NYC 1/25/17 ATIS Time&Money

MiFID II Clock Sync Lessons Learned

FSMLabs
Enterprise Real-Time®

Q4/2017

- Financial trading firms in EU, UK, US and worldwide tried to get in under the wire for RTS 25 requirements of MiFID II
- FSMLabs worked with a variety of firms in from tier 1 global banks to small proprietary trading shops.

Some were way ahead of the curve.
Some were **not** ahead of the curve.

FSMLabs
Enterprise Real-Time®

OUTLINE OF TALK

- Quick look at some general lessons from implementation efforts
- A tour of the key technical issue: understanding what clock sync technology can and cannot do.

GENERAL LESSON 1: SUPPLY AND DEMAND APPLIES TO BANKS

- Most critical obstacle is people
  - Firms with high quality technologists who understand business and engineering issues and have institutional weight saved a lot of money.
  - Lesson: pay market rates. Google and Facebook want the same people.
- Large companies with complex legacy management structures have a hard time doing major company wide technical changes. Not surprising.
- There are way too many old systems out there and VMs are making this worse.
- Things went surprisingly well. This is a solvable problem, even in a hurry.

GENERAL LESSON 1: THIS STUFF IS HARD.

- Most critical obstacle is people
- Firms with high quality technologists who understand business and engineering issues and have institutional weight saved a lot of money.
- **Lesson**: pay market rates. Google and Facebook want the same people

FSMLabs
Enterprise Real-Time®

GENERAL LESSON 2: IT'S JUST ENGINEERING, NOT MAGIC.

- Huge variety of firms did major upgrades of their clock sync successfully in an unreasonably short time.
- We saw a lot of firms that delayed because they thought it was going to be impossible.
- **Lesson:** Avoid heroic efforts by making timely investments in technology.

FSMLabs
Enterprise Real-Time®

GENERAL LESSON 3: MODERNIZE YOUR PLANT

- Kind of stunning how much obsolete and underpowered tech is out there.
- An 8 year old GPS clock in London cannot reliably sync software clocks in a data center in Chicago over a high latency connection.
- **Lesson**: don’t waste time trying to make it work if it can’t work.

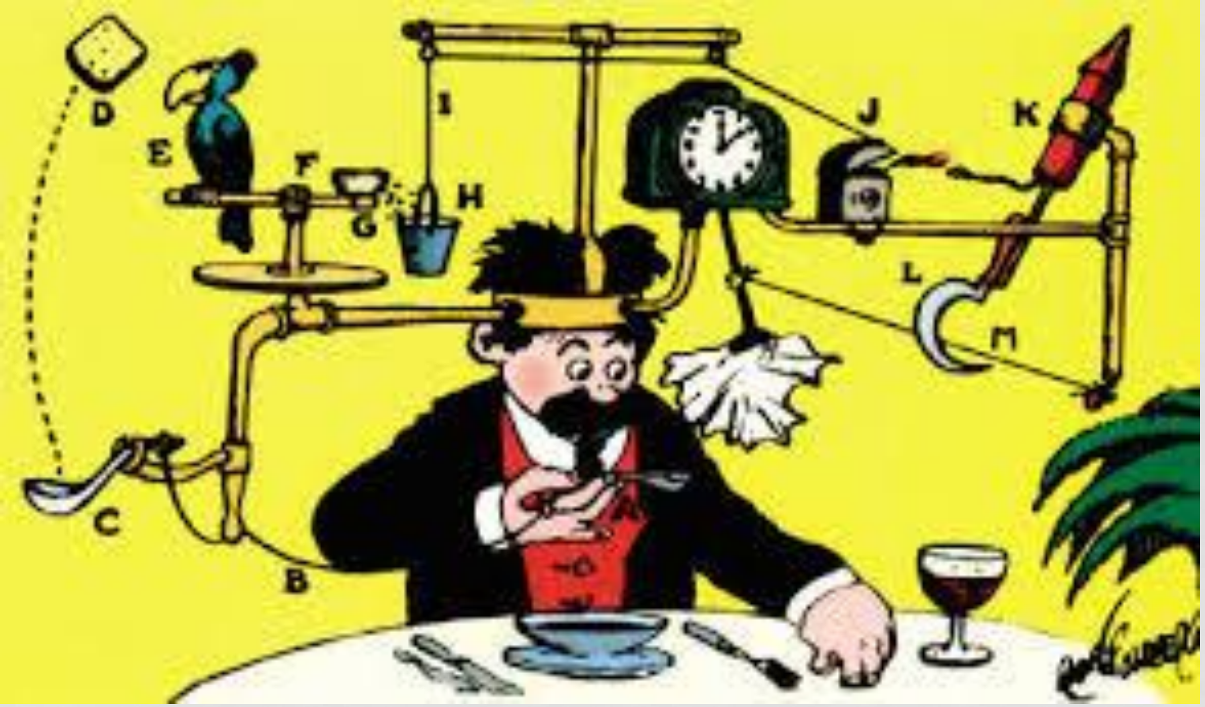

GENERAL LESSON 4: USE COMMON SENSE

- TimeKeeper Compliance collects data from all visible clock sync clients both directly and via GMs, creates searchable database, produces reports. Something like that is essential for regulatory compliance.
- Not unusual to have terrabytes of data.
- Some customers tried to run it with many clients on low resource VMs with limited storage and processing.
- **Lesson**: Don’t take a sheep cart to a tractor pull.

MAINTAINING COMPLIANCE AND PROVING IT REQUIRES AN UNDERSTANDING OF WHAT CLOCK SYNC CAN AND CANNOT DO.

The key questions about measurements

1. “Is it telling the truth as it knows it?”
2. “Does it have enough information?”
3. “Where does the information come from?”
4. “Is that information logged securely?”

FSMLabs
Enterprise Real-Time®

SKEPTICISM.

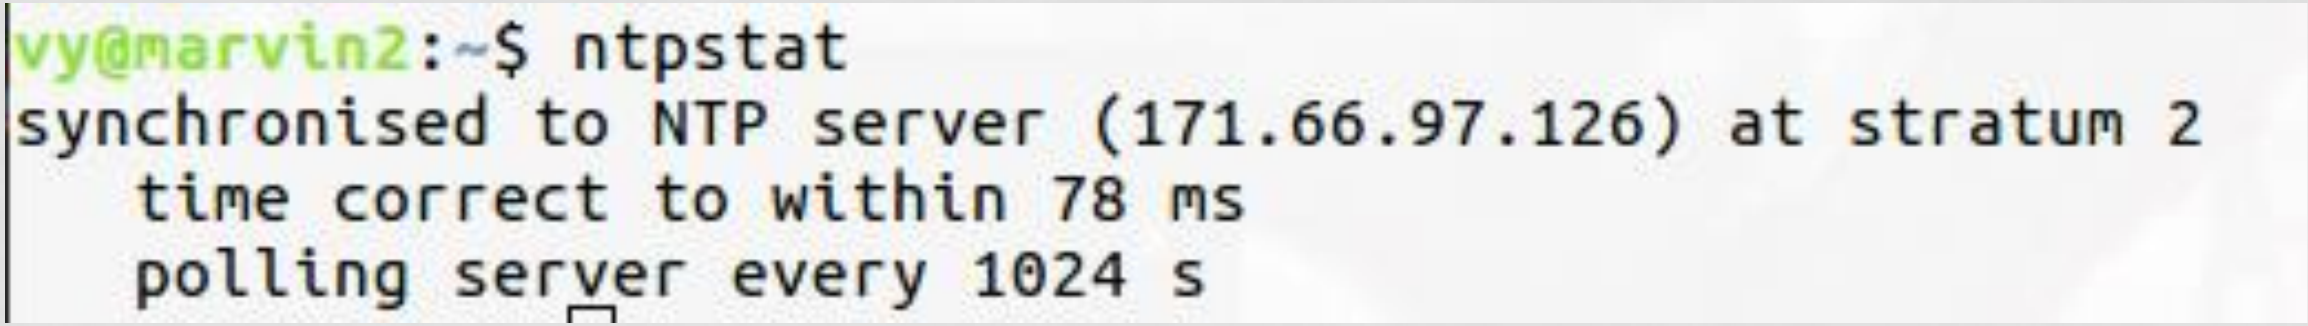

```
vy@marvin2:~$ ntpstat
synchronised to NTP server (171.66.97.126) at stratum 2
   time correct to within 78 ms
   polling server every 1024 s
```

The second line could be true or not, nobody knows.

