OVERVIEW OF TRILL - Transparent Interconnection of Lots of Links

http://tools.ietf.org/wg/trill/

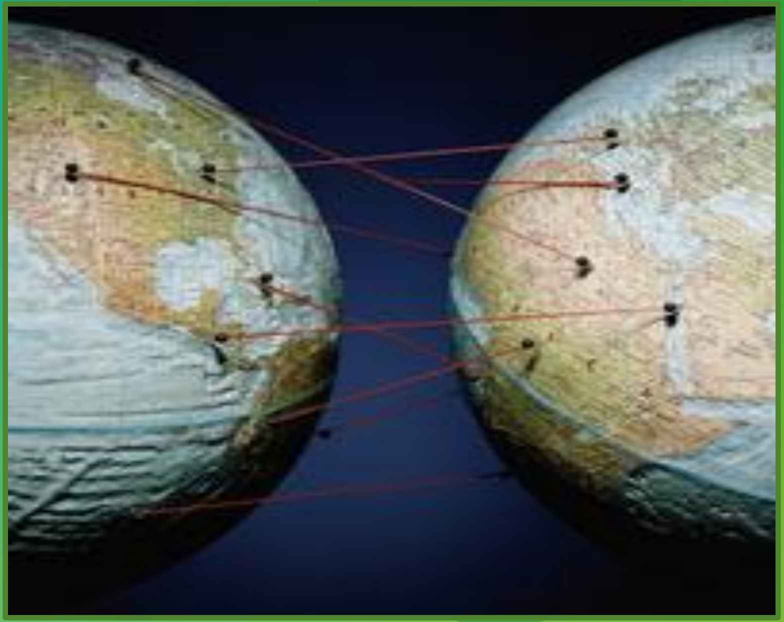

MENOG14, Dubai, 2014

Abdelrahman Ibrahim

Sr. Solution Architect – Cisco

AGENDA

- Current L2 and DC deployment limitations.
- Need for TRILL and Goal
- What is TRILL
- TRILL and RBirdges Key Features
- Use Cases for TRLL
- Unicast and Multicast Handling
- TRILL Control Plane Overview
- TRILL Data Plane Overview
- TRILL Standardization Status

Why Layer 2 in the Data Center?

- Because customers request it!
  - Some protocols rely on the functionality
  - Simple, almost plug and play
  - No addressing
  - Required for implementing subnets
  - Allows easy server provisioning
  - Allows virtual machine mobility

Current Data Center Design

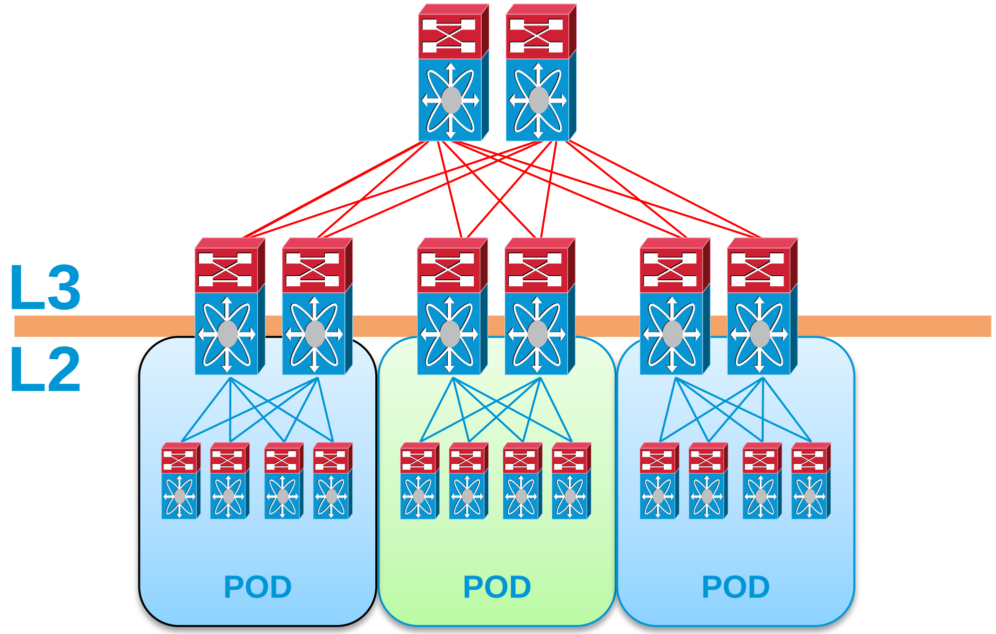

- L2 benefits limited to a POD

Possible Solution for End-to-End L2?

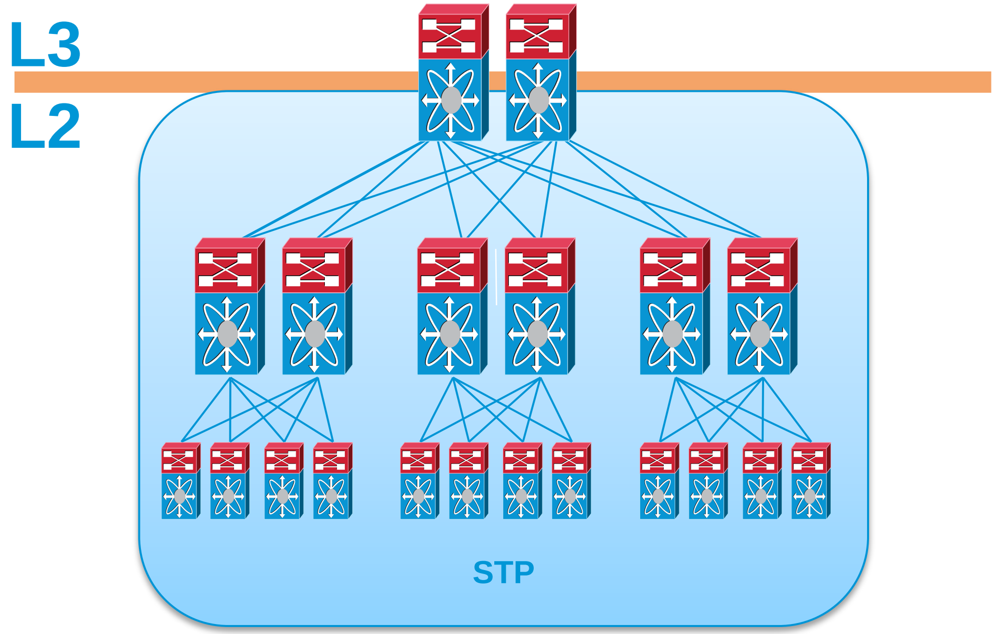

- Just extend STP to the whole network

Typical Limitations of L2

- Local STP problems have network-wide impact, troubleshooting is difficult
- STP provides limited bandwidth (no load balancing)
- STP convergence is disruptive
- Tree topologies introduce sub-optimal paths
- MAC address tables don’t scale
- Flooding impacts the whole network

Spanning Tree and Blocked Links

- Branches of trees never interconnect (no loop!!!)

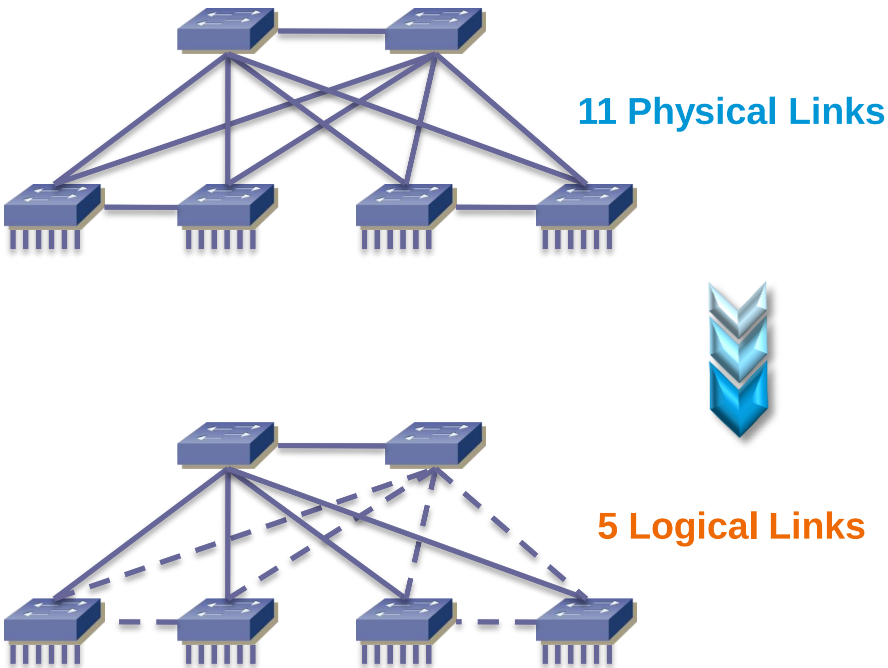

- Spanning Tree Protocol (STP) uses the same approach to build loop-free L2 logical topology
- **Over-subscription ratio exacerbated by STP algorithm**

GOAL

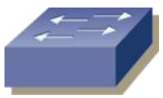

Switching

- Easy Configuration
- Plug & Play
- Provisioning Flexibility

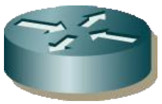

Routing

- Multi-pathing (ECMP)
- Fast Convergence
- Highly Scalable

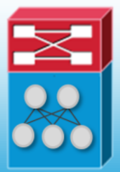

TRILL Fabric

"TRILL brings Layer 3 routing benefits to flexible Layer 2 bridged Ethernet networks"

TRILL Feautues

- Transparency
- Plug & Play
- Virtual LANs
- Frame Priorities
- Data Center Bridging
- Virtualization Support
- Multi-pathing
- Optimal Paths
- Rapid Fail Over
- The safety of a TTL
- Options

What Is a TRILL?

- Externally, a Fabric looks like a single switch
- Internally, a protocol adds Fabric-wide intelligence and ties the elements together.
  This protocol provides in a plug-and-play fashion:
  - Optimal, low latency connectivity any to any
  - High bandwidth, high resiliency
  - Open management and troubleshooting
- TRILL provides additional capabilities in term of scalability and L3 integration

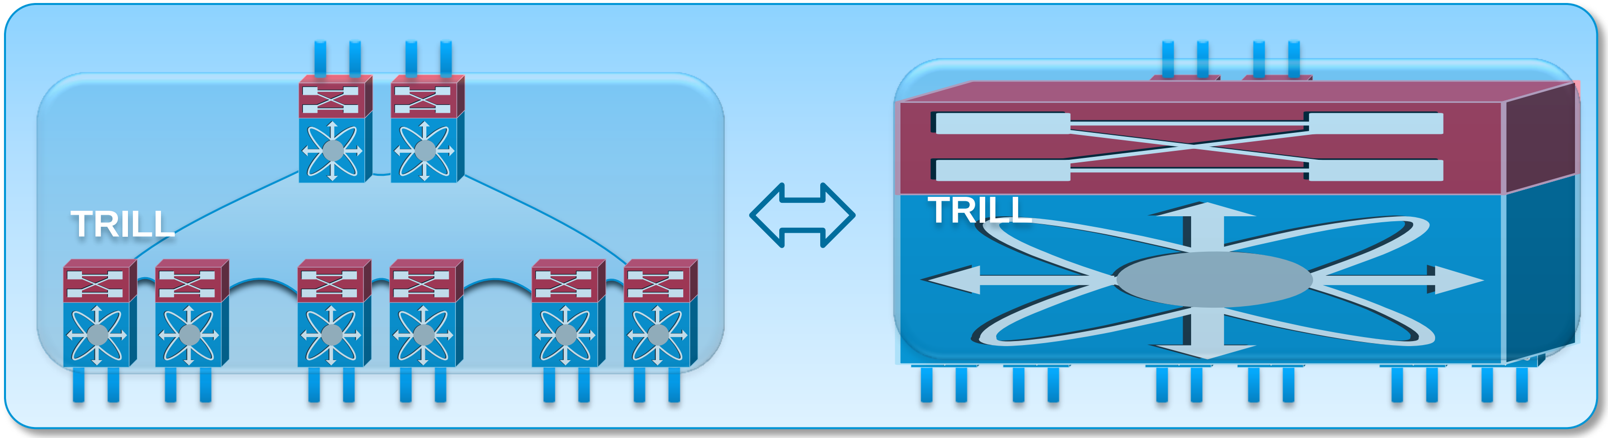

Learning of MAC addresses takes place at the Edge

TRILL IETF Approach to Shortest Path Bridging

- TRILL
  (TRansparent Interconnect of Lots of Links)

  http://www.ietf.org/html.charters/trill-charter.html

- Main areas addressed by TRILL:

  Provide Shortest Path and
  Equal Cost Multi-Pathing for traffic

  Be Plug-n-Play

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.

