

---

# 802.1AS Tutorial

Kevin B. Stanton  
Intel Corporation

2008-11-13

[kevin.b.stanton@ieee.org](mailto:kevin.b.stanton@ieee.org)

New edition with equation corrections

# Abstract

---

**This presentation provides an overview of clock and time synchronization across a heterogeneous bridged LAN using IEEE p802.1AS**

# Agenda

---

- **802.1AS goals / use cases**
- Overview
- Grand Master selection
- Time propagation
  - Media-independent
  - Media-dependent

# 802.1AS goals

---

- **Accurate**
  - 802.1AS time at any two LAN-attached stations accurate to within +/- 500ns
    - If one station thinks it's 09:57:55 AM PST, all others agree to within 500ns
    - Assume 7 network hops
    - Assume 100PPM crystals, 1PPM / second maximum frequency drift
  - Media-transport clocks can be ported using the 802.1AS clock with little or no averaging/filtering
    - E.g. S/PDIF, AES3, HDMI, etc.
- **Plug and play**
  - Grand Master clock is selected automatically
    - New GM communicates its time & frequency offset relative to previous GM
    - If GM experiences step change, communicates the magnitude
    - New clock stabilizes in a fraction of a second
  - Clock tree reconfigures automatically
- **One clock for the entire LAN**
  - LAN-agnostic architecture
    - Follows IEEE 1588-2008 with lots of options removed, simplifications
  - LAN-specific measurements
    - 802.3 measurements follow IEEE 1588-2008
    - 802.11 measurements follow IEEE 802.11v "TimeMeasurement"
    - Annex describes measurements for Coordinated Shared Networks (CSNs)

# These are **NOT** 802.1AS goals

---


- **NO!:** Improve the latency of media packets
- **NO!:** Improve the delivery jitter of media packets
- **NO!:** Improve latency or delivery jitter of 802.1AS packets
- **NO!:** Provide clock which all media must be synchronized to
- **NO!:** Solve time synchronization over 802.3 only (or .11 only)

# Wander (MTIE)

## Detail of End-to-End Jitter Application and Wander Requirements

- Uncompressed SDTV (SDI signal)
- - - Uncompressed HDTV (SDI signal)
- - - MPEG-2, after netwk transport (Ref. Pts. D and E)
- - - MPEG-2, no netwk transport (Ref. Pts. B and C)
- · - Digital Audio, Consumer Interfaces (S/P-DIF)
- · · Digital Audio, Professional Interfaces (AES3)

Network Interface MTIE Masks for Digital Video and Audio Signals



<http://www.ieee802.org/1/files/public/docs2006/avb-garner-avb-requirements-summary-r2-060215.pdf>

# Jitter (UI = Unit Interval)

## Detail of End-to-End Jitter Application and Wander Requirements

Requirement	Uncompressed SDTV	Uncompressed HDTV	MPEG-2, with network transport	MPEG-2, no network transport	Digital audio, consumer interface	Digital audio, professional interface
Wide-band jitter (UIpp)	0.2	1.0	50 $\mu$ s peak-to-peak phase variation requirement (no measurement filter specified)	1000 ns peak-to-peak phase variation requirement (no measurement filter specified)	0.25	0.25
Wide-band jitter meas filt (Hz)	10	10			200	8000
High-band jitter (UIpp)	0.2	0.2			0.2	No requirement
High-band jitter meas filt (kHz)	1	100			400 (approx)	No requirement
Frequency offset (ppm)	$\pm 2.79365$ (NTSC) $\pm 0.225549$ (PAL)	$\pm 10$	$\pm 30$	$\pm 30$	$\pm 50$ (Level 1) $\pm 1000$ (Level 2)	$\pm 1$ (Grade 1) $\pm 10$ (Grade 2)
Frequency drift rate (ppm/s)	0.027937 (NTSC) 0.0225549 (PAL)	No requirement	0.000278	0.000278	No requirement	No requirement

SAMSUNG Electronics

IEEE 802.1 AVB TG / IEEE 1588  
February 21, 2006

10

<http://www.ieee802.org/1/files/public/docs2006/avb-garner-avb-requirements-summary-r2-060215.pdf>

# Use cases

---

- **Distributed, networked audio**
  - Left/right stereo sync for consumer “WiFi speakers”: +/-10us
  - AES11: 5% of sample period → +/- 0.5us time delta
  - [Note: This is not one-hop budget, but full LAN path budget]
- **Media clock porting**
  - Satellite receiver streaming content over home LAN
    - Must match rates – cannot tell satellite to slow down/speed up
- **Studio: House clock or GenLock**
- **Other uses of accurate time synchronization**
  - Manufacturing, instrumentation, telecommunications, etc.



# Agenda

---

- 802.1AS goals / use cases
- **Overview**
- **Grand Master selection**
- **Time propagation**
  - Media-independent
  - Media-dependent

# High level architecture

---

- **Grand Master selection**
  - GM-capable stations advertise that fact via ANNOUNCE messages
  - If station hears from station with “better” clock, does not send ANNOUNCE
    - Settable “Priority” field can override clock quality
    - MAC address is tie breaker
  - Bridges drop all inferior ANNOUNCE messages
    - Forward only the best
  - Last one standing is Grand Master for the LAN
    - GM is the root of the 802.1AS timing tree
    - GM periodically sends the current time
- **Propagation of time**
  - Bridges in the tree propagate time toward the leaves
    - Taking queuing delay into account (aka “Residence Time”)

# 802.1AS is not alone

---

- **802.1AS is one of three 802.1 AVB draft standards**
- **802.1AS over Ethernet (802.3) qualifies as a Profile of IEEE 1588-2008**
  - 802.1AS simplifies 1588 and defines synchronization over other media (e.g. 802.11)

