Use Best Practices for Network Time Protocol

Contents

- Introduction
- Prerequisites
  - Requirements
  - Components Used
- Background Information
- Terminology
- Overview
  - Device Overview
  - NTP Overview
- NTP Design Criteria
  - Association Modes
    - Client/Server Mode
    - Symmetric Active/Passive Mode
    - Broadcast and/or Multicast Mode
    - Set NTP Leap Second
  - 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
    - Create the NTP Design
    - Create a Seed File
    - Baseline NTP Performance Parameters
  - Iterative Tasks
    - Maintain the Seed File
    - Execute the NTP Node Scan
    - Review the NTP Node Reports
- Data Identification
  - General Data Characteristics
  - SNMP Data Identification
    - Cisco NTP MIB System Group
    - Cisco NTP MIB Peer Group - Peers Variable Table

Data Collection

SNMP Data Collection

Data Presentation

NTP Critical Node Report

NTP Interesting Node Report

NTP Configuration Report

Related Information

Introduction

This document describes the best practices for designing Network Time Protocol.

Prerequisites

Requirements

Cisco recommends that you have knowledge of this topic:

- Network Time Protocol
- Clock Technology and Public Time Servers

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

Internet Protocol (IP) based networks have quickly advanced 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 that are targeted at important metrics and measurement capabilities to characterize network behavior. The foundation of many metric methodologies is the measurement of time.

Network time synchronization, to the degree required for modern performance analysis, is an essential exercise. Based in the business models, and the services that are provided, the characterization of network performance is considered an important competitive service differentiator In these cases, great expense is incurred when you deploy network management systems and direct engineering resources to analyze the collected performance data. However, if proper attention is not given to the often-overlooked principle of time synchronization, those efforts are ineffective.

This document describes a hypothetical process definition for network management function management for the Network Time Protocol (NTP). You can use this article as a hypothetical procedure and an informational example. This can be customized by an organization to meet internal objectives.

The information provided by this document is presented in several major sections:

- The Terminology section provides general definitions of terms around 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 can 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 when a user manually logs into the system to issue the CLI command and record the output.
- The Data Presentation section provides report format examples of how the data can be presented.

Terminology

- **Accuracy**—The proximity of the clock absolute value to the offset of zero.
- **Accurate**—When a clock offset is zero at a particular moment in time.
- **Drift**—The measurement in the variation of skew, or the second derivation of the clock offset with respect to time.
- **Joint resolution**—When clocks are compared, it 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 time stamps generated by one clock subtracted 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 Tc and the true time is Tt, then the clock offset is Tc - Tt.
- **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 C1, when two clocks, C1 and C2, are compared. For example, clock C2's offset relative to C1 at a particular moment is Tc2 - Tc1, the instantaneous difference in time reported by C2 and C1.
- **Resolution**—The smallest unit by which a clock time is updated. Resolution is defined in terms of seconds. However, resolution is relative to the clock 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 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 enabled, the calendar is periodically updated from NTP, and this compensates 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 next.

- 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 noncompliant servers than an NTP client and must only be used in situations where strong authentication is not required.

The system clock provides time to the services listed next.

- 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 can be available only for display purposes and cannot

be redistributed.

NTP Overview

NTP is designed to synchronize the time on a network of machines. NTP runs over the User Datagram Protocol (UDP), with 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 primary servers can 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 dependent on the network conditions between the NTP server and the client. The other situation occurs when the router communicates to a bad NTP server (for example, NTP server with large dispersion); the router also increases the poll interval. 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 that runs NTP automatically chooses the machine with the lowest stratum number that it is configured to communicate with 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 next 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 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.

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

The communications between machines that run NTP (associations) are usually statically configured. Each machine is given the IP address of all machines with which it must form associations. Accurate timekeeping is made possible by NTP messages exchanged 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 recommended 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 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 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 by other means. Other machines then synchronize to that machine through NTP.

NTP Design Criteria

Each client in the synchronization subnet, which can 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 time must 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 agree, 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 must aim to provide each client with three or more sources of lower stratum time, several of these can only provide backup service and can 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 does not access directly, can also provide good backup service.

NTP generally prefers lower stratum servers to higher stratum servers unless the lower stratum server time is significantly different. The algorithm is able to detect when a time source is likely to be extremely inaccurate, or insane, and to prevent synchronization in these cases, even if the inaccurate clock is at a lower stratum level. And it can never synchronize a device to another server that is not synchronized itself.

In order to declare if server is reliable, it needs to pass many sanity check, such as:

- Implementations must include sanity timeouts which prevent trap transmissions if the monitoring program does not renew this information after a lengthy interval.
- Additional sanity checks are included for authentication, range bounds, and to avoid use of very old data.
- Checks have been added to warn that the oscillator has gone too long without update from a reference source.
- The peer.valid and sys.hold variables were added to avoid instabilities when the reference source changes rapidly due to large dispersive delays under conditions of severe network congestion. The peer.config, peer.authenticable, and peer.authentic bits were added to control special features and simplify configuration.

If at least one of those checks fail, the router declares it insane.

Association Modes

The next 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 with the server command and with the domain name server (DNS) name or address specified. 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 then adjust the local clock as needed. 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, and each operates 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.

Configuration of an association in client mode, is usually indicated by a server declaration in the configuration file and indicates that you want to obtain time from the remote server, but that one does not want 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. If one of the peers loses all reference sources or simply ceases operation, the other peers automatically reconfigure so that time values can flow from the current peers to all the others in the queue. In some contexts, this is described as a *push-pull* operation, in that the peer either pulls or pushes the time and values based on the particular configuration.

The configuration of 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 that involve 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 that operate as a mutually redundant group. In these modes, the servers in the group members arrange the synchronization paths for maximum performance, based on network jitter and propagation delay. If one or more of the group members fail, the residual members automatically reconfigure as required.

A peer is configured in symmetric active mode with the **peer** command and when the DNS name or address of the other peer is specified. 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 must 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, and this allows all operative 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 that involve one or a few servers and a potentially large client population. A broadcast server is configured with the **broadcast** command and a local subnet address. A broadcast client is configured with the **broadcastclient** command, which allows 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 must always be authenticated.

Set NTP Leap Second

You can use the **ntp leap {add | delete}** command in order to insert a leap second. There are options to add or delete leap seconds. There are two constraints for this to occur:

- Clock must be in sync state.
- The command is accepted only within the month before the leap is to happen. It cannot set leap if the current time is before 1 month of the occurrence of the leap.

After you set it, the leap second gets added or deleted to the last second as shown here:

```
<#root>

NTP leap second added :
Show clock given continuously
vl-7500-6#show clock
23:59:58.123 UTC Sun Dec 31 2006
vl-7500-6#show clock
23:59:58.619 UTC Sun Dec 31 2006
vl-7500-6#show clock
23:59:59.123 UTC Sun Dec 31 2006
vl-7500-6#show clock
23:59:59.627 UTC Sun Dec 31 2006

<< 59th second occurring twice

vl-7500-6#show clock
23:59:59.131 UTC Sun Dec 31 2006
vl-7500-6#show clock
23:59:59.627 UTC Sun Dec 31 2006
vl-7500-6#show clock
00:00:00.127 UTC Mon Jan 1 2007
vl-7500-6#show clock
```

```
00:00:00.623 UTC Mon Jan 1 2007
```

NTP Architecture

These three structures are available for NTP architecture:

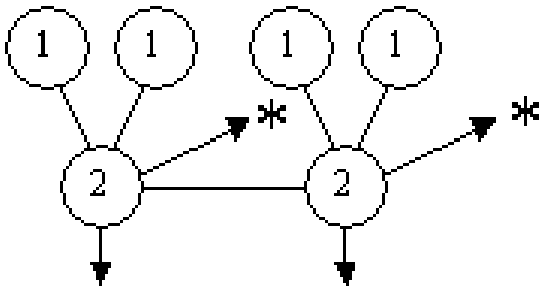
- Flat peer structure
- Hierarchical structure
- Star structure

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 user (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.

