User Perception of 10GBASE-T Training time/Time-To-Link

IEEE P802.3bq 40GBASE-T Task Force

Dave Chalupsky, Intel
Pete Cibula, Intel

PHY Baseline Proposal Ad Hoc – 27 February 2014

What is Time-To-Link (TTL)?

- Time-To-Link (TTL): A system performance metric that characterizes and measures PHY behavior through autonegotiation and the 10GBASE-T startup sequence
  - Defined in 802.3 Clause 28, “Physical Layer link signaling for Auto-Negotiation on twisted pair” and 802.3 Clause 55, Subclause "55.4.2.5.14 Startup sequence“)

- One of two primary performance measures (along with BER) used to characterize 10GBASE-T physical layer link interoperability

Autonegotiation

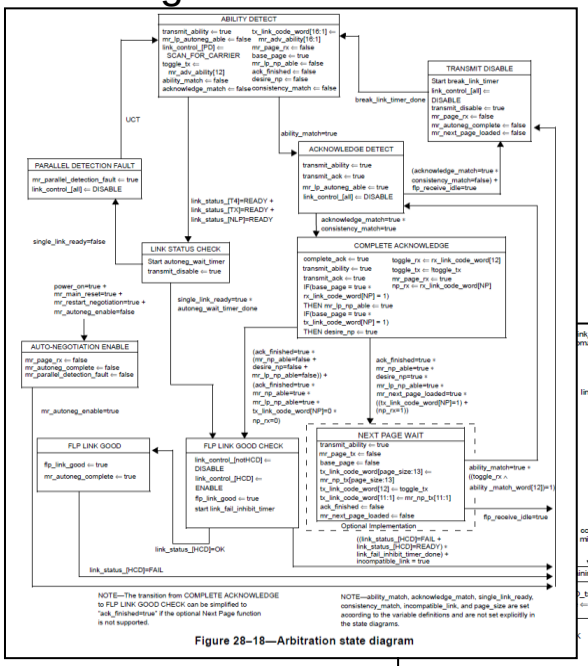

Figure 28–18—Arbitration state diagram

10Gb Startup

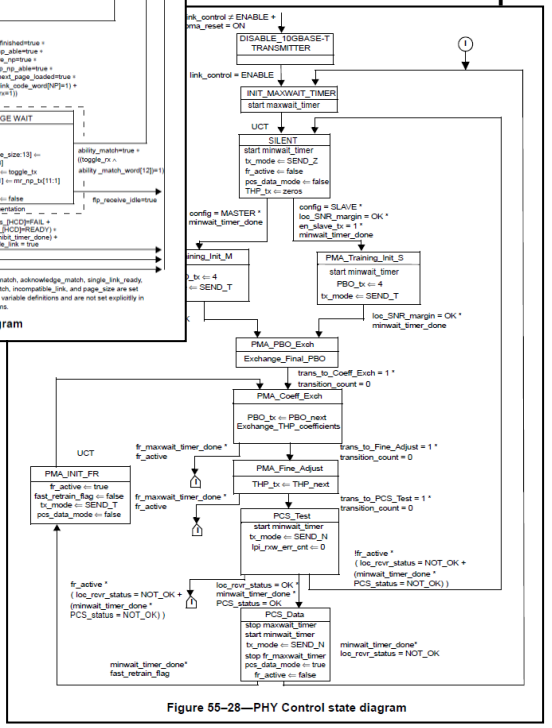

Figure 55–28—PHY Control state diagram

Why is it Important?