The translation is: **You told me the clock from 171.66.97.126 is good. I can keep OS clock to within 78ms of what I estimate is the time on that clock if my calculations on network packet delay and local oscillator frequency (if any such estimates) are good**.

HOW CAN YOU TEST THE CLOCK IN THE OS AGAINST THE EXTERNAL CLOCK?.

?

The second line could be true or not, nobody knows.

The translation is: **You told me the clock from 171.66.97.126 is good. I can keep OS clock to within 78ms of what I estimate is the time on that clock if my calculations on network packet delay and local oscillator frequency (if any such estimates) are good**.

FSMLabs
Enterprise Real-Time®

6 YEARS: OUR BASIC APPROACH HAS WORKED WELL

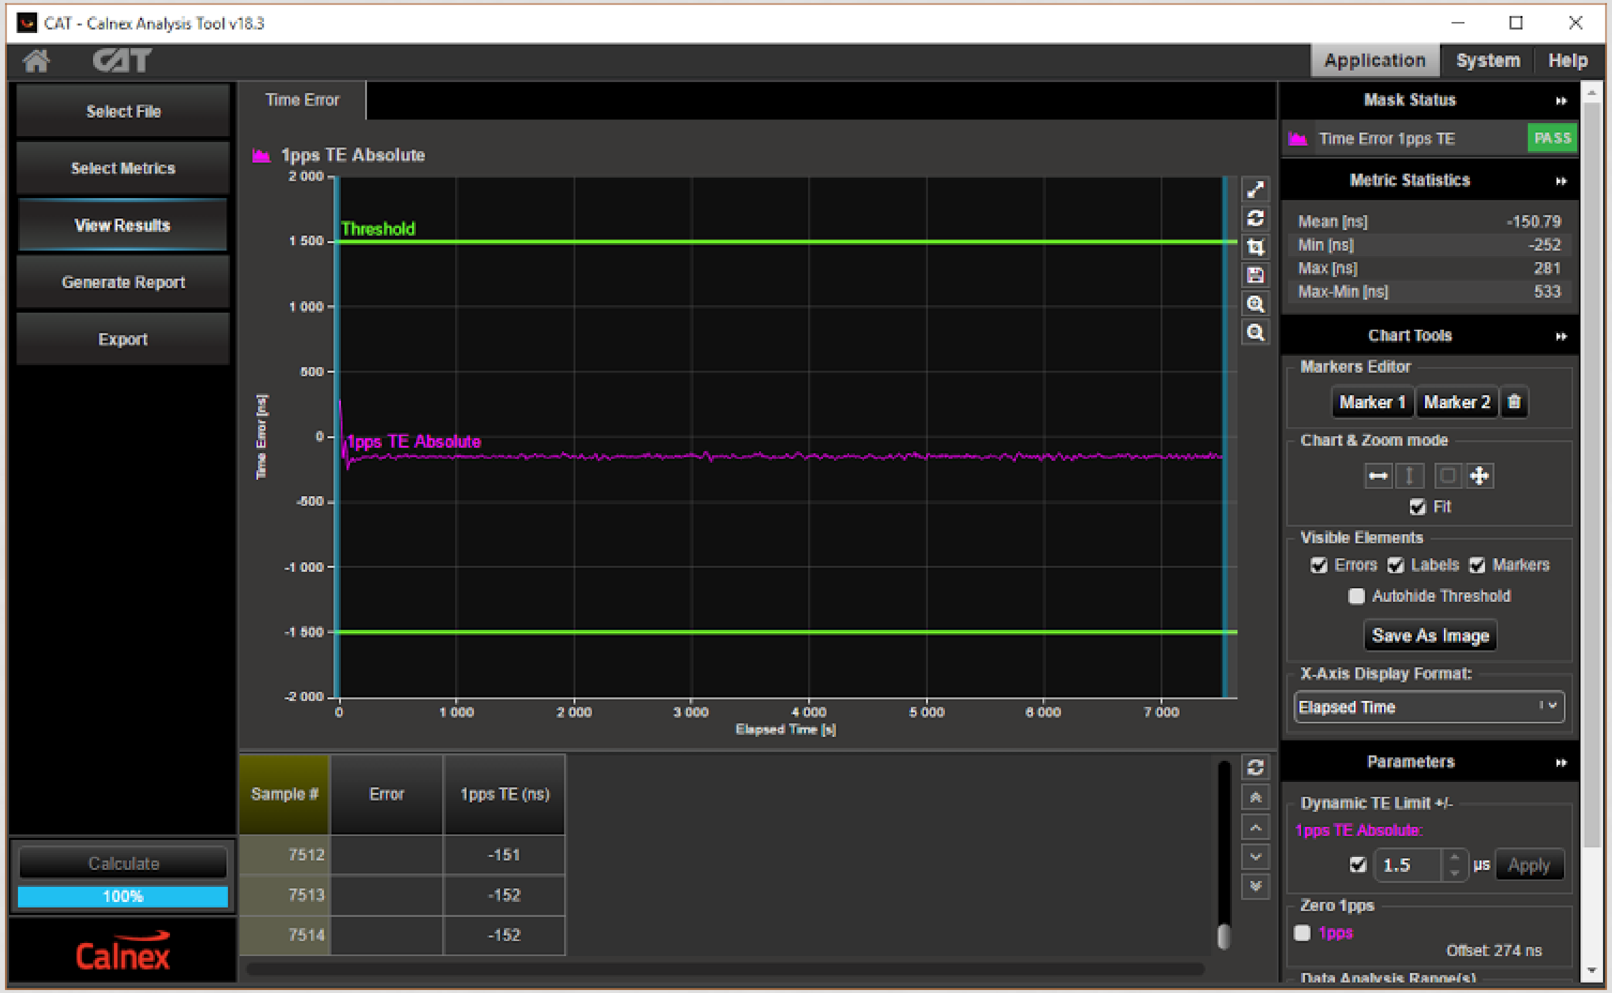

Time Error 1pps TE PASS

Mean [ns]	-150.79
Min [ns]	-252
Max [ns]	281
Max-Min [ns]	533

☑ Fit

☑ Errors ☑ Labels ☑ Markers

□ Autohide Threshold

X-Axis Display Format: Elapsed Time

Sample #	Error	1pps TE (ns)
7512		-151
7513		-152
7514		-152

Dynamic TE Limit +/-

1pps TE Absolute:

☑ 1.5 µs

Zero 1pps

□ 1pps

Offset: 274 ns

FSMLabs
Enterprise Real-Time®

WORKS FOR GPS CLOCK FAILOVER AS WELL AS CLIENT FAILOVER

This clock has 2 PTP sources and GPS(green) all within 200ns
(40 Gbps network)