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.

What/Why/Who TRILL?

- Basically a simple idea:
  - Encapsulate native frames in a transport header providing a hop count.
  - Route the encapsulated frames using IS-IS.
  - Decapsulate native frames before delivery.

Why IS-IS For TRILL?

- The IS-IS (Intermediate System to Intermediate System) link state routing protocol was chosen for TRILL over OSPF (Open Shortest Path First), the only other plausible candidate, for the following reasons:

  - IS-IS runs directly at Layer 2. Thus no IP addresses are needed, as they are for OSPF, and IS-IS can run with zero configuration.
  - IS-IS uses a TLV (type, length, value) encoding which makes it easy to define and carry new types of data.

Routing versus Bridging

- Routing only sends data out a port when it receives control messages on that port indicating this is safe and routing has a TTL for safety.
  - If control messages are not received or not processed, it "fails safe" and does not forward data.
- Bridging (Spanning Tree Protocol) forwards data out all ports (except the one where the data was received) unless it receives control messages on that indicate this is unsafe. There is no TTL.
  - If control messages are not received or not processed, it "fails unsafe", forwards data, and can melt down due to data loops.

Optimum Point-to-Point Forwarding

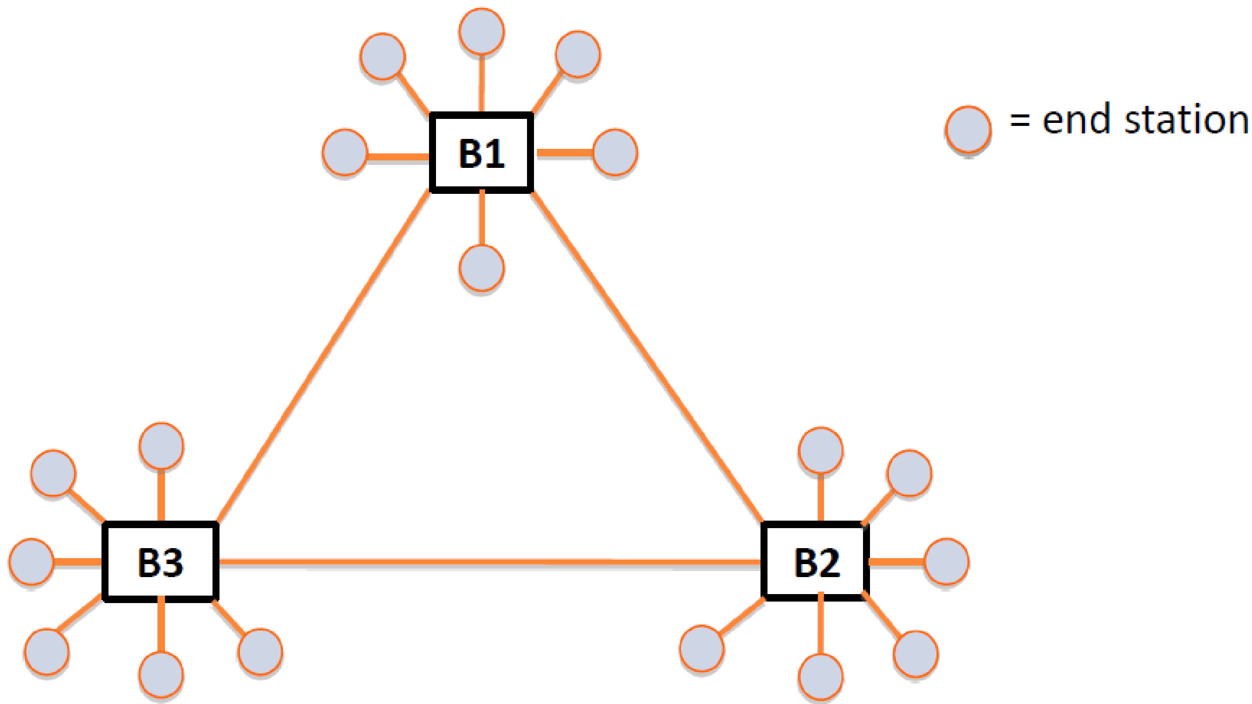

A three bridge network

Optimum Point-to-Point Forwarding

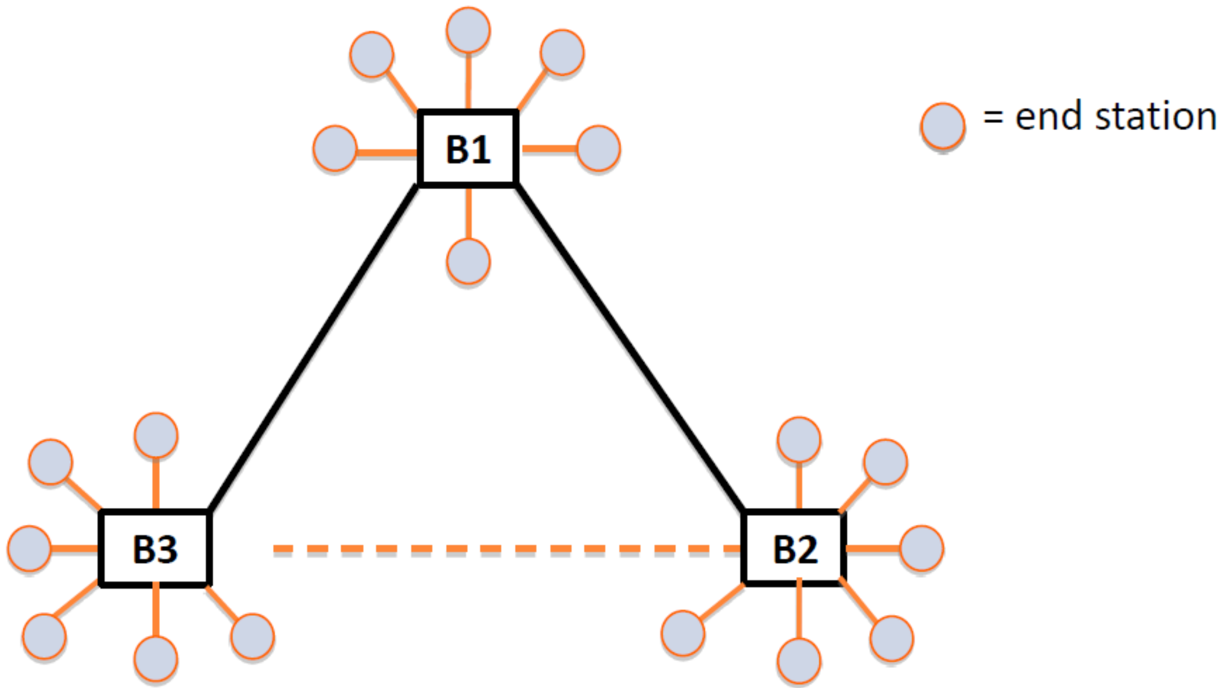

Spanning tree eliminates loops by disabling ports

Optimum Point-to-Point Forwarding

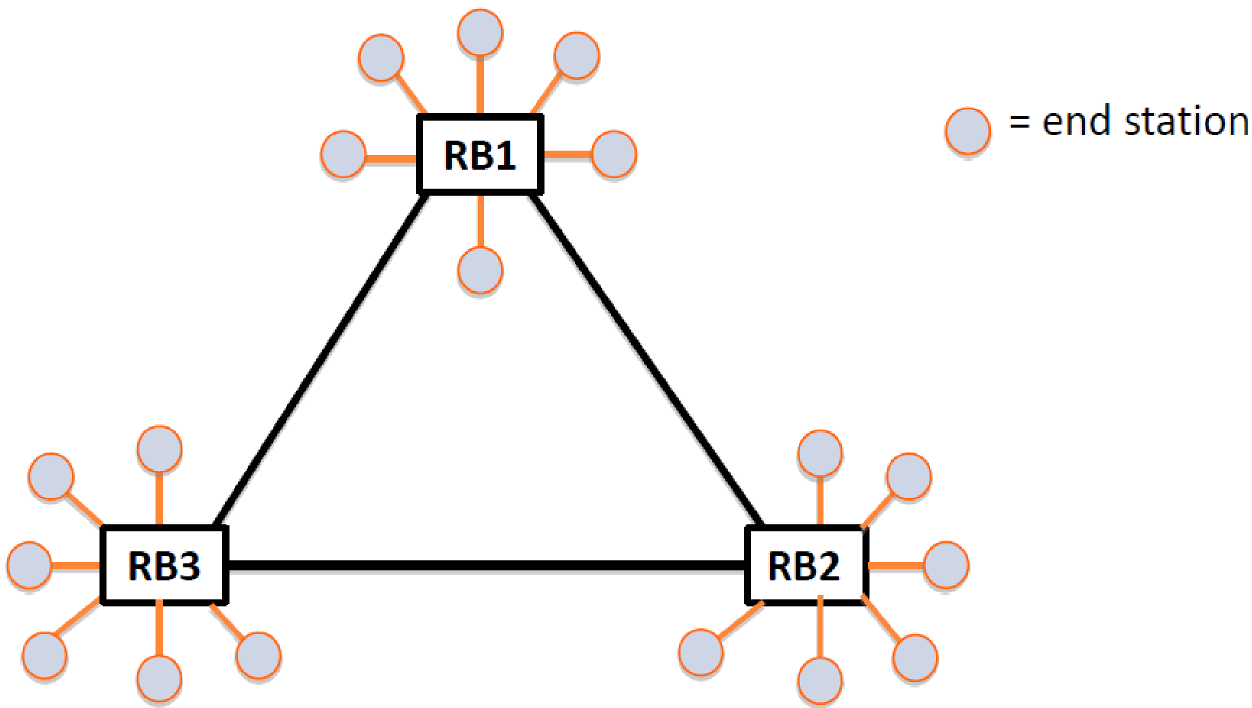

A three RBridge network: better performance using all facilities

Multi-Pathing (Unicast)

