ETHERNET FABRICS 101

- Jon.Hudson@Brocade.com
  - Global Solutions Architect

© 2012 Brocade Communications Systems, Inc.
Proprietary Information. Improper Use with result in Flogging

AGENDA

Why?

What?

TRILL

SPB

And Update from IETF83

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Why Fabrics?

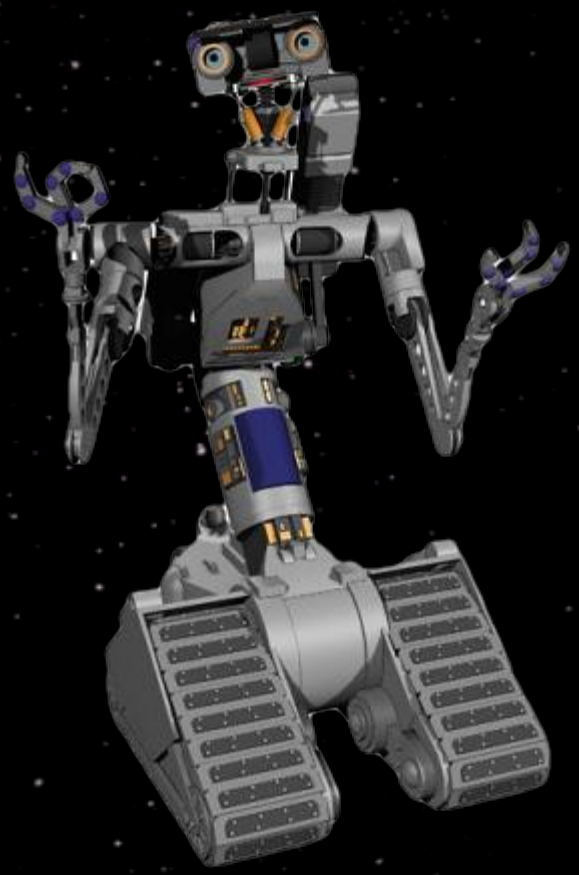

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Fabrics: Proven Stage of Evolution

