

Table of Contents

<u>Network Time Protocol: Best Practices White Paper</u>	1
<u>Introduction</u>	1
<u>Background Information</u>	1
<u>Terminology</u>	2
<u>Overview</u>	3
<u>Device Overview</u>	3
<u>NTP Overview</u>	4
<u>NTP Design Criteria</u>	5
<u>Association Modes</u>	5
<u>NTP Architecture</u>	7
<u>Clock Technology and Public Time Servers</u>	8
<u>Example NTP Deployments</u>	8
<u>WAN Time Distribution Network</u>	8
<u>High Stratum Campus Time Distribution Network</u>	11
<u>Low Stratum Campus Time Distribution Network</u>	12
<u>Process Definitions</u>	13
<u>Process Owner</u>	14
<u>Process Goals</u>	14
<u>Process Performance Indicators</u>	15
<u>Process Inputs</u>	15
<u>Process Outputs</u>	15
<u>Task Definitions</u>	15
<u>Initialization Tasks</u>	15
<u>Iterative Tasks</u>	17
<u>Data Identification</u>	18
<u>General Data Characteristics</u>	18
<u>SNMP Data Identification</u>	19
<u>Data Collection</u>	24
<u>SNMP Data Collection</u>	24
<u>Data Presentation</u>	24
<u>NTP Critical Node Report</u>	24
<u>NTP Interesting Node Report</u>	25
<u>NTP Configuration Report</u>	25
<u>Related Information</u>	26

Network Time Protocol: Best Practices White Paper

Introduction

Background Information

Terminology

Overview

- Device Overview

- NTP Overview

NTP Design Criteria

- Association Modes

- NTP Architecture

- Clock Technology and Public Time Servers

Example NTP Deployments

- WAN Time Distribution Network

- High Stratum Campus Time Distribution Network

- Low Stratum Campus Time Distribution Network

Process Definitions

- Process Owner

- Process Goals

- Process Performance Indicators

- Process Inputs

- Process Outputs

Task Definitions

- Initialization Tasks

- Iterative Tasks

Data Identification

- General Data Characteristics

- SNMP Data Identification

Data Collection

- SNMP Data Collection

Data Presentation

- NTP Critical Node Report

- NTP Interesting Node Report

- NTP Configuration Report

Related Information

Introduction

Internet Protocol (IP) based networks are quickly evolving from the traditional *best effort* delivery model to a model where performance and reliability need to be quantified and, in many cases, guaranteed with Service Level Agreements (SLAs). The need for greater insight into network characteristics has led to significant research efforts being targeted at defining metrics and measurement capabilities to characterize network behavior. The foundation of many metric methodologies is the measurement of time.

Background Information

Network time synchronization, to the degree required for modern performance analysis, is an essential exercise. Depending on the business models, and the services being provided, the characterization of network

performance can be considered an important competitive service differentiator. In these cases, great expense may be incurred deploying network management systems and directing engineering resources towards analyzing the collected performance data. However, if proper attention is not given to the often-overlooked principle of time synchronization, those efforts may be rendered useless.

This document describes a hypothetical process definition for conducting network management functions for the Network Time Protocol (NTP). It is intended that this hypothetical procedure be used as an informational example and customized by an organization to assist in meeting internal objectives.

The information provided by this paper is presented in several major sections, which are described below.

The Terminology section provides general definitions of terms concerning time synchronization.

The Overview section provides background information on network element hardware related to system time, a technological overview of NTP, and key design aspects for the NTP architecture.

The Example NTP Deployment section provides NTP deployment examples with sample configurations for WAN, high stratum campus, and low stratum campus time distribution networks.

The Process Definitions section provides an overview of the process definitions used to accomplish NTP management. The process details are described in terms of goals, performance indicators, inputs, outputs, and individual tasks.

The Task Definitions section provides detailed process task definitions. Each task is described in terms of objectives, task inputs, task outputs, resources required to accomplish the task, and job skills needed for a task implementer.

The Data Identification section describes data identification for NTP. Data identification considers the source of the information. For example, information may be contained in the Simple Network Management Protocol (SNMP) Management Information Base (MIB), in Syslog generated log files, or by internal data structures that can only be accessed by the command line interface (CLI).

The Data Collection section describes the collection of the NTP data. The collection of the data is closely related to the location of the data. For example, SNMP MIB data is collected by several mechanisms such as traps, Remote Monitoring (RMON) alarms and events, or polling. Data maintained by internal data structures is collected by automatic scripts or by a user manually logging into the system to issue the CLI command and recording the output.

The Data Presentation section provides report format examples of how the data may be presented.

Terminology

- **Accuracy** The proximity of the clock's absolute value to the offset of zero.
- **Accurate** When a clock's offset is zero at a particular moment in time.
- **Drift** The measurement in the variation of skew, or the second derivation of the clock's offset with respect to time.
- **Joint resolution** When comparing clocks, this is the sum of the resolutions of C1 and C2. The joint resolution then indicates a conservative lower bound on the accuracy of any time intervals computed by subtracting time stamps generated by one clock from those generated by the other.
- **Node** Refers to an instantiation of the NTP protocol on a local processor. A node can also be referred to as a device.

- **Offset** The difference between the time reported by a clock and the true time as defined by Coordinated Universal Time (UTC). If the clock reports a time T_c and the true time is T_t , then the clock's offset is $T_c - T_t$.
- **Peer** Refers to an instantiation of the NTP protocol on a remote processor connected by a network path from the local node.
- **Relative offset** The notion of true time is replaced by the time as reported by clock C_1 , when comparing how two clocks, C_1 and C_2 , compare. For example, clock C_2 's offset relative to C_1 at a particular moment is $T_{c2} - T_{c1}$, the instantaneous difference in time reported by C_2 and C_1 .
- **Resolution** The smallest unit by which a clock's time is updated. Resolution is defined in terms of seconds. However, resolution is relative to the clock's reported time and not to true time. For example, a resolution of 10 milliseconds means that the clock updates its notion of time in 0.01 second increments and does not mean that this is the true amount of time between updates.

Note: Clocks can have very fine resolutions and still be inaccurate.

- **Skew** A clock's frequency difference, or first derivative of its offset with respect to time.
- **Synchronize** When two clocks are accurate with respect to one another (relative offset is zero), they are synchronized. Clocks can be synchronized and still inaccurate in terms of how well they tell true time.

