

7	Clock Synchronization model for a bridged local area network.....	2
7.1	Peer calibration approach to network synchronization.....	2
7.1.1	Independent device clocks.....	2
7.1.2	Comparison to other synchronization approaches.....	2
7.1.3	Peer calibration mechanism.....	3
7.1.4	Peer calibration protocol stack.....	4
7.2	Overview of the synchronization services.....	4
7.2.1	Network event service.....	5
7.2.2	Global clock service.....	6
7.3	Overview of the protocols.....	7
7.3.1	Partner clock calibration protocol.....	7
7.3.1.1	Timestamp translation service.....	8
7.3.1.2	Link delay service.....	9
7.3.1.3	Required 802.3 partner clock calibration protocol.....	10
7.3.1.4	Required 802.11v partner clock calibraton protocol.....	14
7.3.2	Network event protocol.....	17
7.3.3	Global clock protocol.....	18
7.3.3.1	Network event payload.....	19
7.3.3.2	Precedence field.....	19
7.3.3.3	Clock entity behavior.....	20
7.3.3.4	Slave clock adjustment.....	20
7.3.4	Selection of the protocol parameters.....	21
7.3.4.1	Peer clock calibration parameters.....	21
7.3.4.2	Global clock update operational parameters.....	22

7 Clock Synchronization model for a bridged local area network

The clock synchronization model includes (1) synchronization services available to clients, and (2) protocols which support those services. The model defined for this standard is based on a peer calibration approach.

7.1 Peer calibration approach to network synchronization

7.1.1 Independent device clocks

Each node participating in an IEEE 802.1AS network implements a local timescale, i.e. a device clock, and references this timescale in the 802.1AS protocol messages it transmits. These device clocks are independent, and typically are derived from local free-running crystal oscillators. In contrast to some other approaches, the device clock is not adjusted or servoed to follow a master reference.

Operation of the 802.1AS protocol establishes a set of pairwise calibrated relationships among the device clocks, and uses these relationships to create a shared understanding of time throughout the network. These calibration relationships are maintained on a link-by-link basis, in parallel, without reference to any other link, to a master clock or to a preferred direction of information transfer. At this level the model is “democratic”: every device clock is equally valid, and simply defines a way of representing time which is convenient for protocol operations occurring within that device.

In addition to the free-running device timescale, a node may implement a global clock service defined in this standard; some 802.1AS protocol messages reference this global timescale as well as the device timescale. The global clock service at a node may be visualized as a variable-frequency controlled oscillator, slaved to a global master clock. The actual implementation of the global clock service may use a hardware clock servo or may be algorithmic.

7.1.2 Comparison to other synchronization approaches

The synchronization approach used in these protocols is best understood in comparison with other approaches to time distribution.

A basic model for synchronizing distributed clocks is to emit a master signal (e.g. the once-per-minute *beep* from radio station WWV) which is broadcast to all slave devices and updates them. Each slave maintains a controlled clock which “freewheels” between update signals. If the master signal propagated instantaneously and directly, this mechanism would be sufficient to maintain precise synchronization.

When propagation delay and indirect (multi-hop) communications enter the picture, two alternative concepts are often used to extend the above basic model:

1. Measurement of the propagation delays and relay delays (residence times) along the path taken by the master sync signal on its path to each slave, and compensation for this cumulative delay at the slave (e.g. Transparent Clock method in IEEE 1588v2); or
2. Implementing a slave clock at each relay node, and using that slave clock to regenerate and transmit fresh sync signals to stations on the far side of the relay (e.g. Boundary Clock method in IEEE 1588v2).

Both of these concepts can be effective, and have characteristic strengths and weaknesses. However the peer calibration approach used in this standard implements a third, distinct, concept in dealing with propagation delay and indirect communication scenarios.

7.1.3 Peer calibration mechanism

An important element of the peer calibration implementation is the “cross time stamp” or “time affiliation datum”. A cross time stamp is useful where several timescales coexist; it is a statement “Time T_A on timescale A is the same instant as time T_B on timescale B.” Cross timestamps allow us to make statements *about* time which are true independent of *when* the statement is made or received. This is a fundamental conceptual reorientation from methods based on precisely capturing the arrival of a sync update (“it is now 2:15 – *beep*”).

In the peer calibration approach, global time is distributed from a master station A, through intermediate nodes B and C to end station D, as follows:

1. Station A (master) initiates a clock update cycle, creating a cross-timestamp message “Master clock update: T_{MASTER} is T_A at A”
2. Station B forwards it, rewriting the cross timestamp, saying “Master clock update: T_{MASTER} is T_B at B.” It can do this because it maintains a calibrated relationship between timebases B and A.
3. Station C forwards it, rewriting the cross timestamp, saying “Master clock update: T_{MASTER} is T_C at C.” It can translate from T_B to T_C because it maintains a calibrated relationship between timebases C and B.
4. Station D receives it, and translates T_C to T_D , obtaining an internal representation of “Master clock update: T_{MASTER} is T_D at D”. This station now has the information it requires to update its clock: the master time T_{MASTER} corresponding to a locally-known time T_D .

Note that the update information at station D refers to some instant in the past, but the information is nonetheless accurate, even with no knowledge of how long it took to relay the message from A to D. In this way it is distinct from the alternative 1 in clause 7.1.2. Also note that no intermediate slave clock is used at nodes B or C; in this way it is distinct from alternative 2 in clause 7.1.2. The principle advantage of the peer calibration approach is that it is free from the clock cascade phenomenon, in which a small error in a

slave clock's estimation of master time is amplified by a gain factor as the protocol passes the master timescale from station to station. A chain of N cascaded clocks exhibits gain peaking, in which errors scale by k^N . Because an intermediate node in the peer calibration approach never uses an estimate of master time (either to regenerate a sync signal or to measure residence time) there is no cascade; errors in a chain of N stations accumulate additively rather than by power law.

7.1.4 Peer calibration protocol stack

An overall view of the services and protocols for clock synchronization in an IEEE 802.1AS system are illustrated in Figure 7-1.

