NOKIA

Nokia Service Router Linux
7220 Interconnect Router
7250 Interconnect Router
7730 Service Interconnect Router

Release 26.3

Network Synchronization Guide

Table of contents

© 2026 Nokia.
Use subject to Terms available at: www.nokia.com/terms.

1 About this guide

This document describes the network synchronization features for the Nokia Service Router Linux (SR Linux).

This document is intended for network technicians, administrators, operators, service providers, and others who need to understand how the router is configured.

> **Note:** This manual covers the current release and may also contain some content that will be released in later maintenance loads. See the *SR Linux Release Notes* for information on features supported in each load.

Configuration and command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

1.1 Precautionary and information messages

The following are information symbols used in the documentation.

> **DANGER:** Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.

> **WARNING:** Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.

> **Caution:** Caution indicates that the described activity or situation may reduce your component or system performance.

> **Note:** Note provides additional operational information.

> **Tip:** Tip provides suggestions for use or best practices.

1.2 Conventions

The Nokia SR Linux documentation uses the following command conventions:

- **Bold** type indicates a command that the user must enter.
- Input and output examples are displayed in `Courier` text.
- A vertical bar (|) indicates a mutually exclusive argument.
- Square brackets ([ ]) indicate optional elements.

- Braces ({ }) indicate a required choice. When braces are contained within square brackets, they indicate a required choice within an optional element.
- *Italic* type indicates a variable.

The following table outlines platform grouping conventions used in the SR Linux documentation suite.

> **Note:** Some platforms in the 7250 IXR support mixed systems. For more information about mixed system support, see "Chassis types" in the *Configuration Basics Guide*.

Table 1: Platform grouping legend

Platform group	Description
7215 IXS	7215 IXS-A1
7220 IXR	All 7220 IXR platforms
7220 IXR-D*x*	7220 IXR-D1, 7220 IXR-D2, 7220 IXR-D2L, 7220 IXR-D3, 7220 IXR-D3L, 7220 IXR-D4, 7220 IXR-D5
7220 IXR-H*x*	7220 IXR-H2, 7220 IXR-H3, 7220 IXR-H4, 7220 IXR-H4-32D, 7220 IXR-H5-32D, 7220 IXR-H5-64D, 7220 IXR-H5-64O
7250 IXR[1]	7250 IXR platforms
7250 IXR Gen 2	7250 IXR-6, 7250 IXR-10
7250 IXR Gen 2c+	7250 IXR-6e with IMM2, 7250 IXR-10e with IMM2, 7250 IXR-X1b, 7250 IXR-X3b
7250 IXR Gen 3	7250 IXR-6e with IMM3, 7250 IXR-10e with IMM3, 7250 IXR-18e, 7250 IXR-X4
7250 IXR-6e/10e (mixed system)	7250 IXR-6e (mixed system)[2], 7250 IXR-10e (mixed system)[2]
7730 SXR	7730 SXR-1-32D, 7730 SXR-1d-32D, 7730 SXR-1x-44S

---

[1] References to the 7250 IXR platform group may be appended with (including mixed systems) or (excluding mixed systems) to indicate mixed system support.

[2] References to this platform as part of 7250 IXR (mixed system) indicate mixed system support of 7250 IXR Gen 2c+ (IMM2) and 7250 IXR Gen 3 (IMM3). That is, the 7250 IXR-6e and 7250 IXR-10e can hold and support both IMM2 and IMM3 at the same time.

2 What's new

Table 2: What's new in Release 26.3.1

Topic	Location
Network synchronization support added for 7730 SXR-1-32D.	Network synchronization Central frequency synchronization subsystem

3 Network synchronization

This guide describes network synchronization capabilities available with SR Linux. These capabilities include physical layer frequency distribution via Synchronous Ethernet and packet based distribution of time using the precision time protocol (PTP) of IEEE 1588.

SR Linux supports network synchronization on the following platforms only:

- 7220 IXR-D5, model 3HE17735AB (w/Timing)

- 7250 IXR-6e/10e platforms with the following timing modules:

    – 3HE17011AC CPM — 7250 IXR-6e/10e CPM4 w/ RoT+SYNC

    – 3HE17006AB IMM — 7250 IXR 36x400G QSFP-DD IMM2+SYNC

- 7250 IXR-X1b/X3b

- 7730 SXR-series platforms: 7730 SXR-1x-44S, 7730 SXR-1d-32D, and 7730 SXR-1-32D

In the past, physical layer frequency distribution was required to ensure proper operation of the SDH and SONET network interfaces of a network. That requirement is now replaced by the need to support delivery of the following:

- a reference frequency for mobile base stations to tune their carrier frequencies

- an accurate frequency for the delivery of time using PTP

Note: SR Linux supports various network synchronization options. SR Linux network element ports can support both synchronous Ethernet and PTP/1588. See Supported synchronization options for more information.

The network time architecture in the following figure shows how network synchronization is commonly distributed in a hierarchical PTP topology at the physical layer.

Figure 1: Conventional network timing architecture (North American nomenclature)
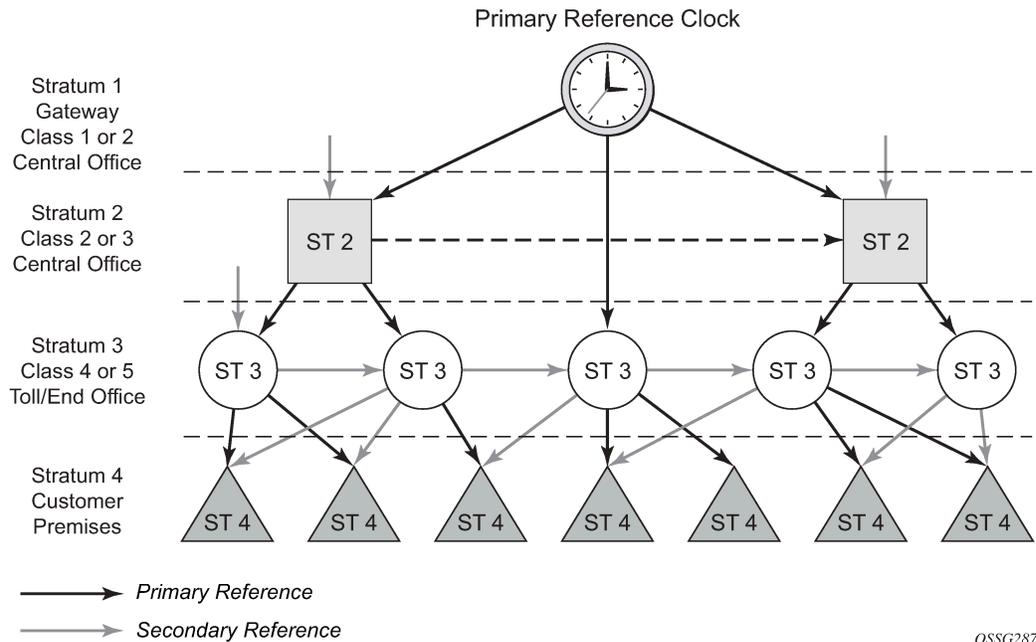
This architecture provides the following benefits:

- limits the need for high-quality clocks at each network element
- only requires reliable and accurate replication of the input to remain traceable to its reference
- uses reliable physical media to provide transport of the timing signal
- does not consume any bandwidth and requires limited additional processing

The synchronization network is designed so a clock always receives timing from a clock of equal or higher stratum level or quality level. This ensures that if an upstream clock has a fault condition (for example, it loses its reference and enters a holdover or free-run state) and begins to drift in frequency, the downstream clock is able to follow it. For greater reliability and robustness, most offices and nodes have at least two synchronization references that can be selected in priority order (such as primary and secondary).

Additional resiliency can be provided by the ability of the node clock to operate within prescribed network performance specifications in the absence of any reference for a specified period. A clock operating in this mode is said to hold the last known state over (or holdover) until the reference lock is once again achieved. Each level in the timing hierarchy is associated with minimum levels of network performance.

Each synchronization-capable port can be independently configured to transmit data using the node reference timing.