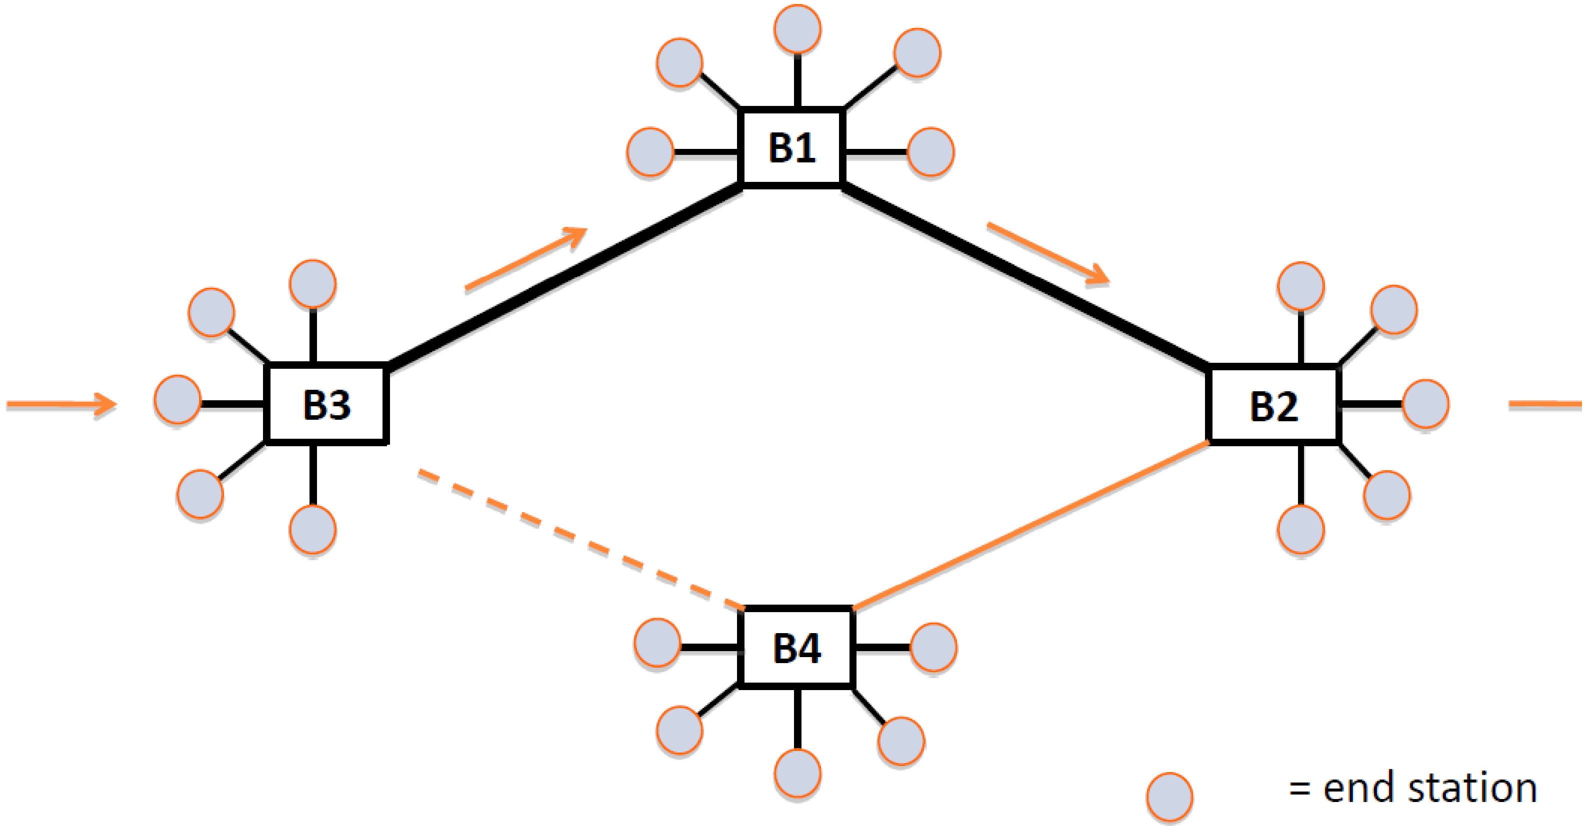

Bridges limit traffic to one path

Multi-Pathing (Unicast)

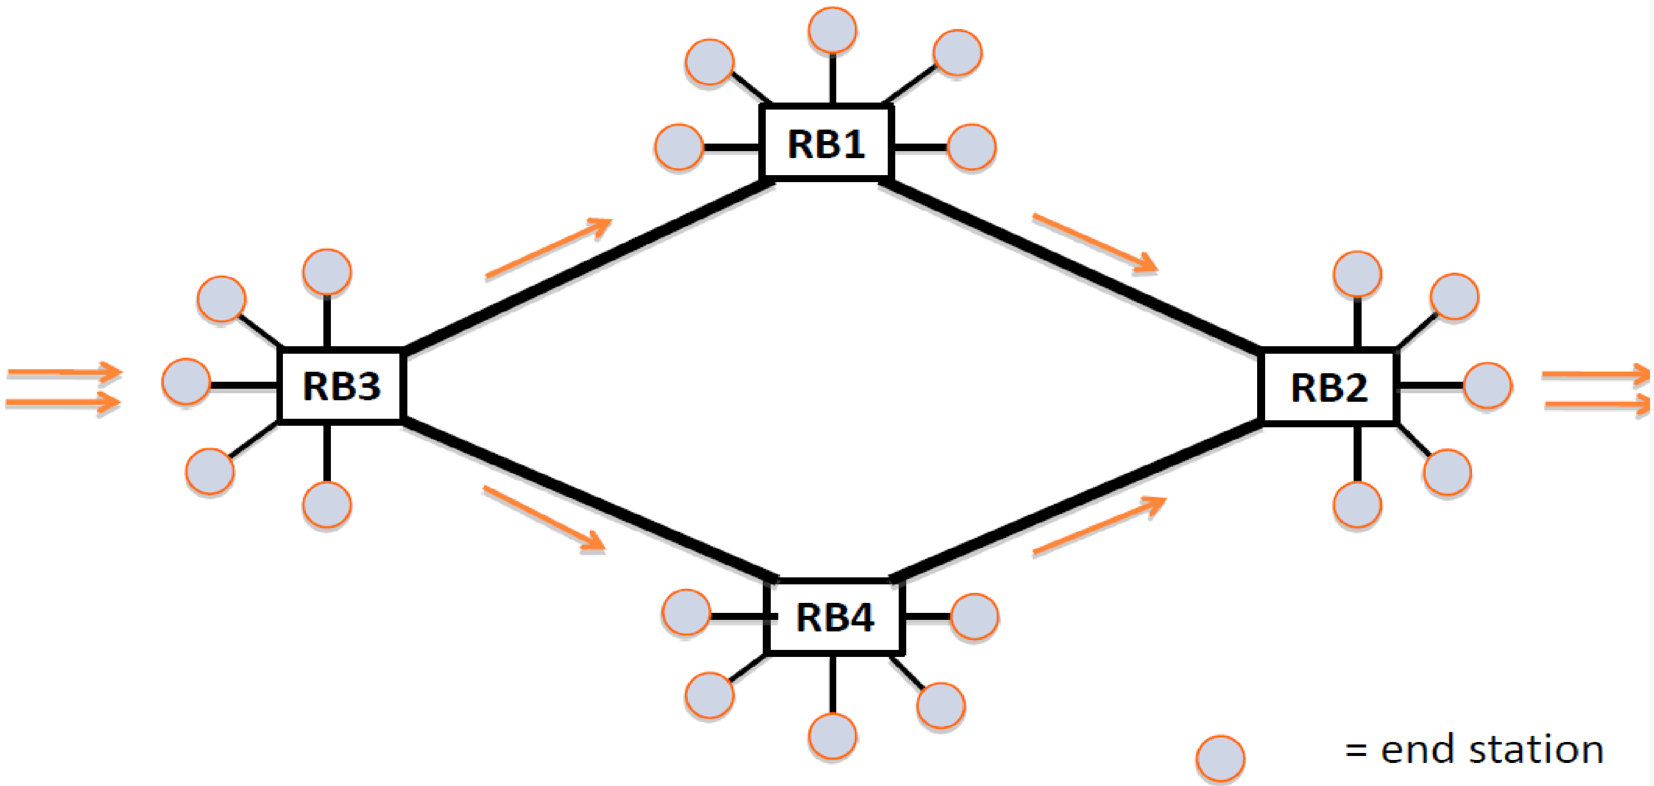

RBridges support
multi-path for higher throughput

Multi-Pathing (Unicast)

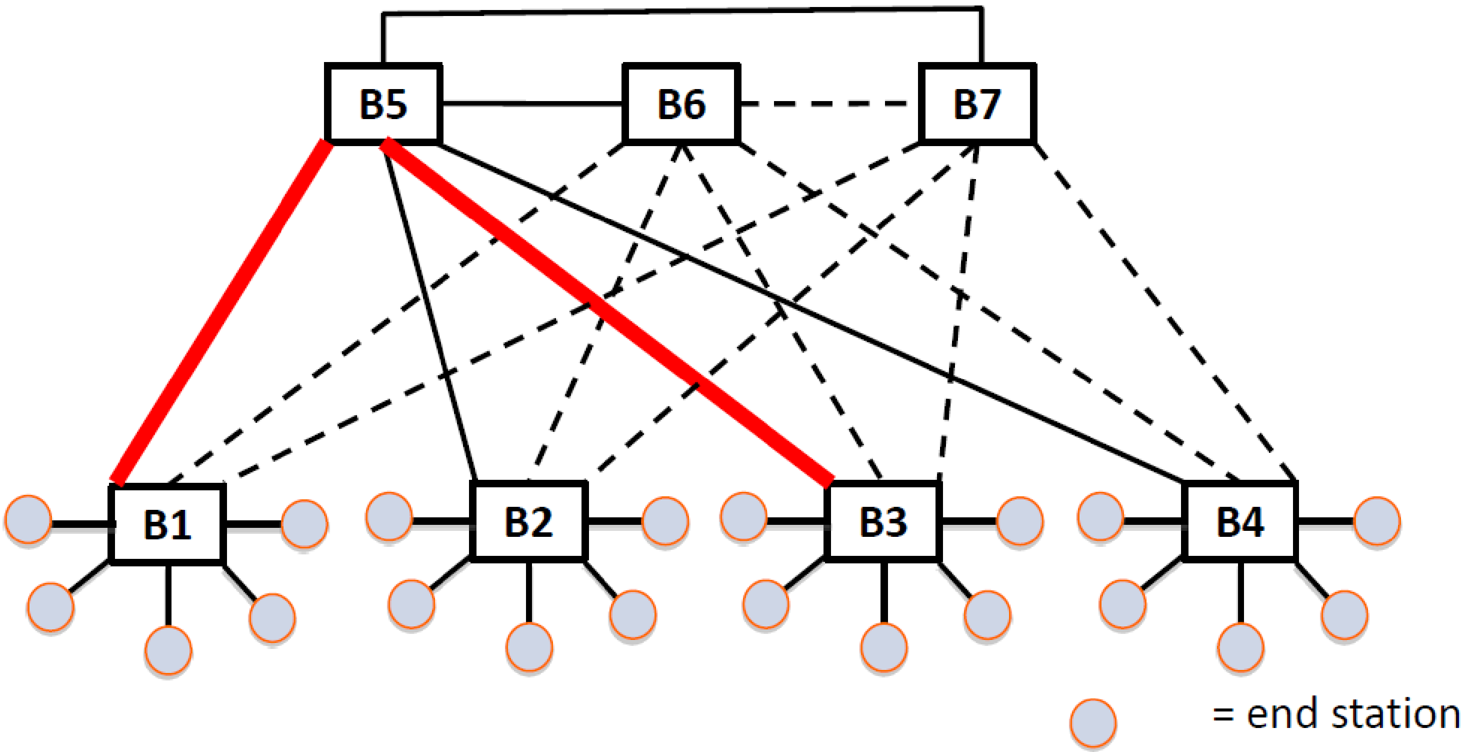

Bridges limit traffic to one path

Multi-Pathing (Unicast)