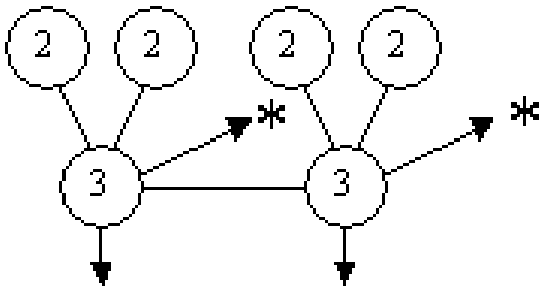
A scalable NTP architecture has a hierarchical structure as seen in the next diagram.

Internet Primary Servers
(Stratum 1)



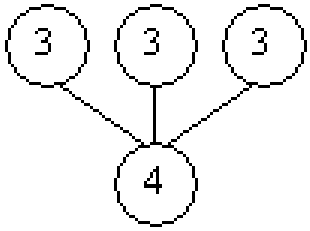
* = to buddy in another subnet

Campus Secondary Servers
(Stratum 2)

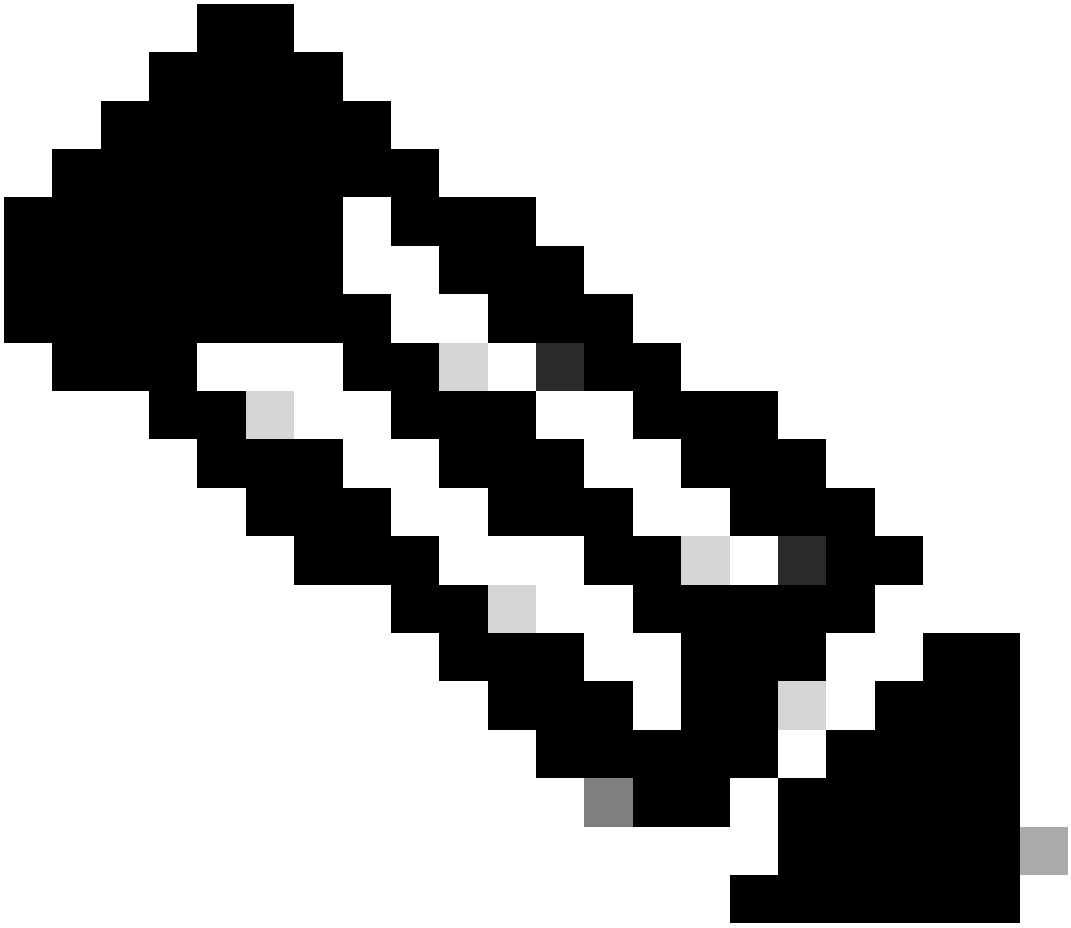


* = to buddy in another subnet

Department Servers
(Stratum 3)



Workstations (Stratum 4)



Note: A series of graphs which show a scalable and hierarchical NTP deployment. The first graph shows two NTP stratum 2 devices, each of which is connected to two stratum 1 devices (shown in the previous diagram of the stratum 2 devices) and a buddy in another subnet indicated by an asterisk. Additionally. each stratum 2 device has an arrow pointing down. The second graph has the same layout, but with stratum 2 devices where the stratum 1 devices were and stratum 3 devices where the stratum 2 devices were. The third graph has one stratum 4 device connected to three stratum 3 devices. In summary, the picture shows a topology where each device is connected to 2-3 devices with a stratum one lower (better) than its own.

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.

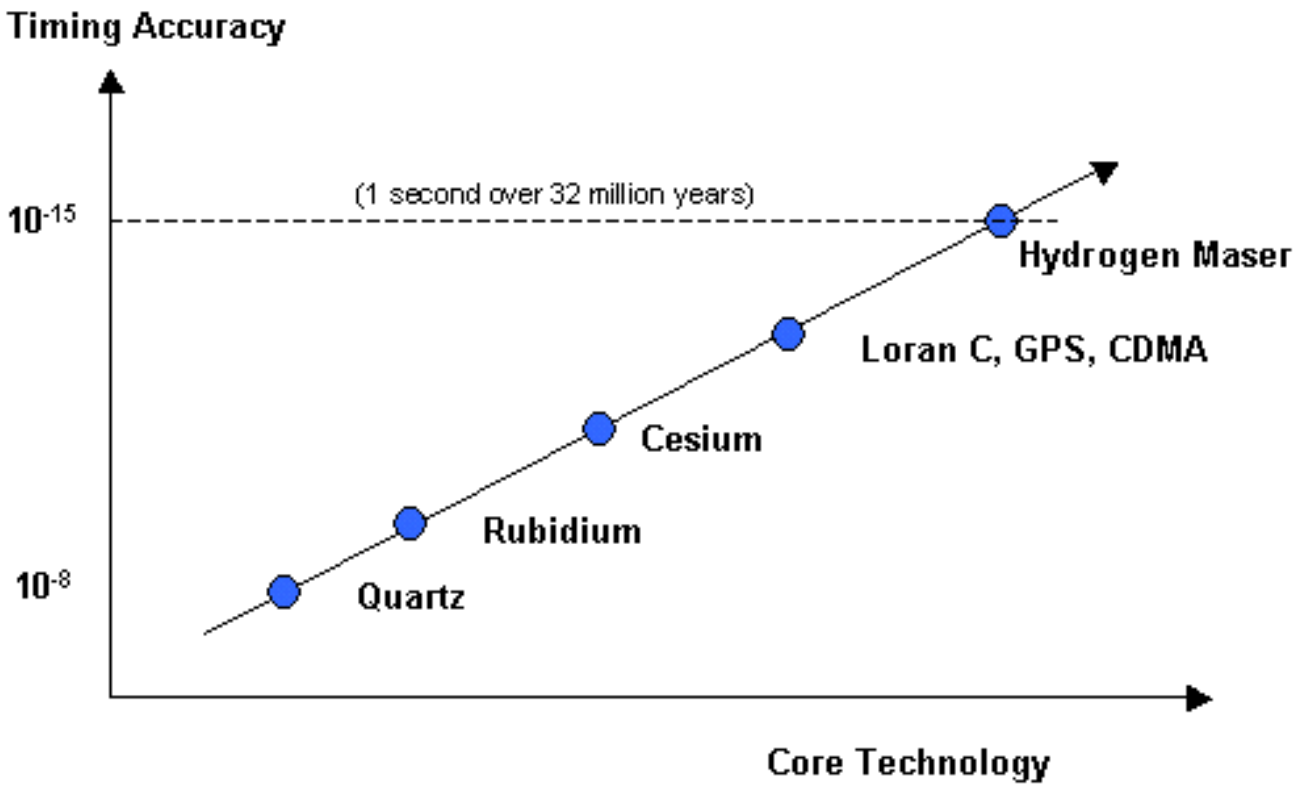
Clock Technology and Public Time Servers

