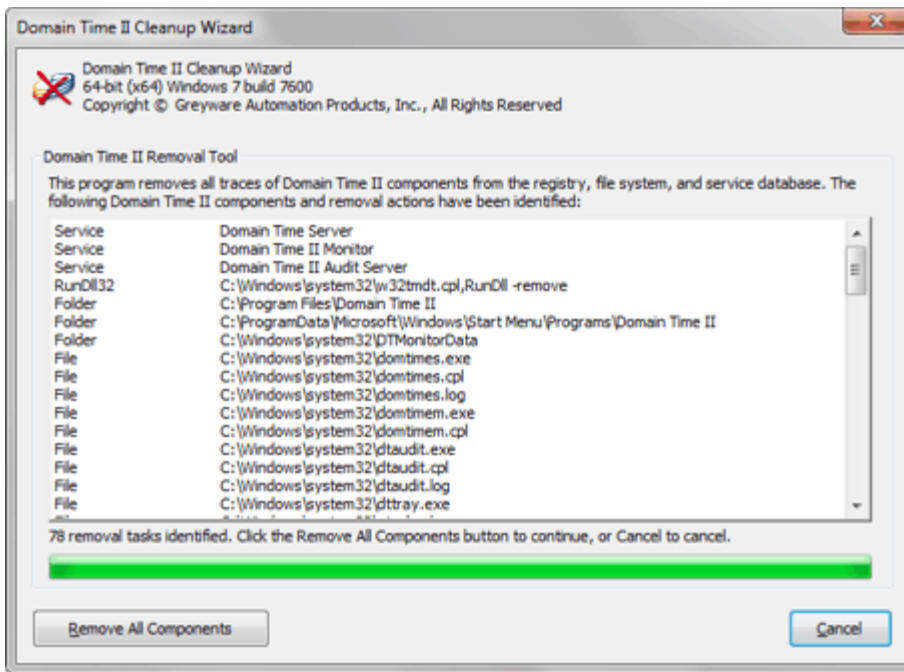
The PTP Master Monitor Display [Click for larger size]

This tool derives the delta values from the announce, sync, and sync followup packets from each master, but it doesn't perform any delay measurement. Therefore, the displayed delta will not be as accurate as you would see from a slave synchronized to the Master. However, it does provide a good eyeball glance at the state of your masters.

Domain Time Removal Tool (DTClean)

DTClean is a utility that completely removes all traces of Domain Time II programs and registry settings from your system. This tool is included with Server, Client, and Manager. With Server and Client the tool is located in the distribution file folders; with Manager it is located in the Program Files\Domain Time II folders

DTClean should be used with care, since it removes all configuration settings as well as program executables. If you are upgrading to a newer version of Domain Time, you



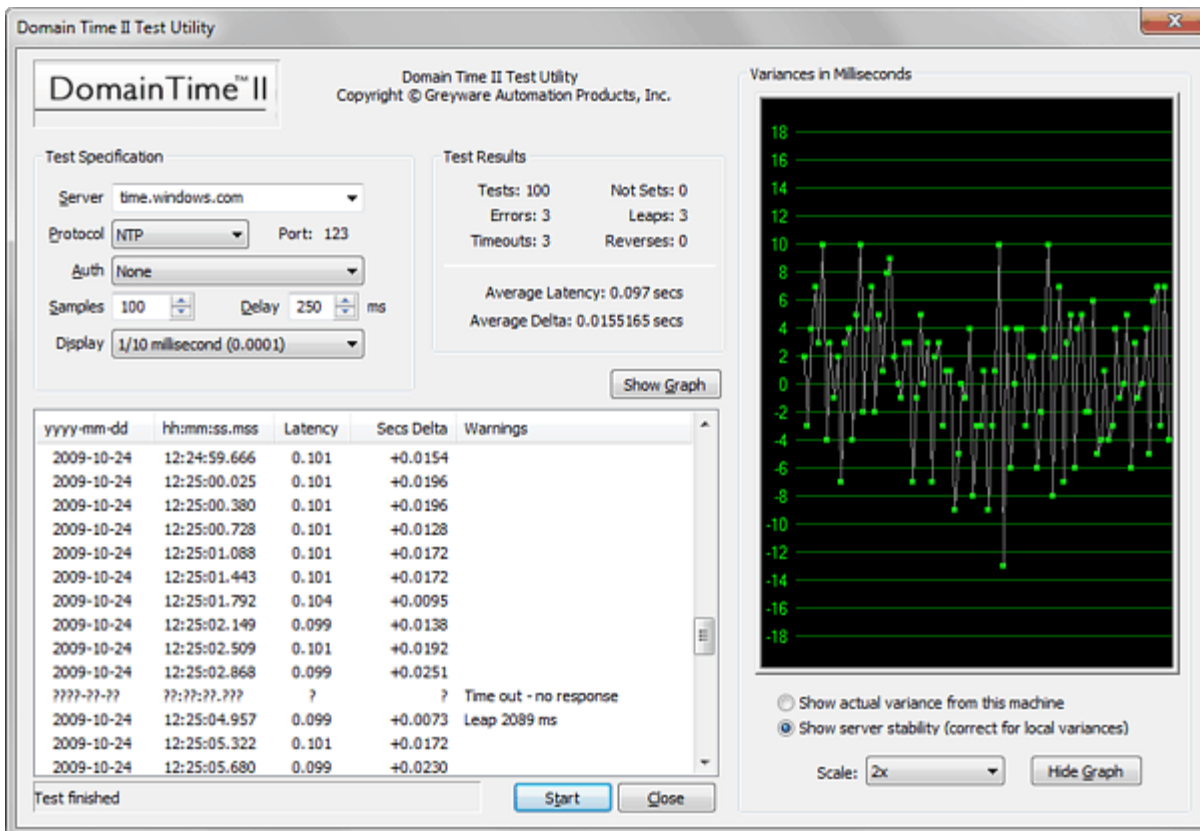
Domain Time II DTClean Utility [\[Click for larger size\]](#)

should use the [Setup program](#) or [Domain Time II Manager](#) instead.

DTClean keeps a log of the components it removes, and you may save a copy of the log file for troubleshooting purposes or to supply to technical support if requested.

DTTest

Use this utility to test the clock stability of any time server. Use it to determine which servers to use as time sources, or to troubleshoot accuracy issues.



Domain Time II Test Utility [\[Click for larger size\]](#)

To test a time server:

- Enter the server name or IP address of the time server you want to test in the **Server** field.
- Use the **Proto** drop-down list to select the time protocol to use for the test (this protocol must be running on the server being tested).
- Click the **Start Button** to begin the test.

You may also want to adjust how many times

and how rapidly to test each server by adjusting the **Poll Interval** and **Number of Tests** items. Different poll rates affect can affect how much detail you see in the server's response characteristics. You may want to compare a very rapid sample rate to the results from a fairly slow sample to see if the server has resolution or response issues when under rapid load.

Hint: If you will be testing against a Domain Time II Server, you will want to temporarily disable the Denial-of-Service protection on the Server. If you don't, Server will interpret rapid test rates as a Denial-of-Service attack and stop responding to your tests.

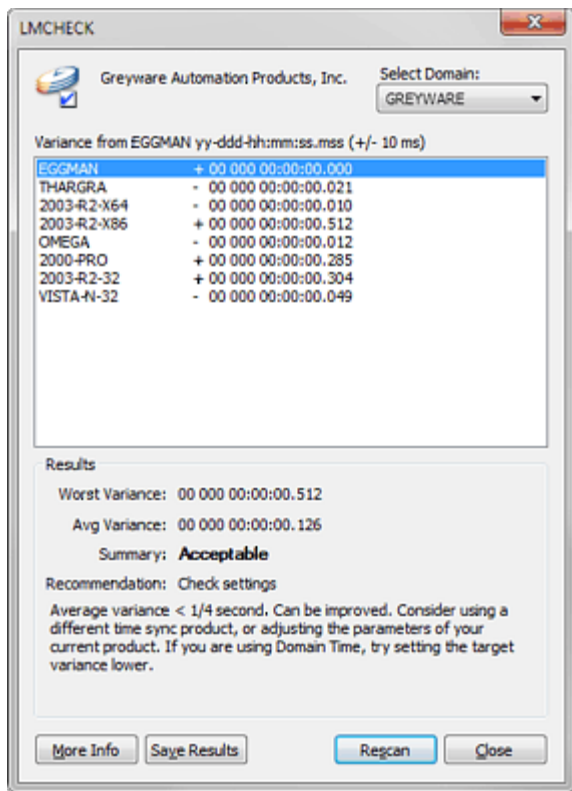
The test will show a running list and a real-time graph showing of the amount of latency detected in the network connection, and also how large a variance exists between your local system clock and the server being tested.

Since both the local machine and the remote system clocks and protocols have some built-in inaccuracies, the values displayed will fluctuate occasionally. However, you should be able to see an overall trend in multiple tests - stable clocks will show a fairly consistent variance, unstable clocks will have constantly varying values.

You can adjust the scale of the graph to show the graph in proper perspective to the accuracy you are expecting to achieve.

LMCheck

Use LMCheck to obtain a quick variance report and save the results to a file. Use this tool to do a quick & dirty check of network synchronization on a network that doesn't already have Domain Time II installed.



[Domain Time II LMCheck Utility](#) [Click for larger size]

- ▶ Nothing to install -- remote machines only have to be running Windows (XP or later)
- ▶ Just run the 32-bit or 64-bit version of LMCHECK.EXE from any Windows machine.

Notes:

- Target machines must be running Microsoft Networking (with NetBIOS-enabled) and respond to NetRemoteTOD queries. NetBT is disabled by default on most current versions of Windows and must be re-enabled to use LMCheck.
- On recent versions of Windows, you must run the program as Administrator (right-click and choose *Run As Administrator* from the context menu).
- The variance report generated by LMCheck cannot be as detailed or as accurate as variance reports provided by the [Domain Time II Manager](#), the [Monitor Service](#), the [DTCheck utility](#), or [Domain Time II Audit Server](#), each of which use much more accurate time protocols and sampling methods to measure the time differentials.

The Domain Time LMCheck test tool lets you roughly assess the current time of Windows machines on your network quickly and easily. It uses the built-in LAN Manager NetRemote TOD (Time of Day) function to check the time on all the machines in the browse list.

Click the **Start** button to perform the scan. Click the **Save Results** button to pull the results up in Notepad so that you may save them wherever you wish.

Time variances from the machine on which you run LMCheck are calculated and displayed, taking into account any network latencies. You may select the domain you wish to scan from the drop-down list.

Although it is included as part of the licensed Domain Time II Management Tools, LMCheck itself is freeware, and can be downloaded separately and freely distributed as long as the program is unmodified.

Domain Time II Remote CPL (DTRCPL)

Use the Remote CPL utility to quickly connect to a Domain Time II Server or Full Client and change its Control Panel Applet settings. This is a useful utility when all you need to do is change a control panel applet setting and you don't need the full power of Domain Time II Manager.

Choose a machine running Domain Time II Server or Full Client from the drop-down list, browse list, or enter its IP address, DNS name, or NETBIOS name into the Machine field. If the connection is successful, you will be presented with a locally-running version of the remote machines' Domain Time II Control Panel Applet. You will then be able to make all the configuration changes you would if you were actually using the remote machine (with the exception of running Time Source tests).

The DTRCPL utility is subject to the same requirements as Domain Time II Manager in order to connect to and control a remote system:

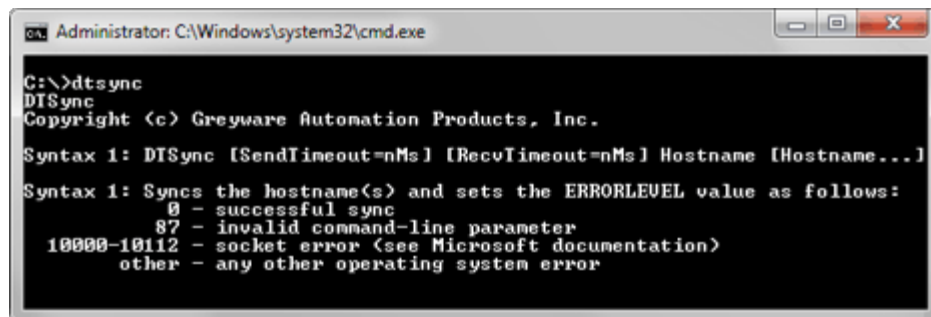
- Your network must be a correctly-configured Windows network, i.e. configured with working name resolution (DNS, WINS, NetBIOS, etc.), correct and functioning Active Directory (if used), working inter-domain trusts, etc.
- Your network must pass both UDP and TCP network traffic sent to destination port 9909. Switches and firewalls must pass this traffic bi-directionally, since traffic will originate either from Manager or the remote machines. Your network must pass this traffic, regardless of what time protocols are used to actually synchronize the time.

Note: As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

- The remote machine must respond to PING requests from the connecting machine.
- The connecting Domain Time program, utility, or service must be run using credentials with sufficient privileges to connect to and write files to the administrative shares on the remote machine using Microsoft Networking (Domain Admin if the target is a domain member, Local Machine Administrator if the target is in a workgroup).
- The Remote Registry Service must be running on the remote systems and its registry keys must be accessible to the connecting program.

DTSync

Use this utility to trigger a sync on specified machines from the command line.




```
Administrator: C:\Windows\system32\cmd.exe
C:\>dtsync
DTSync
Copyright (c) Greyware Automation Products, Inc.

Syntax 1: DTSync [SendTimeout=nMs] [RecvTimeout=nMs] Hostname [Hostname...]

Syntax 1: Syncs the hostname(s) and sets the ERRORLEVEL value as follows:
          0 - successful sync
          87 - invalid command-line parameter
          10000-10112 - socket error (see Microsoft documentation)
          other - any other operating system error
```

[Domain Time II DTSync utility](#) [Click for larger size]



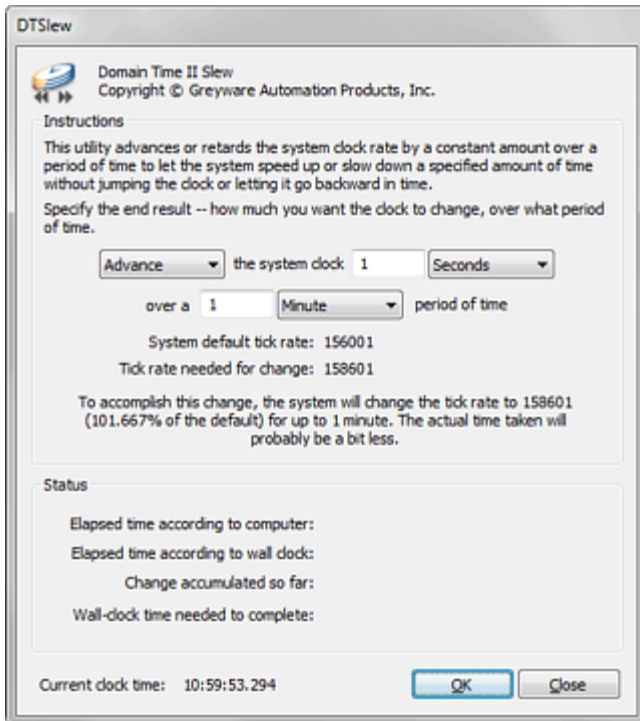
[Domain Time II Remote CPL](#) [Click for larger size]

Run **DTSync** from a command prompt to see a list of all the available parameters and options.

DTSync allows you to specify timeouts and to set the ERRORLEVEL variable so you can create robust batch files to reliably trigger synchronization, even across WAN links.

DTSlew

This utility allows you to smoothly slew the local clock by large amounts.



Domain Time II DTSlew Utility [Click for larger size]

stepping).

Use this utility to move the local clock forward or backward by the amount you specify. The clock will be advanced or retarded using slewing, so you can make the change smoothly with no clock stepping or backwards clock movement.

This is useful if you have to manually change the time on machines running critical services that must have smooth forward clock movement at all times. DTSlew also allows you to make larger changes than would normally be possible by Domain Time Server or Client.

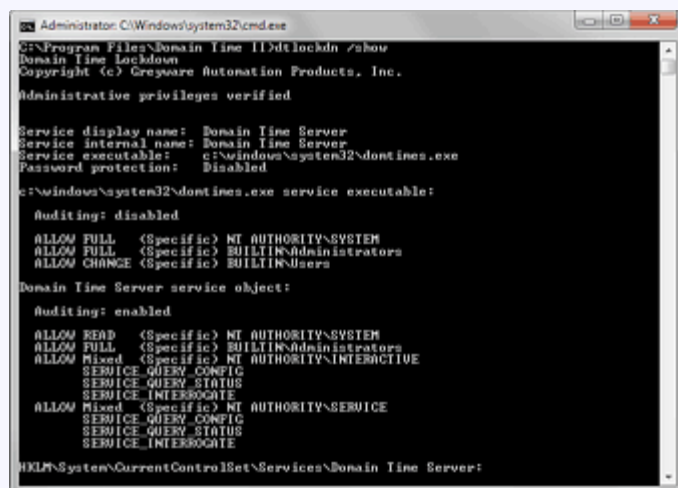
The rate of change is limited to the maximum amount of slewing possible by the hardware on the motherboard. DTSlew will not allow you to select a rate that is outside of these limits.

Note: You will need to stop the Domain Time Server or Client service before running DTSlew in order to prevent conflicts over clock ' control.

Do **NOT** attempt to serve time from a machine running DTSlew, this will cause unpredictable results on your clients as they attempt to track with the time server (such as unexpected

Domain Time Lockdown (DTLockDN)

Domain Time Lockdown is a command-line tool for system administrators to use to help secure (harden) their Domain Time installations.



Domain Time II DTLockDN Utility [Click for larger size]

service.

- The operating system's services database

The operating system maintains an internal database of service objects, including their current status, their permissions, and their settings. Most of this information is stored in the registry under

- Who needs it?

Domain Time Lockdown is useful to system administrators as part of an overall company-wide security policy.

- What does it do?

Domain Time Lockdown lets you set permissions for

- The Domain Time service object

The service object is the handle presented by the operating system to programs wanting to control the service. Just like files or other objects, the service object may have permissions associated with it. Service object permissions control who is allowed to stop, start, query, or configure the

HKLM\System\CurrentControlSet\Services. Ordinary users do not have permission to modify these settings. This area is where the operating system keeps the name of the service executable file, the restart on failure options, the startup type, and so forth.

- The Domain Time parameters stored in the registry

Domain Time keeps its configuration in HKLM\Software\Greyware*product*, where *product* may be either Domain Time Client or Domain Time Server. Information in this area controls what Domain Time does once it is running as a service (time sources, how often to check, system timings, logging options, and all other settings).

- The Domain Time service executable (domtimec.exe or domtimes.exe)

The main service executable lives in the system32 directory. Administrators (and often users) have rights in both the containing folder and the individual files. If users have the right to add or delete files in the folder, they can also delete or rename the service executable, even if the executable file itself is restricted to read-only or has a specific deny ACE protecting it from deletion. The only way to prevent a user who has delete rights for the folder from deleting an individual file is to add a null ACE (effectively remove all permissions). Therefore, unlike the other objects, when you set a user or group to have only READ access, the program will actually remove all access from the executable file for that user or group.

- Aren't the default permissions sufficient?

In most circumstances, yes. Non-administrative users typically don't have the ability to stop, remove, or even install services. They may have limited abilities to control what the running service does, or trigger it to take certain actions—these options vary by the service, and Microsoft and other vendors typically use sensible defaults to help ensure that only administrators can change vital settings.

However, home users (and even some business users) may use an administrative account as their primary logon. Security experts strongly discourage this practice, and Microsoft's own UAC has taken steps to help mitigate the dangers of logging on this way, but nevertheless it is not uncommon for ordinary users to find themselves with full administrative control over their machines, perhaps without even realizing it.

Other accounts or groups sometimes have unintended privileges. On regular workstations, the Power Users group typically has additional control over services. On Domain Controllers, the Server Operators group has similar privileges. Individual accounts or other groups may also be configured to have extended privileges using system or domain policies.

- How does it work?

Domain Time Lockdown edits or replaces the [access control lists](#) to restrict control access and optionally enable auditing. It can also set the service to restart automatically if killed. (The Microsoft property page for service control only allows setting the restart time on the order of minutes; Domain Time Lockdown lets you set a restart time in milliseconds.)

Domain Time Lockdown only supports READ or FULL permissions. READ permissions are required in order for users to query the service, see the current settings, and operate the computer normally. FULL permissions include all READ permissions plus the ability to stop, remove, upgrade, or configure the service.

For example, you could use Domain Time Lockdown to grant FULL permissions to the built-in Administrator account while granting only READ permissions to the built-in Administrators group. This would allow anyone logged in as the local built-in Administrator to control the service, while other members of the Administrators group (including Domain Admins if the machine is a member of a domain) could only view the settings.

There is no predefined hardening for a service, because what access you need to restrict and what access you need to allow is dependent on your network's policies and configuration.

- Syntax

```
dtlockdn [service="Service Display Name"] [options...]
```

Options containing embedded spaces must be enclosed in quotation marks.

If you do not specify `service=` and a service name, the program will look for either Domain Time Client or Domain Time Server (whichever) is installed. If you do specify a service name, it may be any installed service on the machine. We do not support using this program on services other than Domain Time Client or Domain Time Server.