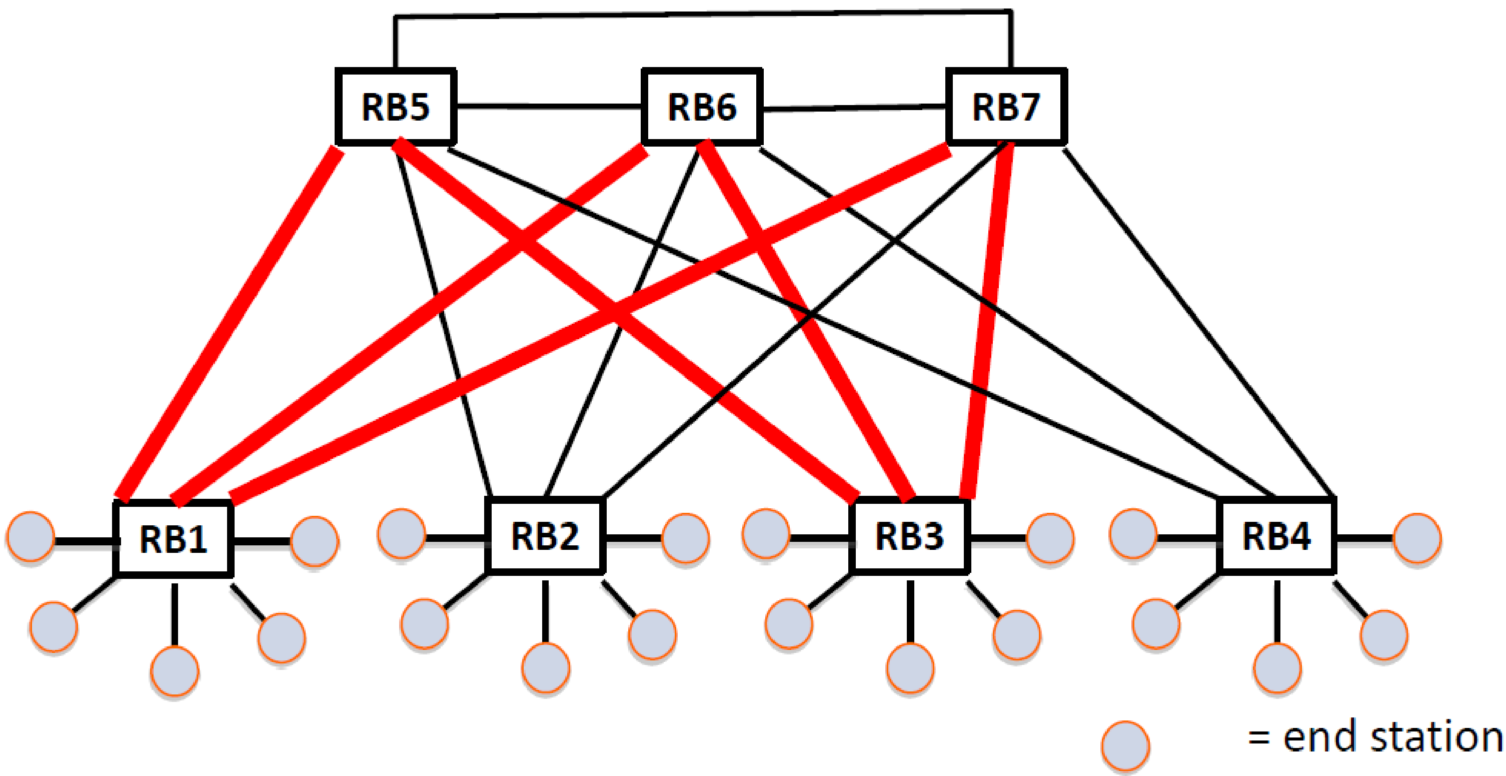

RBridges support
multi-pathing for higher throughput

Multi-Pathing (Multi-destination)

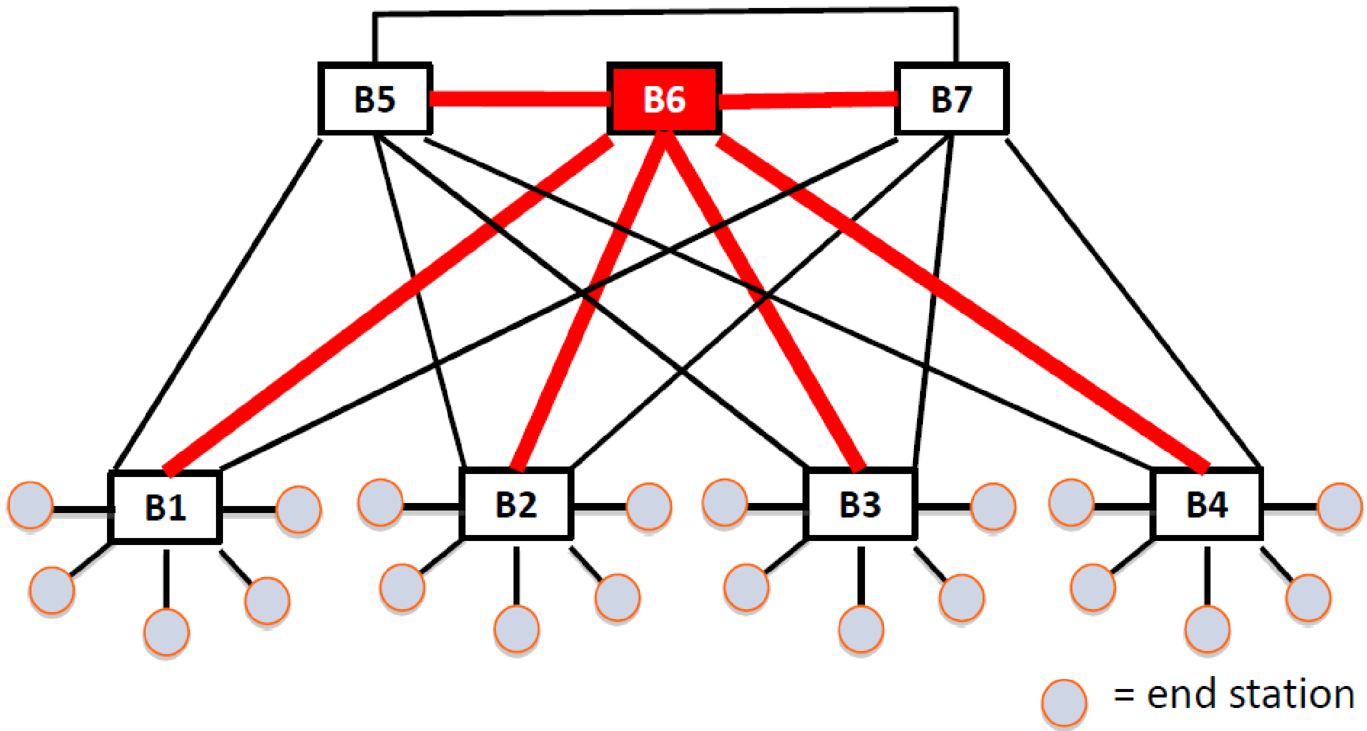

Spanning tree yields a single bi-directional tree for flooding multi-destination frames limiting bandwidth

Multi-Pathing (Multi-destination)

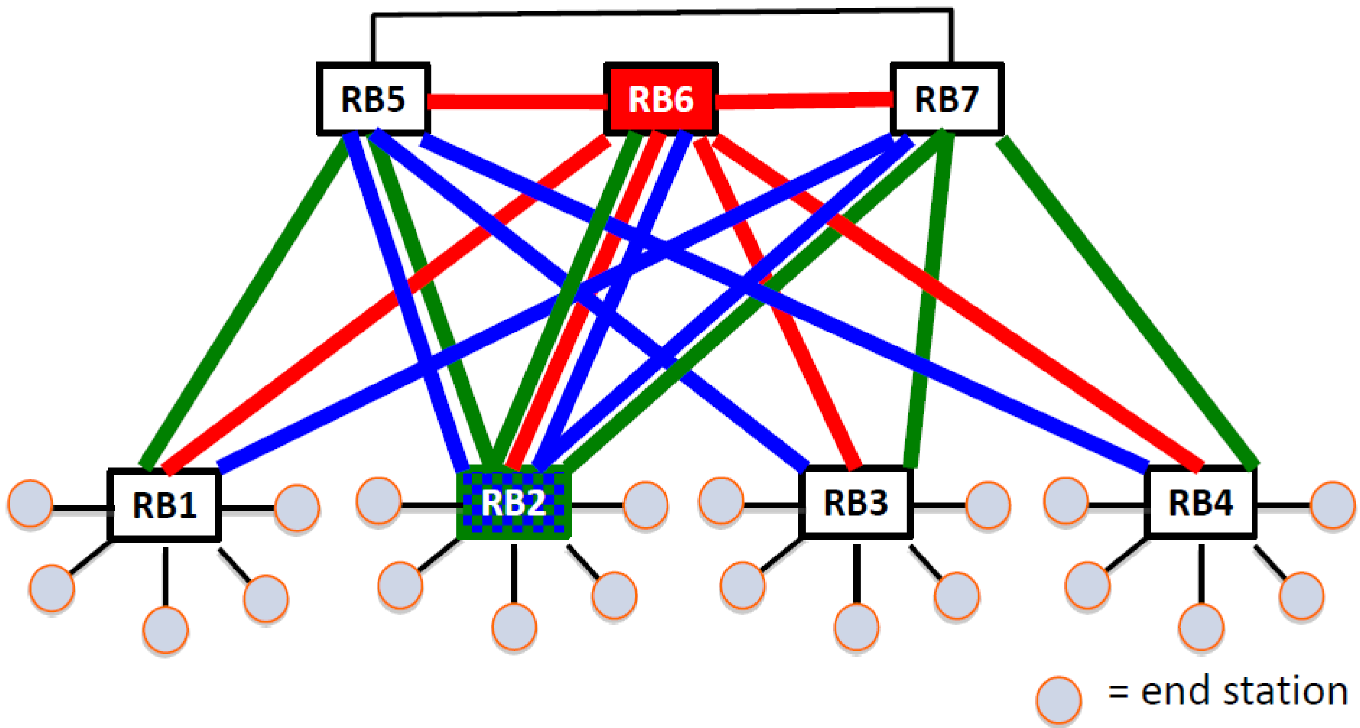

RBridges support multiple distribution trees. The encapsulating Rbridge chooses which to use. RB2 can split multi-destination traffic over three trees.