Overview

Device Overview

The heart of the time service is the system clock. The system clock runs from the moment the system starts and keeps track of the current date and time. The system clock can be set from a number of sources and, in turn, can be used to distribute the current time through various mechanisms to other systems. Some routers contain a battery-powered calendar system that tracks the date and time across system restarts and power outages. This calendar system is always used to initialize the system clock when the system is restarted. It can also be considered as an authoritative source of time and redistributed through NTP if no other source is available. Furthermore, if NTP is running, the calendar can be periodically updated from NTP, compensating for the inherent drift in the calendar time. When a router with a system calendar is initialized, the system clock is set based on the time in its internal battery-powered calendar. On models without a calendar, the system clock is set to a predetermined time constant. The system clock can be set from the sources listed below.

- NTP
- Simple Network Time Protocol (SNTP)
- Virtual Integrated Network Service (VINES) Time Service
- Manual configuration

Certain low-end Cisco devices only support SNTP. SNTP is a simplified, client-only version of NTP. SNTP can only receive the time from NTP servers and cannot be used to provide time services to other systems. SNTP typically provides time within 100 milliseconds of the accurate time. In addition, SNTP does not authenticate traffic, although you can configure extended access lists to provide some protection. An SNTP client is more vulnerable to misbehaving servers than an NTP client and should only be used in situations where strong authentication is not required.

The system clock provides time to the services listed below.

- NTP
- VINES time service
- User **show** commands
- Logging and debugging messages

The system clock keeps track of time internally based on UTC, also known as Greenwich Mean Time (GMT). You can configure information about the local time zone and daylight savings time so that the time is displayed correctly relative to the local time zone. The system clock keeps track of whether the time is authoritative or not. If it is not authoritative, the time will be available only for display purposes and will not be redistributed.

NTP Overview

NTP is designed to synchronize the time on a network of machines. NTP runs over the User Datagram Protocol (UDP), using port 123 as both the source and destination, which in turn runs over IP. NTP Version 3 RFC 1305 is used to synchronize timekeeping among a set of distributed time servers and clients. A set of nodes on a network are identified and configured with NTP and the nodes form a synchronization subnet, sometimes referred to as an overlay network. While multiple masters (primary servers) may exist, there is no requirement for an election protocol.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. An NTP client makes a transaction with its server over its polling interval (from 64 to 1024 seconds) which dynamically changes over time depending on the network conditions between the NTP server and the client. No more than one NTP transaction per minute is needed to synchronize two machines.

NTP uses the concept of a stratum to describe how many NTP hops away a machine is from an authoritative time source. For example, a stratum 1 time server has a radio or atomic clock directly attached to it. It then sends its time to a stratum 2 time server through NTP, and so on. A machine running NTP automatically chooses the machine with the lowest stratum number that it is configured to communicate with using NTP as its time source. This strategy effectively builds a self-organizing tree of NTP speakers. NTP performs well over the non-deterministic path lengths of packet-switched networks, because it makes robust estimates of the following three key variables in the relationship between a client and a time server.

- Network delay
- Dispersion of time packet exchanges A measure of maximum clock error between the two hosts.
- Clock offset The correction applied to a client's clock to synchronize it.

Clock synchronization at the 10 millisecond level over long distance wide-area networks (WANs) (2000 km), and at the 1 millisecond level for local-area networks (LANs), is routinely achieved.

NTP avoids synchronizing to a machine whose time may not be accurate in two ways. First of all, NTP never synchronizes to a machine that is not synchronized itself. Secondly, NTP compares the time reported by several machines, and will not synchronize to a machine whose time is significantly different than the others, even if its stratum is lower.

The communications between machines running NTP (associations) are usually statically configured. Each machine is given the IP address of all machines with which it should form associations. Accurate timekeeping is made possible by exchanging NTP messages between each pair of machines with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each machine can be configured to send or receive broadcast messages. However, the accuracy of timekeeping is marginally reduced because the information flow is one-way only.

The time kept on a machine is a critical resource and it is strongly recommend that you use the security features of NTP to avoid the accidental or malicious setting of incorrect time. The two security features available are an access list-based restriction scheme and an encrypted authentication mechanism.

Cisco's implementation of NTP supports the stratum 1 service in certain Cisco IOS software releases. If a release supports the **ntp refclock** command, it is possible to connect a radio or atomic clock. Certain releases of Cisco IOS support either the Trimble Palisade NTP Synchronization Kit (Cisco 7200 series routers only) or the Telecom Solutions Global Positioning System (GPS) device. If the network uses the public time servers on the Internet and the network is isolated from the Internet, Cisco's implementation of NTP allows a machine to be configured so that it acts as though it is synchronized through NTP, when in fact it has determined the time using other means. Other machines then synchronize to that machine through NTP.

NTP Design Criteria

Each client in the synchronization subnet, which may also be a server for higher stratum clients, chooses one of the available servers to synchronize to. This is usually from among the lowest stratum servers it has access to. However, this is not always an optimal configuration, because NTP also operates under the premise that each server's time should be viewed with a certain amount of distrust. NTP prefers to have access to several sources of lower stratum time (at least three) since it can then apply an agreement algorithm to detect insanity on the part of any one of these. Normally, when all servers are in agreement, NTP chooses the best server in terms of lowest stratum, closest (in terms of network delay), and claimed precision. The implication is that, while one should aim to provide each client with three or more sources of lower stratum time, several of these will only be providing backup service and may be of lesser quality in terms of network delay and stratum. For example, a same-stratum peer that receives time from lower stratum sources the local server doesn't access directly, can also provide good backup service.

Association Modes

The following sections describe the associating modes used by NTP servers to associate with each other.

- Client/Server
- Symmetric Active/Passive
- Broadcast

Client/Server Mode

Dependent clients and servers normally operate in client/server mode, in which a client or dependent server can be synchronized to a group member, but no group member can synchronize to the client or dependent server. This provides protection against malfunctions or protocol attacks.

Client/server mode is the most common Internet configuration. It operates in the classic remote-procedure-call (RPC) paradigm with stateless servers. In this mode, a client sends a request to the server and expects a reply at some future time. In some contexts, this would be described as a poll operation, in that the client polls the time and authentication data from the server. A client is configured in client mode by using the server command and specifying the domain name server (DNS) name or address. The server requires no prior configuration.

In a common client/server model, a client sends an NTP message to one or more servers and processes the replies as received. The server interchanges addresses and ports, overwrites certain fields in the message, recalculates the checksum, and returns the message immediately. Information included in the NTP message allows the client to determine the server time with respect to local time and adjust the local clock accordingly. In addition, the message includes information to calculate the expected timekeeping accuracy and reliability, as well as select the best server.