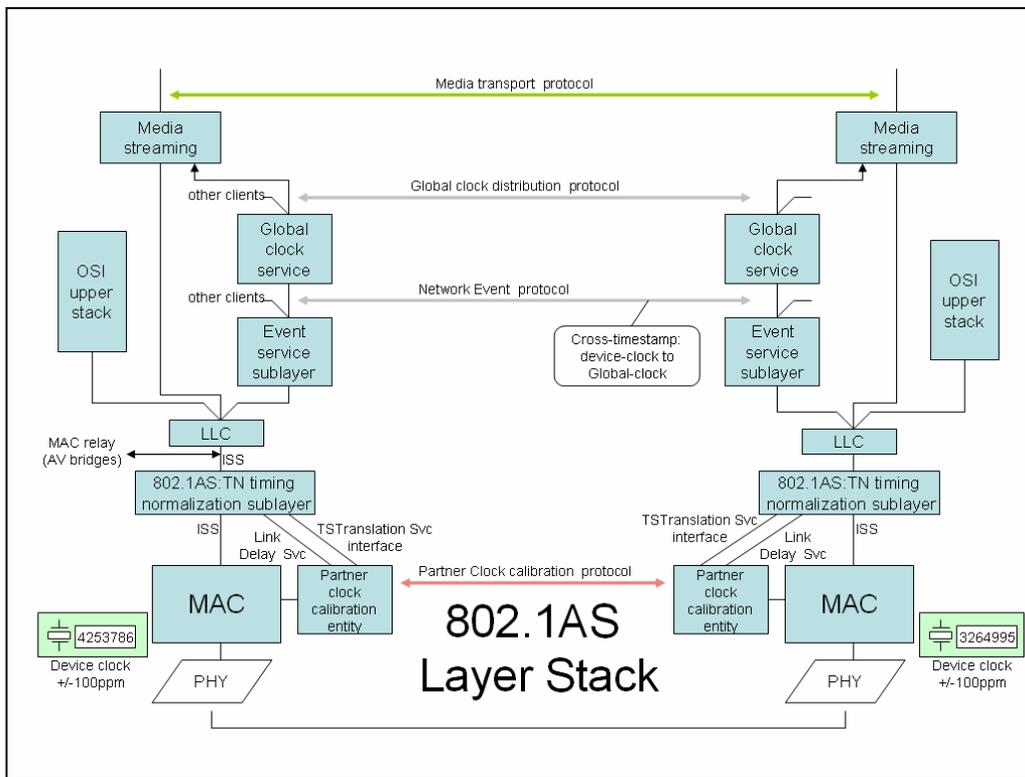


Figure 7-1. Protocol Stack

7.2 Overview of the synchronization services

This standard defines two synchronization services available to clients connected to an IEEE 802 bridged local area network: the *network event service* and the *global clock service*. The network event service is the more fundamental of the two, and assumes only that each client has a locally available timebase as described in 7.1.1. The global clock service is a derived service which provides every client access to a common timescale synchronized to a single grandmaster clock.

The successful establishment of the calibrated time relationships described in clause 7.1, within a set of connected 802.1AS-compliant nodes, renders those nodes capable of delivering the network event service. Information may be transmitted over the network event service using unicast or multicast addressing.

7.2.1 Network event service

The network event service consists of a request primitive and an indication primitive. The service interface exists in the context of a local device timescale.

The semantics of each NetworkEvent primitive is modeled on the IEEE 802.2 DL_UNITDATA primitive:

<pre>NetworkEvent.request (Source MAC Address; Destination MAC Address; Local Event Time; Cumulative Link Delay; Payload; Priority;)</pre>	<pre>NetworkEvent.indication (Source MAC Address; Destination MAC Address; Local Event Time; Cumulative Link Delay; Payload; Priority;)</pre>
--	---

Presentation of a request primitive at one node in an 802.1AS network results in the presentation of a corresponding indication primitive at zero or more connected nodes, as directed by the Destination MAC Address field. Each indication primitive contains the same field values as the request primitive, except that

1. the Local Event Time field is adjusted so as to represent the identical instant of time in the timescale context of the indication, and
2. the Cumulative Link Delay field is replaced by the accumulated sum of link delays (not bridge residence delays) involved in delivering the message.

The Cumulative Link Delay field in the NetworkEvent.request primitive is ignored.

The payload field is of variable length and is opaque to the Network Event service layer. However this standard defines the first octet of the payload to be an Event Type identifier, assigned as follows:

Event Type value	Meaning
0x00	Global clock service protocol ver. 1 (clause 7.2.2)
0x01-0x7F	Reserved
0x80-0xFF	Unassigned (user specified)

Table 7-1. Network event payload types

Client applications may find user-specified payload types to be suitable for sensor reports, actuator control, alternative time distribution schemes, and other purposes outside the scope of the global clock service described in clause 7.2.2.

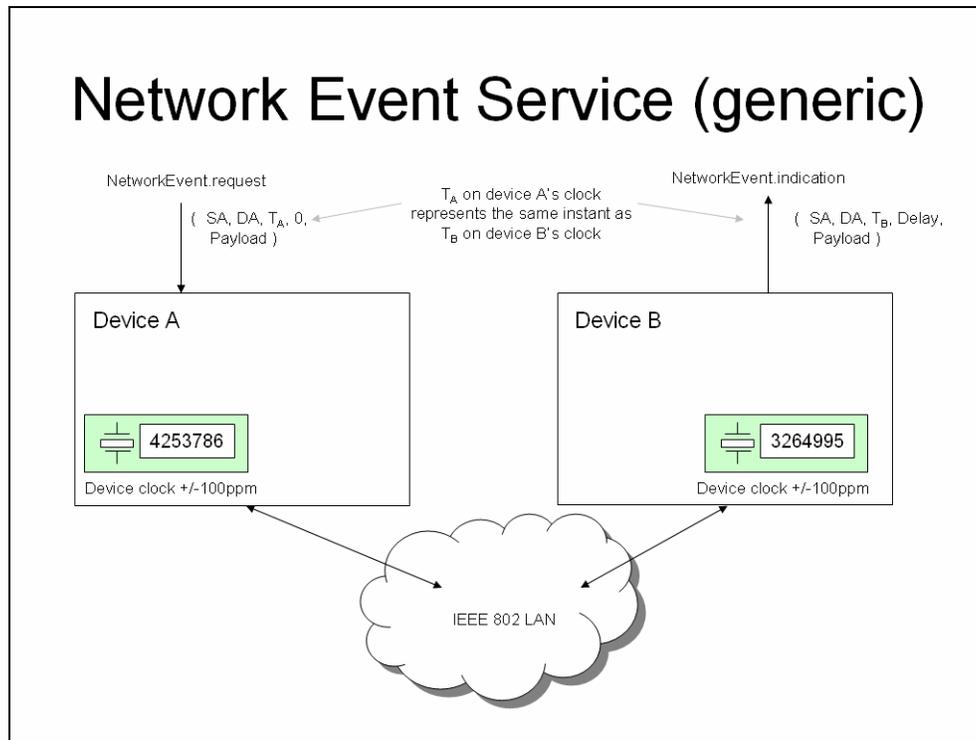


Fig. 7-2. Network Event Service

7.2.2 Global clock service

The global clock service is intended to provide a synchronized timing service throughout a bridged local area network. Its function is comparable to the NTP (network time protocol) service commonly provided in IP (internet protocol) environments.