TRILL Basics

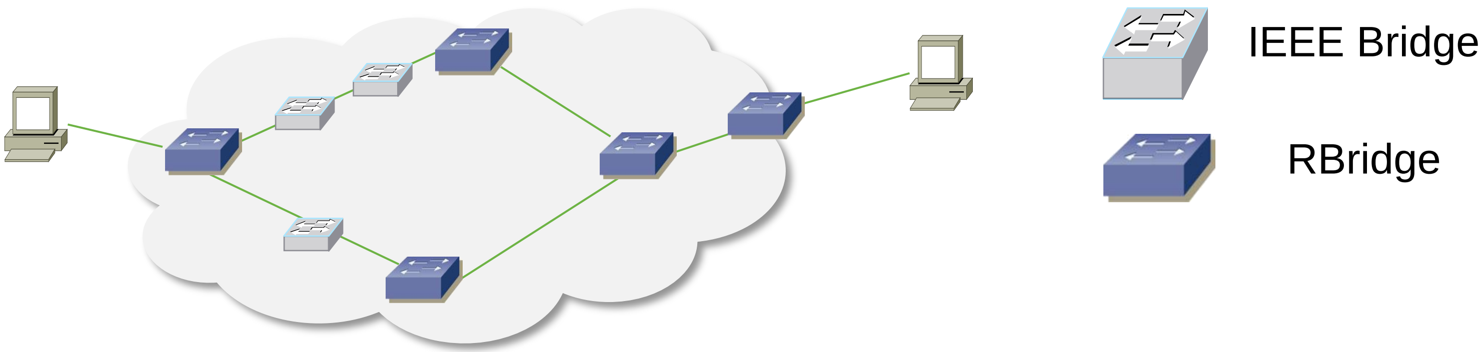

- **A TRILL Network is comprised of Routing Bridges (RBridges/RBs).**
- Each RBridge is uniquely identified by a 'nickname' or rbridge-id (auto-created from ISIS system id)

  RBs can be connected by 802.1 LANs or

  RBs can be connected by simple P2P links (incl. PPP – see RFC 6361)

- **Architecturally, RBridges run "on top" of an 802.1 bridged network similarly to Routers**

  RBridges may be interconnected by classical 802.1Q bridges:
  Allows for gradual migration of existing networks

  RBridges do not participate in xSTP, and drop BPDUs if they are received

26

TRILL

Principles of Operation

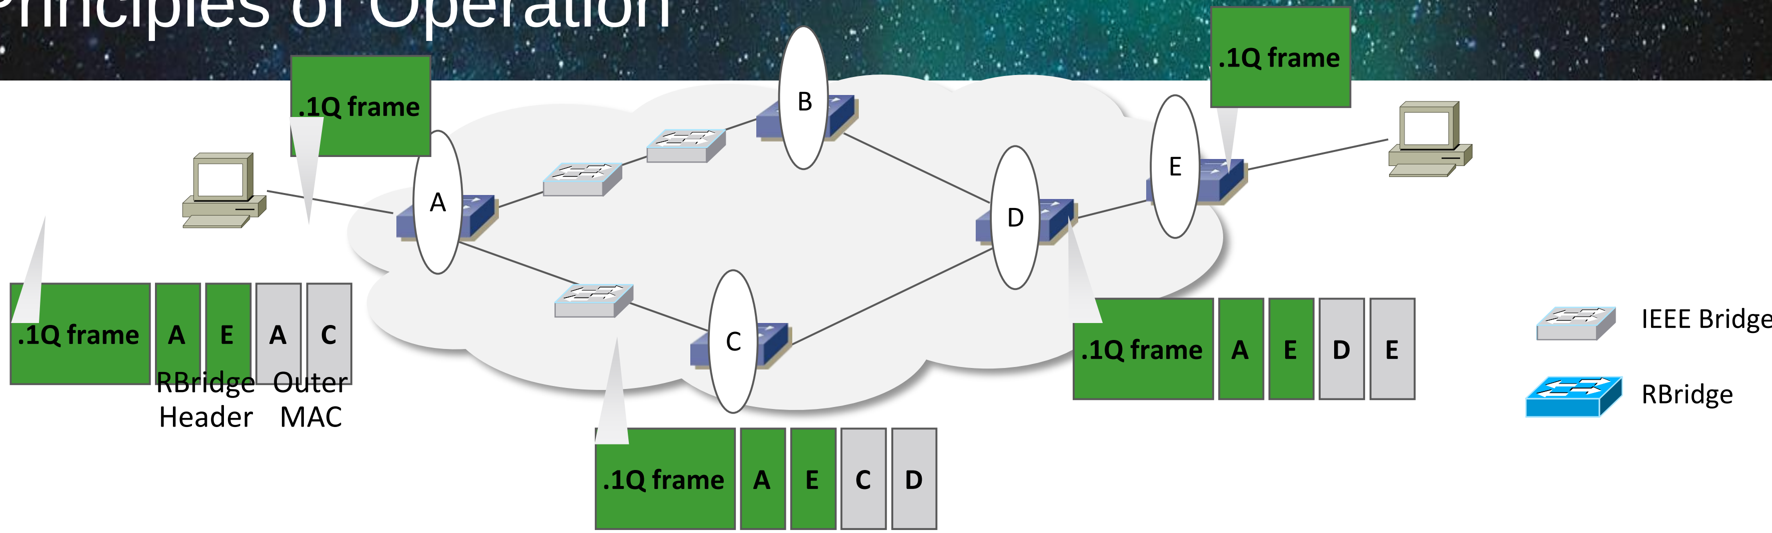

- Frames are encapsulated with the RBridge addresses and further encapsulated with originating rbridge and next hop rbridge MAC address
  - Header fields differ from 802.1ah
  - Headers are swapped hop by hop (similar to routing)
- RBridges learn what MAC addresses are on their edge ports using general dataplane learning and MAY advertise them other RBridges
  - Remote mac-address-to-rbridge binding learning: hardware or control plane
- Unknown unicast /multicast/broadcast frames flooded along pre-calculated distribution tree(s)

TRILL Forwarding

- RBridges use ISIS for discovery and to synchronize Link State Databases
- TRILL uses these Link State Database to
  - Compute pair wise bidirectional paths for unicast (per node and/or per VLAN) between all Rbridges
  - For multicast, distribution trees are calculated rooted at (potentially) every rbridge ;
    trees are given an rbridge-id/nickname as well
- TRILL adds to standard IS-IS
  - Ships in the night with other protocols using ISIS
  - TRILL Hellos
    - Find out whether nodes are on a LAN or P2P link
    - Designated Rbridge (DRB) Election
    - Root-Bridge-IDs
  - See also: RFC 6165 (Extensions to IS-IS for Layer-2 Systems)

TRILL Forwarding (Cont.)

- Edge RBridges learn End Station MAC addresses in the data plane and associate them with RBridge nicknames
  - End Station Address Distribution Information (ESADI)
  - Optional ESADI capability allows RBridges to distribute End Station MAC addresses in IS-IS.
- TRILL Addressing and Forwarding
  - Outer header resembles 802.1ah—MAC-SA and MAC-DA interpreted differently: used to send frame to the next Rbridge, much like a router
  - Provider shim header with TTL field, Ingress and Egress Rbridge Nicknames, and Multicast Flag
  - Rbridges will use the RB-Nicknames (16-bit RBridge-IDs) as well as a Multicast Flag (M-bit) to forward frames
  - If M = 0 : egress rbridge-id is used for unicast forwarding
  - If M = 1 : egress rbridge-id denotes the pre-calculated distribution tree to forward the multicast frame

TRILL—Ethernet Data Encapsulation

Outer Ethernet Header (link specific):

Outer Destination MAC Address (RB2)	
Outer Destination MAC Address	Outer Source MAC Address
Outer Source MAC Address (RB1)	
Ethertype = IEEE 802.1Q	Outer.VLAN Tag Information

- **Outer-VLAN Tag Information:** This is used only if two RBridges communicate across a standard 802.1Q network

TRILL Header:

Ethertype = TRILL	V	R	M	Op-Length	Hop Count
Egress (RB2) Nickname	Ingress (RB1) Nickname				

- **V:** Version
- **M:** Multi-destination; indicates if the frame is to be delivered to a single or multiple end stations
- **Opt-Length:** >0 if an Option field is present
- **Hop Limit:** Similar to TTL
- **RBridge Nickname:** Not the MAC address of the Rbridge, but the a TRILL ID for the RBridge (Egress Nickname used differently if M = 1)

Inner Ethernet Header:

Inner Destination MAC Address	
Inner Destination MAC Address	Inner Source MAC Address
Inner Source MAC Address	
Ethertype = IEEE 0x8100	Inner.VLAN Tag Information
Ethertype = IEEE 0x893B	Inner.VLAN second part

- **Multicast tree pruning:** Requires inspection of *customer Destination MAC Address* and *customer VLAN*
- **In case of Fine Grain Labeling: Second VLAN tag** (see draft-ietf-trill-fine-labeling)