Servers that provide synchronization to a sizeable population of clients normally operate as a group of three or more mutually redundant servers, each operating with three or more stratum 1 or stratum 2 servers in

client/server modes, as well as all other members of the group in symmetric modes. This provides protection against malfunctions in which one or more servers fail to operate or provide incorrect time. The NTP algorithms are engineered to resist attacks when some fraction of the configured synchronization sources accidentally or purposely provide incorrect time. In these cases, a special voting procedure is used to identify spurious sources and discard their data. In the interest of reliability, selected hosts can be equipped with external clocks and used for backup in case of failure of the primary and/or secondary servers, or communication paths between them.

Configuring an association in client mode, usually indicated by a server declaration in the configuration file, indicates that one wishes to obtain time from the remote server, but that one is not willing to provide time to the remote server.

Symmetric Active/Passive Mode

Symmetric active/passive mode is intended for configurations where a group of low stratum peers operate as mutual backups for each other. Each peer operates with one or more primary reference sources, such as a radio clock, or a subset of reliable secondary servers. Should one of the peers lose all reference sources or simply cease operation, the other peers automatically reconfigure so that time values can flow from the surviving peers to all the others in the clique. In some contexts this is described as a *push-pull* operation, in that the peer either pulls or pushes the time and values depending on the particular configuration.

Configuring an association in symmetric-active mode, usually indicated by a peer declaration in the configuration file, indicates to the remote server that one wishes to obtain time from the remote server and that one is also willing to supply time to the remote server if necessary. This mode is appropriate in configurations involving a number of redundant time servers interconnected through diverse network paths, which is presently the case for most stratum 1 and stratum 2 servers on the Internet today.

Symmetric modes are most often used between two or more servers operating as a mutually redundant group. In these modes, the servers in the group members arrange the synchronization paths for maximum performance, depending on network jitter and propagation delay. If one or more of the group members fail, the remaining members automatically reconfigure as required.

A peer is configured in symmetric active mode by using the **peer** command and specifying the DNS name or address of the other peer. The other peer is also configured in symmetric active mode in this way.

Note: If the other peer is not specifically configured in this way, a symmetric passive association is activated upon arrival of a symmetric active message. Since an intruder can impersonate a symmetric active peer and inject false time values, symmetric mode should always be authenticated.

Broadcast and/or Multicast Mode

Where the requirements in accuracy and reliability are modest, clients can be configured to use broadcast and/or multicast modes. Normally, these modes are not utilized by servers with dependent clients. The advantage is that clients do not need to be configured for a specific server, allowing all operating clients to use the same configuration file. Broadcast mode requires a broadcast server on the same subnet. Since broadcast messages are not propagated by routers, only broadcast servers on the same subnet are used.

Broadcast mode is intended for configurations involving one or a few servers and a potentially large client population. A broadcast server is configured using the **broadcast** command and a local subnet address. A broadcast client is configured using the **broadcastclient** command, allowing the broadcast client to respond to broadcast messages received on any interface. Since an intruder can impersonate a broadcast server and inject false time values, this mode should always be authenticated.

NTP Architecture

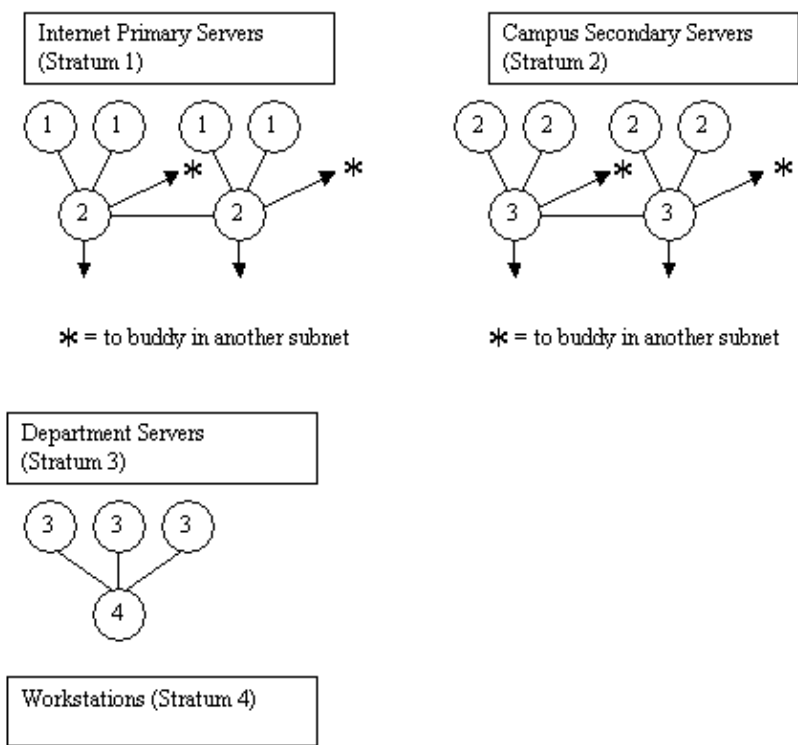
The following three structures are available for NTP architecture.

- Flat peer structure
- Hierarchical structure
- Star structure

In a flat peer structure, all the routers peer with each other, with a few geographically separate routers configured to point to external systems. The convergence of time becomes longer with each new member of the NTP mesh.

In a hierarchical structure, the routing hierarchy is copied for the NTP hierarchy. Core routers have a client/server relationship with external time sources, the internal time servers have a client/server relationship with the core routers, the internal customer (non time servers) routers have a client/server relationship with the internal time servers, and so on down the tree. These relationships are called hierarchy scales. A hierarchical structure is the preferred technique because it provides consistency, stability, and scalability.

A scalable NTP architecture has a hierarchical structure as seen in the diagram below.




In a star structure, all the routers have a client/server relationship with a few time servers in the core. The dedicated time servers are the center of the star and are usually UNIX systems synchronized with external time sources, or their own GPS receiver.

Clock Technology and Public Time Servers

The Internet NTP subnet presently includes over 50 public primary servers synchronized directly to UTC by radio, satellite, or modem. Normally, client workstations and servers with a relatively small number of clients do not synchronize to primary servers. Approximately 100 public secondary servers are synchronized to the primary servers, providing synchronization to a total in excess of 100,000 clients and servers on the Internet. The Public NTP Time Servers lists are updated frequently. There are also numerous private primary and secondary servers not normally available to the public.

In certain cases, where highly accurate time services are required on the private enterprise, such as one-way metrics for Voice over IP (VoIP) measurements, network designers may choose to deploy private external time sources. The diagram below shows a comparative graph of the relative accuracy of the current technologies.