Telco: Humans->P2P->Crossbar->Matrix->Switched Fabrics

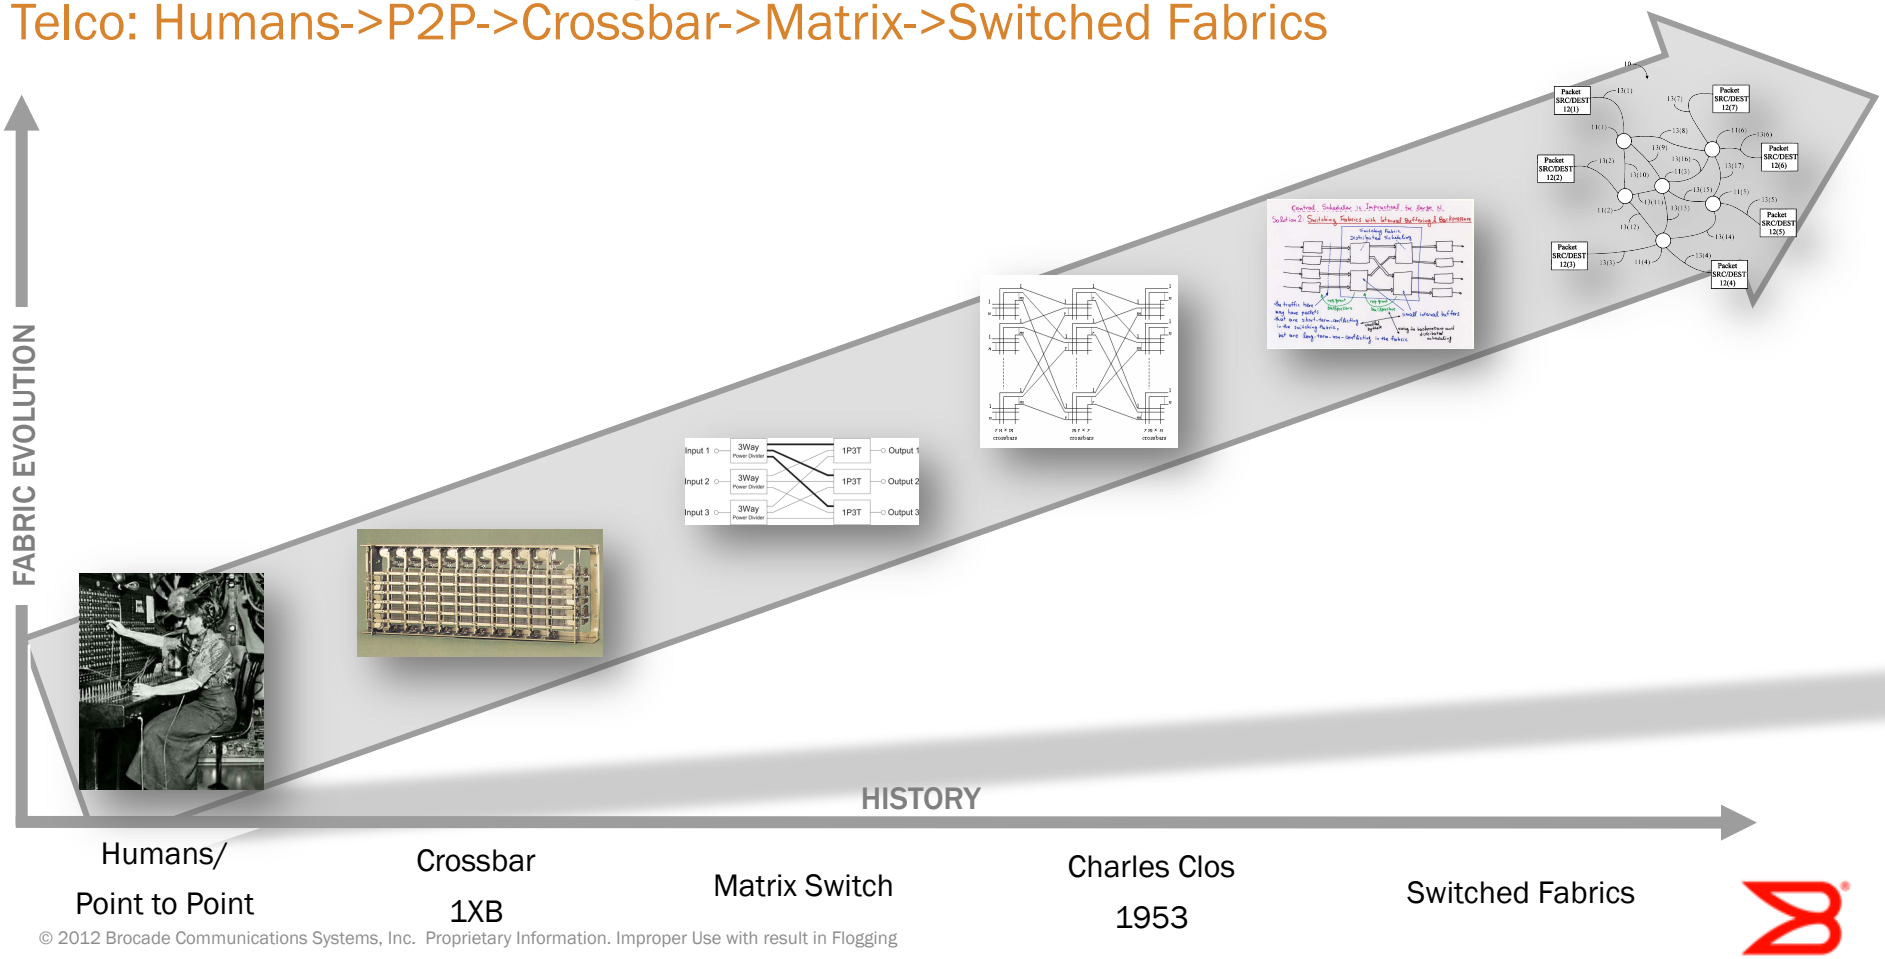

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Fabrics: Proven Stage of Evolution

