Stories from the PTP Battlefront

STAC New York – November 2016

Raymond Russell – CTO

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

PTP Architecture Overview

High-quality distribution of precision time

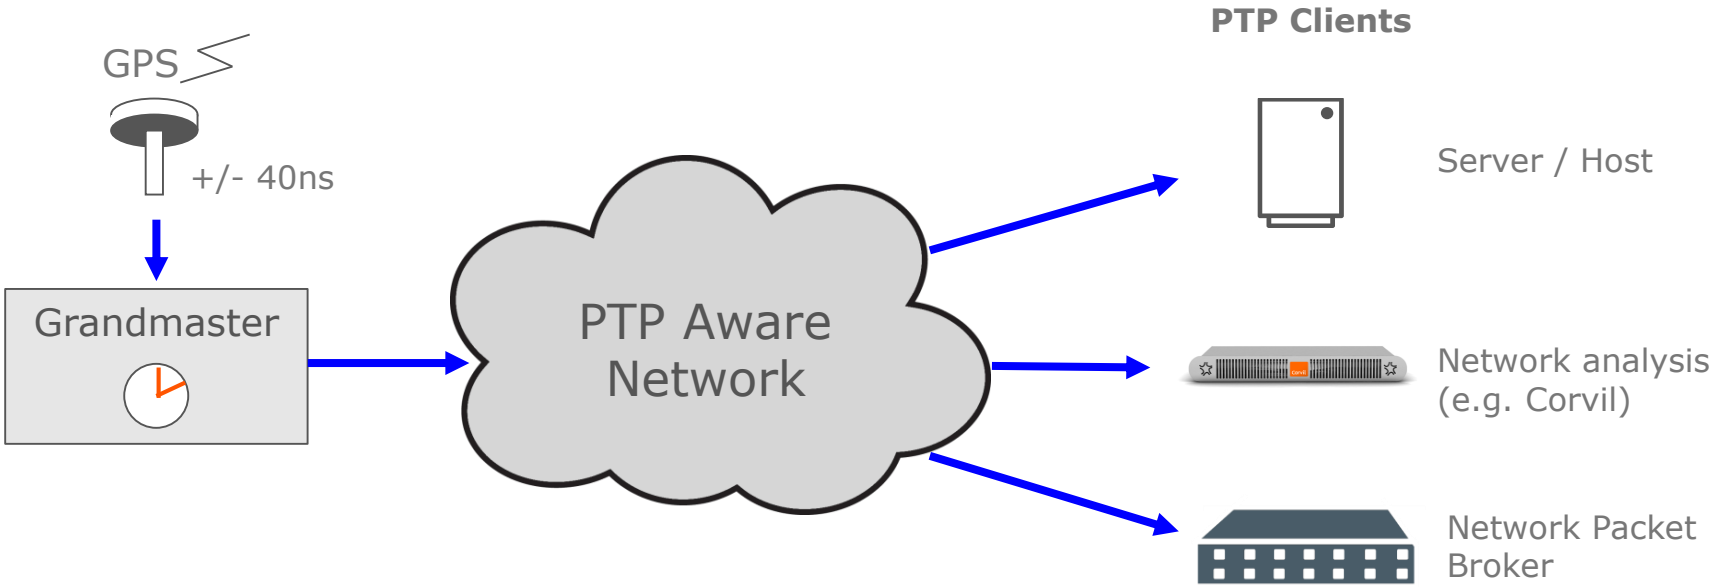

PTP Architecture Review

IEEE 1588-2002 (Version 1)

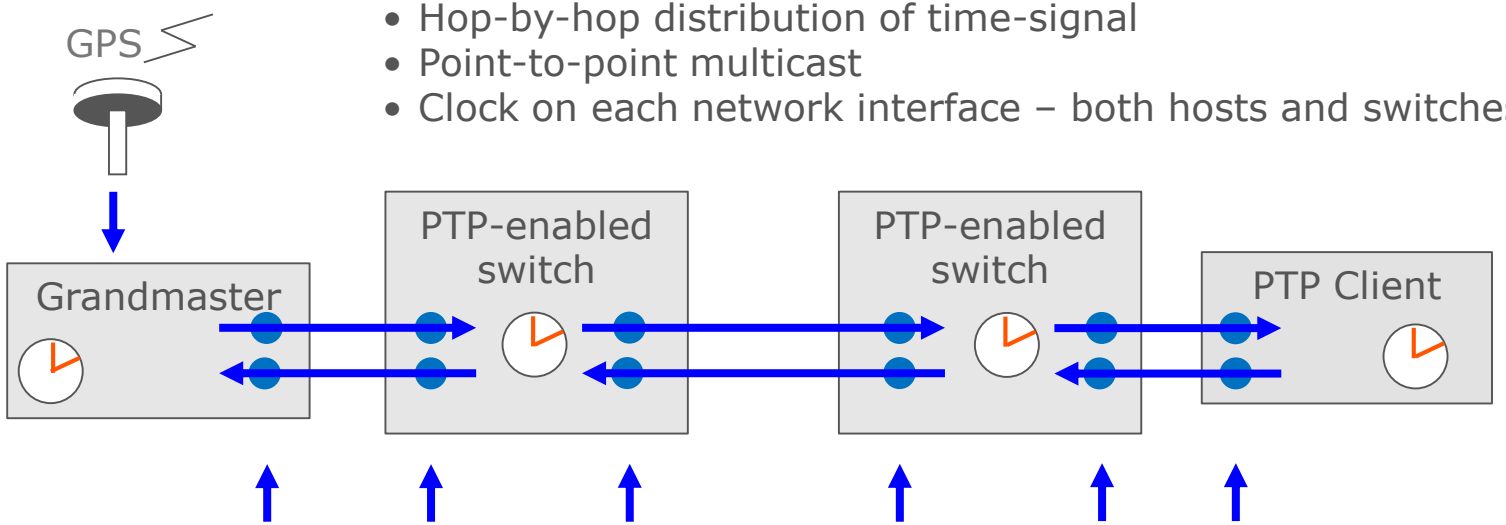

- Hop-by-hop distribution of time-signal
- Point-to-point multicast
- Clock on each network interface – both hosts and switches

PTP Architecture Review

IEEE 1588-2008 (Version 2)

- More transport options – IPv6, Ethernet, etc.; unicast/hybrid mode
- Transparent clocks – timestamps corrected for delay through switch
- Peer-to-peer mode – better scaling