See also: RFC 6325 and RFC 6327

Packet Flow — Known Unicast

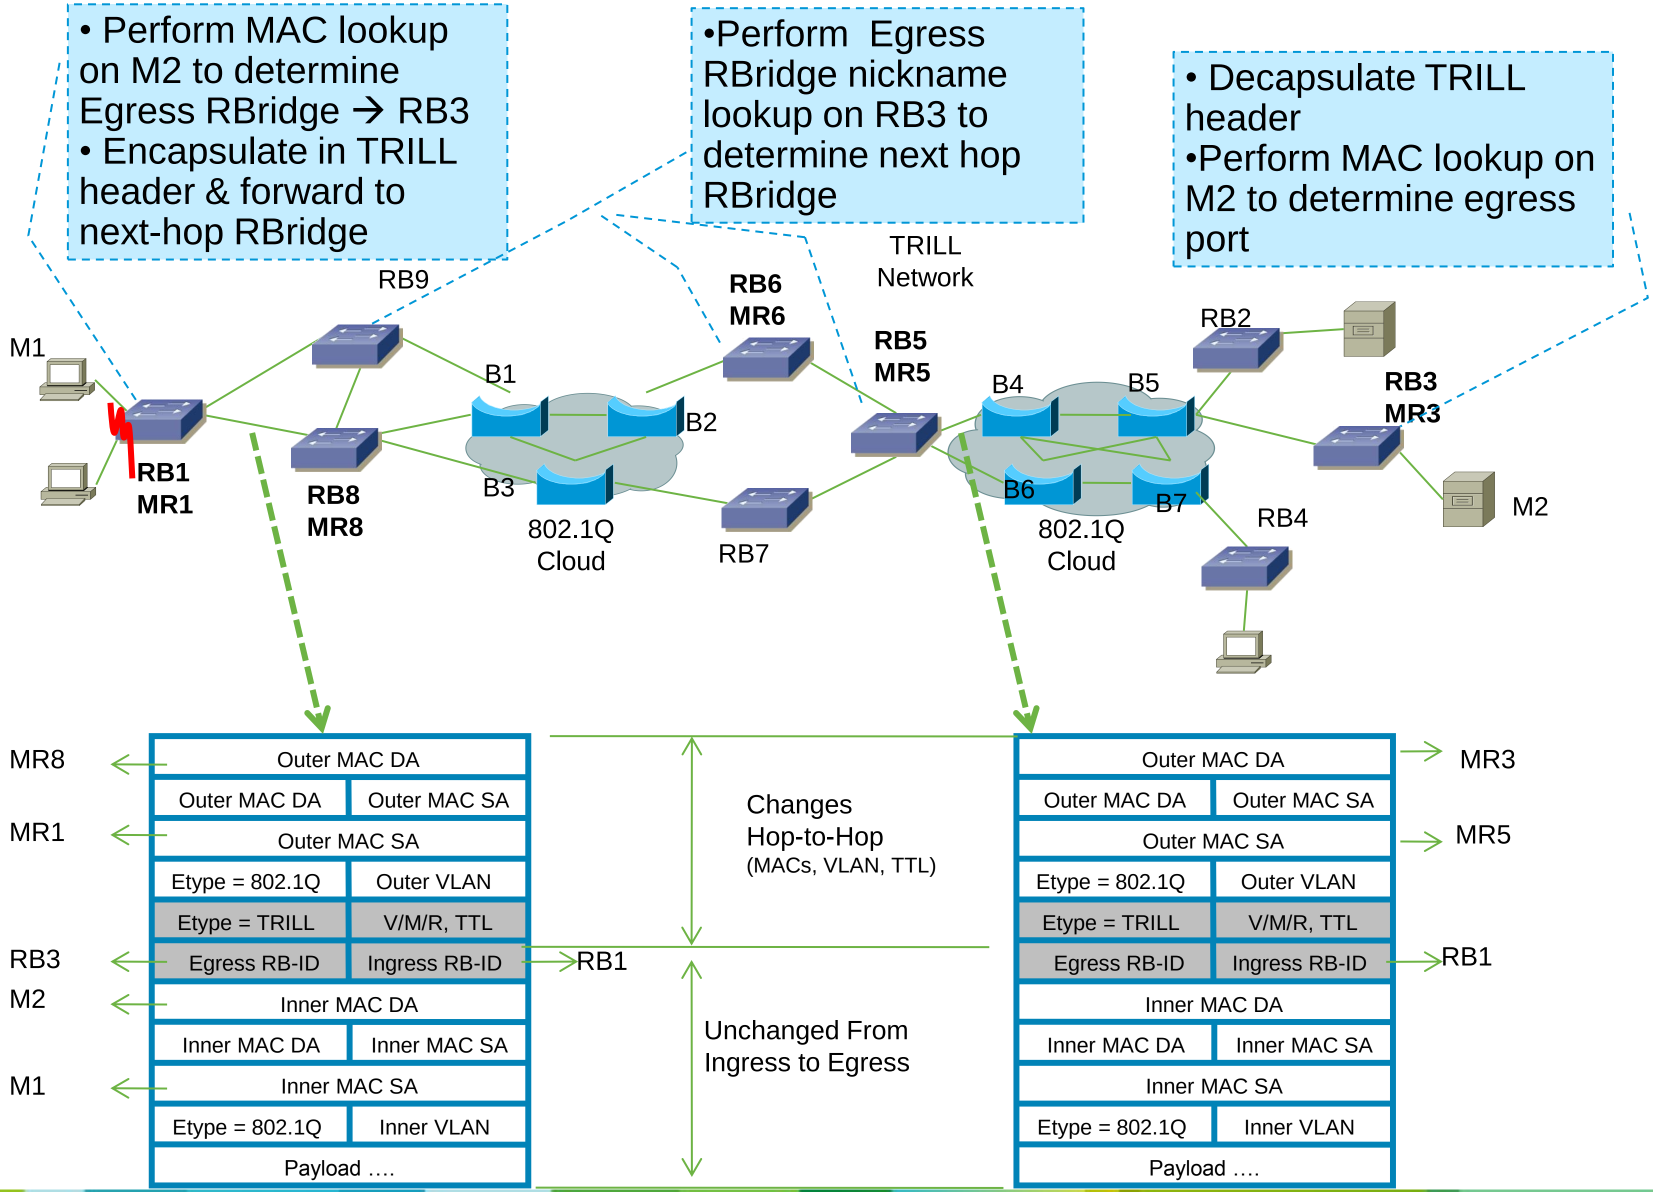

Packet Flow — Multicast/Broadcast/Unknown Unicast

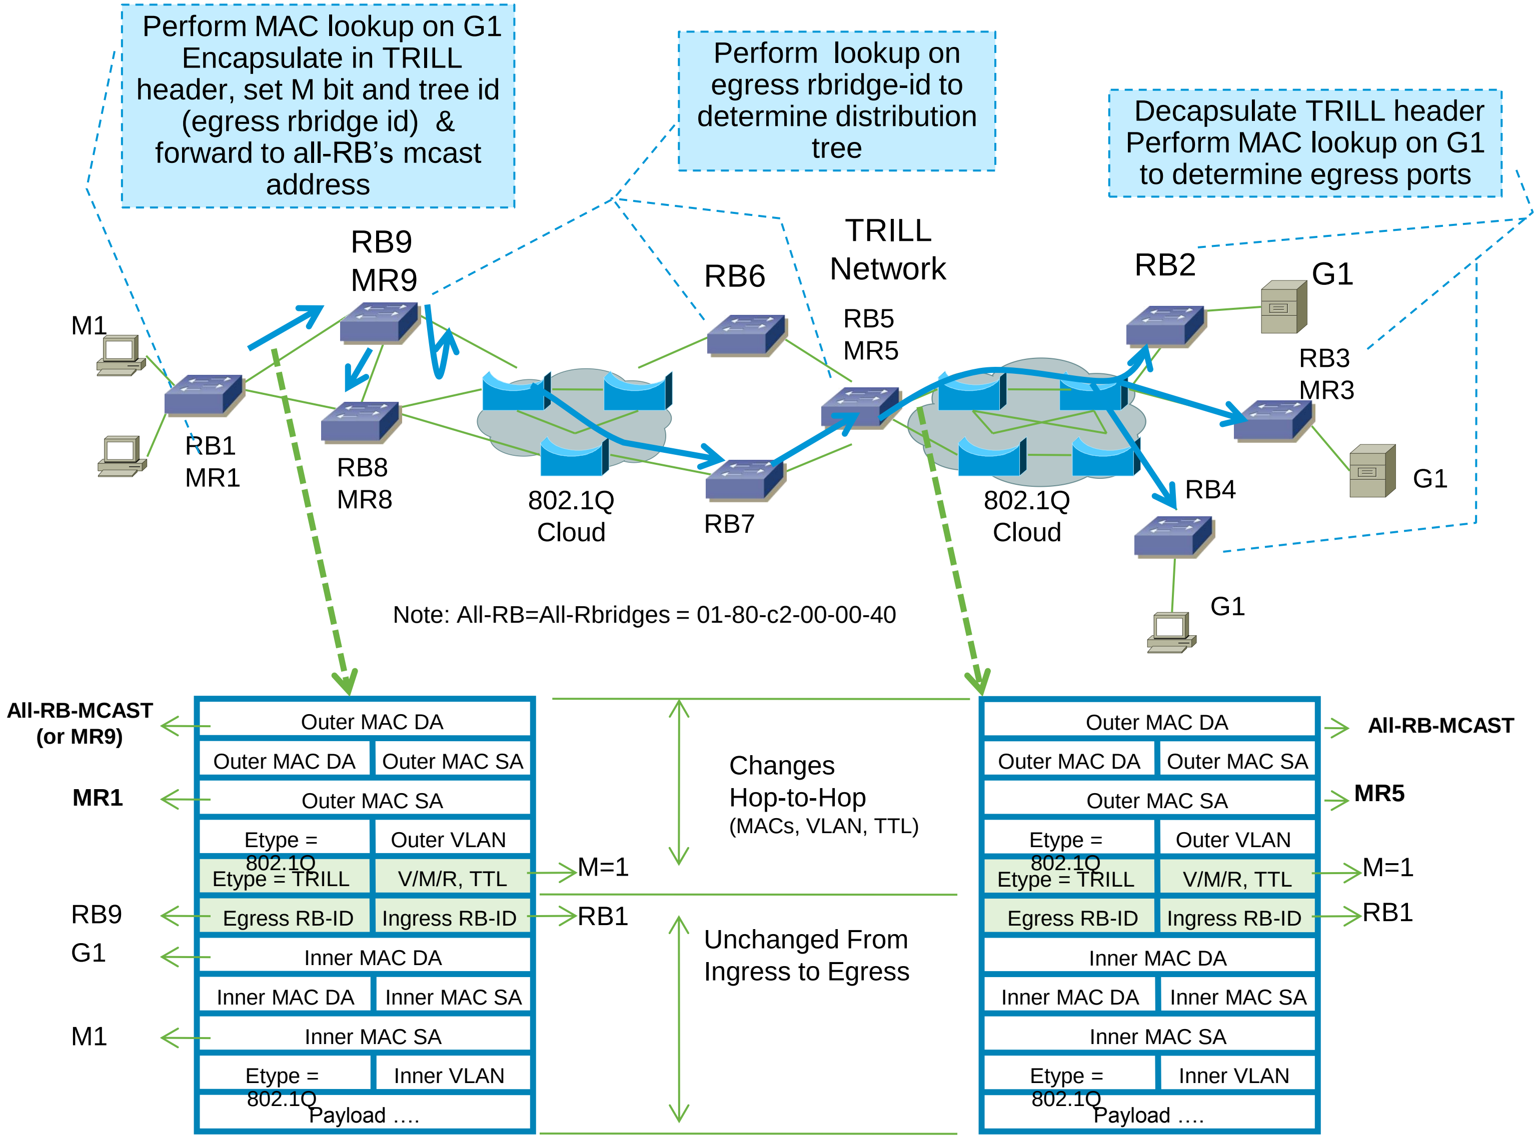

TRILL Benefits

- Shortest path delivery of unicast
- Layer 2 multi-pathing (ECMP) of unicast
- Optimal multicast delivery over shared trees
  - Load-balancing over multiple trees.
  - Per-VLAN/c-group pruning of trees via IGMP/PIM snooping.
- Fast convergence times, Minimal configuration
- Support for Shared Media and P2P links
- Loop Prevention and Mitigation (adds a TTL)
- Support for multi-homing (DRB election)
- Confines MAC Address learning to edge nodes, providing MAC address scalability similar to IEEE 802.1ah (MAC-in-MAC)

Optimal, Low Latency Switching

Shortest path any-to-any

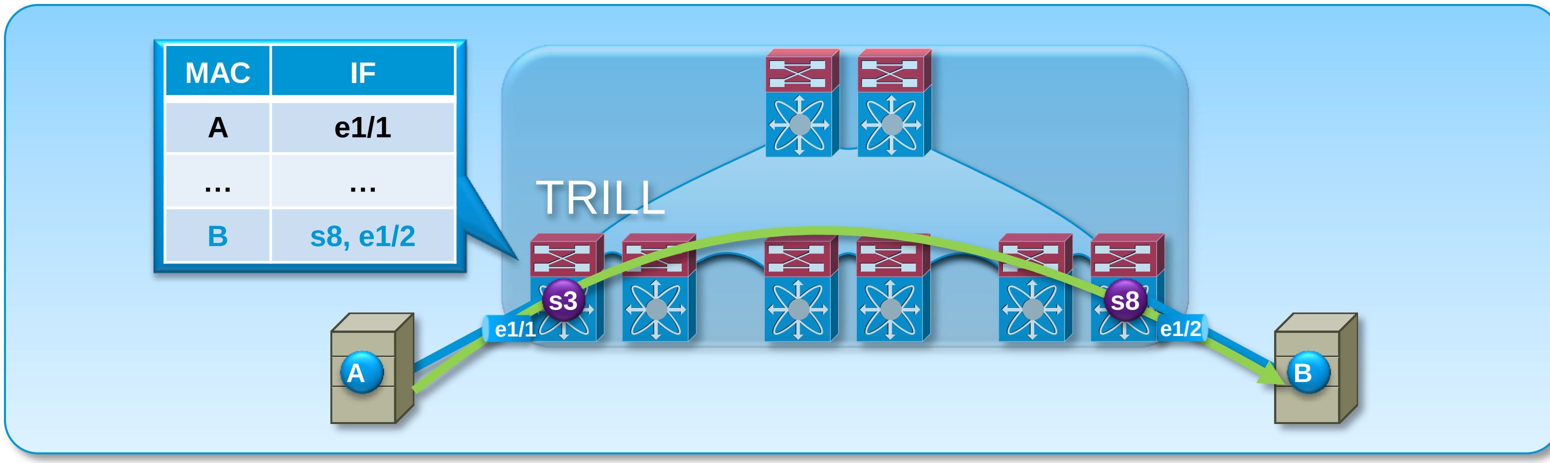

- Single address lookup at the ingress edge identifies the exit port across the fabric
- Traffic is then switched using the shortest path available
- Reliable L2 and L3 connectivity any to any
(L2 as if it was within the same switch, **no STP inside**)

