



Nokia Service Router Linux

Release 24.3

Network Synchronization Guide

3HE 20251 AAAA TQZZA
Edition: 01
March 2024

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

© 2024 Nokia.

Table of contents

1	About this guide	4
1.1	Precautionary and information messages.....	4
1.2	Conventions.....	4
2	What's new	6
3	Network synchronization	7
3.1	Central synchronization subsystem.....	8
3.2	Supported synchronization options.....	9
3.3	Synchronization Status Message.....	10
4	Synchronous Ethernet	11
4.1	Clock source quality level definitions.....	11
4.2	Timing reference selection based on quality level.....	13
4.3	Advanced G.781 features.....	14
5	IEEE 1588 PTP	15
5.1	PTP clock synchronization.....	18
5.2	Port-based timestamping of PTP messages.....	20
5.3	PTP capabilities.....	20
5.4	PTP boundary clock.....	21
5.5	PTP message encapsulations.....	22
5.6	1PPS interfaces.....	22
5.7	Configuration guidelines and restrictions for PTP.....	22
6	Configuration of SyncE references	23
7	Configuration of PTP references	24

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

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.
- An open right-angle bracket indicates a progression of menu choices or simple command sequence (often selected from a user interface). Example: **start** > **connect to**.
- 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.

Generic IP addresses are used in examples. Replace these with the appropriate IP addresses used in the system.

2 What's new

There have been no updates in this document since it was last released.

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 frequency and time using the precision time protocol (PTP) of IEEE 1588.

SR Linux supports network synchronization on the following platform only:

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

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:

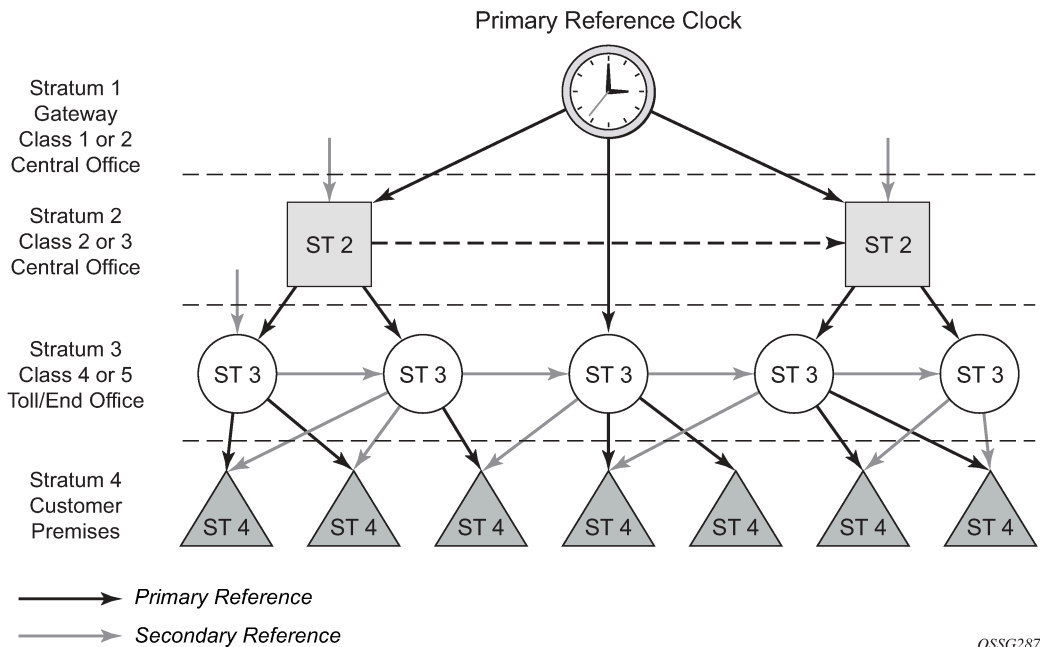
- a reference frequency for mobile base stations to tune their carrier frequencies
- an accurate frequency for the delivery of time using PTP