Specifically for synchronous Ethernet, transmission of a reference clock through a chain of Ethernet equipment requires that all equipment supports synchronous Ethernet. A single piece of equipment that is not capable of performing synchronous Ethernet breaks the chain. Ethernet frames still get through, but downstream devices should not use the recovered line timing because it is not traceable to an acceptable stratum source.

3.1 Central frequency synchronization subsystem

The timing subsystem has a central clock located on the CPM. The timing subsystem performs many of the duties of the network element clock as defined by Telcordia (GR-1244-CORE) and ITU-T G.781.

The central clock uses the available timing inputs to train its local oscillator. The number of timing inputs available to train the local oscillator varies per platform. The priority order of these references must be specified using the **instance 1** and **instance 2** parameters. When the central CPM clock is locked, the clock output is able to drive the clocking on all line cards in the system.

The routers support the selection of node reference using Quality Level (QL) indications. The recovered clock is able to derive its timing from one of the references available on that platform.

The following figure illustrates synchronization reference selection on the 7220 IXR-D5, 7250 IXR 6e/10e, and 7250 IXR-X1b/X3b. It shows how the recovered clock is able to derive the timing from Synchronous Ethernet ports or from a 1588/PTP timeReceiver port.

Figure 2: Logical model of the synchronization reference selection (7220 IXR-D5, 7250 IXR 6e/10e, and 7250 IXR-X1b/X3b)
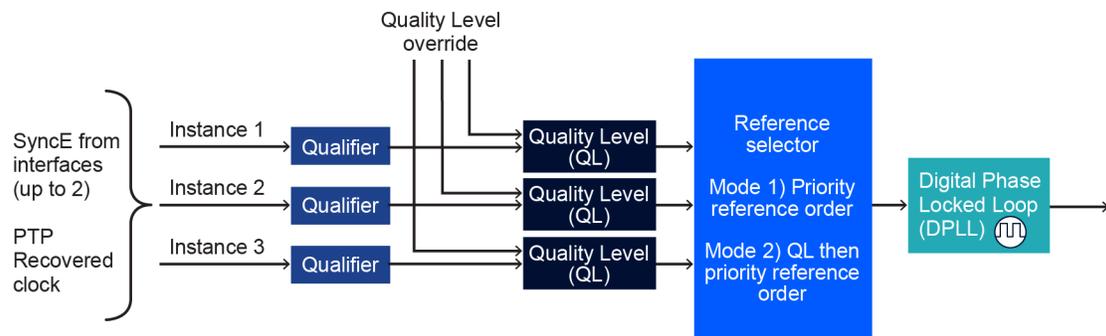
The following figure illustrates synchronization reference selection on the 7730 SXR-1d-32D and 7730 SXR-1-32D. It shows how the recovered clock is able to derive the timing from synchronous input ports from any of the following references:

- Synchronous Ethernet ports
- GNSS RF ports
- SyncE/1588 port on the CPM

Figure 3: Logical mode of the synchronization reference selection (7730 SXR-1d-32D and 7730 SXR-1-32D)
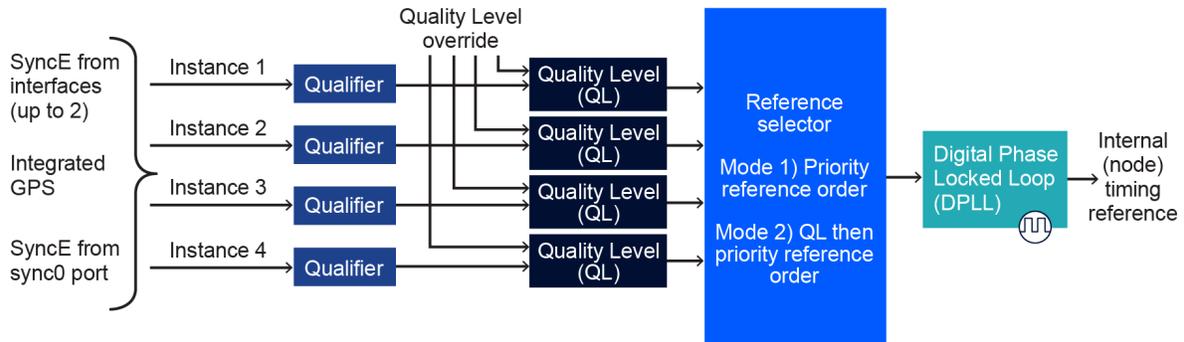
The following figure illustrates synchronization reference selection on the 7730 SXR-1x-44S. It shows how the recovered clock is able to derive the timing from synchronous input ports from any of the following references:

- Synchronous Ethernet ports
- 1588/PTP timeReceiver port
- GNSS RF ports
- SyncE/1588 port on the CPM

Figure 4: Logical mode of the synchronization reference selection (7730 SXR-1x-44s)
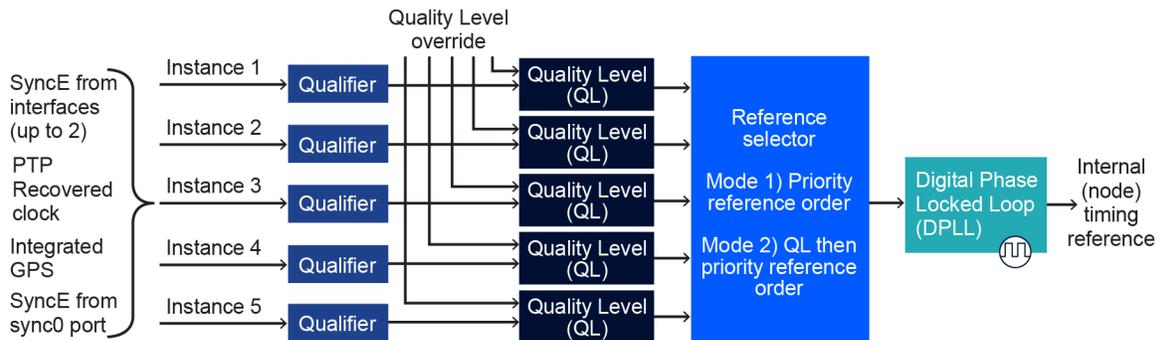
The **revert** setting controls when the central clock can re-select a previously failed reference.

The following table describes the selection followed for two references in both revertive and non-revertive modes.

Table 3: Revertive and non-revertive timing reference switching operation

Status of reference A	Status of reference B	Active reference non-revertive case	Active reference revertive case
OK	OK	A	A

Status of reference A	Status of reference B	Active reference non-revertive case	Active reference revertive case
Failed	OK	B	B
OK	OK	B	A
OK	Failed	A	A
Failed	Failed	Holdover	Holdover
OK	OK	A or B	A

3.1.1 Synchronization Status Message

Synchronization Status Messages (SSM) allow the synchronization distribution network to determine the quality level of the clock sourcing a specific synchronization trail, and to allow a network element to select the best of multiple input synchronization trails. SSM has been defined for synchronous Ethernet for interaction with office clocks, such as BITS or SSUs, and embedded network element clocks.

SSM allows equipment to autonomously provision and reconfigure (by reference switching) their synchronization references, while helping to avoid the creation of timing loops. These messages are particularly useful for synchronization re-configurations when timing is distributed in both directions around a ring.

3.2 Central time synchronization subsystem

SR Linux supports GNSS and PTP to provide accurate time, which is essential for OAM timestamping. NTP does not provide sufficient accuracy for 1-way delay measurements for example. GNSS is only supported on a subset of the SR Linux platforms.

SR Linux does not provide a configuration for time reference selection, as there is a preset fixed priority order. By default, if GNSS is available, it is selected. Otherwise, if PTP is available, it is selected.

Figure 5: Default time reference selection
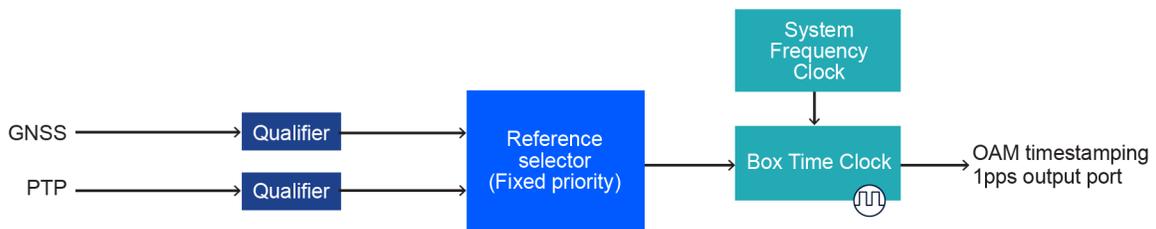
The SR Linux system clock is configured separately and is based on NTP.