Options

<code>/Show</code>	Show current settings; do not make any changes.
<code>/Restart=nnn</code>	Set service to auto-restart if killed after <i>nnn</i> milliseconds.
<code>/NoRestart</code>	Set service to not auto-restart if killed.
<code>/Audit</code> ¹	Enable auditing of unauthorized access.
<code>/NoAudit</code>	Disable auditing of unauthorized access.
<code>/Full="Account"</code> ²	Grant "Account" full control of the service.
<code>/Read="Account"</code> ²	Restrict "Account" to read-only access to the service.
<code>/Revoke="Account"</code> ²	Remove "Account" from the service's access control list.
<code>/Replace</code> ³	Replace permissions instead of merging them.
<code>/ServiceOnly</code>	Apply security only to the service object and executable.
<code>/RegistryOnly</code>	Apply security only to the registry objects.
<code>/Yes</code>	Do not ask for confirmation before making changes. You may use either <code>/Yes</code> or <code>/Y</code> .
<code>/Password="password"</code> ⁴	Set password to lock out subsequent changes. If a password is set, you must provide exactly the same password in the future, or the program will refuse to perform. The only way to clear a password once it has been set is by issuing the <code>/Reset</code> command with the correct password.
<code>/Reset</code> ⁴	Reset the service and registry to default access (read for ordinary users, full control for administrators and the system). If you have set a password using the <code>/password</code> option, you cannot reset the service without providing the correct password again.

1 Enabling auditing with this program sets the appropriate bits in each object's SACL to allow the system to record failed access in the system's security log. If your machine's policy does not have failure auditing enabled for object access, then no entries will appear in the security log.

2 You may specify a username or a group name for *Account*. If the name contains embedded spaces, you must enclose it in quotation marks. You may use plain names, such as *Users*, *"Power Users"*, *Administrator*, or *Joe* to refer to accounts or groups on the local machine. You may also refer to domain users or groups this way. If there is any chance of account name duplication throughout your domain or forest, you should specify the full names: *BUILTIN\Administrators*, *"BIGCORP\Domain Admins"* or other fully-qualified names. In some circumstances, depending on your active directory configuration, you may be able to use the *joe@bigcorp.com* form to specify individual accounts.

3 The program will ensure that the special SYSTEM account always has full control. It is an error to specify SYSTEM as an account on the command line. The program will also ensure that ordinary users and administrators will have the ability to read values they should read, even if you try to `/Revoke` those permissions, or use `/Replace` without specifying all the necessary accounts.

4 Exercise caution when using the optional `/password` option. Once you enter a password, you must provide it again *exactly* the same way in order to use the program again. For example, *MyPassword*, *mypassword*, and *MYPASSWORD* are three different passwords. If your password contains embedded spaces, you must enclose it in quotation marks. The best password contain a mixture of upper-case and lower-case letters, numbers, and punctuation marks. **Passwords are stored using one-way encryption, so we cannot help you recover your password if**

you forget.

Once a password is set, you must provide it for each use of the program thereafter. The only way to clear a password is to use the **/Reset** command, but you must provide the current password to do so. After a reset, you may then set a different password if desired.

Examples

```
dtlockdn /full=Administrator /read=Administrators /replace
```

This example allows the built-in Administrator account to control the service, but blocks all other members of the Administrators group. Any permissions granted by inheritance or prior operations will be replaced.

```
dtlockdn /restart=1000
```

This example changes only the service's auto-resetart time. If the service dies unexpectedly, or is killed using Task Manager or another tool, it will restart in 1000 milliseconds (one second).

```
dtlockdn /full="Domain Admins" /full=Administrator /read=Administrators /replace /restart=1000  
/password="nzlw00Fm_#gadlob88$" /yes
```

This example is similar to the first example, but also grants the group Domain Admins full control, sets the service to restart automatically if killed, sets a password, and suppresses the prompt before executing.

```
dtlockdn /reset /password="nzlw00Fm_#gadlob88$"
```

This example recovers control after permissions have been locked down. The security will be reset to generic defaults, and the password will be removed. Note that if a password hadn't been set, any user with full administrative rights on the machine could have issued the **/Reset** command and then reconfigured the security and perhaps have added a different password.

Most options for Domain Time II Manager and the Management Tools are set using the Domain Time II Manager program. A few advanced options can only be set by changing the registry. This section explains the registry entries used by Domain Time II Manager and the Management Tools.

Caution:

Modifying Registry entries requires basic familiarity with the Windows Registry and its operations. Incorrect changes to the Registry can result in unpredictable, perhaps non-repairable, damage, up to and including a non-bootable system! Have a qualified technician make the changes for you if you are not comfortable with the process. We cannot be responsible for registry problems.

Domain Time II Manager

Domain Time II Manager settings are located in this key:

```
HKEY_LOCAL_MACHINE
  Software
    Greyware
      Domain Time II Manager
        Parameters
```

Value Name: ICMP TTL (hop limit)

Value Type: REG_DWORD

Default Value: 32 (decimal)

Range: 1 to 255 (decimal)

Notes: This value controls the number of router hops that are allowed in an ICMP echo ("ping") request. Domain Time pings machines first to help eliminate long waits for machines that are unreachable. You should only need to adjust this value if you have an LAN/WAN configuration requiring more than the default 32 hops.

Value Name: Service Log Filename

Value Type: REG_SZ

Default Value: %SystemRoot%\System32\dtman.log

Notes: Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename dtman.log in the %SystemRoot%\System32\ folder. The complete path and filename must be specified (i.e. C: \Windows\System32\dtman.log) and the drive specified must be a local drive.

Domain Time II Monitor Service

The Monitor Services has two main registry keys:

Logs and Alert settings are located in this key:

```
HKEY_LOCAL_MACHINE
  Software
    Greyware
      Domain Time II Monitor
        Logs and Alerts
```

Parameter settings are located in this key:

```
HKEY_LOCAL_MACHINE
  Software
    Greyware
      Domain Time II Monitor
        Parameters
```

Value Name:	Service Log Filename
--------------------	----------------------

Value Type:	REG_SZ
--------------------	--------

Default Value:	%SystemRoot%\System32\dtmonitor.log
-----------------------	-------------------------------------

Notes:	Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename dtmonitor.log in the %SystemRoot%\System32\ folder. The complete path and filename must be specified (i.e. C: \Windows\System32\dtmonitor.log) and the drive specified must be a local drive.
---------------	--

Domain Time II Update Server

Domain Time II Update Server settings are located in this key:

```
HKEY_LOCAL_MACHINE
  Software
    Greyware
      Domain Time II Update Server
        Parameters
```

Value Name:	Service Log Filename
--------------------	----------------------

Value Type: REG_SZ

Default Value: %SystemRoot%\System32\dtupdate.log

Notes: Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename dtupdate.log in the %SystemRoot%\System32\ folder. The complete path and filename must be specified (i.e. C: \Windows\System32\dtupdate.log) and the drive specified must be a local drive.

Domain Time II Audit Server

Version 5.2

Domain Time II Audit Server is a robust central data-collection and monitoring service. You can monitor time synchronization, collect historical data, raise alerts, and generate summary reports for your entire network from your Management Workstation.

IMPORTANT: If upgrading from 4.x, please read the [4.x to 5.x Considerations](#) page.

[Installation Instructions](#)

[System Requirements](#)

[Disk Space Estimator](#)

Installation

Audit Server is tightly integrated with [Manager](#) and [Server](#) and **must** be set up on a machine where they both are already installed.

Note: Since many functions of Domain Time II Audit Server depend on accurate time calculations, it should always be run on a physical (not virtual) machine.

- If you have obtained Audit Server in a separate distribution package, use its [Setup](#) program to install Audit Server.
- If the Audit Server files are already present on Domain Time II Manager, you will have the option to install the Audit Server service from the **Audit Server** item on Manager's menu.

Audit Server runs as a background service, so it should already be running after installation. To configure Audit Server, you will need to use Domain Time Manager. Click the **Domain Time Manager** icon in the *Start -> All Programs -> Domain Time //* program folder.

You may also launch the Domain Time II Manager program (and many other installed Domain Time II components) by right-clicking on the Domain Time icon in the System Tray to bring up the context menu.

Network Requirements

Verify that your environment meets the minimum requirements for performing remote operations using Domain Time components. In order to be able to install, upgrade, or configure remote machines:

- Your network must be a correctly-configured Windows network, i.e. configured with working name resolution (DNS, WINS, NetBIOS, etc.), correct and functioning Active Directory (if used), working inter-domain trusts, etc.
- Your network must pass both UDP and TCP network traffic sent to destination port 9909. Switches and firewalls must pass this traffic bi-directionally, since traffic will originate either from Manager or the remote machines. Your network must pass this traffic, regardless of what time protocols are used to actually synchronize the time.

Note: As of Version 5.2.b.20150821, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for detailed information.

- The remote machine must respond to PING requests from the connecting machine.
- The connecting Domain Time program, utility, or service must be run using credentials with sufficient privileges to connect to and write files to the administrative shares on the remote machine using Microsoft Networking (Domain Admin if the target is a domain member, Local Machine Administrator if the target is in a workgroup).

- The Remote Registry Service must be running on the remote systems and its registry keys must be accessible to the connecting program.
- All files from the original distribution for each type of product you want to install (Server, Client, etc.) must be extracted and present on your connecting machine. Setup copies these to the proper locations in the `\Program Files\Domain Time II` folder for you automatically when you install the Management Tools.

Configuration

Configure Reference Time

Before using Domain Time Audit Server, you should make sure you have decided on what time source(s) to use to act as Reference Time. Audit Server uses the Reference Time for comparisons when measuring time deltas (variances) and for recording timestamps in logs and alerts.

Note: Audit Server shares Manager's settings for this option.

Reference Time is configured by selecting *Options -> Network Options -> Reference Time Sources...* from the Manager menu.

Important: Stable reference time is critical to obtaining trustworthy variance data from your network. Choose sources that are known to be reliable and available over low-latency connections.

Reference Time Sources

Select how the service should determine the reference time. The reference time is the time against which other machines are measured.

Reference Clock Type: Specify a list of servers

☒ Analyze time samples and choose the best, or average equally good samples (recommended)

Server Name or IP	Protocol	Auth	Reps	Delay
<input checked="" type="checkbox"/> nist1.symmetriccom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp1.symmetriccom.com	NTP	None	3	512
<input checked="" type="checkbox"/> ntp2.symmetriccom.com	NTP	None	3	512

[Reference Time Source Selection](#) [\[Click for larger size\]](#)

The **Reference Clock Type:** Use this machine's clock list gives you multiple options for obtaining reference time:

- **Use this machine's clock**

The local machine's clock is used as the reference. Use this setting if you have the Domain Time Server running on this machine set to synchronize using PTP. Otherwise, only use this setting if the local time on your machine is being well-corrected by a reliable process, either by Domain Time or another source, such as an internal GPS clock card.

- **Use this machine's sources**

When selected, Audit Server will use the same time sources used by the Domain Time Server or Client installed on the local machine. This is an excellent option if you have already configured the local Server or Client to obtain time from reliable sources using the NTP or DT2 protocols.

- **Specify a list of servers**

Use this option to specify the exact machines you want to use for your reference time.

- **Discover DT2 server(s)**

- **Discover NTP server(s)**

■ Discover any available server(s)

The auto-discovery options allows Audit Server to locate available servers of the selected type on the network. Discovery will use all discovered servers if the *Analyze all listed servers and choose the best...* checkbox is checked, otherwise it will use the first discovered server.

Note: To avoid the possibility of inadvertently using a free-running local clock, the discovery process will not use the local machine, even if the local machine is a time server.

Analyze time samples and choose the best, or average equally good samples (recommended)

This controls whether Audit Server applies advanced analysis algorithms to the collected time samples.

When this box is checked, Audit Server contacts all of the listed servers to collect a group of time samples. It then performs statistical analysis on the collected samples to determine the reliability and uses the most reliable samples to derives the correct time.

See the [About Time Samples](#) sidebar for more information and rule-of-thumb suggestions on acquiring time samples.

If you are collecting multiple samples, checking this box will almost always improve your reference time's accuracy and reliability.

If this box is unchecked, no comparative analysis among samples is performed. In addition, the list of time servers to query becomes a fallback-only list. In other words, Audit Server will only contact first listed time server. This server will always be used unless it is unavailable, at which point the next listed server will be used. If that server is unavailable, the next server in the list will be tried, etc. When the first listed server becomes available again, the Server will revert to using it exclusively.

Audit Configuration

How to configure and run Audits.

Audit List

As of version 5.1, Audit Server shares Domain Time II Manager's view of the network and is completely integrated with it. Manager contains the master database from which machines can be selected for audit.

Machines appearing in the Manager lists can be audited for the following information:

- Machines running Domain Time II can be fully audited, including collection of synchronization (drift) logs and full statistical audit records.
- Machines answering NTP queries can be partially audited, with limited statistical audit records (variance, last time source, etc.). Audit Server can create drift logs based on measured offsets at audit time.
- As of version 5.2.b.20170101, PTP masters and slaves may also be audited. Audit Server can create drift logs based on available delta measurement (masters) and/or reported deltas from available management messages.

Alerts can be raised for the above systems at audit time. In addition, machines running Domain Time II v5.1 and later can also provide real-time alerts. See the [Alerts](#) page for more information.

To manually select machines for audit:

- Select a machine list in Domain Time II Manager and be sure the **Audited** column is displayed in the Details pane. (Use *View -> Add/Remove Columns* item on the Manager menu, or right-click the column header and be sure *Audited* is selected on the context menu.)
- Click the *Audited* entry for your machine to change it to "Yes," or right-click the machine item and choose *Enable Audit* from the context menu.
- You can enable auditing on multiple machines at once by highlighting the machines you want in the Details pane and choosing *Enable Audit* from the right-click context menu.

Machines can also be automatically added to and pruned from the Audit list. See the [Audit List Management](#) item on the *Audit Server -> Advanced* menu for details.

Machines on the Audit list will be contacted by Audit Server when an audit collection run is initiated. Audit collection runs can be scheduled or manually triggered.


Quick Start

To audit Windows machines:

- Install Domain Time Client or Server.
- Toggle the "Audited" column entry to your desired [Audit Group](#) for your selected Domain Time machines from either the [Domain Time Nodes](#) or the [Domains and Workgroups](#) lists. (Note: even though Domain Time machines may also appear on the [PTP Nodes](#) list, do not choose to audit them from that list.)
- Enable [Post-Audit Alerts](#), if desired.
- Enable [Real-Time Alerts](#) on Domain Time, if desired.
- Enable collection of [Synchronization logs](#) from Domain Time, if desired.

To audit Linux machines via ntpd or chronyd:

- Be sure ntpd or chronyd is set to respond to standard NTP time requests (act as an NTP time server).
- Add the Linux machine(s) to the [NTP Nodes](#) list.
- Toggle the "Audited" column entry to your desired [Audit Group](#) for your Linux machines from the [NTP](#)



Select Columns

[Click for larger size]

[Nodes](#) list.

- Enable [Post-Audit Alerts](#), if desired.
- Enable collection of [NTP Server Drift logs](#) either at audit time or on a regular polling schedule, if desired.


To audit ptpd or other PTP machines via PTP Monitor:

- Note: Do not use PTP Monitor to audit Domain Time machines. Use the [Windows instructions](#) above.
- Carefully read and configure PTP Monitor per the [documentation](#).
- Toggle the "Audited" column entry to "Yes" for your selected PTP node(s) from the [NTP Nodes](#) list.
- Enable collection of [PTP Node Drift logs](#), if desired.

Audit Schedule

Audit Server can run audit collections from your Audit list on a schedule you specify. Multiple times and schedules can be configured.

To configure your audit schedules, pull up the Schedule dialog by clicking *Audit Server -> Schedule...* from the Domain Time Manager menu.



[Audit Schedule](#) [Click for larger size]

You can Add, Edit, or Delete schedules from this dialog. The schedules can be displayed in either 12 or 24 hour format.

Note: Schedule times and dates always refer to the time on the local machine.

When you click Add or Edit, you'll be presented with the **Audit Schedule Editor** where you can select the day(s) and time(s) of your audits, and whether you want your audit to repeat on a regular interval.

[Set Audit Schedule](#)

Enabled

If this schedule is enabled, an audit will run on the day(s) and time(s) you specify. The time of day refers to local time on this computer.

Every: Mon Sat
 Tue Sun
 Wed
 Thu
 Fri

Time: repeating every

Trigger an Audit

You can start an audit run manually from Manager either by choosing the *Audit Server -> Audit Now* menu item or by pressing F11.

Audits can also be triggered remotely using the [DTCheck](#) command-line tool that is included with Server and Client using this syntax:

```
DTCHECK \\machine_name_or_ip_address /cmd="Audit Trigger"
```

Audit Tasks

Audit Server can perform certain tasks before and after an audit run. To configure these tasks, choose the *Audit Server -> Audit Tasks* menu item.

Pre-Audit Tasks

- Trigger synchronization of audited machines before beginning the audit
- Wait for all synchronizations to complete
- Proceed with audit even if some machines have not fully synchronized yet
- Scan the network before contacting individual audited machines
- Use multicast to locate DT2 machines that may have changed IPs or names

The **Trigger synchronization of audited machines before beginning the audit** option will cause Audit Server to send a synchronization trigger to audited machines before auditing them. You can indicate whether Audit Server should pause until machines are all synchronized before proceeding (which can cause the audit to take more time to complete) or to proceed with the audit immediately.

This option can be used to satisfy some regulatory requirements (such as the FINRA requirement that machines be synchronized before the start of trading and at fixed periods throughout the day). However, triggering a synchronization before the audit may not give you an accurate picture of the actual state of your machines from the audit records, since all machines will be recently synchronized at the time of the audit. You will need to use [Synchronization \(drift\) logs](#) and Real-Time Alerts to determine the actual status of machines between audits.

If **Scan the network before contacting individual audited machines** is checked, Audit Server will use Manager's scan settings to collect data by multicast/broadcast before attempting to check with each machine. If unchecked, Audit Server will skip the initial scan. In general, you should leave this enabled.

The **Use multicast to locate DT2 machines that may have changed IPs or names** function helps keep your Manager database up-to-date. You should leave this enabled unless you have a completely static network configuration.

Post-Audit Tasks

- Send a summary of the audit results by email
- Send errors as an attachment
- Auto-generate a textual version of the audit results
- Reset the statistics on audited machines
- Delete audit result files that are more than days old.

Send a summary of the audit results by email will cause Audit Server to send a summary email after each audit (see sample below).

Check the **Send errors as an attachment** checkbox if you want to include details on machines that have errors.

Domain Time II Audit Summary

This is an automated summary from Domain Time II Audit Server. An audit has completed. 10 machines were checked during this audit.

- Audit Server: EGGMAN
- Audit Time: Mon 30 Nov 2009 00:08:51
- Audit Status: **Passed**
- Audit Errors: 0
- Audit Summary:
 - 0 reported clock not set
 - 1 did not respond this audit
 - 0 had variance in excess of 5000 milliseconds
 - 0 did not set the clock for over 1440 minutes
 - 10 reported no alerts
 - 0 real-time alerts have not been acknowledged

Auto-generate a textual version of the audit results

When checked, Audit Server will expand a text version of the audit results in the Audit Results folder. This is similar information you see when clicking the **Print Details** button when looking at audit results using the [Audit Viewer](#). Note, these files will not display by default in the **Audit Server -> Audit Results** list in Manager's left-hand pane. To view them, right-click the "Audit Results" label and choose "Open containing folder" from the context menu.

Reset the statistics on audited machines

This option will reset all statistics on the remote machine after each audit.

Note: This will also clear out the synchronization (drift) log on the remote machine. This option is not reversible; once the stats and drift logs are cleared, they are deleted permanently. In most cases, you should not use this option since you cannot recover the data if sync log collection fails on the Audit Server for any reason. You should only use this option if you are successfully collecting synchronization logs on the Audit Server so that you retain the data. See the [Synchronization Logs](#) section for more information.

Delete audit result files that are more than days old.

This option will trim your audit results list after they reach the specified age. The most recent Audit Results are shown in the **Audit Server** category of Manager's Tree pane.

Note: Records can accumulate at a rapid pace, particularly if you are auditing many machines on a frequent schedule and you can easily fill up your existing storage. You should arrange to archive off this data to other media if you want to preserve the records indefinitely rather than letting them accumulate without limit. Audit records are found in the folder specified on the [Audit Server -> Advanced -> Data Folders...](#) menu item.

Alerts and Audit Groups

Domain Time II Audit Server can raise various alerts based on information collected during collection runs and from real-time data provided by Server and Client. In order for Audit Server to provide alerts, you need to configure the alert thresholds and the type of alerts desired. Select *Audit Server -> Alerts and Audit Groups -> Configure* from the Manager menu.

Audit Groups

As of v5.2.b.20180601, you can set alerts and notifications for up to eight custom-designated groups. On older versions, only one global set of alerts and notifications can be configured.

Each Audit Group has its own individual set of alert thresholds and notification lists. Once you've defined an audit group, you can assign any machine shown in Manager [nodes lists](#) to the group by clicking on the **Audited** field to select the desired group name.

Note: When upgrading from older versions of Audit Server, any machines previously set to be audited will be automatically assigned to Group 1. You may then re-assign them to any group you choose.

On the Audit Groups page, you can also set basic Real-Time Alert global defaults, Real-Time Alert defaults for machines that are set to unaudited, and other Advanced Real-Time Alert configuration options.

[Audit Group Configuration](#) [Click for larger size]

Audit Groups Click the button corresponding to the audit group you want to configure. See [detailed instructions](#) below:

Default audit group sets the group to which machines are auto-added if they are added with auditing enabled.

Double-check anomalous test results by sending a follow-up unicast (applies to all groups)

For speed and efficiency, Audit Server first requests audit results by broadcast/multicast. Check this box if not all your machines respond reliably to initial scans (you can see this in the Audit Server log). Although enabling this function is more robust, it may significantly slow down audits if you have a large number of non-responding machines, since timeouts

are invoked for each audited machine that does not respond.