High Bandwidth, High Resiliency

Equal Cost Multi-Pathing

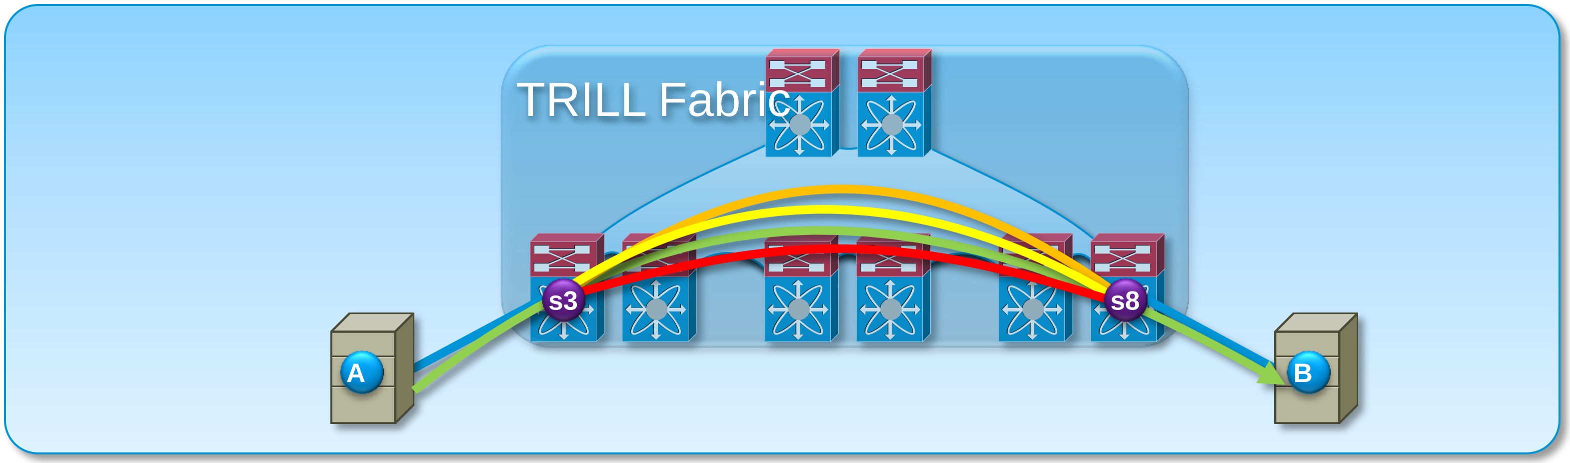

- Multi-pathing (up to 256 links active between any 2 devices)
- Traffic is redistributed across remaining links in case of failure, providing fast convergence

TRILL is Scalable

Conversational Learning

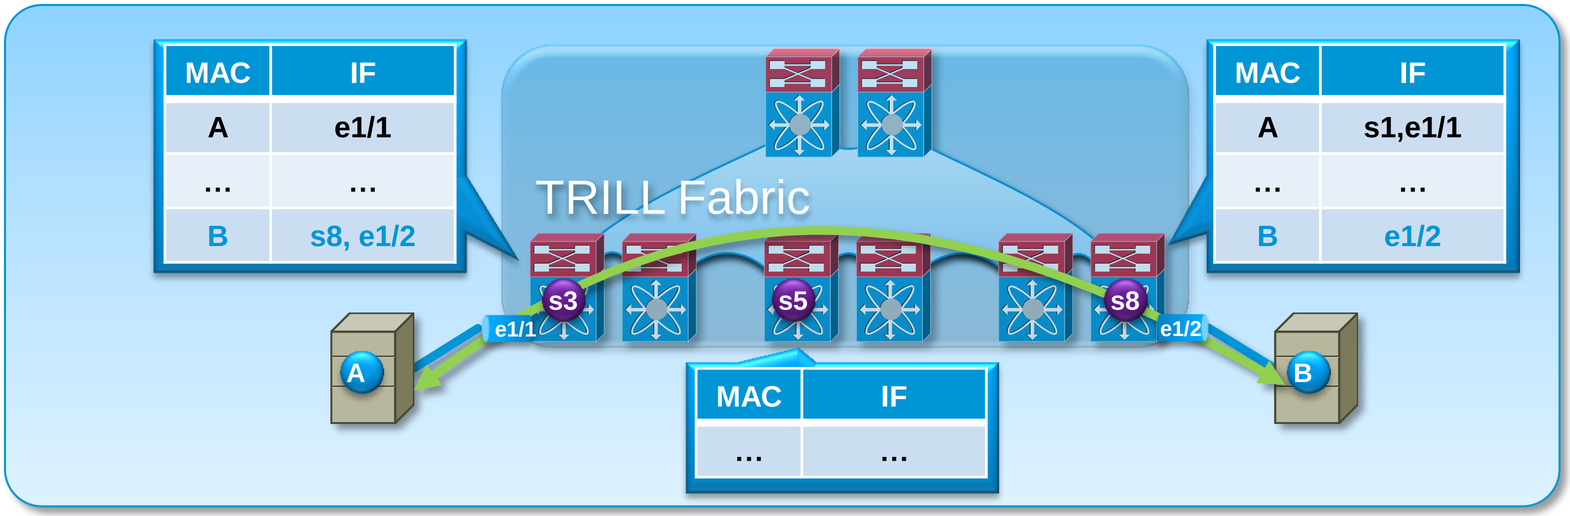

- Per-port MAC address table only needs to learn the peers that are reached across the fabric

A virtually unlimited number of hosts can be attached to the fabric

Layer 2 integration

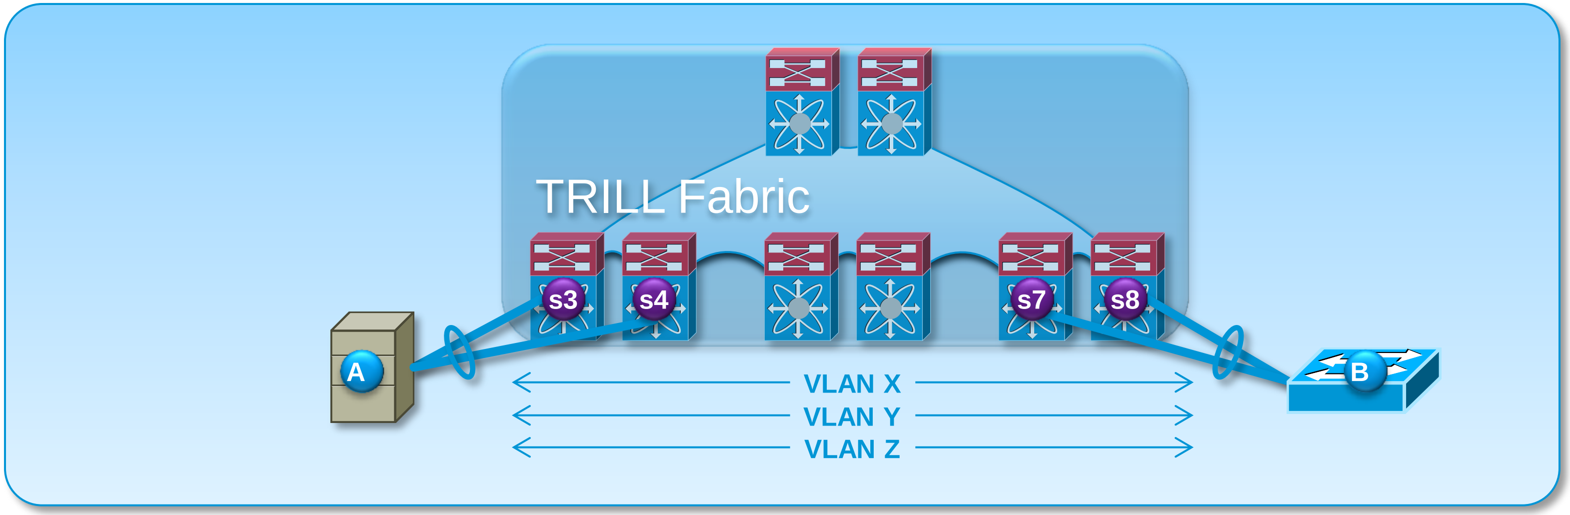

- Allows extending VLANs with no limitation (no risks of loop)
- Devices can be attached active/active to the fabric using IEEE standard port channels and without resorting to STP

Loop Mitigation with TRILL

Time To Live (TTL) and Reverse Path Forwarding (RPF) Check

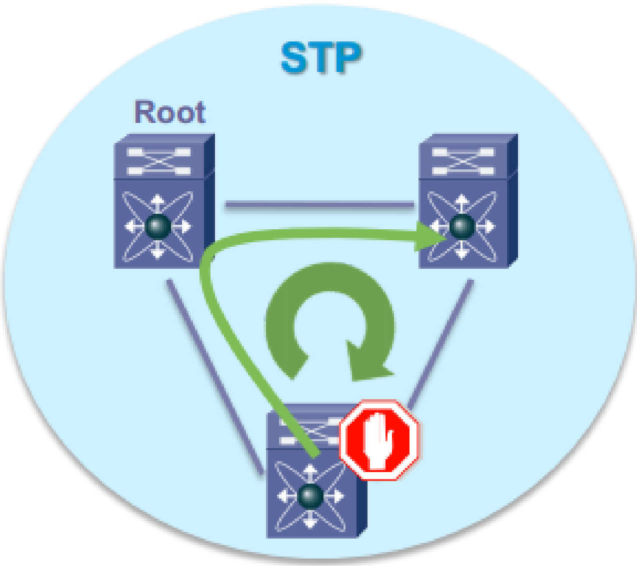

- The control protocol is the only mechanism preventing loops
- If STP fails →loop
  - No backup mechanism in the data plane
  - Flooding impacts the whole network

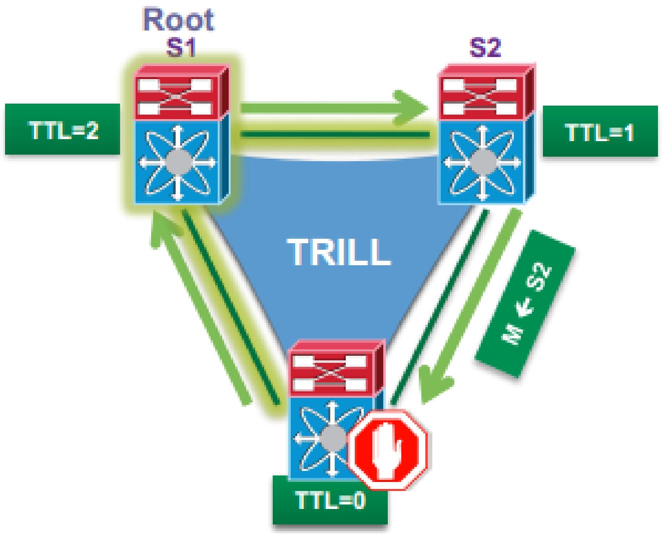

- TTL in TRILL header
- RPF Check for multi-destination traffic
- The data plane is protecting against loops

Key TRILL Features

- Optimum point-to-point forwarding
- Multi-Pathing of both unicast and multi-destination traffic
- The safety of a TTL
  - TTL is implemented in the data plane and provides safety even if the control plane has crashed
- Support for VLANs and frame priorities
- Distribution of multi-cast optimized based on VLAN and multi-cast group
- Full jumbo frame support including jumbo routing messages
- Efficient routing computations

Some Other TRILL Features

- Compatible with classic bridges. RBridges can be incrementally deployed into a customer bridged LAN.
- Compatible with existing IPv4 and IPv6 Routers. RBridges are as transparent to routers as bridges are.
- Unicast forwarding tables at transit RBridges scale with the number of RBridges, not the number of end stations. Transit RBridges do not learn end station addresses.
- MTU feature lets TRILL safely make use of larger campus MTUs and would support traffic engineered routes based on MTU.