3.3 Supported synchronization options

The following table summarizes the synchronization options that are available.

Table 4: Supported synchronization options

Synchronization option	Supported platform	Notes
SyncE with SSM	7220 IXR-D5; model 3HE17735AB (w/Timing) 7250 IXR-6e/10e (with timing modules) 7250 IXR-X1b/X3b 7730 SXR	Supported for frequency only
1588/PTP with port-based timestamps	7220 IXR-D5; model 3HE17735AB (w/Timing) 7250 IXR-6e/10e (with timing modules) 7250 IXR-X1b/X3b 7730 SXR	Supported for time and frequency
1pps interface	7220 IXR-D5; model 3HE17735AB (w/Timing) 7250 IXR-6e/10e (with timing modules) 7730 SXR	Supported for output only
GNSS	7730 SXR only	See GNSS.
SyncE/1588 port	7730 SXR only	

4 Synchronous Ethernet

Traditionally, Ethernet-based networks use the physical layer transmitter clock derived from an inexpensive +/-100 ppm crystal oscillator and the receiver locks onto it. There is no need for long-term frequency stability because the data is packetized and can be buffered. For the same reason, there is no need for consistency between the frequencies of different links.

However, you can derive the physical layer transmitter clock from a high-quality frequency reference by replacing the crystal with a frequency source traceable to a primary reference clock. This does not affect the operation of any of the Ethernet layers, for which this change is transparent. The receiver at the far end of the link locks onto the physical layer clock of the received signal and gains access to a highly accurate and stable frequency reference. Then, in a manner similar to conventional hierarchical network synchronization, this receiver can lock the transmission clock of its other ports to this frequency reference, and a fully time-synchronous network can be established.

The advantage of using synchronous Ethernet (SyncE), compared with methods that rely on sending timing information in packets over an unclocked physical layer, is that it is not influenced by impairments introduced by the higher levels of the networking technology (packet loss, packet delay variation). Therefore, the frequency accuracy and stability may exceed those of networks with unsynchronized physical layers.

SyncE allows operators to gracefully integrate existing systems and future deployments into a conventional industry-standard synchronization hierarchy. The concept behind SyncE is analogous to SONET/SDH system timing capabilities. It allows you to select any (optical) Ethernet port as a candidate timing reference. The recovered timing from this port can then be used to time the system. You can then ensure that any system output is locked to a stable, traceable frequency source.

Supported reference interfaces

Any Ethernet interface can be configured as the underlying reference interface for an instance, with the following exceptions:

- management port
- SFP+ ports 33 and 34 (on 7220 IXR-D5)

SyncE interface limitations

On the 7730 SXR-1d-32D, only one port of each odd/even numbered port pair (port 1/2, port 3/4, port 5/6, and so on) can be configured as a frequency timing source.

4.1 Timing reference selection based on quality level

For a synchronous Ethernet interface that supports Ethernet Synchronization Message Channel (ESMC) or for PTP, the clock class provides a quality level value to indicate the timing source of the far-end transmitter. These values provide input to the selection process on the nodal timing subsystem. This selection process determines which input to use to generate the signal on the SSM egress ports and the reference to use to synchronize the nodal clock, as follows:

- For the reference inputs, if the interface configuration supports the reception of a QL over SSM or ESMC, the quality level value is associated with the timing derived from that input.

- For the reference inputs, if they are synchronous Ethernet ports and the ESMC is disabled, the quality level value associated with that input is QL-UNKNOWN.

- For the reference inputs, if the interface configuration supports the reception of a QL over SSM or ESMC, but the quality level value received over the interface is not valid for the type of interface, the quality level value associated with that input is QL-INVALID.

- For the reference inputs, if they are external synchronization ports, the quality level value associated with the input is QL-UNKNOWN.

- For the reference inputs, if they are synchronous Ethernet ports and the ESMC is enabled, but no valid ESMC Information PDU has been received within the previous 5 seconds, the quality level value associated with that input is QL-FAILED.

- For GNSS reference input, the quality level is PRS/PRC if a frequency is successfully recovered; otherwise, the quality level is QL-FAILED.

- If you have configured an override for the quality level associated with an input, the node displays both the received and override quality level value for the input. If no value has been received, the associated value displays instead.

After the quality level values have been associated with the system timing inputs, the reference inputs and the external input timing ports are processed by the system timing module to select a source for the SSU. This selection process is as follows:

- Before an input can be used as a potential timing source, it must be enabled using the **ql-selection** command. If **ql-selection** is disabled, the priority order of the inputs for the Synchronous Equipment Timing Generator (SETG) is defined using the **priority** parameter configured under each **instance**.

- If **ql-selection** is enabled, the priority of the inputs is calculated using the associated quality level value of the input and using the **priority** parameter configured under each **instance**. The inputs are ordered by the internal relative quality level based on their associated quality level values. If two or more inputs have the same quality level value, they are placed in order based on where they appear in the priority. The priority order for the SETG is based on both the reference inputs and the external synchronization input ports.

- After a prioritized list of inputs is calculated, the SETG and the external synchronization output ports are configured to use the inputs in their respective orders.

- After the SETG and external synchronization output ports priority lists are programmed, the highest-qualified priority input is used. To be qualified, the signal is monitored to ensure that it has the expected format and a frequency that is within the pull-in range of the SETG.

4.2 Advanced G.781 features

The central clock of the node supports several advanced features of the ITU-T G.781 standard. These include the specification of minimum acceptable QL values for the input references and the specification of a Wait To Restore timer for input references. These features allow for more options in the management of the synchronization topology.

5 GNSS

The SR Linux routers with network synchronization support can run frequency synchronization using a Layer 1 interface such as synchronous Ethernet. In cases where these methods are not possible, or where accuracy cannot be ensured for the service, you can deploy a GNSS receiver as a synchronous timing source. GNSS data is used to provide network-independent frequency and ToD synchronization.

The following table lists the support by platform for GPS and Galileo references using an integrated GNSS RF port.

Table 5: GNSS timing references supported on SR Linux platforms

SR Linux platform	Integrated GNSS timing reference	
	GPS	**Galileo**
7730 SXR	✓	✓

An SR Linux chassis equipped with a GNSS receiver and an attached GNSS antenna can be configured to receive frequency traceable to Stratum-1 (PRC/PRS). The GNSS receiver provides a synchronization clock to the SSU in the router with the corresponding QL for SSM. This frequency can then be distributed to the rest of the router from the SSU as configured with the following commands:

- **system sync freq-references instance priority**
- **system sync freq-clock ql-selection**

The GNSS reference is qualified only if the GNSS receiver is in a position hold state and has a frequency successfully recovered.

If GNSS signal loss or jamming result in the unavailability of timing information, the GNSS receiver automatically prevents output of clock or synchronization data to the system, and the system can revert to alternate timing sources.

5.1 GNSS failure with APTS

When the G.8275.2 profile is used for GNSS-enabled 7220 IXR platforms, the APTS capability frequently measures and stores the delay offset between the GNSS time and a backup PTP session time. If a GNSS failure occurs, the backup PTP session automatically becomes the selected reference for time and frequency, and the stored delay offset value is added to or subtracted from the backup PTP session time to keep time and phase for the router as accurate as possible.

When GNSS has recovered and is stable, the system automatically switches back to GNSS for time and frequency reference, and backup PTP monitoring and delay measurement resumes.

Related topics

ITU-T G.8275.2 profile and APTS

6 IEEE 1588 PTP

✏️ **Note:** The IEEE 1588 Working Group has introduced the terms timeTransmitter and timeReceiver as alternatives to the former master/slave terminology. This document uses these new terms.

PTP is a timing-over-packet protocol defined in the IEEE 1588 Version 2 of the standard, published in 2008. The standard was revised to Version 2.1 in 2019.