Backplanes: P2P->Crossbar->Matrix->Switched Fabrics

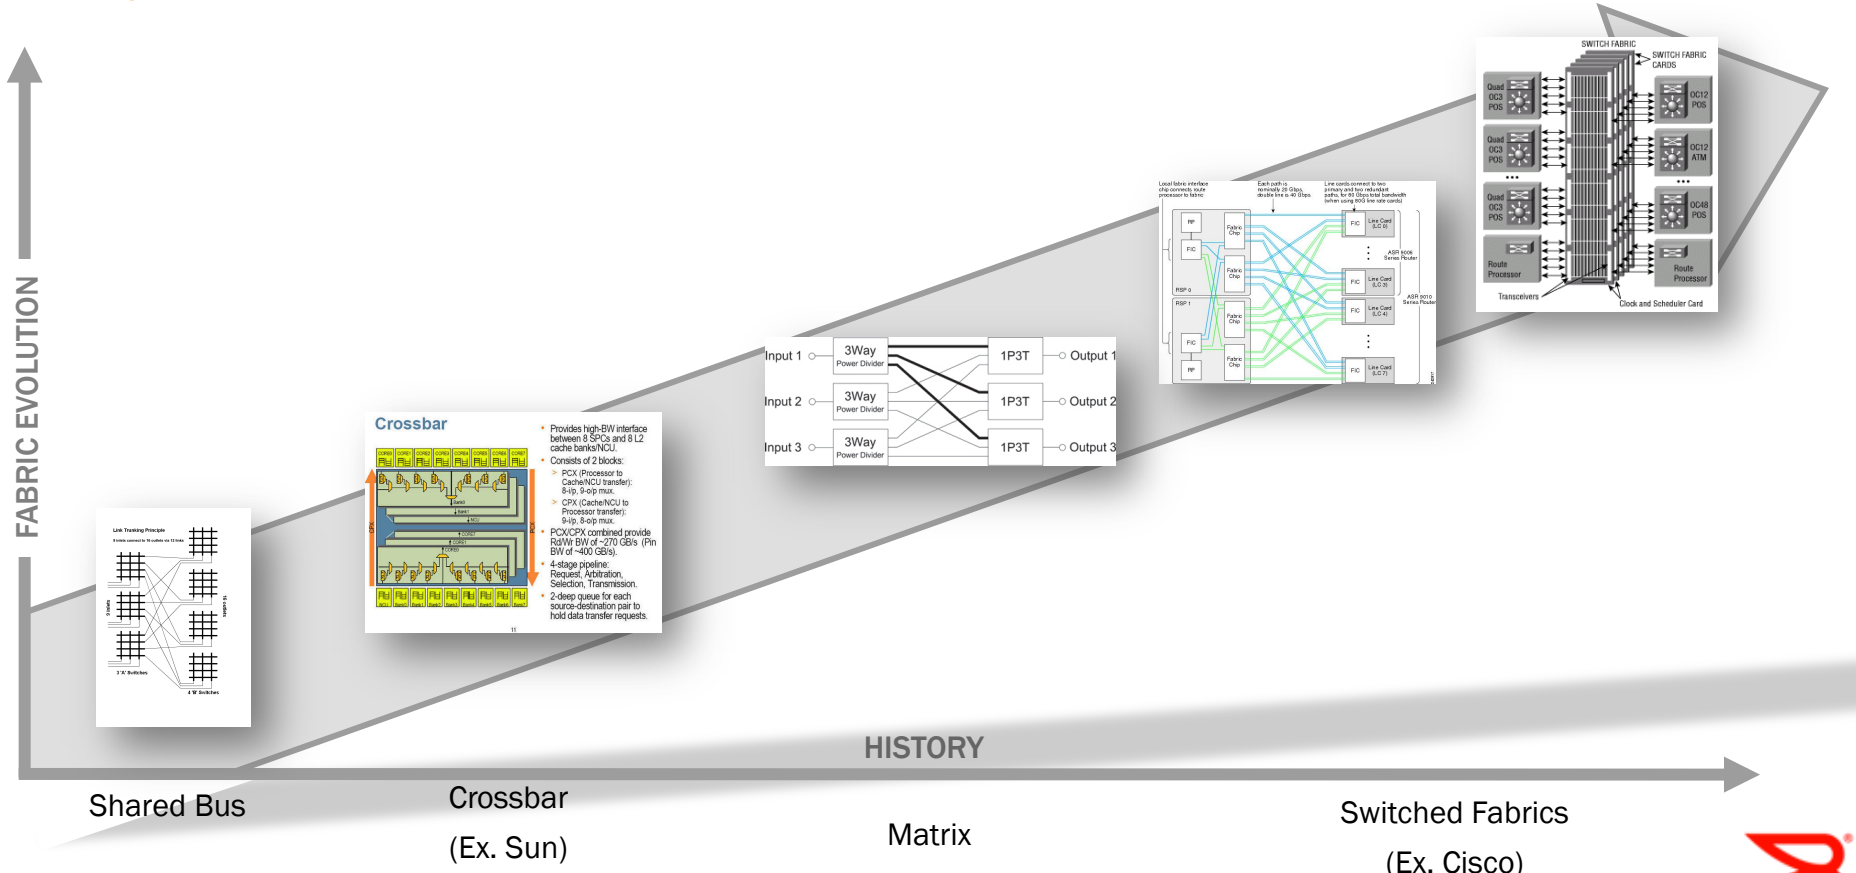

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Fabrics: Proven Stage of Evolution