At any time, a single node is selected as the “grandmaster clock”, providing a master timescale which is distributed to all participating 802.1AS nodes in a network. As the network configuration changes (e.g. if the current grandmaster becomes disconnected), a different node may be selected to provide the grandmaster clock function.

An instance of the global clock service presents an interface by which a client may discover the current time. This interface is unaffected by whether the current grandmaster clock is located at a local or distant node in the network.

The global clock provides an “Event/Timestamp” interface. The interface consists of an event request primitive, a timestamp indication primitive, and a status parameter. The event request primitive has no internal fields: the presentation of the request primitive

alone constitutes an event. Presentation of an event request primitive by the client results, a short time later, in the presentation of a timestamp indication primitive to the client, containing a numerical value representing the global time at which the event was presented.

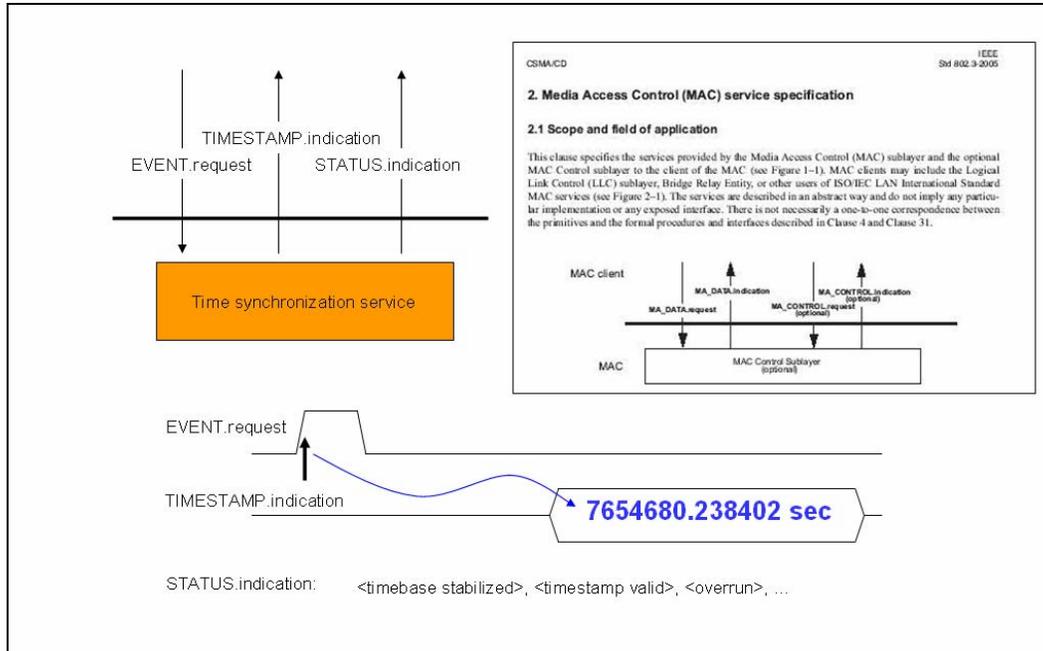


Figure 7-3. Global clock interface

7.3 Overview of the protocols

Two fundamental protocols support the network event service:

1. the partner clock calibration protocol, and
2. the network event protocol.

The global clock service is supported by the network event service, using a particular payload field. The payload contains information from the current grandmaster clock, including a current clock-time reading.

7.3.1 Partner clock calibration protocol

The partner clock calibration protocol is a symmetric protocol executed between two entities at opposite ends of a communication link. Each entity – a partner clock calibration entity – is associated with a single network port of an end device or a bridge. Each port operates in the context of a local timescale (device clock) as described in clause 7.1. In all device time fields transmitted with the 802.1AS protocol, time values are expressed as unsigned integers with the least significant bit representing 2^{-40} sec. Thus

the maximum representable time for 8-octet fields (e.g. ingress/egress timestamps) is 2^{24} sec, or about 194 days. As device timescale epochs are arbitrary, all algorithms should be designed to accommodate rollover of timestamp values.

The primary purpose of the partner clock calibration protocol is for each entity to establish a relationship between its own timescale and the timescale of its partner entity. Based on this relationship, the partner clock calibration entity is able to convert between timescales; specifically it can translate a timestamp expressed in the distant timescale into a corresponding timestamp in its own timescale. This capability is expressed as a TTS (timestamp translation service – clause 7.3.1) interface to clients within the device protocol stack.

An additional purpose of the partner clock calibration protocol is to determine the transmission delay over the attached link. This capability is expressed as the ability to update a cumulative delay value, accessible through an LDS (link delay service) interface.

Typically a partner clock calibration entity will internally track its partner's device clock in terms of an *offset* term and a *rate* term, and will time-filter them to reduce noise. These terms describe the performance of the distant clock by comparison to the local timescale only; there is no concept of a global timescale involved. It is possible to use more sophisticated clock models (e.g. those including a *rate drift* term), the result is simply to provide a more accurate timestamp translation service under certain application conditions.

The operation of the partner clock calibration entity has some similarity to a PLL (phase locked loop) slaved to the frequency and phase of the link partner's device clock. However a slave clock is only capable of answering the question “what time is it *now* on my partner's clock?”, while the partner clock calibration entity answers questions of the form “what time was it when my partner's clock read 2:15?” The timestamp translation service describes the latter type of query, and the corresponding answer.