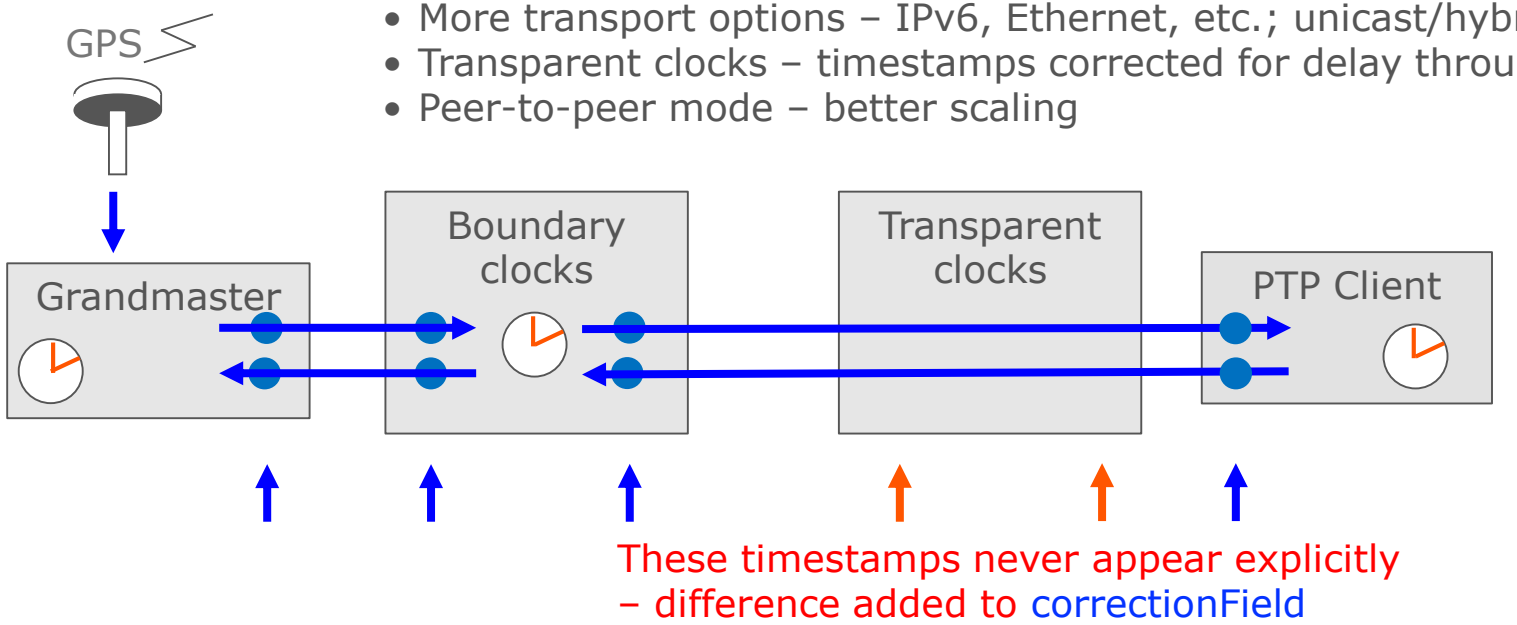

- Most implementations today typically use Boundary Clocks & Multicast
- Discussion today will focus upon BC & MC

PTP Operation

Slave calculates offset from Master

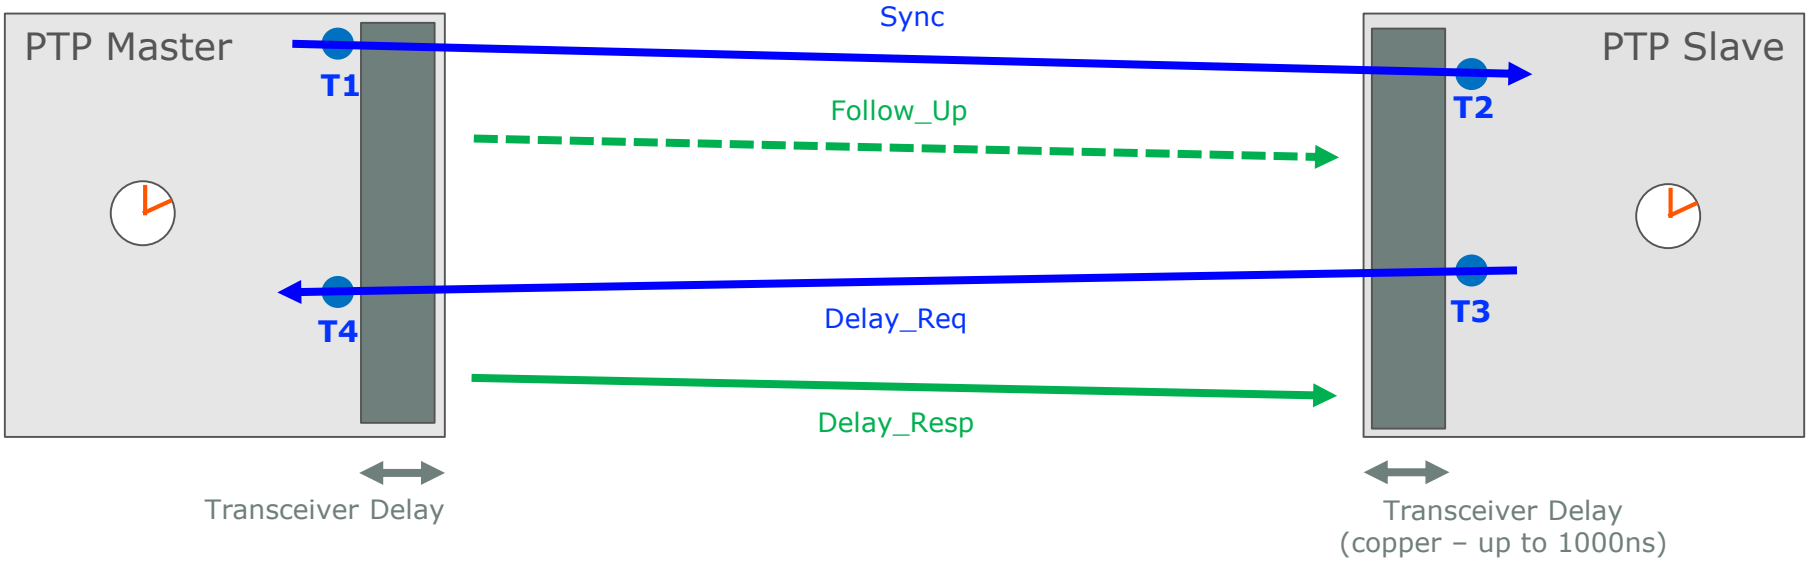

Master to Slave difference = T2-T1
Slave to Master difference = T4-T3

One way latency = (Master to Slave diff+ Slave to Master diff) / 2
Offset = Master to Slave diff – One way latency