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 and provide synchronization to a total in excess of 100,000 clients and servers on the

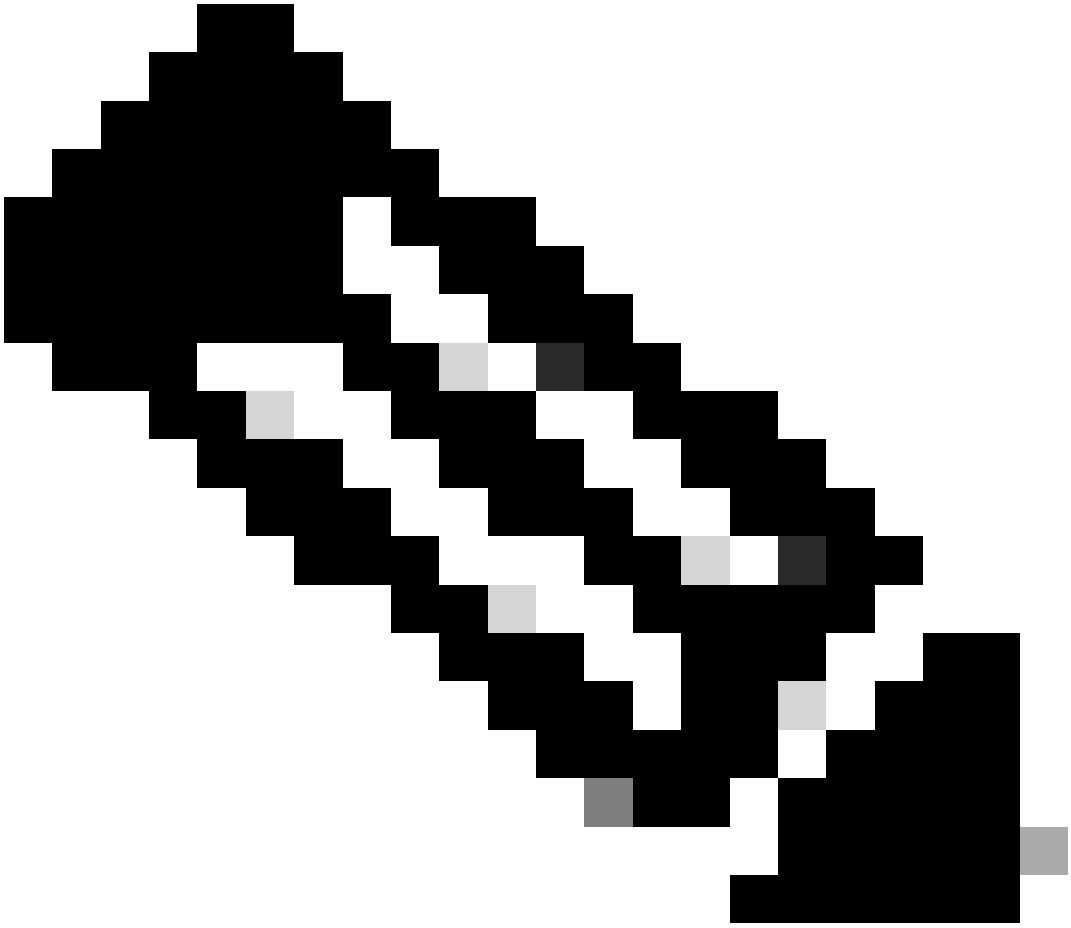
Internet. The Public NTP Time Servers lists are updated frequently. There are also numerous private primary and secondary servers not normally available to the public.

Note: PIX and ASA cannot be configured as an NTP server, but they can be configured as an NTP client.

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 can choose to deploy private external time sources. The next diagram shows a comparative graph of the relative accuracy of the current technologies.



Comparative Graph



Note: A graph which presents increasingly precise ways of keeping time from quartz (10 to the negative 8th power) to hydrogen maser (10 to the negative 15th power). The latter indicates approximately 1 second in accuracy loss over 32 million years. The other methods listed between these two (from least to most accurate) are Rubidium, Cesium, Loran C, GPS, and CDMA. The last three (Loran C, GPS, and CDMA) are listed together.

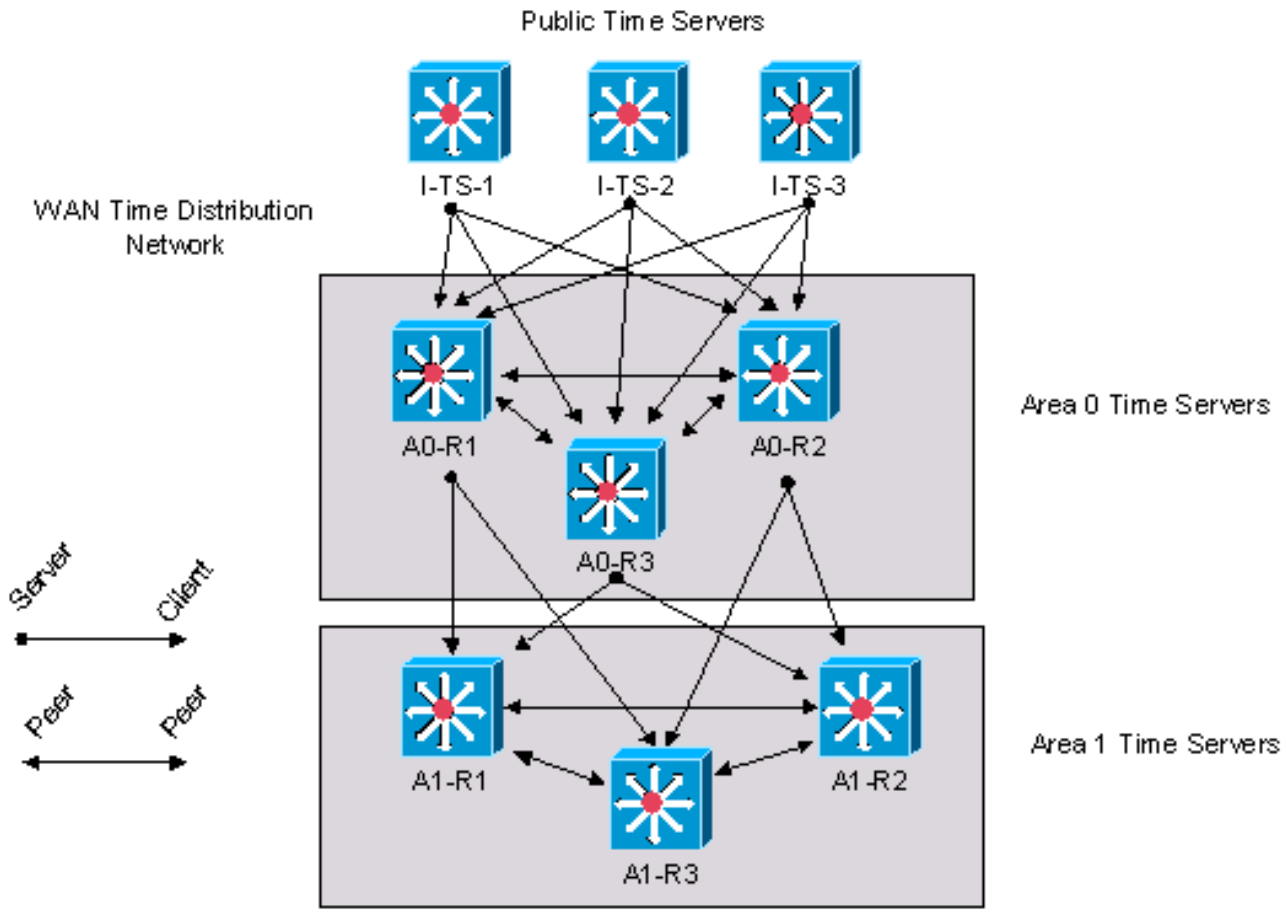
Until recently, the use of external time sources have not been widely deployed in enterprise networks due to the high cost of quality external time sources. However, as the Quality of Service (QoS) requirements increase and the cost of the time technology continues to decrease, external time sources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