Enable processing of Real-Time Alerts (applies to all groups and logs)

Enable/disable overall Real-Time Alerting.

Only raise Real-Time Alerts for audited nodes (applies to all groups)

Typically, Real-Time Alerts may be raised for machines whether or not they are set to be audited. This setting restricts them to audited nodes.

Keep Real-Time Alert history logs

Max size: KB (min 1, max 65536)

When enabled, Audit Server keeps a separate log of all Real-Time Alert activity (found in the `\Program Files\Domain Time II\RT Alert History` folder)

Click the button to set Real-Time Alert parameters for unaudited machines. These settings are very similar to the Real-Time Alert settings for the individual audited groups (See [detailed instructions](#) below). Email alerts for these machines can only be directed to the global email distribution list.

Click the button to set additional parameters for Real-Time Alerts on the **Advanced Real-Time Alert Configuration** dialog page. These options are covered in the [Advanced Settings](#) section at the end of this document:

Configure Audit Group

You may configure the individual alerts configuration for each Audit Group by clicking on its corresponding button. This brings up the **Configure Alerts** dialog for the selected group.

Audit Group 1

The settings on this page pertain only to this audit group. You should give it a descriptive name. The group name appears in logs, audit results, and alerts.

Group name: (max 15 characters)

Post-Audit Alerts

Audit Server can alert you when the deltas on your network exceed the tolerances you specify. Select the conditions to trigger an alert after an audit completes:

☒ A node's time is off by or more ☐ seconds ☒ ms ☐ us

☒ A machine's clock has not been set for or more minutes

☐ An audited machine fails to respond for or more audits

If a post-audit email alert is generated for a node in this audit group, send a copy to:

Real-Time Alerts

Audit Server can raise an alert between audits if a machine reports either that it cannot set the clock, or that it corrected an excessive delta.

☒ Raise alert upon receipt of a real-time alert from a Domain Time node if it cannot set its clock, or if a correction exceeds ☐ seconds ☒ ms ☐ us

☐ Do not count startup corrections as excessive, regardless of magnitude

If a Domain Time node reports that it has lost sync with its PTP master:

☐ Ignore it

☒ Treat it as a warning (auto-resets when master regained)

☐ Treat it like any other error (requires acknowledgement)

If a Real-Time alert email is generated for a node in this audit group, send a copy to:

[Configure Alerts dialog for an Audit Group](#) [\[Click for larger size\]](#)

Use the **Group Name** field to set the audit group name. This is the name is used in displays, reports, and alerts.

Post-Audit Alerts

These alerts are raised after a scheduled or manually-triggered [Audit Run](#). These thresholds are used to determine which machines raise alerts.

A node's time is off by **secs** **ms** **µs**

Any audited machine with a time delta exceeding this value (as compared to the [Reference Time](#)) will raise an alert.

A machine's clock has not been set for **or more minutes**

Any audited machine that has not set its time more recently than this value will raise an alert.

An audited machine fails to respond for **or more audits**

An alert is raised if an audited machine hasn't responded for this number of audits.

If a post-audit email alert is generated for a node in this audit group, send a copy to:

You may enter a custom list of email addresses to which this alert should be sent. **Note:** basic email settings must first be configured using the global [Email Configuration](#) setup (detailed below). This custom list is in addition to the default recipients specified in the global Email Configuration. Global recipients always receive all emails. Individual recipients specified in each group's settings will only receive alert emails pertaining to that one group. If you leave this field blank, alerts will be sent to the default list of recipients (if enabled).

Real-Time Alerts

Domain Time Servers and Clients can send Real-Time Alert data to Audit Server during each time check/statistics roll-up event scheduled on the [Timings](#) page). See the Configure Nodes for Real-Time Alerts sidebar on the right for setup instructions. This data can be evaluated and used to raise an alert based on the threshold value you specify in this section.

Raise alert upon receipt of a real-time alert from a Domain Time node if it cannot set its clock, or if a correction exceeds secs
ms µs

Do not count startup corrections as excessive, regardless of magnitude
You may instruct Audit Server to not raise an alert based on the first Real-Time response after a component restarts. This prevents spurious alerts during service startup, since the first correction of the clock is often very large.

If a Domain Time node reports that it has lost sync with its PTP master:
This setting controls how to handle alerts when slaves lose their PTP Master:

- Ignore it
- Treat it as a warning (auto-resets when master regained)
- Treat it like any other error (requires acknowledgement)

If a Real-Time alert email is generated for a node in this audit group, send a copy to:

You may enter a custom list of email addresses to which this alert should be sent. Note: basic email settings must first be configured using the global [Email Configuration](#) setup (detailed below). This custom list overrides the default email distribution recipients specified in the global Email Configuration. If you leave this field blank, alerts will be sent to the default list of recipients (if enabled).

Configure Nodes for Real-Time Alerts

Each Server and Client must be configured to send Real-Time Alert Data to the Audit Server before alerts can be generated. This can be done by:

- Configuring the [Audit Server Real-Time Alerts](#) section of the Status Reports property page on the Server or Client Control Panel applet.
- Using [Active Directory](#) policies
- Selecting machines on the Details Pane of Manager's Real-Time Alerts category and right-clicking to choose Enable Real-Time Alerts from the context menu.

Real-Time Alerts appear in the Real-Time Alerts category of Domain Time Manager. Alerts persist until they are dismissed from Manager (by right-clicking the machine's name in the Real-Time Alerts display and choosing Reset Alert item from the context-menu. Real-Time Alerts will also be sent using email if configured. Machines still in alert status at the time of an audit run will also be summarized in Audit summary emails.

Alert Actions

You may also set Manager to play sounds when Real-Time Alert status changes. See the [Options -> Appearance and Interface -> Interface Options](#) menu item.

Alert Actions

Audit Server can raise an alert in several different ways. Choose the kinds of alerts you want to receive.

- ☐ Record details in the Event Viewer log
- ☐ Send an email alert notice
- ☐ Send an SNMP trap
 - ☐ Send a post-audit "All Okay" trap

Community:

Server:

Choose your desired alert method in this section. You can also enable/disable these items directly on the Audit Server -> Alerts and Audit Groups menu.

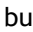
The SNMP alerts and Email items require additional configuration.

SNMP Configuration

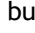
Enter the SNMP community name and password used by your Network Management System (NMS), as well as its DNS name or IP address. Your community name and password must match the one in use by the receiving system.

Best Practices for SNMP include using a unique community name and hard-to-guess password on production systems. The default community public should only be used for initial testing. Although Domain Time only sends outgoing trap information and is therefore not susceptible to SNMP remote control vulnerabilities, you should still be mindful of SNMP security for the benefit of your other SNMP devices.

The Domain Time MIB File

Domain Time comes with a MIB file that you can use to compile on your SNMP monitoring system so that your traps are interpreted correctly. The MIB text file is generated when you click the  button on the SNMP Traps property page of the Server or Client Control Panel applet so you don't need to worry about locating it in some obscure installation folder or having online access.

Email Configuration

Click the  button to configure your Email Settings (or choose Audit Server -> Email Setup from the Manager menu).

You must configure these email settings before Audit Server can send notification emails.

Set the From address

Email Setup

From To CC BCC RTAlerts Summaries Server Queue

Sender's Name: Doman Time II Alerts

Sender's Email: dialerts@ourdomain.com

Email Format: Plain text then HTML

Plain Text

Plain text then HTML

HTML then plain text

Send Test Email... OK Cancel

Email Setup From and Format Selection [\[Click for larger size\]](#)

Specify the From: email address that will appear on the notification emails. You can also specify the format and MIME part order of the emails:

- Plain Text
- Text part followed by HTML part
- HTML part followed by Text part

Choose the format that provides the best compatibility with your email system.

Set the TO/CC/BCC distribution lists

Email Setup

From To CC BCC RTAlerts Summaries Server Queue

Recipient's Name	Recipient's Email Address
IT Helpdesk	helpdesk@ourdomain.com
Lead Engineer	lead.engineer@ourdomain.com
On Call Engineer	oncall.engineer@ourdomain.com

Add Delete Edit

Send Test Email... OK Cancel

Email Recipients List [\[Click for larger size\]](#)

Use the To, CC, and BCC tabs to add the email addresses of your desired recipients.

Set the Outgoing Server

Email Setup

From To CC BCC RTAlerts Summaries **Server** Queue

SMTP Server:

Server Port: STARTTLS if supported ▼

Auth Username: (optional)

Auth Password: (optional)

Outgoing SMTP Server Settings [\[Click for larger size\]](#)

Enter the server address and account login information required for Audit Server to send outgoing mail through your SMTP server.

Optional: Real-Time Alert and Daily Summary distribution lists

Email Setup

From To CC BCC **RTAlerts** Summaries Server Queue

Recipient's Name	Recipient's Email Address
Alerts	alerts@ourdomain.com

☐ Use To/CC/BCC list for RTAlerts

Real-Time Alerts/Summaries Distribution List Settings [\[Click for larger size\]](#)

As of version 5.2.b.20160922, Audit Server has the ability to send Real-Time Alert and Daily Summary emails to a different distribution list than the addresses used in the TO/FROM/BCC settings. To use this feature, uncheck the Use To/CC/BCC list... checkbox and enter the email addresses you want to use for the distribution list. If enabled, Real-Time Alerts and/or Daily Summaries will only go to the addresses listed here, they will no longer be sent to the TO/CC/BCC address lists.

Send Test Email

Once you have entered all of the above information, click the Send Test Email button to generate a test email.

If your test email encounters any errors, you'll receive a pop-up window showing the entire SMTP conversation so you can easily troubleshoot the problem:



The email could not be delivered. The system error code was error 8225:
A protocol error occurred,
and the last SMTP result code was 535.

A log of the SMTP conversation follows:

```
I: Looking up smtp.greymware.com
I: Connected to 71.252.193.51:25
R: 220-smtp.sff.net
R: 220-Greyware Mailman 1.5.b.20090107R
R: 220 ready
S: EHLO Fleegman
R: 250-smtp.sff.net hello Fleegman
(pool-173-74-57-68.dl1stb.fios.verizon.net [173.74.57.68])
R: 250-SIZE 15728640
R: 250-AUTH=LOGIN PLAIN
R: 250-AUTH LOGIN PLAIN
R: 250 HELP
S: AUTH PLAIN
R: 334 send authentication
S: SVRIZWxwZGVzawBJVEhkbHBkZXNrAHNwYWFjZQ==
R: 535 <ITHelpdesk> mailbox does not exist
E: Protocol error: expected 235 but got 535
```

OK

Send Test Email, Showing SMTP Error [\[Click for larger size\]](#)

Check the email queue to troubleshoot delivery issues

Email Setup

From To CC BCC RTAlerts Summaries Server Queue

Queue Folder: C:\Windows\Temp\GWServiceSMTP\Domain Time II Audit Serv

SMTP Trace: Disabled

Open... Browse...

Send Test Email... OK Cancel

Email Queue Settings and Email Logs [\[Click for larger size\]](#)

This page contains the settings for the email queue and email logs.

The Queue Folder: specifies the location of the folder where outgoing emails are queued. The **email.log** trace file is also kept in this folder.

Note: In most cases, you will not need to adjust this location. If you do decide to change the folder location, you must pick a location on a local disk (not a networked share) with sufficient permissions (Change) granted to the Audit Server service account so that it is able to manage the queues.

Use the **SMTP Trace:** Disabled drop-down list to select the level of detail you want to keep in the **email.log** trace file. In general, you should only enable the trace file if you are troubleshooting an email delivery issue. Otherwise, your **email.log** file may grow to an unmanageable size over time.

Use the **Open** or **Browse** buttons to open the queue folder and locate the **email.log** file, which is a plain text file you can open in any editor, such as Notepad.

Advanced Configuration: Email-related registry settings

Depending on your email server configuration, you may also need to adjust these additional settings in the Windows registry.

Email registry settings are located in the **HKEY_CLASSES_ROOT\Gap\GWServicessmtp** key.

TLSSkipCertErrors (REG_DWORD)

Introduced in v5.2.b.20140922 with default=0 (ignore no errors). As of v5.2.b.20160711, the default changed to 0x311 (accept certs that are self-signed, expired, or have the wrong CN)

If this value is zero, the server cert must pass all tests. If the value is non-zero, it is a bitmask specifying which particular types of errors may be ignored. See [Microsoft's documentation](#) for a list of certificate errors that may be ignored. Use a logical **OR** to combine multiple values.

- **0x00000080** - Ignore errors associated with certificate revocation
- **0x00000100** - Ignore errors associated with an unknown (or self-signed) certificate authority
- **0x00000200** - Ignore errors associated with wrong use of a certificate
- **0x00001000** - Ignore errors associated with an invalid/mismatched common name
- **0x00002000** - Ignore errors associated with an expired certificate

You may set the value to 0x10000000 in order to regain strict certificate checking, 0x0000FFFF to disable certificate checking altogether, or any combination of the above values.

TLSSkipProtocols (REG_DWORD)

Introduced in v5.2.b.20160711. This is a bitmask of acceptable encryption protocols. The default value is 0x0AA0. Use a logical **OR** to combine multiple values.

- **0x00000002** - PTC1 (not recommended)
- **0x00000008** - SSL2 (not recommended)
- **0x00000020** - SSL3 (not recommended, but included in default for backward compatibility)
- **0x00000080** - TLS 1.0 (not recommended, but included in default for backward compatibility)
- **0x00000200** - TLS 1.1
- **0x00000800** - TLS 1.2
- **0xFFFFFFFF** - any available protocol (not recommended)

FQDN (REG_SZ)

Introduced in v5.2.b.20160711. This value contains the name to use during SMTP envelope negotiations; specifically, it is the name presented as the HELO or EHLO name immediately after receiving the server's greeting.

In previous versions, the name used was the sending machine's fully-qualified host name. However, workgroup members or machines just starting may only have a bald hostname available. This new value is set the first time an email is sent, and used thereafter for all subsequent emails. If a fully-qualified name is not discoverable, then Domain Time will use either a dotted-quad IP enclosed in brackets, or the computer name followed by .smtp.local. RFC 2821 section 4.1.1.1 requires one of these two forms. You may change the name if your particular email server requires an externally-verifiable DNS name to be presented.

As of v5.2.b.20170522, you may also customize the subject lines of your alert emails by making a change in the registry. See the [SMTP](#) section of the registry documentation.

IP Restrictions

Click the button to restrict which systems are allowed to contact Audit Server.

IP Restrictions

You may limit Audit Server's scope by specifying a list of IP address ranges. In most cases, you should limit the scope to your local subnets.

These restrictions apply only to auto-adding newly discovered machines and to recognizing the receipt of real-time alerts. A machine already on the audit list will be audited regardless of its IP address.

☐ No restrictions
☒ Permit only listed ranges
☐ Deny any in listed ranges

First IP in range	Last IP in range
192.168.10.0	192.168.10.255

[Remove](#)

First IP in range: . .
 Last IP in range: . .

[Add](#)

[OK](#) [Cancel](#)

IP Restrictions dialog [\[Click for larger size\]](#)

The **IP Restrictions** dialog applies to both machines sending [Real-Time Alerts](#) and also which machines are available to be [Auto-added to the Audit List](#).

You can both permit and deny access from IP ranges. To restrict a single IP address, enter the same IP address for both the First and Last range items.

Advanced Real-Time Alert Settings

Coalesce

Raise alert immediately

Group alerts and send no more often than once every minutes

These selections allow you to group your alerts together to prevent being overwhelmed by immediate alerts, or to receive them individually as they occur.

Record Backlog

If Audit Server is busy or the service is stopped, an alert backlog can develop. In general, old real-time alerts aren't real-time any more, so Audit Server will ignore all but the most recent ones.

Max backlog: records (range 1-10,000)

The **Max backlog** setting controls how many older queued alerts should be displayed when a backlog occurs. You shouldn't have to adjust this value unless your server is extremely busy and real-time alerts are regularly being dropped in normal use. If you set this value too large, you may have stale data appearing when a machine is rebooted.

Alert Sharing and the Alert Viewer

Alert Sharing

Audit Server can forward received alerts or status changes to the

DTAlert program for real-time desktop display of individual and overall status on multiple machines.

Alert sharing enabled Auto-Manage Windows Firewall
TCP Port (default port is 9910)

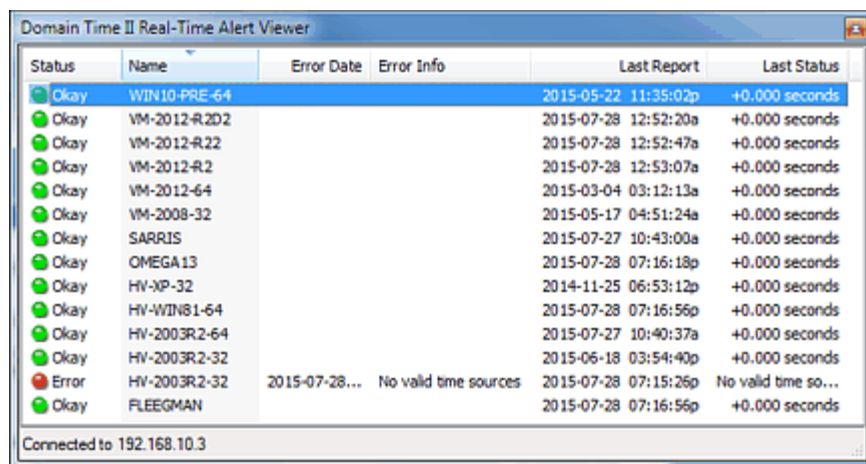
If this option is enabled, you may monitor the status of your Real-Time alerts using the Domain Time II Alert Viewer (see below). This port is also used if you are using the Audit Server [Standby Mode](#). If you are not using the Alert Viewer or Standby Mode, you may disable this option.

Force Auto-Manage Windows Firewall

As of Version 5.2.b.20150828, Domain Time supports automatic management of the Windows Firewall to allow access to the required time protocol and control ports. See [Auto-Manage Windows Firewall Settings](#) for a detailed explanation.

Domain Time II Alert Viewer applet (DTALERT.EXE)

Audit Server includes a handy **Alert Viewer** applet that can display the alert status from any/all Audit Servers on your network on any Windows desktop (XP and above) you'd like. It also gives you a handy customizable desktop clock display. This allows you to have a visual indicator of the status your entire time network on your desktop, or any other system where that information would be useful. It also makes an excellent desktop clock.



Status	Name	Error Date	Error Info	Last Report	Last Status
Okay	WIN10-PRE-64			2015-05-22 11:35:02p	+0.000 seconds
Okay	VM-2012-R2D2			2015-07-28 12:52:20a	+0.000 seconds
Okay	VM-2012-R22			2015-07-28 12:52:47a	+0.000 seconds
Okay	VM-2012-R2			2015-07-28 12:53:07a	+0.000 seconds
Okay	VM-2012-64			2015-03-04 03:12:13a	+0.000 seconds
Okay	VM-2008-32			2015-05-17 04:51:24a	+0.000 seconds
Okay	SARRIS			2015-07-27 10:43:00a	+0.000 seconds
Okay	OMEGA13			2015-07-28 07:16:18p	+0.000 seconds
Okay	HV-XP-32			2014-11-25 06:53:12p	+0.000 seconds
Okay	HV-WIN81-64			2015-07-28 07:16:56p	+0.000 seconds
Okay	HV-2003R2-64			2015-07-27 10:40:37a	+0.000 seconds
Okay	HV-2003R2-32			2015-06-18 03:54:40p	+0.000 seconds
Error	HV-2003R2-32	2015-07-28...	No valid time sources	2015-07-28 07:15:26p	No valid time so...
Okay	FLEEGMAN			2015-07-28 07:16:56p	+0.000 seconds

Connected to 192.168.10.3

[Domain Time Alert Viewer](#) [\[Click for larger size\]](#)

The Alert Viewer applet program is named DTALERT.EXE. The program comes in both 32 and 64-bit versions. If your Domain Time II Manager is 64-bit, the 64-bit version will be located in the **C: \Program Files\Domain time II** folder, and the 32-bit version will be in the **C: \Program Files\Domain time II\i386** folder. Conversely, if your Domain Time II Manager is 32-bit, the 32-bit version will be in the **C: \Program Files\Domain time II** folder, and the 64-bit version will be in the **C: \Program Files\Domain time II\AMD64** folder.

You may copy the DTALERT.EXE file to any machine you'd like (be sure to copy the correct 32 or 64-bit version to match the type of machine), and then run it to configure the clock display and current alert status of your Audit Server(s). You may run as many copies of DTALERT.EXE on various machines as you need.

The program will display the current date and time on your desktop along with a colored flag representing the current alert level. The Alert Viewer shows the current overall alert status present on your monitored Audit Servers. The flag next to the clock will change color to reflect the worst reported status of any monitored system (green, yellow, or red). A white flag indicates the software cannot contact any Audit Servers. As of 5.2.b.20170101, you can also enable audio alerts to be notified by a sound when the status changes.

Double-click any part of the clock display to show the alert status of the individual machines providing Real-Time alerts to the monitored Audit Server(s).

Note: The software only reports alert status. To reset or configure alerts, you must use the Manager on the the Audit Server machine(s) actually collecting the [Real-Time alerts](#).

To configure the program options, run DTALERT.EXE and right-click on any part of the clock to display the context menu.

