

MARIST

School of Computing Science and Mathematics



Assuring Time Accurate Transactions

Precision Time Protocol for Finance and Banking



Robert Lynch – robert.lynch3@marist.edu

MARIST

School of Computing Science and Mathematics



Robert Lynch

Undergraduate Information Technology Student

Casimer DeCusatis

Faculty Advisor



Greg Lacey

Joint Study Project Manager


Paul Wojciak

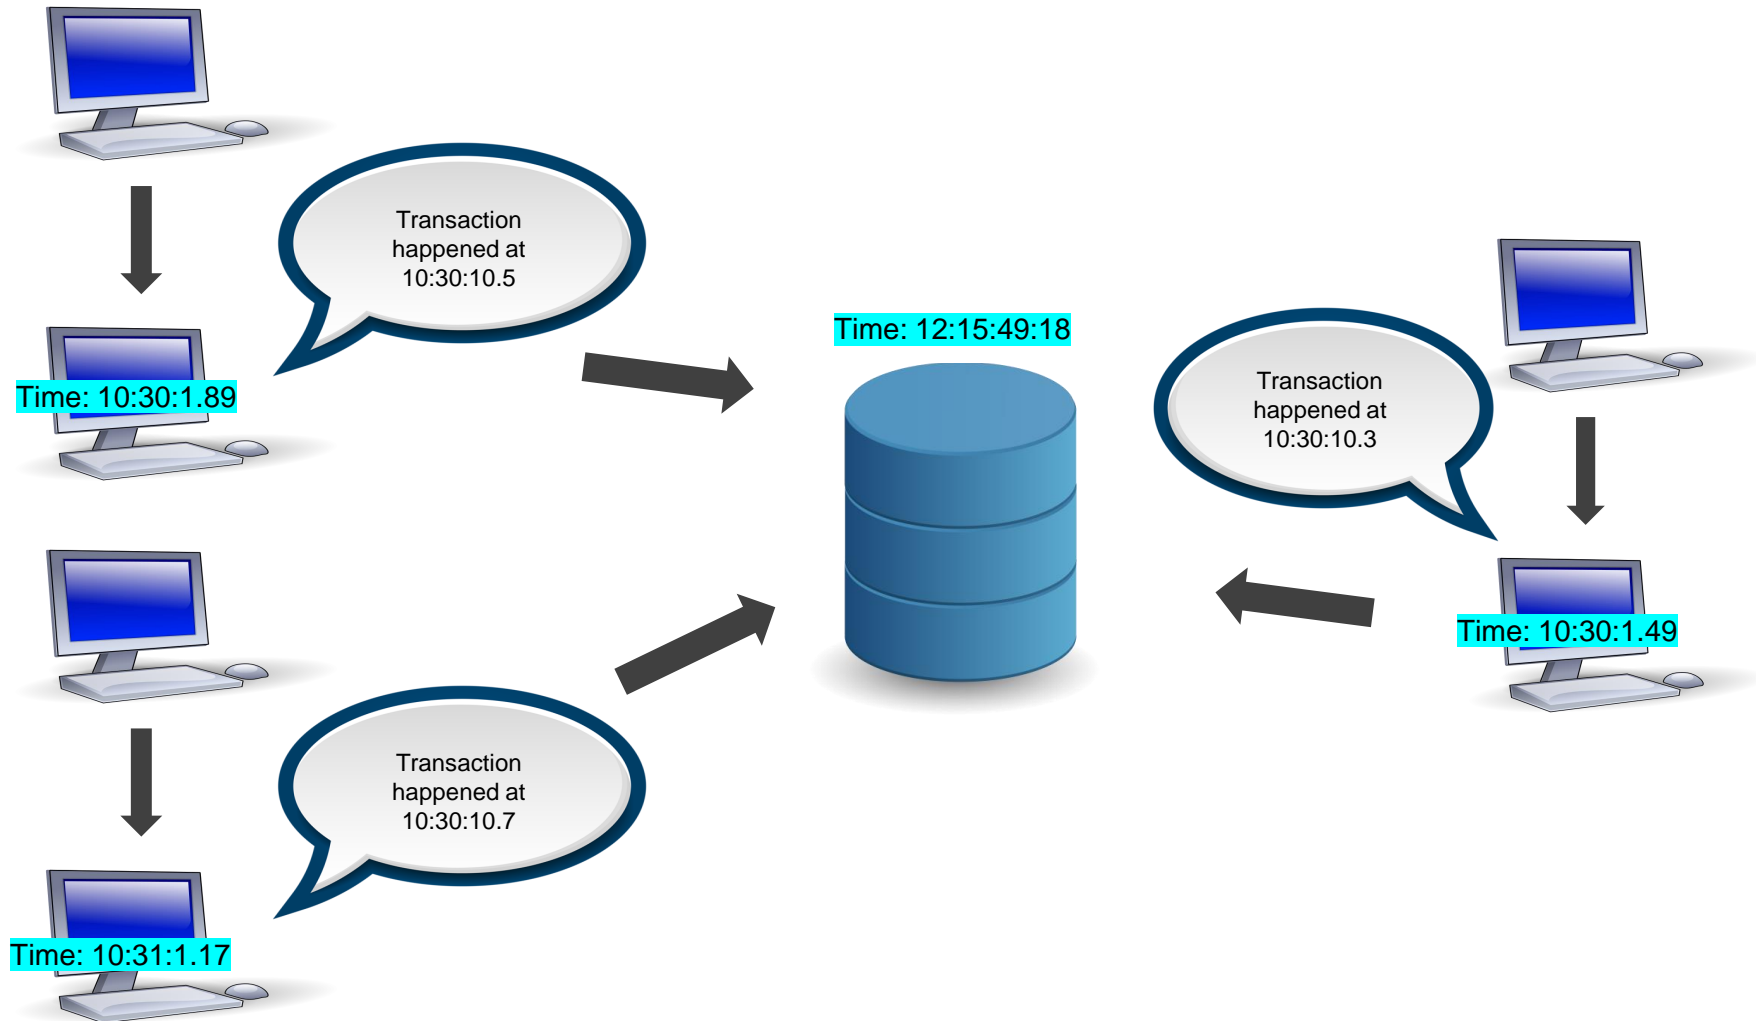
Senior Technical Staff Member

John Houston

Senior Software Engineer









- What can we do?

- Sync our clocks!
 - Network Time Protocol (NTP)
 - Precision Time Protocol (PTP)



What are the big differences?



PTP

- 1 nanoseconds (ns)
 - .000000001
- Hardware Timestamping
- Hierarchical approach