Offset = ((T2-T1)-(T4-T3))/2

PTP Operation

Slave calculates offset from Master

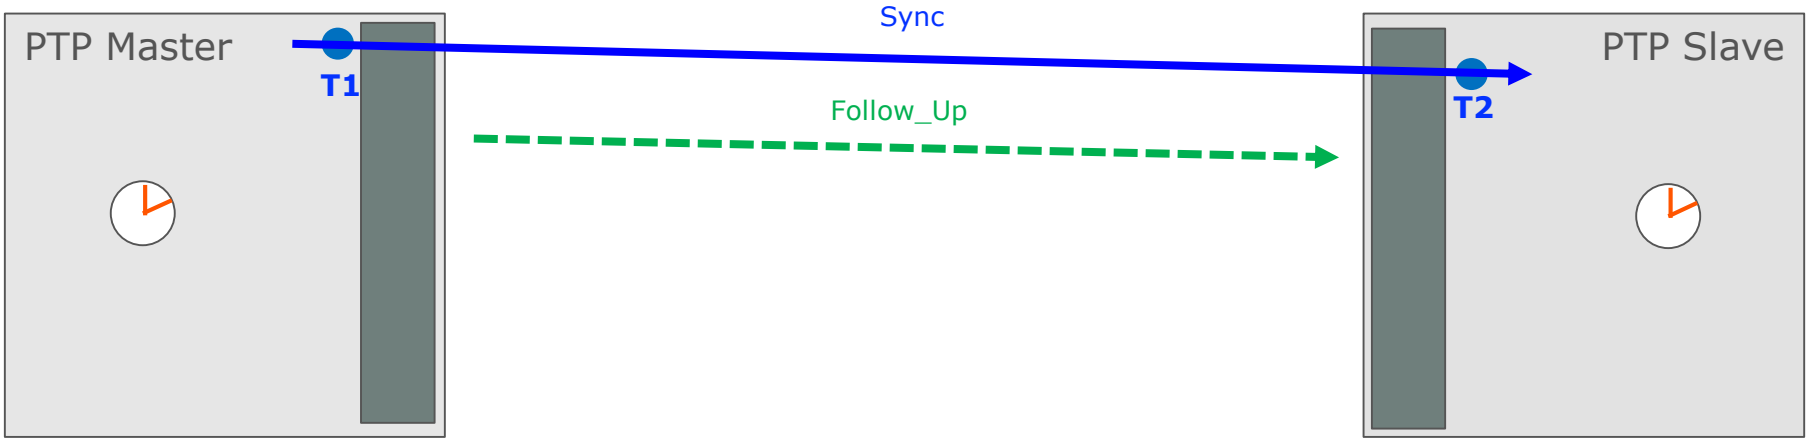

Q. Why the Follow_Up message?

A. Hard to write accurate timestamp into packet as you send it – just send it, timestamp it, and follow up later

Contrast with NTP

Purely host-based software

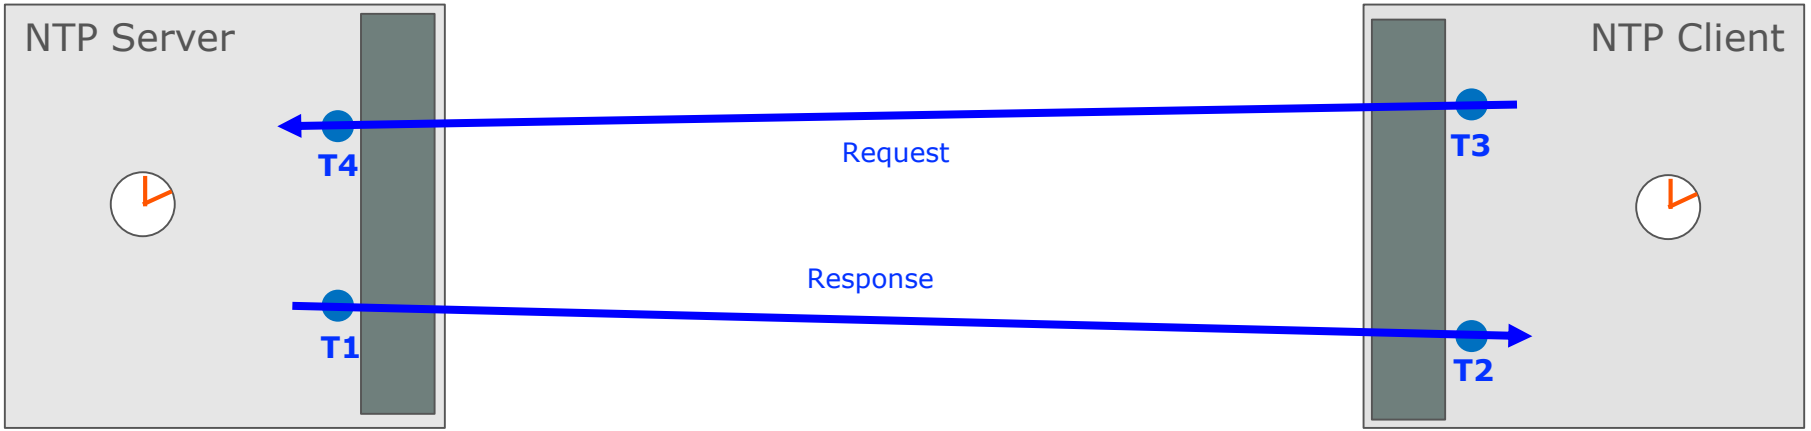

Q. What can go wrong?

A. Jitter makes timestamps stale, inaccurate

- Congestion in the network
- Processing in the host stack

UDP ping-pong between hosts

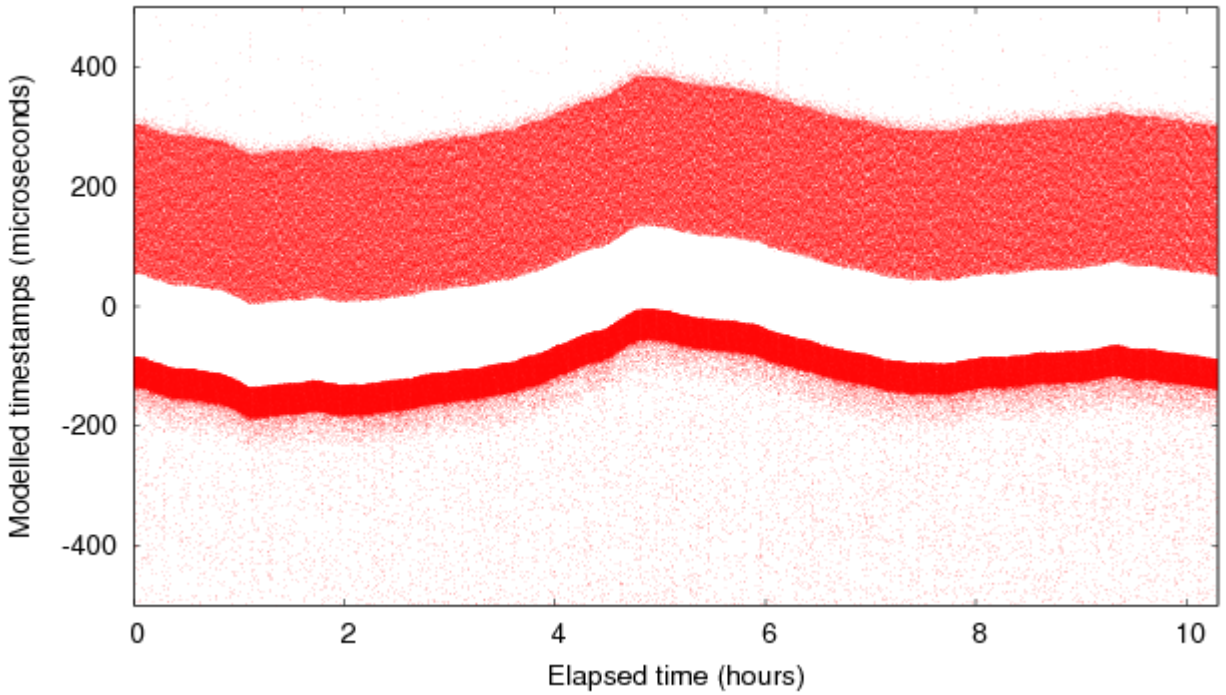

Data points:

- Upper: T4 – T3

  (+client-to-server)
- Lower: T1 – T2

  (–server-to-client)

- Different stack delays in each direction
- Clock wander

- Hardware timestamping on host NICs eliminates noise – easier to correct for offset, drift, wander

PTP Operation

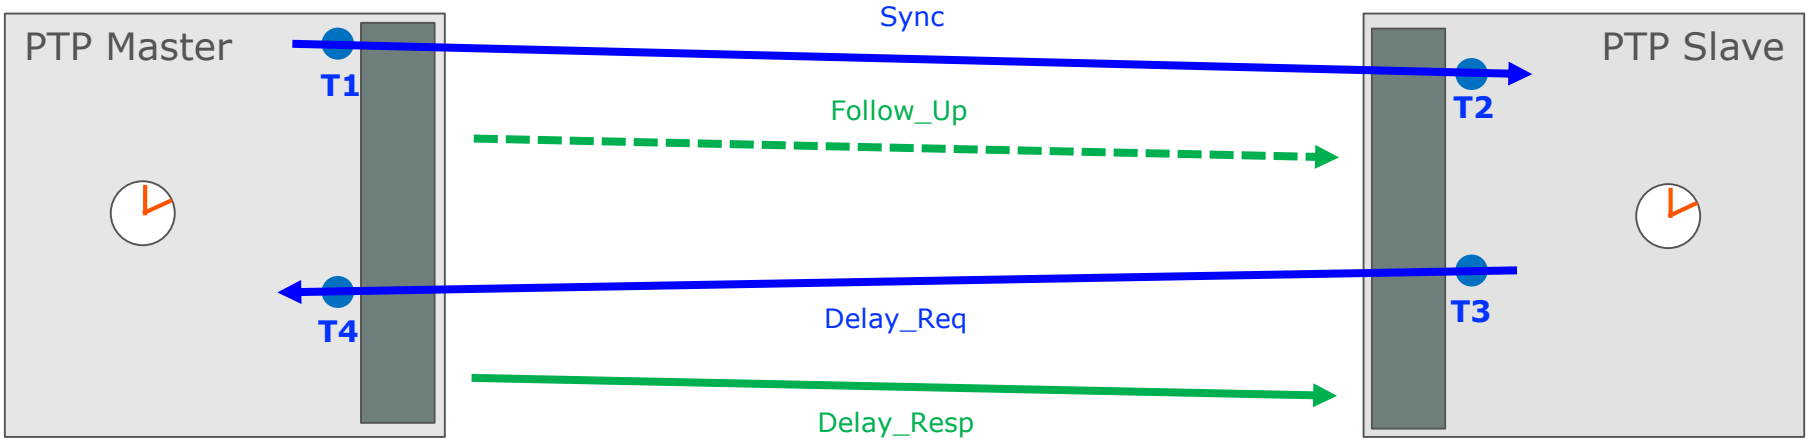

- Hardware timestamping helps eliminate network jitter
- Decoupling of Sync messages from Delay_Req