Note: Various network synchronization options are supported in SR Linux. The 7220 IXR-D5 can support both synchronous Ethernet and PTP/1588 on its ports. 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 synchronization subsystem

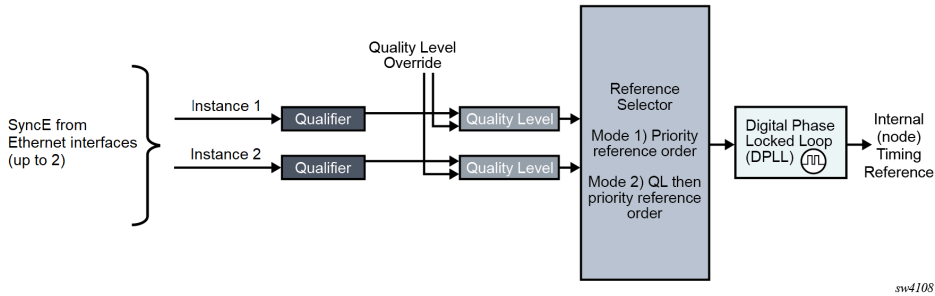
The timing subsystem has a central clock located on the CPM. The timing subsystem performs many of the duties of the network element clock as defined by Telcordia (GR-1244-CORE) and ITU-T G.781.

The central clock uses the available timing inputs to train its local oscillator. The 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. It shows how the recovered clock is able to derive the timing from Synchronous Ethernet ports.

Figure 2: Logical model of the synchronization reference selection



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 1: 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
Failed	OK	B	B
OK	OK	B	A
OK	Failed	A	A
Failed	Failed	Holdover	Holdover
OK	OK	A or B	A

3.2 Supported synchronization options

The following table summarizes the synchronization options that are available.

Table 2: Supported synchronization options

Synchronization option	Supported platform	Notes
SyncE with SSM	7220 IXR-D5; model 3HE17735AB (w/Timing)	—
1588/PTP with port-based timestamps	7220 IXR-D5; model 3HE17735AB (w/Timing)	—

Synchronization option	Supported platform	Notes
1pps interface	7220 IXR-D5; model 3HE17735AB (w/Timing)	Supported for output only

3.3 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.

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 unlocked 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
- two SFP+ ports (33 and 34)

SyncE interface limitations

SFP+ Ports 33 and 34 do not support SyncE (frequency or SSM).

4.1 Clock source quality level definitions

Clock source quality levels have been identified for tracking network timing flow. These levels make up all the defined network deployment options described in Recommendation G.803 and G.781. The Option I network is a network developed on the original European SDH model. The Option II network is a network developed on the North American SONET model.

In addition to the QL values received over SSM of an interface, the standards also define additional codes for internal use. These include the following:

- QL INVx is generated internally by the system if and when an unallocated SSM value is received, where x represents the binary value of this SSM. With SR Linux, all these independent values are assigned as the single value of QL-INVALID.
- QL FAILED is generated internally by the system if and when the terminated network synchronization distribution trail is in the signal fail state.

There is also an internal quality level of QL-UNKNOWN. This is used to differentiate from a received QL-STU code, but is equivalent for the purposes of QL selection.

The following table lists the synchronization message coding and source priorities for SSM received.

Table 3: Synchronization message coding and source priorities — SSM received

SSM value received on port		Internal relative quality level
SyncE interface in SDH mode	SyncE interface in SONET mode	
0010 (prc)	0001 (prs)	1. Best quality
	0000 (stu)	2
	0111 (st2)	3
0100 (ssua)	0100 (tnc)	4
	1101 (st3e)	5
1000 (ssub)		6
	1010 (st3/eec2)	7
1011 (sec/eec1)		8. Lowest quality qualified in QL-enabled mode
	1100 (smc)	9
		10
	1110 (pno)	11
1111 (dnu)	1111 (dus)	12
Any other	Any other	13. QL_INVALID
		14. QL-FAILED
		15. QL-UNC

The following table lists the synchronization message coding and source priorities for SSM transmitted.