- **Clock** - These context menu items let you control the appearance and function of the desktop clock display. You can set attributes such as font, color, background, opacity, etc.
- **Status** - These settings control the display of alert data from your selected Audit Server(s).
 - **Visible** - sets whether the Real-Time Alert Viewer status windows is open and visible. This window displays the status of all machines reporting real-time alerts to your selected Audit Server(s). You can toggle whether this window is open by double-clicking on any part of the clock display.
 - **Servers** - This is where you tell the viewer which Audit Server(s) you want it to monitor for alerts. Enter the DNS Name or IP address of each Audit Server. Note that **Alert Sharing** over port 9910 TCP must also be enabled on each listed Audit Server (see above).
 - **Date/Time Format** - This lets you set the format for all dates and times displayed on the status viewer.
- **Start at Logon** - When this item is checked, the Alert Viewer will automatically load whenever you log in.

Auto-Acknowledgement of Resolved Alerts

Auto-Acknowledgement of Resolved Alerts

Real-time alerts that resolve themselves change from red to yellow, and normally stay yellow to let you see that an error had occurred. You may acknowledge warnings using Manager, or have Audit Server do it automatically.

Auto-Acknowledge enabled

Wait: minutes after last error occurred.

If this option is enabled, Audit Server will automatically acknowledge Real-Time Alert warnings for machines that are not still in an error state. The machines will return to green status after the period of time you specify here. If unchecked, machines that had an issue will stay in the yellow warning state until manually acknowledged using Domain Time Manager.

Data Collection

Audit Server can collect a variety of data from audited systems. It can then present this data to you in different ways to help monitor the health of your network time, interface with other systems, and for regulatory compliance purposes.

Audit Data Types

Audit Server collects two main types of data from audited systems. You can view any type of data collected by Audit Server by clicking on the "Audit Server" category in the left-hand column of [Manager's interface](#)

- ▶ **Audit Results Records** - a snapshot record of information about the audited system taken at the time of the audit. It contains information such as current time on the target, it's last time source, etc. [Daily Reports](#) can be generated based on the Audit Result Records during each audit run to summarize and/or export data.
- ▶ **Synchronization (Drift) logs** - a running log of time deltas (either reported or measured, depending on the type of audited time client and protocol).

Ouput

Domain Time can output collected data in several formats. Collected Synchronization logs may be viewed using the [Drift Viewer](#) and/or converted to CSV or text formats. The Audit Logs from audit runs can be output to [Syslog](#) servers. And, the [Daily Reports](#) function allows you to create a customized report of your audit runs and results. You may also send [email summaries](#) of audit runs.

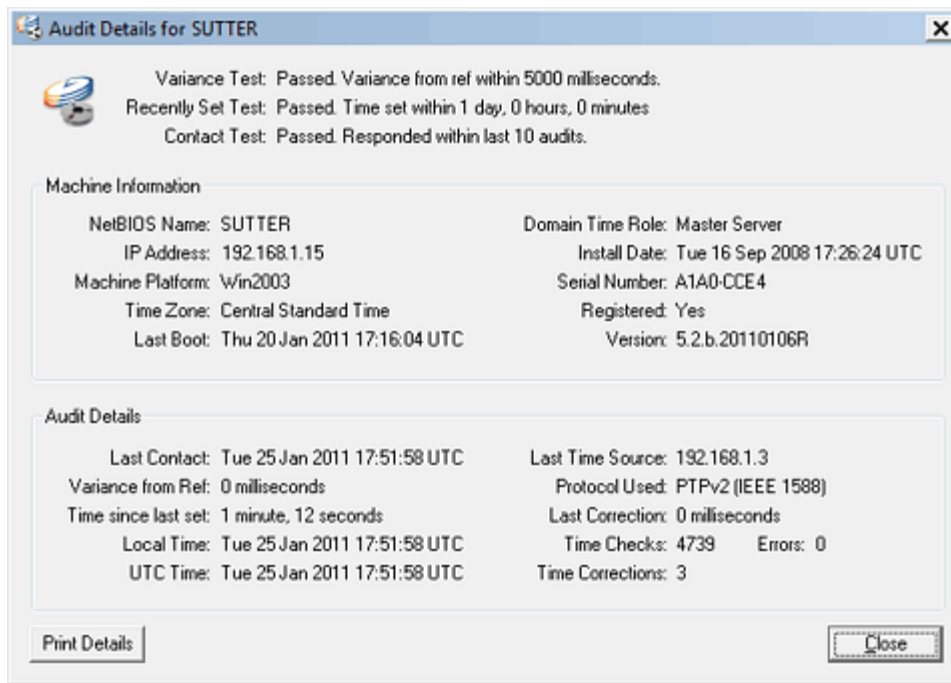
Audit Results Records

Audit Results Records are highly compact collections of data collected from [audited systems](#) during an audit run.

Audit Results records can be gathered from

- Domain Time systems on Windows
- domtimed daemons running on 'nix systems
- From machines running NTP daemons that respond to time request packets
- From machines running PTP (IEEE 1588-2008/2019) that respond to management messages via [PTP Monitor](#)

Domain Time machines will provide more complete statistics on their operation than do NTP sources, but both contain enough data for auditing/alerting purposes.



[Audit Result Record Details](#) [\[Click for larger size\]](#)

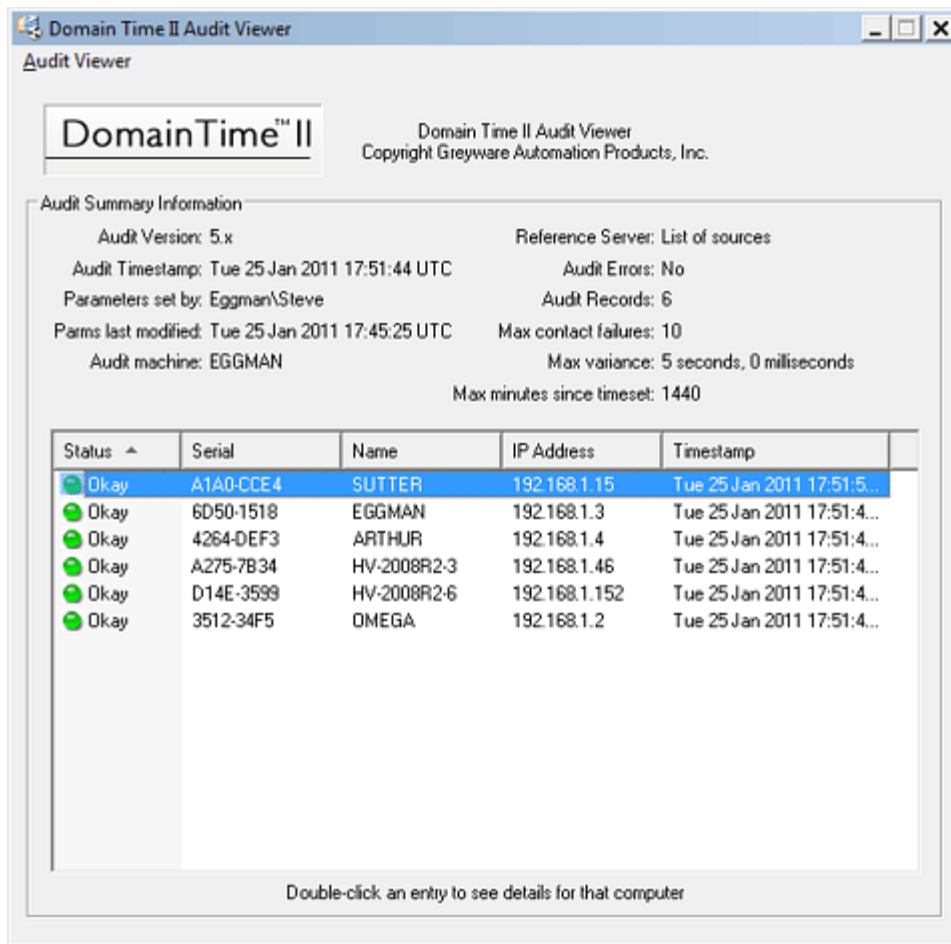
The data from all audited machines during an audit run are collected into audit results files kept in the Audit Data Cache folder. The folder location is specified on the Audit Server -> Advanced -> Data Folders... menu (by default, **C:\Program Files\Domain Time II\Audit Data Cache**). Each audit run results in a new file.

Disk requirements

Although each individual audit record file is highly compact, the size of each individual audit results file depends on how many machines are included. If unattended (and unlimited by the Audit Tasks setting described above), the folder can grow to contain a very large amount of data. You should plan to regularly archive this data off to your normal archival storage.

Please see the [Audit Disk Space Estimator](#) page to calculate your disk space requirements for storing audit data.

You may configure Audit Server to limit the growth of the Audit Data Cache by deleting all Audit Results Records over a certain age. This setting is found on the *Audit Server* -> *Audit Tasks* menu item.



Audit Data Viewer [\[Click for larger size\]](#)

Audit Results Records are viewed using a utility program called Audit Data Viewer (DTREADER.EXE-launched automatically when you view results through Manager). The DTReader utility is associated with files having the extension .dtad (DT Audit Data) during the installation of Audit Server. However, it can be used to view Audit Record Results files on other systems. Simply copy the program to any machine from which you'd like to view audit records.

Note: DTREADER.EXE does not function on Windows Server Core systems. To view sync logs collected by Audit Server on Server Core systems, you must copy the DTREADER.EXE utility to a non-Core system and use it from there to view the .dtad audit records through a network share on the Core machine.

Synchronization (Drift) Logs

Audit Server can collect or generate several types of synchronization logs from audited machines into a central location where they can be reviewed, maintained, or archived off for data retention purposes.

Logs collected from Domain Time Servers and Clients:

- **Domain Time II Synchronization logs** - a running log of the results of each successful DT2 or NTP time synchronization (or sample aggregation if using PTP) by the Domain Time Server or Client.

Domain Time II Synchronization logs are only available from Domain Time Clients or Servers. The logs are kept locally on each Server or Client, but are copied and appended to the Audit Server data store during an audit run. The time deltas contained in these records show the results of time corrections applied by Client or Server to match its configured time source(s). This information is reported by the audited machine itself and reflects its perspective of time accuracy as compared to its sources.

Machines selected to be audited on either the *Domains and Workgroups* or *Domain Time Nodes* lists in Manager will

collect this type of synchronization log. Domain Time Synchronization log filenames begin with the Domain Time Serial Number and have the file extension **. dt**

- **Domain Time II PTP Offset Synchronization logs** - a running log of the reported offset between the Domain Time Slave and its PTP Master (using the PTP protocol).

Domain Time II PTP Offset Synchronization logs are only available from Domain Time Clients or Servers. The logs are kept locally on each Server or Client, but are copied and appended to the Audit Server data store during an audit run. The time deltas contained in these records show the results of time corrections applied by Client or Server to match its PTP master. This information is reported by the audited machine itself and reflects its perspective of time accuracy as compared to its sources.

Machines selected to be audited on either the *Domains and Workgroups* or *Domain Time Nodes* lists in Manager will collect this type of synchronization log (if also enabled on the Synchronization Log configuration dialog). Domain Time PTP Offset Synchronization log filenames begin with the Domain Time Serial Number and have the file extension **_ptp. dt**

Notes:

Domain Time II Synchronization Logs can only be collected from Domain Time II Server and Clients version 3.1 and later. Domain Time II PTP Offset Logs may only be collected if all components (Audit Server/Manager and Client/Server) are version 5.2.b.2015037 or later.

Both of these logs are limited in size on the Client or Server and older data scrolls off over time. Using Audit Server to collect this information allows you to preserve this data for audit trail and archival purposes. Note, in order to have a complete central record, you must Audit the machines often enough to collect the data before it scrolls off on the individual machines.

Connecting to Domain Time versions prior to 5.2, the Audit Server must use credentials with sufficient rights to connect to the administrative shares on the remote systems to collect drift logs. Current versions obtain the data using direct communication over Port 9909 UDP/TCP.

Drift Logs generated by tracking other systems:

- **NTP Server Drift logs** - a running log of the time deltas of audited NTP machines measured at the time of each audit run or collected on a set polling schedule.

NTP Server Drift logs can be collected from any machine that can be configured to respond to standard NTP time requests (such as ntpd on Linux). Drift files for each audited NTP machines are created/appended to during audit runs or on a set polling schedule. The time deltas contained in these records show the measured difference between the NTP timestamp replies and Manager's configured [Reference Time](#) source(s).

Machines selected to be audited on Manager's *NTP Nodes* list will collect this type of synchronization log. NTP Server Drift log filenames begin with **NTP Server** and end in **_ntp. dt**

- **PTP Node Drift logs** - a running log of the deltas (either reported or measured) of audited PTP Nodes.

As of Domain Time version 5.2.b.20170101, PTP Node Drift logs can be collected from any machine that is discovered by [PTP Monitor](#). Drift files for each audited node are created/appended to during audit runs.

When collecting PTP master data, the delta reported is the measured difference between the Master's announced time and Manager's configured [Reference Time](#) source(s). When collecting PTP slave data, the delta reported is the reported offset between the slave and its master. See the [Offset measurement](#) section of the PTP Monitor documentation for details.

Machines selected to be audited on the *PTP Nodes* list will collect this type of synchronization log (if also enabled on the

Synchronization Log configuration dialog).

Note:

The auditing of PTP Nodes is a separate function from other types Audit Server auditing. The "Audited" setting's column in Manager's display for PTP Monitor is independent of the "Audited" settings on the *Domains & Workgroups*, *NTP Nodes*, *Domain Time Nodes*, or *Real-Time Alerts* displays. Enabling/Disabling auditing on the PTP Monitor display will not change the audit settings on the other pages, and vice versa.

PTP Nodes Drift log filenames begin with **PTP Node** and end in **_ptp.dt**

Synchronization Log Collection Settings

Use the *Audit Server -> Synchronization Logs -> Configure* menu item to bring up the Synchronization Configuration Dialog. Alternately, you may right-click the *Audit Server\Synchronization Logs* item in Manager's Tree and choose *Configure...* from the context menu.

Configure Synchronization Logs

Synchronization Log Collection

Audit Server can collect synchronization logs (drift graph data) from each audited Domain Time machine. For audited NTP Servers or PTP Nodes, Audit Server can create drift data points.

☒ Enabled

☒ Foreground - collection must finish before audit completes
☐ Background - collection finishes independent of scheduled audits

☒ Run background collection every 30 minutes (range 5-1440)
☒ Collect NTP drift on same schedule as other drift records
☐ Collect NTP drift every 120 seconds (range 10-3600)

☒ Collect PTP sample data from audited Domain Time machines
☐ Collect PTP node data from audited PTP Monitor masters
☐ Collect PTP node data from audited PTP Monitor slaves

Conversions

☐ None
☒ CSV - save drift records in Daily Drift .csv files
☐ TXT - expand each binary .dt file to a text file (deprecated)

☒ Limit size of collected binary (.dt) synchronization log files

☐ Allow any number of records per machine
☒ Never keep more than 5000 records (range 1-604800)
☐ Allow records of any age
☒ Delete records older than 45 days (range 1-5000)

Log filename format: Default (recommended)

[Synchronization Configuration Dialog](#) [\[Click for larger size\]](#)

Foreground - collection must finish before audit completes

Background - collection finishes independent of scheduled audits

Run background collection periodically, not just at audit time

These choices determine whether Audit Server will collect the sync logs in a separate thread from the audit run itself. Collecting sync logs from each audited machine can take an extended amount of time, particularly if you have a large number of machines to audit. Choosing **Background** allows collection of the basic audit data very quickly, and then the collection of the sync logs can complete in the background. Running the collection in the background periodically can make collection even more efficient.

Collect NTP drift on same schedule as other drift records

Collect NTP drift every seconds (range 10-3600)

These settings (new as of v5.2.b.20171113) allow you to set a schedule of how often audited NTP Nodes are polled for their offset data, allowing you to create NTP sync logs with more data points than would normally be collected from

regularly scheduled Audits alone. This can present a more comprehensive picture of NTP Node synchronization.

Collect PTP sample data from audited Domain Time machines

Available on version 5.2.b.20150307 or higher. If the target machine is synchronizing using PTP, its PTP offset logs can be collected at the same time as the regular drift synchronization logs. PTP offset log collection is subject to the same limits and schedule as regular drift log collection.

Collect PTP node data from audited PTP Monitor masters

Available on version 5.2.b.20170101 or higher. If enabled, PTP masters selected to be audited on the *PTP Nodes* list in Manager will have a drift file created for them. PTP master node offsets are calculated with each received Sync or Sync/Followup (typically once per second). Data is buffered internally by Audit Server server before being written to the file. Log collection is therefore more or less continuous for PTP master nodes, but it still subject to the same limits as regular drift log collection.

Collect PTP node data from audited PTP Monitor slaves

Available on version 5.2.b.20170101 or higher. If enabled, PTP slaves selected to be audited on the *PTP Nodes* list in Manager will have a drift file created for them. A new data point is generated with each commanded or scheduled sweep (typically once every 30 seconds), and are buffered internally by Audit Server before being written to the file. Log collection for PTP slaves is subject to the same limits as regular drift log collection, but does not follow the same schedule.

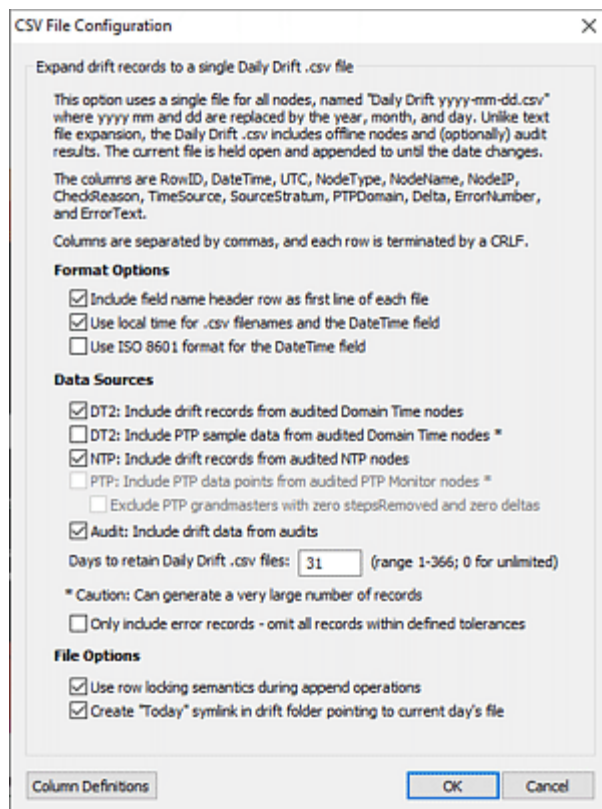
Conversions

As of version 5.2.b.20180303, Audit Server has two options for auto-converting binary .dt sync logs into text formats suitable for import into other database or analysis programs: CSV or TXT conversion. Note, the TXT method is now deprecated and should be avoided.

CSV - save drift records in Daily Drift .csv files

Due to the limitations of .txt file conversions, most notably the lack of data points for offline nodes and the high CPU/disk usage required for conversion, Domain Time now supports a Daily Drift CSV file as the preferred alternative to expanding individual .dt files to .txt files.

A Daily Drift file is named *Daily Drift yyyy-mm-dd.csv*. As the name implies, Audit Server creates one Daily Drift file per day. You may configure Audit Server to keep Daily Drift files forever, or to keep only the last *n* days' worth of files. Each file is organized as a comma-separated value file, with each row representing a single data point.



CSV File Expansion [\[Click for larger size\]](#)

Include field name header row as the first line of each file

If enabled, each Daily Drift file will begin with a header: *RowID, DateTime, UTC, NodeType, NodeName, NodeIP, CheckReason, TimeSource, SourceStratum, PTPDomain, Delta, ErrorNumber, ErrorText* (spaces after commas added in this file so word-wrap can occur; there are no spaces after commas in the .csv file itself). The exact syntax and meaning for each column is described fully on Manager's CSV settings dialog. Click the [Column Definitions](#) button to see the full definition of the columns.

Use local time for .csv filenames

If selected, Audit Server will use local time for the *yyyy-mm-dd* part of the filename. If not selected, Audit Server will use UTC for the *yyyy-mm-dd* part of the filename. Audit Server switches to a new file at midnight, either local time or UTC, based on this option. Additionally, this option controls the value of the DateTime field and the value of the UTC field within the .csv file. The default format for the DateTime field is *yyyy-mm-dd hh:mm:ss*. If you are using local time, the DateTime field will be local time, and the UTC field will be *N*. If using UTC, the DateTime field will be UTC, and the UTC field will be *Y*.

Use ISO 8601 format for the DateTime field.

If selected, the DateTime field format will be either *yyyy-mm-ddThh:mm:ssZ* if using UTC, or *yyyy-mm-ddThh:mm:ss±HH:MM* if using local time.

You may configure the data sources for the Daily Drift file:

DT2: Include drift records from audited Domain Time Nodes

If selected, drift records from audited Domain Time nodes will be included. This is the data you see on any individual machine when you click the [Drift...](#) button on the Control Panel applet. Domain Time generates one drift record at each check interval, so, for example, if you are checking the time once a minute, Domain Time will generate 1440 data points per day.

DT2: Include PTP sample data from audited Domain Time Nodes

If selected, PTP sample data from audited Domain Time nodes will be included. Domain Time nodes generate one

data point per PTP sync, so this category can produce a large amount of data. This data you see on any individual Domain Time machine when you click the [Graph](#) link on the Control Panel applet. If your PTP master is sending one sync per second, Domain Time will generate 86,400 records per day.

NTP: Include drift records from audited NTP nodes

If selected, drift data from audited NTP nodes will be included. A row will be added for each audited NTP node at each NTP drift collection period. If an audited NTP node is offline, a row will be added with the appropriate error number and description.