Until recently, the use of external time sources have not been widely deployed in enterprise networks due to the high cost of quality external time sources. However, as the Quality of Service (QoS) requirements increase and the cost of the time technology continues to decrease, external time sources for enterprise networks are becoming a viable option.

Example NTP Deployments

WAN Time Distribution Network

In the diagram below, a corporate autonomous system (AS) obtains time information from three public time servers. The corporate AS is shown as Area 0 and Area 1 time servers. In this example, the NTP hierarchy follows the Open Shortest Path First (OSPF) hierarchy. However, OSPF is not a prerequisite for NTP. It is only used as an illustrative example. NTP may be deployed along other logical hierarchical boundaries such as an Enhanced Interior Gateway Routing Protocol (EIGRP) hierarchy or the standard Core/Distribution/Access hierarchy.



The following is the Cisco IOS configuration for device A0-R1 in the above diagram.

```

clock timezone CST -5
clock summer-time CDT recurring

!--- This router has a hardware calendar.
!--- To configure a system as an
!--- authoritative time source for a network
!--- based on its hardware clock (calendar),
!--- use the clock calendar-valid global
!--- configuration command. Notice later that
!--- NTP will be allowed to update the calendar
!--- and Cisco IOS will be configured to be an
!--- NTP master clock source.
!--- Cisco IOS will then obtain its clock from
!--- the hardware calendar.

clock calendar-valid

!--- This allows NTP to update the hardware
!--- calendar chip.

ntp calendar-update

!--- Configures the Cisco IOS software as an
!--- NTP master clock to which peers synchronize
!--- themselves when an external NTP source is
!--- not available. Cisco IOS will obtain the
!--- clock from the hardware calendar based on
!--- the previous line. This line will keep the

```

```
!--- whole network in Sync even if Router1 loses  
!--- its signal from the Internet. Assume, for  
!--- this example, that the Internet time servers  
!--- are stratum 2.
```

```
ntp master 3
```

```
!--- When the system sends an NTP packet, the  
!--- source IP address is normally set to the  
!--- address of the interface through which the  
!--- NTP packet is sent.  
!--- Change this to use loopback0.
```

```
ntp source Loopback0
```

```
!--- Enables NTP authentication.
```

```
ntp authenticate  
ntp authentication-key 1234 md5 104D000A0618 7  
ntp trusted-key 1234
```

```
!--- Configures the access control groups for  
!--- the public servers and peers for additional  
!--- security.
```

```
access-list 5 permit <I-TS-1>  
access-list 5 permit <I-TS-2>  
access-list 5 permit <I-TS-3>  
access-list 5 permit <A0-R2>  
access-list 5 permit <A0-R3>  
access-list 5 deny any
```

```
!--- Configures the access control groups for the  
!--- clients to this node for additional security.
```

```
access-list 6 permit <A1-R1>  
access-list 6 permit <A1-R2>  
access-list 6 permit <A1-R3>  
access-list 6 deny any
```

```
!--- Restricts the IP addresses for the peers  
!--- and clients.
```

```
ntp access-group peer 5  
ntp access-group serve-only 6
```

```
!--- Fault tolerant configuration polling for 3 NTP  
!--- public servers, peering with 2 local servers.
```

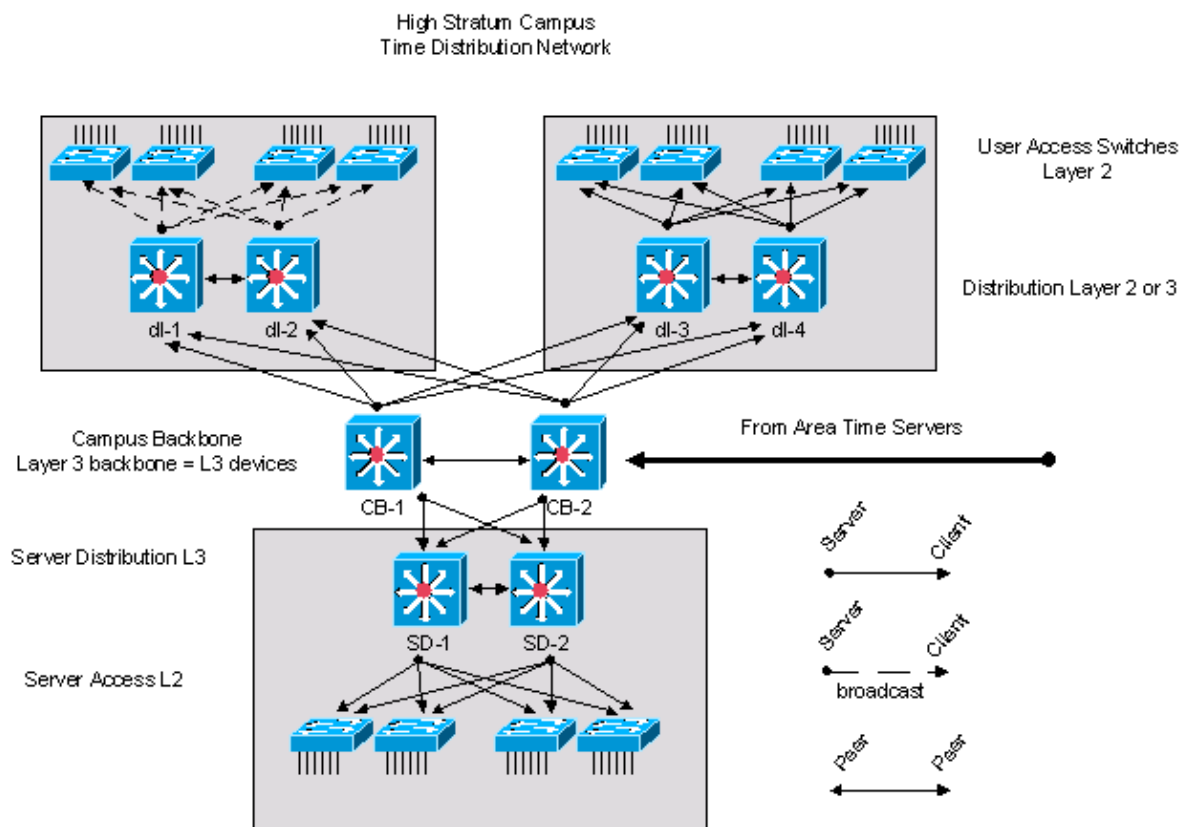
```
ntp server <I-TS-1>  
ntp server <I-TS-2>
```

```
ntp server <I-TS-3>
ntp peer <A0-R2>
ntp peer <A0-R3>
```

High Stratum Campus Time Distribution Network

The previous section described a WAN time distribution network. This section moves one step down in the hierarchy to discuss time distribution on a high stratum campus network.