- Server networking drivers must meet 3rd-party certifications
- Example - Windows Hardware Quality Labs (WHQL) testing & certification “devfund”
  - A series of “device fundamentals” tests to evaluate the compatibility, reliability, performance, security and availability of a device in Windows OS
  - Includes many automated driver stress tests that execute multiple device resets
  - Long link times appear as a “failure” to these tests, which expect a link in 3s-4s based on 10Mb/100Mb/1Gb PHY performance

Server device fundamentals requirements

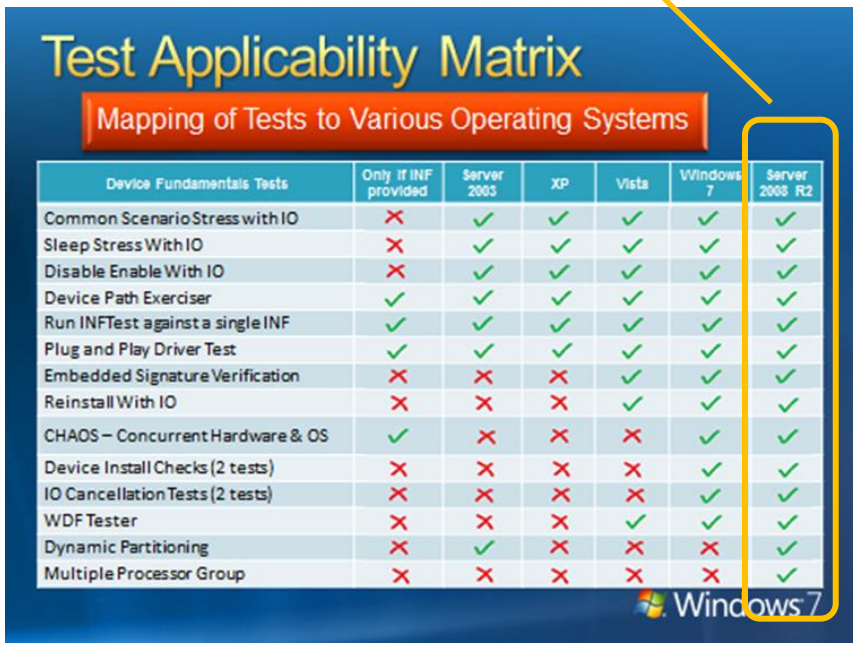

Device Fundamentals Tests	Only if INF provided	Server 2003	XP	Vista	Windows 7	Server 2008 R2
Common Scenario Stress with IO	✗	✓	✓	✓	✓	✓
Sleep Stress With IO	✗	✓	✓	✓	✓	✓
Disable Enable With IO	✗	✓	✓	✓	✓	✓
Device Path Exerciser	✓	✓	✓	✓	✓	✓
Run INFTest against a single INF	✓	✓	✓	✓	✓	✓
Plug and Play Driver Test	✓	✓	✓	✓	✓	✓
Embedded Signature Verification	✗	✗	✗	✓	✓	✓
Reinstall With IO	✗	✗	✗	✓	✓	✓
CHAOS – Concurrent Hardware & OS	✓	✗	✗	✗	✓	✓
Device Install Checks (2 tests)	✗	✗	✗	✗	✓	✓
IO Cancellation Tests (2 tests)	✗	✗	✗	✗	✓	✓
WDF Tester	✗	✗	✗	✓	✓	✓
Dynamic Partitioning	✗	✓	✗	✗	✗	✓
Multiple Processor Group	✗	✗	✗	✗	✗	✓

Source: Device Fundamentals Overview Presentation at lhv_devfund.pptx

- Long TTLs (>6s) can lead to device certification failures!

Link Interoperability Measurements

- Representative Link Interoperability metrics associated with TTL
  - Time-To-Link (Time to achieve link after link initiation event)
  - # Link Attempts (Number of attempts for each link)
  - # Link Drops (Number of link drops observed after link is established)
  - Clock Recovery (Master/Slave resolution)
  - TTL Distribution (% of links by link time)
  - Speed Downshift/Downgrade (Resolved speed if other than 10Gb/s)

Example: TTL Distribution and Master/Slave Resolution by Channel Length

- Example of 10GBASE-T TTL measured from 2m to 115m channels (9,790 links)
- TTL across 2m-100m
  - Average TTL = 7.6s
  - Average time in AN = 5s
  - Average time in training = 2.6s
- Note apparent loop timing trend towards MASTER preference with increasing channel length
- Very long TTLs (>15s) at 100m+ channels are associated with downshifts to 1Gb link speed