Table 4: Synchronization message coding and source priorities — SSM transmitted

Internal relative quality level	SSM values to be transmitted by interface of type	
	SyncE interface in SDH mode	SyncE interface in SONET mode
1. Best quality	0010 (prc)	0001 (PRS)
2	0100 (ssua)	0000 (stu)
3	0100 (ssua)	0111 (st2)
4	0100 (ssua)	0100 (tnc)
5	1000 (ssub)	1101 (st3e)
6	1000 (ssub)	1010 (st3/eec2)
7	1011 (sec/eec1)	1010 (st3/eec2)
8. Lowest quality qualified in QL-enabled mode	1011 (sec/ eec1)	1100 (smc)
9	1111 (dnu)	1100 (smc)
10	1111 (dnu)	1111 (dus)
11	1111 (dnu)	1110 (pno)
12	1111 (dnu)	1111 (dus)
13. QL_INVALID	1111 (dnu)	1111 (dus)
14. QL-FAILED	1111 (dnu)	1111 (dus)
15. QL-UNC	1011 (sec/eec1)	1010 (st3/eec2)

4.2 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.
- 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.

4.3 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 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, which was published in 2008.

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 can potentially provide better performance.



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.

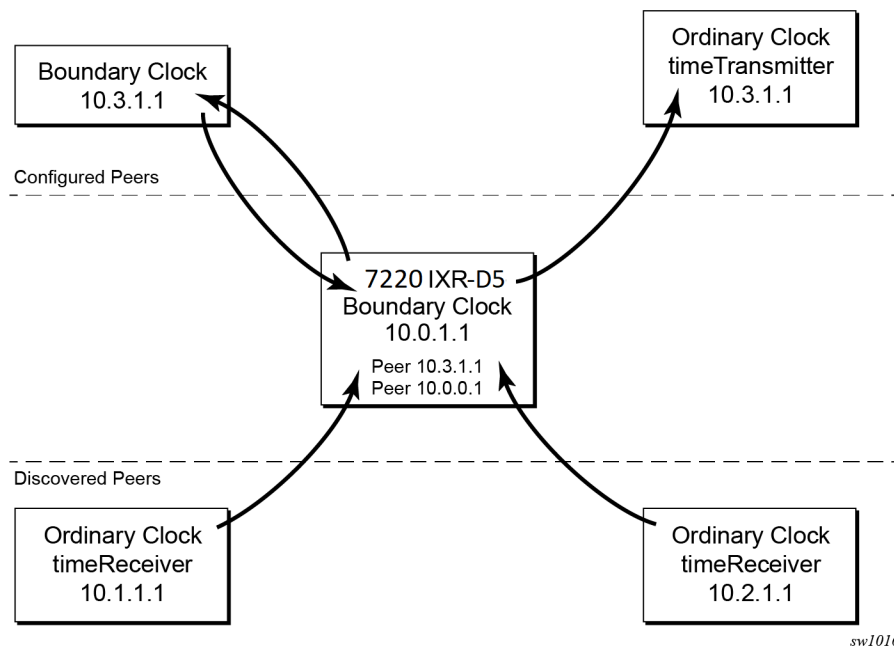
The following table describes PTP clock type support.

Table 5: 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

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 is based on multicast Ethernet transported through Ethernet ports.


Figure 3: 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 a 7220 IXR-D5 configured as a boundary clock. The 7220 IXR-D5 has three ports configured for multicast Ethernet communications. Port 1/2/1 of the 7220 IXR-D5 shows a connection where there are two neighbor clocks connecting to a single port of the 7220 IXR-D5 through an end-to-end transparent clock.