NTP


- 1 milliseconds (ms)
 - .001
- Software Timestamping
- Multiple server approach




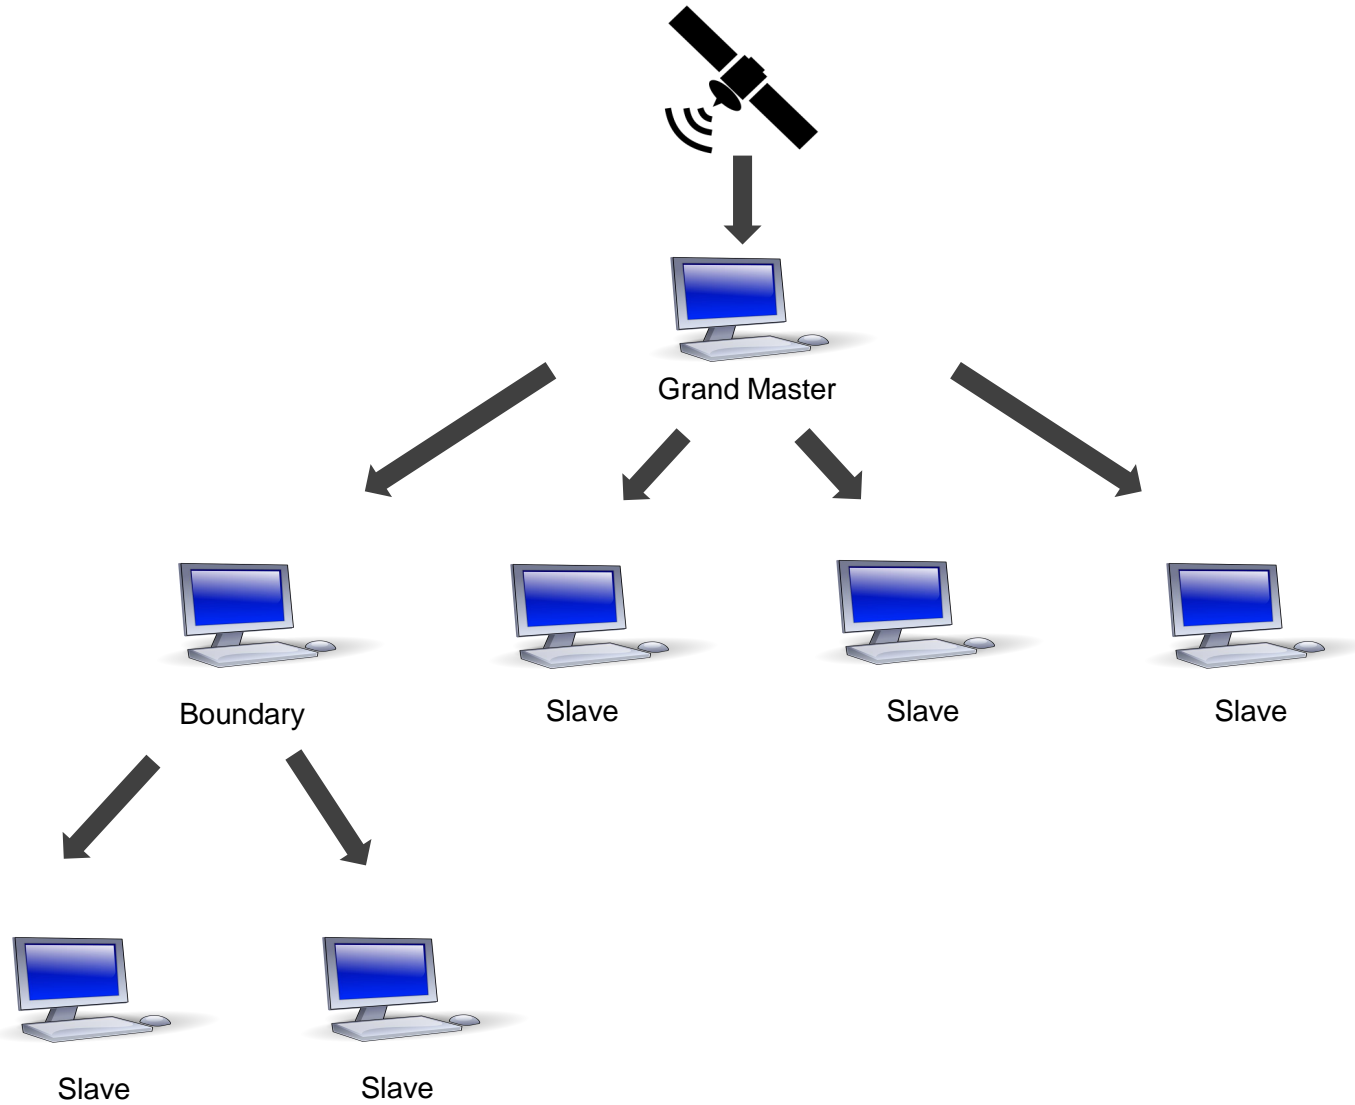
Why isn't NTP enough

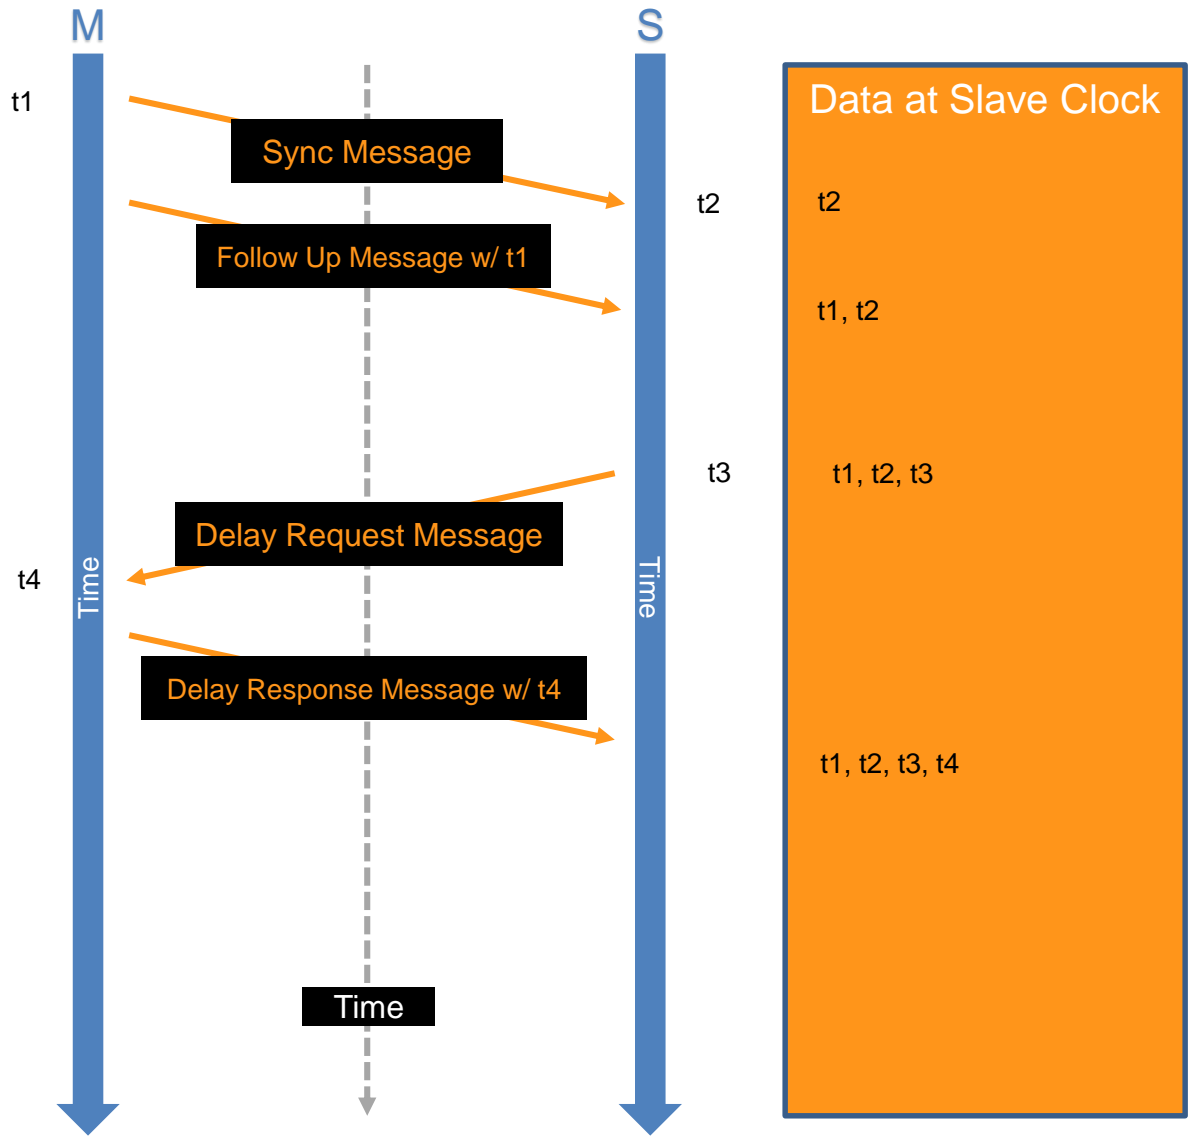
“European Securities Market Authority (ESMA) will introduce Regularity Technical Standard Number 36 on MiFID II, meaning all high-frequency trading deals **need to be timestamped within one microsecond.**”