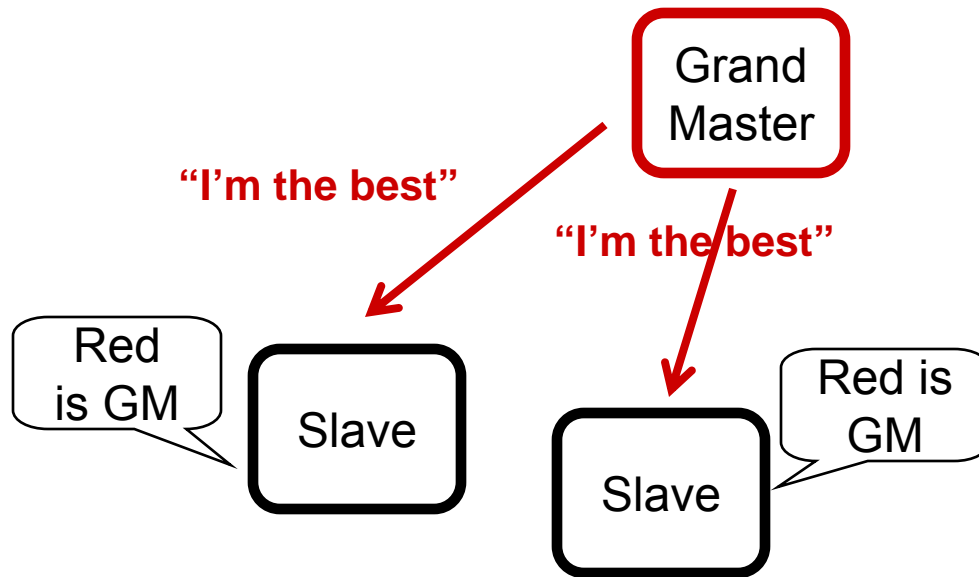
# Agenda

---

- 802.1AS goals / use cases
- Overview
- **Grand Master selection**
- Time propagation
  - Media-independent
  - Media-dependent