Figure 4: Ethernet multicast ports



al_0527.1

The ordinary clock timeTransmitter, ordinary clock timeReceiver, and boundary clock support PTP operation over multicast Ethernet.

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 7220 IXR-D5 supports the following profile:

- G.8275.1 (*g8275dot1-2014*)

This is the only profile supported on the 7220 IXR-D5.

When the 7220 IXR-D5 receives 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 *g8275dot1-2014*, 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 the normal case 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 *g8275dot1-2014*.

Table 6: Local clock parameters: *g8275dot1-2014*

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 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 6 248 — 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

5.1 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. The 7220 IXR-D5 supports Ethernet with multicast only. 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.

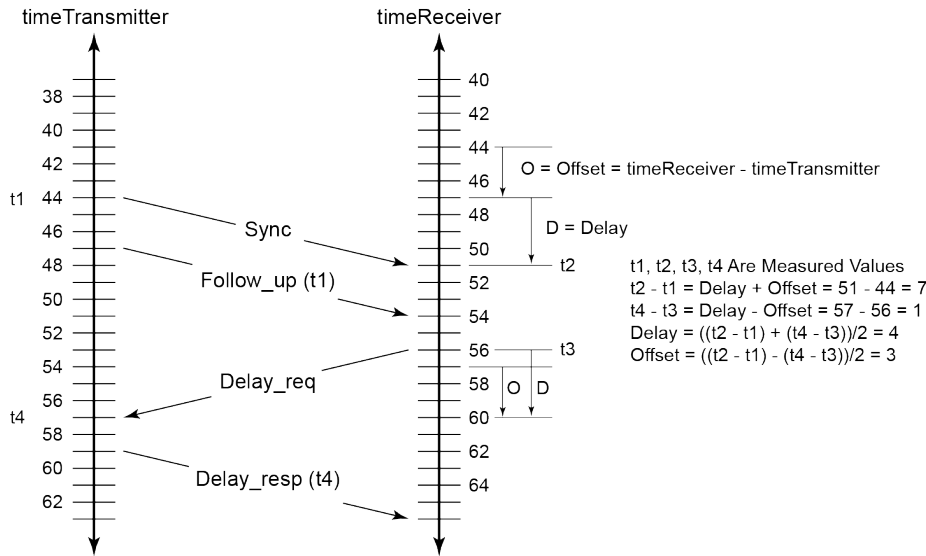


Note: The 7220 IXR-D5 supports only one-step timeTransmitter port operation. All node types can operate timeReceiver ports that receive from a one-step or two-step timeTransmitter port.

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 5: PTP timeReceiver clock and timeTransmitter clock synchronization timing computation




OSSG732

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 6: Logical model for using PTP/1588 for network synchronization



5.2 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.

5.3 PTP capabilities

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

Table 7: PTP encapsulation type and PTP profile support

Platform	PTP encapsulation type	PTP profile
	Ethernet (multicast only)	G.8275.1
7220 IXR-D5; model 3HE17735AB (w/Timing)	✓	✓

For each PTP message type exchanged between the router and an external 1588 clock, a unicast session 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 8: 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


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

5.4 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 7: PTP boundary clock



5.5 PTP message encapsulations

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

No other encapsulations are supported.

5.6 1PPS interfaces

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



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

5.7 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 PTP is enabled. 1PPS interface signals can be used after the system is configured to use PTP as a reference and is locked to PTP.

6 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 system sync
  system {
    sync {
      freq-references {
        instance 1 {
          admin-state enable
          underlying-interface ethernet-1/32
          priority 2
        }
        instance 2 {
          admin-state enable
          underlying-interface ethernet-1/1
          priority 3
        }
      }
      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 on interface ethernet-1/32.

```
# info interface ethernet-1/32 ethernet sync
interface ethernet-1/32 {
  ethernet {
    sync {
      ssm {
        admin-state enable
      }
    }
  }
}
```

7 Configuration of PTP references

Example: Using PTP references

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

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


Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)