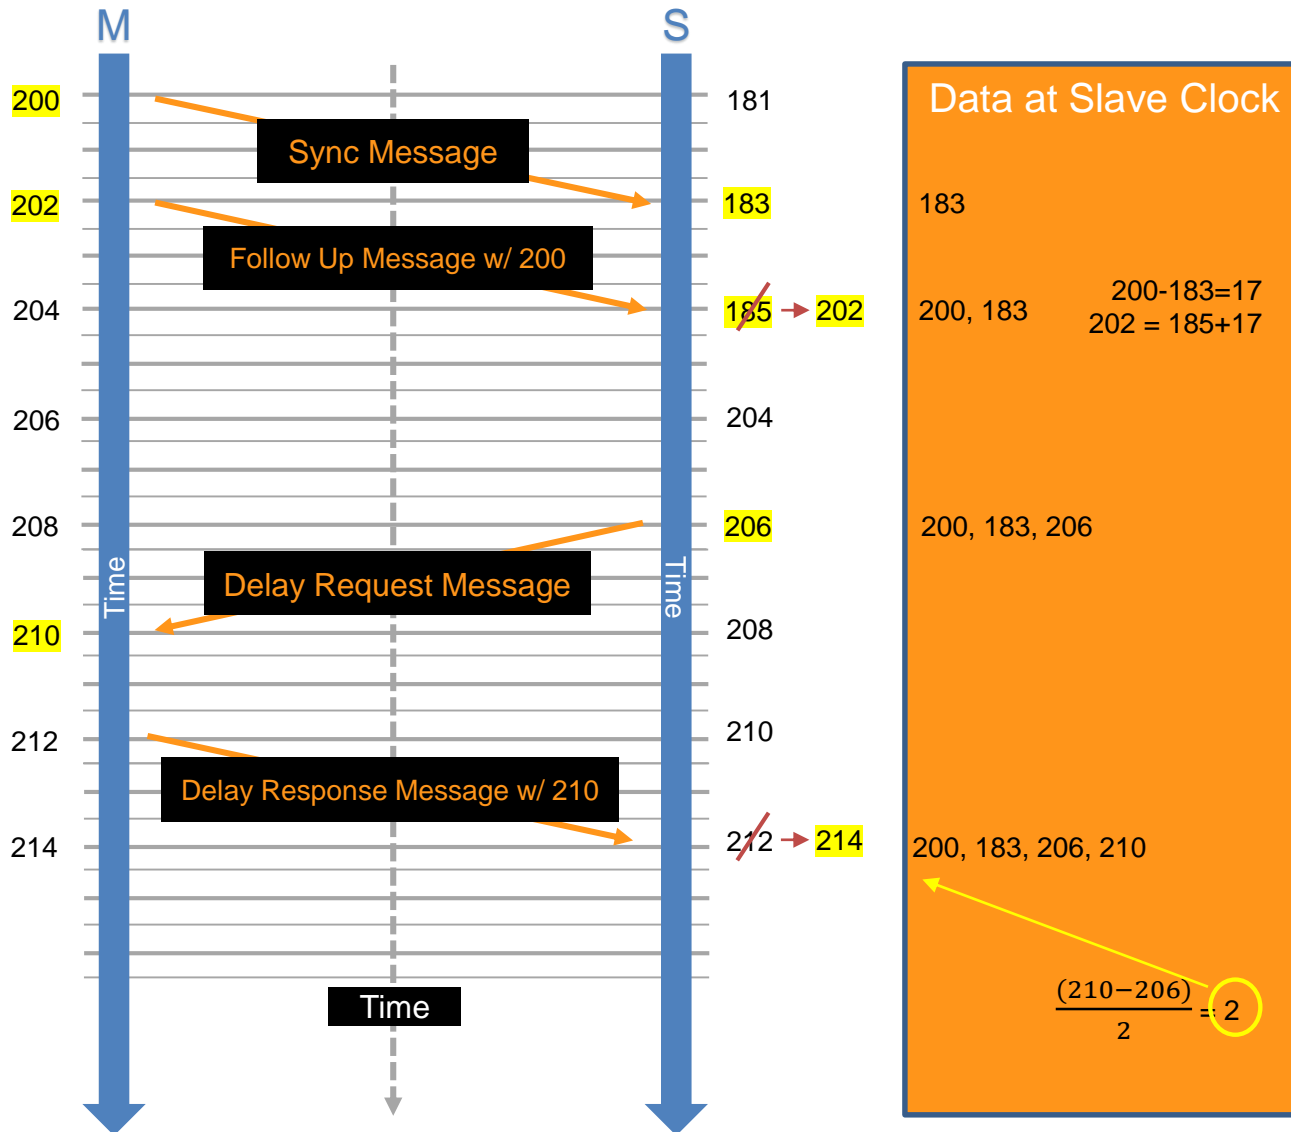
SEC Rule 613














Our Setup

- 2x IBM x3550 M3
 - Intel x540-AT2 PCI NICs
- G8264 PTP Switch
 - Adjusts the residence time
- Ubuntu Linux 17.1
- LinuxPTP 1.8
 - PTP4L






Our Setup

Servers on individual subnet and VLAN. No routing to the subnet.