Example NTP Deployments

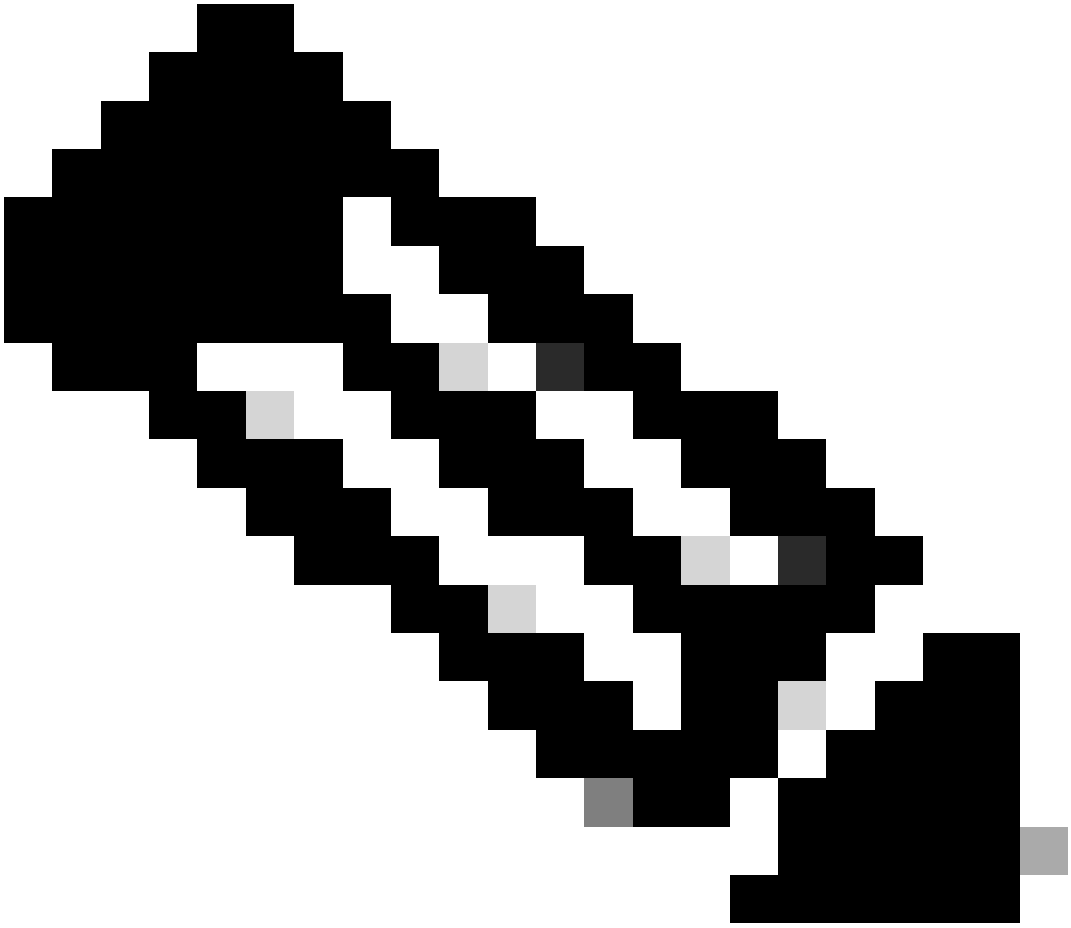
WAN Time Distribution Network

In the next diagram, 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 describes the Open Shortest Path First (OSPF) hierarchy. However, OSPF is not a prerequisite for NTP. It is only used as an illustrative example. NTP can be deployed along other logical hierarchical boundaries such

as an Enhanced Interior Gateway Routing Protocol (EIGRP) hierarchy or the standard Core/Distribution/Access hierarchy.



WAN Time Distribution Network



Note: A diagram which lays out an NTP topology that spans multiple networks. Three devices in Area 1 (OSPF) are peers of each other and clients of servers in Area 0. Three devices in Area 0 are peers of each other, clients of public time servers, and servers of clients in Area 1. The Public Time Servers are only shown as servers for clients in Area 0.

This example is the Cisco IOS configuration for device A0-R1 as shown in the previous 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 can be allowed to update the calendar
!--- and Cisco IOS can be configured to be an
!--- NTP master clock source.
```

```
!--- Cisco IOS can then obtain its clock from
!--- the hardware calendar.


clock calendar-valid


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


ntp update-calendar


!--- 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 can obtain the
!--- clock from the hardware calendar based on
!--- the previous line. This line can 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 <AO-R2>
access-list 5 permit <AO-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