The primary difference when considering time distribution on a high stratum campus network is the potential usage of the broadcast association mode. As described earlier, the broadcast association mode simplifies the configurations for the LANs, but reduces the accuracy of the time calculations. Therefore, the trade-off in maintenance costs must be considered against accuracy in performance measurements.



The high stratum campus network, shown in the diagram above, is taken from the standard Cisco Campus network design and contains three components. The campus core consists of two Layer 3 devices labeled CB-1 and CB-2. The server component, located in the lower section of the figure, has two Layer 3 routers labeled SD-1 and SD-2. The remaining devices in the server block are Layer 2 devices. In the upper left, there is a standard access block with two Layer 3 distribution devices labeled dl-1 and dl-2. The remaining devices are Layer 2 switches. In this client access block, the time is distributed using the broadcast option. In the upper right, there is another standard access block that uses a client/server time distribution configuration.

The campus backbone devices are synchronized to the area time servers in a client/server model.

The configuration for the dl-1 Layer 3 distribution devices is shown below.

```
!--- In this case, dl-1 will be a broadcast server  
!--- for the Layer 2 LAN.
```

```
internet Ethernet0  
ntp broadcast
```

```
clock timezone CST -5  
clock summer-time CDT recurring
```

```
!--- When the system sends an NTP packet, the  
!--- source IP address is normally set to the  
!--- address of the interface through which the  
!--- NTP packet is sent.  
!--- Change this to use loopback0.
```

```
ntp source Loopback0
```

```
!--- Enables NTP authentication.
```

```
ntp authenticate  
ntp authentication-key 1234 md5 104D000A0618 7  
ntp trusted-key 1234
```

```
!--- Configures the access control groups for  
!--- the public servers and peers for  
!--- additional security.
```

```
access-list 5 permit <CB-1>  
access-list 5 permit <CB-2>  
access-list 5 permit <dl-2>  
access-list 5 deny any
```

```
!--- Restricts the IP addresses for the peers  
!--- and clients.
```

```
ntp access-group peer 5
```

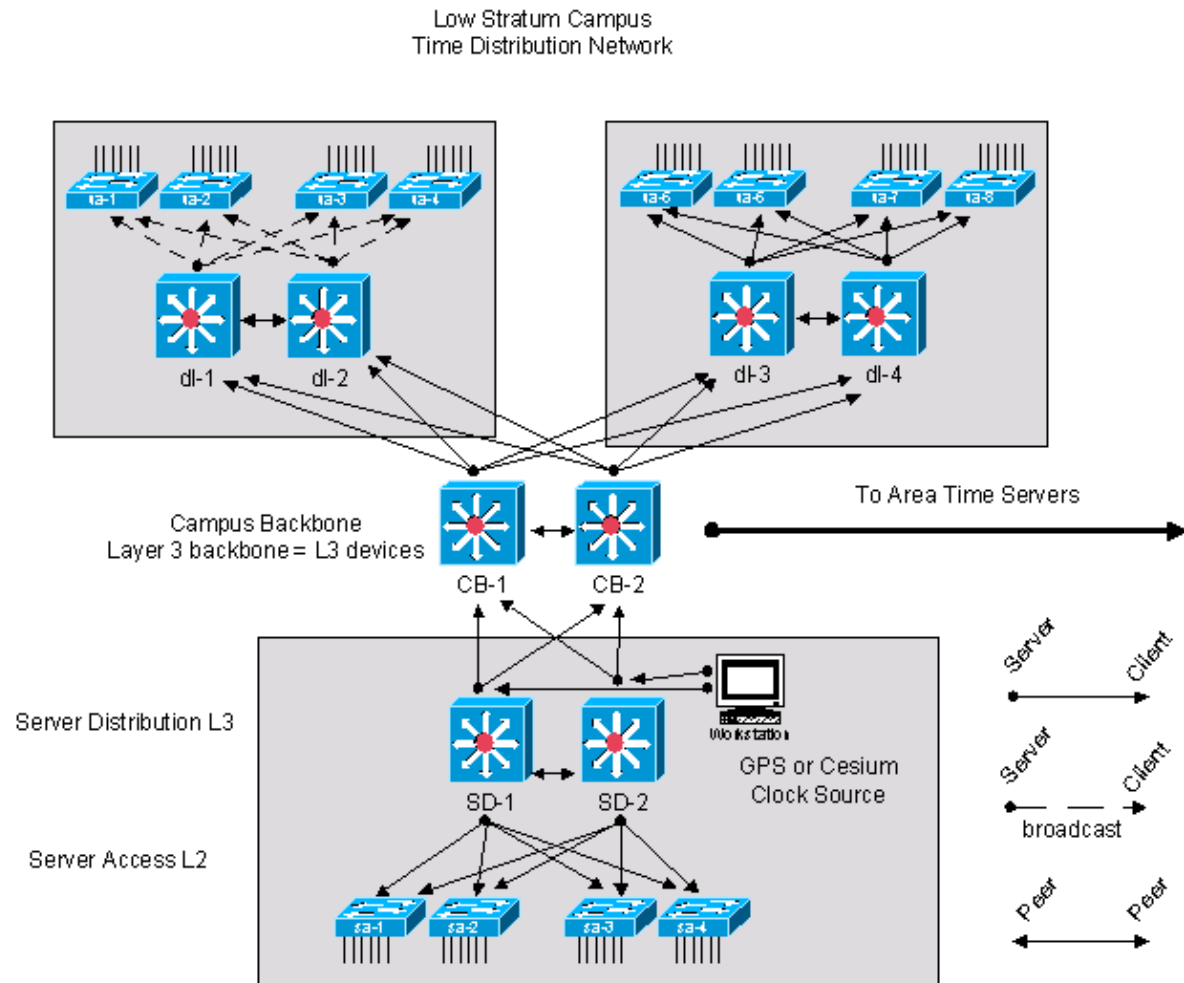
```
!--- Fault tolerant configuration polling 2  
!--- local time servers and 1 local peer.
```

```
ntp server <CB-1>  
ntp server <CB-2>  
ntp peer <dl-2>
```

Low Stratum Campus Time Distribution Network

In the diagram below, a GPS or Cesium time source is provided at the central data center for the low stratum campus network. This provisions a stratum 1 time source on the private network. If there are multiple GPS or Cesium time sources located in the private network, then the time distribution in the private network should be modified to take advantage of the available time sources.