PTP provides the capability to synchronize network elements to a stratum-1 clock, PRC traceable source, or PRTC traceable source over a network that may or may not be PTP-aware. PTP is a standards-based protocol, has low bandwidth requirements, can transport frequency, time, and phase, and generally provides better performance than NTP.

✏️ **Note:** Many applications do not need time alignment but only phase alignment. However, phase is derived from time and so, for the remainder of the document, references to time with regards to PTP imply both phase and time.

The following are the basic types of PTP clocks:

- ordinary clock
- boundary clock
- end-to-end transparent clock
- peer-to-peer transparent clock

The boundary clock and ordinary clock can be used for frequency and time distribution, depending on the selected profile.

The following table describes PTP clock type support.

Table 6: PTP clock support

| Platform	Ordinary clock		Boundary clock	Transparent clock	
	timeReceiver	time Transmitter		End-to-end	Peer-to-peer
7220 IXR-D5; model 3HE17735AB (w/ Timing)	Not supported	Not supported	✓	Not supported	Not supported
7250 IXR-6e/10e (with timing modules)	Not supported	Not supported	✓	Not supported	Not supported
7250 IXR-X1b/X3b	Not supported	Not supported	✓	Not supported	Not supported
7730 SXR	Not supported	✓ (see note)	✓	Not supported	Not supported

✏️ **Note:** To enable ordinary clock timeTransmitter, configure a boundary clock with all PTP ports set to **master-only true**.

The ordinary clock timeTransmitter, ordinary clock timeReceiver, and boundary clock communicate with neighboring IEEE 1588 clocks. These neighbor clocks can be ordinary clock timeTransmitters, ordinary clock timeReceivers, or boundary clocks. The communication can be based on either unicast IP sessions transported through IP interfaces or multicast Ethernet transported through Ethernet ports.

For unicast IP sessions, the external clocks are labeled "peers". There are two types of peers: configured and discovered. An ordinary clock timeReceiver or a boundary clock should have configured peers for each PTP neighbor clock from which it may accept synchronization information. The router initiates unicast sessions with all configured peers. An ordinary clock timeTransmitter or boundary clock accepts unicast session requests from external peers. If the peer is not a configured peer, it is considered a discovered peer. An ordinary clock timeTransmitter or boundary clock can deliver synchronization information toward discovered peers. The following figure shows the relationship of various neighbor clocks that use unicast IP sessions to communicate with an SR Linux router configured as a boundary clock with two configured peers.

Figure 6: Peer clocks
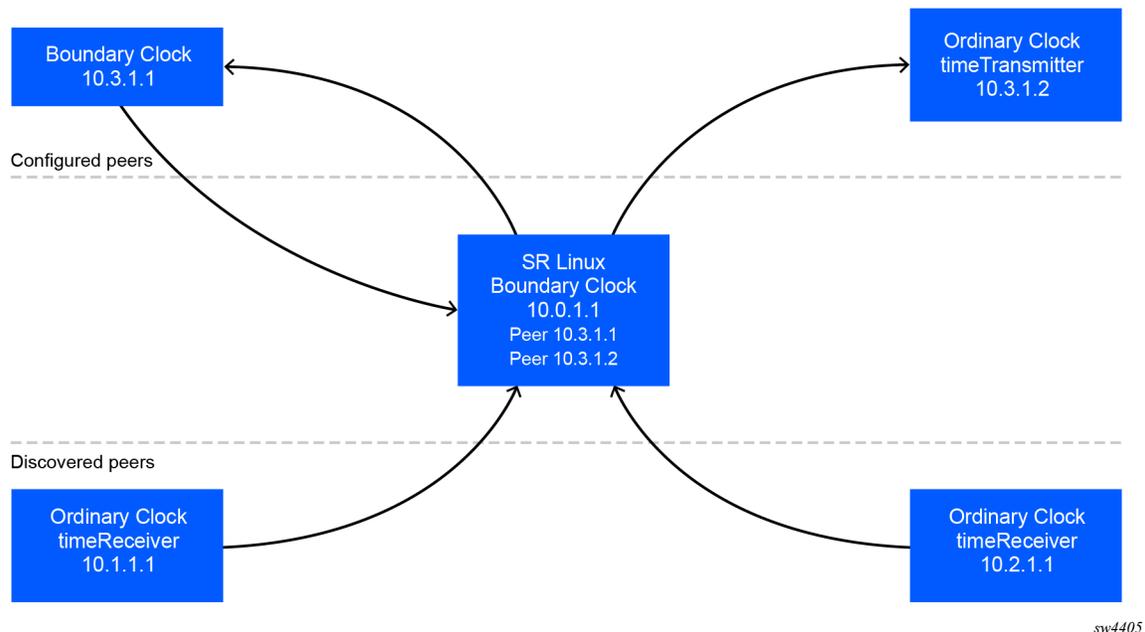
For multicast Ethernet operation, the router listens for and transmits PTP messages using the configured multicast MAC address. Neighbor clocks are discovered when messages are received through an enabled Ethernet port. An ordinary clock timeTransmitter, ordinary clock timeReceiver, and a boundary clock support more than one neighbor PTP clock connecting into a single port. This may be encountered with the deployment of an Ethernet multicast LAN segment between the local clock and the neighbor PTP ports using an end-to-end transparent clock or an Ethernet switch. The Ethernet switch is not recommended because of the introduction of packet delay variation (PDV) and the potential degradation of performance, but it can be used if appropriate for the application.

The following figure shows the relationship of various neighbor clocks using multicast Ethernet sessions to communicate with an SR Linux router configured as a boundary clock. The SR Linux router has three ports configured for multicast Ethernet communications. Port 1/2/1 of the SR Linux router shows a connection where two neighbor clocks connect to a single port of the SR Linux router through an end-to-end transparent clock.

Figure 7: Ethernet multicast ports
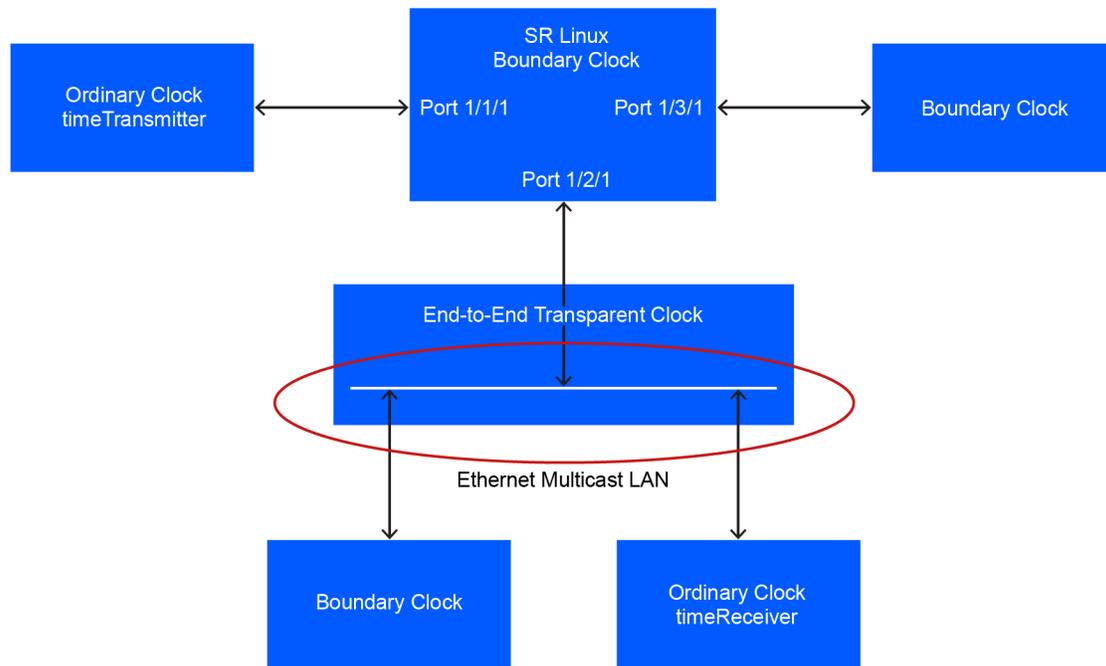
The ordinary clock timeTransmitter, ordinary clock timeReceiver, and boundary clock support PTP operation over unicast IPv4, unicast IPv6, or multicast Ethernet at any one time.