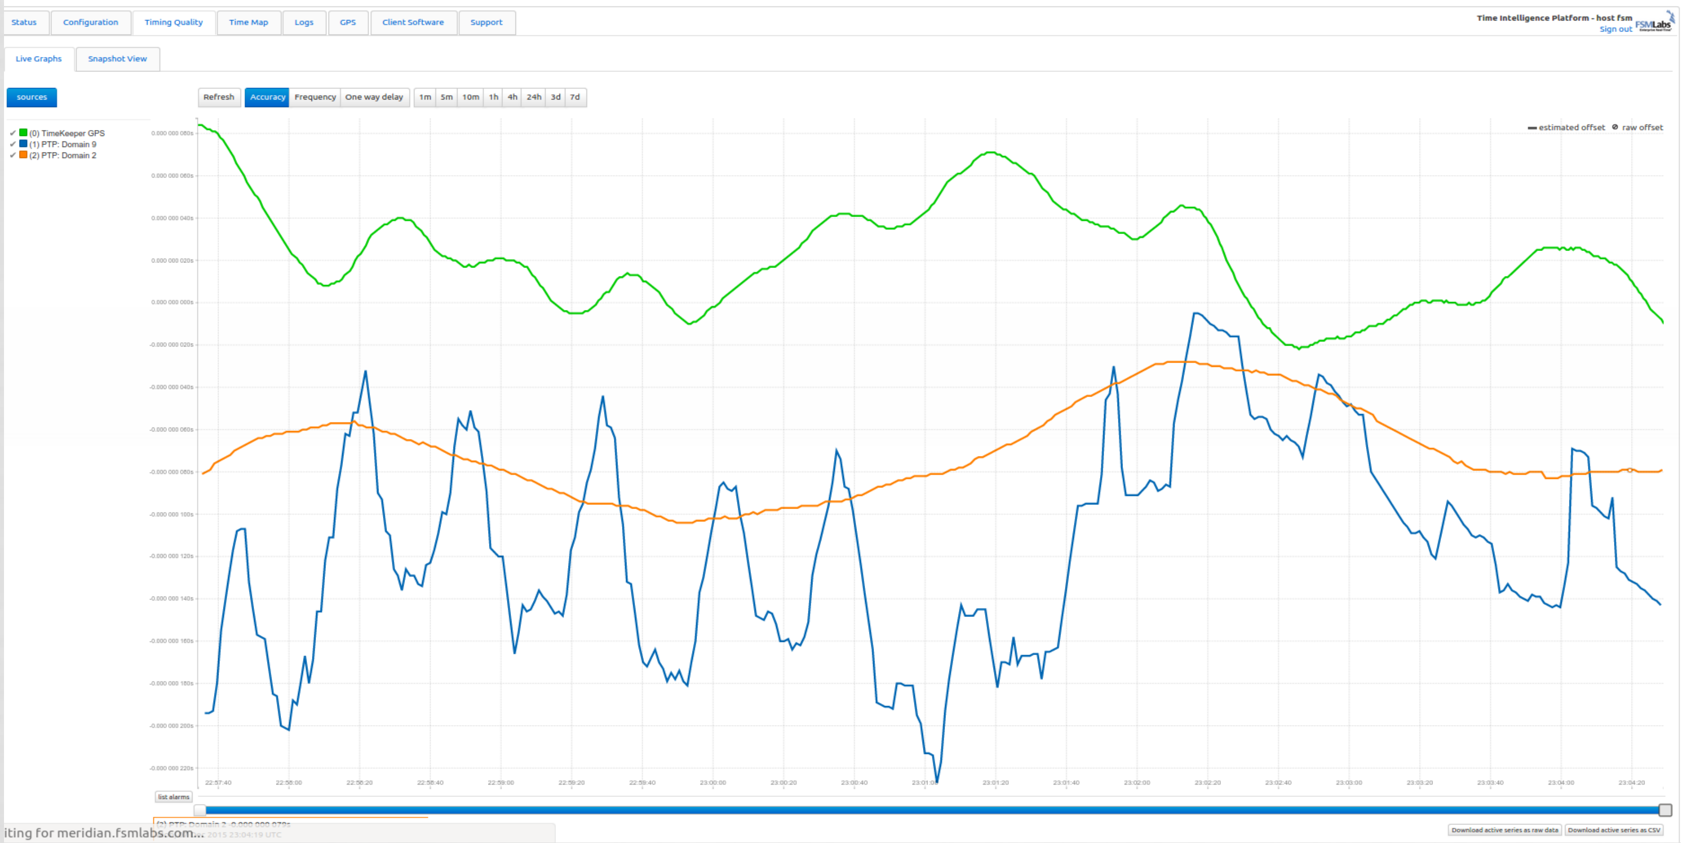

Boundary clocks/Stratum Servers also benefit

SYSTEMS WITH SOURCE CHECK SURVIVED MANY FAILURES

- GPS Clock that silently lost GPS and switched to NTP backup, repeatedly.
- Overeager Network Security that cut off part of PTP protocol.
- Month premature Leap Second jumps.
- Terrible Switch Boundary clocks that never worked.
- High speed WAN connect with 12 microseconds asymmetry not detected by network admins.

OTHER SURVIVED FAILURES

- Lightning strikes on GPS antennas
- NTPd servers changing their own sources
- Broken or malfunctioning GPS clocks
- Bad oscillators (or overheated ones) on boundary clocks
- GPS spoofing and jamming
- Misconfigured routers/switches
- Failed terrestrial PTP sources
- Lost PTP multicasts on switch restarts

SURPRISES

- Failures are more common than expected
  - Especially jamming and GPS reception
- Many systems had no reliable sources at all – so failover was not an issue
- How quickly customers got used to it working and ignored failure signals since the system just recovered

ADAPTATIONS

- Due to prevalence of systems with only one or zero reliable time sources, original failover could oscillate between weak sources. Fixed.
- Needed to radically improve diagnostics to help solve problems, find configurations (customers wanted fault tolerance and then repair)
  - Map of network time distribution
  - Deep diagnostic of GPS signal