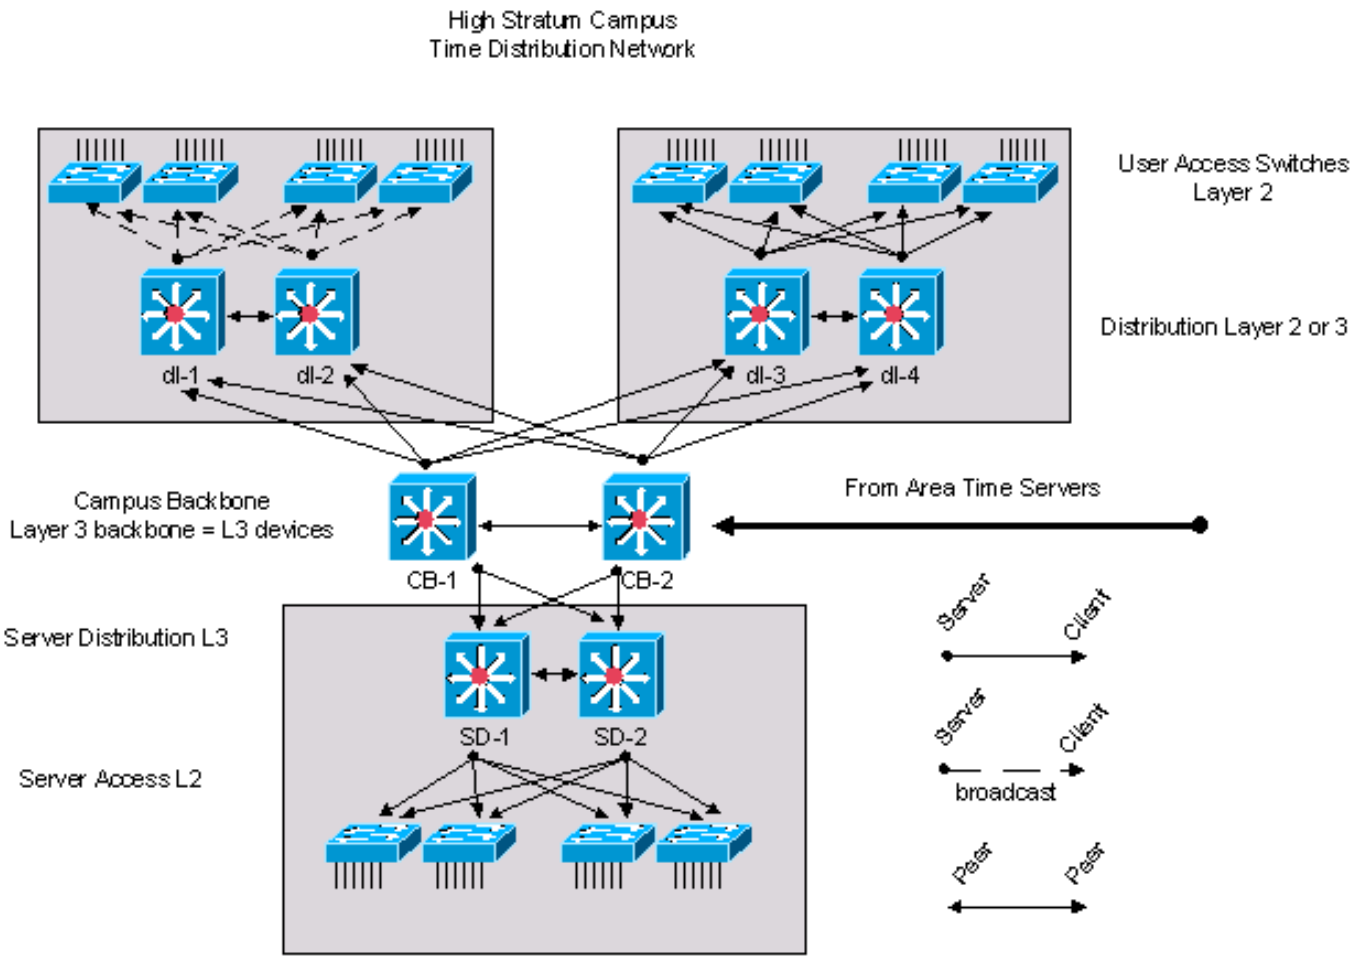

!--- Fault tolerant configuration polling for 3 NTP
!--- public servers, peering with 2 local servers.


ntp server <I-TS-1>
ntp server <I-TS-2>
ntp server <I-TS-3>
ntp peer <A0-R2>
ntp peer <A0-R3>
```

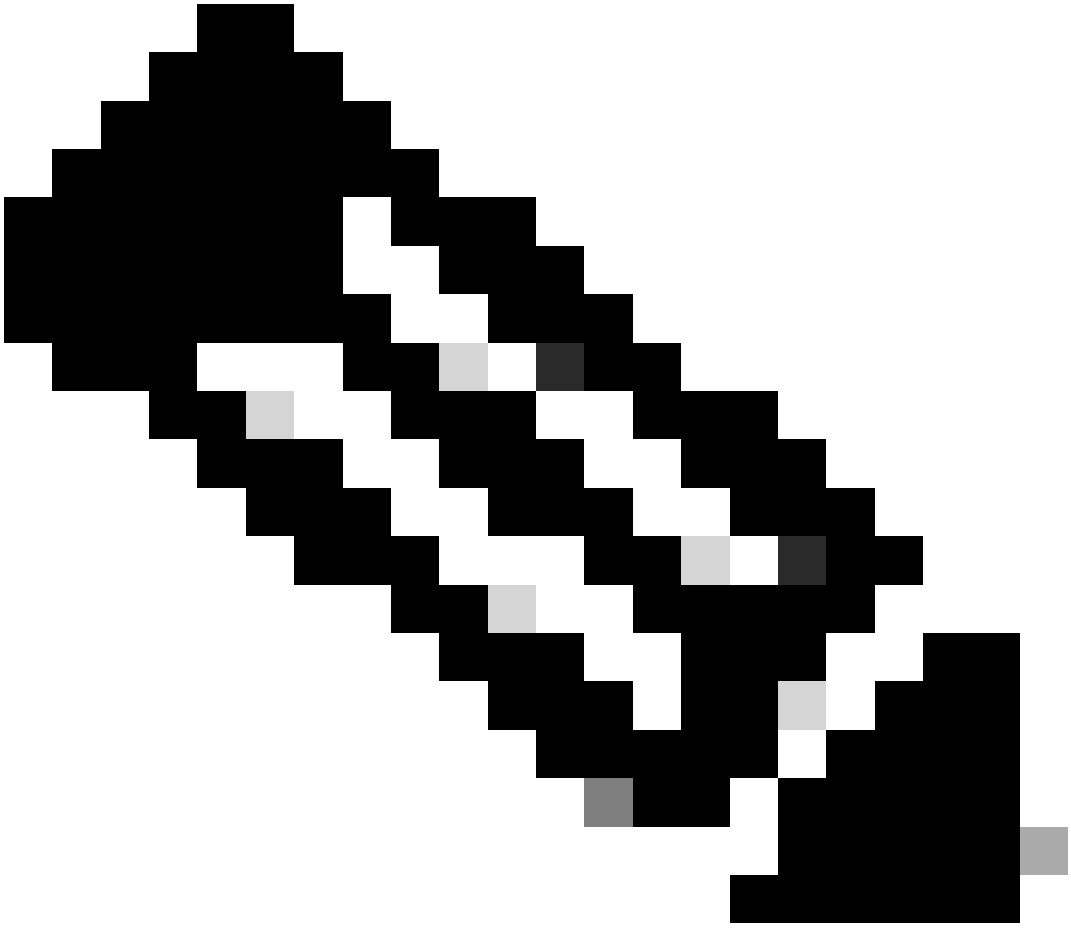
High Stratum Campus Time Distribution Network

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 considered primary difference for time distribution on a high stratum campus network is the potential use 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.



High Stratum Campus Time Distribution Network



Note: A diagram titled High Stratum Campus Time Distribution Network which includes a generic three-tier topology (backbone, distribution, access). Access switches are clients of distribution switches, distribution switches are clients of backbone switches, and backbone switches are clients of area time servers (not pictured). The distribution switches are divided into pairs, and have a peer relationship only with the other switch in the pair. The two backbone switches are also peers with each other. Four access switches (in the top left) are shown as broadcast clients with dotted arrows, while all other client-server and peer-peer relationships are non-broadcast.

The high stratum campus network, shown in the previous diagram, 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 other 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 rest of the devices are Layer 2 switches. In this client access block, the time is distributed with 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.

This is the configuration for the dl-1 Layer 3 distribution devices:

```
!--- In this case, dl-1 can 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