In general, the same principles and configurations apply as with the previous examples. The primary difference in this case is that the root of the synchronization tree is a private time source rather than a public time source from the Internet. This changes the design of the time distribution network to take advantage of the high accuracy private time source. The private time source is distributed throughout the private network using the principles of hierarchy and modularity that have been described in the previous sections.

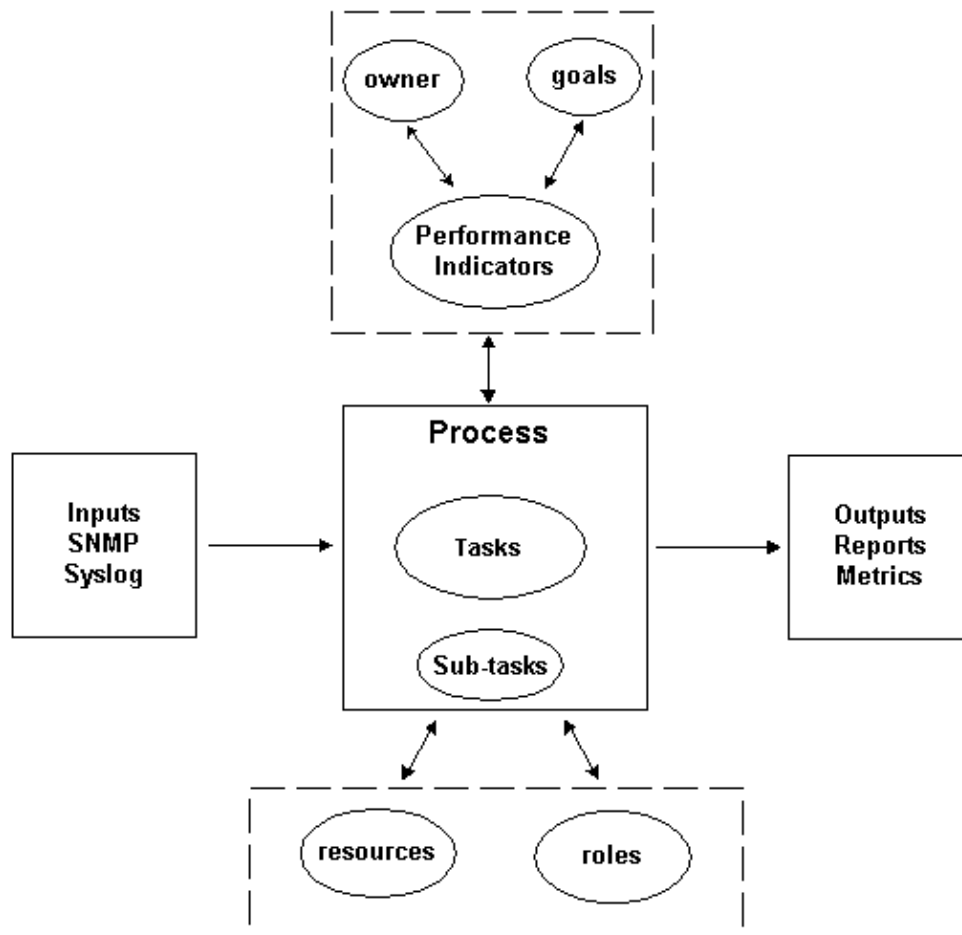


Process Definitions

A process definition is a connected series of actions, activities, and changes performed by agents with the intent of satisfying a purpose or achieving a goal.

Process control is the process of planning and regulating, with the objective of performing a process in an effective and efficient way.

Graphically, this is shown in the diagram below.



The output of the process has to conform to operational norms that are defined by an organization and are based on business objectives. If the process conforms to the set of norms, the process is considered effective since it can be repeated, measured, managed, and it contributes to the business objectives. If the activities are carried out with a minimum effort, the process is also considered efficient.

Process Owner

Processes span various organizational boundaries. Therefore, it is important to have a single process owner who is responsible for the definition of the process. The owner is the focal point for determining and reporting if the process is effective and efficient. If the process fails to be effective or efficient, the process owner drives the modification of the process. Modification of the process is governed by change control and review processes.

Process Goals

Process goals are established to set the direction and scope for the process definition. Goals are also used to define metrics that are used to measure the effectiveness of a process.

The goal of this process is to provide criteria to be documented during the NTP design phase, and to provide an audit capability for a deployed NTP architecture ensuring long-term compliance with the intended design.

Process Performance Indicators

Process performance indicators are used to gauge the effectiveness of the process definition. The performance indicators should be measurable and quantifiable. For instance, the performance indicators listed below are either numeric or measured by time.

- The length of time required to cycle through the entire process.
- The frequency of execution required in order to proactively detect NTP issues before they impact users.
- The network load associated with the execution of the process.
- The number of corrective actions recommended by the process.
- The number of corrective actions implemented as a result of the process.
- The length of time required to implement corrective actions.
- The backlog of corrective actions.
- The errors in troubleshooting or problem diagnosis attributed to NTP related issues.
- The number of items added, removed, or modified in the seed file. This is an indication of accuracy and stability.

Process Inputs

Process inputs are used to define criteria and prerequisites for a process. Many times, identification of process inputs provides information on external dependencies. A list of inputs related to NTP management is provided below.

- NTP design documentation
- NTP MIB data collected by SNMP polling

Process Outputs

The process outputs are defined as follows:

- NTP configuration reports defined in the Data Presentation section of this paper
- NTP corrective actions

Task Definitions

The following sections define the initialization and iterative tasks associated with NTP management.

Initialization Tasks

Initialization tasks are executed once during the implementation of the process and should not be executed during each iteration of the process.

Create the NTP Design

In verifying prerequisite tasks, if it is determined that any one of the tasks is not implemented or does not provide sufficient information to effectively serve the needs of this procedure, this fact should be documented by the process owner and submitted to management. The table below outlines the prerequisite initialization tasks.

Prerequisite Task	Description
Task objectives	Create a detailed design document for the NTP architecture that meets design requirements and cost objectives
Task inputs	<ul style="list-style-type: none"> • Design technical and economic requirements • Existing network design documentation • Criteria defining required aspects to be recorded in the design to enable management functions • IT application deployment information • Performance monitoring requirements
Task output	NTP design documentation
Task resources	Network engineer architect Network operations architect
Task roles	Network design technical approval by Engineering and Operations reviewers Network design costs approved by responsible budget manager

Create a Seed File

The NTP management process requires the use of a seed file to remove the need for a network discovery function. The seed file records the set of routers that are governed by the NTP process and is also used as a focal point to coordinate with the change management processes in an organization. For example, if new nodes are entered into the network, they need to be added to the NTP seed file. If changes are made to the SNMP community names because of security requirements, those modifications need to be reflected in the seed file. The table below outlines the processes for creating a seed file.