Optimal PTPv1 Deployment

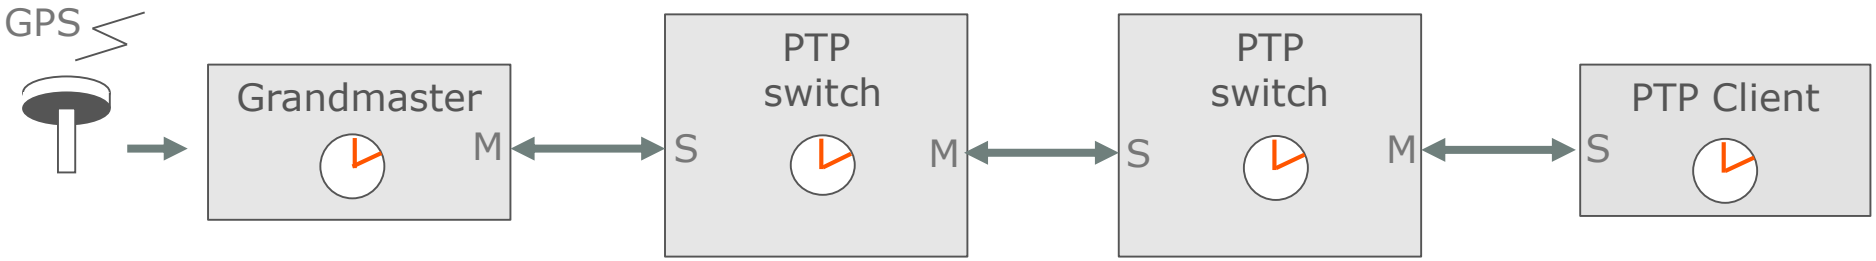

- Point to point Master–Slave relationship across each link
- PTP Switches do not forward PTP multicast messages
- Timestamping performed in HW on interface
  - Offset between Master & Slave is accurately calculated
  - Not affected by other traffic on network link

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

Scenario 1

Client not receiving PTP sync messages

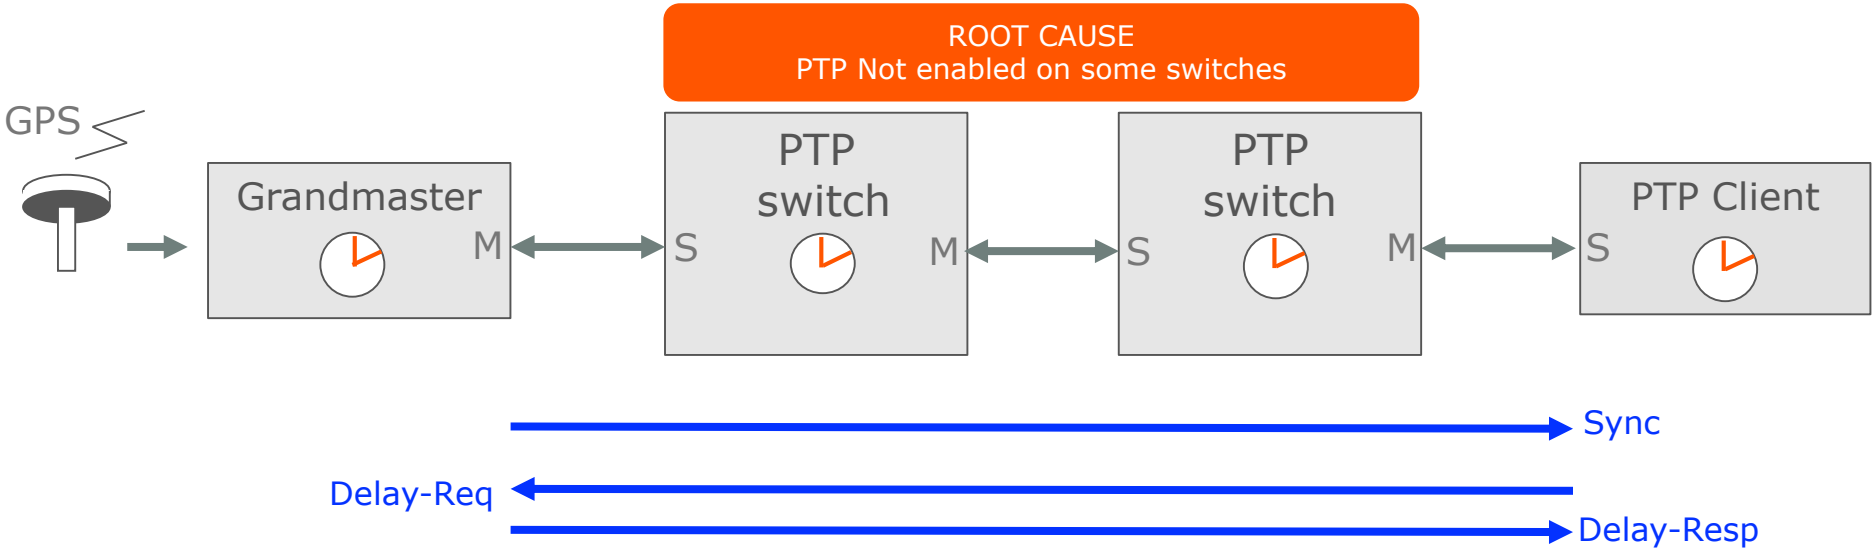

AHA!
PTP uses multicast…

Ok, so let's enable Multicast routing

Seems to be working now…

Or is it?

Scenario 2

Same fundamental issue. PTP not enabled on switches
But client is receiving sync. Why?

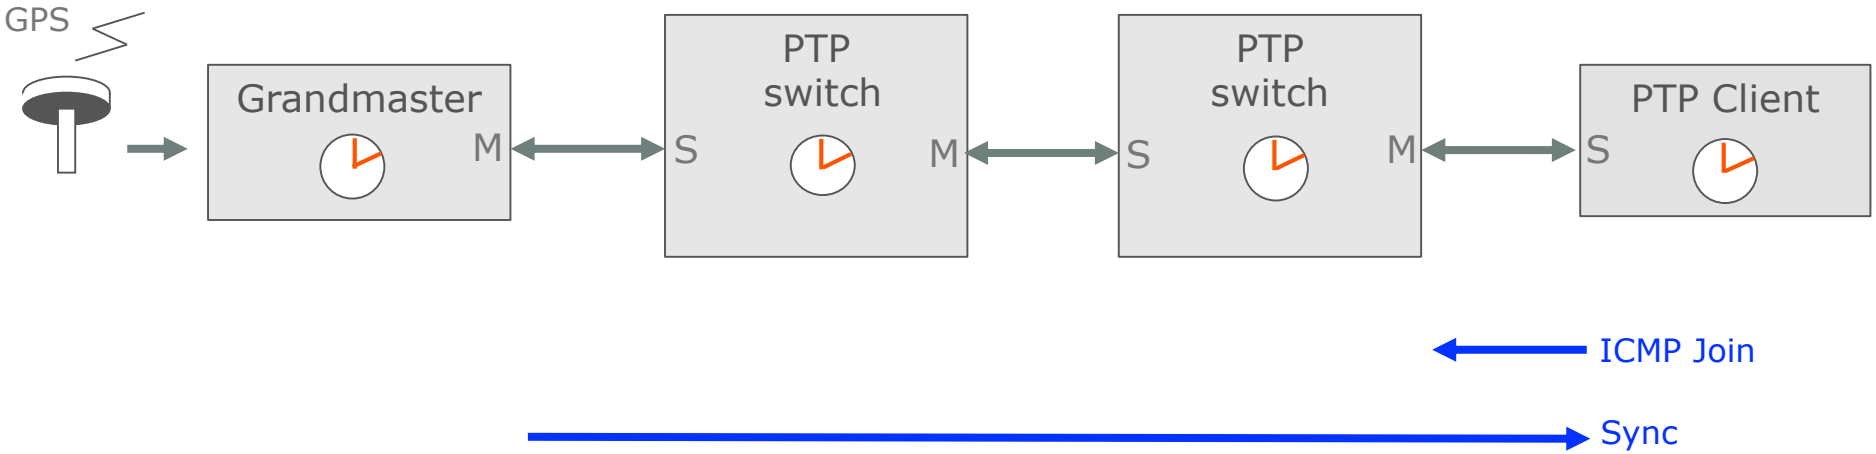

- Although seemingly working, accuracy will be severely compromised
- Not realising ROI of expensive hardware
- Could go undetected for some time