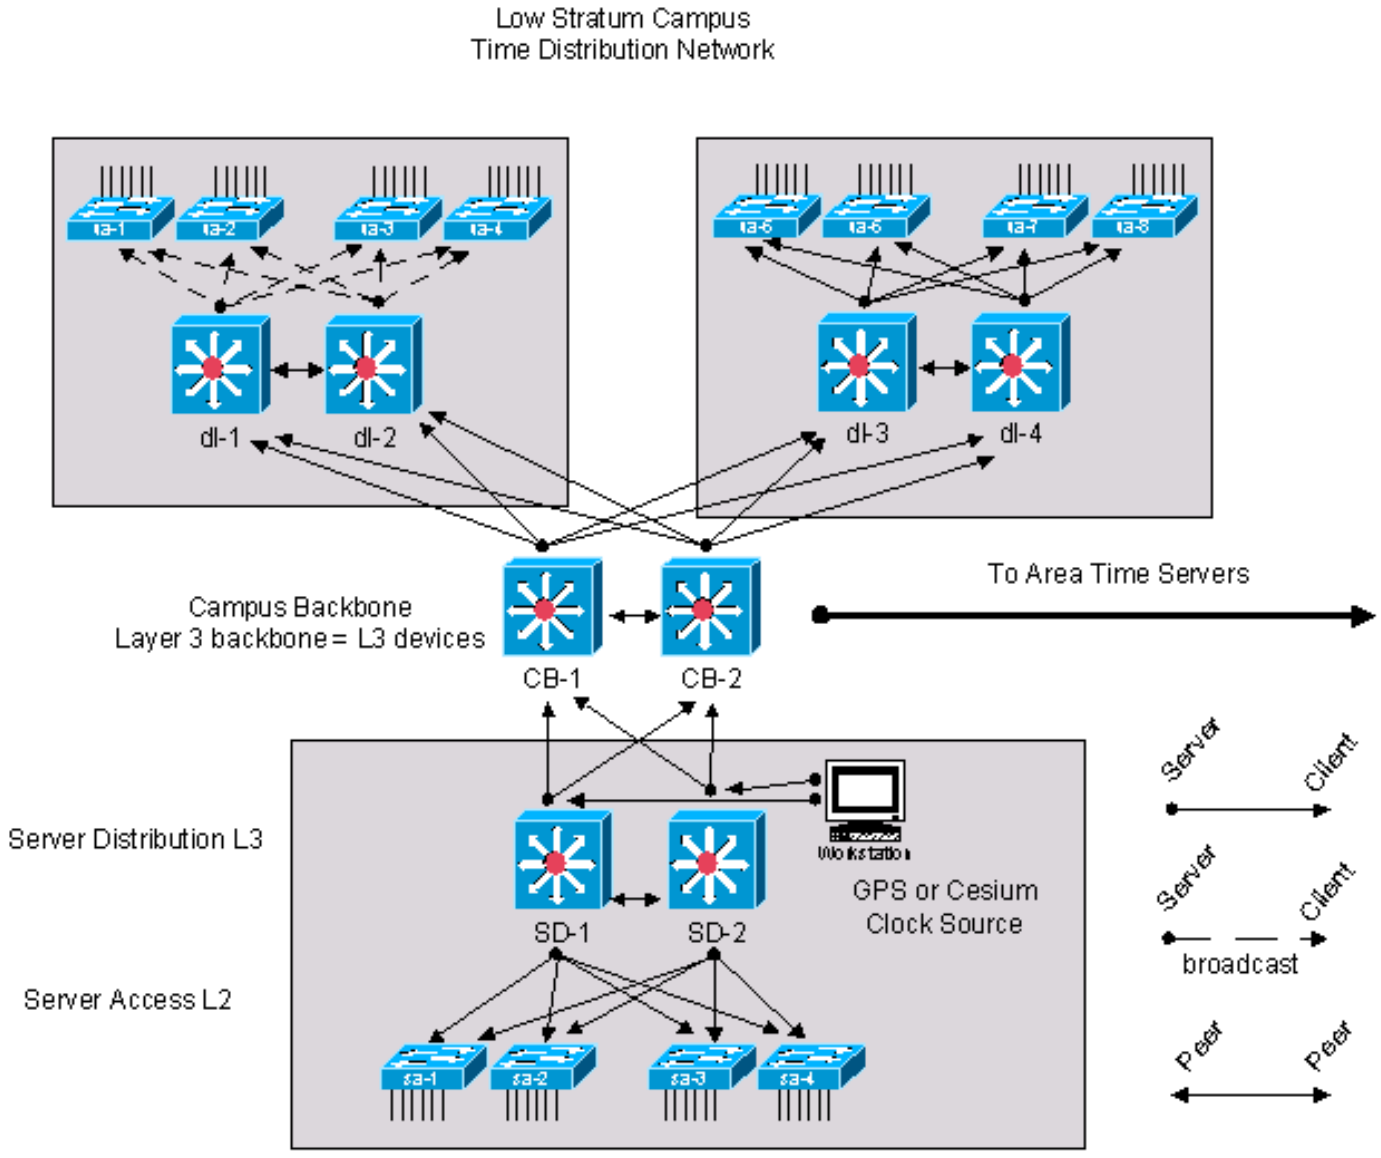

!--- Fault tolerant configuration polling 2
!--- local time servers and 1 local peer.


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

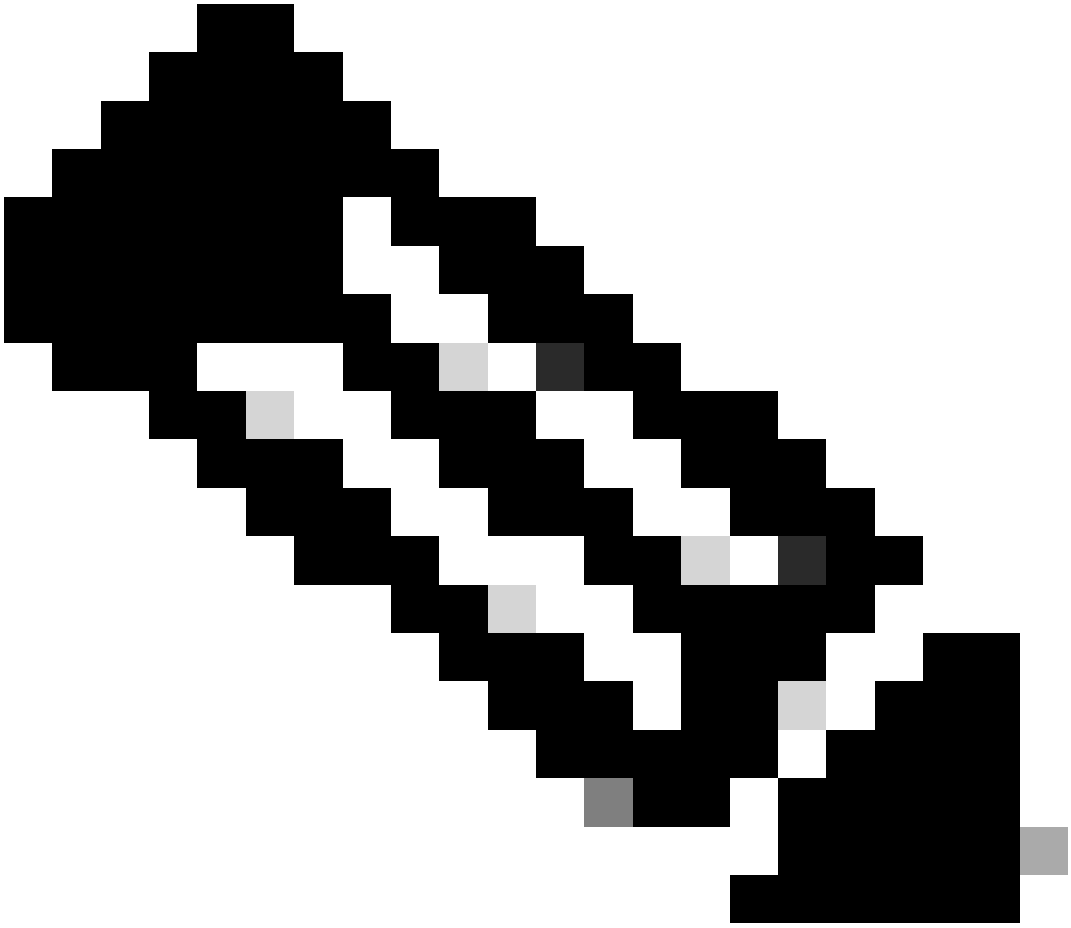
Low Stratum Campus Time Distribution Network

In the diagram next, 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 must 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 with the principles of hierarchy and modularity that have been described in the previous sections.