PCI: PCI-BUS -> PCI-X Crossbar -> PCIe Switched Fabric

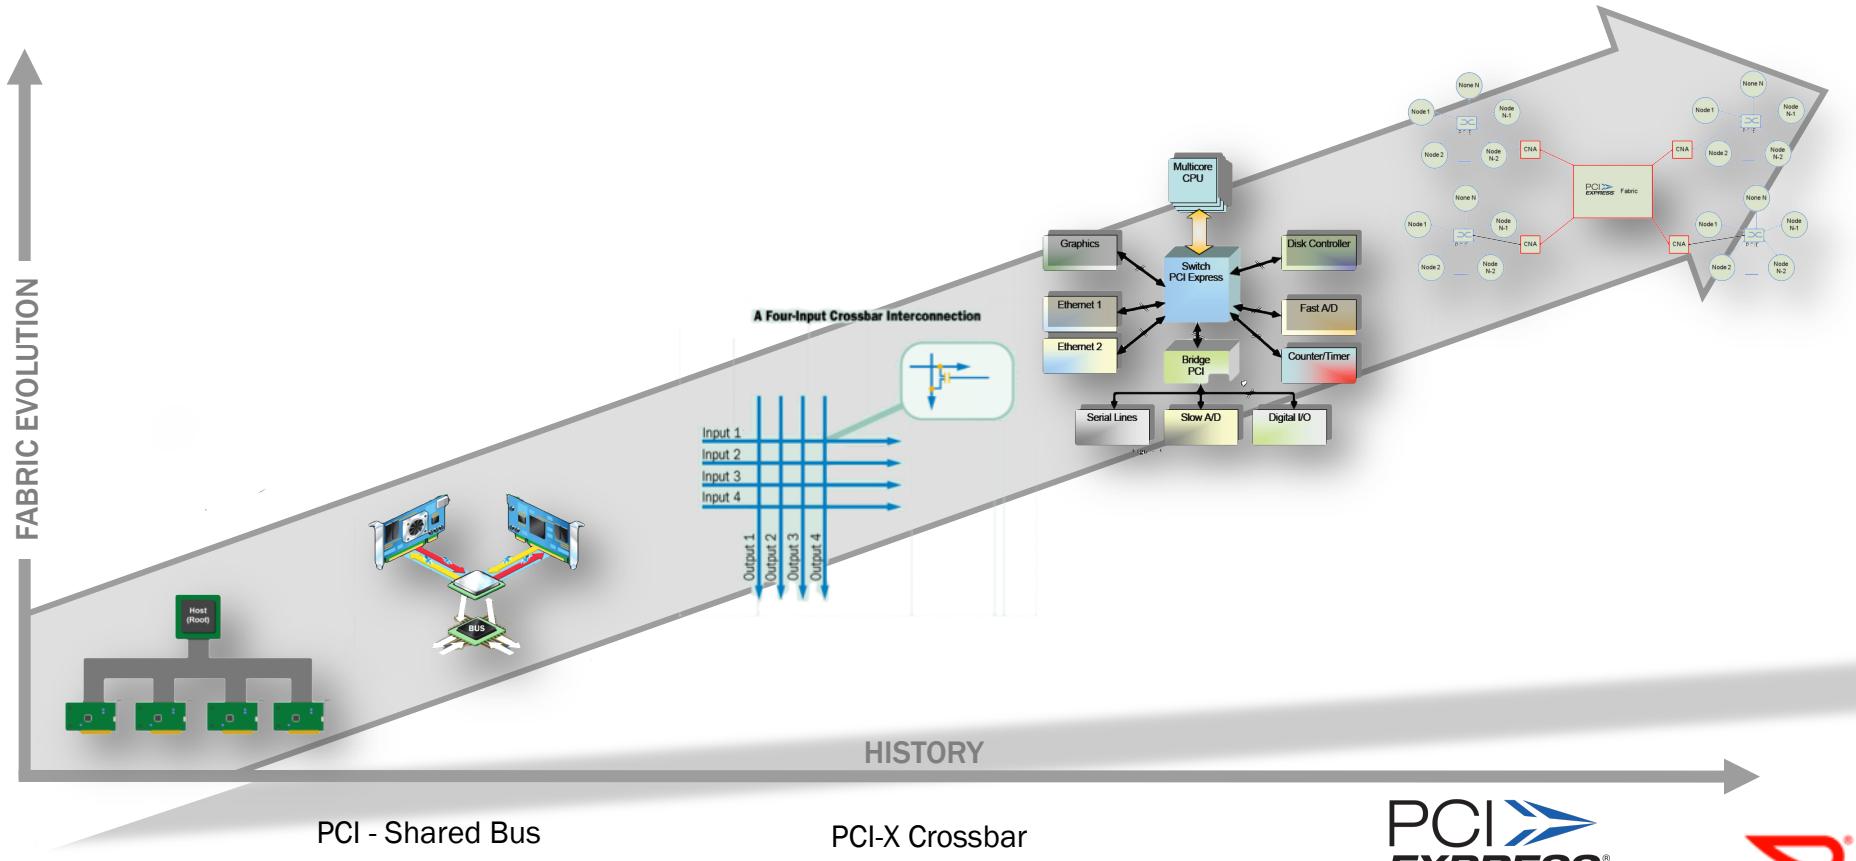

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Scaling Virtual Server Environments

Technical Challenges Today

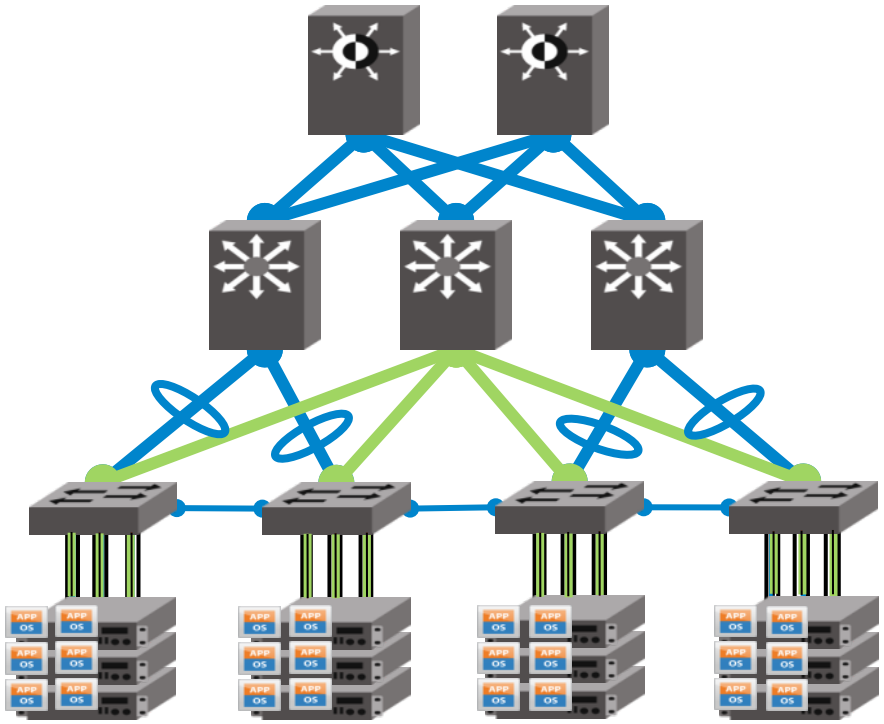

Layer 2: only 1 active path

- STP disables other paths

40% of links not used

Increase utilization using MSTP (spanning tree per VLAN)

- Increases complexity
- Creates multiple single-path networks; limits sphere of mobility