The IEEE 1588 standard includes the concept of PTP profiles. These profiles are defined by industry groups or standards bodies that define how IEEE 1588 is used for an application.

The SR Linux routers support the following profiles:

*   G.8275.1 (*itug8275dot1*)

    This profile is restricted to PTP over Ethernet and restricts PTP ports to be physical interfaces.

*   G.8275.2 (*itug8275dot2*)

    This profile is restricted to PTP over IP and restricts PTP ports to be PTP peers.

When the SR Linux routers receive an Announce message from one or more configured ports, it executes a Best timeTransmitter Clock Algorithm (BTCA) to determine the state of communication between it and its neighbors. The system uses the BTCA to create a hierarchical topology allowing the flow of synchronization information from the grandmaster clock through the network to all boundary and timeReceiver clocks. Each profile has a dedicated BTCA.

If the **profile** setting for the clock is *itug8275dot1*, the precedence order for the best timeTransmitter selection algorithm is very similar to that used with the default profile. It ignores the **priority1** parameter, includes a **local-priority** parameter, and includes the ability to force a port to never enter timeReceiver state (**master-only**). The precedence is as follows:

1. clockClass

2. clockAccuracy

3. PTP variance (offsetScaledLogVariance)

4. priority2

5. localPriority

6. clockIdentity

7. stepsRemoved from the grandmaster

Note: If the two clocks being compared have a clockClass of less than 128, the comparison of the clockIdentity is skipped. Skipping the comparison of the clockIdentity is common because the vast majority of clocks used as grandmasters advertise a clockClass of less than 128.

The following table describes the local parameter settings when the **profile** setting for the clock is *itug8275dot1*.

Table 7: Local clock parameters: itug8275dot1

Parameter	Value and description
clock-identity	Chassis MAC address following the guidelines of 7.5.2.2.2 of IEEE 1588
clock-class	6 — local clock is using a time reference from a GNSS receiver
	7 — local clock is in holdover after losing time reference from the local GNSS receiver for no more than ten minutes
	135 - local clock is configured as a boundary clock and in holdover; the boundary clock was previously locked to a grandmaster clock with a clock class of 6 and is still considered to be within the holdover specification
	165 — local clock is configured as a boundary clock and in holdover; the boundary clock was previously locked to a grandmaster clock with a clock class of 6 and is considered to be outside of the holdover specification
	248 — local clock is configured as boundary clock
clock-accuracy	0xFE — unknown
	0x21 — when using a time reference from a GNSS receiver
offset-scaled-log-variance	0xFFFF — not computed
	0x4e5d (1.8E-15) — when using a time reference from a GNSS receiver

If the **profile** setting for the clock is *itug8275dot2*, the precedence order for the best timeTransmitter selection algorithm is very similar to that used with the *itug8275dot1* profile. It ignores the **priority1** parameter, includes a **localPriority** parameter, and includes the ability to force a port to never enter timeReceiver state (**master-only**). The precedence is as follows:

1. clockClass

2. clockAccuracy

3. PTP variance (offsetScaledLogVariance)

4. priority2

5. localPriority

6. clockIdentity

7. stepsRemoved from the grandmaster

Note:
If the two clocks being compared have a clockClass of less than 128, the comparison of the clockIdentity is skipped. Skipping the comparison of the clockIdentity is the normal case because the vast majority of clocks used as grandmasters advertise a clockClass of less than 128.

The following table describes the 7220 IXR local parameter settings when the **profile** setting for the clock is *itug8275dot2*.

Table 8: Local clock parameters: itug8275dot2

Parameter	Value and description
clockIdentity	Chassis MAC address following the guidelines of 7.5.2.2.2 of IEEE 1588
clockClass	6 — local clock is using a time reference from a GNSS receiver
	7 — local clock is configured as a grandmaster clock and in holdover after losing time reference from the local GNSS receiver for no more than ten minutes
	135 - local clock is configured as a boundary clock in holdover; the boundary clock was previously locked to a grandmaster clock with a clock class of 6 and is still considered to be within the holdover specification
	165 — local clock is configured as a boundary clock in holdover; the boundary clock was previously locked to a grandmaster clock with a clock class of 6 and is considered to be outside of the holdover specification
	248 — local clock is configured as a grandmaster clock or boundary clock in free-run mode
	255 — local clock is configured as an ordinary clock timeReceiver
clockAccuracy	0xFE — unknown
	0x21 — when using a time reference from a GNSS receiver
offsetScaledLogVariance	0xFFFF — not computed
	0x4e5d (1.8E-15) — when using a time reference from a GNSS receiver

See ITU-T G.8275.2 profile and APTS for more information about the G.8275.2 profile.

There is a limit to the number of external PTP clocks from which the ordinary clock timeReceiver or boundary clock requests unicast service. There is also a limit to the number of external PTP clocks to which the ordinary clock timeTransmitter or boundary clock grants unicast service. An association where the boundary clock has a symmetric relationship with another boundary clock (they both have the other as a configured peer) consumes a request and a grant unicast service in each router.

There are limits to the maximum transmitted and received event message rates supported in the router. Each unicast IP service established consumes a portion of one of the unicast message limits. When either limit is reached, additional unicast service requests are refused by sending a grant response with zero in

the duration field. See the scaling guide for the appropriate release for the specific unicast message limits related to PTP.

Multicast messages are not considered when validating the unicast message limit. When multicast messaging on Ethernet ports is enabled, you need to monitor the PTP load to ensure that the load does not exceed the limits. The **show platform control detail** command identifies the load of the PTP software process. If the capacity usage reaches 100%, the PTP software process on the router is at its limit of transmitting and receiving PTP packets.

Because you cannot control the amount of PTP messaging that is being received over the Ethernet ports, use the following statistics commands to identify the source of the message load. Statistics shown for PTP ports are associated with Ethernet-encapulated PTP (G.8275.1 profile) and statistics for PTP peers are associated with IP-encapulated PTP (G.8275.2 profile).

- **info from state system sync ptp instance** *<instance>* **default-ds statistics** to display aggregate statistics
- **info from state system sync ptp instance** *<instance>* **port-ds-interface-list** *<port-index>* **statistics** to display total messages for a specific PTP port
- **info from state system sync ptp instance** *<instance>* **port-ds-cfg-ip-list** *<port-index>* **statistics** to display statistics for configured IP peers
- **info from state system sync ptp instance** *<instance>* **port-ds-dsc-ip-list** *<port-index>* **statistics** to display statistics for discovered IP peers

The following figure shows the unicast negotiation procedure performed between a timeReceiver and a peer clock that is selected to be the timeTransmitter clock. The timeReceiver clock requests Announce messages from all peer clocks but only requests Sync and Delay_Resp messages from the clock selected to be the timeTransmitter clock.