Switch is not in production, very little traffic.





Our Tests

5 minutes, replicated for confirmation


- Baseline
 - Direct connected
 - Through a switch
- Stress Tests
 - CPU Saturation
 - NIC Saturation
- Response Tests
 - Failover



Baseline

Direct Connection

Direct Attached PTP




Spikes of +/- 1 ns



Baseline


G8264 Switch





Baseline

G8264 Switch



***Note Different Scales

— PTP Transparent On



What do we see?

- PTP off \rightarrow \pm ~90 nanoseconds offset
- PTP Transparent \rightarrow \pm ~58 nanoseconds offset

Our baseline is \pm ~58 nanoseconds offset




Stress Tests

- CPU Load
- NIC Load



Stress Tests

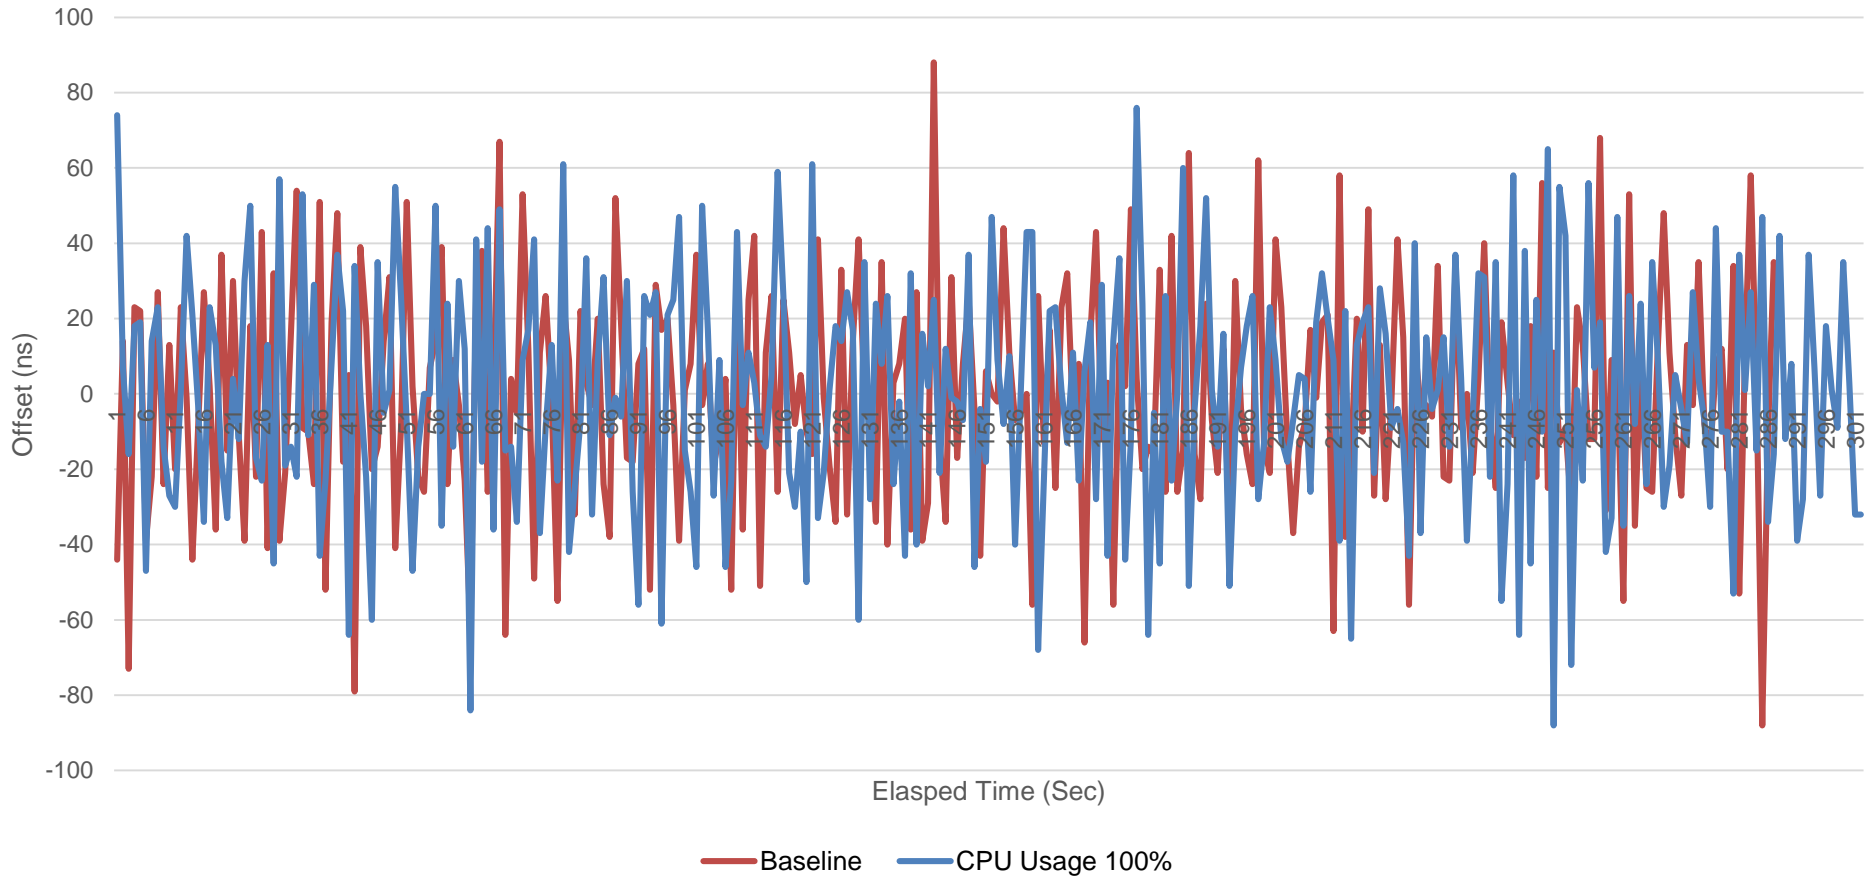
CPU Load





Stress Tests