Prerequisite Task	Description
Task objectives	<p>Create seed file that identifies three categories of network devices</p> <ol style="list-style-type: none"> 1. Critical devices Polled on a frequent basis for configuration information 2. Interesting devices Polled less frequently 3. All NTP enabled devices Polled the least amount
Task inputs	NTP design documentation Network topology documentation
Task output	Seed file
Task resources	Design criteria that will be used to identify and prioritize the nodes involved in the NTP

Baseline NTP Performance Parameters

Several of the parameters available for monitoring the NTP network exhibit some normal expected variations. The process of baselining is used to characterize the normal expected variations and to set thresholds that define unexpected or abnormal conditions. This task is used to baseline the variable set of parameters for the NTP architecture. For a more detailed discussion of baselining techniques see the Baseline Process: Best Practice White Paper.

Process	Description
Task objectives	Baseline variable parameters
Task inputs	Identify variable parameters cntpSysRootDelay cntpSysRootDispersion cntpPeersRootDelay cntpPeersRootDispersion cntpPeersOffset cntpPeersDelay cntpPeersDispersion
Task outputs	Baseline values and thresholds
Task resources	Tools for collecting SNMP data and calculating
Task role	baselines Network Engineer NMS Engineer

Iterative Tasks

Iterative tasks are executed during each iteration of the process and their frequency is determined and modified in order to improve the performance indicators.

Maintain the Seed File

The seed file is critical for the effective implementation of the NTP management process. Therefore, the current state of the seed file must be actively managed. Changes to the network that impact the contents of the seed file need to be tracked by the NTP management process owner.

Process	Description
Task objectives	Maintain accuracy of the seed file
Task inputs	Information on network changes
Task outputs	Seed file
Task resources	Reports, notifications, meetings concerning changes
Task role	Network Engineer NMS Engineer

Execute the NTP Node Scan

Collect information on critical, interesting, and configuration scans defined by this procedure. Run these three scans at different frequencies.

Critical nodes are devices that are seen as very important to the performance collection data points. The critical node scan is executed often, for example, hourly, or on a demand basis before and after changes. Interesting nodes are devices that are deemed important to the overall integrity of the NTP architecture, but may not be in the time synchronization tree for critical performance data collection. This report is executed periodically, for example, daily or monthly. The configuration report is a comprehensive and resource intense report that is used to characterize the overall NTP deployment configuration against design records. This report is executed less frequently, for example, monthly or quarterly. An important point to consider is that the frequency that the reports are collected can be adjusted based on the observed stability of the NTP architecture and business needs.

Process	Description
Task objective	Monitor NTP architecture
Task input	Network device data
Task output	Reports
Task resources	Software applications to collect data and produce reports
Task role	Network engineer

Review the NTP Node Reports

This task requires that the critical, interesting, and configuration reports are reviewed and analyzed. If issues are detected, then corrective actions should be initiated.

Process	Description
Task inputs	Scan reports
Task outputs	Stability analysis Corrective actions
Task resources	Access to network devices for further investigation and verification
Task role	Network engineer

Data Identification

General Data Characteristics

The following table describes data that is considered interesting for analyzing the NTP architecture.

Data	Description
Nodes ID	A device that has NTP configured
Peers	The configured peers for the device

Synchronization source	The selected peer for synchronization
NTP configuration data	Parameters used to judge the consistency of the NTP design
NTP quality data	Parameters used to characterize the quality of the NTP associations

SNMP Data Identification

Cisco NTP MIB System Group

The NTP SNMP data is defined by the Cisco-NTP-MIB. For current information regarding the releases that support this MIB, use the CCO Feature Navigator tool and select the MIB Locator option. This tool is accessed through the TAC Tools for Voice, Telephony and Messaging Technologies page.

The system group in the Cisco NTP MIB provides information for the target node running NTP. The target node is the destination of the SNMP queries.

Object Name	Object Description
cntpSysStratum	<p>The stratum of the local clock. If the value is set to 1, a primary reference, then the Primary-Clock procedure described in Section 3.4.6, in RFC-1305 is invoked.</p> <p>::= { cntpSystem 2 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.2</p>
cntpSysPrecision	<p>Signed integer indicating the precision of the system clock, in seconds, to the nearest power of two. The value must be rounded to the next larger power of two. For instance, a 50-Hz (20 ms) or 60-Hz (16.67 ms) power-frequency clock is assigned the value -5 (31.25 ms), while a 1000-Hz (1 ms) crystal-controlled clock is assigned the value -9 (1.95 ms).</p> <p>::= { cntpSystem 3 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.3</p>
cntpSysRootDelay	<p>A signed fixed-point number indicating the total round-trip delay in seconds, to the primary reference source at the root of the synchronization subnet.</p>

	<p>::= { cntpSystem 4 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.4</p>
cntpSysRootDispersion	<p>The maximum error in seconds, relative to the primary reference source at the root of the synchronization subnet. Only positive values greater than zero are possible.</p> <p>::= { cntpSystem 5 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.4</p>
cntpSysRefTime	<p>The local time when the local clock was last updated. If the local clock has never been synchronized, the value is zero.</p> <p>::= { cntpSystem 7 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.7</p>
cntpSysPeer	<p>The current synchronization source containing the unique association identifier cntpPeersAssocId of the corresponding peer entry in the cntpPeersVarTable of the peer acting as the synchronization source. If there is no peer, the value is zero.</p> <p>::= { cntpSystem 9 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.9</p>
cntpSysClock	<p>The current local time. Local time is derived from the hardware clock of the particular machine and increments at intervals depending on the design used.</p> <p>::= { cntpSystem 10 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.1.10</p>

Cisco NTP MIB Peer Group – Peers Variable Table

The peer group of the Cisco NTP MIB provides information on the peers of the target node.