Scenario 3

Path Delay not working

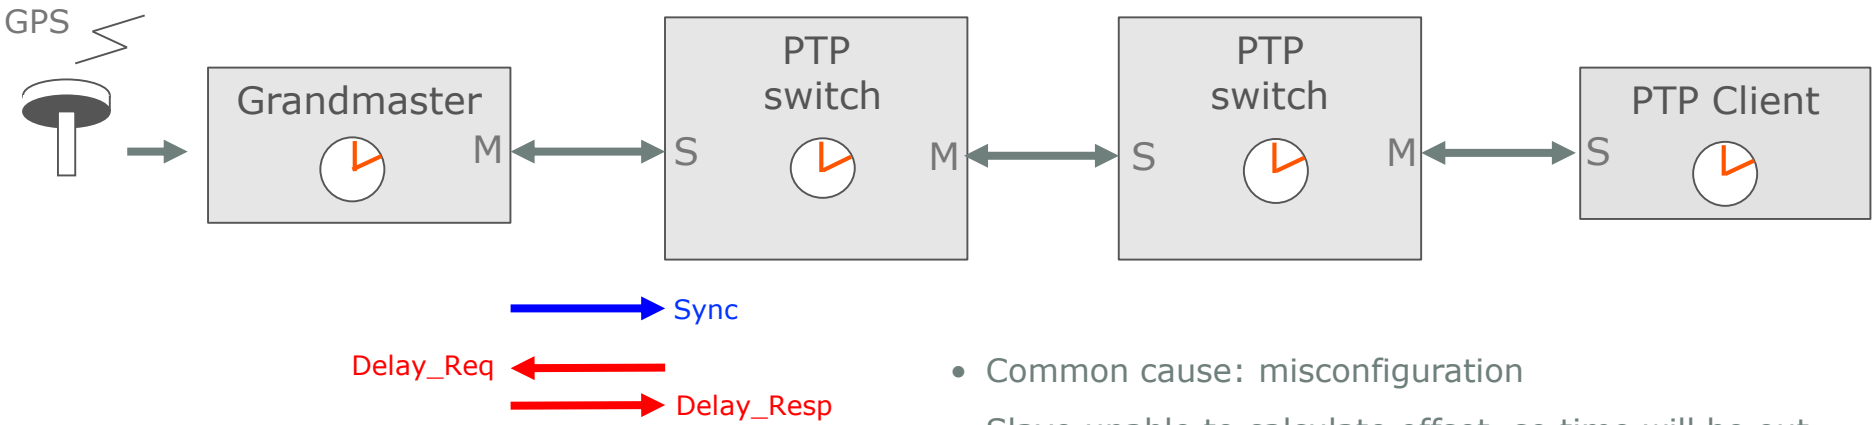

- Common cause: misconfiguration
- Slave unable to calculate offset, so time will be out
- Default action is often to assume zero offset

Latencies not accounted for:

- Propagation - 5ns per metre
- Transceiver delay (large with copper – up to 1us)

Scenario 4

Introduction of a non-PTP switch

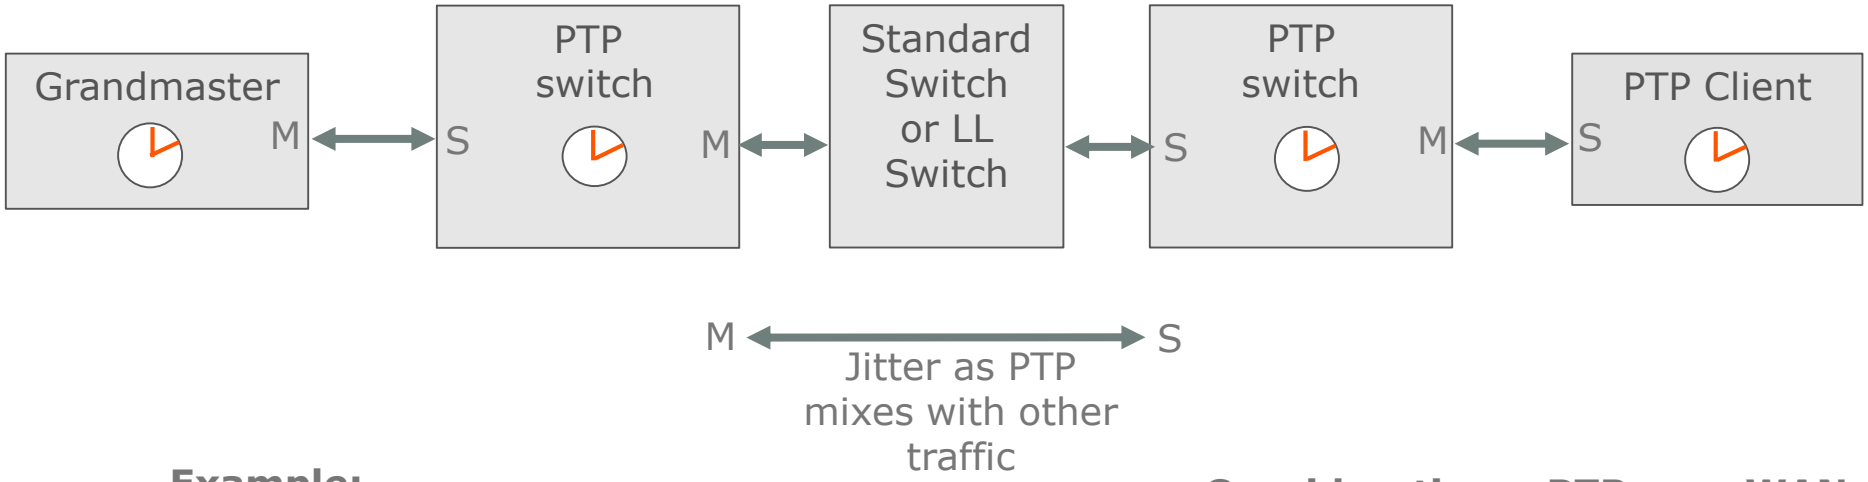

Example:

- Link speed 1Gbps
- Takes 12us to serialize 1500 bytes

- Queuing of PTP packets has a huge impact on achievable accuracy.

Consideration – PTP over WAN:

- Dark fibre vs Ethernet service?

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

GPS – Accurate but vulnerable

- Physical Disturbances
  - Cable cut
  - Signal blocked
- Weather
- Solar Storms
- Sabotage
- Bugs

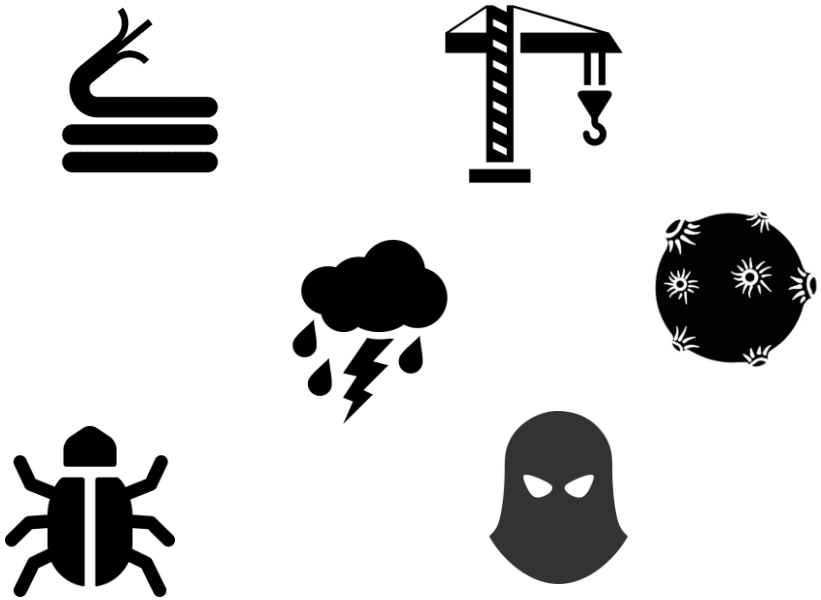

Grandmaster

Loss of GPS signal

- What happens?
- BMCA where resilience is provided
- GM goes into holdover and continues providing time
  - Rate of drift and quality of time is dependent upon the specification (and price) of GM
- Example 0.1 ppm = 100ns per second
  - After 1 minute, off by 6us
  - After 1 hour, off by over 1/3ms!