The optimum technique for performing partner clock calibration depends on the specific medium constituting the communications link and the application environment in which it operates. This standard defines one technique which must be supported on IEEE 802.3 network ports, and a second technique which must be supported on IEEE 802.11 network ports, in order for equipment to interoperate. The protocol defined for IEEE 802.3 ports is suitable when a full-duplex point-to-point link directly connects two devices implementing it. It is permissible for vendors to implement additional, perhaps proprietary techniques, and such implementations are assured to support the network event and global clock services, provided the corresponding protocol entities present a functionally compatible TTS client service meeting the error budget limits of clause 7.3.4.1.

7.3.1.1 Timestamp translation service

The TTS, or timestamp translation service, is an abstract internal service interface defined as an aid to modeling the interaction of the partner clock calibration protocol and the network event protocol. The TTS is not directly accessible to applications and there is no implication that actual device designs should or should not employ such an interface in their implementation.

The TTS interface is associated with a device port and exists in the context of a local timescale and a partner timescale. It consists of two primitives, TTS_partner_time.request and TTS_local_time.indication. Each primitive contains a single semantic element:

Device_Timestamp

Upon presentation of a TTS_partner_time.request primitive containing a device timestamp *partner-time*, the partner clock calibration entity will respond to its client with a TTS_local_time.indication primitive. The TTS_local_time.indication primitive will contain a timestamp *local-time* which is an estimate of the local device clock value corresponding to the same instant referenced by the *partner-time* value in the partner's timescale.

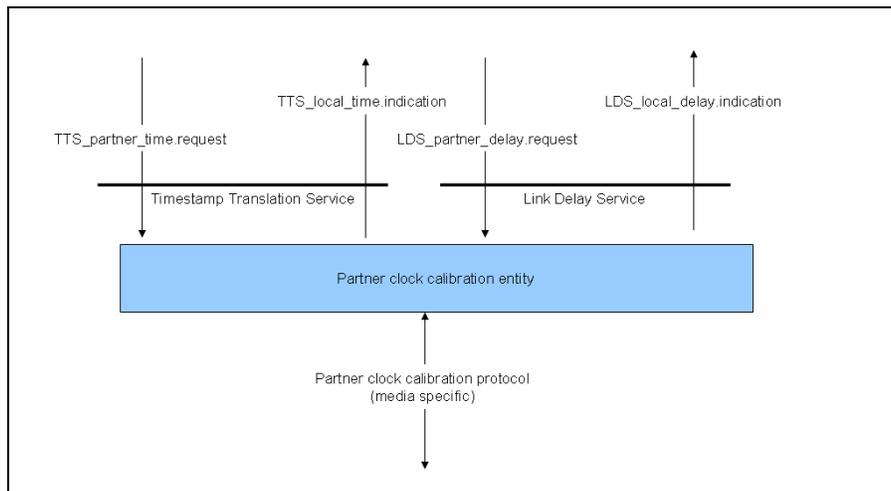


Figure 7-4. TTS and LDS interfaces.

7.3.1.2 Link delay service

The LDS, or link delay service, is an abstract internal service interface defined as an aid to modeling the interaction of the partner clock calibration protocol and the network event protocol. The LDS is not directly accessible to applications and there is no implication that actual device designs should or should not employ such an interface in their implementation.

The LDS interface is associated with a device port and exists in the context of a local timescale and a partner timescale. It consists of two primitives, LDS_partner_delay.request and LDS_local_delay.indication. Each primitive contains a single semantic element:

Cumulative_Delay

Upon presentation of an LDS_partner_delay.request primitive containing an element *pre-link-delay*, the partner clock calibration entity will respond to its client with an LDS_local_delay.indication primitive. The LDS_local_delay.indication primitive will contain *post-link-delay* which is the sum of the corrected *pre-link-delay* and the estimated link delay, expressed in the local device timescale. The *pre-link-delay* datum requires correction because it is expressed in the partner's timescale.

7.3.1.3 Required 802.3 partner clock calibration protocol

The defined 802.3 partner clock calibration protocol is a symmetrical protocol, in which each link endpoint periodically transmits a LocalSync frame to the other. The transmission times of the two entities need not be synchronized; in fact the protocol operates correctly if the two partner entities transmit at substantially different rates.

For each LocalSync frame transmitted or received, the port registers a timestamp representing the instant at which the frame passed a particular reference point. The protocol assumes that the transit delay from the sender's reference point to the receiver's reference point is fixed (or very slowly varying), and that the transit delay is equal in both directions.

The measured timestamps at each link partner are shared with the other partner by means of LocalSync messages. On this basis, each partner can calculate an estimate of the other partner's clock behavior.

The LocalSync frame contains 14 fields:

Field	Size octets	Function
Source Address	6	Source MAC
Destination Address	6	Assigned link-local MAC 01-80-C2-00-00-0E
Ethertype	2	0x88F7 assigned to 802.1AS
Function	1	Function code distinguishing LocalSync frame (gets timestamp)
Version	1	Protocol version = 0x00
Sequence	1	Message sequence number
ET ₁	6	Egress timestamp, message 1
IT ₁	6	Ingress timestamp, message 1
ET ₂	6	Egress timestamp, message 2
LogVariance	2	Nominal variance of sender's timestamps, typically based on clock granularity
LogMaxRate	2	Maximum tolerable incoming LocalSync message rate
Fill	21	Fill/padding Ignored if Version = 0x00
FCS	4	Frame Check Sequence
	64	

Table 7-2. LocalSync frame.

The timestamps in the message are reported in the context of the link partners' respective device timescales. At the issuer of the LocalSync frame, "message 1" refers to the most recently *received* LocalSync frame for which the issuer has both egress and ingress timestamps available. Being a received frame, the ingress timestamp for message 1 has been generated locally, while the egress timestamp is known by reception of an earlier LocalSync message. "Message 2" refers to the most recently *transmitted* LocalSync frame, i.e. the frame with a Sequence field one less than the current frame.

Upon receiving a LocalSync frame i , the peer clock calibration entity may confirm that the received sequence number has advanced by one (i.e. no packets were lost) and retrieve the previous ingress timestamp, which is associated with message 2 and is identified as IT_2 . It may then compute as follows:

$$local-point_i = (ET_1 + IT_2)/2$$

$$partner-point_i = (ET_2 + IT_1)/2$$

These two timestamps represent the same instant in the local and partner device timescales, respectively, assuming symmetrical link propagation delay.