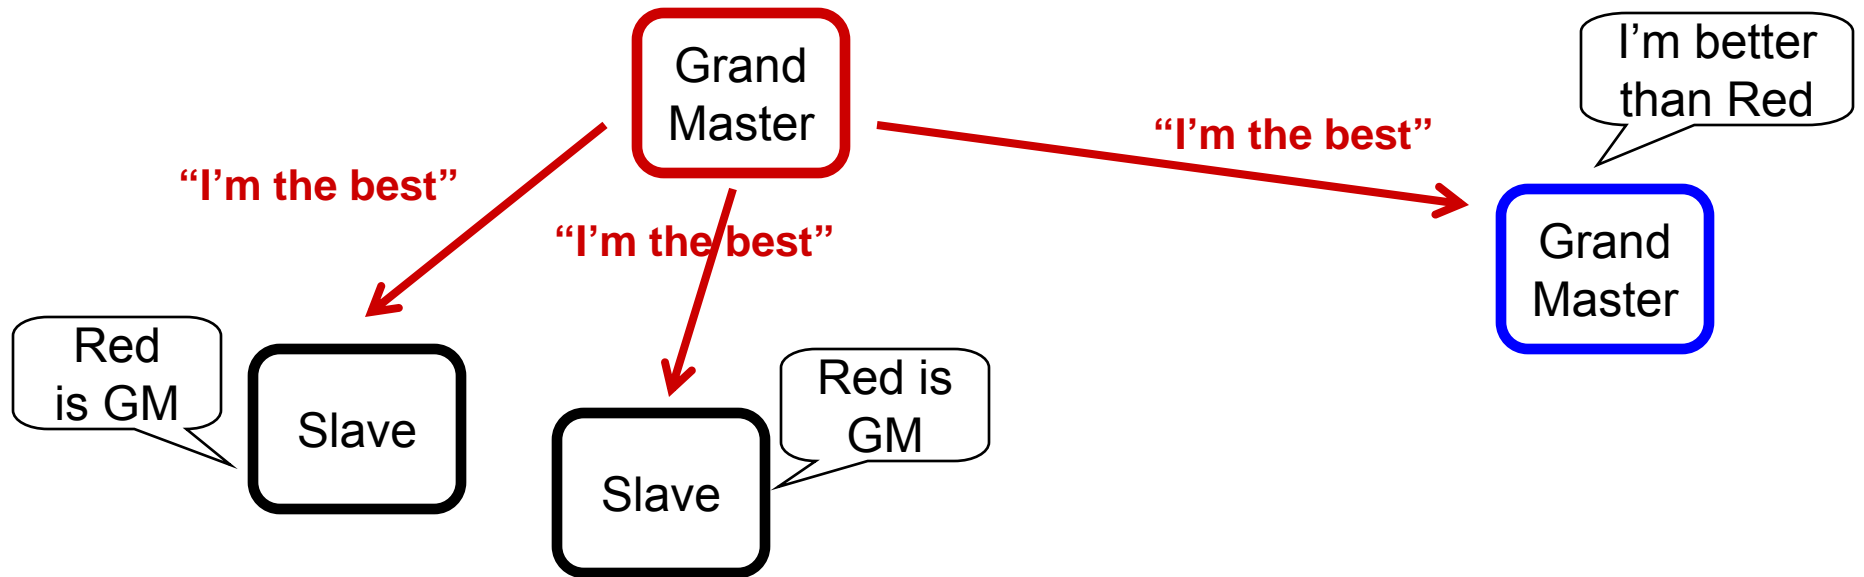
# Grand Master selection – Steady state

---



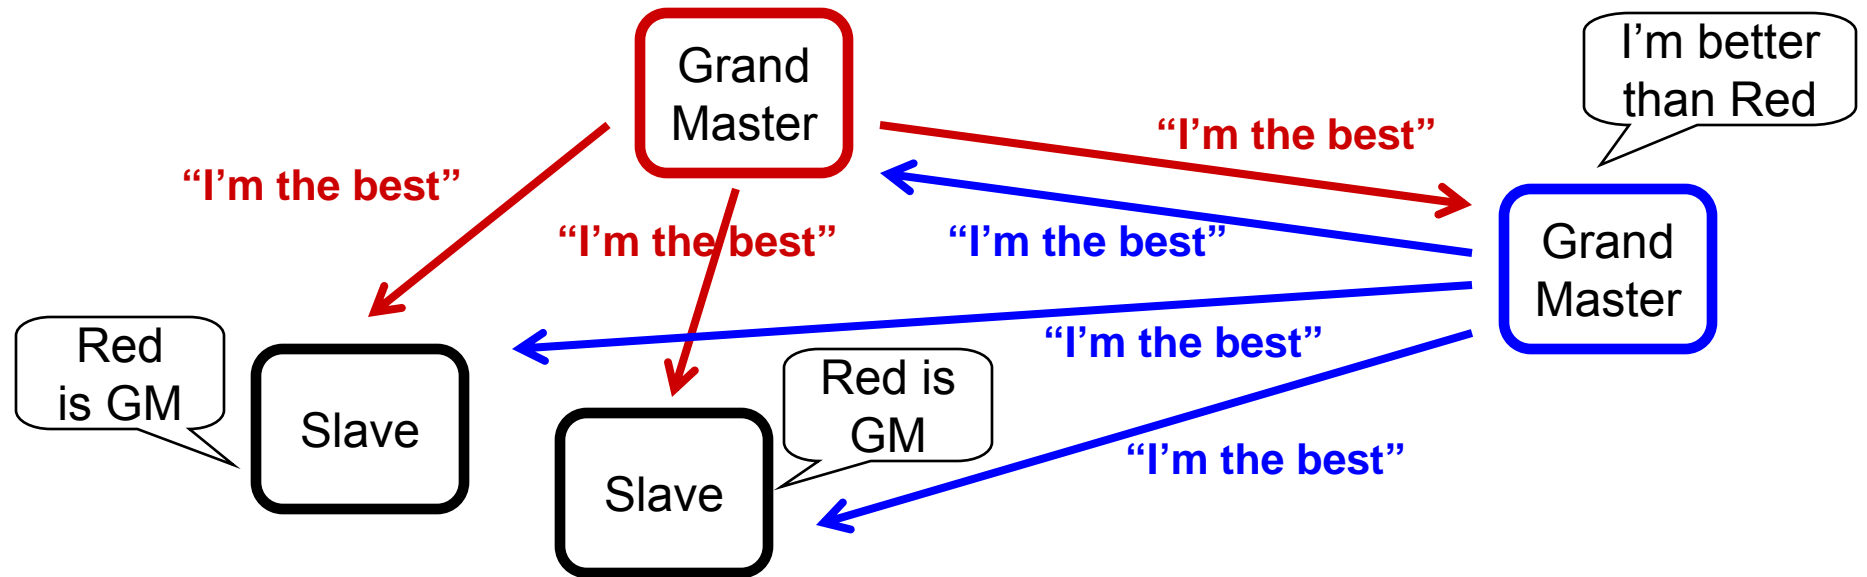
- **Steady state: The GM sends ANNOUNCE**
- **Everyone determines best ANNOUNCE info (including self)**
  - ...using simple bitwise compare
- **Everyone knows their role**
- **Life is good**

# Grand Master selection – New, better GM



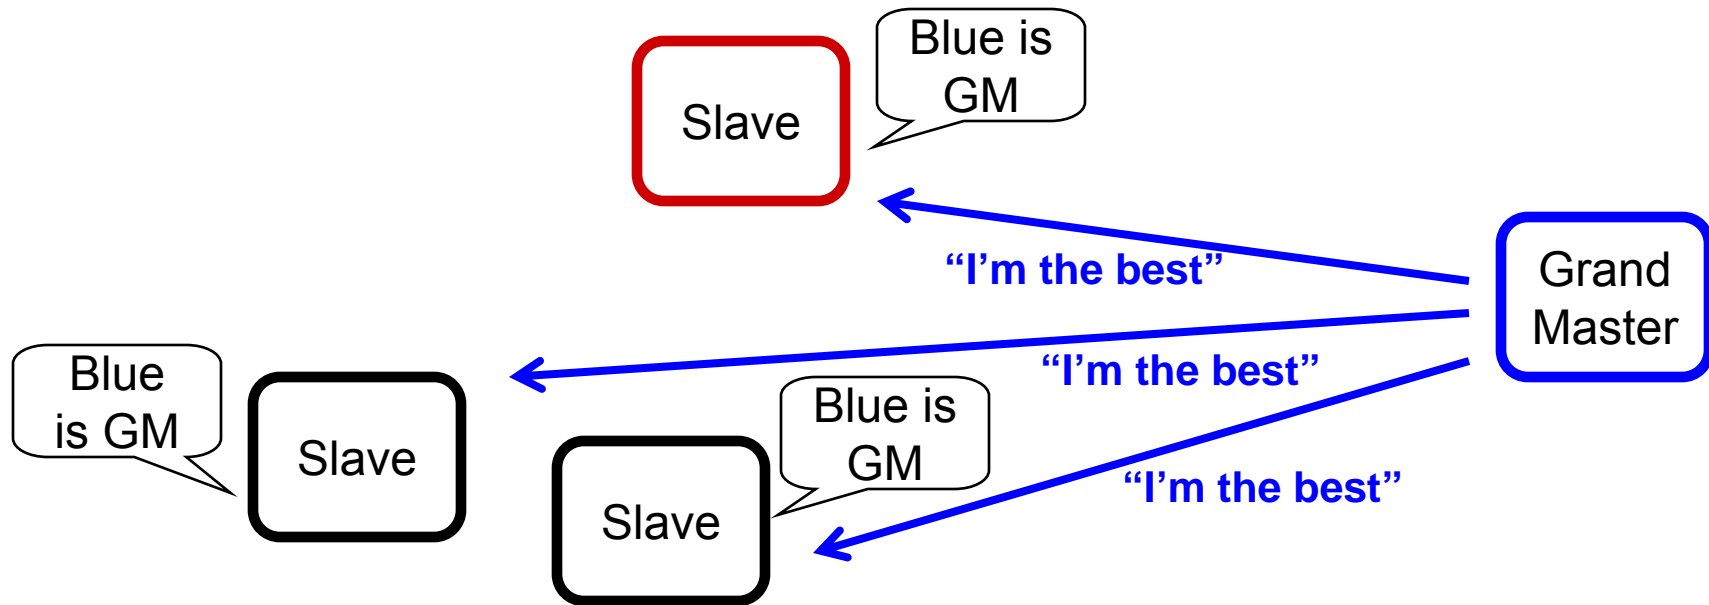
- **Blue station with better clock appears**

# Grand Master selection – New, better GM



- **Blue station with better clock appears**
- **Blue sends ANNOUNCE**
- **If stations receive multiple ANNOUNCE messages, ignore the inferior ones**