Figure 8: Message sequence between the PTP timeReceiver clock and PTP timeTransmitter clock
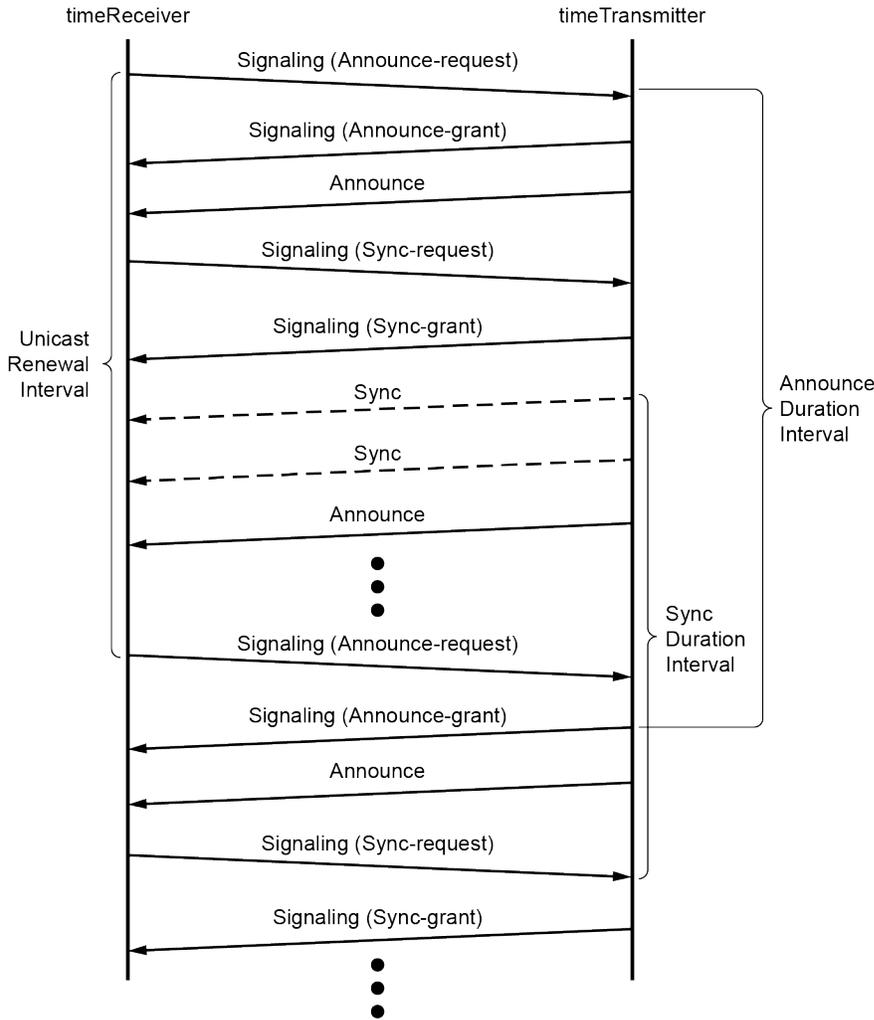
6.1 PTP clock synchronization

The IEEE 1588 standard synchronizes the frequency and time from a timeTransmitter clock to one or more timeReceiver clocks over a packet stream. This packet-based synchronization standard defines transport to use UDP/IP with unicast or Ethernet with multicast. See PTP capabilities for information about supported PTP encapsulation types.

As part of the basic synchronization timing computation, several event messages are defined for synchronization messaging between the PTP timeReceiver clock and PTP timeTransmitter clock. A one-step or two-step synchronization operation can be used, with the two-step operation requiring a follow-up message after each synchronization message.

Note: The SR Linux routers support only one-step timeTransmitter port operation. SR Linux routers can operate timeReceiver ports that receive from a one-step or two-step timeTransmitter port.

During startup, the PTP timeReceiver clock receives the synchronization messages from the PTP timeTransmitter clock before a network delay calculation is made. Before any delay calculation, the delay is assumed to be zero. A drift compensation is activated after several synchronization message intervals occur. The expected interval between the reception of synchronization messages is configurable.

The following figure shows the basic synchronization timing computation between the PTP timeReceiver clock and PTP best timeTransmitter. This figure illustrates the offset of the timeReceiver clock referenced to the best timeTransmitter signal during startup.

Figure 9: PTP timeReceiver clock and timeTransmitter clock synchronization timing computation
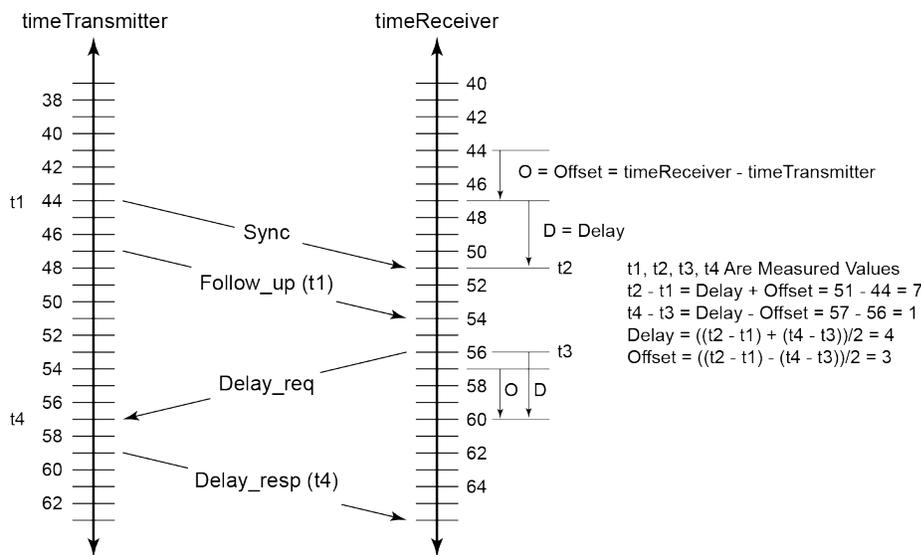
When using IEEE 1588 for distribution of a frequency reference, the timeReceiver calculates a message delay from the timeTransmitter to the timeReceiver based on the timestamps exchanged. A sequence of these calculated delays contain information of the relative frequencies of the timeTransmitter clock and timeReceiver clock but has a noise component related to the PDV experienced across the network. The timeReceiver must filter the PDV effects to extract the relative frequency data, then adjust the timeReceiver frequency to align with the timeTransmitter frequency.

When using IEEE 1588 for distribution of time, SR Linux uses the four timestamps exchanged using the IEEE 1588 messages to determine the offset between the SR Linux time base and the external timeTransmitter clock time base. SR Linux determines the offset adjustment and, in between these adjustments, maintains the progression of time using the frequency from the central clock of the node. This allows time to be maintained using a Synchronous Ethernet input source even if the IEEE 1588 communications fail. When using IEEE 1588 for time distribution, the central clock should at a minimum have the PTP input reference enabled.

The following figure shows a logical model for using PTP/1588 for network synchronization.

Figure 10: Logical model for using PTP/1588 for network synchronization
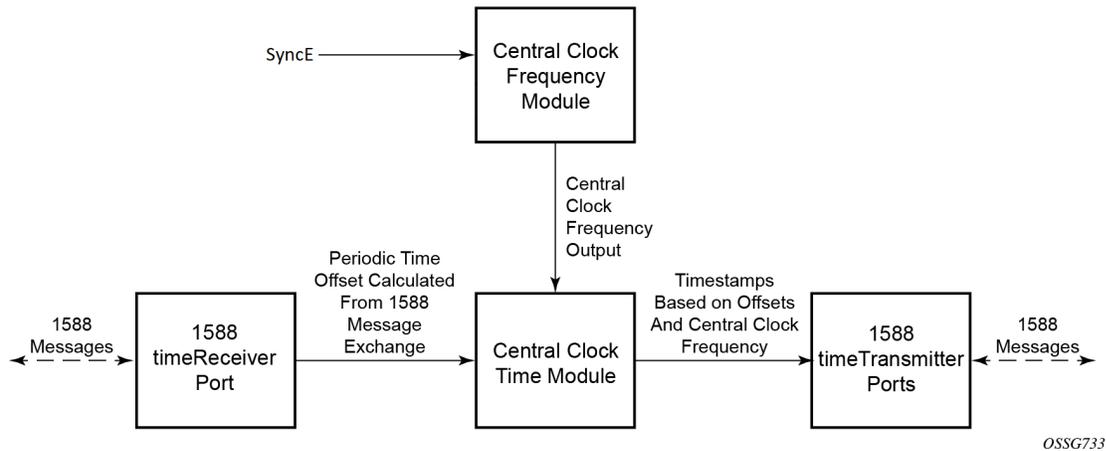
6.2 Performance considerations

Although IEEE 1588 can be used on a network that is not PTP-aware, the use of PTP-aware network elements (boundary clocks) within the packet switched network improves synchronization performance by reducing the impact of PDV between the grandmaster clock and the timeReceiver clock. When IEEE 1588 is used to distribute high accuracy time, such as for mobile base station phase requirements, the network architecture requires the deployment of PTP awareness in every device between the grandmaster clock and the mobile base station timeReceiver.

In addition, performance is improved by the removal of any PDV caused by internal queuing within the boundary clock or timeReceiver clock. This is accomplished with hardware that can detect and time stamp the IEEE 1588 packets at the Ethernet interface. This capability is referred to as port-based time stamping. The SR Linux platforms that are 1588-capable support port-based time stamping (see IEEE 1588 PTP).

> **Note:** 7730 SXR platforms equipped with a GNSS receiver can function as a grandmaster clock.

