Network Time Protocol (NTP)

NTP Overview

NTP synchronizes timekeeping among a set of distributed time servers and clients. This synchronization allows events to be correlated when system logs are created and other time-specific events occur.

NTP uses the User Datagram Protocol (UDP) as its transport protocol. All NTP communication uses Coordinated Universal Time (UTC). An NTP network usually receives its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP distributes this time across the network. NTP is efficient; no more than one packet per minute is necessary to synchronize the two machines to within a millisecond of each other.

NTP uses the concept of a "stratum" to describe how many NTP "hops" away a machine is from an authoritative time source. A "stratum 1" time server typically has an authoritative time source (such as a radio or atomic clock, or a GPS time source) directly attached, a "stratum 2" time server receives its time via NTP from a "stratum 1" time server, and so on.

NTP avoids synchronizing to a machine whose time may not be accurate, in two ways. First, NTP never synchronizes to a machine that isn't synchronized itself. Second, NTP compares the time reported by several machines and doesn't synchronize to a machine whose time is significantly different than the others, even if its stratum is lower. This strategy effectively builds a self-organizing tree of NTP servers.

The Cisco implementation of NTP doesn't support stratum 1 service; in other words, it's not possible to connect to a radio or atomic clock (for some specific platforms, however, you can connect a GPS time-source device). We recommend that time service for your network be derived from the public NTP servers available in the IP Internet.

If the network is isolated from the Internet, the Cisco implementation of NTP allows a machine to be configured so that it acts as though it's synchronized via NTP, when in fact it has determined the time using other means. Other machines can then synchronize to that machine via NTP.

Several manufacturers include NTP software for their host systems, and a publicly available version for systems running UNIX and its various derivatives is also available. This software also allows UNIX-derivative servers to acquire the time directly from an atomic clock, which would then propagate time information along to Cisco routers.

The communications between machines running NTP (known as *associations*) are 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.

The Cisco implementation of NTP supports two ways that a networking device can obtain NTP time information on a network:

- By polling host servers

- By listening to NTP broadcasts

In a LAN environment, NTP can be configured to use IP broadcast messages. As compared to polling, IP broadcast messages reduce configuration complexity, because each machine can simply be configured to send or receive broadcast or multicast messages. However, the accuracy of timekeeping is marginally reduced because the information flow is one-way only.

An NTP broadcast client listens for broadcast messages sent by an NTP broadcast server at a designated IPv4 address. The client synchronizes the local clock using the first received broadcast message.

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

When multiple sources of time (VINES, hardware clock, manual configuration) are available, NTP is always considered to be more authoritative. NTP time overrides the time set by any other method.

Preventing Issues due to GPS Week Number Rollover (WNRO)

- If there are no GPS sources in the NTP source chain or server chain, there's no impact of GPS Week Number Rollover (WNRO).

- GPS WNRO affects only the system clock and not user traffic.

- Contact your GPS manufacturer to fix the GPS source for this condition.

To mitigate the impact of GPS sources that are subject to GPS WNRO perform the following optional workarounds:

- If the GPS source has been identified to be a cause of potential disruption on April 6, 2019 (or after), configure ntp master in the Cisco that is device connected to this source, and its clock on the Stratum 1 device to isolate it. This configuration enables the device to present its own clock for synchronization to downstream NTP clients.

Note The usage of ntp master command as mentioned above is only a workaround to this condition. Use this command until the GPS source-related conditions are resolved, and to prevent the distribution of incorrect clock values throughout the network.

- Configure multiple NTP servers (ideally 4, but more than 3) at the Stratum 2 level of the network, to enable NTP clients at Stratum 2 level to get clock from more than one Stratum 1 server. This way, WNRO affected Stratum 1 servers are staged to be marked as 'false ticker' or 'outlier' clock sources as compared to other non-WNRO affected Stratum 1 servers.

Configure NTP

Choose the Method to Obtain NTP Time Information

Configuring Poll-Based Associations

Note No specific command enables NTP; the first NTP configuration command that you issue enables NTP.

You can configure the following types of poll-based associations between the router and other devices (which may also be routers):

- Client mode
- Symmetric active mode

The client and the symmetric active modes should be used when NTP is required to provide a high level of time accuracy and reliability.