TIME MAP – SHOWING ALL SOURCES AND CLIENTS

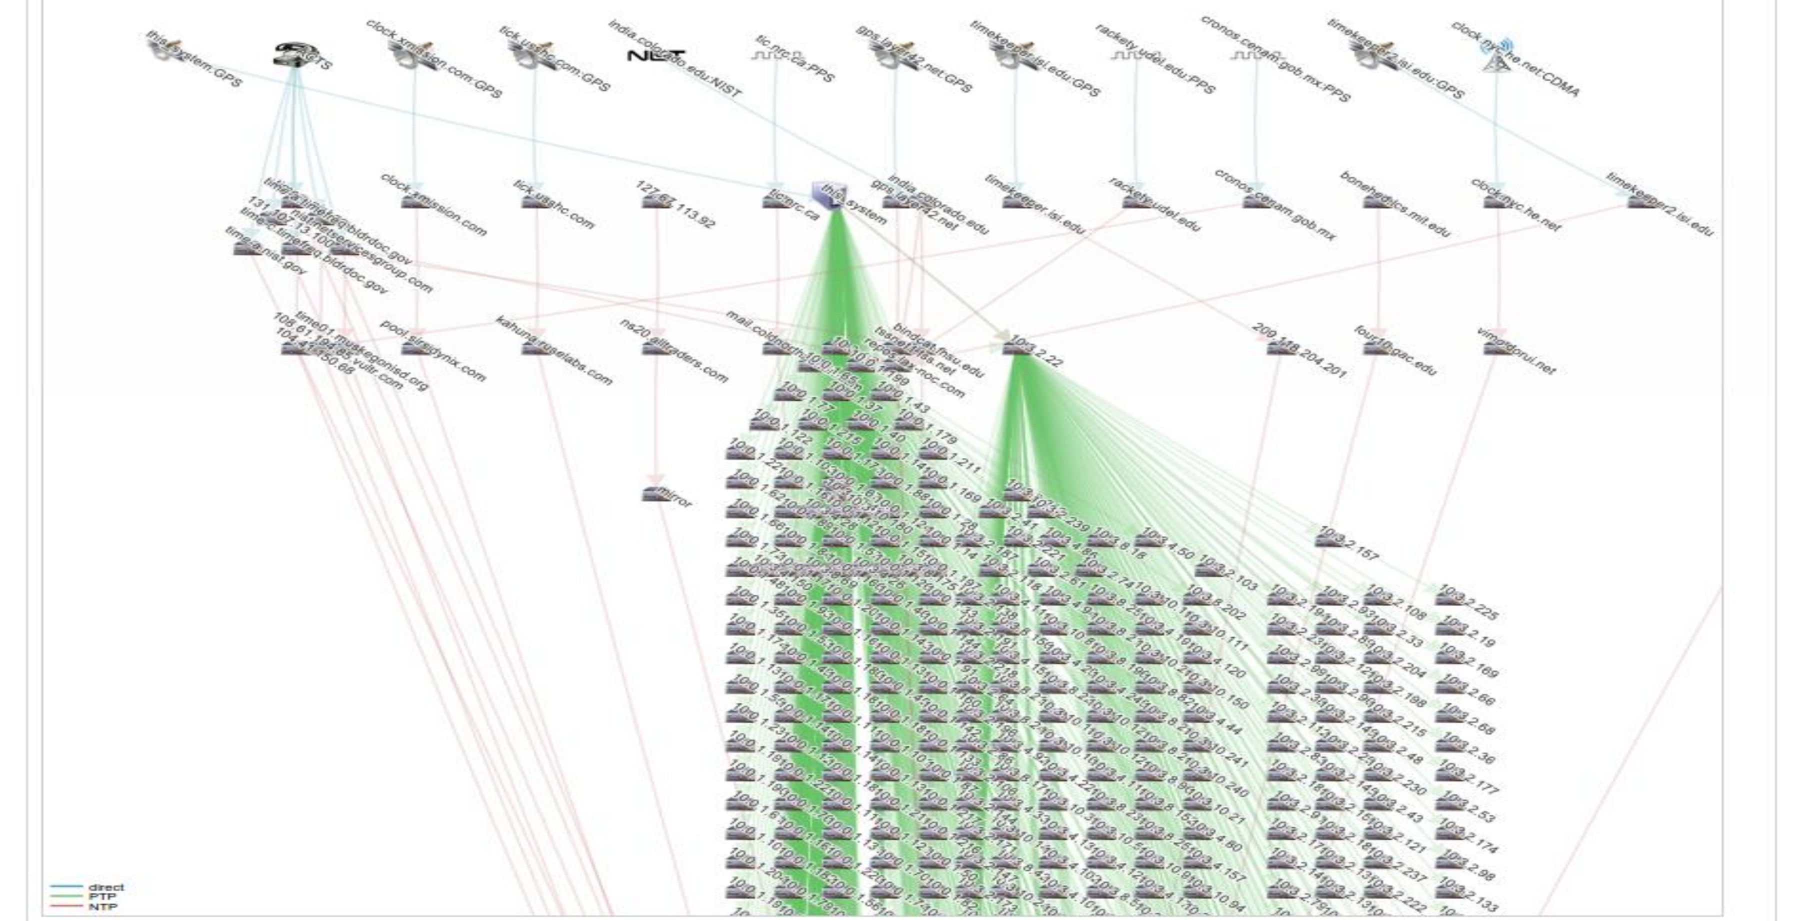

Green is PTP, red is NTP, blue is source

TIME MAP – OFTEN SHOWED “INDEPENDENT” SOURCES WERE NOT.

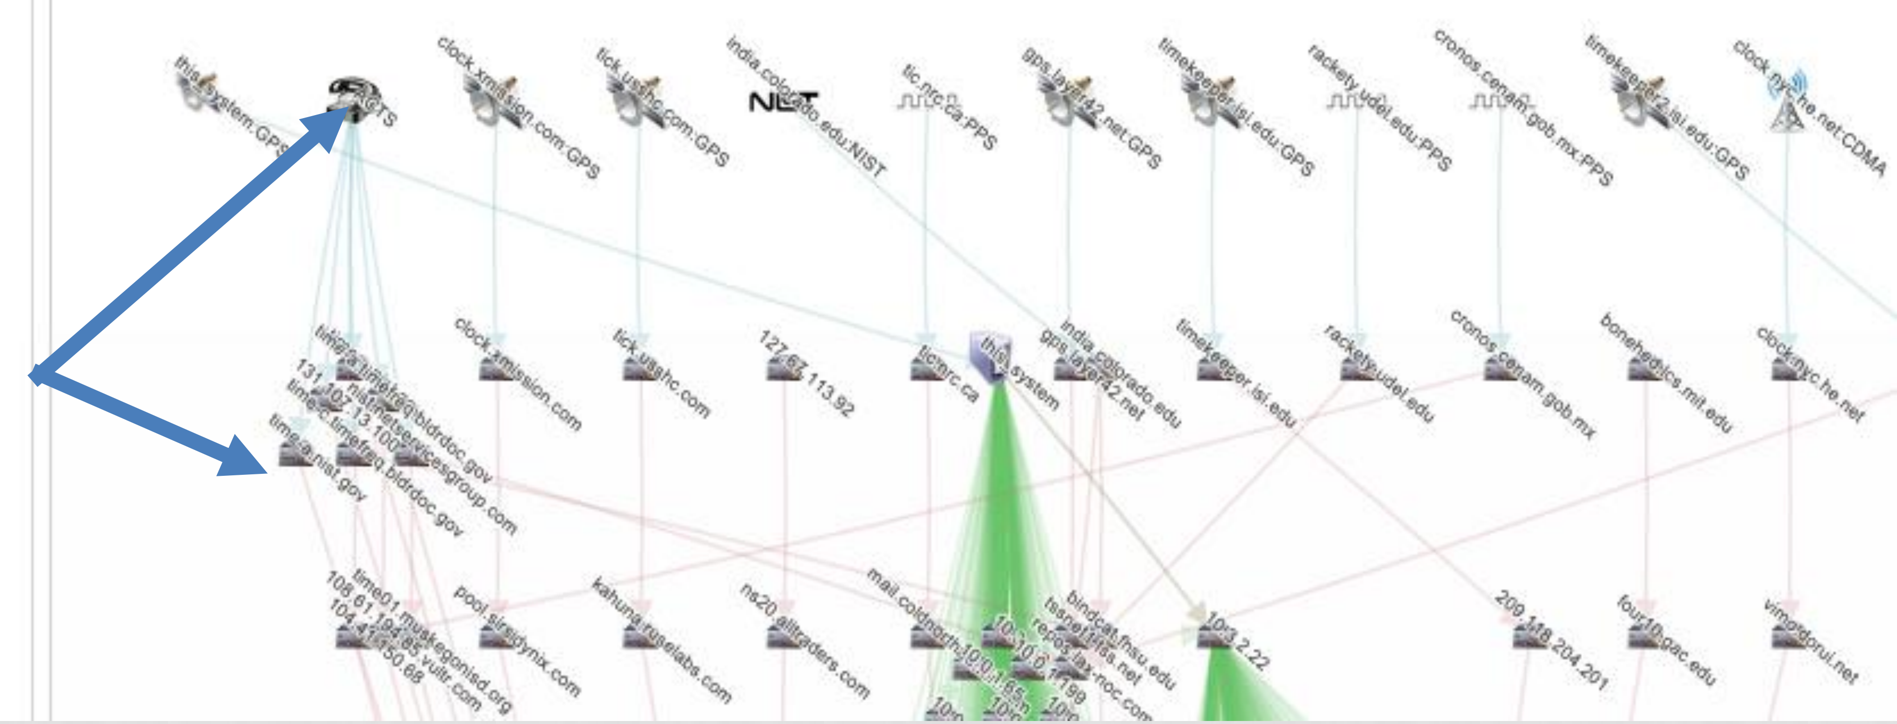

Several “independent” sources taking data from same “ACTS” modems – bad sign.