Link failure

- STP reconvergence—network is down
- Broadcast storms stress network

Layer 3 as an alternative

- Greater complexity; higher cost
- VM mobility limited to rack

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

TRILL-Based Networking

Transparent Interconnection of Lots of Links

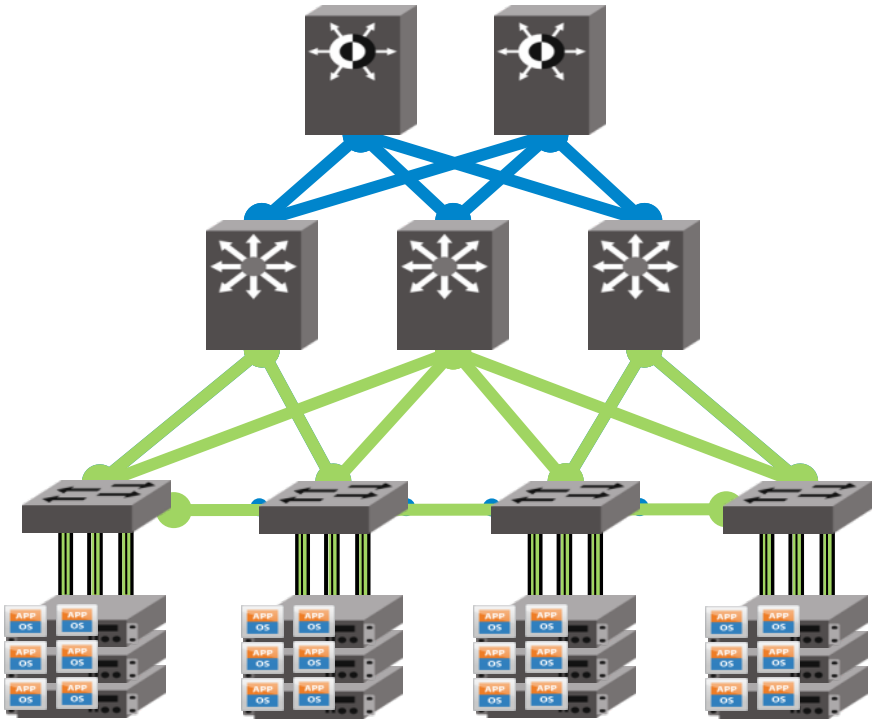

Layer 2: multiple active paths

- TRILL calculates multiple routes

100% of links used

ECMP handles even distribution

- ASIC-level frame-based load balancing

Link failure

- Within a trunk triggers no reconvergence
- Total path failure has very quick recovery

TOR Layer 3 complexity NOT needed

- Simple flat Layer 2 networks
- Fast low-latency switching in hardware

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

BROCADE

Inspired by a Real Life Incident

In November 2002, Beth Israel Deaconess Hospital in Boston, Massachusetts, had a total network meltdown:

- Their network took four days of heroic efforts to be restored to an operational state! In the mean time the staff was reduced to using paper and pencil.
- Beth Israel Deaconess had grown by acquiring various clinics and just plugged all those bridged networks together.
- The article in Boston's primary newspaper specifically mentioned "Spanning Tree Protocol" as the problem!
- Radia Perlman, who invented spanning tree over 25 years ago, decided it was time to come up with a better way.

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper use will result in Flogging

What/Why/Who TRILL?

TRILL

TRansparent Interconnection of Lots of Links

- TRILL WG Charter
  - http://www.ietf.org/dyn/wg/charter/trill-charter.html
- A standard specified by the IETF (Internet Engineering Task Force)
- TRILL Working Group co-chaired by
  - Donald E. Eastlake 3rd, Huawei Technologies
  - Erik Nordmark, Cisco Systems

RBridge – Routing Bridge

- Device that implements TRILL

RBridge Campus –

- A network of RBridges, links, and any intervening bridges, bounded by end stations / layer 3 routers.

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

What/Why/Who TRILL?

TRILL Provides

- –transparent forwarding using optimal paths with zero configuration,
- –safe forwarding even during routing transients,
- –support for multi-pathing for unicast and multicast traffic, and
- –improved scalability.

Who invented TRILL?

Radia Perlman of Intel, a major contributor to link-state routing, and the inventor of DECnet Phase V from which IS-IS was copied, as well as the inventor of the Spanning Tree Protocol.

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

EVOLUTION OF TRILL

1. Radia Perlman's idea is accepted by the IETF and the TRILL WG is formed. Basic idea is shortest path transparent frame routing using IS-IS and encapsulation with a hop count.
2. Basic idea unchanged + improved data plane address learning & VLAN support
3. Basic idea unchanged + improved data plan address learning & VLAN support + MTU robustness
4. To Come: continued additive enhancements with OAM, etc.

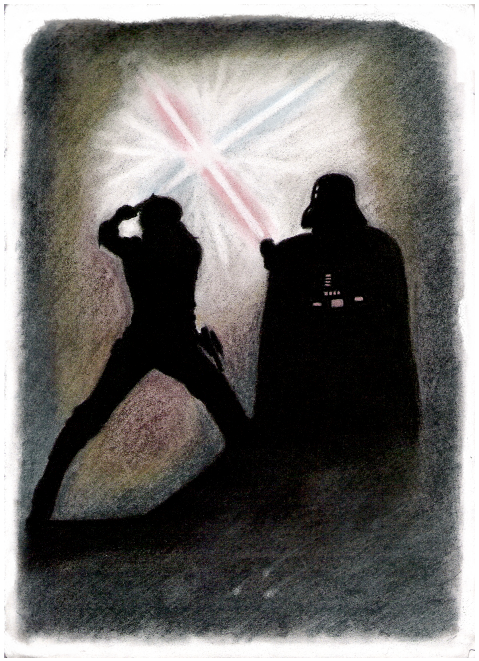

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Evolution of SPB

1. Radia Perlman's idea are rejected by IEEE 802.1. They say there isn't a problem, TRILL is a terrible idea, spanning tree is good, routing sucks, and hop counts (TTLs) are evil.
2. Whoops, there is a problem. They start 802.1aq for spanning tree based shortest path bridging. Still say TRILL is terrible, routing sucks, hop counts are evil.
3. Whoops, spanning tree doesn't hack it. They copy a little of using IS-IS and nicknames from TRILL but don't actually do routing. Still say TRILL is a terrible idea and hop counts are evil.
4. Whoops, we can't multipath enough. Try to multipath more. Link agreement protocol etc. is a kludge. Try to find some way to add hop counts to SPB. Still say TRILL is a terrible idea.