PTP: Include PTP data points from audited PTP Monitor nodes

If selected, and if PTP Monitor is enabled, drift data points from audited PTP nodes will be included. For PTP masters, this is one data point for each Sync. For PTP slaves, this is one data point for each PTP sweep. The interval between sweeps is configurable. This category can produce a very large number of records. For example, if you are sweeping the network once per minute, *each* slave will generate 1440 data points per day. If you are monitoring two PTP masters, each sending Syncs once per second, *each* master will generate 86,400 records per day. As with NTP nodes, a non-responding PTP node will generate a record with the appropriate error number and description. A record is also generated when a PTP master downgrades its time quality, changes its leap second advertisement, or goes offline.

Audit: Include drift data from audits

If selected, one record will be generated for each audited node, regardless of type, immediately after an audit completes. This information is essentially the same as the audit summary, but configured to match the format of the other data sources. If a node is offline, or exceeds your chosen tolerance, the record will include an error code and description.

Only include error records - omit all records within defined tolerances

Introduced in version v5.2.b.20190701, this option only includes synchronization log events that exceed the Post-Audit Alert tolerances specified in each audit group (see [Alerts and Audit Groups](#)). Use this option with care, since all other synchronization events will be omitted from the Daily .CSV file.

■ **Daily Drift CheckReason Entries**

The CheckReason field of the Daily Drift contains information about why each time check occurred. Here are possible entries:

CheckReason	Explanation
AuditSignal	Triggered by Audit Server
ClockChange	Sync because another user or process has changed the clock
ControlPanel	Triggered from the Domain Time applet
DT2BCast	Time received by DT2 protocol broadcast
DT2MCast	Time received by DT2 protocol multicast
Heartbeat	Triggered by a broadcast/multicast heartbeat
IndieSignal	Triggered by a Domain Time Server in the Independent Server Role
LeapSecond	Triggered after a leap second adjustment
MasterSignal	Triggered by a Domain Time Server in the Master Role
NewIPAddress	Sync due to the OS signal of a change in the IP address stack

NTPBCast	Time received by NTP protocol broadcast
NTPMCast	Time received by NTP protocol multicast
PowerResume	Sync due to the OS signal of a resume from sleep/hibernation
PTP	Time received by PTP
PTPFirstTS	Time set to match the first received PTP timestamp
PTPMaster	Sync due to a change of PTP Master
PTPSlave	Sync due to becoming a PTP Slave
PTPSweep	Sync triggered by a PTP Monitor sweep
Scheduled	Sync occurred as scheduled by the applet
SlaveSignal	Triggered by a Domain Time Server in the Slave Role
Startup	Sync due to startup of the Domain Time service
SyncTrigger	Triggered by a command from a Domain Time component
TimezoneChng	Sync due to a signal from the OS of a change to the timezone
Training	Sync occurred during an accelerated clock training sequence
VeracityCheck	Response when auditing an NTP Node

■ Daily Drift Error Codes

An error code of zero means no error. Otherwise, the error code represents the reason for failure, and the error text describes the problem. If the error code is non-zero, the Delta field will always be 0.0000000 (no delta information available), and the TimeSource field will be "" (an empty quoted string). When you are processing the .csv file, make sure that you don't accept the delta of 0.0000000 as valid if the error code is non-zero.

All error codes are standard Microsoft error codes, and you may look them up online if you are not familiar with them. For example, 10064 is the Winsock error WSAEHOSTDOWN ("Host is down") and will usually appear with the error text "Node offline."

A special case is error code 1246, which translates to Microsoft's ERROR_CONTINUE ("Continue with work in progress"). This code is used for informational messages when a PTP master upgrades its quality of service or changes its leap second flags. The error text will describe the change, but 1246 should not be considered an error.

■ Daily Drift CSV Operation

Audit Server caches records internally, then flushes them to disk approximately once every two minutes. The current Daily Drift file is held open by Audit Server during the entire 24-hour period of data collection. Other processes may read from, but not modify, an open Daily Drift file.

By default, Audit Server creates a symbolic link in the drift folder named **Today** (with no file extension). This is a soft link pointing to the current day's file. You may disable this functionality by unchecking the checkbox on the Daily Drift CSV settings dialog. The Today link is primarily useful for automated processes that import the .csv file into a dbms or other storage and that only want today's information.

If you are importing Daily Drift files into your dbms or other system throughout the day, use the RowId field to know whether or not a record has already been imported. The RowId field begins at 1 and increments by 1 for each row.

If you are collecting DT2 drift, or DT2 PTP sample data, be aware that the collection schedule may affect the previous day's Daily Drift file. For example, if you are collecting data once per hour, the data collected at 12:15am will contain records from the previous day. Audit Server will put yesterday's records in yesterday's file, and today's records in today's file. Audit Server will open and close the Daily Drift files as required.

Audit Server normally requests only new records from audited Domain Time nodes. It remembers the last collection, and asks for records newer than that. When you change an unaudited Domain Time node to audited, or when you add a node that has never been audited before, Audit Server will request all existing records. The returned data may span several weeks, depending on how often the Domain Time node is configured to check the time. This can create a proliferation of Daily Drift .csv files from the recent past. The RowId field will continue from when each Daily Drift file was last extended, even if the Daily Drift file itself has since been deleted.

■ Best Practices

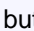
You should set the drift folder to a local disk with sufficient storage to hold all the .dt files and the .csv files. A Daily Drift .csv file may be hundreds of megabytes. The default location is C:\Program Files\Domain Time II\Synchronization Logs. You may reconfigure the location from within Manager. Use the Audit Server menu, Advanced, Data Folders dialog to change the location. Manager will create the new folder(s) as necessary, and move the data from the old folder(s) to the new one(s). Manager briefly stops Audit Server to ensure that all data is flushed to disk and no files are held open, performs the move, and then restarts Audit Server.

For your dbms to consume the .csv information, we recommend that you run your import procedure against each Daily Drift .csv file (file mask "Daily Drift ???-??-???.csv") in the drift folder, using the RowId to determine if any row for a particular day is new. You may delete (or archive) old .csv files after importing them. If Audit Server needs to recreate a deleted Daily Drift file, the RowId field will continue from where the deleted file left off.

You cannot delete the current day's Daily Drift file. Audit Server allows the current file to be read, but not modified or deleted. Before appending new records, Audit Server locks the portion of the file beyond the current end, writes the new data, and then releases the lock. If this locking behavior interferes with your import procedure, you may disable it by unchecking the checkbox on the Daily Drift CSV setting dialog.

TXT - expand each binary .dt file to a text file (deprecated)

Enabling this function will cause Audit Server to create a text file version of the binary sync log collection file(s). The text files will be named and formatted according to the settings indicated. This function is deprecated as of v5.2.b.20180303.

Note: This process is very expensive in both disk and processing overhead. You can easily end up with gigabytes of textual data, taking excessively long times to add even a single new data point. You should only use this option if you require a text file be kept for a specific purpose, since the text files are dramatically larger than the binary files. Normally, you would use the **View Logs** function described below to view the binary files in a more friendly graphical format and generate a text file only if necessary by clicking the  button on the Drift Graph display.

IMPORTANT: As of v5.2.b.20171113, Audit Server will no longer attempt to create textual versions of drift files with more than 64K records (approximately 6MB). If this occurs, a warning message will appear in Audit Server's log. To avoid issues, you should ensure your sync logs do not exceed this number of records (see **Limit size of collected Synchronization Logs** section above). Consider archiving off the sync logs (and any .txt expansions) to another location on a regular basis.

In more recent versions (as of v5.2.b.20180303), you may use the DTDRIFT.EXE program to edit the size of or repair the .dt binary drift logs. See the [Viewing/Managing Collected Logs](#) section below.

Limit size of collected Synchronization Logs

You should restrict log size by limiting the number of records kept per machine (older records are rolled off to make room for new entries), and/or by deleting all records over a certain age.

Audit Server provides two ways to limit the size of .dt files.

- **Never keep more than n records**

If selected, Audit Server will periodically prune the earliest records from each .dt file.

- **Delete records older than n days**

If selected, Audit Server will periodically examine each .dt file and discard records older than n days

You may use both 1. and 2. together, or neither of them. If you use neither, the .dt file will keep growing as new data points are added. In versions of Domain Time prior to 5.2.b.20171113, there was no cap on the size. Starting with version 5.2.b.20171113, Audit Server enforces a cap of 604,800 records (enough for one data point per second for a full week). When a .dt file grows beyond this size, Audit Server archives the oldest records by placing them in the \Synchronization Logs\Archives subfolder, allowing the original .dt file to begin growing again.

Disk requirements

Although the binary synchronization logs files are quite efficient at recording individual delta events, the overall disk space needed depends on how many machines are being collected and how often events are being recorded in each type of file.

- **Domain Time II Synchronization Logs**

A data point is written for each successful time synchronization (or PTP aggregation). The overall schedule for these is set by the "Timings" settings on the Server or Client, however, other events can trigger additional synchronizations. Examine a representative sample of machines in normal operation to determine the number of records you'll require.

- **Domain Time II PTP Offset Logs**

A data point is written for each received Master sync packet. The schedule for this is determined by the PTP Master sync schedule (often 1/sec).

- **NTP Drift Logs**

A data point is written each time an audit is run. The schedule for this is set by Audit Server.

- **PTP Nodes Drift Logs**

A data point is written each time an audit is run. The schedule for this is set by Audit Server.

IMPORTANT:

As of v5.2.b.20171113, the graphical viewer for synchronization logs will not open files larger than 604800 records. You should limit the size of these files and regularly archive them to ensure they stay under this limit.

As of v5.2.b.20171113, the maximum number of records kept by Audit Server is 604800, enough for one data point per second for one full week (just over 12MB of binary data). If you have disabled the size limit for drift files, a file that grows larger than the maximum will be moved into an "Archives/yyyymmdd" subfolder, and then the file restarted. A warning message will appear in Audit Server's text log. If the file cannot be archived, an error message will appear in the text log and the older data will be lost.

If unattended (and unrestricted by the [Limit size of collected synchronization logs](#) setting), the folder can grow to contain a very large amount of data. You should plan to regularly archive this data off to your normal archival storage. "Regularly" may mean several times a day, depending on your data collection rates.

Please see the [Audit Disk Space Estimator](#) page to calculate your disk space requirements for storing synchronization logs.


Log filename format:: Default (recommended)

Sets the way sync log filenames are constructed. The default format is **Serial - Name**

Viewing/Managing Collected Logs

You may view collected Synchronization logs by expanding the Audit Server item in the Manager tree and clicking on the *Synchronization Logs* item. You may also choose the *Audit Server -> Synchronization Logs -> Open Containing Folder* menu item.

- Synchronization logs are collected in the folder specified for them on the *Audit Server -> Advanced -> Data Folders...* menu item.
- View the collected logs in graphical format by choosing *Audit Server -> Synchronization Logs -> View Drift Graphs...* from the menu. Filenames for PTP records will end in "_ptp", otherwise they are standard drift log files.



The Drift Graph [\[Click for larger size\]](#)

Synchronization Logs are viewed using a utility program called DTDRIFT.EXE (launched automatically when you view results through Manager). The DTDrift utility is associated with files having the extension .dt during the installation of Domain Time Server, Client, and/or Audit Server. However, it can be used to view Synchronization Log files on other systems. Simply copy the program to any machine from which you'd like to view audit records.

Note: The DTDRIFT.EXE program does not function on Windows Server Core systems. To view sync logs collected by Audit Server on Server Core systems, you must copy the DTDRIFT.EXE utility to a non-Core system and use it from there to view the .dt synchronization logs through a network share on the Core machine.

As of v5.2.b.2017113, the graphical viewer for drift graphs will not open files larger than 604800 records. Ensure you are limiting your Sync logs to less than this size.

The DTDRIFT.EXE program can also run as a command-line program. Use **DTDRIFT -help** from the command line to see all of the options. One of the things DTDRIFT.EXE can do is convert binary .dt files to their .txt equivalents, just as if you had opened the files in the graphical viewer and clicked the **Raw Data** button. As of v5.2.b.20180303, it can also

be used to chop large binary .dt files in to manageable chunks or repair damaged .dt files.


Added as of v5.2.b.20180101:

- **-convert** [-localtime] filespec. -localtime is optional. If not supplied, UTC will be used. filespec may be either a fully-qualified path and filename, or a path with *.dt (no other file extensions are supported). If the path or filename has spaces, you must enclose it in quotation marks. For example, dt drift -convert "C:\Program Files\Domain Time II\Synchronization Logs*.dt" will convert each .dt file in the named folder to its .txt equivalent. The original .dt file is not altered.

Added as of v5.2.b.20180303:

- **-chop** command-line parm. It must be followed by the full path to a .dt file, or use the wildcard path*.dt (much the same as for -convert). While -convert will read a .dt file and create the corresponding text version, -chop will split the .dt file into chunks named foo_Part001.dt, foo_Part002.dt, etc.
- **-repair** command-line parm. It must be followed by the full path to a .dt file, or use the wildcard path*.dt. The -repair switch examines the file(s) for invalid entries and removes them. Note: This is a prophylactic function; no .dt file has ever become corrupted.
- **-csv** command-line parm. This switch is only valid with -convert, and may optionally be combined with -noheader. The switches must be followed by the full path to a .dt file, or use the wildcard path*.dt. Example: dt drift -convert -csv -noheader "d:\drift files*.dt" or dt drift -convert -csv C:\myfile.dt. The .csv file(s) will be created in the same folder as the .dt file(s).

Hint: Double-click on any part of the graph to bring up limit markers handy for seeing the range of deltas displayed. Note you need to be zoomed in enough to see actual variations in the graph.

Click the  button to see the underlying statistical data and individual records used to create the graphical display.

- If you have chosen to expand the binary logs to text files (see the configuration option above), you can view the text versions by choosing *Audit Server -> Synchronization Logs -> View Text Reports...* from the menu.

Daily Reports

Daily Reports are summary results files using a user-specified format, created during each audit run from audit record data. They are particularly useful for exporting audit data to external programs.

You set up Daily Reports using the *Audit Server -> Daily Reports -> Configure* menu item.

Daily Report

Audit Server can keep a report of all audit successes and failures during a 24-hour period. Each audit during the day appends to the daily report. A new report is started at the first audit after midnight.

Because the daily report covers a one-day period, the filename is composed from the date. You may configure the filename format.

Specify the date format string and extension (.txt, .html, etc.) for the daily report filenames:

.txt Available tokens: d, dd, ddd, dddd, M, MM, MMM, MMMM, Y, YY, YYYY
Sample filename: 2016-10-16.txt

Daily Report Format

Optional Headers

- ☒ Comments
- ☒ Local time of audit
- ☐ UTC time of audit
- ☒ Field names

You may either type directly in the box above, or select items from the list below:

%%	(Percent)
\	(Backslash)
\n	(Line Feed)
\r	(Carriage Return)
\r\n	(CRLF)
\t	(Tab)
AuditStampVersion	Audit stamp version number
AverageInfo	List of servers used for averaging
Checks	Number of time checks (whether or not correction made) since startup
ContactFailures	Number of consecutive contact failures

☒ Times from fields in Daily Report show as local time

Audit Server Daily Report Configuration Dialog [\[Click for larger size\]](#)

When enabled, Audit Server will create a special summary log of audit records each day in the folder specified for Daily Reports on the *Audit Server -> Advanced -> Data Folders...* menu item. Click the *Audit Server -> Daily Reports -> View* menu item to browse through the existing reports.

Notes:

Daily Audit Summary Logs only include information from audit records; they do not include information from the Synchronization Logs.

The **View Logs** button displays the contents of the Daily Report Summary collection folder using the Explorer shell which does not function on Windows Server Core systems. Use Notepad to view the files manually or view them from a remote machine using any text reader.

A new summary log file will be created each day. Any audits performed during that day will be appended to the log.

Daily Reports are particularly useful if you are using your own log file collection and analysis program and need the audit record information to appear in a particular format to be imported correctly.

You may specify the date format and extension to be used in creating filenames. The default extension is .txt, however, as of version 5.2.b.20160922, you may specify .htm or .html which will wrap the output in minimal HTML tags sufficient for viewing with a browser.

The **Daily Report Format** section is where you specify how data will appear in the log. You can specify the format of the header used before the records as well as the format of the records themselves.

The format string entered in the text field indicates the order of data variables (keywords surrounded by the % character) which represent specific data collected from the audited machine, special characters (such as \r representing a carriage return), and delimiters (if any) used to create each line of the log file. You can preview the effect of your settings by clicking the **Show Example** button.

For example the format string:

%Status%, %MachineName%, %IP%, %DST%, %TimeZone%\r\n

results in a log file entry with this format:

```
#
# Audit results from audit performed at 17:00:00 UTC
#
# Status,MachineName,IP,DST,TimeZone\r\n
OK,DC_2,172.10.1.12,Y,Central Daylight Time
OK,PDC,172.10.1.10,Y,Central Daylight Time
OK,NTP Server,192.43.244.18,?,Unknown
```

Note that the entry for the NTP server in the example above shows ? in the DST and Unknown in the TimeZone fields. This information is only available from Domain Time II components.

These are the items that can be included in the format string:

Delimiters

You may specify any text you want to use between variables in the format string.

Special Characters

\n	line feed
\r	carriage return
\t	tab character
\\	backslash character
%%	percent sign character

Data Variables

%Status%	Whether or not the machine was audited successfully Returns OK or Err
%AuditStampVersion%	Audit stamp version number
%ContactFailures%	Number of consecutive contact failures
%SecsSinceLastSet%	Number of seconds since time was last set
%Variance%	Variance from reference at time of audit
%LastContact%	Time this machine was last contacted
%SerialNumber%	Machine's serial number
%LastProtocol%	Name of last time protocol used to set the time
%LocalTime%	Local time (adjusted for timezone and dst) at time of audit
%UTC%	UTC time at time of audit
%LastVariance%	Variance last time machine corrected its time

%Corrections%	Number of time corrections since last startup
%Checks%	Number of time checks (whether or not correction made) since startup
%Errors%	Number of times machine encountered an error while checking the time
%InstallDate%	Time this machine's client was installed
%UnixTime%	Time (in seconds) at time of audit (usually matches LocalTime)
%LastSet%	Time machine last corrected its time
%LastStartup%	Time machine last started the time service
%LastSource%	Most recently-used time source
%TimeZone%	Time zone (for example, "Eastern Standard Time")
%Version%	Version number of time software on machine
%MachineName%	Machine's NetBIOS name
%DNSName%	Machine's DNS name (if available)
%IP%	Machine's last-known IP address
%DST%	Y if machine is known to be applying Daylight Savings Time correction N if machine is known to NOT be applying DST correction ? if machine's treatment of Daylight Savings Time is unknown
%Role%	Machine's Domain Time II role (client, server, etc)
%Registered%	Y if software is registered N if software is an evaluation copy (or not a Domain Time component)
%OS%	Name of architecture, operating system, and OS version
%AverageInfo%	List of servers used for averaging (if available)

Audit Disk Space Estimator

The information on this page will help you estimate how much disk storage to allow for Audit Server's operations. There are two main types of data that can be collected by Audit Server:

- **Audit Records** are snapshot summary records with details about each audited machine taken at the time the Audit is performed.
- **Synchronization Logs** are collected historical records of each individual time check performed by a Domain Time Client or Server. They may consist of the Drift and/or PTP Offset logs.

Audit Record Disk Space Calculator

This calculator lets you estimate the storage space required for your audit records. Note that this calculation only estimates space for audit records; see the section below for information on estimating synchronization (drift) logs space.

1. Enter the **number** of machines you want to audit, and how often you will be auditing them.

How many machines per audit?	
How many audits per day?	
How many days per week?	

2. Then click the **Calculate** button....

Total Audit Records per Week:	30 (45.0 KB)
Total Audit Records per Month:	129 (194.0 KB)

You can configure Audit Server to automatically delete older Audit Records to prevent unlimited growth of disk usage on the [Audit Tasks](#) dialog.

Synchronization (Drift) Log Estimation

When Audit Server is configured to collect Synchronization Logs, synchronization records will be collected from each audited machine into log files on the Audit Server. Each time an audit is performed, new synchronization records are appended to the stored log files.

Each individual sync event in the log takes ~20 bytes of log space.

Therefore, the size of each machine's sync log depends on how often that particular machine is synchronizing its time. The total amount of disk space required can be calculated by determining how often a machine is synchronizing per day and multiplying that number times the size required to store each sync event.

For example, collecting logs from 10 machines that synchronize themselves an average of once an hour would use ~ 4k per day.

24 sync events/day x 20 bytes log space = 480 bytes per machine.
480 bytes/machine x 10 machines = 4800 bytes (4 kb)

However, logs can grow much larger if you are auditing more machines and these machines are synchronizing more often. If you audit 100 machines that are synchronizing themselves every minute, you would generate almost 3 meg of logs per day.

$1440 \text{ sync events/day} \times 20 \text{ bytes log space} = 28,800 \text{ bytes (28k) per machine.}$

$28 \text{ kbytes/machine} \times 100 \text{ machines} = 2.8 \text{ Megabytes}$

You can set a limit on the overall size of your synchronization logs in the Audit Server Synchronization Log [configuration settings](#).

PTP Monitor

As of version 5.2.b.20161230, Audit Server includes the ability to directly monitor and audit PTP nodes, including Domain Time (of course), as well as, appliance-type grandmasters, network switches, ptpd daemons on Linux and other platforms, and most other PTP implementations.

- ▶ [Usage](#)
- ▶ [Considerations](#)
- ▶ [Requirements](#)
- ▶ [Basic Operation](#)
- ▶ [Configuration](#)
- ▶ [Discovery Sweeps](#)
- ▶ [Limitations](#)
- ▶ [Other Information](#)

IMPORTANT: This is an Experimental Feature.