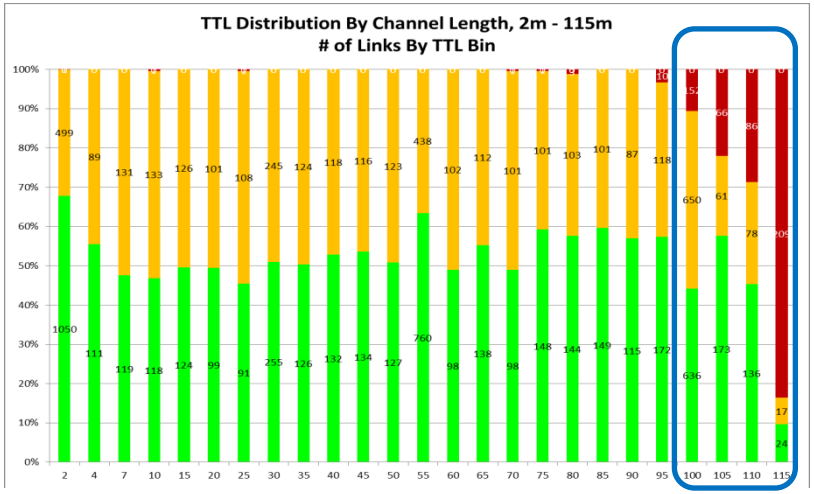

Green: TTL <= 7s; Yellow: 7s < TTL <=15s; Red: TTL >15s

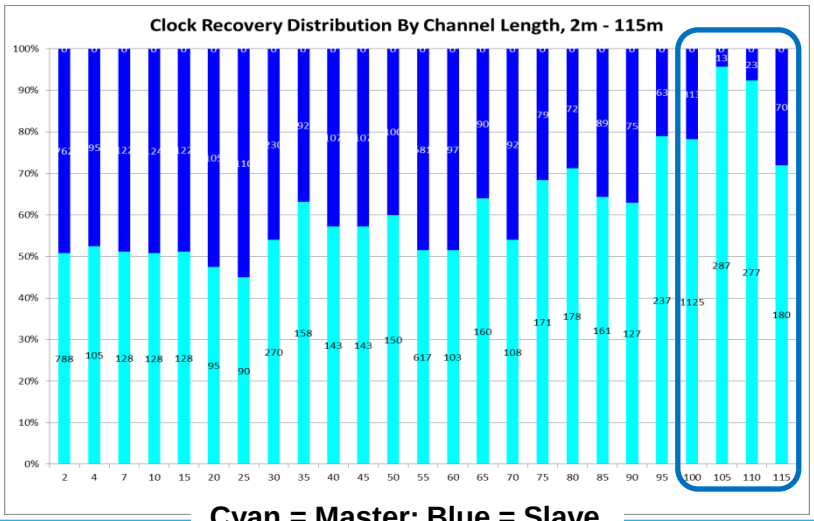

Cyan = Master; Blue = Slave

Time-To-Link Levers?

