The “WHY” Behind a 3m NO-FEC Copper Cable Solution

Joel Goergen – Cisco Systems

IEEE802.3 July 2015 Plenary
Version r7

Supporters

- Brad Booth – MicroSoft
- Mike Andrewartha – MicroSoft
- Upen Reddy – Cisco Systems
- Vineet Salunke – Cisco Systems
- Ali Ghiasi – Ghiasi Quantum LLC
- Scott Sommers – Molex
- Tom Palkert – Molex
- Chris Roth – Molex
- Scott Irwin – MoSys, Inc.
- Vasu Parthasarathy - Broadcom

Overview

- Rack Definition, Cable Routing, and Cable length
- A case for no FEC
- Summarizing a simplified solution

Cisco Rack Series R42610 / Based on EIA-310-D

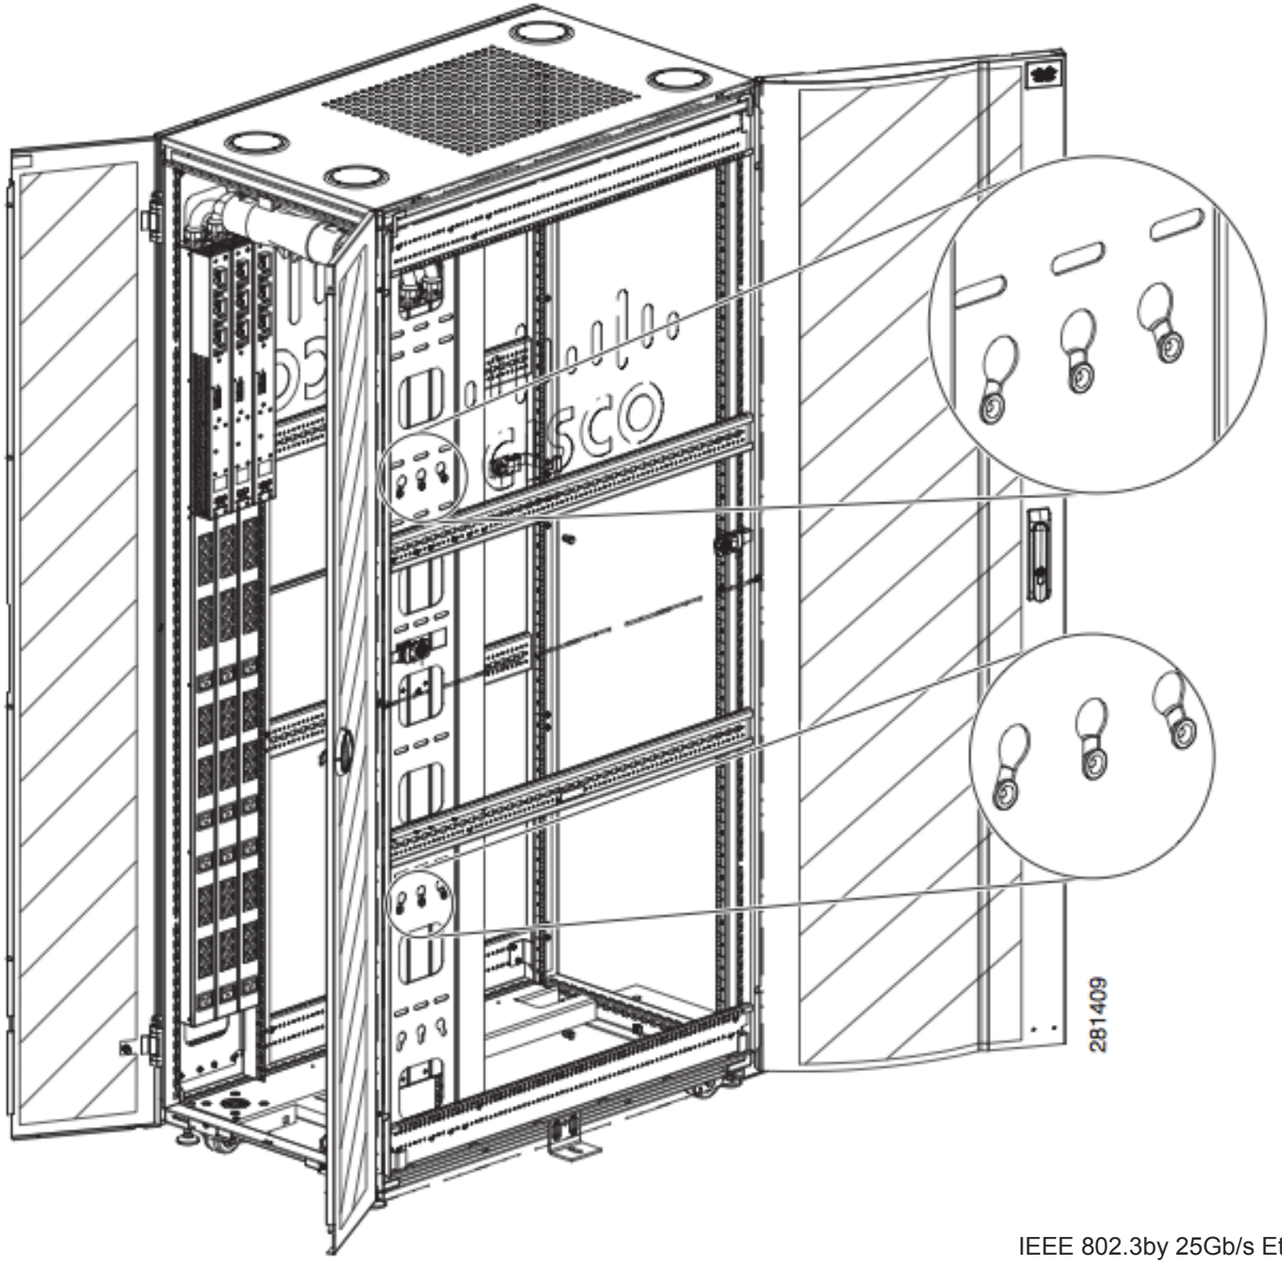

- H x W x D
  2000mm x 610mm x 1102mm
- Side Panels
- Equipment Mounting Capacity
  42RU
- Static Load Capacity
  2100 lb / 954 kg
- Application – well suited for hooded systems and other types of hot aisle / cold aisle containment.
- EIA-310-D type racks are well suited for the new California building codes introduced on Jan 1, 2015, requiring all installations to support hot aisle / cold aisle containment.