When a networking device is operating in the client mode, it polls its assigned time serving hosts for the current time. The networking device then picks a host from all the polled time servers to synchronize with. Because the relationship that is established in this case is a client-host relationship, the host doesn't capture or use any time information sent by the local client device. This mode is most suited for file-server and workstation clients that aren't required to provide any form of time synchronization to other local clients. Use the **server** command to individually specify the time-serving hosts that you want your networking device to consider synchronizing with and to set your networking device to operate in the client mode.

When a networking device is operating in the symmetric active mode, it polls its assigned time-serving hosts for the current time and it responds to polls by its hosts. Because this is a peer-to-peer relationship, the host also retains time-related information about the local networking device that it's communicating with. This mode should be used when there are several mutually redundant servers that are interconnected via diverse network paths. Most stratum 1 and stratum 2 servers on the Internet today adopt this form of network setup. Use the **peer** command to individually specify the time-serving hosts that you want your networking device to consider synchronizing with and to set your networking device to operate in the symmetric active mode.

When the router polls several other devices for the time, the router selects one device with which to synchronize.

Note To configure a peer-to-peer association between the router and another device, you must also configure the router as a peer on the other device.

You can configure multiple peers and servers, but you can't configure a single IP address as both a peer and a server at the same time.

To change the configuration of a specific IP address from peer to server or from server to peer, use the **no** form of the **peer** or **server** command to remove the current configuration before you perform the new configuration. If you don't remove the old configuration before performing the new configuration, the new configuration doesn't overwrite the old configuration.

Step 1 Form a server association with another system.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# server 172.19.69.1 minpoll 8 maxpoll 12
```

This step can be repeated as necessary to form associations with multiple devices.

Step 2 Form a peer association with another system.

```
Router(config-ntp)# peer 192.168.22.33 minpoll 8 maxpoll 12
source hundredGigE 0/0/0/1
Router(config-ntp)# end
```

This step can be repeated as necessary to form associations with multiple systems.

Note To complete the configuration of a peer-to-peer association between the router and the remote device, the router must also be configured as a peer on the remote device.

Step 3 Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 server 172.19.69.1 minpoll 8 maxpoll 12
 peer 192.168.22.33 minpoll 8 maxpoll 12 source HundredGigE0/0/0/1
!
```

Configuring Broadcast-Based NTP Associations

In a broadcast-based NTP association, an NTP server propagates NTP broadcast packets throughout a network. Broadcast clients listen for the NTP broadcast packets propagated by the NTP server and don't engage in any polling.

Broadcast-based NTP associations should be used when time accuracy and reliability requirements are modest and if your network is localized and has many clients (more than 20). Broadcast-based NTP associations are also recommended for use on networks that have limited bandwidth, system memory, or CPU resources. Time accuracy is marginally reduced in broadcast-based NTP associations because information flows only one way.

Use the **broadcast client** command to set your networking device to listen for NTP broadcast packets propagated through a network. For broadcast client mode to work, the broadcast server and its clients must be located on the same subnet. The time server that is transmitting NTP broadcast packets must be enabled on the interface of the given device using the **broadcast** command.

Use the **broadcast** command to set your networking device to send NTP broadcast packets.

Note No specific command enables NTP; the first NTP configuration command that you issue enables NTP.

Step 1 Configure the specified interface to receive NTP broadcast packets.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# broadcastdelay 2
```

```
Router(config-ntp)# interface HundredGigE 0/2/0/0
Router(config-ntp-int)# broadcast client
```

Note   Go to the next step to configure the interface to send NTP broadcast packets.

Step 2   Configure the specified interface to send NTP broadcast packets.

```
Router(config-ntp-int)# broadcast destination 10.50.32.149
Router(config-ntp-int)# end
```

Note   Go to the previous step to configure the interface to receive NTP broadcast packets.

Step 3   Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 interface HundredGigE0/2/0/0
  broadcast client
  broadcast destination 10.50.32.149
 !
 broadcastdelay 2
!
```

Configuring NTP Access Groups

Note   No specific command enables NTP; the first NTP configuration command that you issue enables NTP.

The access list-based restriction scheme allows you to grant or deny certain access privileges to an entire network, a subnet within a network, or a host within a subnet.

The access group options are scanned in the following order, from least restrictive to most restrictive:

1. **peer**—Allows time requests and NTP control queries and allows the system to synchronize itself to a system whose address passes the access list criteria.

2. **serve**—Allows time requests and NTP control queries, but doesn't allow the system to synchronize itself to a system whose address passes the access list criteria.