Using a previous LocalSync frame $i-N$, N steps prior to frame i , the peer clock calibration entity may compute the frequency ratio between its own device clock and its partner's:

$$rate_i = (local-point_i - local-point_{i-N}) / (partner-point_i - partner-point_{i-N})$$

In order to reduce the effect of measurement jitter, the peer clock calibration entity will perform temporal low-pass filtering on the *local-point*, *partner-point*, and *rate* parameters. The timestamp translation service may then convert *partner-time* to *local-time* according to:

$$local-time = (partner-time - partner-point) \cdot rate + local-point$$

As an example, a single-pole discrete-time IIR (infinite impulse response) low-pass filter with z -transform response

$$F(z) = 0.02 / (1 - 0.98 \cdot z^{-1})$$

may be applied identically to the sequences *local-point_i*, *partner-point_i*, and *rate_i*.

The LogVariance field is informational and describes the expected variance of timestamp errors from the sending device, expressed as a 16-bit signed integer according to

$$\text{LogVariance} = 256 \cdot \log_2(\text{variance}),$$

with *variance* expressed in seconds². For example, if the errors have a +/- 20ns uniform statistical distribution (25MHz clock granularity), the variance (mean squared error) is 133 ns² and LogVariance is -5848. This field may be used by the recipient in implementation-specific ways, such as optimizing the filter behavior or message rate of the protocol. Note that this field does *not* characterize the performance of the device timescale, but rather of the timestamping process as a measure of ingress and egress events on that timescale. It is expected that a device will have a fixed design-dependent value for this field, influenced by both clock granularity and any timing uncertainty between PHY and timestamp capture. A device should provide an accurate value in this field; if no accurate value is provided a value of 32767 ($2^{15}-1$) must be placed in this field.

The LogMaxRate field is informational and describes the maximum rate at which the sending device will correctly process incoming LocalSync messages, expressed as a 16-bit signed integer according to

$$\text{LogMaxRate} = 256 \cdot \log_2(\text{message-rate}),$$

with *message-rate* expressed in frames per second. For example, if the maximum rate is 100 messages per second, the LogMaxRate value is 1701. A device must provide a value no less than <<specified default message rate, currently 100/sec>>. This field may be used by the recipient in implementation-specific ways, such as optimizing the filter behavior or message rate of the protocol. A peer clock calibration entity may send LocalSync frames at any rate up to that expressed in its partner's LogMaxRate field. The conditions under which various message rates may be generated are implementation-specific and outside the scope of this standard.

In order to estimate the link delay, the partner clock calibration entity may process the timestamps described above as follows:

$$local-interval_i = (IT_2 - ET_1)$$

$$partner-interval_i = (ET_2 - IT_1)$$

$$delay_i = (local-interval_i - rate \cdot partner-interval_i)/2$$

and applying temporal filtering to derive a smoothed *delay* parameter. Note that the intermediate variables *local-interval_i* and *partner-interval_i* may be either positive or negative. The *delay* parameter is used in providing the link delay service, according to:

$$post-link-delay = pre-link-delay \cdot rate + delay .$$

The maximum cumulative link delay value represented in the 802.1AS protocol suite is approximately 256 seconds ($2^{-40} \cdot (2^{48} - 1)$). If the computed *post-link-delay* value is larger than this, the value $2^{-40} \cdot (2^{48} - 1)$ must be returned at the link delay service interface.

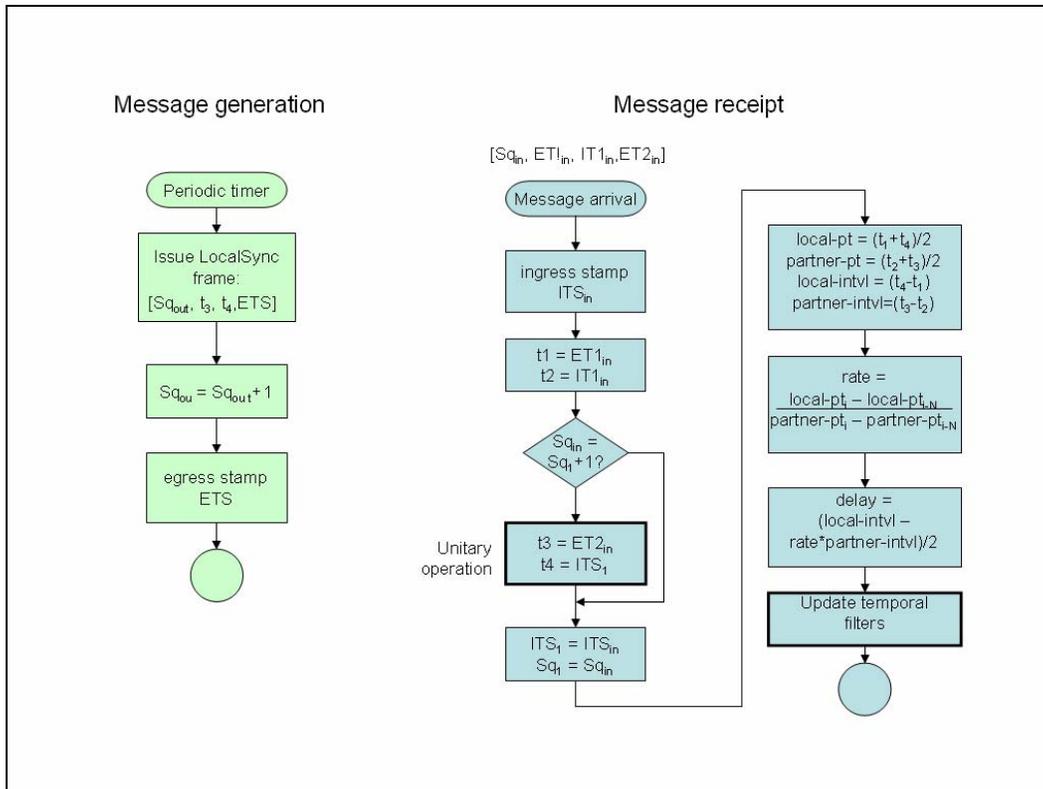


Figure 7-5. Partner clock calibration protocol: message processing.