SECOND KEY DIAGNOSTIC NEEDED TO HELP WITH GPS INTERFERENCE OR JAMMING ISSUES.

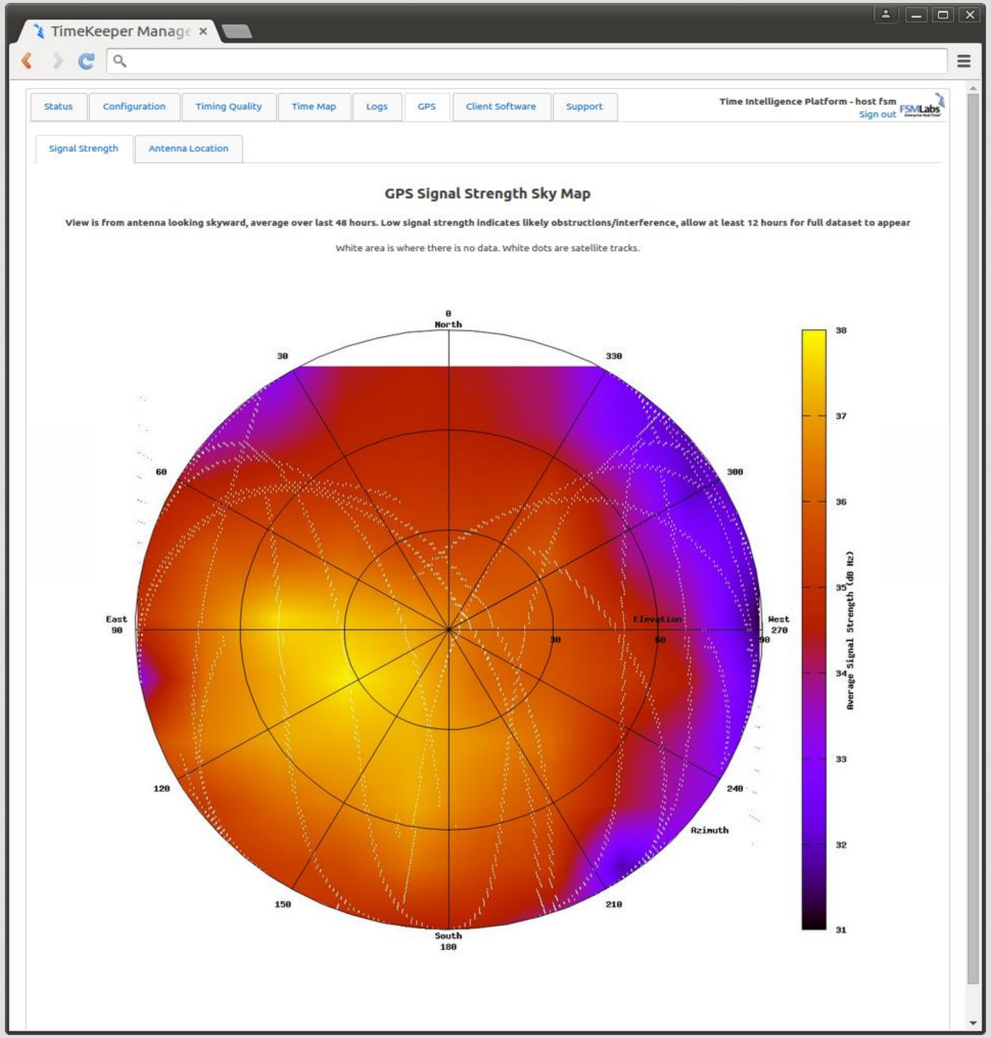

GPS Signal Strength Sky Map

View is from antenna looking skyward, average over last 48 hours. Low signal strength indicates likely obstructions/interference, allow at least 12 hours for full dataset to appear

White area is where there is no data. White dots are satellite tracks.

Build specialized heat map from GPS signal data so show composite picture of signal strength. Purple areas show blocked reception.

NEW MEXICO DURING A JAMMING TEST AT WHITE SANDS.