It may or may not work as expected on your network. See the [Requirements](#) and [Limitations](#) sections below.

Usage

- Use this list primarily for monitoring non-Windows PTP Masters and other PTP devices that cannot be audited directly via the [Domain Time II Nodes](#) or [Domains and Workgroups](#) lists in Manager.
- Although Domain Time machines running PTP will also show up in the [PTP Nodes](#) list, you should not audit them from here. You will get more accurate monitoring by auditing Domain Time machines from the [Domain Time II Nodes](#) or [Domains and Workgroups](#) lists and setting Audit Server to collect synchronization logs from them. See [Synchronization Logs](#) for details.
- Linux machines running PTP can either be monitored and audited using PTP Monitor (from the [PTP Nodes](#) list) or by using ntpd or chronyd as a reporting agent for time, in which case you can monitor and audit from the [NTP Nodes](#) list. As described below, PTP monitoring is more complex and subject to a number of limitations. If you are unable to use PTP Monitor against a Linux system for any reason, you should use the ntpd reporting method instead. To use ntpd to monitor machines synchronizing by PTP, check your ntpd man pages on how to set ntpd to run but not synchronize the clock, i.e. by adding lines like these to your ntp.conf file:

```
server 127.0.1.0
fudge 127.0.1.0 stratum 2
```

In this configuration, PTP will set the Linux system clock and ntpd will merely report that time. You can then add the Linux machines to the [NTP Nodes](#) list for auditing.

Considerations

- PTP is primarily a time-distribution methodology, and monitoring of non-master nodes is not always possible. Support for management messages is optional in the specifications, and some manufacturers may leave it out entirely, or implement only portions of the specification. You may use the [PTPCheck](#) utility to test the capability of any PTP node to respond to management messages.
- PTP monitoring uses a mixture of multicast and unicast to gather data. Contact with each node is inherently unreliable, and success will depend on both your network configuration and the capabilities of the node being monitored.
- PTP Monitoring is quite network-intensive, with large amounts of multicast traffic. Make every effort to limit the scope of your scans using the Hops limits and Domains to Monitor section of the configuration. The number of machines that can be effectively monitored depends largely on your network capabilities to handle this extra traffic. You can monitor many more

systems using the direct monitoring methods for Windows (Domain Time Nodes) and Linux (NTP Nodes) methods mentioned above.

- PTP monitoring uses many more message types and functions than normal PTP synchronization traffic. Many PTP-capable switches and routers have had bugs that have prevented these messages from being propagated correctly to all machines. Be sure to update your switches and routers to the latest firmware. If you are having problems seeing machines using PTP Monitor, try swapping out the PTP-aware device with a standard network switch and see if that resolves the issue. If so, you'll need to obtain a fix from the hardware vendor.

Requirements

You must install or upgrade to Audit Server version 5.2.b.20161215 or later (Server, Manager, and Audit Server must be installed and running the same version). You must then enable PTP Monitor from Manager's menu (*Audit Server -> PTP Monitor -> Enable*) or by right-clicking the **PTP Nodes** label in Manager tree and choosing *Configure...* from the context menu.

- Only IEEE1588-2008/2019 (PTPv2.0 & v2.1) nodes may be monitored. PTPv1 messages are not supported, and PTPv3 has yet to be released. For the remainder of this discussion, "the standard" refers to IEEE1588-2008.
- PTP Monitor uses Layer 3 (UDP) only. It cannot detect or interact with nodes that only use lower-level transports.
- Nodes to be monitored must be visible on the wire from the machine running PTP Monitor. If your network is segregated to prevent multicast traffic from crossing boundaries, then PTP Monitor will be limited to those nodes visible from its own location. You may run multiple instances of PTP Monitor on different networks to overcome this limitation.
- Generally, nodes being monitored must respond to both multicast and unicast PTP management messages, as defined in the standard. The exception is Masters may be partially identified even without management messages, via their Announce and Sync messages. As of version 5.2.b.20170922, PTP Monitor may be configured to use multicast management messages for node discovery and also for delay measurement against discovered Masters. Note, this improves compatibility at a cost of more network traffic. You may use the [PTPCheck](#) utility to test the capability of any PTP node to respond to management messages.

PTP Monitor uses multicast to discover nodes and to collect information that all nodes have in common. By default, PTP Monitor then uses unicast directly to each node to collect further information and monitor changes. As of v5.2.b.20170922, all discovery may be configured to be via multicast instead of unicast.

PTP Monitor sends its multicast discovery packets to each multicast-capable interface on the machine hosting PTP Monitor. If the interfaces change dynamically (for example, connecting a VPN, changing an IP address, enabling/disabling an adapter, or plugging/unplugging a cable), PTP Monitor will dynamically reconfigure itself to use only the "up" interfaces. It will also attempt to rejoin the IEEE 1588-specified multicast groups as interfaces come and go. This behavior means that PTP Monitor will likely be able to "see" all of the nodes reachable from the host machine; however, you must still configure routers, switches, or firewalls to allow the traffic as needed.

Domain Time includes a very useful tool for testing to be sure you are able to receive management messages. Use the [PTPCheck](#) utility to verify that management messages are passing across your subnet boundaries and through your switches, routers, and boundary clocks correctly.

PTP Monitor sends its unicast followups according to its host's routing table. Unlike multicasts, unicast routing is managed by the operating system. If your network consists of VPN-linked subnets, you may need to adjust the routing table for each subnet to ensure the proper gateway is used to reach each node. (If you can "ping" a node from the command line, the routing table is correct.) It does not make sense to send duplicate unicasts over each interface, since in most cases, the operating system will ensure the correct gateway is selected, and the redundant packets would either be dropped or report network-unreachable errors.

- PTP Monitor uses a default IPv4 TTL (and IPv6 Hopcount) of 1 for multicasts. You may adjust this value in order to monitor nodes outside of your local LAN. However, you may also need to adjust the TTL/Hopcount on the monitored nodes in order for the replies to reach PTP Monitor. PTPd uses a default TTL of 64 for everything except Peer-to-Peer messages. You may adjust this by editing `/etc/ptpd2.conf` and adding (or editing) the line `ptpengine: multicast_ttl=n` where *n* is the TTL you want, then restarting the PTPd daemon. (This information is true for the official PTPd version 2.3.1; ports, customized

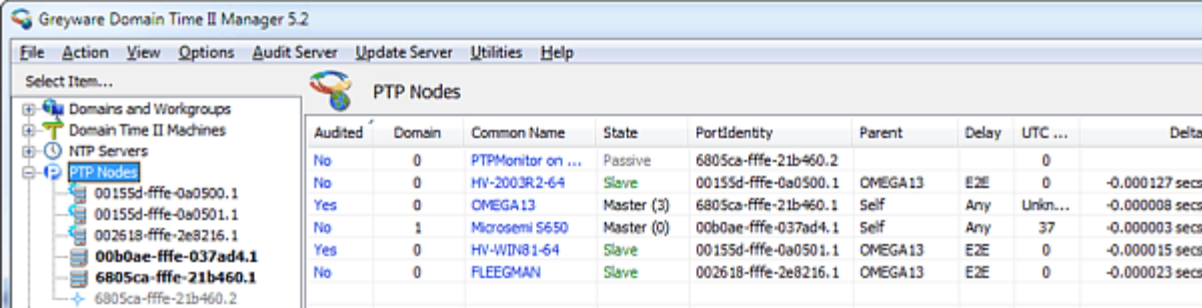
versions, or newer versions may work differently.) Domain Time allows you to adjust the IPv4 TTL and IPv6 Hopcount from the Control Panel applet's Network/Broadcasts and Multicasts page. Changes take effect immediately. We recommend you set the TTL/Hopcount to the lowest number possible for your network requirements.

Note: PTP Monitor cannot detect or interact with nodes using the Telecom Profile (negotiated unicast). Such nodes are forbidden by the standard from joining multicast groups. Even if they could respond to unicast queries, there is no way to discover a node's PortIdentity when it uses the Telecom Profile.

- Due to the amount of network traffic generated by PTP, the practical number of PTP Nodes that may be monitored from a single Audit Server is lower than when monitoring using other protocols (such as DT2 or NTP). The exact number is difficult to quantify, since it depends on your particular network design and the capabilities of the network devices and host machine to handle large numbers of UDP packets. PTP time accuracy and monitoring will suffer dramatically from delayed or dropped UDP packets, so keep a close eye on your network metrics to watch for high/spiky latencies or packet queueing. For this reason, Audit Server has a built-in limitation on the number of PTP nodes it will track. By default, the limit is 2000 nodes. Please contact support if your situation warrants a higher limit.
- Although you may operate PTP Monitor with Manager set to use DT2 or NTP sources for its [reference clock](#), it is highly recommended to have the Domain Time Server on your PTP Monitor machine set to be a PTP slave of your best local hardware PTP grandmaster. Domain Time Server would then be steering the local clock to precisely match the master. The correct reference clock setting for Manager would then be to use the local machine's clock.

Basic Operation

You may view the active status of your PTP network using the **PTP Nodes** section of [Domain Time II Manager](#).



Audited	Domain	Common Name	State	PortIdentity	Parent	Delay	UTC ...	Delta
No	0	PTPMonitor on ...	Passive	6805ca-fffe-21b460.2			0	
No	0	HV-2003R2-64	Slave	00155d-fffe-0a0500.1	OMEGA13	E2E	0	-0.000127 secs
Yes	0	OMEGA13	Master (3)	6805ca-fffe-21b460.1	Self	Any	Unkn...	-0.000008 secs
No	1	Microsemi 5650	Master (0)	00b0ae-fffe-037ad4.1	Self	Any	37	-0.000003 secs
Yes	0	HV-WIN81-64	Slave	00155d-fffe-0a0501.1	OMEGA13	E2E	0	-0.000015 secs
No	0	FLEEGMAN	Slave	002618-fffe-2e8216.1	OMEGA13	E2E	0	-0.000023 secs

Manager's PTP Monitor Display [\[Click for larger size\]](#)

Displayed nodes may be selected to be included in scheduled Domain Time II Audit Server [audit list](#), and will therefore raise the same alerts and be included in the same reports during audit scans as other audited protocols such as DT2 or NTP. PTP nodes may be auto-added to the Audit List using [Audit List Management](#).

Tracking via PTP is provided in addition to monitoring by other protocols. For example, a PTP appliance that also serves NTP may appear in both the **NTP Nodes** list and in the **PTP Nodes** list. A Domain Time Server acting as a PTP Master may appear in those two lists plus the **Domain Time II Nodes** list. Each protocol has its own advantages and disadvantages, and it may occasionally be useful to monitor multi-protocol nodes by more than one protocol.

For example, if you have a Domain Time Server acting as a PTP master, auditing it as a PTP Node will show you the quality of time being served by the master, whereas the regular Domain Time II auditing will show you how well that machine is tracking its own sources. However, in general you should not audit machines on both the PTP Nodes list and from the Domain Time II Nodes list. Use the Domain Time II method if available.

PTP Monitor can monitor multiple PTP domains. If you have more than one logical PTP network (a "domain") sharing the same wire, nodes will discard messages from any domain except their own. However, PTP Monitor can see them all, and track both Masters and Slaves in multiple domains. You may limit the domains being monitored from Manager's configuration pages for PTP Monitor. The [configuration](#) dialog page lets you specify individual domain numbers, ranges, or both.

PTP Monitor can collect synchronization logs (drift files) from audited PTP nodes (masters and slaves). These two new drift

files types are separate from the drift files collected by other protocols and are described in detail on the [Synchronization Logs](#) section of the Audit Data documentation.

Notes:

If a node in the PTP Nodes list is a Domain Time machine, you may remote-control it by double-clicking either the portIdentity or the IP address, or by right-clicking anywhere on the line and choosing Control Panel from the context menu.

If a node is running the Linux domtimed daemon, double-clicking will show the statistics. Otherwise, remote control is not available from the PTP Nodes list.

The Audit Server [Real-Time Alert](#) feature is not provided for PTP Nodes. However, Domain Time II Servers or Clients can provide that functionality independently of PTP Monitor.

The auditing of PTP Nodes is a separate function from other types Audit Server auditing. The "Audited" setting's column for PTP Monitor is independent of the "Audited" settings on the *Domains & Workgroups*, *NTP Nodes*, *Domain Time Nodes*, or *Real-Time Alerts* displays. Enabling/Disabling auditing on the PTP Monitor display will not change the audit settings on the other pages, and vice versa.

Configuration

PTP Monitor is configured using the PTP Monitor Configuration dialog. You launch the dialog from either Manager's menu (*Audit Server -> PTP Monitor -> Configure*) or by right-clicking the **PTP Nodes** label in Manager tree and choosing *Configure...* from the context menu.

PTP Monitor Configuration

Monitor Enabled

IPv4 Enabled (required)

IPv6 Enabled (optional)

Keep a node online if it responds to some, but not all, management messages

Operating Domain:

Range 0-127, default 0

Boundary Hops:

Range 1-64, default 1

Multicast Hops/TTL:

Range 1-64, default 1

Domains to Monitor:

List of domains, or ranges

Examples: 0,1,4 monitors only domains 0,1, and 4

0,1-5,127 monitors domains 0,1,2,3,4,5, and 127

0 monitors only domain 0

Sweep at audit time, or when refreshed from Manager

Sweep at a regular interval in the background

Interval:

seconds; range 5-86400, default 30

Sweep follow-ups:

Unicast

Multicast

Master Check

seconds; range 0-65535, default 2 (zero = disabled)

Interval:

Seconds between path delay measurements

Unicast

Multicast

Auto Drop Period:

days; range 1-3650

before removing unresponsive nodes from the list

Log activity messages if Audit Server's log is in trace or debug

mode

Monitor Enabled

IPv4 Enabled (required)

IPv6 Enabled (optional)

PTP Monitor is off by default. When PTP Monitor is enabled, IPv4 operation is always active. Do not enable IPv6 unless it is required. Otherwise you will be duplicating multicast traffic needlessly.

Keep a node online if it responds to some, but not all, management messages

By default, PTP Monitor will show nodes as online if they have responded to all management messages during the previous discovery sweep. However, some PTP nodes do not reply to all management messages (see [Limitations](#)) and would therefore show as offline. Checking this box ensures that nodes responding with partial information will be marked as online.

Operating Domain: Range 0-127, default 0

Allows you to specify the base PTP Domain that PTP Monitor operates on by default.

Boundary Hops: Range 1-64, default 1

Indicates the number of PTP boundaries to cross when discovering machines. Increase this value only if you are on a network segment connected to a PTP Boundary Clock and you wish to discover PTP devices on the other side of the Boundary Clock (such as the actual top-level Grandmaster). Otherwise PTP Monitor will only discover the Boundary Clock on the local segment. Note you may also need to increase the *Multicast Hops/TTL* value in order to hear multicasts from the remote subnet.

Multicast Hops/TTL: Range 1-64, default 1

Indicates the number of router hops a PTP Monitor multicast packet will traverse. This value is also known as multicast TTL (Time To Live). You must set this number large enough to account for all of the router/switch transitions a packet must cross in order to reach your entire network. Note that this value is independent of the [Network Discovery](#) multicast TTL value used for machine discovery by Domain Time Manager itself.

Domains to Monitor: Comma-delimited list of domains, or ranges

This value specifies on which PTP Domains you want PTP Monitor to attempt to discover nodes. PTP Monitor can monitor all possible PTP domains (0-127) simultaneously. You may limit the scope in order to reduce network traffic, or in order to segregate monitoring functions among multiple PTP Monitors.

Set this value to include only the actual PTP domain numbers in use on your network. Discovery sweeps are sent to each specified domain, and can generate a significant amount of unneeded traffic and take a significant amount of time if unused PTP domains are included. You may specify individual domains and/or a range of domains.

Sweep at audit time, or when refreshed from Manager

Sweep at a regular interval in the background

Interval: seconds; range 5-86400, default 30

Use these radio buttons to select the type and rate of [discovery sweeps](#). As mentioned above, discovery sweeps can generate a significant amount of multicast traffic. If you are monitoring a large number of PTP Nodes, you may want to only discover nodes on-demand (or when running an Audit), or use a relatively infrequent automatic scan interval.

As of version 5.2.b.20170922 you may select whether follow-ups are sent via unicast or multicast using the **Sweep follow-ups** radio buttons. Using multicasts creates a significant amount of extra traffic to remote subnets and should only be used if you have devices that do not support unicast management messages.

Master Check Interval: seconds; range 0-65535, default 2 (zero = disabled)

Master clocks are discovered by their announcements independently of discovery scans. PTP Monitor regularly measures the mean path delay between the Master and the PTP Monitor machine to assure accurate time delta information. This selection allows you to specify the number of seconds between delay measurements.

As of version 5.2.b.20170922 you may select whether delay measurement is done using unicast (hybrid mode) or multicast-only using the radio buttons. Using multicasts creates a significant amount of extra traffic to remote subnets and should only be used if you have devices that do not support hybrid mode.

Auto Drop Period: **days**; range 1-3650

Specifies how long unresponsive PTP nodes remain on the PTP Nodes list. Stale PTP nodes in the list result in extra discovery sweep traffic/timeout delay. Use this value to help keep your nodes list current.

Discovery Sweeps

Discovering PTP Nodes on a network is a complex process. Although some information can be gathered passively by listening to PTP traffic, it's necessary to periodically send queries of various sorts to acquire all available data. These exploratory probes are known as Discovery Sweeps.

Sweeps are run based on the setting on the [PTP Monitor Configuration](#) dialog (described above).

PTP Monitor uses both multicast and unicast to obtain state information about nodes. General discovery is done by periodic multicasts; follow-up queries are sent directly to each node using unicast. (As of v5.2.b.20170922, you may send follow-ups using multicast only. See the [Configuration](#) section above.) Most hardware grandmasters, PTPd as of version 2.3, and all Domain Time nodes support mixed message types. This is very similar to the "hybrid" mode used by slaves; see the [Enterprise Profile](#) section of the PTP Profiles documentation for more information.

PTP Monitor does not need to sweep the network in order to discover master nodes; an overheard Announce and subsequent Syncs/Sync Follow-ups are sufficient.

PTPd and other software slave nodes can only be monitored effectively by sweeping the network. If you only require the information as part of an audit, you may let the commanded sweep from Audit Server collect the information. If you require an up-to-date display on Manager, you will need to use F5 Refresh or enable periodic background sweeps.

Domain Time nodes announce state changes and significant events, so sweep is not required for the current status of Domain Time slaves or masters. If all of your PTP client software is running Domain Time, you may disable periodic sweeps altogether. Note that Audit Server will perform a sweep at the beginning of an audit, or when you first open Manager (or hit F5 on Manager's PTP Nodes display). This behavior helps ensure the information from each node is as recent as possible.

Limitations

Although PTP Monitor is able to auto-discover most PTP nodes on a network, there are circumstances that may prevent machines from being detected and/or fully identified.

Management Messages

As noted in the [Requirements](#) section, only nodes that support both multicast and unicast management messages may be monitored (with the exception of Master nodes using multicast Announces and Syncs).

PTP Monitor uses the information from the Clock Description management query to fill in the fields for device name, hard/firm/software versions, and other identifying information. You may edit the "Common Name" and "DNS Name" fields. Please note that all but a very few fields of the Clock Description response are optional, and that many implementations either do not support the message at all, or support only a subset of the information. For example, all PTPd nodes report a device name of "PTPDv2" and a software version number, but no other identifying information.

Management message handling is optional per the standard. Most appliance-type grandmasters, ptpd as of version 2.3, and Domain Time, accept and reply to management queries. However, nodes are not required to handle management messages at all; even those that support management messages may only support a limited subset of message types.

PTP Monitor can detect and monitor nodes acting in the Master state, whether or not they reply to management messages. Note: if the Master being monitored does not support management messages, some of PTP Monitor's

information will be incomplete.

PTP Monitor tracks masters by listening for multicast Announce and Sync messages, and calculating the difference between its own clock and the advertised time. In this sense, PTP Monitor acts like another Slave; that is, delay information is collected periodically, and the timestamps in the Sync messages are used to determine the offset of the master from the PTP Monitor machine. on the network.

Note that this is not as precise of a measurement as a true slave could derive, because the machine is not performing corrections based on the incoming timestamps. Master node offsets are recalculated upon the receipt of each Sync/Sync Follow-up message.

If the Master does not support E2E or P2P unicast delay requests, then you should change to using multicast for delay measurement. Since PTP Monitor is not actually a Slave (does not try to match frequencies), the deltas reported by PTP Monitor will be somewhat higher than those reported by Domain Time Server running on the same machine. This is expected behavior, and is essential to PTP Monitor's functioning because PTP Monitor can track multiple masters at the same time.

Management messages, whether unicast or multicast, are always directed to port 320. Unicast replies to management messages may be sent either to the source port of the request, or to port 320 (this is implementation-dependent). For this reason, PTP Monitor always sends its requests from port 320, to ensure that nodes from different manufacturers will be able to reply no matter how they interpret the standard.

Offset measurement

The PTP standard is primarily a specification for time distribution; in particular, the standard specifies how Slaves should determine an acceptable Master, and how they should interact with that Master, and forbids Slaves from responding to Master-only messages.

For example, Slaves are forbidden from responding to delay requests, so the network distance between PTP Monitor and a Slave cannot be measured directly. Slaves are also forbidden from placing Announces or Syncs on the wire, so the offset between PTP Monitor and a Slave cannot be measured directly, either. Through management messages, Slaves can report their own measurements of their offset and distance from the Master, and which Master they are following.