Cabling Installation – Front to Back

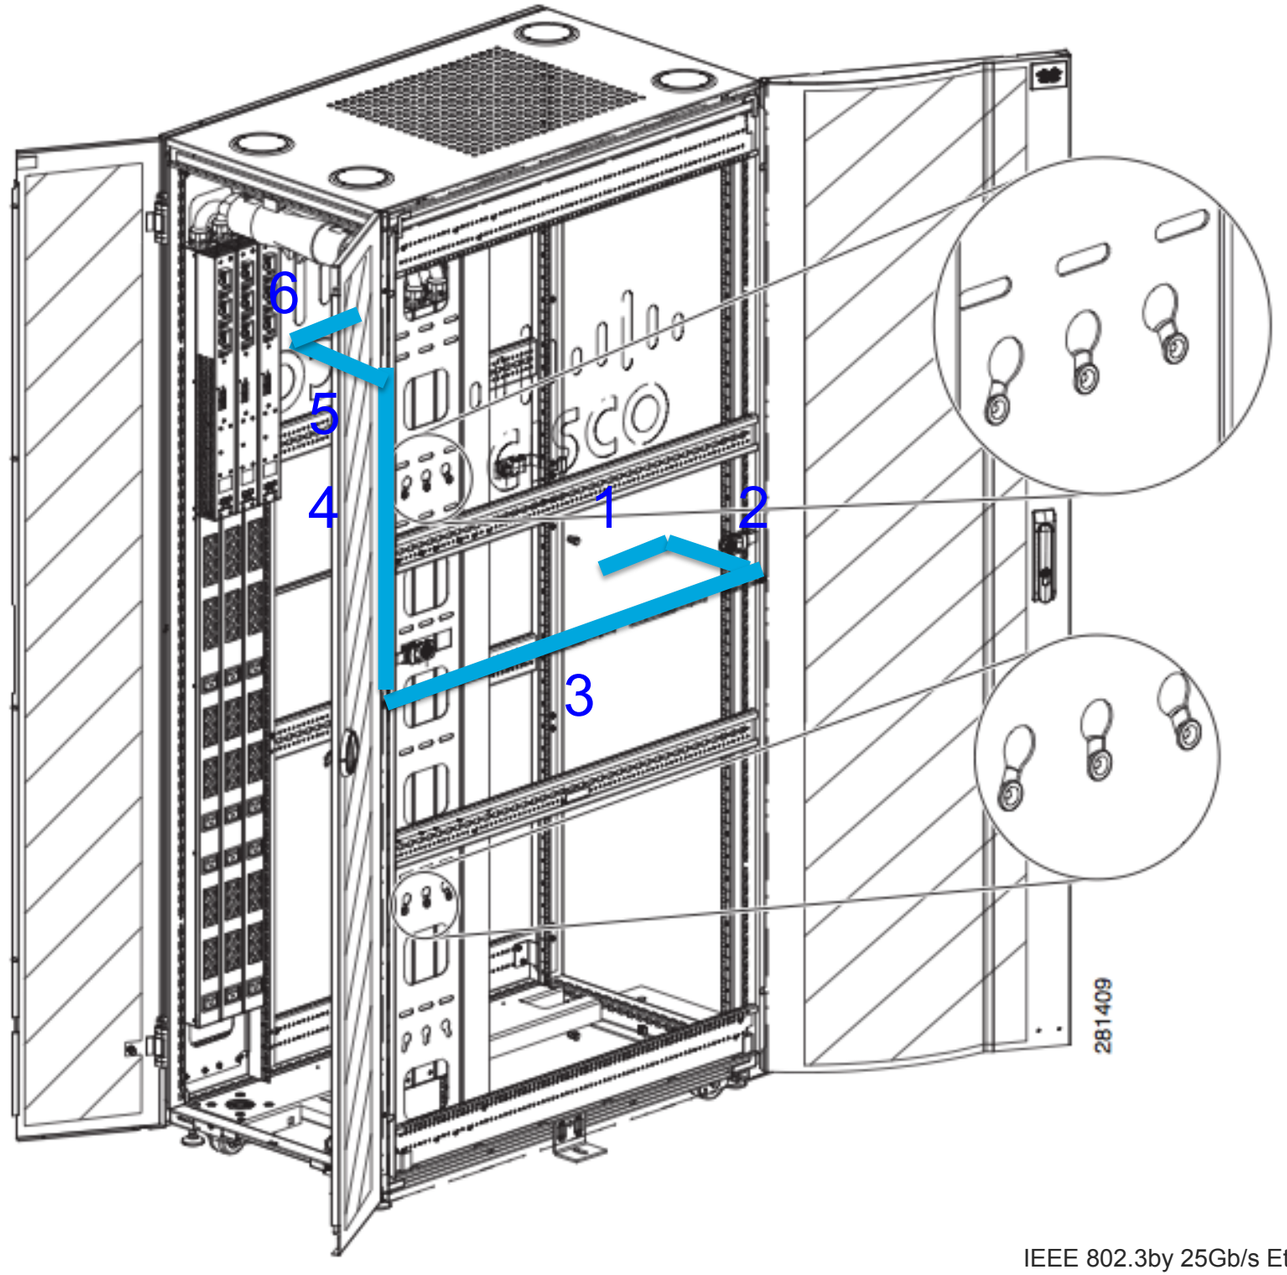

- Cabling is an art form. There is a lot of pride that goes into a professional installation
- Every detail is considered, with focus given to layout, labeling, and debug.
- Consider this common strategy
  - 1 – 152mm
  - 2 – 304mm
  - 3 – 1066mm
  - 4 – 914mm
  - 5 – 304mm
  - 6 – 152mm
- This real life case is 2892mm.