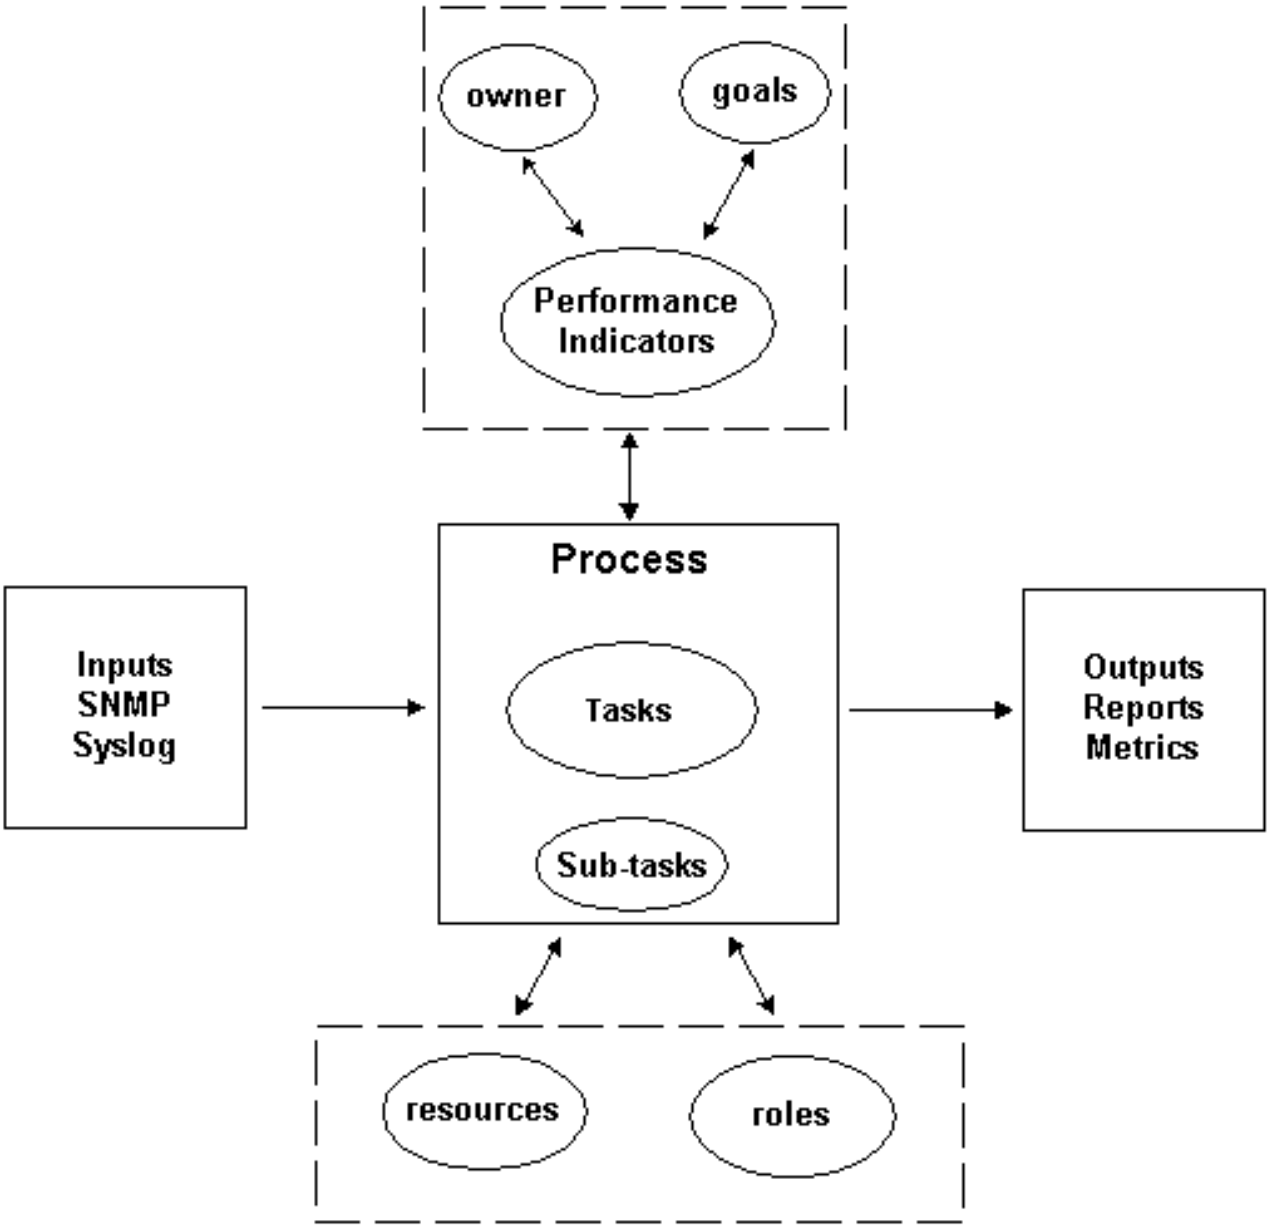
Low Stratum Campus Time Distribution Network



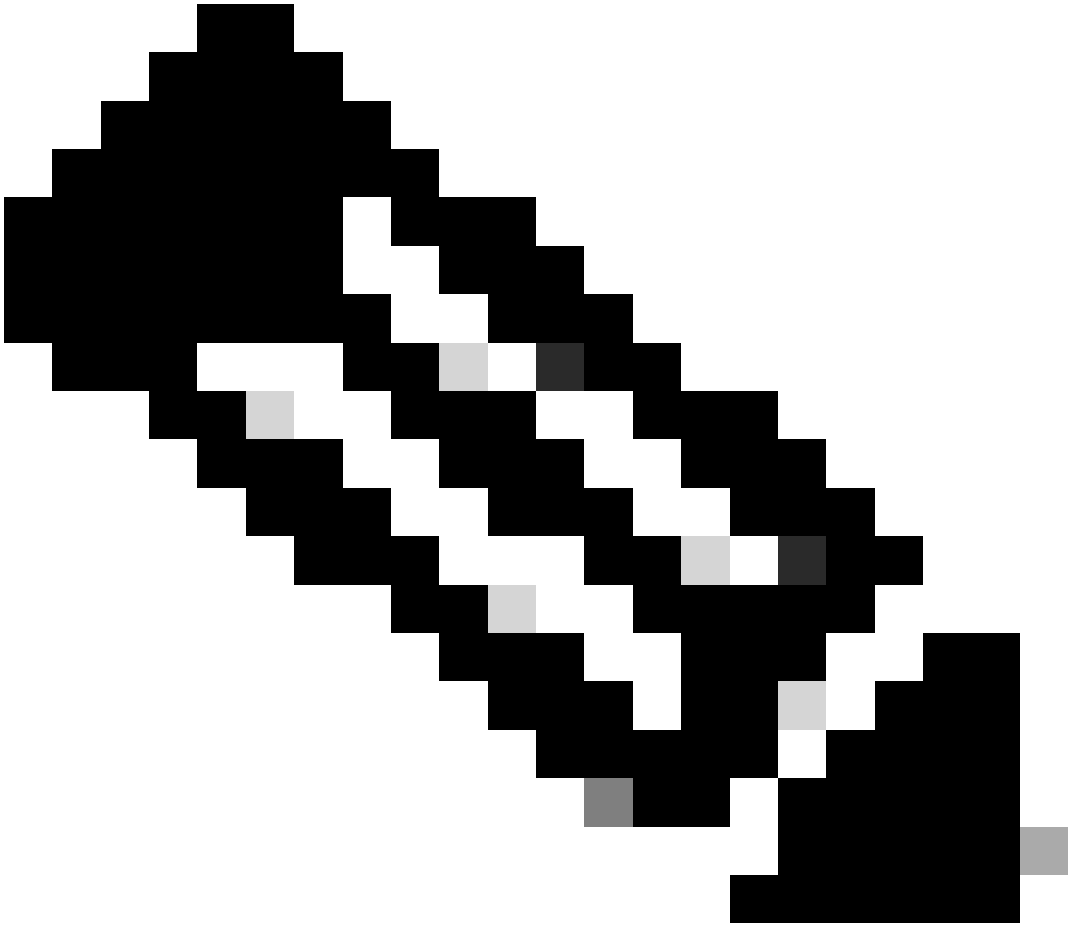
Note: A diagram titled Low Stratum Campus Time Distribution Network which includes a generic three-tier topology (backbone, distribution, access). Two distribution switches have GPS or cesium clocks attached. The access switches directly attached to these distribution switches and the backbone switches are clients of these distribution switches. All other distribution switches on the network are clients of the backbone switches, and the remaining access switches are also clients of their directly-attached distribution switches. Four access switches (in the top left) are shown as broadcast clients with dotted arrows, while all other client-server and peer-peer relationships are non-broadcast.