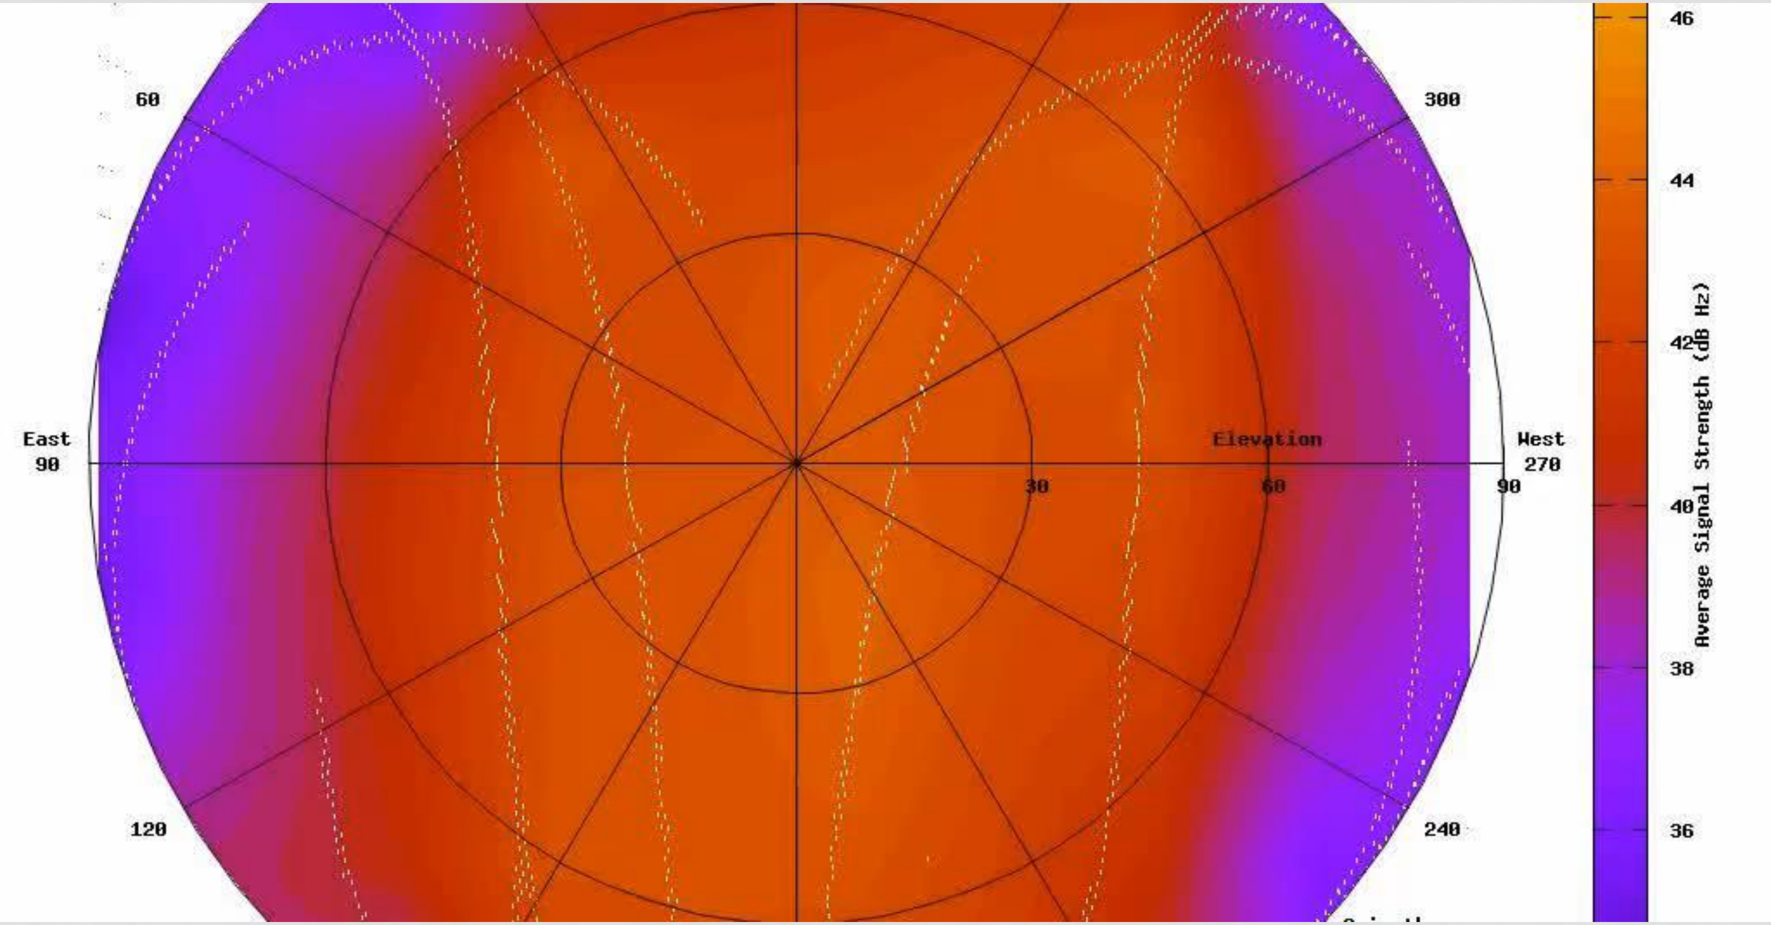

JAMMING AT LD4

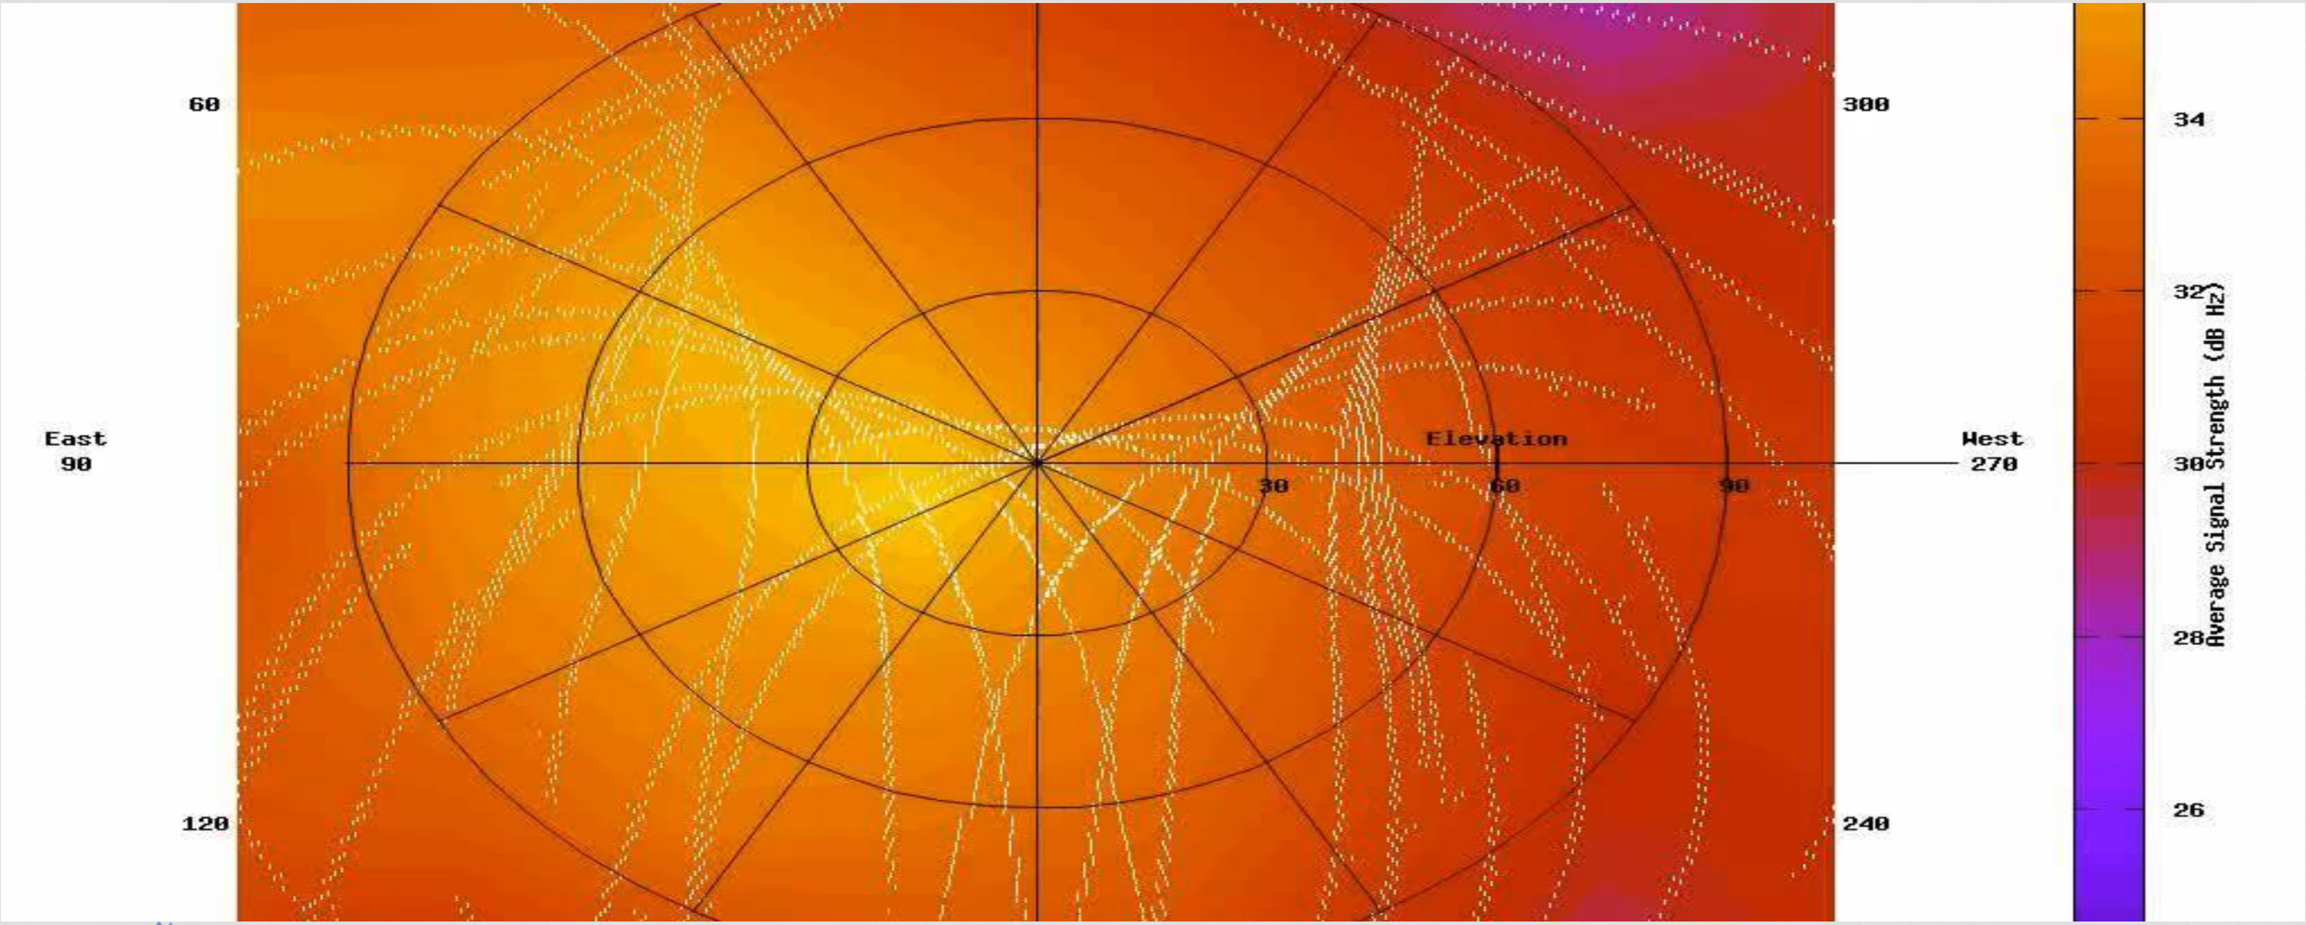

FSMLabs Enterprise Real-Time®

SUMMARY OF THE APPROACH

- **Time protocol agnostic – PTP, PTP-Telecom, NTP, PPS, Bus Card, … all are sources**
- **Multiple sources are essential for**
  - **Fault-tolerance**
  - **Security**
  - **Documentation (e.g. for regulators)**
- **Intelligence in client/slave: time consumer has information and analytics not available to time sources**

FSMLabs
Enterprise Real-Time®

5 YEARS EXPERIENCE – SOME WITH GIANT NETWORKS

- Time distribution is really fragile with many points of failure.
- Existing systems are often terrible.
- Security and fault-tolerance are often indistinguishable.