# Grand Master selection – New, better GM

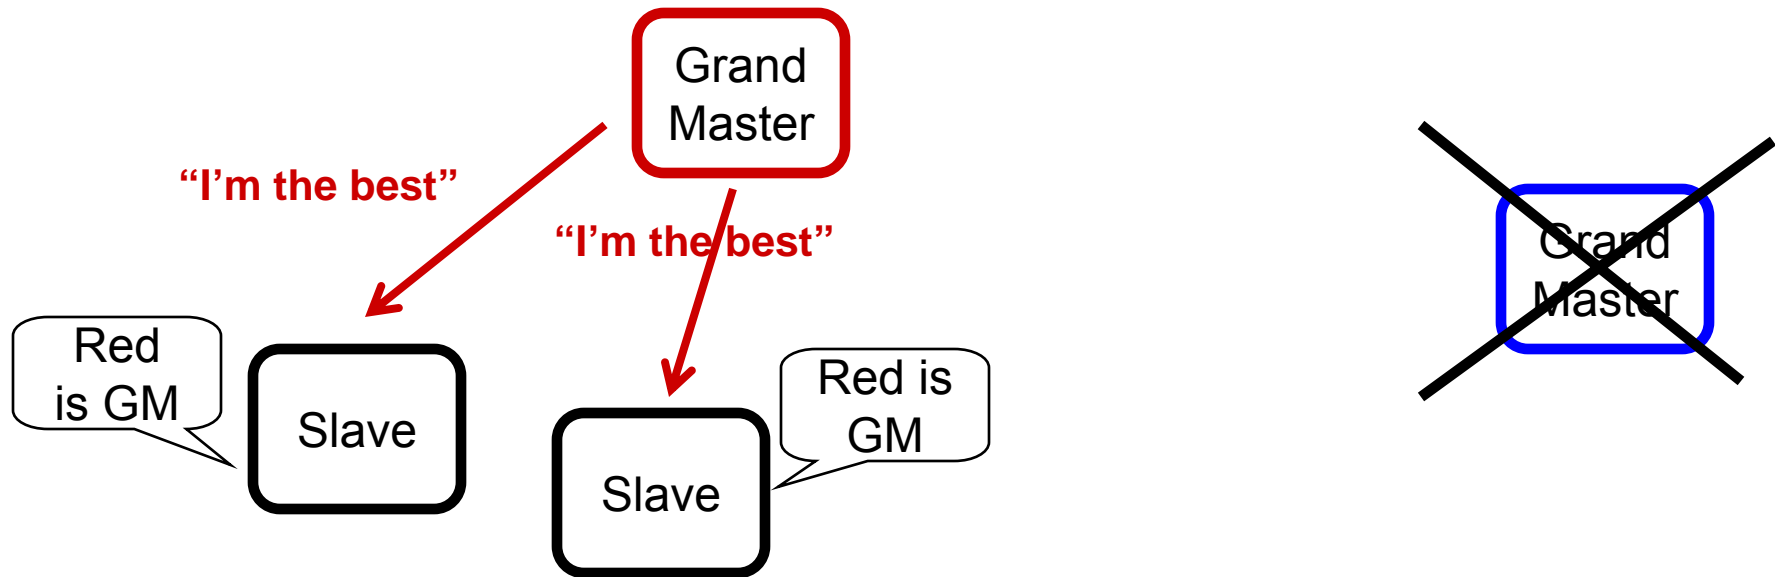


- **Blue station with better clock appears**
- **Blue sends ANNOUNCE**
- **If stations receive multiple ANNOUNCE messages, ignore the inferior ones**
- **Eventually the blue station is the undisputed GM**



# Grand Master selection – Lost GM

---



- If **Blue** disappears, all **GM-capable** stations send **ANNOUNCE** if they believe themselves to be best
- Eventually, only **GM** sends **ANNOUNCE**

# ANNOUNCE comparison

---

- **The credentials passed in ANNOUNCE messages are compared (in order of decreasing importance):**
  - Priority (settable by management)
  - Multiple “**quality of my clock**” fields
  - MAC address

# Agenda


---

- 802.1AS goals / use cases
- Overview
- Grand Master selection
- **Time propagation**
  - Media-independent
  - Media-dependent

# Time domain is propagated to the “subnet”

---

- **Bridges [/APs] relay**
  - The best ANNOUNCE
  - The GM’s time
- **Links can be**
  - 802.3 Ethernet
  - 802.11 WiFi
  - Almost anything
- **Let’s look at one bridge**




# Grand Master selection – Bridges help decide




- **Bridges filter inferior ANNOUNCE (an optimization)**
  - Bridges forward best ANNOUNCE messages
  - Bridges drop all inferior ANNOUNCE messages
  - Benefits:
    - Reduction in network traffic
    - Potentially faster GM selection

# Propagating time through a bridge



- **Time passed from master over a link with frames that are LAN/media-specific**
  - Illustrated abstractly here by a Blue frame
- **Bridge measures how the long Blue frame is in bridge**
  - Called "Residence Time"
- **Also measures & compensates for link delay and rateRatio**

# Example linkDelay measurement




$$\begin{aligned} \text{linkDelay} &= \left( \frac{(t1' - t1) - (t2' - t2)}{2} \right) \\ &= \left( \frac{(25 - 0) - (5005 - 5000)}{2} \right) \\ &= 10 \end{aligned}$$

**If link delay is symmetrical, link delay is 10**

**Note: one *could* use measure neighborRateRatio and linkDelay using the same packets (802.11 does, 802.3 does not)**

# Example neighborRateRatio measurement



$$\begin{aligned} \text{neighborRateRatio} &= \frac{(t1' - t1)}{(t2' - t2)} \\ &= \frac{(10000 - 0)}{(60002 - 50000)} \\ &\approx 0.998 \end{aligned}$$