3. **serve-only**—Allows only time requests from a system whose address passes the access list criteria.

4. **query-only**—Allows only NTP control queries from a system whose address passes the access list criteria.

If the source IP address matches the access lists for more than one access type, the first type is granted. If no access groups are specified, all access types are granted to all systems. If any access groups are specified, only the specified access types are granted.

For details on NTP control queries, see RFC 1305 (NTP version 3).

Step 1   Create an access group and apply a basic IPv4 or IPv6 access list to it.

```
Router# configure
Router(config)# ntp
```

```
Router(config-ntp)# access-group peer peer-acl
Router(config-ntp)# end
```

Step 2   Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 access-group ipv4 peer peer-acl
 broadcastdelay 2
!
```

Configuring NTP Authentication

This task explains how to configure NTP authentication.

The encrypted NTP authentication scheme should be used when a reliable form of access control is required. Unlike the access-list-based restriction scheme that is based on IP addresses, the encrypted authentication scheme uses authentication keys and an authentication process to determine if NTP synchronization packets sent by designated peers or servers on a local network are deemed as trusted, before the time information that it carries along is accepted.

The authentication process begins from the moment that an NTP packet is created. A message authentication code (MAC) is computed using the MD5 Message Digest Algorithm and the MAC is embedded into an NTP synchronization packet. The NTP synchronization packet together with the embedded MAC and key number are transmitted to the receiving client. If authentication is enabled and the key is trusted, the receiving client computes the MAC in the same way. If the computed MAC matches the embedded MAC, the system is allowed to sync to the server that uses this key in its packets.

After NTP authentication is properly configured, your networking device only synchronizes with and provides synchronization to trusted time sources.

Step 1   Define the authentication keys.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# authenticate
Router(config-ntp)# authentication-key 3 md5 clear key1
```

Each key has a key number, a type, a value, and, optionally, a name. Currently the only key type supported is **md5**.

Step 2   Define trusted authentication keys.

```
Router(config-ntp)# trusted-key 3
Router(config-ntp)# commit
```

If a key is trusted, this router only synchronizes to a system that uses this key in its NTP packets.

Step 3   Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 authentication-key 3 md5 encrypted 020D01425A
 authenticate
```

```
trusted-key 3
!
```

Configuring the Source IP Address for NTP Packets

By default, the source IP address of an NTP packet sent by the router is the address of the interface through which the NTP packet is sent. Use this procedure to set a different source address.

Step 1 Configure an interface from which the IP source address.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# source HundredGigE 0/0/0/1
Router(config-ntp)# end
```

Note This interface is used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the **source** keyword in the **peer** or **server** command shown in .

Step 2 Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 authentication-key 3 md5 encrypted 020D01425A
 authenticate
 trusted-key 3
 source HundredGigE0/0/0/1
!
```

Configuring the System as an Authoritative NTP Server

You can configure the router to act as an authoritative NTP server, even if the system isn't synchronized to an outside time source.

Step 1 Make the router an authoritative NTP server.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# master 9
Router(config-ntp)# end
```

Note Use the **master** command with caution. It's easy to override valid time sources using this command, especially if a low stratum number is configured. Configuring multiple machines in the same network with the **master** command can cause instability in time keeping if the machines don't agree on the time.

Step 2 Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 master 9
```

Updating the Hardware Clock

On devices that have hardware clocks (system calendars), you can configure the hardware clock to be periodically updated from the software clock. This is advisable for devices using NTP, because the time and date on the software clock (set using NTP) is more accurate than the hardware clock. The time setting on the hardware clock has the potential to drift slightly over time.

Step 1   Configure the router to update its system calendar from the software clock at periodic intervals.

```
Router# configure
Router(config)# ntp
Router(config-ntp)# update-calendar
Router(config-ntp)# end
```

Step 2   Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 update-calendar
```

Verifying the Status of the External Reference Clock

This task explains how to verify the status of NTP components.

Note   The commands can be entered in any order.