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Ethernet Fabric

Feature:

- Link Speed Agnostic
- Data Center Bridging (DCB)
- TRILL/SPB
- Convergence Ready
- Single Logical Switching

Benefit:

- Flexible to meet Bandwidth Needs
- Lossless Ethernet Capability
- No Spanning Tree
- Suitable for all Applications
- Node and Link Level Redundancy

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

BROCADE

Convergence

Convergence:

- The ability of a single network infrastructure to support the needs of multiple technologies
- An aspect of Ethernet fabrics that allows them to scale to meet the needs of different applications

Networks that are convergence ready support:

Storage —iSCSI, FCoE, FCIP

Voice—VoIP

Video—Streaming or multicast

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

BROCADE

Simple and Automated

Feature:

- Self-forming
- Auto-learning
- Auto-healing
- Automated Migration of Port Profiles (AMPP)
- Hypervisor Aware

Benefit:

- Easier to Deploy & Scale
- Limited Manual Intervention
- Rapid Outage Recovery
- Dynamic Security
- Application Awareness

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

BROCADE

Evolutionary

Feature:

- Hardware-based ISL Load Balancing
- SAN Connectivity
- Dynamic Service Insertion
  - Fabric Extension
  - Security Integration
  - Routing

Benefit:

- Efficient use of Bandwidth
- Investment Protection
- Extended Private Cloud
- Preserve Security Policies

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

BROCADE

STANDARDS, TERMS, AND TECHNOLOGIES

TRILL, SPB, Flat Networks, and Convergence

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use will result in Flogging

TRILL—Transparent Interconnect of Lots of Links

Overview

Devices are Routing Bridges (RBridges)

Data plane uses TRILL protocol

Control plane uses IS-IS Layer 2 link-state routing protocol

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

The Network with Spanning-Tree

Limited Virtualization and Storage Optimization

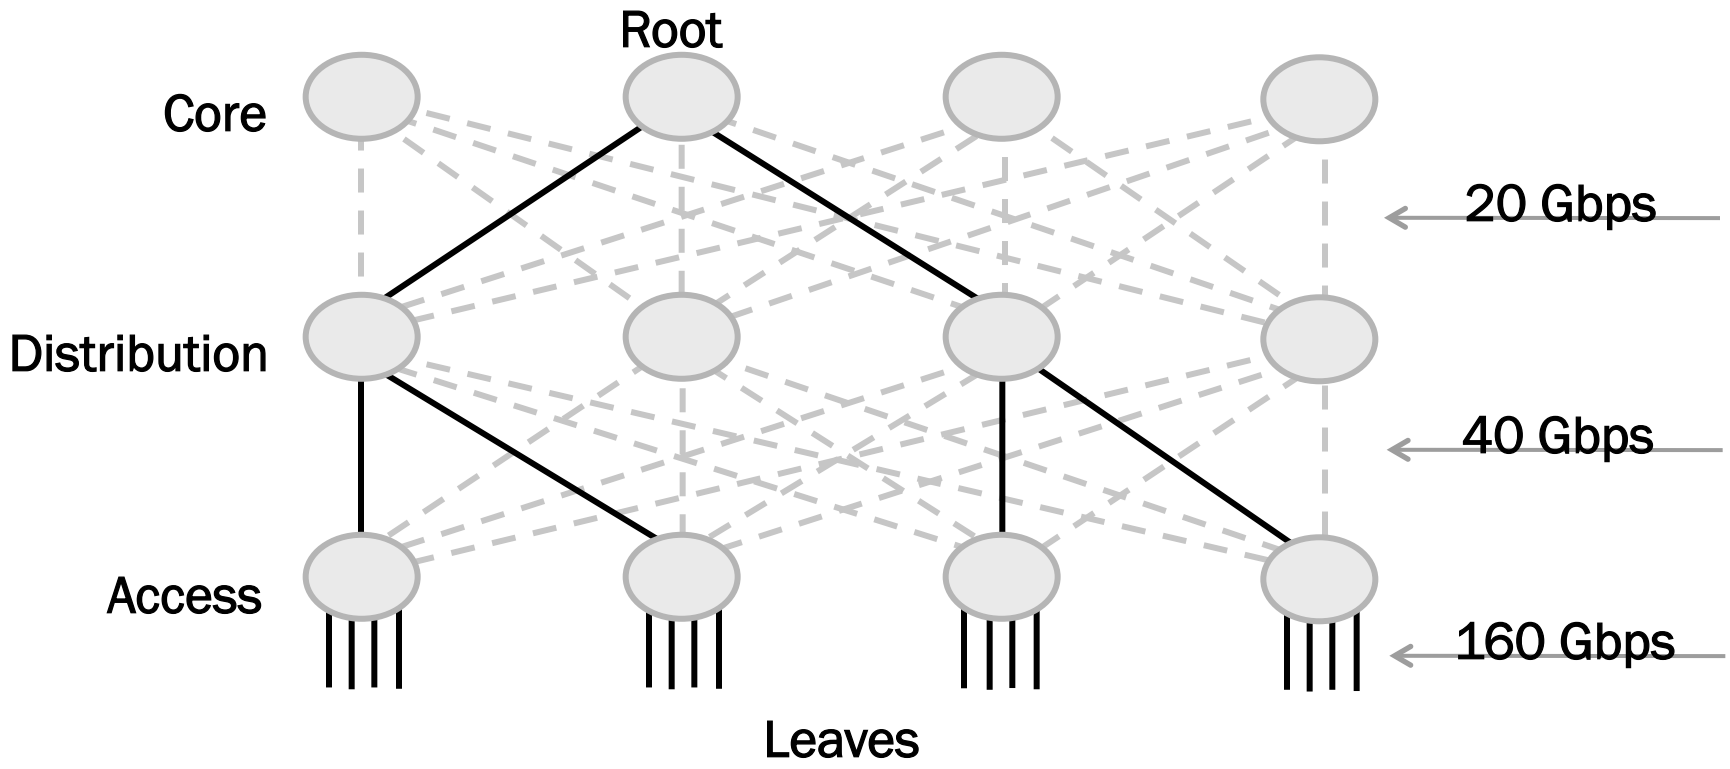

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

The Network with L2MP (TRILL)