→ neighbor is slow by 2000 PPM

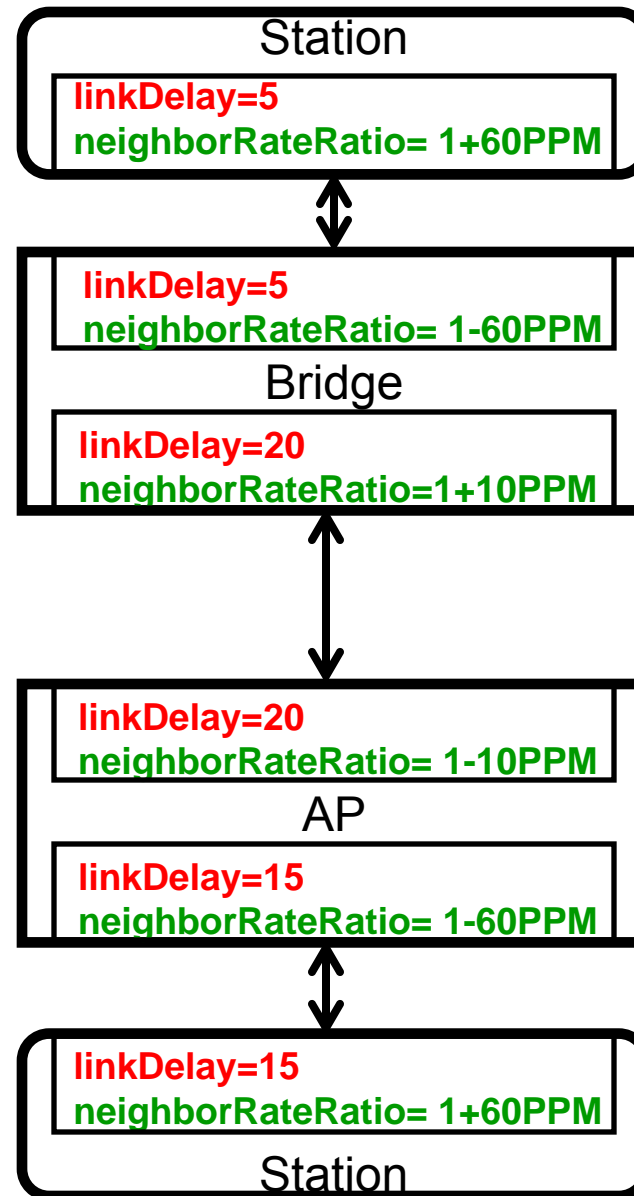
**If link delay is constant, station on left is running 0.02% slower than station on the right**



# Link Delay and Neighbor-Rate-Ratio

## All ports now know

- linkDelay to neighbor
- neighborRateRatio
  - E.g. 1.000060 or 1 + 60PPM



Note: No GrandMaster is needed here

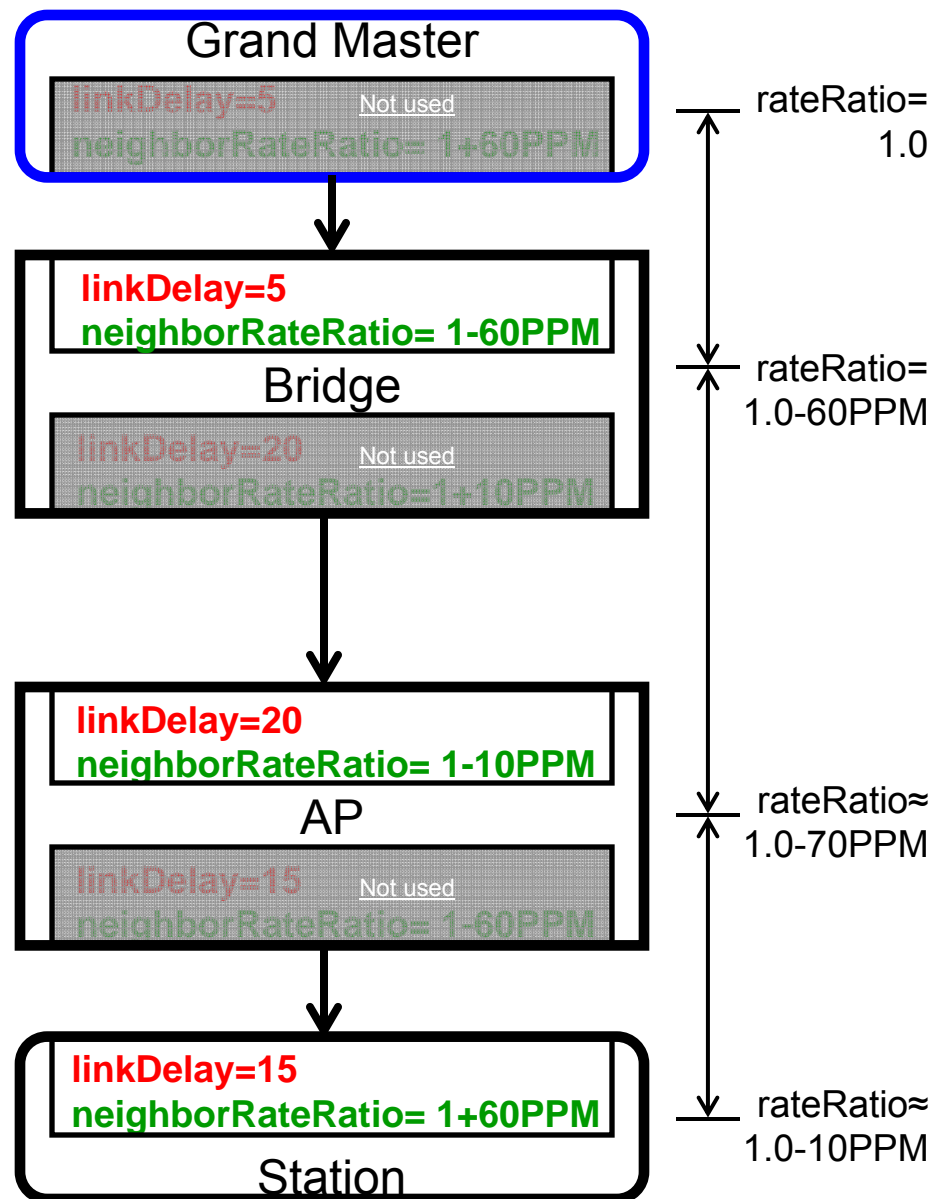
# End-to-end Rate Ratio

When a GM sends “time is X now”, neighborRateRatios accumulated at each hop  
 ...using the equation:


$$\text{rateRatio} += (1 - \text{neighborRateRatio})$$

[It starts out equal to 1.0 in the GM]

Why? When new GM, time stabilizes very quickly



# End-to-end time synchronization



**Grand Master initiates timing message M1 every Sync Interval**

**Each bridge/AP computes delay experienced by M1 in M2:**

delay' = prop' + residenceTime' X rateRatio' (rateRatio is described in a previous slide)

**Slave computes: currentTime = now plus the accumulated delay**

Actually, a 802.3 M1 carries original 'now' and delay separately due to IEEE 1588 legacy; 802.11 carries only their sum