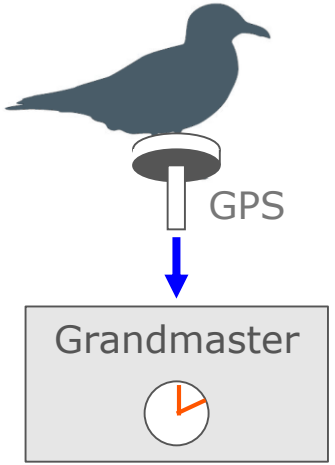

Grandmaster

Inaccurate GPS signal received

- 26th Jan 2016 - 13us GPS error for 12 hours
- Rogue satellite SVN-23 decommissioned
- Error pushed to 15 satellites by ground system software
- Multiple customers logged support tickets reporting PPS errors of 13us

BBC Sign in Menu

NEWS

Home | Video | World | UK | Business | Tech | Science | Magazine

Technology

GPS error caused '12 hours of problems' for companies

By Chris Baraniuk
Technology reporter

4 February 2016 | Technology | Share

THINKSTOCK

System engineers were "called out of bed" over the problems

Grandmaster

TAI vs UTC

- PTP is designed to publish TAI (International Atomic time) plus the UTC offset
- Manages leap seconds – all have been positive so far
- Current offset is 36 seconds
- Next one scheduled for December 31st 2016

- Examples where GM publishes UTC with a zero offset, instead of TAI.
- How will clients handle the next leap second? And how do you test that?

UTC Date	UTC Time	Difference TAI vs. UTC
30/06/1972	23:59:60	11 secs
31/12/1972	23:59:60	12 secs
31/12/1973	23:59:60	13 secs
31/12/1974	23:59:60	14 secs
31/12/1975	23:59:60	15 secs
31/12/1976	23:59:60	16 secs
31/12/1977	23:59:60	17 secs
31/12/1978	23:59:60	18 secs
31/12/1979	23:59:60	19 secs
30/06/1981	23:59:60	20 secs
30/06/1982	23:59:60	21 secs
30/06/1983	23:59:60	22 secs
30/06/1985	23:59:60	23 secs
31/12/1987	23:59:60	24 secs
31/12/1989	23:59:60	25 secs
31/12/1990	23:59:60	26 secs
30/06/1992	23:59:60	27 secs
30/06/1993	23:59:60	28 secs
30/06/1994	23:59:60	29 secs
31/12/1995	23:59:60	30 secs
30/06/1997	23:59:60	31 secs
31/12/1998	23:59:60	32 secs
31/12/2005	23:59:60	33 secs
31/12/2008	23:59:60	34 secs
30/06/2012	23:59:60	35 secs
30/06/2015	23:59:60	36 secs
31/12/2016	23:59:60	37 secs

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

Analyzing time quality downstream

Comparison of multiple time sources and clock modeling

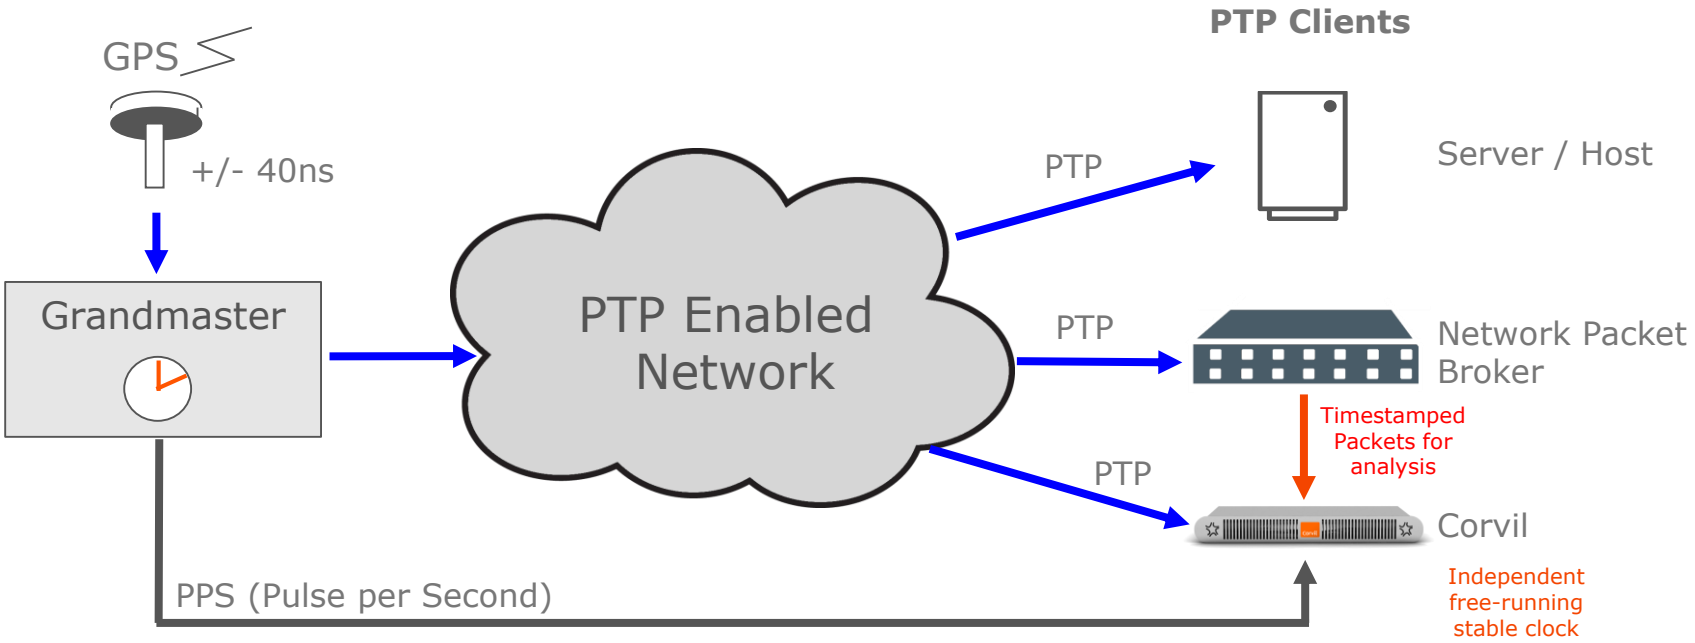

A Selection of memorable issues

Bugs and anomalies discovered in the field

- Various Grandmasters and switches with 1 microsecond of jitter
- Missing delay response messages, resulting in incorrect offset
- PTP Sync to follow-up message delay of up to 5 seconds
- PTP Switch off by 5.5us due to port mis-config 100Mbps / 10Gbps
- 3rd party PTP service providers not meeting SLAs
- Switch that appeared to sync to random offsets from UTC
  - E.g. 18 minutes 19.5 seconds, or 55 minutes
  - Bug: Multiples of 2^40 nanoseconds

Lab test: Holdover Investigation

- Compare elapsed time of free-running clock with GPS
- Plot accumulated difference in microseconds:
- Clock-drift of about -4.7ppm