Keep all links forwarding

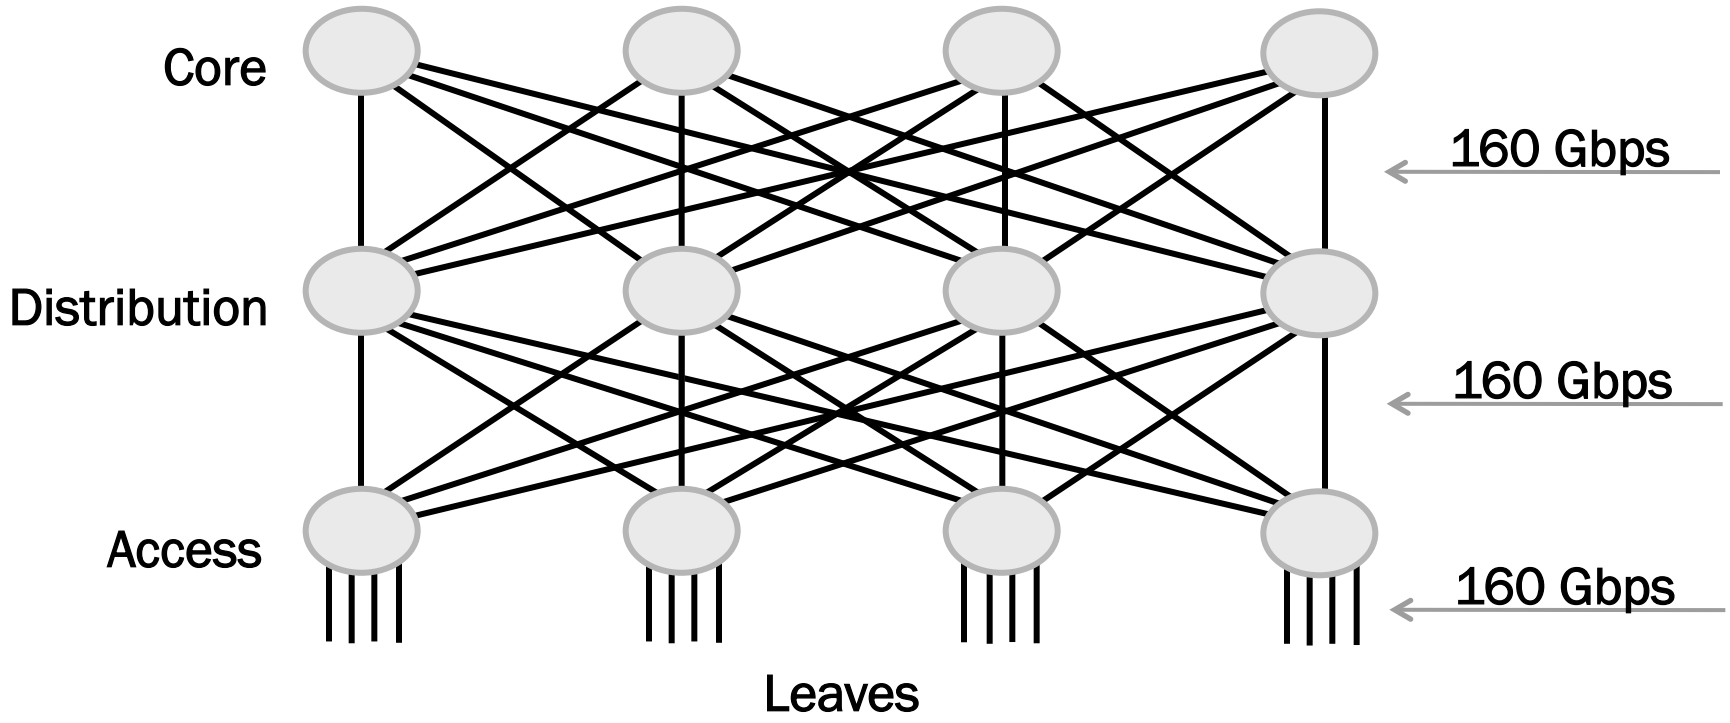

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

SPB—Shortest Path Bridging

Overview

Devices are Ethernet bridges (support 802.1ad stacking, ag OAM, and ah PBB)

Data plane uses MAC-in-MAC

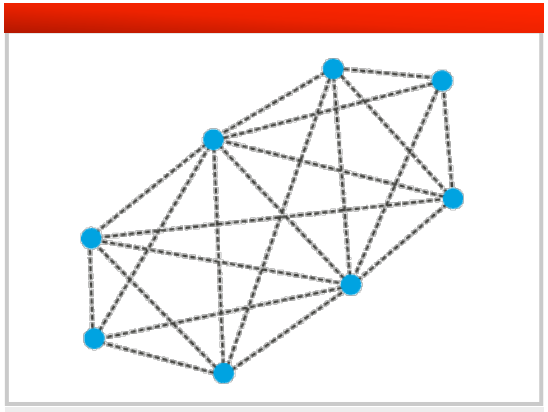

Control plane uses IS-IS Layer 2 link-state routing protocol

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

TRILL and SPB Use of IS-IS

Functions

RBridges and SPB bridges:

- Use link-state Hellos to find each other
- Calculate shortest paths to all other RBridges/ bridges
- Build routing tables

TRILL—Ingress RBridges encapsulate TRILL data; egress RBridges decapsulate TRILL data
SPB—Ingress bridge adds external MAC (destination); egress bridge removes external MAC

Link-state protocols

- Flood configuration information to nodes
- Used for shortest-path calculations
- Distribute configuration database

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Role of Link-State Routing

Discovery and shortest path

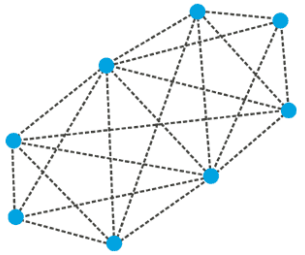

Link-State Routing Protocols Are Used To:

- Discover Ethernet fabric members
- Determine Virtual LAN (VLAN) topology
- Establish Layer 2 delivery using shortest-path calculations
- Nodes tell every node on the network about their closest neighbor
- The nodes distribute only the parts of the routing table containing their neighbors

Link-State Routing Neighbor Information

- Gathered continuously
- The list is flooded to all neighbors
- Neighbors in turn send it to all of their neighbors and so on
- Flooded whenever there is a (routing-significant) change
- Allows nodes to calculate the best path to any other node in the network

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

TRILL vs. SPB

Different approaches to the same problem

Characteristic	TRILL	SPB
Standards Body	IETF (Standard July 2011)	IEEE 802.1aq (soon)
Link-State Protocol	IS-IS (new PDUs)	IS-IS (new PDUs)
Encapsulation	TRILL Header	MAC-in-MAC
Multi-Path Support	Yes	Yes
Loop Mitigation	TTL	RPFC
Packet Flow	Hop by Hop	Symmetric
Configuration Complexity	Easy	Moderate
Troubleshooting	Moderate	Easy (OAM)

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Flat Networking

TRILL and/or SPB allow for large Layer 2-based networks

- Hosts can directly communicate with each other without routers
- Highly interconnected, all paths available, and all links active
- Flat is synonymous with low latency
- Low latency is a fundamental building block for meeting user expectations

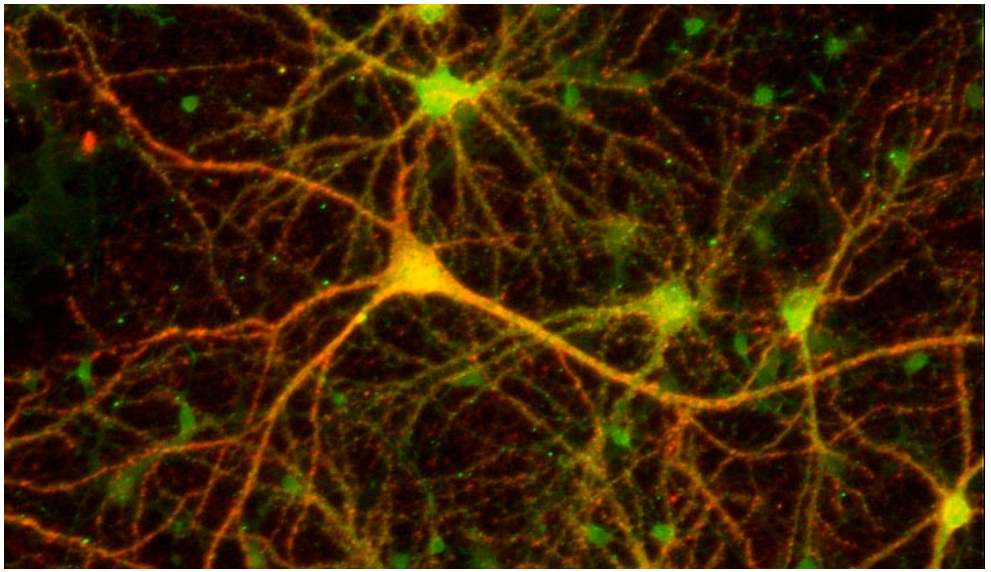

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

TRILL Status

- After a nearly record time in the Editors Queue due to politics and such, TRILL is now an IETF Standard. You may notice the documents listed as "Proposed Standard". After six months it will become a "Draft Standard", then finally a "Standard"
- "RBridges: Base Protocol Specification": RFC 6325
  - Base Protocol, frame and primary function definitions
  - Published as an RFC July 2011
- "Extensions to IS-IS for Layer-2 Systems": RFC 6165
  - IS-IS additions required by both TRILL and 802.1aq
  - Published as an RFC July 2011
- "TRILL Use of IS-IS": RFC 6326
  - IS-IS extensions and changes for TRILL
  - Published as an RFC July 2011
- "RBridges: Adjacency": RFC 6327
  - State machines to clarify the IS-IS extensions
  - Published as an RFC July 2011
- "RBridges: Appointed Forwarders": RFC 6439
  - Appointed Forwarders: Improved documentation of AFs from RFC 6325 4.2.4
  - Provides for a single Rbridge per VLAN to handle among other things Multicast traffic.
  - Published as an RFC November 2011

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

Some Current Work in the TRILL WG

http://tools.ietf.org/pdf/draft-yong-trill-trill-o-mpls-01.pdf

http://tools.ietf.org/pdf/draft-tissa-trill-oam-03.pdf

http://tools.ietf.org/pdf/draft-tissa-trill-cmt-01.pdf

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use with result in Flogging

References and More Information

- IETF TRILL WG:
  - http://datatracker.ietf.org/wg/trill/charter/
- SPB: IEEE 802.1aq – Shortest Path Bridging
  - http://www.ieee802.org/1/pages/802.1aq.html
- RBridges and the IETF TRILL Protocol (NANOG48):
  - http://www.nanog.org/meetings/nanog48/abstracts.php?pt=MTUwNCZuYW5vZzQ4&nm=nanog48
- Shortest Path Bridging – IEEE 802.1aq (NANOG49):
  - http://www.nanog.org/meetings/nanog49/abstracts.php?pt=MTYwNSZuYW5vZzQ5&nm=nanog49
- TRILL and SPB Tutorials, The Great Debate (NANOG50):
  - http://www.nanog.org/meetings/nanog50/agenda.php

© 2012 Brocade Communications Systems, Inc. Proprietary Information. Improper Use Will Result in Flogging