Cabling Installation – Top to Bottom

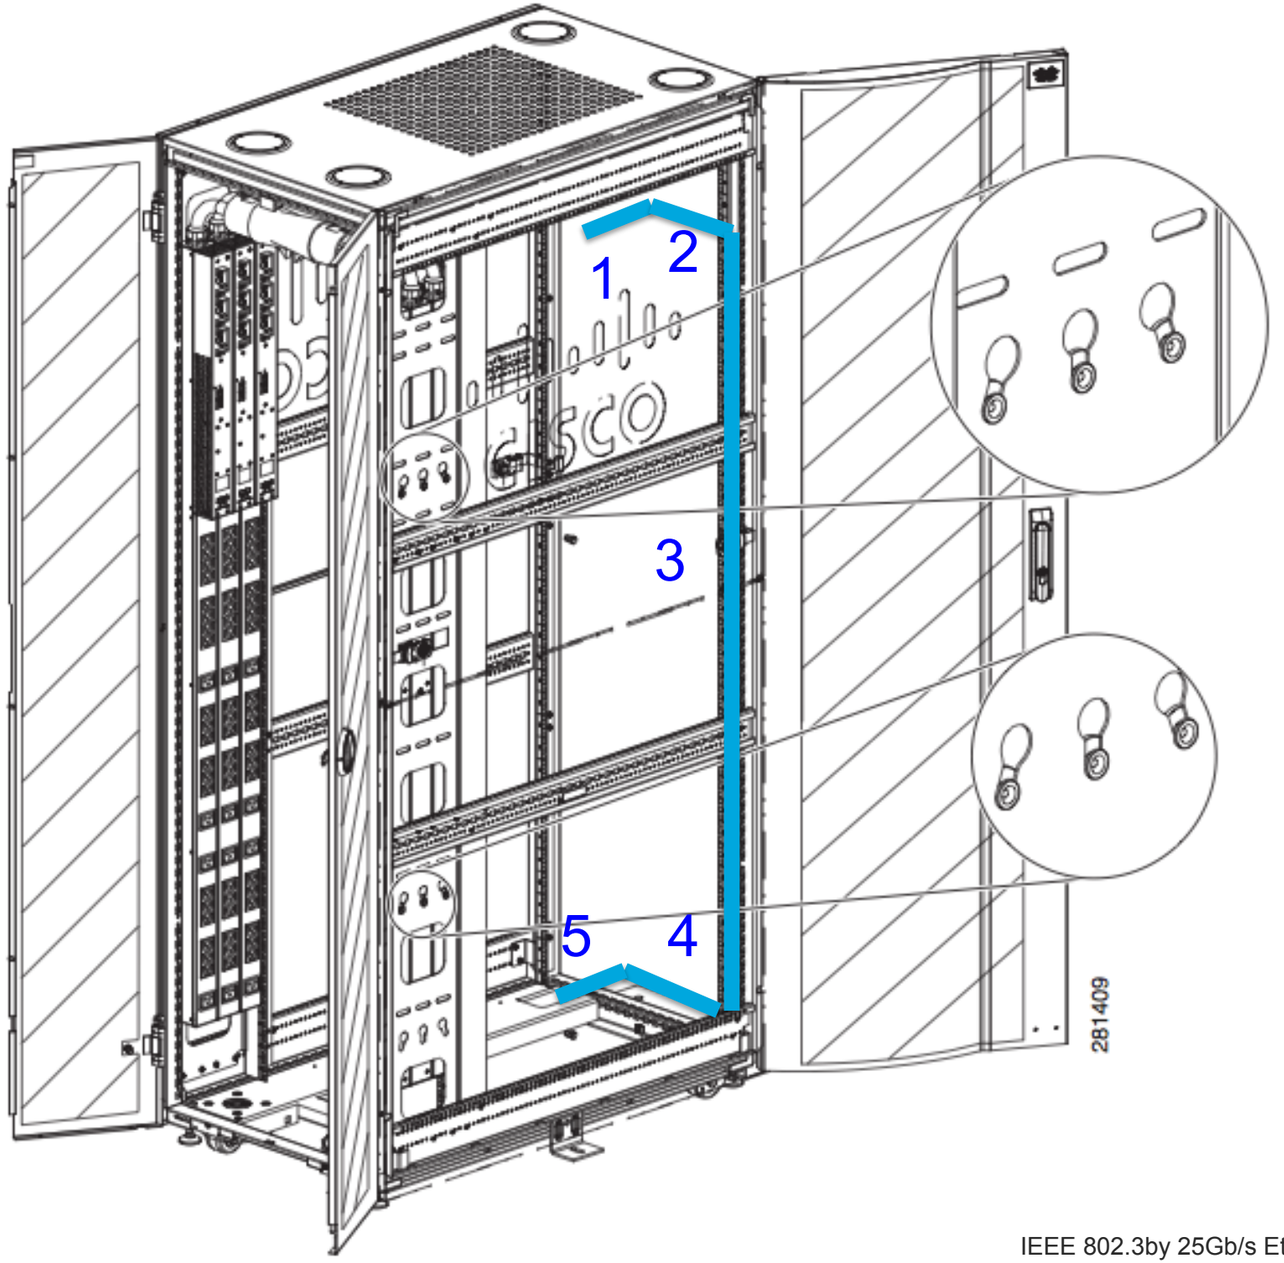

- Consider this common strategy
  - 1 – 152mm
  - 2 – 304mm
  - 3 – 1778mm
  - 4 – 304mm
  - 5 – 152mm
- This real life case is 2690mm.

Cabling Layout Summary

- Racks are becoming deeper to accommodate cooling efficiency in higher power installations. 1102mm depth and deeper are common place.
- We should not be dictating end user implementation by ignoring top-to-bottom and front-to-back applications.
- 2000mm length cable, regardless of signaling and coding, ignores a lot of standard cabling implementations common place in core, data center, and enterprise applications.
- Up to 3000mm length cable should be the specified length in IEEE p802.3by to optimize for broad end user applications, eliminating the CA-N 2000mm cable.
- Eliminating the 2000mm cable prevents the need for a cable dependent setting to eliminate the possibility of cross connecting 3000mm and 2000mm solutions in the rack, since FEC is required for one and not the other.
- Having both FEC and no-FEC solutions in the same rack will significantly increase end user debug time. Eliminating the 2000mm CA-N solution and applying no-FEC or optional FEC to the 3000mm CA-S solution will help to minimize debug time.

FEC Adds Power

- 25Gbps SERDES Cores currently have an amazing receiver structure that has all the knobs already turned on. Removing FEC will not add any significant power to the receiver functions to compensate for the loss in current implementations. This is my experience and knowledge of widely available SERDES cores.
- Generic implementation of the KR FEC in a 25Gbps SERDES is 200mW.
- Generic implementation of the RS FEC in a 25Gbps SERDES is 300mW.