PTP synchronization data for slave nodes is collected from the slave's own measurement of its offset from its master. The slave's information is collected only when a discovery sweep occurs. IEEE 1588-2008 specifies only the type of data to be returned, not its source; this means that the values provided are implementation-dependent. Some nodes may return the most-recently calculated offset prior to correcting for it, others may return the offset after the most recent correction, while still others may supply filtered data. Further, IEEE 1588 does not specify what values should be returned when a node was previously a slave but has lost its master, nor what masters should declare if they are not using a direct connection to a GPS timesource. Therefore, although the information may be present and valid, PTP Monitor will only report on it if the node is currently a slave.

Masters are required by the standard to set their own offset and delay values to zero. Slaves are required to report their Master's time source as their own. If a Master happens to use NTP to obtain its own time, it should report its time source as NTP, but will still report its offset and delay as zero. If a Slave uses PTP to obtain time from that master, it will also show its time source as NTP. The Slave will report its offset and delay from the Master, but not from the Master's source, because the standard has no way to represent a master's offset or delay. Likewise, if the Master is an appliance using GPS, it and all of its Slaves will report a time source of GPS, even though the slaves are using PTP.

PTP Monitor directly measures Masters by processing its Announces and Syncs. It does not attempt to measure Slaves, because the standard does not allow it. For Slaves, the reported offset and delay are accepted at face value.

The standard leaves clock steering as an implementation detail. Thus, there is no way to distinguish between a Slave that has measured its offset and delay but not yet compensated for it, from a Slave that reports its most recent actual synchronization. Each manufacturer is free to report whatever set of values it feels appropriate (including none at all if it fails to respond to management messages).

Time Since Last Set

NTP and DT2 report their time sources, and the time elapsed since they checked. This information is available without examining logs, by virtue of the data contained in the time packets exchanged.

The PTP standard does not define this property, and there is no way of measuring it. A Master is presumed to be always correct at all times, and the clock steering mechanism used by Slaves is implementation-dependent. Therefore, any node that claims to be either Master or Slave is presumed to have set its clock if it responds at all. The Slave may, in fact, be in the middle of a two-hour slew to correct the offset, and it may report either the offset from when it discovered the need to slew, the current offset as the slew occurs, or the expected offset after the slew has completed. It will not report when it discovered the offset or whether a correction has been made.

A node that remains online (responding to management messages), but that reports its state as something other than Master or Slave, allows PTP Monitor to calculate the time since last set as long as the node has been a Master or Slave in the past. For example, if a Slave loses its Master while being monitored and changes from Slave to Listening, PTP Monitor will know how much time has passed since the clock was last a Slave. But if a node remains in the Listening or Passive state, or fails to respond to management messages (perhaps because it's now offline), PTP Monitor has no data, and reports will show "Unknown" as the time since last set.

Other Information

portIdentity

All PTP nodes are required to have a unique portIdentity. The portIdentity is formed from a 64-bit value called the clockIdentity, plus a port number. The clockIdentity is normally formed from the MAC address of the primary adapter, with 0xFFFE inserted in the middle to make a 6-byte field into an 8-byte field. clockIdentity may also use the manufacturer's 3-byte OUI plus a guaranteed-unique remaining 5-byte value.

PTP Nodes use the portIdentity to distinguish among Masters. When a Slave sends a delay request to its Master, it includes its own portIdentity in the request, and the Master sets this as the target for the reply. This scheme allows either multicast or unicast for delay measurement between Slaves and their Masters, since a Slave can pick out its own reply.

Management messages, whether multicast or unicast, must include the target portIdentity. PTP nodes cannot be addressed by their IP address or DNS name alone. PTP Monitor uses a special kind of multicast management message addressed to "all clock, all ports" to gather the portIdentities of the various nodes. (There is no corresponding "all domains" message type, although v3, when released, will likely support it. Domain Time already does. At the moment, the "all clocks, all ports" messages must be duplicated for each monitored domain.) After learning a node's portIdentity, PTP Monitor can then direct unicast requests for further information directly to the node, using the IP address from which it responded. This combination of multicast and unicast allows PTP Monitor to track nodes even if they use DHCP and change IP addresses.

Hardware appliances typically have the portIdentity "baked in" at the moment of manufacture, usually using the manufacturer's OUI. A hardware appliance should never change its portIdentity under normal circumstances, although each Ethernet port on the appliance may have its own portIdentity (usually the same clockIdentity with a different port number).

PTP-aware routers and switches, acting as boundary clocks, typically have a baked-in clockIdentity shared by all ports, and a port number corresponding to each Ethernet port. These values normally never change.

Software nodes, such as ptpd, follow the standard exactly, using the MAC address of the primary adapter. This is usually sufficient to ensure that, from boot to boot, the portIdentity will be unique and constant. However, the concept of "primary adapter" is OS-dependent, and may change either after a reboot or even while running. Most implementations will not change portIdentities while running, even if moved to a different adapter (say, for instance, a live migration from one virtual host to another, or a live change of active adapters). At reboot, however, the node will discover a different MAC address, and begin participating with a new portIdentity.

Domain Time nodes maintain a persistent portIdentity from the moment of first mobilization. If a duplicate node is discovered on the network, Domain Time will switch to the Faulty state and report the duplicate's IP address in the log file. Otherwise, Domain Time will continue using the same portIdentity, even if NICs are swapped or the machine is live-migrated to a new host. Domain Time allows an admin to change the clockIdentity if a conflict is discovered.

Since the only uniquifier for PTP is the portIdentity, historical records are subject to contamination if nodes end up swapping portIdentities. This could happen if you swap NICs between two ptpd nodes, or if you migrate ptpd nodes to different hosts and then reboot them. The standard's only requirement is that all portIdentities must be unique during the operation of the time distribution, as seen from the Master's perspective.

If a node merely changes its portIdentity to a different (but unique) value, PTP Monitor will show the old node offline, and begin tracking the new one. Since synchronization logs are kept by the portIdentity, you may end up with a new log starting for the same node, or (worst case), data from one node being appended to data from another node that formerly used that portIdentity.

If you clone installations, you may end up with multiple nodes having the same portIdentity. The behavior of PTP under these circumstances is undefined. The nodes will probably not be able to synchronize with the Master, and PTP Monitor will not be able to tell them apart. Domain Time checks for this condition at each startup, but detecting duplicates depends on the proper functioning of multicast management messages among non-Master nodes.

Advanced Options

This page describes Audit Server's Advanced Options.

Audit List Management...

Audit Server can add discovered machines to the Audit list and also remove non-responding systems from the list automatically.

Automatic Addition to the Audit List

Check the **Add Domain Time Nodes discovered during audit** checkbox to add any new machines running Domain Time Server, Client, Windows Time Agent, or the domtimed daemon found on the network to the list of audited machines.

Check the **Add NTP Nodes discovered during audit** checkbox to add any newly-discovered NTP daemons to the list of audited machines.

Check the **Add machines discovered by receipt of startup Real-Time Alerts** checkbox (version 5.2.b.20150307 or later) to allow new machines not already in the Manager database to be added upon receipt of a Real-Time Alert upon service start.

For security purposes, Audit Server will not accept Real-Time Alerts from machines that are not already present in the Manager database (appearing in the *Domains and Workgroups* list) by default. However, in some circumstances, such as adding new machines to the network that don't exist in Active Directory, this will prevent a machine set to "Always audit this machine" on its *Status Reports* configuration page from being auto-added (since its Real-Time alerts are being rejected). Enabling this checkbox can bypass this restriction. Use only if required.

If checked, and a previously-unknown machine sends a Real-Time Alert shortly after boot or service restart, Audit Server will attempt to add the machine to the audit list. The sending machine must respond to Audit Server's query before it can be added. Audit Server will only try unknown machines a few times before giving up.

Check the **Add PTP masters discovered by PTP Monitor** checkbox (version 5.2.b.20170101 or later) to add any PTP master servers discovered by PTP Monitor to the audit list.

Check the **Add PTP masters discovered by PTP Monitor** checkbox (version 5.2.b.20170101 or later) (version 5.2.b.20170101 or later) to add any PTP slaves discovered by PTP Monitor to the audit list.

Check the **Add machines that have synchronized with Domain Time II Server** checkbox to add those systems to the list of audited machines.

When checked, Audit Server will automatically add systems to the Audit List by contacting Server(s) and retrieving a list of all machines (ephemera) that have synchronized their time with that server using Domain Time II protocols. Multiple servers may be contacted to obtain their machine lists, if desired.

[Audit List Management](#) [\[Click for larger size\]](#)

This method is a reliable method for populating the Audit List, and it has the added advantage of adding machines that are not currently online. However, it cannot discover any Domain Time II components that are not synchronizing with a Domain Time II Server. Those machines must be discovered using Domain Time Manager list discovery and/or entered manually and added to the list.

Notes:

- The "Adding machines that have synchronized with Server" function requires Domain Time II Server version 3.1 and later.
- Only systems that synchronize with Domain Time Server(s) using the DT2 protocol can be auto-discovered.
- The Audit Server must use credentials with sufficient rights to connect to the administrative share on the remote Server(s). See the [Service Credentials...](#) and [IP Restrictions](#) sections below for details on those settings.
- Machines may also be manually added to the audit list using Domain Time II Manager, either one-at-a-time or in a batch. See the [Select machines to audit with Audit Server](#) section of the "How to Manage Domain Time Remotely" page of the Manager documentation.

Foreground - collection must finish before audit completes

Background - collection finishes independent of scheduled audits

Run background collection periodically, not just at audit time

These choices determine whether Audit Server will collect the server ephemera data in a separate thread from the audit run itself. Collecting ephemera data records from each Server can take an extended amount of time, particularly if you have a large number of synchronization events, since Audit Server must parse each event to determine whether or not it represents a new machine to be added.

Choosing **Background** allows collection of the basic audit data very quickly, and then the collection of the ephemera logs can complete in the background. Running the collection in the background periodically can make collection even more efficient.

Obtain records from this machine only

Specify a list of servers

Collection of the list of machines that synchronize with Domain Time II Server is enabled by default only on the Domain Time II Server on which Audit Server itself is installed. Other Domain Time II Servers will not keep a record of synchronizing machines until you enable data collection on them by entering them in the Server List. You will see a confirmation dialog when the server is successfully added to the list.

Automatic Removal from the Audit List

Stop auditing machines that haven't responded in over **days** will trim the audit list of any machines that have not been contacted in the specified period. Uncheck the box if you do not want to trim the list.

Reset last contact date and failure count when a machine is added manually sets the failure counters to defaults when manually adding machines.

Data Folders...

Choose where Audit Server stores records, reports, and logs.

A screenshot of a settings window with a light beige background. It contains three text input fields stacked vertically. The first field is labeled 'Audit Results', the second 'Daily Reports', and the third 'Synchronization Logs'. Each field has a small icon to its left and a small 'x' icon to its right for clearing the text.

The file locations can be any valid file folder to which the Audit Server service account has sufficient rights to read and write files.

You should specify locations on physically-attached storage so that the background service may access them without interruption. If you change a location, Audit Server will automatically move existing files to the new location for you.

If you must, you may indicate any valid UNC path to store the files on a remote machine, however, be aware that should the remote machine become unavailable for any reason, audit data collected during that period will be irretrievably lost.

IMPORTANT:

Since files in these folders are used to create an audit trail, best practice requires that they must be as secure as possible, and we strongly recommended that the folder be located on a local drive using the NTFS filesystem to accomplish this.

Folder permissions should be set so that only the Audit Server service account (usually System) has Full Control. By default, everyone else should be denied access entirely. If you choose to grant exceptions (such as to export Daily Report files), you should take care to only grant Read-Only rights to the required user/group. You may also wish use operating system auditing to monitor the folders for unauthorized changes.

Service Credentials...

Audit Server needs administrative rights to be able to collect synchronization logs and ephemera discovery records from remote systems. The settings on the *Audit Server -> Advanced -> Credentials...* dialog allow you to specify the account used by Audit Server for this purpose.

A screenshot of the 'Audit Server Credentials' dialog box. The title bar is blue with the text 'Audit Server Credentials'. The dialog has a light beige background and contains two identical sections. Each section starts with the text 'Run the Audit Server service using a Domain Admin account'. Below this text are three labels: 'Domain:', 'Username:', and 'Password:', each followed by a text input field. The first section's input fields are slightly offset to the right compared to the second section's.

You have the choice of having the Audit Server service itself run under the LocalSystem account and supply the administrative access credentials only when performing an audit, or having the service running with the administrative privileges at all times. In general, the first option is preferred. In either case, account details are encrypted in the registry.

Audit Server can access other domains and workgroup members as long as the credentials supplied match an administrative account on the domain (or local machines in the workgroup). If you select a workgroup or domain to which

Audit Server does not have administrative access, the collection will fail and will be noted in the logs.

IP Restrictions...

The *Audit Server -> Advanced -> IP Restrictions...* menu item allows you restrict which systems are allowed to contact Audit Server.

IP Restrictions

You may limit Audit Server's scope by specifying a list of IP address ranges. In most cases, you should limit the scope to your local subnets.

These restrictions apply only to auto-adding newly discovered machines and to recognizing the receipt of real-time alerts. A machine already on the audit list will be audited regardless of its IP address.

☐ No restrictions
☒ Permit only listed ranges
☐ Deny any in listed ranges

First IP in range	Last IP in range
192.168.10.0	192.168.10.255

First IP in range:

Last IP in range:

IP Restrictions dialog [Click for larger size]

The **IP Restrictions** dialog applies to machines sending [Real-Time Alerts](#), which machines are available to be [Auto-added to the Audit List](#), and which machines can provide [synchronization logs](#).

You can both permit and deny access from IP ranges. To restrict a single IP address, enter the same IP address for both the First and Last range items.

Standby Mode

As of version 5.2, Audit Server has the ability to be configured as a "hot-spare" backup to a functioning Audit Server. This function is ideal for use in Disaster Recovery sites or for providing redundancy in normal operation.

How it works

To configure Audit Server replication, a functioning Audit Server on the network is chosen as the Primary Audit Server. This server is considered to be running in Normal Mode. You may then designate one or more other Audit Servers to act as Secondary Server(s) to the Primary by configuring the Standby Mode options and enabling Standby Mode.

Standby Mode

You can configure this machine to be a backup by enabling Standby Mode. When operating in Standby Mode, Audit Server periodically synchronizes all of its settings with the primary server. Replication includes logs, reports, registry settings, and the audit list. If the primary server goes offline, you may disable Standby Mode and this machine will take over running audits, generating reports, and issuing alerts.

NOTE: When Standby Mode is enabled, this machine will not perform audits or collect data. It will only replicate data and settings from the primary. Any changes you make to settings on this machine will be lost at the next replication.

☒ Standby Mode Enabled

Primary Server: (DNS name, NetBIOS, or IP address)

Replicate every: minutes (range 1-99999)

☒ Keep old files and reports, even if they no longer exist on the primary
☐ Delete files and reports that have been expired or deleted on the primary

☐ Automatically resume Normal Mode after: sequential replication failures

Credentials

You must provide credentials for accessing the primary server's registry and file system. The account needs read access to the registry, and change access to the drive shares.

Domain:
 Username:
 Password:

[Standby Mode Options](#) [\[Click for larger size\]](#)

While in Standby Mode, Audit Server will continually replicate all of its settings and data from the Primary Audit Server on the schedule you specify. This results in the Secondary Audit Server being ready to take over audit duties should the Primary become unavailable.

If the Primary is offline, the Secondary Audit Server can be brought online by disabling Standby Mode, either manually or automatically (see below).

IMPORTANT:

When a machine is brought out of Standby Mode, it will immediately begin auditing using the Primary's list of audited machines on the Primary's schedule. If the Primary machine is online (or comes back online) while the Secondary is in Normal Mode, both the original Primary and the Secondary machines will be auditing the same set of machines simultaneously. This can result in collected audit data (such as drift logs) being split unpredictably between machines. To avoid this, take care to be running only one Audit Server in Normal mode at a time.

When you set a machine in Normal mode to Standby mode, replication of all data from the Primary begins immediately. This **will result in irretrievable data loss** of any previously collected data on the Audit Server entering Standby mode.

!!! All data and configuration changes on the Secondary will be overwritten !!!
by the Primary's data when the Secondary enters Standby Mode.

**BE SURE TO BACKUP ALL COLLECTED DATA ON THE SECONDARY MACHINE
TO ARCHIVAL STORAGE BEFORE (RE)ENABLING STANDBY MODE**

Replication Settings

Network port: For replication to occur, you will need to allow traffic over the DT Alert Sharing port (default 9910 TCP) to pass any intervening routers/switches/firewalls. You may enable and configure the port by changing the values on the [Advanced Real-Time Alert Configuration](#) dialog.

Primary Server:

Specify the name or address of the Audit Server that will be the Primary.

Replicate every minutes (range 1-99999)

Indicate how often you want the Secondary to replicate data from the Primary.

Keep old files and reports, even if they no longer exist on the primary

This setting keeps all files replicated from the Primary, even if they are later removed from the Primary. This option maintains the most data, but may result in disk space issues, since data will continue to accumulate.

Delete files and reports that have been expired or deleted on the primary

This setting seeks to make the Secondary match the Primary as closely as possible, including replicating file deletions.

Automatically resume Normal Mode after: sequential replication failures

This option allows the Secondary Audit Server to automatically switch to Normal Mode if the designated Primary Server is unreachable for a certain number of replication attempts, thereby providing auto-failover capability. If this option is unchecked, you must manually disable Standby Mode to change to Normal Mode.

Warning:

Any inability to replicate with the Primary will be considered a failed replication attempt regardless of the actual cause. Network problems, credentials issues, DNS lookup failures, etc. can all result in failed replications, even though the Primary machine may not actually be down.

If you enable this option, you should be prepared for the possibility that the Secondary Audit Server may come online and begin auditing under those circumstances.

There is no ability to automatically "fail-back" when the Primary comes back online. You should try to avoid running both the Primary and Secondary systems in Normal Mode simultaneously (see the **Important** note above).


Credentials

Credentials

Domain:

Username:

Password:



Audit Server in Standby Mode needs administrative rights to be able to replicate settings and data from the Primary Audit Server. Enter the necessary information in this section of the dialog. Typically, you will need to use an account with administrative access to the remote systems.

Click the **Test** button to verify that your credentials are correct and that replication can proceed.

Service Logs

Settings on this page control the Audit Server's [text](#) and [syslog](#) options.

Text log

The Audit Server service keeps a record in text format of its activity, which includes housekeeping and audit run information.

To configure the Audit Server logs, click *Audit Server -> Audit Log File Settings* from the Manager menu.

Text Log

Log Level: Information

Max Size: KB (use zero to mean unlimited size)

Log Roll: Daily at Midnight

Delete old logs

Keep up to old logs

Text Logs are kept in the `%SystemRoot%\System32\` folder. There are two main log files collected when the service is running:

- **dtaudit.log**

This is the currently active service text log file.

If log archiving is enabled (see below), additional archived log files will be created using a `dtaudit.YYYYMMDD.log` naming scheme (i.e. `dtaudit.20090928.log`).

- **dtaudit.startup.log**

A detailed text log of the service startup process. Only data from the latest startup is included.

To view these logs, click the button, which launches the Domain Time Log Viewer.

Log Level

This drop-down chooses what type of entries to include in the log. You can increase or decrease the amount of information logged as needed.

The available levels are (in increasing amount of detail):

- **Disabled**

This switch will only disable the **dtaudit.log** file. The other system log, **dtaudit.startup.log** cannot be disabled.

- **Errors**

Only messages marked as Errors will be logged

- **Warnings**

Logs will include Errors and Warnings

- **Information**

Includes Errors, Warnings, and Information messages.

- **Trace**

Includes all of the above, plus detailed Trace information.

- **Debug**

Includes all available information provided by the service.

CAUTION:

Debug logging will generate a great deal of data, so be sure to only enable it when you need the additional information, and don't forget to turn it off when finished troubleshooting.

Max size

This sets how large the log file is allowed to grow (in kilobytes).

Once the maximum size is reached, the oldest events will be scrolled off to make room for new events. Enter 0 (zero) if you don't want to limit the log size.

It's a good idea to set a log size that will allow you to keep enough history to help you determine the timeframe and scope of any issues you may encounter.

Log Roll

Domain Time can automatically archive the text log on a daily, weekly, or monthly schedule.

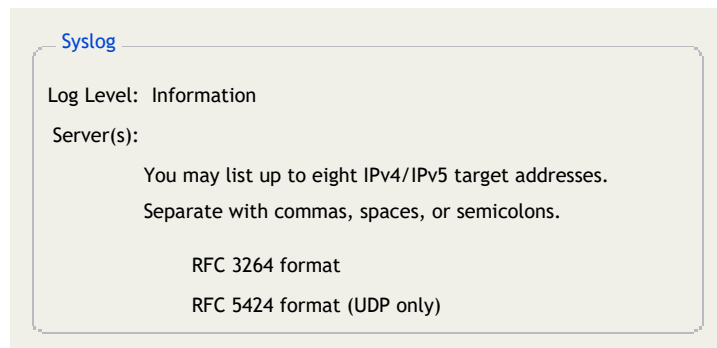
When the log is archived, all existing log events in the **dtaudit.log** file will be written to an archive file named **dtaudit.YYYYMMDD.log** (i.e. dtaudit.20090928.log) and the current log file will then be cleared to accept new data.

You can choose how many archived log files to keep on the machine. When the indicated limit is reached, the oldest log file will be deleted.

Syslog

As of version 5.2.b.20180101, Audit Server can also send a subset of the Audit Server service logs to syslog servers.