Object Name	Object Description
-------------	--------------------

cntpPeersVarTable	<p>This table provides information on the peers with which the local NTP server has associations. The peers are also NTP servers running on different hosts.</p> <p>This is a table of cntpPeersVarEntry</p> <p>::= { cntpPeers 1 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1</p>
cntpPeersVarEntry	<p>Each peers' entry provides NTP information retrieved from a particular peer NTP server. Each peer is identified by a unique association identifier. Entries are automatically created when the user configures the NTP server to be associated with remote peers. Similarly, entries are deleted when the user removes the peer association from the NTP server. Entries can also be created by the management station by setting values for cntpPeersPeerAddress, cntpPeersHostAddress, cntpPeersMode and making the cntpPeersEntryStatus as active (1). At the very least, the management station has to set a value for cntpPeersPeerAddress to make the row active.</p> <p>INDEX { cntpPeersAssocId }</p> <p>::= { cntpPeersVarTable 1 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1</p>
cntpPeersAssocId	<p>An integer value greater than zero that uniquely identifies a peer with which the local NTP server is associated.</p> <p>::= { cntpPeersVarEntry 1 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.1</p>
cntpPeersConfigured	<p>This is a bit indicating that the association was created from</p>

	<p>configuration information and should not be deassociated even if the peer becomes unreachable.</p> <p>::= { cntpPeersVarEntry 2 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.2</p>
cntpPeersPeerAddress	<p>The IP address of the peer. When creating a new association, a value for this object should be set before the row is made active.</p> <p>::= { cntpPeersVarEntry 3 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.3</p>
cntpPeersMode	<p>SYNTAX INTEGER { unspecified (0), symmetricActive (1), symmetricPassive (2), client (3), server (4), broadcast (5), reservedControl (6), reservedPrivate (7) }</p> <p>When creating a new peer association, if no value is specified for this object, it defaults to symmetricActive (1).</p> <p>::= { cntpPeersVarEntry 8 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.8</p>
cntpPeersStratum	<p>The stratum of the peer clock.</p> <p>::= { cntpPeersVarEntry 9 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.9</p>
cntpPeersRootDelay	<p>A signed fixed–point number indicating the total round–trip delay in seconds, from the peer to the primary reference source at the root of the synchronization subnet.</p> <p>::= { cntpPeersVarEntry 13 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.13</p>
cntpPeersRootDispersion	

	<p>The maximum error, in seconds, of the peer clock relative to the primary reference source at the root of the synchronization subnet. Only positive values greater than zero are possible.</p> <p>::= { cntpPeersVarEntry 14 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.14</p>
cntpPeersRefTime	<p>The local time at the peer when its clock was last updated. If the peer clock has never been synchronized, the value is zero.</p> <p>::= { cntpPeersVarEntry 16 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.16</p>
cntpPeersReach	<p>A shift register used to determine the reachability status of the peer, with bits entering from the least significant (rightmost) end. A peer is considered reachable if at least one bit in this register is set to one (object is non-zero). The data in the shift register is populated by the NTP protocol procedures.</p> <p>::= { cntpPeersVarEntry 21 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.21</p>
cntpPeersOffset	<p>The estimated offset of the peer clock relative to the local clock, in seconds. The host determines the value of this object using the NTP clock-filter algorithm.</p> <p>::= { cntpPeersVarEntry 23 }</p> <p>object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.21</p>
cntpPeersDelay	<p>The estimated round-trip delay of the peer clock relative to the local clock over the network path between them, in seconds. The host determines the value of this object using the NTP clock-filter algorithm.</p>

	<pre> ::= { cntpPeersVarEntry 24 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.24</pre>
cntpPeersDispersion	<p>The estimated maximum error of the peer clock relative to the local clock over the network path between them, in seconds. The host determines the value of this object using the NTP clock-filter algorithm.</p> <pre> ::= { cntpPeersVarEntry 25 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.25</pre>

Data Collection

SNMP Data Collection

All of the information required by this procedure can be collected through SNMP queries. In order to parse the data and produce the reports, custom scripts or software programs will have to be developed.

Data Presentation

NTP Critical Node Report

Critical nodes are devices that are important in the synchronization tree of selected performance data collection points. If there is a high revenue VoIP service being monitored and one-way-delay-variation metrics are being collected, then the source and destination nodes where the time stamps are recorded are considered critical nodes.

In this example, the NTP design has been established following an example OSPF hierarchy. Therefore, the reports described below are formatted to group the NTP devices according to the OSPF area of the device. In cases where a node has interfaces in multiple areas, a decision must be made by the report generation software as to which area the node will be listed for report purposes. As mentioned earlier, OSPF is not a prerequisite for NTP. It is only used in this paper as an illustrative example.

Area	Device	Device Data	Value
AreaId #n	DeviceId #1	cntpSysStratum	
		cntpSysPrecision	
		cntpSysRootDelay	
		cntpSysRootDispersion	
		cntpSysRefTime	
		cntpSysPeer	
		cntpSysClock	

DeviceId #n	cntpSysStratum	
	cntpSysPrecision	
	cntpSysRootDelay	
	cntpSysRootDispersion	
	cntpSysRefTime	
	cntpSysPeer	
	cntpSysClock	

NTP Interesting Node Report

The format of the interesting node report is the same as the format for the critical node report. Interesting nodes are nodes that are considered important to the overall NTP architecture, but may not directly contribute to the time synchronization of critical performance monitoring points.

NTP Configuration Report

The configuration report is a comprehensive report that collects information on the overall NTP architecture. This report is used to record and verify the NTP deployment against design records.

Area	Device	Peer	Peer Data	Value	
AreaId #n	DeviceId #n	PeerId #1	cntpPeersAssocId		
			cntpPeersConfigured		
			cntpPeersPeerAddress		
			cntpPeersMode		
			cntpPeersStratum		
			cntpPeersRootDelay		
			cntpPeersRootDispersion		
			cntpPeersRefTime		
			cntpPeersReach		
			cntpPeersOffset		
			cntpPeersDelay		
			cntpPeersDispersion		
			PeerId #n	cntpPeersAssocId	
				cntpPeersConfigured	
cntpPeersPeerAddress					
cntpPeersMode					
cntpPeersStratum					
cntpPeersRootDelay					
cntpPeersRootDispersion					

			cntpPeersRefTime	
			cntpPeersReach	
			cntpPeersOffset	
			cntpPeersDelay	
			cntpPeersDispersion	

Related Information

- **RFC 1305 Network Time Protocol**
 - **RFC 2330 Framework for IP Performance Metrics**
 - **Essential IOS Features Every ISP Should Consider v2.84**
 - **Best Practices for Catalyst 4000, 5000 and 6000 Series Switch Configuration and Management**
 - **Performance Management: Best Practices White Paper**
 - **Baseline Process: Best Practices White Paper**
 - **Technical Support – Cisco Systems**
-

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.