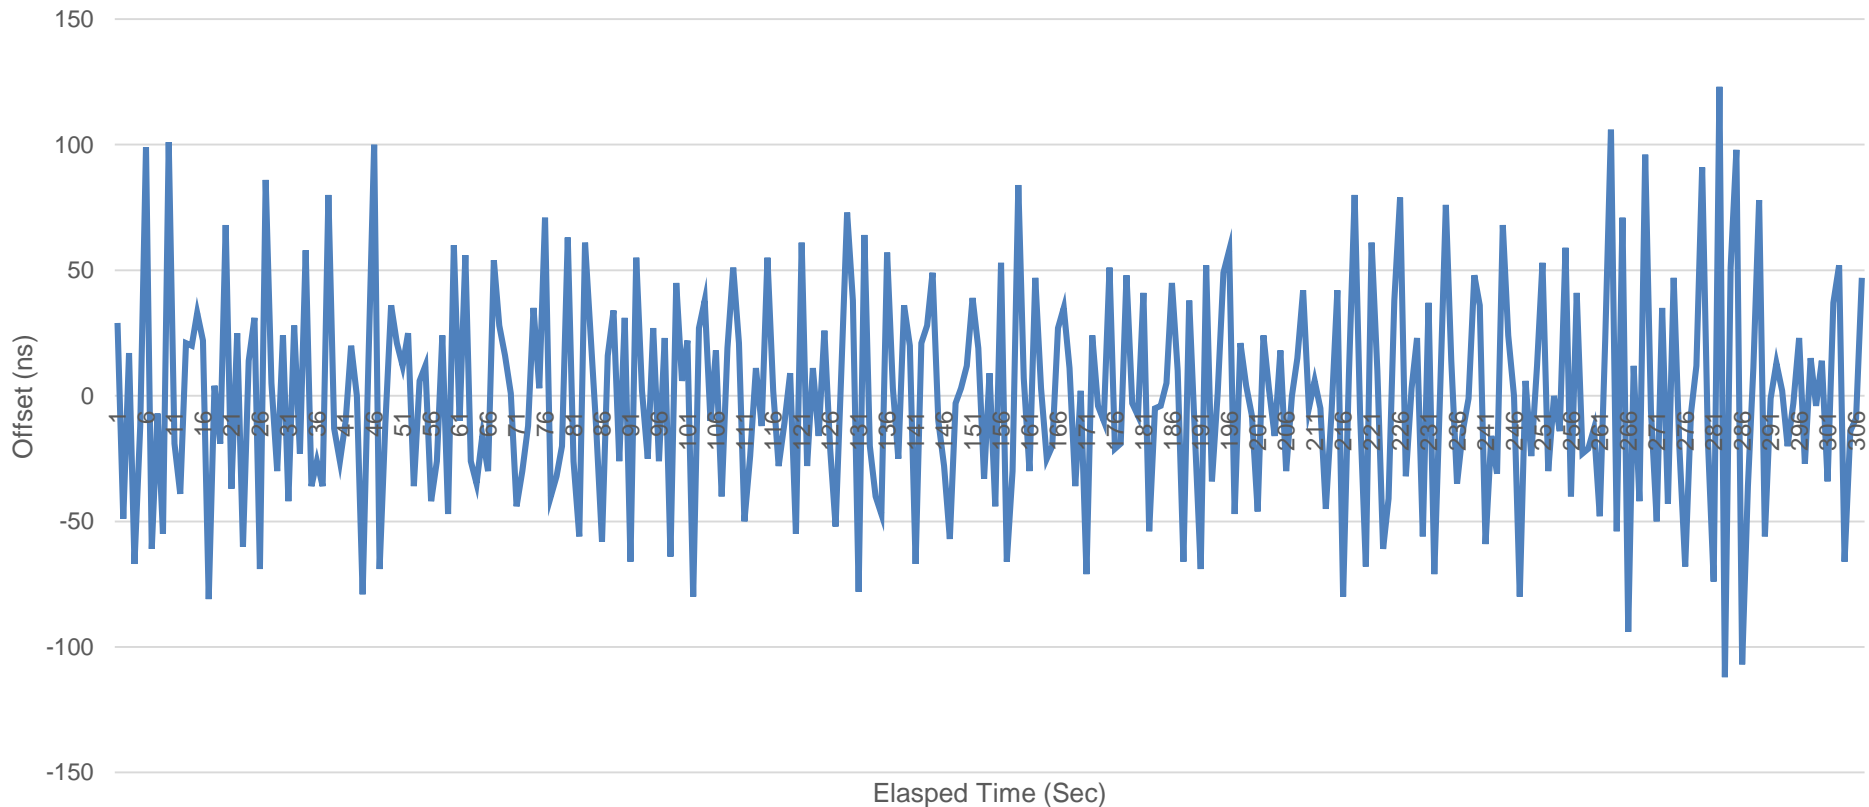
CPU Load





Stress Tests

NIC Load



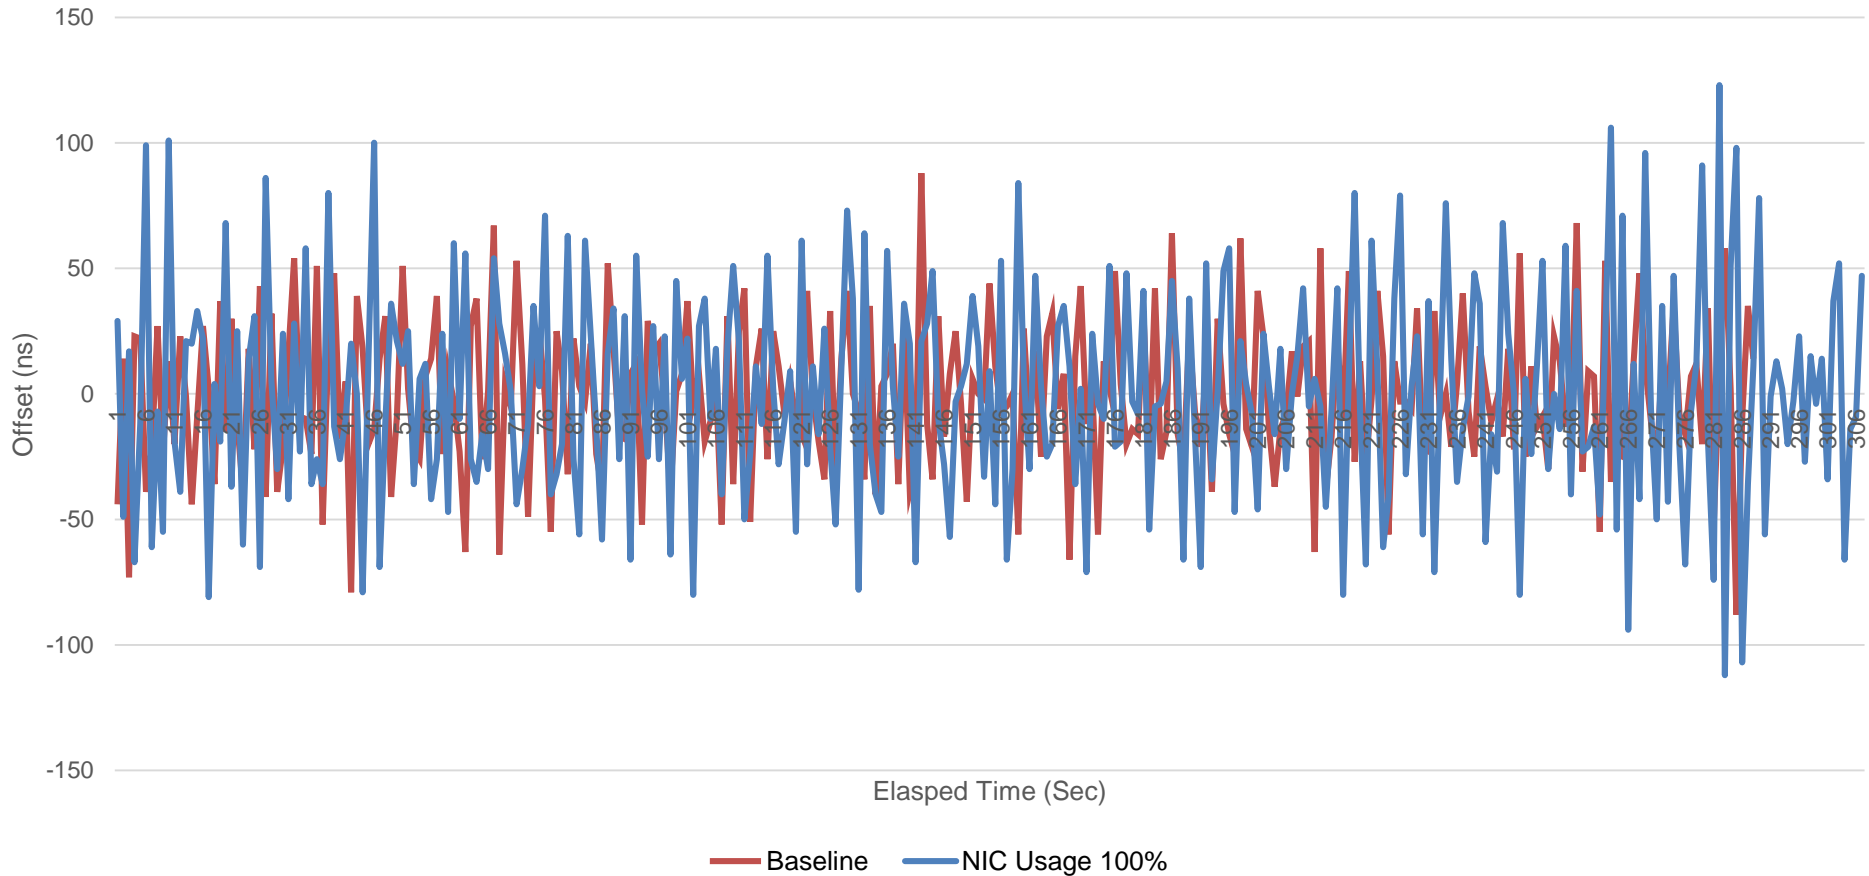
***Note Different Scales

— NIC Usage 100%



Stress Tests

NIC Load



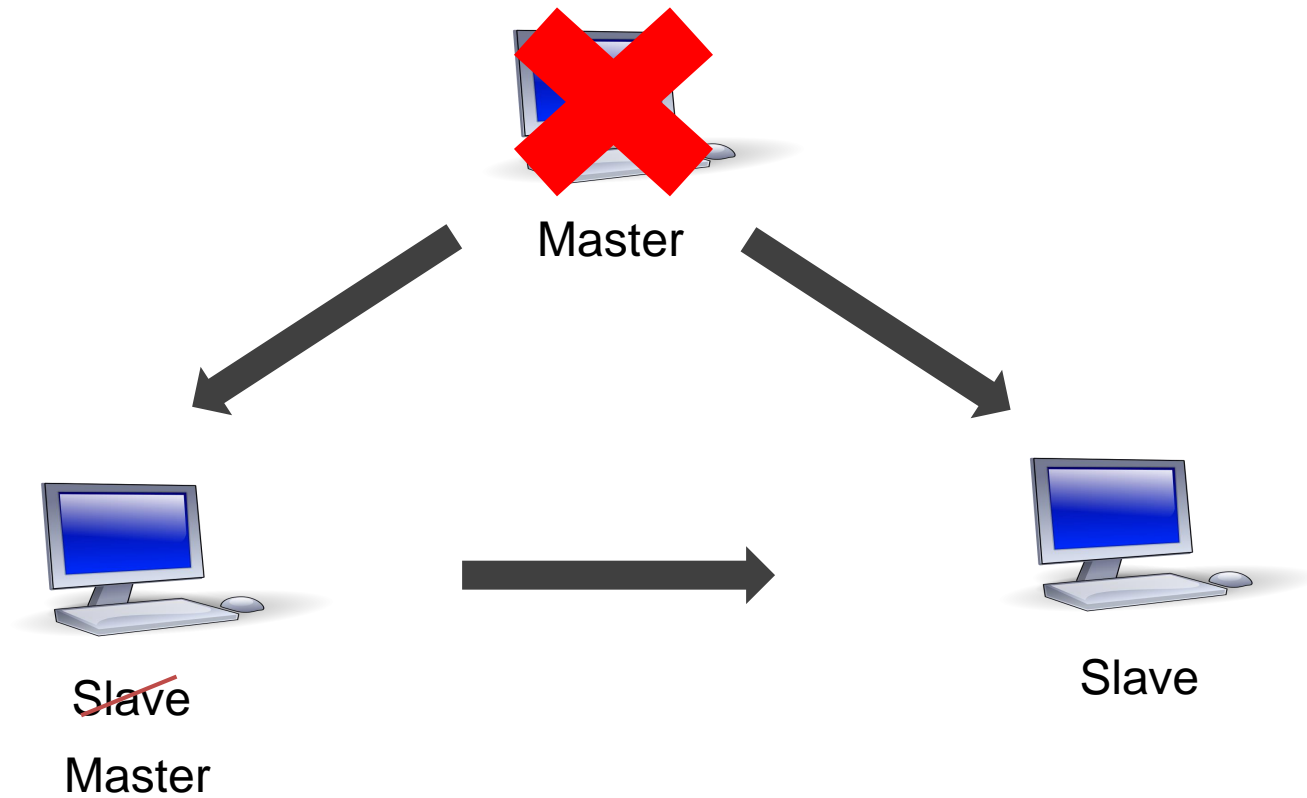


What does this mean?

- CPU has no effect on PTP
- NIC prioritizes PTP traffic



Failover Testing






Failover Testing

Conducted over an hour period with the grandmaster cycling every 2 minutes




Failover Testing





Failover Testing

Slave Time to Resynchronize with other GM





How does this apply to business?

- Accurate Records
- Encryption Tokens
- Time Sensitive information



Future Research

- Security
 - Rogue Clocks
 - Bad Packets
- Span-ability
 - Multisite
 - More clients/larger topology