- Diagnostics is often as important as resilience.
- People are highly inventive about finding ways to break systems.

CONTACT INFO

Victor Yodaiken
FSMLabs, Inc.
11701 Bee Caves Road, Suite 200
Austin, TX 78738
USA
yodaiken@fsmlabs.com

Telephone: 1-512-263-5530

BONUS SLIDE

Multiple Internet NTP sources pre-Leap Second 2015 as Google time servers "slew" off correct time

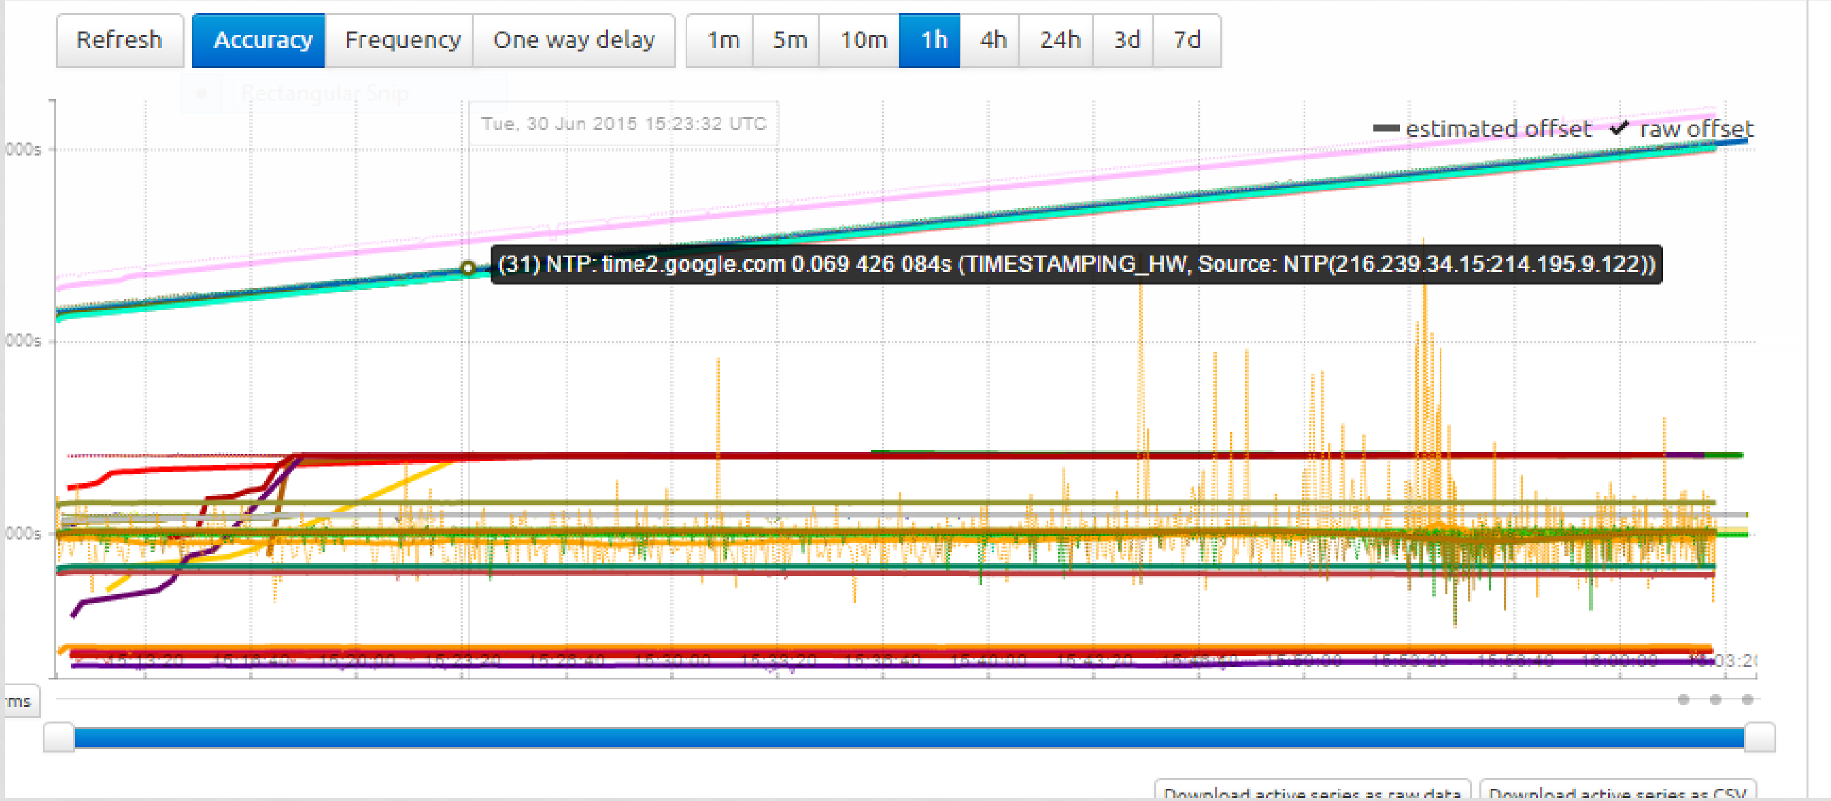

(31) NTP: time2.google.com 0.069 426 084s (TIMESTAMPING_HW, Source: NTP(216.239.34.15:214.195.9.122))

FSMLabs
Enterprise Real-Time®

BONUS SLIDE

Multiple Internet NTP sources Leap Second 2016 as Google time servers "slew" off correct time and others "lurch" backward.