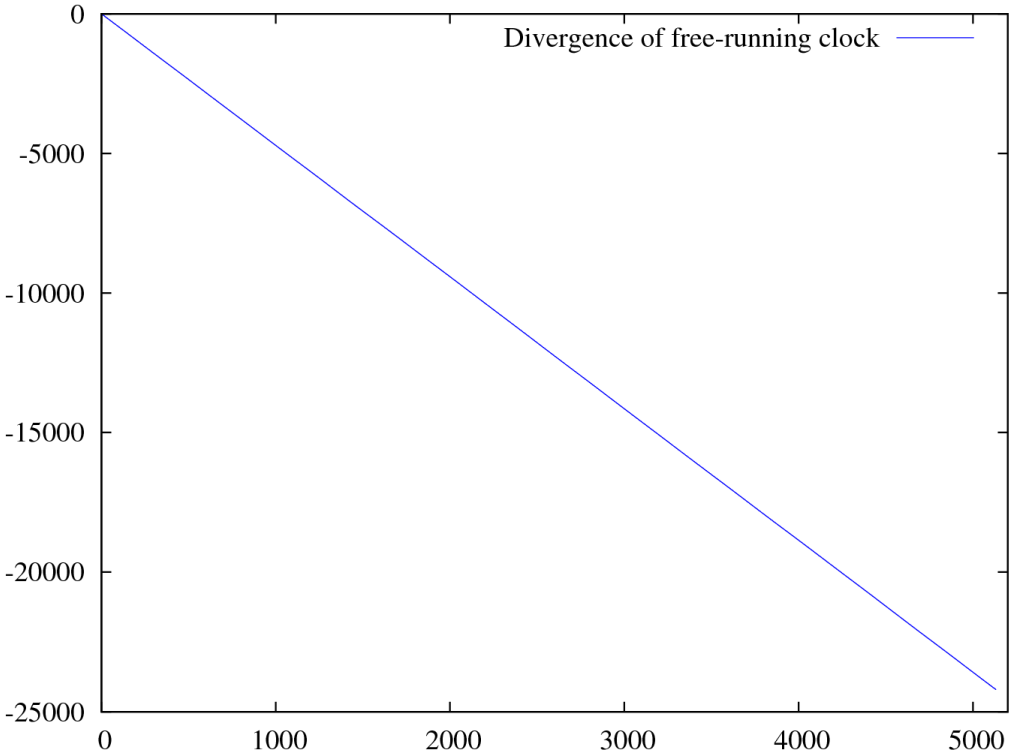

Correct for clock drift

- Linear transformation to remove drift
- Reveals jitter and wander
- Oscillation with 5½ min period?

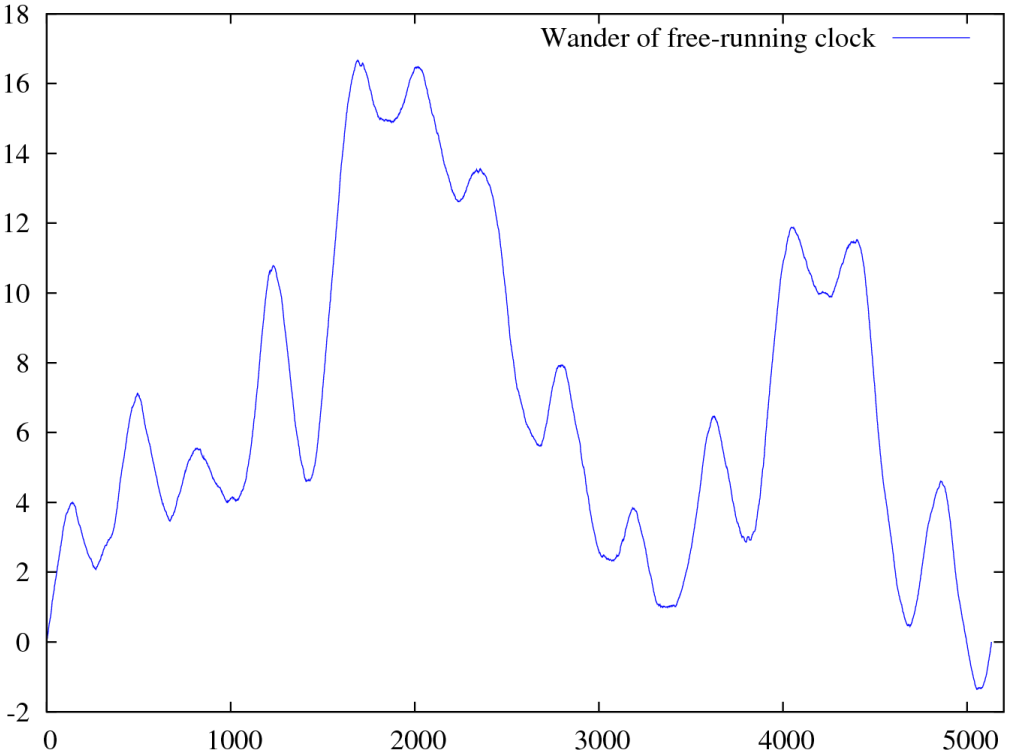

Suspect Environmental Effect

- Air-conditioning in lab cycles roughly every 5½ mins…
- Check changes in clock frequency:

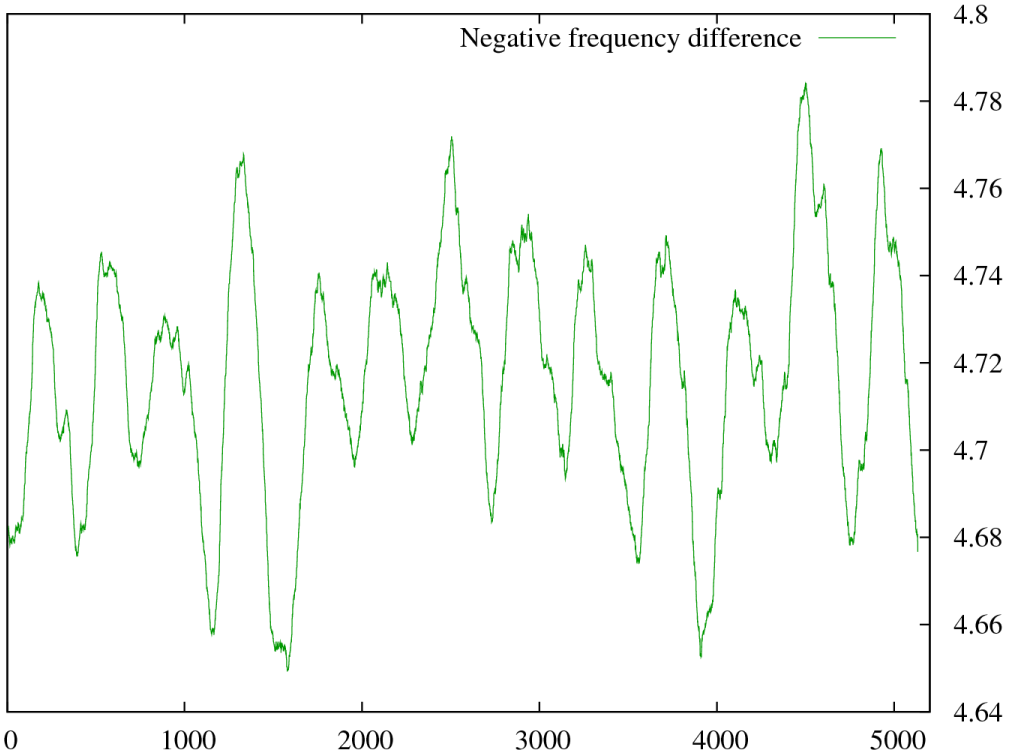

Clock Frequency Variance

Due to Temperature Fluctuation

- Temperature sensitivity:

  -0.4ppm/C

- For every rise in temp. of 1 degree Celsius, clock slows by 0.4ppm

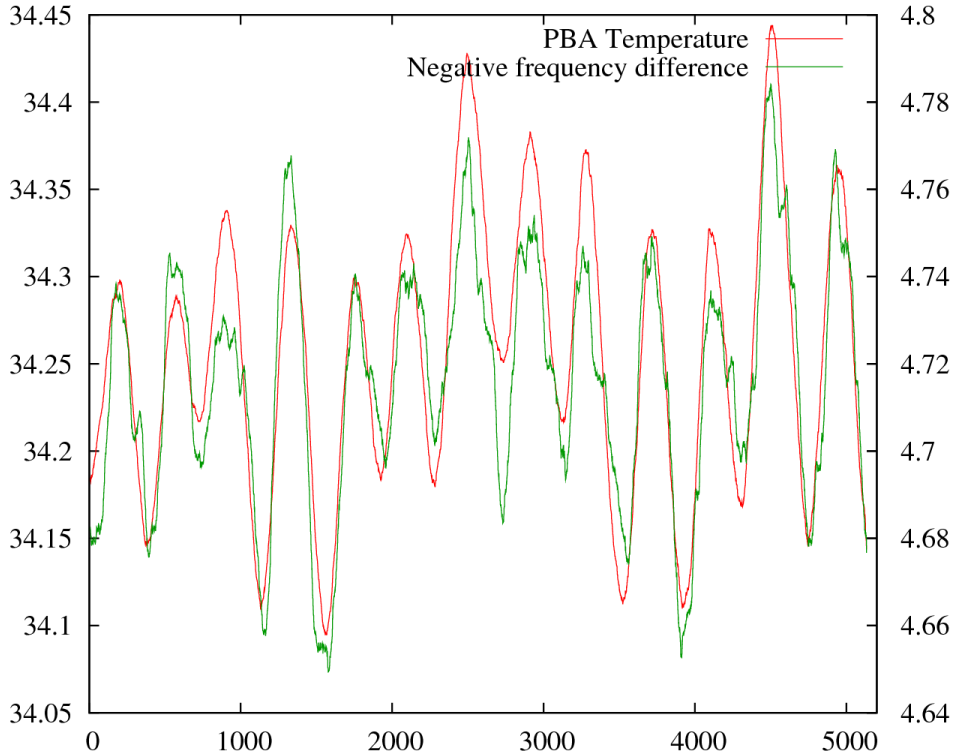

PTP & UTC Traceability for MiFID II

Corvil

Detection

- Track PTP against PPS
- Alert on PTP jitter over specified threshold
- Passively monitor PTP traffic to multiple hosts and validate quality