6.3 Port-based timestamping of PTP messages

To meet the stringent performance requirements of PTP mobile network applications, the 1588 packets must be time-stamped at the ingress and egress ports. This requires the use of 1588 port-based timestamping on the ports handling the PTP messages. This avoids any possible PDV that may be introduced between the port and the CPM. The ability to timestamp in the interface hardware is critical to achieve optimal performance. All PTP-capable ports support port-based timestamping.

6.4 PTP capabilities

The following table describes PTP encapsulation type and PTP profile support.

Table 9: PTP encapsulation type and PTP profile support

Platform	PTP encapsulation type			PTP profile	
	IPv4	IPv6	Ethernet	G.8275.1 (Ethernet only)	G.8275.2 (IPv4/IPv6)
7220 IXR-D5	✓	✓	✓	✓	✓
7250 IXR-6e/10e (with timing modules)	✓	✓	✓	✓	✓
7250 IXR-X1b/X3b	✓	✓	✓	✓	✓
7730 SXR	✓	✓	✓	✓	✓

With the G.8275.2 profile, every PTP message exchange between the router and an external 1588 clock must be established using the Unicast Negotiation procedures. The router allows the configuration of the message rate to be requested from external 1588 clocks. The router also supports a range of message rates that it grants to requests received from the external 1588 clocks.

The following table describes the ranges for the request and grant rates.

Table 10: Message rates ranges

Message type	Rates requested by SR Linux		Rates granted by SR Linux	
	Min	Max	Min	Max
Announce	1 packet per 16 seconds	8 packets per second	1 packet per 16 seconds	8 packets per second
Sync	1 packet per second	64 packets per second	1 packet per second	64 packets per second
Delay_Resp	1 packet per second	64 packets per second	1 packet per second	64 packets per second
(Duration)	300 seconds	300 seconds	1 second	1000 seconds

State and statistics data for each timeTransmitter clock are available to assist in the detection of failures or unusual situations.

6.5 PTP boundary clock

IEEE 1588 can function across a packet network that is not PTP-aware. However, the performance may be unpredictable. PDV across the packet network varies with the number of hops, link speeds, utilization rates, and the inherent behavior of the routers. By using routers with boundary clock functionality in the path between the grandmaster clock and the timeReceiver clock, one long path over many hops is split into multiple shorter segments, allowing better PDV control and improved timeReceiver performance.

This allows PTP to function as a valid timing option in more network deployments and allows for better scalability and increased robustness in specific topologies, such as rings.

Boundary clocks can simultaneously function as a PTP timeReceiver of an upstream grandmaster (ordinary clock) or boundary clock, and as a PTP timeTransmitter of downstream timeReceivers (ordinary clock) and boundary clocks, as shown in the following figure. The time scale recovered in the timeReceiver side of the boundary clock is used by the timeTransmitter side of the boundary clock. This allows time to be distributed across the boundary clock.

Figure 11: PTP boundary clock
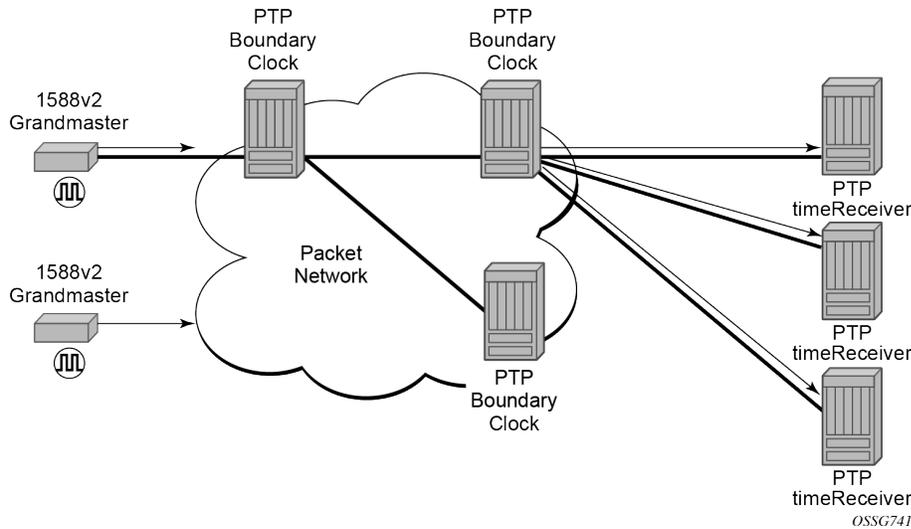
When a 7730 SXR router with an enabled GNSS port is configured with boundary clock functionality, the boundary clock acts as a grandmaster clock. The timeReceiver function stops and the timeTransmitter ports use frequency and time recovered from the GNSS port.

6.6 ITU-T G.8275.2 profile and APTS

SR Linux supports Recommendation ITU-T G.8275.2, which, similar to Recommendation ITU-T G.8275.1, specifies the architecture that allows the distribution of time and phasing. ITU-T G.8275.1 supports full timing support from the network and ITU-T G.8275.2 supports both partial timing support (PTS) and assisted partial timing support (APTS).

When the SR Linux router is configured to use the G.8275.1 or G.8275.2 profile, it uses an alternate BTCA for best timeTransmitter clock selection. This BTCA includes a PTP dataset comparison that is defined in IEEE 1588-2008, but with the following differences:

- the **priority1** attribute value is removed from the dataset comparison

- the **master-only** parameter value must be considered

- multiple active grandmaster clocks are allowed; therefore, the BTCA selects the nearest clock of equal quality

- a port-level **local-priority** attribute value is used to select a timeReceiver port if two ports receive an Announce message. This attribute is used as a tiebreaker in the dataset comparison algorithm if all other previous attributes of the datasets being compared are equal.

- the **local-priority** parameter value is considered for the default dataset

When the PTP clock is configured to use the G.8275.2 profile without GNSS configured, the clock operates using PTS. Only nodes that include an embedded GNSS can operate using APTS.

When the PTP clock is configured to use the G.8275.2 profile and the internal GNSS is configured and operationally up, GNSS is the preferred reference by default and is selected as the source of time for PTP. For extra resilience, APTS can be deployed by configuring a PTP backup over a network with PTP support to a remote grandmaster clock.

During normal operation, the local GNSS source is the reference for time and frequency. If the GNSS source fails, the PTP backup is automatically used to keep the clocks synchronized and stable. The delay offset value, which is calculated while the GNSS source was up, is applied to the PTP backup to keep time and phase as accurate as possible. See GNSS failure with APTS for more information about APTS functionality during a GNSS failure.

The following table describes the mapping between ITU-T G.8275.2 and PTP clock types. T-BC-A and T-TSC-A clocks are applicable to APTS.

Table 11: Mapping between ITU-T G.8275.2 and PTP clock types

Clock type from ITU-T G.8275.2	Description	Clock type from IEEE 1588
T-GM	timeTransmitter ordinary clock (clock with a single PTP port; cannot be a timeReceiver from another PTP clock)	Ordinary clock
	timeTransmitter boundary clock (clock with multiple PTP ports; cannot be a timeReceiver from another PTP clock)	Boundary clock[3]
T-BC-P (partial)	Boundary clock (may become a grandmaster clock, or may be a timeReceiver from another PTP clock)	Boundary clock
T-BC-A (assisted partial)	Boundary clock assisted by a local time reference that is used as a primary source of time (may become a grandmaster clock, or may be a timeReceiver to another PTP clock)	Boundary clock[4]
T-TSC-P (partial)	Always timeReceiver; single-port ordinary clock	Ordinary clock
	PTP clock at the end of the PTP synchronization chain; multiple port clock	Boundary clock[3]
T-TSC-A (assisted partial)	Always timeReceiver; single-port ordinary clock assisted by a local time reference that is used as a primary source of time	Ordinary clock[4]

---

[3] As defined by IEEE 1588, a clock with multiple PTP ports is a boundary clock.

[4] Examples of a local time reference include PRTC or a GNSS-based time source.

Clock type from ITU-T G.8275.2	Description	Clock type from IEEE 1588
	PTP clock at the end of the PTP synchronization chain; multiple-port clock assisted by a local time reference that is used as a primary source of time	Boundary clock[3, 4]

6.7 PTP message encapsulations

The SR Linux implementation of PTP over Ethernet provides support for PTP within raw Ethernet (no VLAN header) encapsulation.