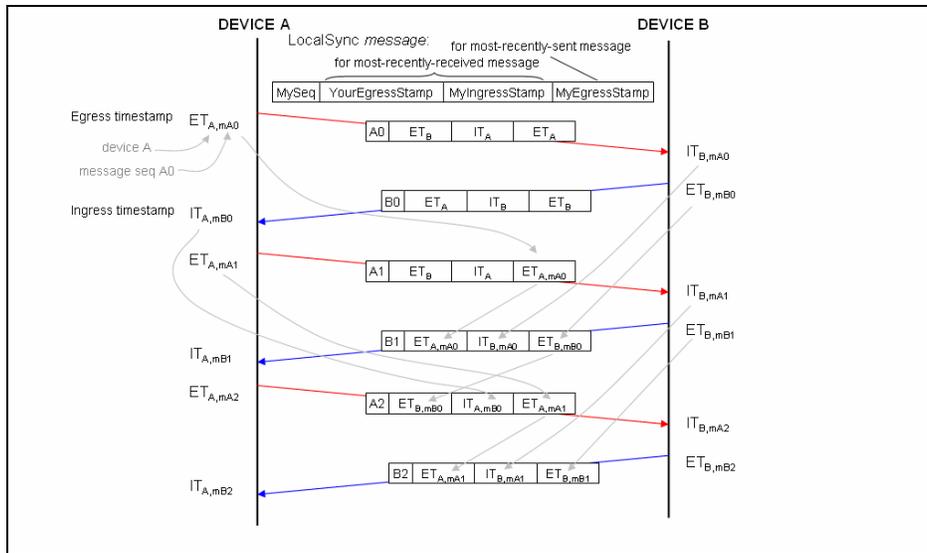


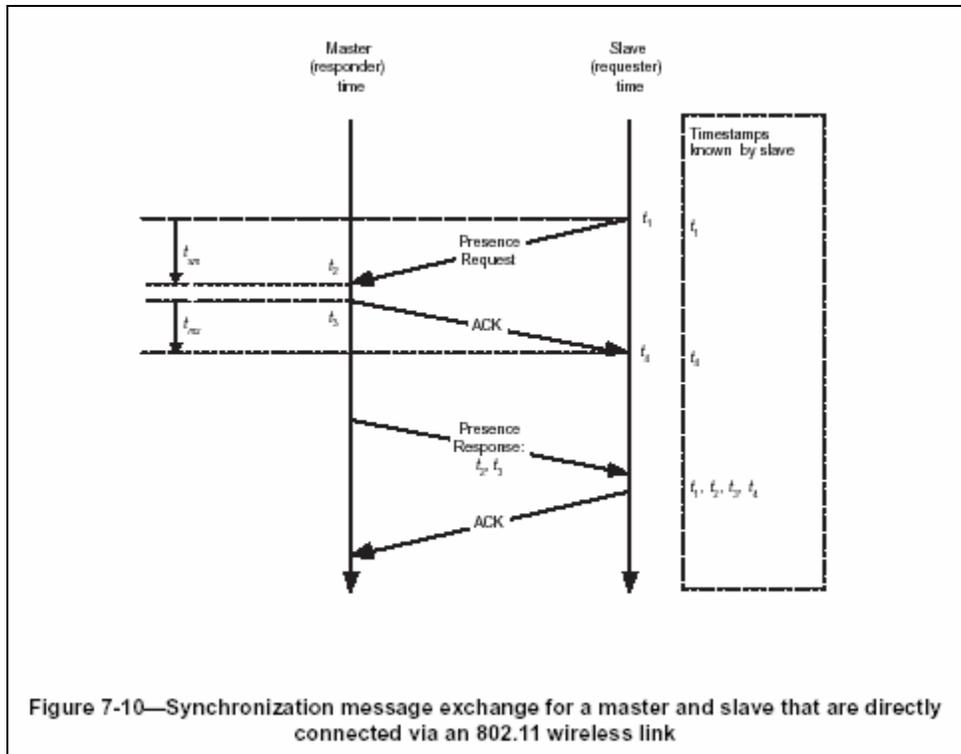
Figure 7-6. Partner clock calibration protocol: message exchange.

<<need modeling to optimize coefficients for jitter/wander vs lag>>

<<so far this is suitable for point to point. could probably be extended to work with shared media by paying attention to Source Address field>>

7.3.1.4 Required 802.11v partner clock calibration protocol

The 802.11 partner clock calibration protocol is based on the presence request messages defined in IEEE 802.11v.



The message exchange pattern is:

- a) The requester sends a Presence Request message to the responder and notes the time, t_1 , at which it was sent
- b) The responder receives the Presence Request message and notes the time of reception, t_2
- c) The responder returns an ACK message to the requester, and notes the time, t_3 , at which it was sent
- d) The requester receives the ACK message and notes the time of reception, t_4
- e) The responder communicates the times t_2 and t_3 to the requester in a Presence Response message
- f) The requester sends an ACK message to the responder to acknowledge that it received the Presence Response message

At the conclusion of this exchange of messages, the requester possesses all four timestamps (t_1 , t_2 , t_3 , and t_4). The requester's peer clock calibration entity may then compute parameters for the current exchange i as follows:

$$local-point_i = (t_1 + t_4)/2$$

$$partner-point_i = (t_2 + t_3)/2$$

These two timestamps represent the same instant in the local and partner device timescales, respectively, assuming symmetrical link propagation delay.

Using a previous presence message exchange $i-N$, N steps prior to message i , the peer clock calibration entity may compute the frequency ratio between its own device clock and its partner's:

$$rate_i = (local-point_i - local-point_{i-N}) / (partner-point_i - partner-point_{i-N})$$

In order to reduce the effect of measurement jitter, the peer clock calibration entity will perform temporal low-pass filtering on the *local-point*, *partner-point*, and *rate* parameters. The timestamp translation service may then convert *partner-time* to *local-time* according to:

$$local-time = (partner-time - partner-point) \cdot rate + local-point$$

As an example, a single-pole discrete-time IIR (infinite impulse response) low-pass filter with *z*-transform response

$$F(z) = 0.02 / (1 - 0.98 \cdot z^{-1})$$

may be applied identically to the sequences *local-point_i*, *partner-point_i*, and *rate_i*.

<<need modeling to optimize coefficients for jitter/wander vs lag>>

The 802.11v presence message exchange is not symmetrical, so only the requester is able to perform these computations. The presence exchange is therefore executed periodically in both directions, with the link partners adopting alternative roles as requester and responder. By this means, both peer clock calibration entities are enabled to offer the timestamp translation service within their respective devices' protocol stacks.