Step 1   Display the status of NTP associations.

```
Router# show ntp associations
address         ref clock     st  when  poll reach  delay  offset    disp
 ~172.19.69.1      .AUTH.            16     -  1024    0    0.00   0.000   15937
 ~192.168.22.33    .AUTH.            16     -  1024    0    0.00   0.000   15937
*~127.127.1.1      .LOCL.             9    51    64   37    0.00   0.000  438.28
 * sys_peer, # selected, + candidate, - outlayer, x falseticker, ~ configured
```

Step 2   Display the status of NTP.

```
Router# show ntp status
Clock is synchronized, stratum 10, reference is 127.127.1.1
nominal freq is 1000000000.0000 Hz, actual freq is 1000000000.0000 Hz, precision is 2**24
reference time is E8CE945C.8E2A8B07 (15:01:48.555 UTC Mon Oct  9)
clock offset is 0.000 msec, root delay is 0.000 msec
root dispersion is 63.52 msec, peer dispersion is 63.40 msec
loopfilter state is 'FREQ' (Drift being measured), drift is 0.0000000000 s/s
system poll interval is 64, last update was 9 sec ago
```

```
authenticate is enabled, panic handling is disabled,
hostname resolution retry interval is 1440 minutes.
```

Disabling NTP Services on a Specific Interface

NTP services are disabled on all interfaces by default.

NTP is enabled globally when any NTP commands are entered. You can selectively prevent NTP packets from being received through a specific interface by turning off NTP on a given interface.

Step 1 Disable NTP services on the specified interface using one of the following commands:

- **interface**
- **no interface**

```
Router# configure
Router(config)# ntp
Router(config-ntp)# interface HundredGigE 0/0/0/1 disable
Router(config-ntp)# end
```

or

```
Router# configure
Router(config)# ntp
Router(config-ntp)# no interface HundredGigE 0/0/0/1
Router(config-ntp)# end
```

Step 2 Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 interface HundredGigE0/0/0/1
  disable
 !
```

FQDN for NTP Server

NTP on Cisco IOS XR Software supports configuration of servers and peers using their Fully Qualified Domain Names (FQDN). While configuring, the FQDN is resolved via DNS into its corresponding IPv4 or IPv6 address and is stored in the running-configuration of the system. NTP supports FQDN for both IPv4 and IPv6 protocols. You can configure FQDN on default VRF.

Starting Cisco IOS XR Software Release 7.9.1 you can configure FQDN in nondefault VRF also.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
FQDN for NTP Server on Nondefault VRF	Release 7.9.1	You can now specify a Fully Qualified Domain Name (FQDN) as the hostname for NTP server configuration over nondefault VRFs. FQDNs are easy to remember compared to numeric IP addresses. Service migration from one host to another can cause a change in IP address leading to outages. Prior releases allowed FQDN handling in only default VRFs.

Configure FQDN for NTP server

Configuring FQDN on NTP Server on Default VRF

Step 1 Use the **ntp server** command with the FQDN name to configure FQDN on default VRF. You don't need to specify the VRF name.

In the following example, *time.cisco.com* is the FQDN.

```
Router# configure
Router(config)# ntp server time.cisco.com
Router(config)# commit
```

Step 2 Verify the configured NTP profile details.

```
Router# show running-config ntp
ntp
 server 192.0.2.1
!
```

Step 3 Verify that an NTP association has come up using the **show ntp associations** command.

```
Router# show ntp associations

      address         ref clock      st  when  poll reach  delay  offset    disp
~192.0.2.1        173.38.201.67   2    42   128    3  196.06  -14.25  3949.4
* sys_peer, # selected, + candidate, - outlayer, x falseticker, ~ configured
```

Configuring FQDN on NTP Server on Nondefault VRF

Before you begin

- Configuration must exist for DNS resolution over that specific VRF.

• The server must be reachable.

Step 1 FQDN must be reachable from the router to configure it as an NTP server or peer. You can use the **ping** command and verify that FQDN is reachable.

In the following example, *time.cisco.com* is the FQDN and *vrf_1* is the VRF over which it's reachable.

```
Router# ping time.cisco.com vrf vrf_1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.0.2.1 timeout is 2 seconds:
```

Step 2 When you have confirmed that FQDN is reachable, you can configure FQDN to be used as an NTP server/peer.

```
Router# configure
Router(config)# ntp server vrf vrf_1 time.cisco.com minpoll 4 maxpoll 4 iburst
Router(config)# commit
```

Note If the FQDN you're trying to configure isn't reachable, the CLI treats it as invalid input.

Step 3 Verify the configured NTP profile details.

```
Router# show running-config ntp
 ntp
  server vrf vrf_1 192.0.2.1 minpoll 4 maxpoll 4 iburst
 !
```

Step 4 Verify that an NTP association has come up using the **show ntp associations** command.