# Measurements summarized

---

**Accurate local egress and ingress timestamps required in the MAC/PHY**

**From the timestamps, the following are computed:**

- **Link delay (per slave port)**
- **Next-neighbor rate ratio (PPM offset to link partner)**

**And timing measurement propagation results in measurement of or computation**

- **Residence time (per transmitted time measurement frame)**
- **rateRatio (PPM relative to the GM)**

**These measurements described in 802.1AS/D4.0:**


- Clause 11: point-to-point full duplex links (i.e. Modern Ethernet)
- Clause 12: IEEE 802.11 links

# Agenda

---

- 802.1AS goals / use cases
- Overview
- Grand Master selection
- **Time propagation**
  - Media-independent
  - Media-dependent
    - 802.11 links
    - 802.3 links
    - Other CSN links

# The 802.11v TIMINGMSMT protocol



**First exchange takes a measurement**

**Subsequent exchanges also pass timestamps from prior measurement**

**Free-running counter used for timestamps**

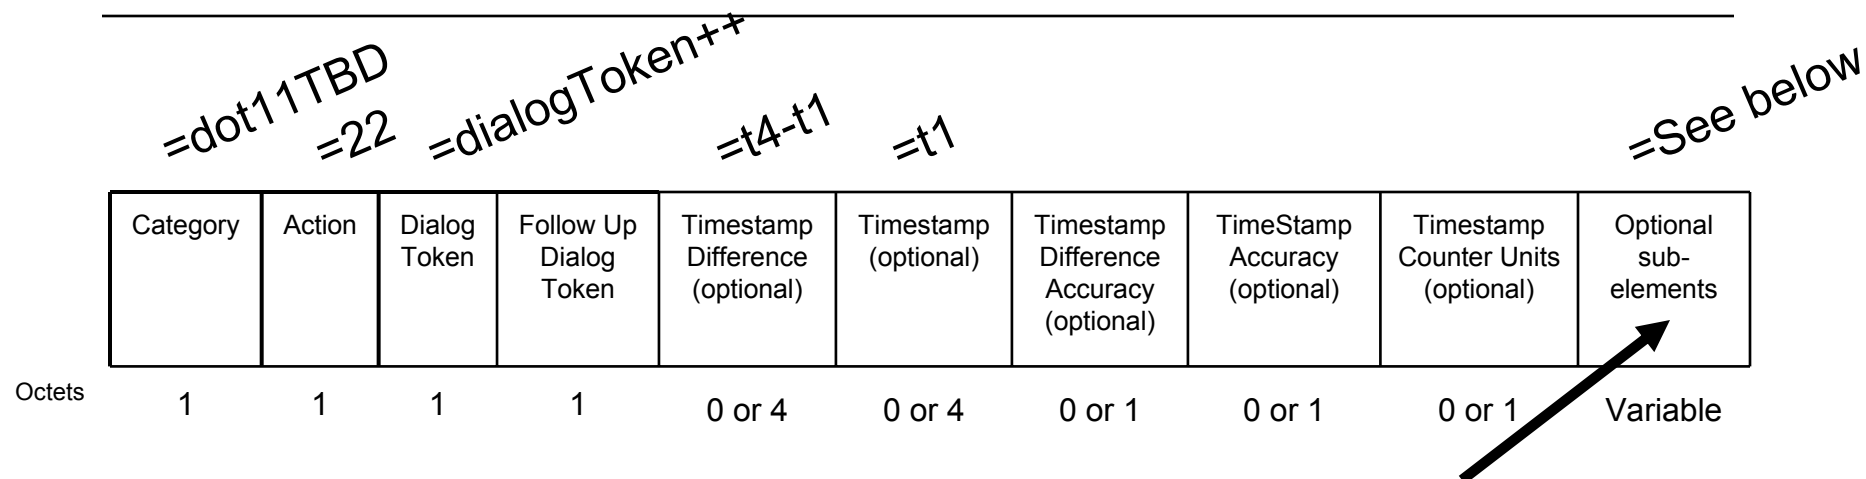
**Allows us to compute:**

$$\text{neighborRateRatio} = \frac{(t1' - t1)}{(t2' - t2)}$$

$$\text{linkDelay} = \frac{[(t4 - t1) - (t3 - t2)]}{2}$$

**[note: rateRatio is also applied]**

# The 802.11v TIMINGMSMT action frame



Optional Sub-element field set equal to “Vendor Specific” (SubElement ID = 221), carries:

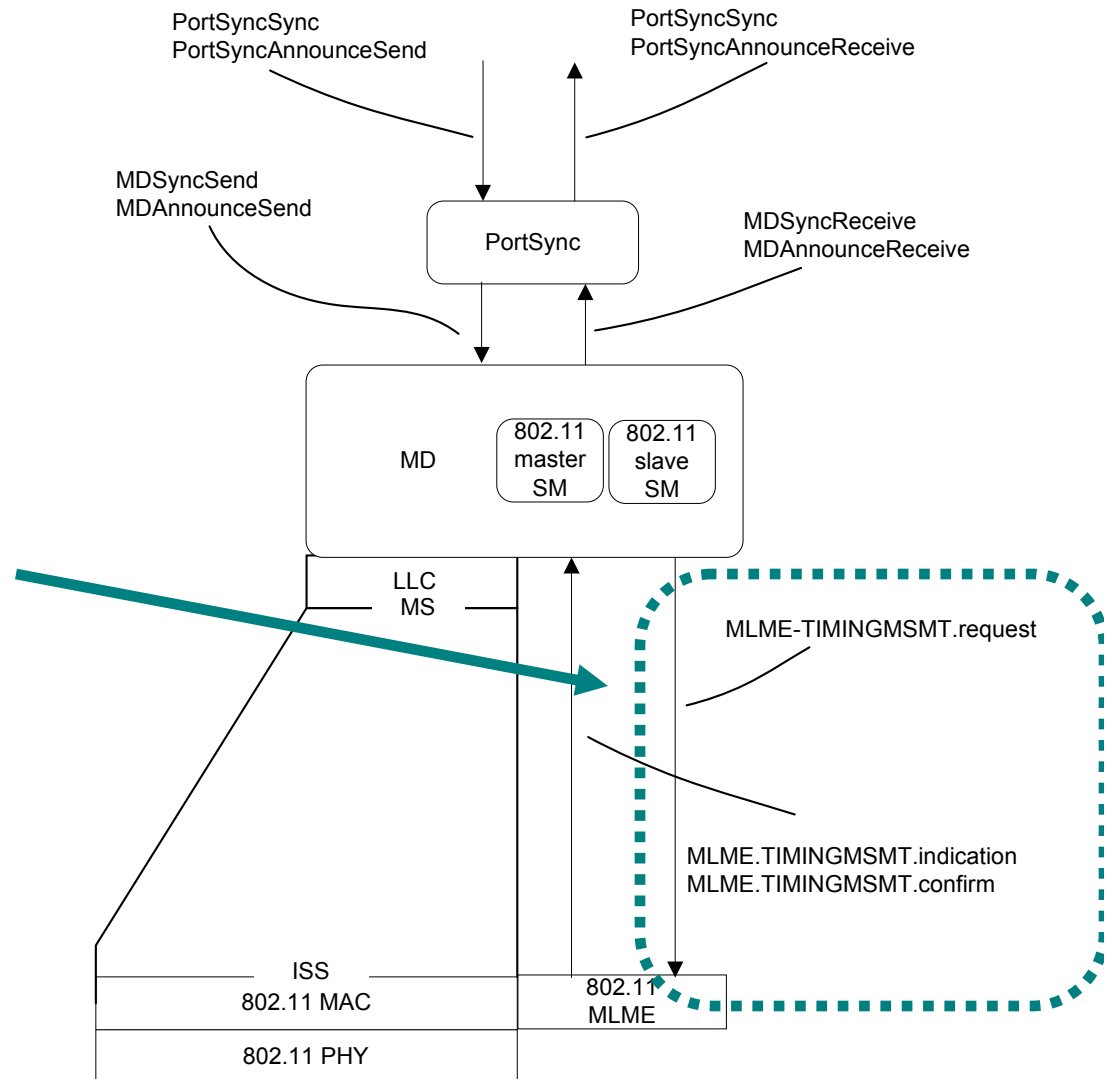
1. **sourcePortIdentity** (10B)
2. **logMeanMessageInterval** (1B)
3. **preciseOriginTimestamp** (10B) --AKA “Now”
4. **rateRatio** (?B -- ‘double’)
5. **gmTimeBaseIndicator** (2B)
6. **lastGmPhaseChange** (12B)
7. **lastGmFreqChange** (4B)

# 802.11 links in the 802.1AS architecture

**PortSync is per-port,  
media-independent**

**MD is per-port media  
dependent**

**802.11v defines these  
MLME.TIMINGMSMT  
primitives**






# 802.3 protocol (step 1 of 2)

## Measure link delay:

1. Requester schedules PathDelayReq for transmission
  - Using master's free-running clock
2. As it passes out the PHY, t1 is captured
  - Using master's free-running clock
3. Time t2 captured as passes from PHY to MAC
  - Using slave clock
4. Responder schedules PathDelayResp for transmission
  - Using slave clock
5. Timestamps t3 and t4 captured
  - Using local free-running clock
6. PathDelayRespFollowup carries t2 and t3 to requester

If link delay is fixed & symmetric:

$$\text{link\_delay} = \frac{[(t4-t1) - (t3-t2)]}{2}$$



# 802.3 protocol (step 2 of 2)


## Synchronize clocks

1. Master schedules Sync for Tx
2. As it passes out the PHY, t1 captured
  - Using master's free-running clock
3. Time t2 captured as passes from PHY to MAC
  - Using slave clock
4. FOLLOWUP carries t1 to slave

If link delay is fixed & symmetric:

Slave's clock offset  
=  $t2 - t1 - \text{link\_delay}$

Note: APs & bridges do this too, communicate the 'residence time' per Sync in the Followup frame