<<add description of delay computation for 802.11v links>>

7.3.2 Network event protocol

The network event protocol consists of NetworkEvent frames passed through the network from the issuer of the network event request primitive to the recipients of the network event indication primitive. The NetworkEvent frame consists of eleven fields:

Field	Size octets	Function
Source Address	6	Source MAC
Destination Address	6	Destination MAC (unicast, multicast, or broadcast)
Ethertype	2	0x88F7 assigned to 802.1AS
Function	1	Function code distinguishing Network Event frame (no ingr./egr. timestamp)
Version	1	Version identifier = 0x00
Event Timestamp	6	Timestamp (sender's timescale)
Cumulative Link Delay	4	Accumulated delay measurement
Payload size	2	Length (octets) of Payload field
Payload	N	Payload (variable length)
Fill (optional)	M	Fill/padding Ignored when version = 0x00
FCS	4	Frame Check Sequence
	32+N+M	

Table 7-3. Network event frame.

Presentation of a NetworkEvent.request primitive at a device causes it to generate a network event protocol frame with the parameters given in the primitive. The Cumulative Link Delay field is initialized to zero prior to transmission.

The Event Timestamp field and the Cumulative Link Delay field are valid in the context of the device transmitting the frame. A device receiving such a frame addressed to itself must translate the timestamp into its own device timescale before presenting the semantic content of the frame as a network event indication. In addition, the device must adjust the Cumulative Delay field by translating it into its own device timescale and adding the incoming link delay. The timestamp translation and cumulative delay adjustments are provided by the partner clock calibration entity (clause 7.3.1).

A bridge device is responsible for forwarding network event frames to its attached ports in accordance with the model of IEEE 802.1D and IEEE 802.1Q specifications. However an 802.1AS bridge is also responsible for adjusting the Event Timestamp and Cumulative Link Delay fields in each outgoing NetworkEvent frame so that they are valid in the device timescale context of the outgoing port, and for maintaining a correct FCS field. When all bridge ports share the same device clock context, the adjustment may be

achieved by applying the timestamp translation service and link delay service at the ingress port, then forwarding the corrected frames through the bridge fabric as specified by existing 802.1 standards. A bridge device complying with the initial version of this standard should ignore the Version field on incoming frames and shall set the Version field of outgoing frames to 0x00. (Note: this behavior is intended to provide forward compatibility with subsequent versions which make use of the Fill field for semantic content.)

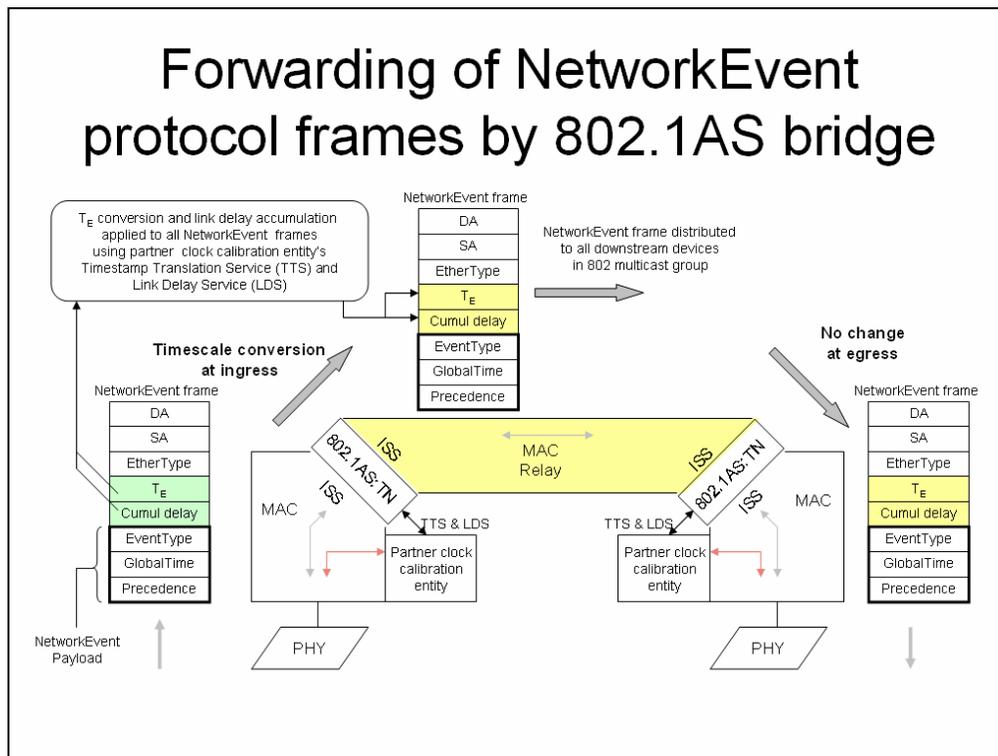


Figure 7-8. Network Event frame forwarding.

7.3.3 Global clock protocol

The global clock protocol is layered above the network event protocol, and uses a specific form of network event primitives to implement a grandmaster clock distribution system within a bridged local area network. In typical use, the grandmaster distributes global time by means of broadcast network events offered to the network event service. The protocol also supports group-multicast or unicast addressing for those applications requiring only a subset of the connected stations to participate in the service.

Selection of a single “best clock” to serve as grandmaster in a network is performed automatically on the basis of a Grandmaster precedence parameter assigned to each clock entity. The grandmaster precedence parameter encodes accuracy and other features of the clock in such a way that a lower value of the parameter represents a superior choice as master clock.

7.3.3.1 Network event payload

The Global clock protocol defines a specific structure for the payload field of network event primitives exchanged in the protocol. That structure contains the Event Type assigned to the global clock protocol (0x00), an origin timestamp from the grandmaster, and a precedence field used for selection of which clock should adopt the grandmaster role. When the global clock payload is incorporated into a network event frame, the full frame is as follows:

Field	Size octets	Function
Source Address	6	Source MAC
Destination Address	6	Destination MAC (unicast, multicast, or broadcast)
Ethertype	2	0x88F7 assigned to 802.1AS
Function	1	Function code distinguishing Network Event frames
Version	1	Version identifier = 0x00
Event Timestamp	6	Timestamp (sender's timescale)
Cumulative Link Delay	4	Accumulated delay measurement
Payload size	2	= 25 (octet count of Payload)
Payload		
Event Type	1	Assigned value 0x00
Grandmaster Time	10	Origin time stamp (GM timescale)
GM Precedence	14	Precedence for GM selection
Fill (optional size, minimum 7 octets)	7+M	Fill/padding Ignored when Version = 0x00
FCS	4	Frame Check Sequence
	64+M	

Table 7-4. Network event frame carrying global clock payload.