SR Linux platforms that support PTP over IPv4 and IPv6 support the following encapsulations:

- PTP within UDP/IPv4 within raw Ethernet (no VLAN header)
- PTP within UDP/IPv6 within raw Ethernet (no VLAN header)
- PTP within UDP/IPv4 within dot1q Ethernet (one VLAN header)
- PTP within UDP/IPv6 within dot1q Ethernet (one VLAN header)

No other encapsulations are supported.

6.8 1PPS interfaces

The 1PPS interfaces are supported in SR Linux. These interfaces output the recovered signal from the system clock when PTP is enabled. 1PPS interfaces are supported on the following platforms:

- 7220 IXR-D5
- 7250 IXR-6e/10e (with timing modules)
- 7730 SXR

Note: 1PPS signals must not be used when PTP is not enabled. Additionally, these interfaces are output interfaces.

6.9 Configuration guidelines and restrictions for PTP

- The PTP timeReceiver and timeTransmitter capabilities are available on all Ethernet ports, except the management port and 7220 IXR-D5 SFP+ ports 33 and 34
- 1PPS interface signals are enabled when both PTP and the 1 PPS interface are enabled. 1PPS interface signals can be used after the system is configured to use PTP as a reference and is locked to PTP.

7 Configuration of SyncE references

Example: Using SyncE references

The following example shows the configuration of the frequency clock with two SyncE input references.

```
--{ candidate shared default }--[  ]--
# info with-context system sync
    system {
        sync {
            freq-references {
                instance 1 {
                    admin-state enable
                    priority 2
                    source {
                        interface ethernet-1/32
                    }
                }
                instance 2 {
                    admin-state enable
                    priority 3
                    source {
                        interface ethernet-1/1
                    }
                }
            }
            freq-clock {
                ql-input-threshold prs
                network-type sonet
                ql-selection true
                wait-to-restore 1
                revert true
            }
        }
    }
```

Example: Enabling SSM

The following example enables SSM (support of ESMC frames) on interface ethernet-1/32.

```
--{ candidate shared default }--[  ]--
# info with-context interface ethernet-1/32 ethernet synce
    interface ethernet-1/32 {
        ethernet {
            synce {
                ssm {
                    admin-state enable
                }
            }
        }
    }
```

8 Configuration of PTP references

Example: Using PTP references

The following example shows a PTP boundary clock configuration using profile itug8275dot1 with one upstream timeTransmitter and one downstream timeReceiver.

```
--{ candidate shared default }--[  ]--
# info with-context system sync ptp
    system {
        sync {
            ptp {
                instance 1 {
                    default-ds {
                        instance-type bc
                        instance-enable true
                        priority2 128
                    }
                    port-ds-interface-list 1 {
                        admin-state enable
                        master-only false
                        source {
                            interface ethernet-1/6
                        }
                    }
                    port-ds-interface-list 2 {
                        admin-state enable
                        source {
                            interface ethernet-1/21
                        }
                    }
                }
            }
        }
    }
```

Example: PTP boundary clock using itug8275dot2 profile

The following example shows a PTP boundary clock configuration using profile itug8275dot2 with one configured peer.

```
--{ candidate shared default }--[  ]--
# info with-context system sync ptp
    system {
        sync {
            ptp {
                ptp-profile itug8275dot2
                timing-source-net-inst ptp-ni
                instance 1 {
                    default-ds {
                        instance-type bc
                        instance-enable true
                    }
                    port-ds-cfg-ip-list 123 {
                        admin-state enable
                        peer {
                            ip-address 10.1.1.1
```

```
                }
              }
            }
          }
        }
      }
```
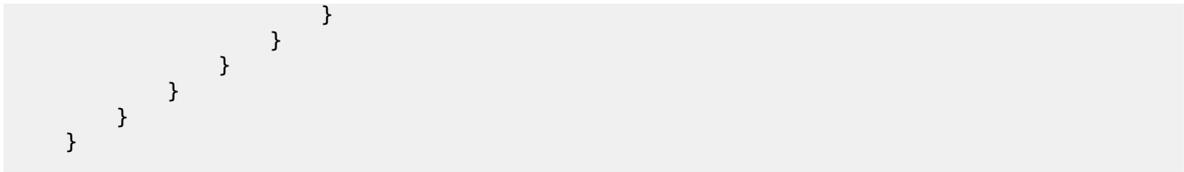
When one configured peer and one discovered peer exist on the same node, the **info from state system sync ptp** command displays entries under the **port-ds-cfg-ip-list** and **port-ds-dsc-ip-list** contexts, as shown in the following example.

The output in this example has been cropped to show only the most relevant leaves.

```
--{ candidate shared default }--[  ]--
# info from state with-context system sync ptp
    system {
        sync {
            ptp {
                ptp-profile itug8275dot2
                timing-source-net-inst ptp-ni
                instance 1 {
                    default-ds {
                        instance-type bc
                        instance-enable true
                        clock-identity 0x205E97FFFEAF107C
                        number-ports 2
...
                    parent-ds {
                        grandmaster-identity 0x18C300FFFE4A2C00
                        grandmaster-priority1 128
                        grandmaster-priority2 128
                        parent-port-identity {
                            clock-identity 0x18C300FFFE4A2C00
                            port-number 1
                        }
                        grandmaster-clock-quality {
                            clock-class 6
                            clock-accuracy 33
                            offset-scaled-log-variance 20061
                        }
                        protocol-address {
                            network-protocol udp-ipv4
                            ip {
                                network-instance ptp-ni-toGm
                                ip-address 10.1.1.1
                            }
                        }
                    }
...
                    port-ds-cfg-ip-list 123 {
                        ptp-port-number 1
                        admin-state enable
                        port-state slave
                        best-master true
                        parent-clock true
...
                        clock-identity 0x18C300FFFE4A2C00
                        port-number 1
                        grandmaster-identity 0x18C300FFFE4A2C00
                        grandmaster-priority1 128
                        grandmaster-priority2 128
                        steps-removed 0
                        last-rx-interface ethernet-1/32
```

```
                                last-tx-interface ethernet-1/32
                                peer {
                                    ip-address 10.1.1.1
                                }
...
                    port-ds-dsc-ip-list 1 {
                        ptp-port-number 2
                        port-state master
                        log-min-delay-req-interval -6
                        log-announce-interval 1
                        log-sync-interval -6
                        major-version-number 2
                        minor-version-number 0
                        peer {
                            network-instance ptp-ni-toSc
                            ip-address 10.4.4.4
                        }
...
                    }
                }
            }
        }
```

9 Network synchronization tools commands

The following table describes the tools commands related to network synchronization that allow you to clear PTP statistics and reset the PTP clock algorithm.

Table 12: PTP tools commands

Command	Definition
tools system sync ptp instance *instance* **clear-statistics**	Clears all statistics for the PTP instance.
tools system sync ptp instance *instance* **default-ds freq-recovery-engine statistics clear**	Clears all PTP statistics for the PTP frequency recovery engine.
tools system sync ptp instance *instance* **default-ds statistics clear**	Clears PTP statistics and event counters in the default data set.
tools system sync ptp instance *instance* **default-ds time-recovery-engine statistics clear**	Clears all PTP statistics for PTP time recovery engine.
tools system sync ptp instance *instance* **inactive-peers clear**	Clears information related to inactive PTP peers.
tools system sync ptp instance *instance* **port-ds-configured-peer** *port-index* **statistics clear**	Clears all PTP statistics in this PTP port data set for configured peers.
tools system sync ptp instance *instance* **port-ds-discovered-peer** *port-index* **statistics clear**	Clears all PTP statistics in this PTP port data set for discovered peers.
tools system sync ptp instance *instance* **port-ds-interface** *port-index* **statistics clear**	Clears all PTP statistics in this PTP port data set for interfaces.
tools system sync ptp instance *instance* **recovery-engine reset**	Resets the PTP clock algorithm and restarts the recovery process.

Customer document and product support

Customer documentation
Customer documentation welcome page

Technical support
Product support portal

Documentation feedback
Customer documentation feedback