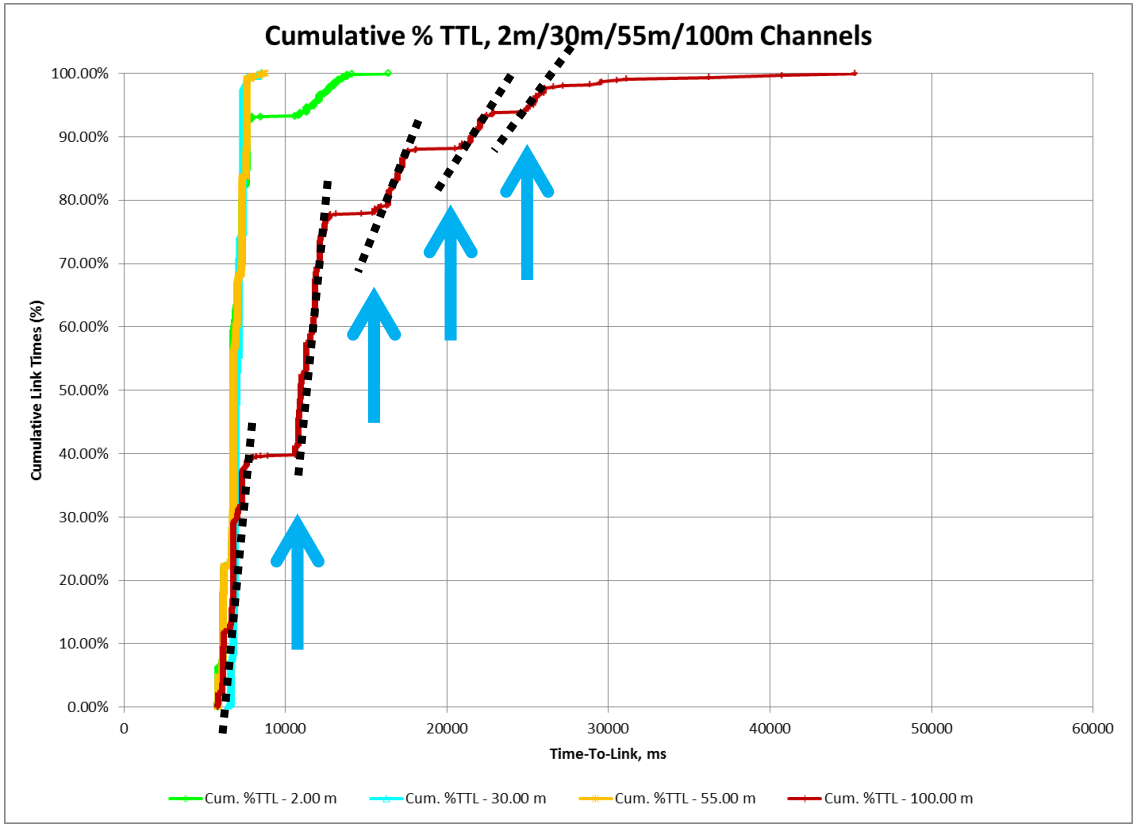

Suspected "Retrain"

Suspected "Retry"

- TTL is a combination of both autonegotiation and 10Gb startup behavior
  - Two sources of variability? "Retrain" (variability through 55.4.6.1) and "Retry" (return to 28.3.4)
  - Longest TTLs typically driven by multiple passes through the Clause Arbitration state diagram after failed training attempts

Observations from 10GBASE-T

- Channel topologies significantly affect the channel solutions realized by PHY DSP systems
  - “Peaky” impairments (return loss, crosstalk) appear to be a factor in link-trial-to-link-trial variability in the system solution
  - Transition region between RL/crosstalk-driven to IL-driven solutions
  - Channel lengths near 10GBASE-T PBO transitions
- PHY-specific responses to channel characteristics drive variability in autonegotiation and training time
  - Loop timing/clock recovery resolution
  - Time spent in 10GBASE-T startup states
- May have implications for both system performance and end-user experience
  - Potential to affect product time-to-market and customer ease-of-use

Considerations for 40GBASE-T

- Can autonegotiation and 40GBASE-T startup times be improved to be consistently less than or equal to 6s?
  - Improved loop timing?
  - Changes in 10GBASE-T startup state timing?
    - Example – Simple PBO scheme similar to that proposed in Wu_01a_0214_802.3bq_adhoc.pdf
  - Others?

Thank You!

Test Channels

- Focused channel selection using multiple cable types and lengths
  - 2m, 4m, 7m, 30m, 55m, 90m and 100m are “standard” channels for both TTL and BER
  - Other channel lengths (typically 5m increments) are used to check for consistent link behavior over a range of PHY channel solutions (different PBOs, operating margin, delay/delay skew, etc.)
- Includes direct connection, 2-connector, and 4-connector topologies
- Test channel matrix will (of course) be modified for 40GBASE-T