7.3.3.2 Precedence field

The global clock frame contains a GM precedence field. The GM precedence field contains 6 fields totalling 14 octets, with field interpretation as defined in IEEE 1588v2.

	Field	Size octets	Function
MSB	Priority1	1	Primary clock priority
	Class	1	Clock accuracy class
	Time Source	1	Source of time
	Variance	2	Rated variance of clock
	Priority2	1	Secondary clock priority
LSB	Clock ID	8	Unique clock identifier
		14	

Table 7-5. Precedence field.

7.3.3.3 Clock entity behavior

A clock entity may be in one of three states: master, slave, or backoff. The entity generates network event requests only when in master state.

When a global clock entity receives a network event indication message, it processes the contents as follows:

1. If the Event Type does not have the assigned value 0x00, the message is ignored.
2. If the Grandmaster precedence field in the message is inferior to the clock entity's own, the message is ignored.
3. Otherwise (i.e. if the Grandmaster precedence field in the message is equal to or superior to the clock entity's own), the message is considered "qualified"; the clock entity enters (or remains in) slave mode and its clock is adjusted, treating the message as a cross time stamp affiliating the Event Time field of the message with the Grandmaster time field of the payload.

If a global clock entity in the slave state does not receive any qualified network event indications for a timeout period P (currently specified at 300msec), it enters the backoff state and initializes its backoff timer. The random backoff period (currently specified as 0-100ms) is intended to minimize the threat of a "multicast storm" when a current grandmaster goes off-line. If the clock entity still has not received a qualifying network event at the expiration of the backoff period, it enters master state and immediately begins to issue network event request primitives at regular intervals (currently specified at 100msec).

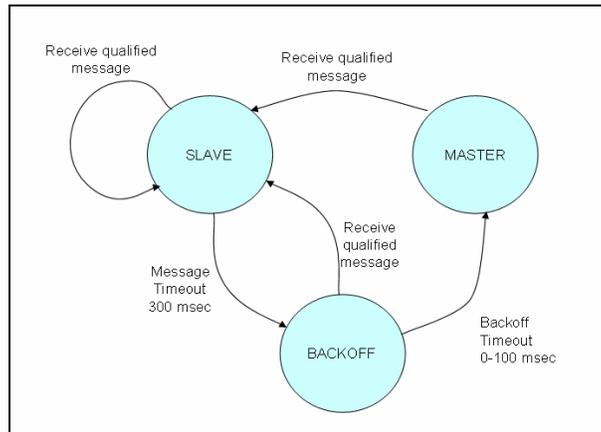


Figure 7-9. Clock Entity state diagram.

7.3.3.4 Slave clock adjustment

A global clock entity in slave state responds to a qualified network event indication by adjusting its internal state. Considering the global clock service interfaces (clause 7.2.2) at both grandmaster and slave, the purpose of the slave clock adjustment is to minimize the deviation between the timescales of the grandmaster and the slave as presented to their respective clients.

This standard does not specify a particular algorithm for slave clock adjustment. As an example, the partner clock calibration method of clause 7.3.1.3 may be adapted to this task, where the Event Time and Grandmaster Time fields of a global clock protocol message i are treated as $partner-point_i$ and $local-point_i$ variables, respectively. At the global clock entity's client interface, an Event.request primitive causes a snapshot to be taken using the local device timescale, the clock entity converts that to the global timescale using the algorithm described for timestamp translation service, and delivers that converted value in a Timestamp.indication primitive.

7.3.4 Selection of the protocol parameters

The primary protocol parameters are

1. the repetition rate for peer clock calibration protocol messages,
2. the filter characteristics applied in the peer clock calibration entities,
3. the repetition rate for global clock update messages, and
4. the filter characteristics applied in the global clock entities.

In combination with the device-dependent characteristics (i.e. timestamp granularity, clock phase noise, clock drift), these parameters determine the achievable system performance in terms of synchronization accuracy, jitter/wander, and settling time.

<<this outlines a design methodology; we need to run the numbers and simulate the results>>

7.3.4.1 Peer clock calibration parameters

The design methodology for this protocol is to assign an error budget to each peer clock calibration entity such that a 7-hop chain will accumulate a predictable maximum error; then global clock entities can be designed to meet particular application requirements based on this maximum error in their update message stream. Also other applications making use of the network event service experience predictable behavior. The peer clock error budget is expressed as an MTIE mask <<TBD>>. The design of the peer clock calibration message rate and filtering algorithm to meet this error budget will rely on the device-dependent characteristics noted above, viz. timestamp granularity, phase noise, clock drift. A standard message rate is associated with the 802.3 and 802.11v protocols defined in this standard <<TBD>>. Nonetheless it is permissible for vendor-specific protocols to deviate from this rate or to negotiate the rate dynamically depending on application requirements, e.g. for power-up initialization, transient recovery, or enhanced accuracy.

7.3.4.2 Global clock update operational parameters

The optimal repetition rate for global clock update messages is independent of the device and media characteristics, provided the error budget assigned at the timestamp translation service is met. Thus a single global clock update rate used throughout a mixed-link-type network is appropriate. The rate specified in this standard <<100msec TBD>> is intended to support media clock reconstruction to the jitter/wander requirements of consumer and professional audiovisual applications.

In some alternative time distribution protocols, increasing the update rate for comparable messages (e.g. 1588 Sync messages) is expected to improve accuracy and dynamic response. Each such message is associated with some timing uncertainty, dependent on the hardware implementation and operating environment; by averaging the effect a large number of messages over a short period the uncertainty can be reduced.

In contrast, in the peer calibration approach, the system timing uncertainty comes primarily from calibration errors in the timestamp translation model, and increasing the rate of global clock messages (i.e. repeating the translation operation more often) has little effect on the accuracy of time transfer.