To configure the Audit Server Syslog output, click *Audit Server* -> *Audit Syslog Settings* from the Manager menu.



The screenshot shows a configuration window titled "Syslog" with a light beige background. Inside the window, the "Log Level" is set to "Information". Below this, the "Server(s):" field is empty, with a note stating "You may list up to eight IPv4/IPv5 target addresses. Separate with commas, spaces, or semicolons." At the bottom, there are two radio buttons for the log format: "RFC 3264 format" (which is selected) and "RFC 5424 format (UDP only)".

When enabled, Audit Server can send the service activity information to syslog server(s). You can set the level of information sent independently of the settings for the text service log.

You may list up to eight targets on the Syslog Server line. Only IP addresses are supported. Separate targets with commas, spaces, or semicolons.

Use the radio buttons to select the RFC format that conforms to your syslog server.

Most options for Domain Time II Audit Server are set using the Audit Server configuration items in Domain Time II Manager. A few advanced options can only be set by changing the registry. This section explains the registry entries used by Domain Time II Audit Server.

Caution:

Modifying Registry entries requires basic familiarity with the Windows Registry and its operations. Incorrect changes to the Registry can result in unpredictable, perhaps non-repairable, damage, up to and including a non-bootable system! Have a qualified technician make the changes for you if you are not comfortable with the process. We cannot be responsible for registry problems.

The Domain Time II Audit Server settings are located in these keys (click the names to jump to details):

[HKEY_LOCAL_MACHINE](#)
 [Software](#)
 [Greyware](#)
 [Domain Time II Audit Server](#)
 [Auth](#)
 [Logs and Alerts](#)
 [Parameters](#)
 [PTPMonitor](#)
 [Standby](#)

Auth

Authentication settings are located in this key:

[HKEY_LOCAL_MACHINE](#)
 [Software](#)
 [Greyware](#)
 [Domain Time II Audit Server](#)
 [Auth](#)

You should not make manual changes to this key or its subkeys unless instructed by Technical Support.

Logs and Alerts

Logs and Alert settings are located in this key:

[HKEY_LOCAL_MACHINE](#)
 [Software](#)
 [Greyware](#)
 [Domain Time II Audit Server](#)

Logs and Alerts
Audit Data Collection
Daily Reports
Drift Collection
Ephemera Collection
Event Viewer
RTAlert history
SMTP
SNMP
Text

Audit Data Collection Key

Value Name: Multicast by Serial Number to locate non-responding audited machines

Value Type: REG_SZ

Default Value: *True*

Options: *True or False*

Notes: Controls whether Audit Server tries to locate non-responding audited machines by sending a DT2 multicast using the audited machine's serial number. Change this value to "False" if your network does not allow multicast to eliminate unnecessary timeouts. Changes take effect at the next scan; you do not need to restart any services.

Value Name: Retries on Contact Failure (range 1-5, default 1)

Value Type: REG_DWORD

Default Value: 1

Range: 1-5

Notes: Controls how many times Manager and Audit Server will retry to contact a machine that fails to respond to unicast followups. Multiple retries are generally a bad idea, since if a machine fails to respond to a unicast once, it is unlikely to respond a second later. However, in some circumstances, such as auditing remote offices across a poor WAN or dealing with a congested switch or router that is discarding packets, you may want to increase the retries. Changes take effect at the next scan; you do not need to restart any services.

SMTP Key

Value Name: Email Subject Audit Alert

Value Type: REG_SZ

Default Value: Domain Time II Audit Alert

Notes: Sets the email subject line for Audit Alerts. You may customize this using your own text.

Value Name: Email Subject Audit Summary

Value Type: REG_SZ

Default Value: Domain Time II Audit Summary

Notes: Sets the email subject line for Audit Summaries. You may customize this using your own text.

Value Name: Email Subject Real-Time Alert

Value Type: REG_SZ

Default Value: Domain Time II Real-Time Alert

Notes: Sets the email subject line for Real-Time Alerts. You may customize this using your own text.

Parameters

General settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time II Audit Server
Parameters

Value Name: Service Log Filename

Value Type: REG_SZ

Default Value: %SystemRoot%\System32\dtaudit.log

Notes: Sets the location and name of the service log file. If this value is not present or is blank, the log file will be created with the default filename dtaudit.log in the %SystemRoot%\System32\ folder. The complete path and filename must be specified (i.e. C:\Windows\System32\dtaudit.log) and the drive specified must be a local drive.

Standby

PTP Monitor settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware

Domain Time II Audit Server
PTPMonitor

You should not make manual changes to this key or its subkeys unless instructed by Technical Support.

Standby

Standby Mode settings are located in this key:

HKEY_LOCAL_MACHINE
Software
Greyware
Domain Time II Audit Server
Standby

You should not make manual changes to this key or its subkeys unless instructed by Technical Support.

Domain Time II Agent

Version 4.2

Domain Time II Windows Time Agent is a special Control Panel applet that allows you to easily configure the Microsoft Windows Time Service instead of using the w32tm command-line utility or manually changing registry settings. Windows Time Agent solves several of the major drawbacks of using Windows Time, such as determining how the service is actually configured, and whether or not it is really working.

Note: This software is designed for Windows XP and later. Although it will run on Windows 2000, many of its functions and tab pages will be unavailable.

Windows Time Agent does not synchronize the time; it configures and monitors the Windows Time Service.

Windows Time Agent provides visual indicators of the clock status so you can tell at a glance whether the machine is is synchronized, what time sources are in use, and how accurate the clock is. The Windows Time Agent also shows clock drift data in scalable graphical displays so you can see how your clock is performing over time.

The Windows Time Agent works splendidly as a stand-alone utility. However, when combined with Domain Time II [Audit Server](#) it becomes even more powerful. Audit Server can collect the clock drift data from Windows Time Agents to add to its audit trail, plus it can raise alerts if any machine's clock is not synchronized. Now you can know exactly how Windows Time is performing across your entire network.

[Installation Instructions](#)

[System Requirements](#)

Installation

The utility is included in the Domain Time II Server, Client, and Manager distribution files.

Important Note: The Windows Time Agent is installed by default with version 4.1 Domain Time II Server and Clients. As of version 5.1, Agent is **not** installed when Domain Time II Server or Client is installed.

Upgrading an existing v4.1 machine to v5.1 or later will not remove an installed Windows Time Agent, but will upgrade it to the latest version*. A previously-installed Agent may be disabled on the **Advanced** property page of the Domain Time Server or Client applet (v5.1 or later).

* The version and build date of Windows Time Agent is independent of the main Domain Time II suite, so the displayed version may vary from other Domain Time II components.

Installation/Upgrade

- To install the Domain Time II Windows Time Agent from Domain Time II distribution files:
 - Copy the **w32tmdt.cpl** program directly into your %system%/System32 folder. The 32-bit version is located in the /i386 folder; the 64-bit version is in the /AMD64 folder of the distribution files.
 - Start the Domain Time II Windows Time Agent applet from the icon in the Windows Control Panel.

Note: On systems with User Account Control (UAC) enabled, you must *Shift+Right Click* and choose **Run As...** from the context menu to launch the Control Panel applet. On Windows Server

Core, type in **w32tmtdt. cpl** on the command line.

Removal

Do not just delete the w32tmtdt.cpl file from the %system%/System32 folder - the program should always be uninstalled.

- Use the **Programs and Features (Add/Remove Programs** on earlier versions of Windows) utility from the Control Panel to remove the program.

Note: If Windows Time Agent was installed before Server or Client was installed, removing Server or Client will not remove it. You must remove Windows Time Agent after removing Server or Client.

Client Tab

The **Client** tab page displays the settings you can make to the time client function of the Windows Time service. If Domain Time II Server or Client is installed, all options on this page will be disabled, since the Domain Time II service will be obtaining the time and managing the system clock.

The screenshot shows the 'Client' tab of the Windows Time Agent settings window. The window title is 'Windows Time Agent Version 4.2 Copyright © Greyware Automation Products, Inc.'. The 'Time Client' section explains that Windows Time can synchronize the clock with an external time source and asks the user to select how they want it to find time sources. There are five radio button options: 'Disabled', 'NoSync - NTP client enabled but not managing the clock', 'NT5DS - Use domain hierarchy to find a time server', 'AllSync - Use domain hierarchy, then fall back to specified servers', and 'NTP - Use specified servers and modes'. The 'NTP' option is selected, and a 'Special Interval' of 604800 seconds is entered. Below this is a table with four columns: 'Server IP or DNS name', 'Sync Mode', 'Stratum', and 'Current Variance'. The table lists two servers: 'time.windows.com' (Stratum -, Timed out) and 'time.nist.gov' (Stratum 1, +0.001 secs). There are five empty rows for additional servers. At the bottom, a green status icon indicates 'Clock synchronized to within 1 second', and a 'Sync Now' button is present. The 'Currently selected time source' is '192.43.244.18 (time.nist.gov)'. The window has 'OK', 'Cancel', and 'Apply' buttons at the bottom.

	Server IP or DNS name	Sync Mode	Stratum	Current Variance
1.	time.windows.com	Special Interval	-	Timed out
2.	time.nist.gov	Special Interval	1	+0.001 secs
3.		Special Interval		
4.		Special Interval		
5.		Special Interval		
6.		Special Interval		

[Windows Time Agent: Client tab](#) [\[Click for larger size\]](#)

The Windows Time service can be set to a variety of methods to obtain the time and set the local clock. Use the radio buttons to select the method you prefer. Consult the Microsoft Windows Time documentation for a description of each of these options.

Disabled

The client portion of Windows Time service will not be loaded.

NoSync

The client portion of the Windows Time service is loaded, but does not attempt to obtain the time or synchronize the clock. (Note, on Windows Server 2003, the service takes control of the NTP port 123 UDP, the Windows XP & 2000 versions do not).

NT5DS

This is the mode selected by default when a machine is a member of an Active Directory domain. It uses Active Directory to discover a time server (called the "inbound time partner"), then uses either SNTP or LAN Manager to retrieve the time. NT5DS mode does not allow you to specify the server, and on XP or above, the server must be a Windows Domain Controller that provides signed time packets.

AllSync

This mode uses either the NT5DS and NTP methods to synchronize with a server. It tries NT5DS first, and if that fails (perhaps because your machine isn't a member of a domain, or your machine's logon server is not available), it falls back

to the list of manually-configured NTP/SNTP servers.

NTP

Special Interval: seconds

If Special Interval is selected as the NTP sync mode (see below), this field allows you to specify the number of seconds to wait between attempts to synchronize the system clock. This setting corresponds to the SpecialPollInterval value in the [HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient] registry key.

Use the time source fields to enter the IP address or DNS name of NTP time sources you want the Windows Time service to synchronize with.

The **Sync Mode** dropdown box allows you to select the NTP sync mode to use:

- **Special Interval:** As mentioned above, synchronizes using the Special Poll Interval setting
- **Fallback Only:** Specifies this entry should be only used for fallback in the event other methods fail
- **Symmetric Active:** Sends NTP sync requests using symmetric active packets (not compatible with all NTP servers)
- **NTP Client:** Sends NTP sync requests using standard NTP Client packets

The button will instruct the Windows Time service to resynchronize with its time source. The sync may not occur immediately, it happens according to the internal schedule of the time service.

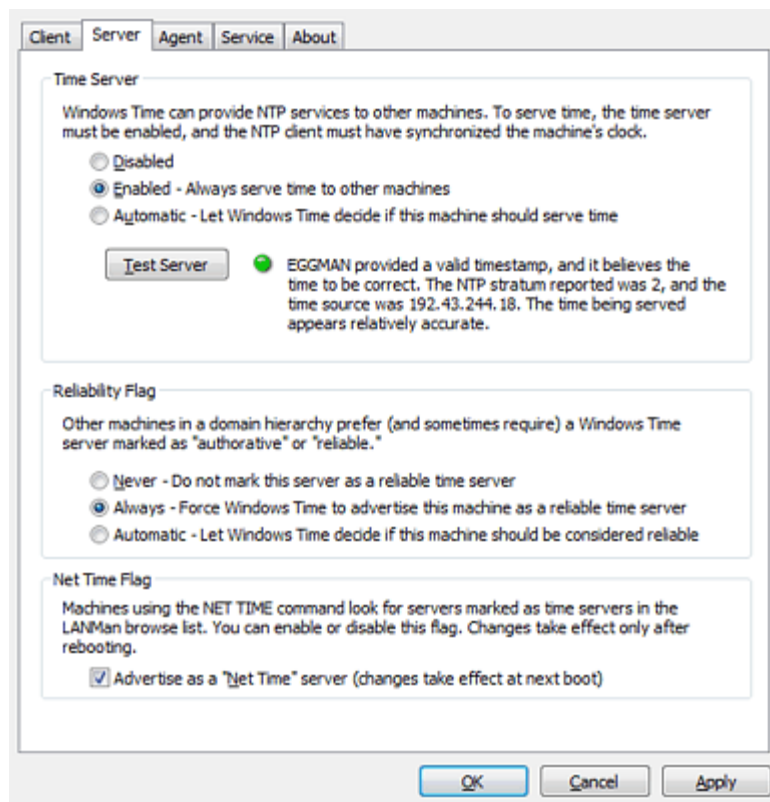
Visual status indicators

The applet will constantly display the current clock variance of each selected source, as well as an indicator light showing the sync status of the local clock. The currently selected time source (inbound time partner) is displayed as well.

Note: The variances shown are being displayed in the foreground by the applet itself. They will appear whether or not the Windows Time Service is running.

Server Tab

The **Server** tab page allows you to set characteristics of the time server function of the Windows Time service.



Windows Time Agent: Server tab [\[Click for larger size\]](#)

The **Time Server** section allows you to specify whether the Windows Time service will act as an NTP server.

Disabled - The Windows Time service will not act as a time server.

Enabled - The Windows Time service will always act as a time server.

Automatic - The Windows Time service will act as a time server if the machine is a domain controller; it will not serve time if the machine is a domain member or stand-alone system.

Click the [Test Server](#) button to test the NTP time server function. The display will indicate if the server is active, what the time source was used, stratum information, and whether it is correctly serving time.

NOTE: This test does not necessarily test the Windows Time NTP server. It merely tests the machine for an NTP server and reports the results. If another program is serving NTP on the machine, the Test Server button will show a response, even if the Windows Time service itself is stopped and disabled.

The **Reliability Flag** section lets you change how Windows Time indicates the reliability of the time being served.

Never - The Windows Time service will not be announced as a reliable time server.

Enabled - The Windows Time service will always be announced as a reliable time server.

Automatic - The Windows Time service will be announced as a reliable time server according to its position in the

domain heirarchy.

"A computer that is configured to be a reliable time source is identified by time clients as the root of the Windows Time service. The root of the Windows Time service is the authoritative server for the domain and typically is configured to retrieve time from an external NTP server or hardware device. Other time servers can be configured as a reliable time source to optimize how time is transferred throughout the domain hierarchy. If a domain controller is configured to be a reliable time source, the Net Logon service announces that domain controller as a reliable time source when it logs on to the network. When other domain controllers look for a time source to synchronize with, they select a reliable source first, if one is available."

--- from [How to configure an authoritative time server in Windows Server 2003](#)

The **Net Time Flag** section lets you determine if Windows Time will announce itself as a LanMan time server.

Advertise as a "Net Time" server... - If this box is checked, clients using the NET TIME command to synchronize will be able to find and use this server. This really has nothing to do with the Windows Time service, but is here for your convenience.

Agent Tab

The **Agent** tab page let you enable special data collection and reporting functions to report on the activity of the Windows Time service.

Note: This tab page is unavailable on Windows 2000.

The collected information can be displayed on real-time clock drift graphs or gathered centrally by the [Domain Time II Audit Server](#) to create an audit trail of the time sync history of this machine. Domain Time II Audit Server can also monitor the sync status of the Windows Time service via the Windows Time Agent and raise administrative alarms if the time is not synchronized.

Read more about [Domain Time II Audit Server](#).

Agent Options

The Windows Time Agent collects synchronization statistics by checking the clock's accuracy at regular intervals. By default, Agent uses the last-used NTP time source as the reference clock. If Agent cannot determine the last-used time source, it tries each of the manually-configured time sources until it finds a reference clock.

You may override this behavior by specifying a particular time server to be used as the reference time source.

☒ Enable the Windows Time Agent

Reference NTP Server: (leave blank to use the defaults)

Check Interval: minutes

☒ Log a warning in Event Viewer if unable to verify the time

☒ Log an error in Event Viewer if verified time is off by more than seconds

History

Recent synchronization history. This list shows the time servers used by Windows Time to synchronize the clock. The most-recently-used server is on the top.

15 Nov 2009 08:37:30	192.43.244.18	time.nist.gov
15 Nov 2009 08:19:42	192.43.244.18	time.nist.gov
15 Nov 2009 08:09:11	192.43.244.18	time.nist.gov

Windows Time Agent: Agent tab [\[Click for larger size\]](#)

The **Agent Options** section lets you configure the reference clock and logging options of the Windows Time Agent.

Enable the Windows Time Agent

Enables or disables the Agent data collection and reporting functions. If disabled, no logging or drift-clock graphing is possible, and the Agent will not respond to Audit Server audit queries.

Reference NTP Server:

Use this field to specify an NTP server to use as the reference clock for variance calculations. The Agent will compare other time sources and the local clock to this server. If the field is left blank, the Agent will use the last NTP time source used by the Windows Time client as the reference clock.

Check Interval: **minutes**

The Agent will check the sync status of the service at the frequency set in this field.Data points will be added to the drift graphs according to the same schedule.

Log a warning in Event Viewer if unable to verify the time


Logs sync errors in the Windows Application Event Log.

Log a warning in Event Viewer if verified time is off by seconds

Logs errors in the Windows Application Event Log when the clock drifts outside this range.


The **History** section displays a list of the most-recently-used time servers with which the Windows Time service synchronized.

You can also display the clock drift graph by clicking the button.



Windows Time Service clock drift [\[Click for larger size\]](#)

It's quite likely by examining this data, you'll find the Windows Time service isn't keeping your clocks as synchronized as you need them to be.You'll definitely want to take a look at our [Domain Time II time synchronization software](#).It will keep your clocks looking like this:



Domain Time II clock drift [\[Click for larger size\]](#)

Service Tab

Use the Domain Time II Windows Time **Service** tab to stop and start the Windows Time service and enable debug logging.

Windows Time Agent: Service tab [\[Click for larger size\]](#)

The **Service Status** section lets you start or stop the Windows Time service, and set the service startup flag (Automatic, Manual, or Disabled).

The **Service Debug Log File** section allows you to enable an extremely detailed debug log of all Windows Time service activity. This log can be extremely useful in tracking down problems when the Windows Time service is not synchronizing correctly.

Enable Debug Log File

Enables or disables debug logging. You must stop and restart the Windows Time service after changing this setting.

Log File Path and Name:

Use this field to specify the path and filename to use for the debug log file. This path must be on a local hard drive.


Note: On Windows Vista and above, the Windows Time service doesn't have sufficient NTFS file system rights to create its own debug log file (it merely creates "Access Denied" errors in the System event log). You must therefore manually add *Modify* permissions for the **Local Service** account to the folder you select for your debug log file.

Log File Maximum Size: (megabytes)

This sets the maximum size for the debug log. The log can grow quite quickly, so be sure you specify enough space to capture events for the time period you want to examine.

Log File Entries:

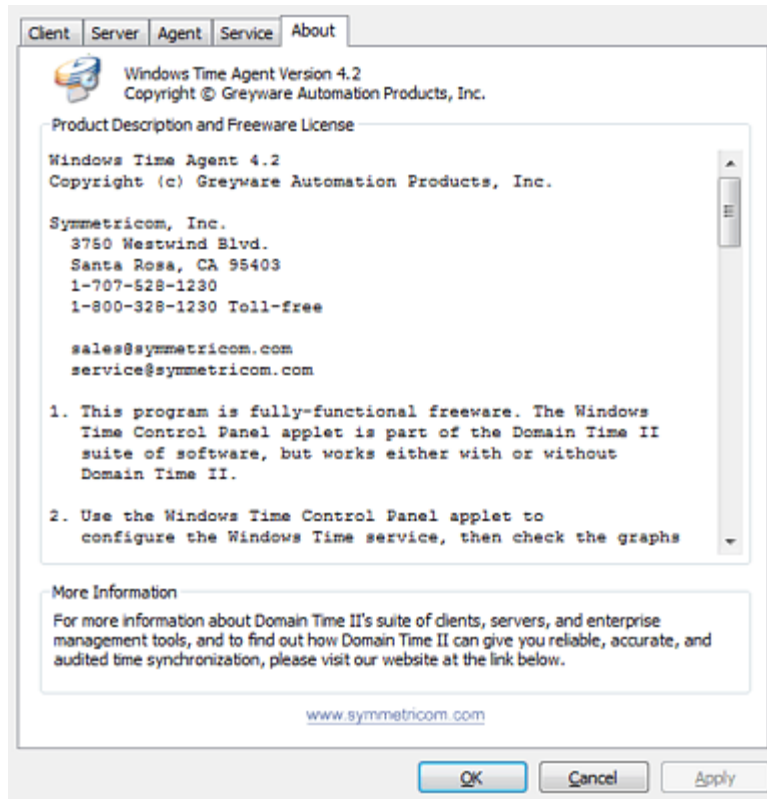
This field specifies what level of detail you want included in the debug log. **0-300** is the maximum level of detail. Microsoft documentation indicates a setting of **0-116** will give basic details.

Click the  button to view the log in real-time.

About Tab

Domain Time II Agent Version 4.2

The **About** tab displays a short description of the program and the Freeware License.



Windows Time Agent: About tab [Click for larger size]