- 100,000 ports with an increase in .2W is 20kW
  - 175,200kWh per year
  - $19,272 per year using $0.11 per PG&E kWh … rates vary across the world from $0.07 to $0.33 USD.

FEC adds Latency

- CL108, RS-FEC @ 25G = 250 ns block delay (5280 bits) = 250 ns.

  http://www.ieee802.org/3/25GSG/public/adhoc/architecture/ran_081214_25GE_adhoc.pdf

  Referred to CL91 100G RS-FEC from IEEE 802.3bm (striped over 4 lanes).

  CL108 is the new clause in IEEE P802.3by, single lane 25G RS FEC
- CL74, KR FEC @ 25G = 82 ns block delay (2112 bits) = 82 ns.

  http://www.ieee802.org/3/25GSG/public/adhoc/architecture/ran_081214_25GE_adhoc.pdf
- Adding 82ns into the port path is similar to adding 500 inches of circuit board trace.
- Gate count is a factor in asic design, but mostly ignored as it is likely any generic asic core will include those gates with zero latency and zero power adds to maximize core use. This is not true for FPGA designs where the used gates could be redirected for other logic use.
- There are applications in high performance computing, trade, and finance that require optimized Ethernet solutions with minimum MAC to MAC latency.

Summary

- 25Gbps SERDES core technology has robust receiver technology. Adding FEC is not likely to increase that.
- FEC adds power … system designers and end users would like to minimize power.

Summary

- Offering two primary solutions, where the CA-L specification requires FEC, and the CA-S specification has no FEC or an optional FEC.
- Power (and thermal from it) is a primary concern for systems designers and end users.
- Cable length has to cover a broad range of in-rack end user implementations.

Thank you.