Reporting

- Time series of PTP accuracy for audits
- Explicit UTC sync Y/N flag and alerting
- Continuous Sync health reported with order records

Agenda

Precision Time Protocol
A Brief Recap & How it Works

Common Deployment Gotchas
A look at PTP Switches

Upstream Considerations
GPS & Grandmaster Clocks

The Downstream Perspective
Observing Time Quality

Wrap up & Summary
Take Aways & Best Practices

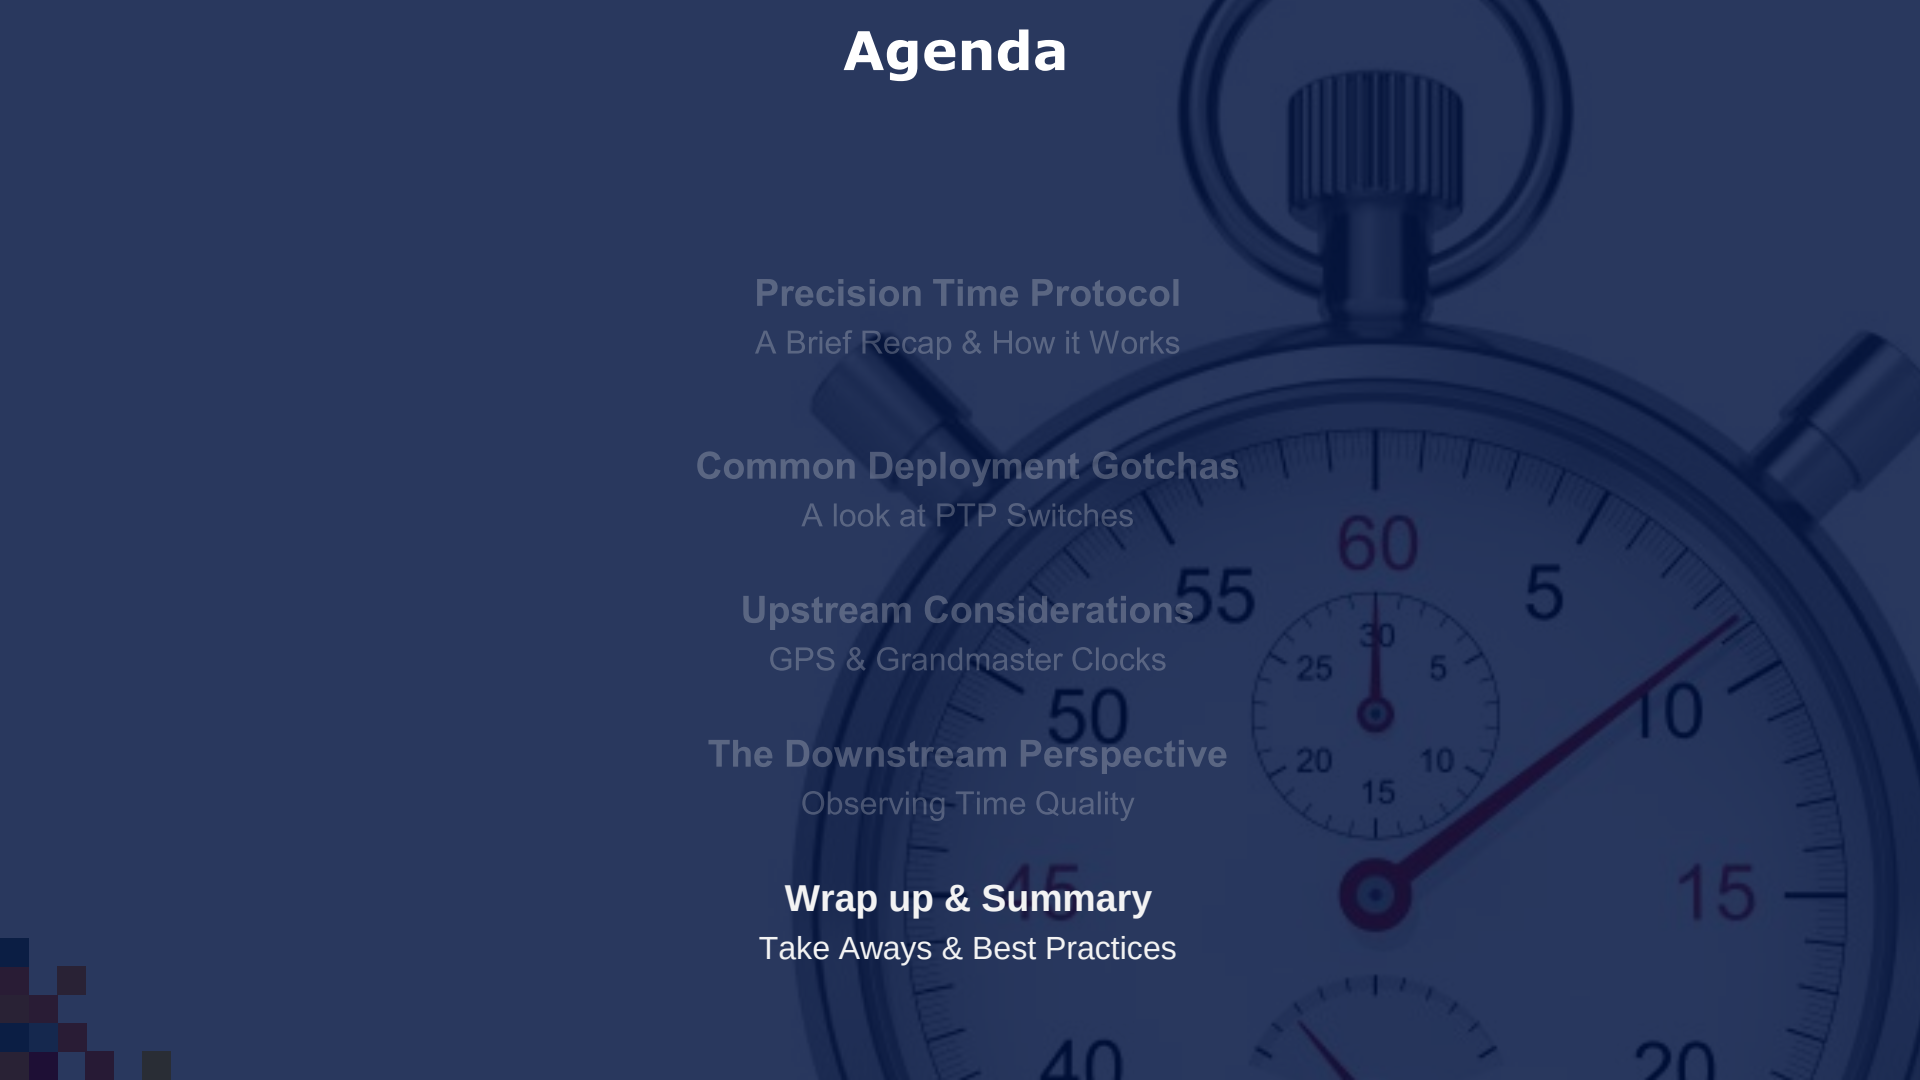

Summary

- With Boundary Clock implementation, accuracy is achieved with a Master/Slave across each physical link.
- Multicast forwarding (across switch) is NOT required
- PTP Does NOT require a dedicated network (just PTP aware switches)
- Ensure PTP switches are enabled
- Use multiple time sources to spot anomalies
- Bugs do exist, so expect the unexpected


Corvil
Thank You