Process Definitions

A process definition is a connected series of actions, activities, and changes performed by agents that intend to satisfy a purpose or achieve a goal. Process control is the process of planning and regulating, with the objective to perform a process in an effective and efficient way. This is shown in the next diagram.



Series of Processes



Note: A diagram which specifies the meaning of a process used by this document. There are five regions. The left region has a solid border. It contains Inputs, SNMP, and Syslog. There is a one-way arrow from the left region to the center region. The right region also has a solid border. It contains Outputs, Reports, and Metric. There is a one-way arrow from the center region to the right region. The top region has a dotted border. It contains owner, goals, and Performance Indicators. There are circles with solid borders around all three. There are two-way arrows between (a) owner and Performance Indicators (b) goals and Performance Indicators and (c) the top region and the center region. The bottom region also has a dotted border. It contains resources and roles. There are circles with solid borders around both. There are two-way arrows which appear to connect resources and roles with the center region, but they stop at the border of the bottom region. The center region has a solid border and a heading which reads Process. It also contains one each of Tasks and Sub-tasks. Each has a solid circular border. Tasks has more blank space inside the circle than any other item in the graph.

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 that determines and reports 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 to ensure 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 must be measurable and quantifiable. For instance, the performance indicators listed next 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 next.

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

Process Outputs

The process outputs are defined as:

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

Task Definitions

The next 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 must not be executed during each iteration of the process.

Create the NTP Design

When you verify 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 must be documented by the process owner and submitted to management. The next table 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	• Design technical and economic requirements • Current network design documentation • Criteria that defines 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 next table outlines how to create a seed file.

Prerequisite Task	Description
Task objectives	Create seed file that identifies three categories of network devices: 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 can be used to identify and prioritize the nodes involved in the NTP architecture.

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.

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 that collect SNMP data and calculate baselines.

Task role	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 that concern 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 cannot 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 must 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 next table describes data that is considered significant when you analyze 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 about 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 that runs NTP. The target node is the destination of the SNMP queries.

Object Name	Object Description
cntpSysStratum	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. ::= { cntpSystem 2 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.2
cntpSysPrecision	Signed integer that indicates 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). ::= { cntpSystem 3 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.3
cntpSysRootDelay	A signed fixed-point number that indicates the total round-trip delay in seconds, to the primary reference source at the root of the synchronization subnet. ::= { cntpSystem 4 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.4
cntpSysRootDispersion	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. ::= { cntpSystem 5 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.4
cntpSysRefTime	The local time when the local clock was last updated. If the local clock has never been synchronized, the value is zero. ::= { cntpSystem 7 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.7
cntpSysPeer	The current synchronization source that contains 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. ::= { cntpSystem 9 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.9
cntpSysClock	The current local time. Local time is derived from the hardware clock of the particular machine and increments at intervals based on the design used. ::= {

	cntpSystem 10 } object identifier = .1.3.6.1.4.1.9.9.168.1.1.10

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	This table provides information on the peers with which the local NTP server has associations. The peers are also NTP servers that run on different hosts. This is a table of cntpPeersVarEntry ::= { cntpPeers 1 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1
cntpPeersVarEntry	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 sets the cntpPeersEntryStatus as active (1). At the very least, the management station has to set a value for cntpPeersPeerAddress to make the row active. INDEX { cntpPeersAssocId } ::= { cntpPeersVarTable 1 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1
cntpPeersAssocId	An integer value greater than zero that uniquely identifies a peer with which the local NTP server is associated. ::= { cntpPeersVarEntry 1 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.1
cntpPeersConfigured	This is a bit that indicates that the association was created from configuration information and must not be disassociated even if the peer becomes unreachable. ::= { cntpPeersVarEntry 2 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.2
cntpPeersPeerAddress	The IP address of the peer. When a new association is created, a value for this object must be set before the row is made active. ::= { cntpPeersVarEntry 3 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.3
cntpPeersMode	SYNTAX INTEGER { unspecified (0), symmetricActive (1), symmetricPassive (2), client (3), server (4), broadcast (5), reservedControl (6), reservedPrivate (7) } When a new peer association is created, if no value is specified for this object, it defaults to symmetricActive (1). ::= { cntpPeersVarEntry 8 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.8
cntpPeersStratum	The stratum of the peer clock. ::= { cntpPeersVarEntry 9 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.9
cntpPeersRootDelay	A signed fixed-point number that indicates the total round-trip delay in seconds, from the peer to the primary reference source at the root of the synchronization subnet. ::= { cntpPeersVarEntry 13 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.13
cntpPeersRootDispersion	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. ::= { cntpPeersVarEntry 14 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.14
cntpPeersRefTime	The local time at the peer when its clock was last updated. If the peer clock has never been synchronized, the value is zero. ::= { cntpPeersVarEntry 16 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.16
cntpPeersReach	A shift register used to determine the reachability status of the peer, with bits that enter 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. ::= { cntpPeersVarEntry 21 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.21

cntpPeersOffset	The estimated offset of the peer clock relative to the local clock, in seconds. The host determines the value of this object that uses the NTP clock-filter algorithm. ::= { cntpPeersVarEntry 23 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.21
cntpPeersDelay	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 that uses the NTP clock-filter algorithm. ::= { cntpPeersVarEntry 24 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.24
cntpPeersDispersion	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 that uses the NTP clock-filter algorithm. ::= { cntpPeersVarEntry 25 } object identifier = .1.3.6.1.4.1.9.9.168.1.2.1.1.25

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 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 that is monitored and one-way-delay-variation metrics are 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 next an example OSPF hierarchy. Therefore, the reports described next are formatted to group the NTP devices by 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 can be listed for report purposes. As mentioned earlier, OSPF is not a prerequisite for NTP. It is only used in this document 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 cannot 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 Cisco IOS Features Every ISP must Consider v2.84**
- **Cisco Technical Support & Downloads**