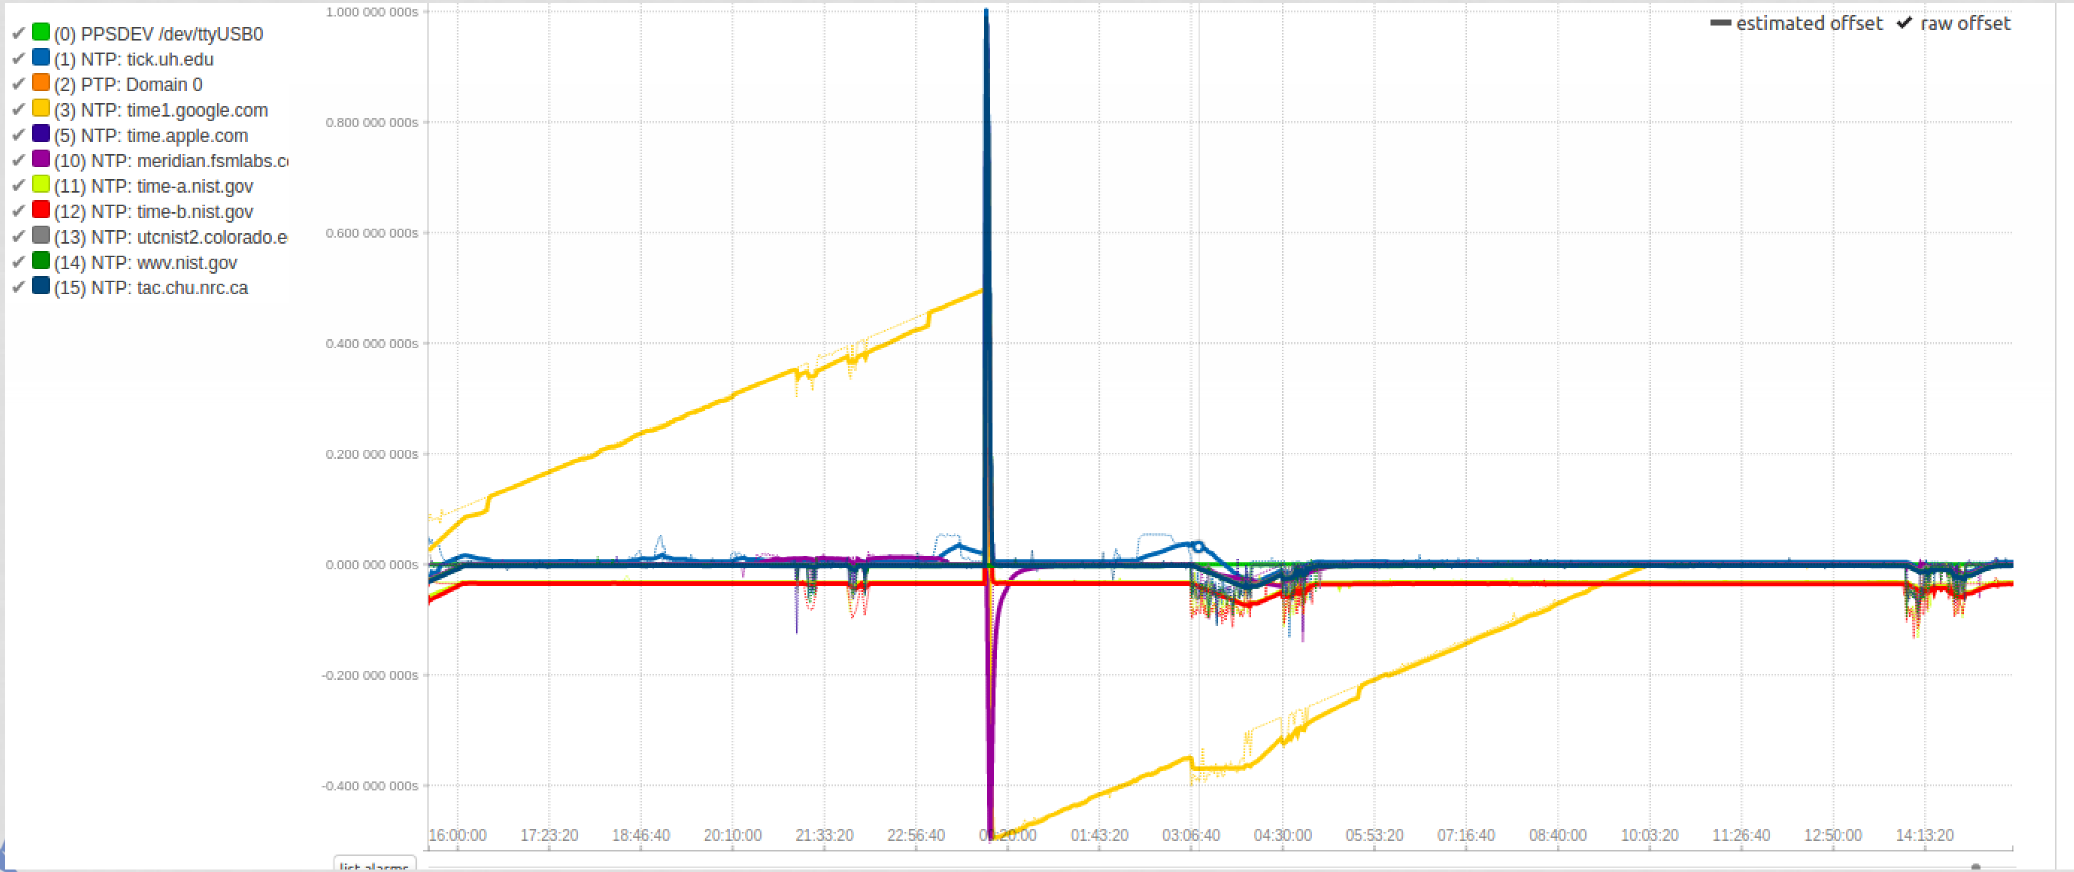
✓ 🟩 (0) PPSDEV /dev/ttyUSB0
✓ 🟦 (1) NTP: tick.uh.edu
✓ 🟧 (2) PTP: Domain 0
✓ 🟨 (3) NTP: time1.google.com
✓ 🟪 (5) NTP: time.apple.com
✓ 🟪 (10) NTP: meridian.fsmlabs.c
✓ 🟩 (11) NTP: time-a.nist.gov
✓ 🟥 (12) NTP: time-b.nist.gov
✓ ⬛ (13) NTP: utcnist2.colorado.e
✓ 🟩 (14) NTP: www.nist.gov
✓ 🟦 (15) NTP: tac.chu.nrc.ca
estimated offset
✓ raw offset
1.000 000 000s
0.800 000 000s
0.600 000 000s
0.400 000 000s
0.200 000 000s
0.000 000 000s
-0.200 000 000s
-0.400 000 000s
16:00:00
17:23:20
18:46:40
20:10:00
21:33:20
22:56:40
00:20:00
01:43:20
03:06:40
04:30:00
05:53:20
07:16:40
08:40:00
10:03:20
11:26:40
12:50:00
14:13:20
list alarms