# Coordinated Shared Network (CSN) protocol


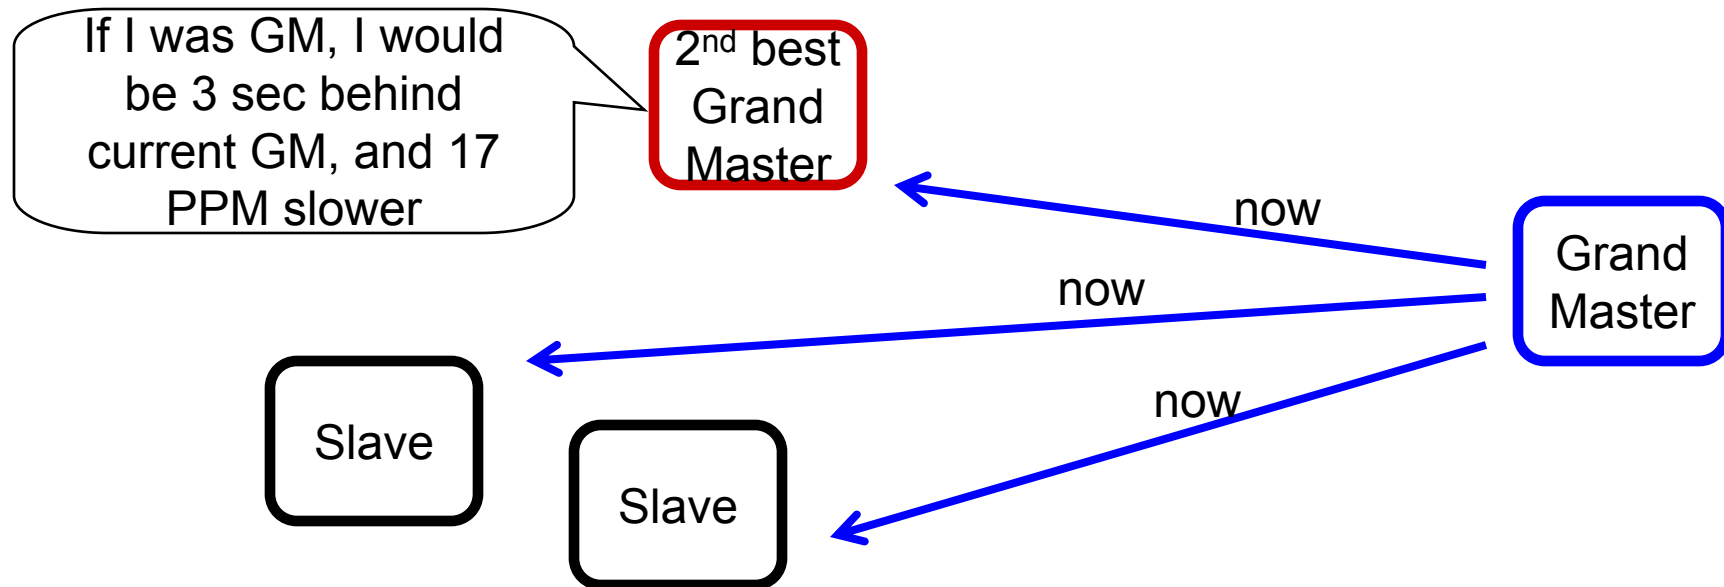


Figure 7: Sync Messages Propagation with CSN Network Reference Clock

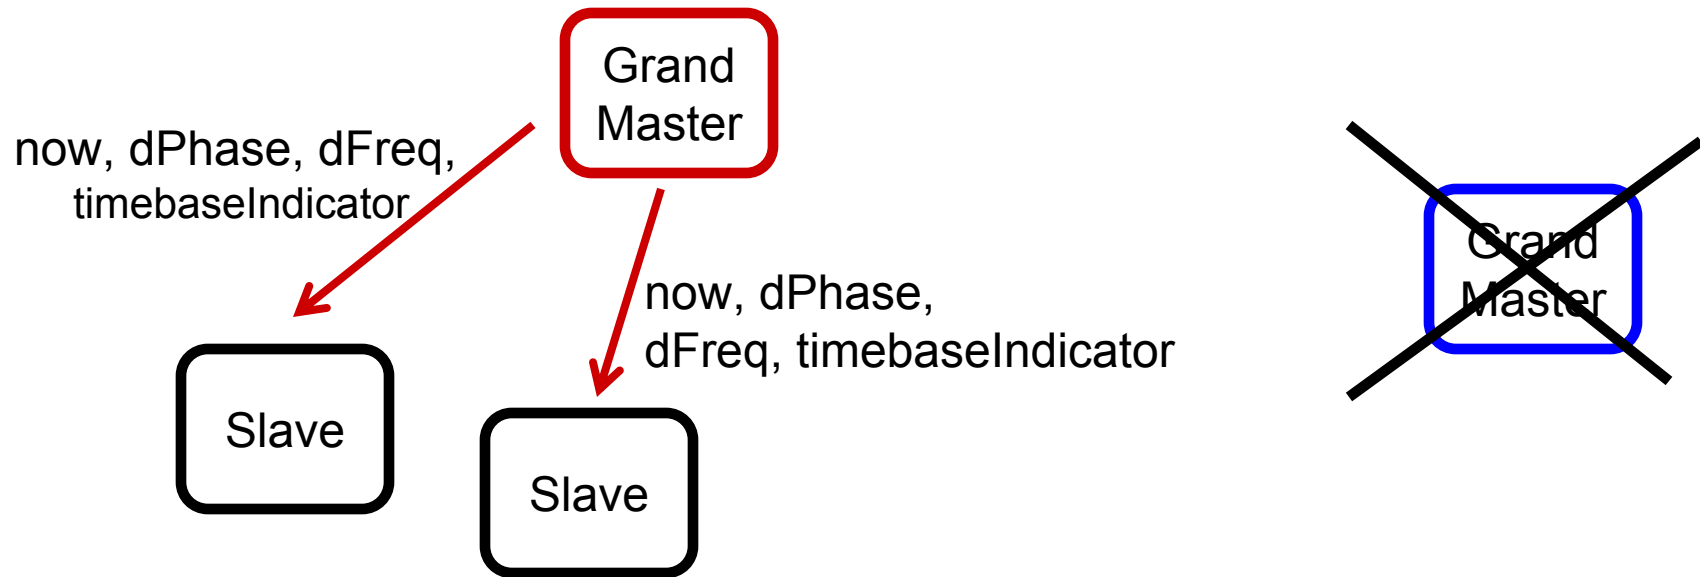
- The above taken from current proposal for .1AS Annex
  - <http://www.ieee802.org/1/files/public/docs2008/as-phkl-csn-appendix-0809-d03.pdf>
- MoCA's 2.0 requirements include "should support" AVB

# Handling of discontinuities



- **All potential GMs know their local time offset and natural frequency offset (rateRatio) from the current GM**
- **Time offset can be close to zero (unless know real time from GPS or other)**

# Handling of discontinuities – new GM



- **New GM always communicates extra info with ‘now’**
  - TimebaseIndicator (increments each discontinuity)
  - Last phase change
  - Last frequency change
- **GM also does this if it experiences a step change**
  - E.g. GPS regains lock or user manually sets time in the GM

# Relating 802.1AS time to media clock

---

- **“Presentation Time” placed in media packet header**
  - Time is from 802.1AS
  - Three proposed methods
    - Layer-2: IEEE 1722 (based on IEC-61883)
    - Layer-3/RTP: IEEE 1733
    - Layer-3/TCP/HTTP: TBD – need simple addition to DLNA
- **Samples or video held in client until Presentation Time**
- **Two consecutive media frames are sufficient for determining the proper media rate**
  - As opposed to a few minutes today—faster resume time, lower audio jitter, lower “play” latency, less buffering

---

# **B A C K U P**

# **Differences in functionality between .11 & .3**

---

- **Fractional nanoseconds missing from .11**
- **Master 802.11 port doesn't know LinkDelay**
- **Multipath may increase timestamp error variation**



# Vendor specific extensions

---

## Use an 802.1 OUI for the 802.1AS-specific TLV

### 7.3.2.26 Vendor Specific information element

The Vendor Specific information element is used to carry information not defined in this standard within a single defined format, so that reserved information element IDs are not usurped for nonstandard purposes and so that interoperability is more easily achieved in the presence of nonstandard information. The information element is in the format shown in Figure 92 and requires that the first 3 octets of the information field contain the OUI of the entity that has defined the content of the particular Vendor Specific information element. The length of the information field ( $n$ ) is  $3 \leq n \leq 255$ . The OUI field shall be a public OUI assigned by the IEEE. It is 3 octets in length. The length of the vendor-specific content is  $n-3$  octets.




Figure 92—Vendor Specific information element format

# Notes

---