```
Router# show ntp associations
 address          ref clock     st  when  poll reach  delay  offset    disp
~192.0.2.1 vrf vrf_1
                  173.38.201.115  2    14    16    37  179.10  13.492  16.680
* sys_peer, # selected, + candidate, - outlayer, x falseticker, ~ configured
```

NTP-PTP Interworking

Starting Cisco IOS XR Software Release 7.11.1, NTP-PTP interworking provides the ability to use Precision Time Protocol (PTP), and other valid time of day (TOD) sources such as Data over Cable Service Interface Specification (DOCSIS) Timing Interface (DTI) and Global Positioning System (GPS), as the time source for the operating system in the units of nanosec level accuracy. PTP is capable of achieving nanosecond-level accuracy, while NTP is typically only accurate to within milliseconds. By using PTP as a reference clock, NTP can improve its accuracy and meet the needs of applications that require high precision timing.

Before the support of NTP-PTP interworking, only backplane time was supported for the operating system time of the router.

Table 2: Feature History Table

Feature Name	Release Information	Feature Description
NTP-PTP Interworking	Release 7.11.1	We have improved NTP synchronization and reliability to achieve nanosecond-level accuracy for applications that require high-precision timing. This is achieved by enabling NTP-PTP interworking which allows the use of PTP as the reference clock. As in previous releases, the NTP client continues to support polling NTP protocol-based external time servers to synchronize the local system clock and achieve accuracy within the millisecond range.

NTP-PTP interworking also provides the means to communicate status changes between PTP and NTP processes. It also supports the unambiguous control of the operating system time and backplane time in the event of bootup, switchovers, or card and process failures.

With NTP-PTP interworking, NTP is less likely to lose synchronization. As, PTP is more robust to network delays and disruptions than NTP. So, if there's a problem with the network, PTP can still maintain accurate synchronization.

Configuring NTP-PTP Interworking

Before you begin

- Ensure that PTP is enabled, before configuring NTP-PTP Interworking.

- For PTP, the Grandmaster (GM) gets the clock from a GPS/GNSS reference clock:

    - If the PTP-NTP feature is enabled on a GM node, verify that the GM gets clock reference from the FPS/GNSS clock reference and then configure the CLI on the GM node.

    - If the PTP-NTP feature is enabled on a Boundary Clock (BC) node, ensure that the GM gets clock reference from the FPS/GNSS clock reference and then configure the CLI on the BC node.

    - If the PTP-NTP feature is enabled on a Transparent Clock (TC) node, ensure that the GM gets the clock reference from the FPS/GNSS clock reference, and the BC node gets the clock from that GM node, the TC node gets the clock from the BC node, and then configure the CLI on the TC node.

- If the GM isn't connected to any GPS/GNSS reference clock, the default PTP clock is set to Jan 1, 1970.

Step 1 You can configure NTP-PTP Interworking in any of the following ways:

- Setting NTP Primary Reference Clock as PTP

```
Router# configure
Router(config)# ntp
Router(config-ntp)# master primary-reference-clock
Router(config-ntp)# commit
```

• Configuring NTP Server with IP address

The following example shows an NTP configuration to allow the system clock to be synchronized by time server hosts at IP address 198.51.100.1. You can take the IP address of a neighboring PTP interface.

```
Router# configure
Router(config)# ntp server 198.51.100.1
Router(config-ntp)# commit
```

Step 2   Verify the NTP status using the **show ntp status** command.

```
Router# show ntp status

Clock is synchronized, stratum 1, reference is 198.51.100.1
nominal freq is 1000000000.0000 Hz, actual freq is 101341889.2967 Hz, precision is 2**24
reference time is 8497CD13.A6AEB9DA (00:02:27.651 UTC Tue Jun 30 1970)
clock offset is -0.077 msec, root delay is 0.000 msec
root dispersion is 3937.89 msec, peer dispersion is 3937.74 msec
loopfilter state is 'CTRL' (Normal Controlled Loop), drift is 0.0000088676 s/s
system poll interval is 64, last update was 4 sec ago
authenticate is disabled, panic handling is disabled,
hostname resolution retry interval is 1440 minutes
```

Step 3   Verify that an NTP association has come up using the **show ntp associations** command.

```
Router# Show ntp associations
  address         ref clock     st  when  poll reach  delay  offset    disp
*~198.51.100.1    .PTP.          0    -    64    0   0.00   0.000   16000
```