Peering: Are RBridges Bridges or Routers?

- They are obviously Bridges because
  - RBridges deliver unmodified frames from the source end station to the destination end station
  - RBridges can operate with zero configuration and auto-configure themselves
  - RBridges provide the same restriction of frames to VLANs and support frame priorities as IEEE 802.1Q-2011 bridges do
  - RBridges, by default, learn MAC addresses from the data plane

Peering: Are RBridges Bridges or Routers?

- They are obviously Routers because
  - RBridges route frames using a routing protocol rather than forwarding with the spanning tree protocol
    - RBridges swap the outer addresses on each RBridge hop from the ingress RBridge to the egress RBridge
    - RBridges decrement a hop count in TRILL frames on each hop
  - RBridges optionally learn MAC addresses by distribution through the control plane
  - RBridges can act based on IP multicast control messages (IGMP, MLD, and MRD) and restrict the distribution of IP derived multicast frames

Peering: Are RBridges Bridges or Routers?

- Really, they are a new species, between IEEE 802.1 bridges and routers:

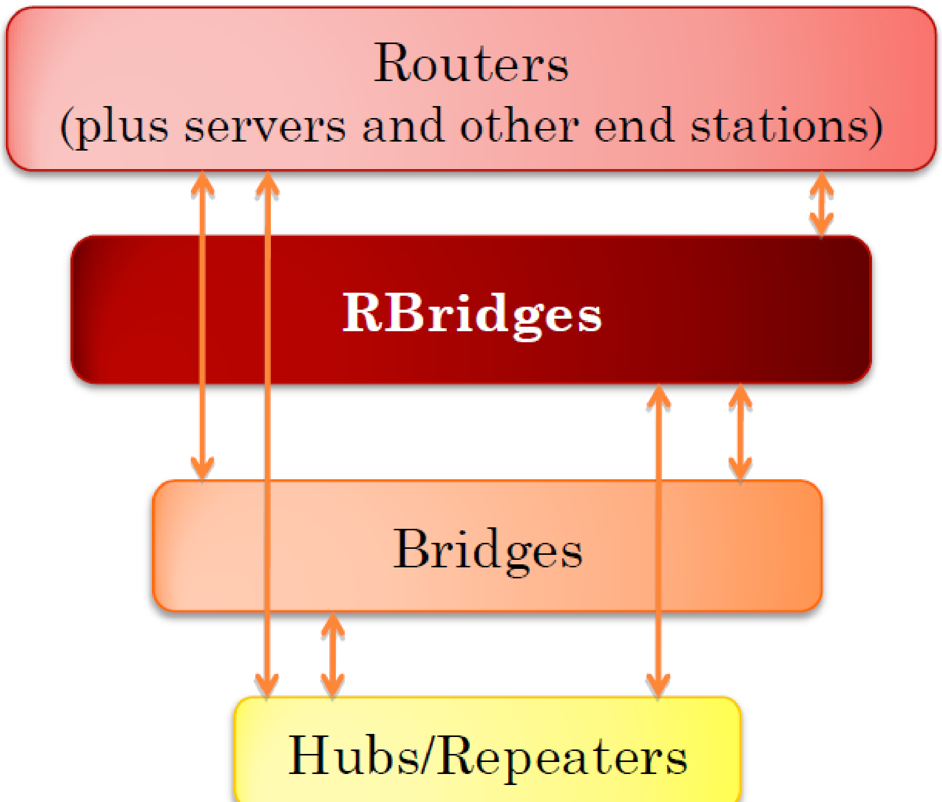

TRILL Silicon

- Here are six publicly known independent silicon implementations of the TRILL Fast Path. In some cases the vendor has multiple different chips supporting TRILL.
  - Broadcom – merchant silicon
  - Brocade – products
  - Cisco – products
  - Fulcrum – merchant silicon
  - Marvell – merchant silicon
  - Mellanox – merchant silicon

TRILL Open Source

Three Open Source Implementations

1. **Oracle: TRILL for Solaris**

   TRILL ships as part of Solaris 11

2. **VirtuOR:** www.virtuor.fr

   http://sourceforge.net/p/opentrill/wiki/Home/

3. TRILL Port to Linux:
   National University of Sciences and Technology
   Islamabad, Pakistan

   Muhammad Mohsin Sardar
   mohsin.sardar@seecs.edu.pk

   http://wisnet.seecs.nust.edu.pk/projects/trill/index.html

TRILL Use Cases

TRILL use Case

- “Acme Power Plant” Process Control
  - Large process control commonly uses Ethernet
  - Some process control protocols interpret network interruption >1 second as equipment failure
  - Even Rapid Spanning Tree Protocol can take >3 second to recover from root bridge failure
  - Core RBridges reduce/eliminate spanning tree

1. Acme Power Plant

Bridged Process Control Network

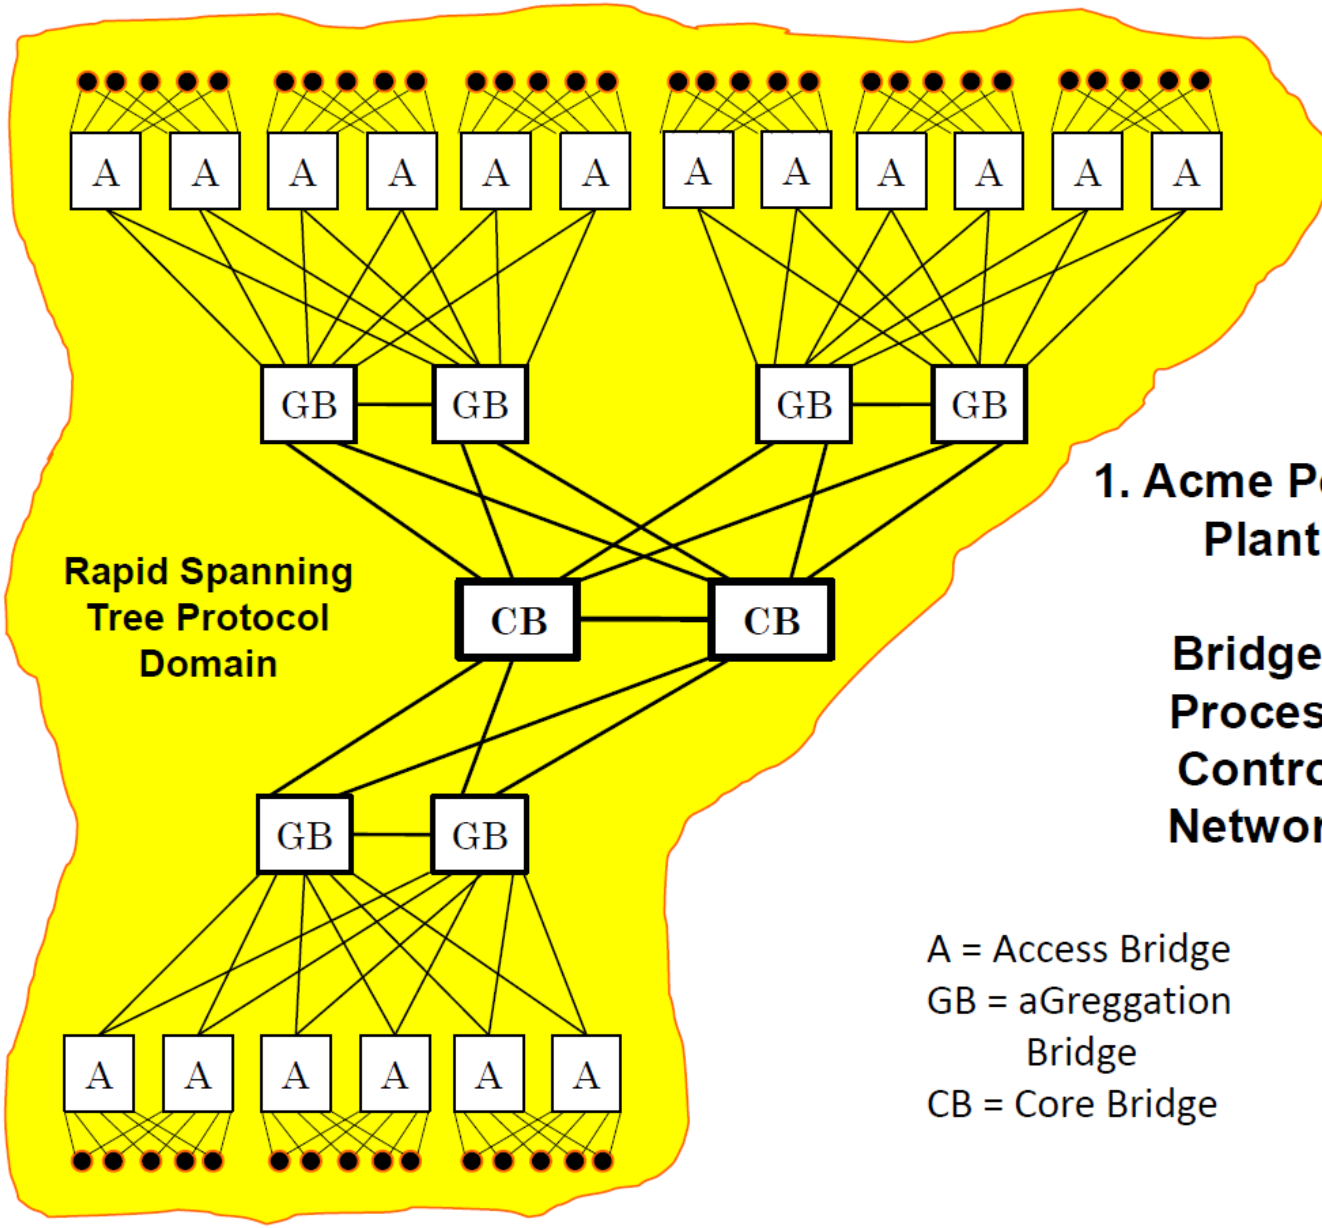

A = Access Bridge
GB = aGreggation Bridge
CB = Core Bridge

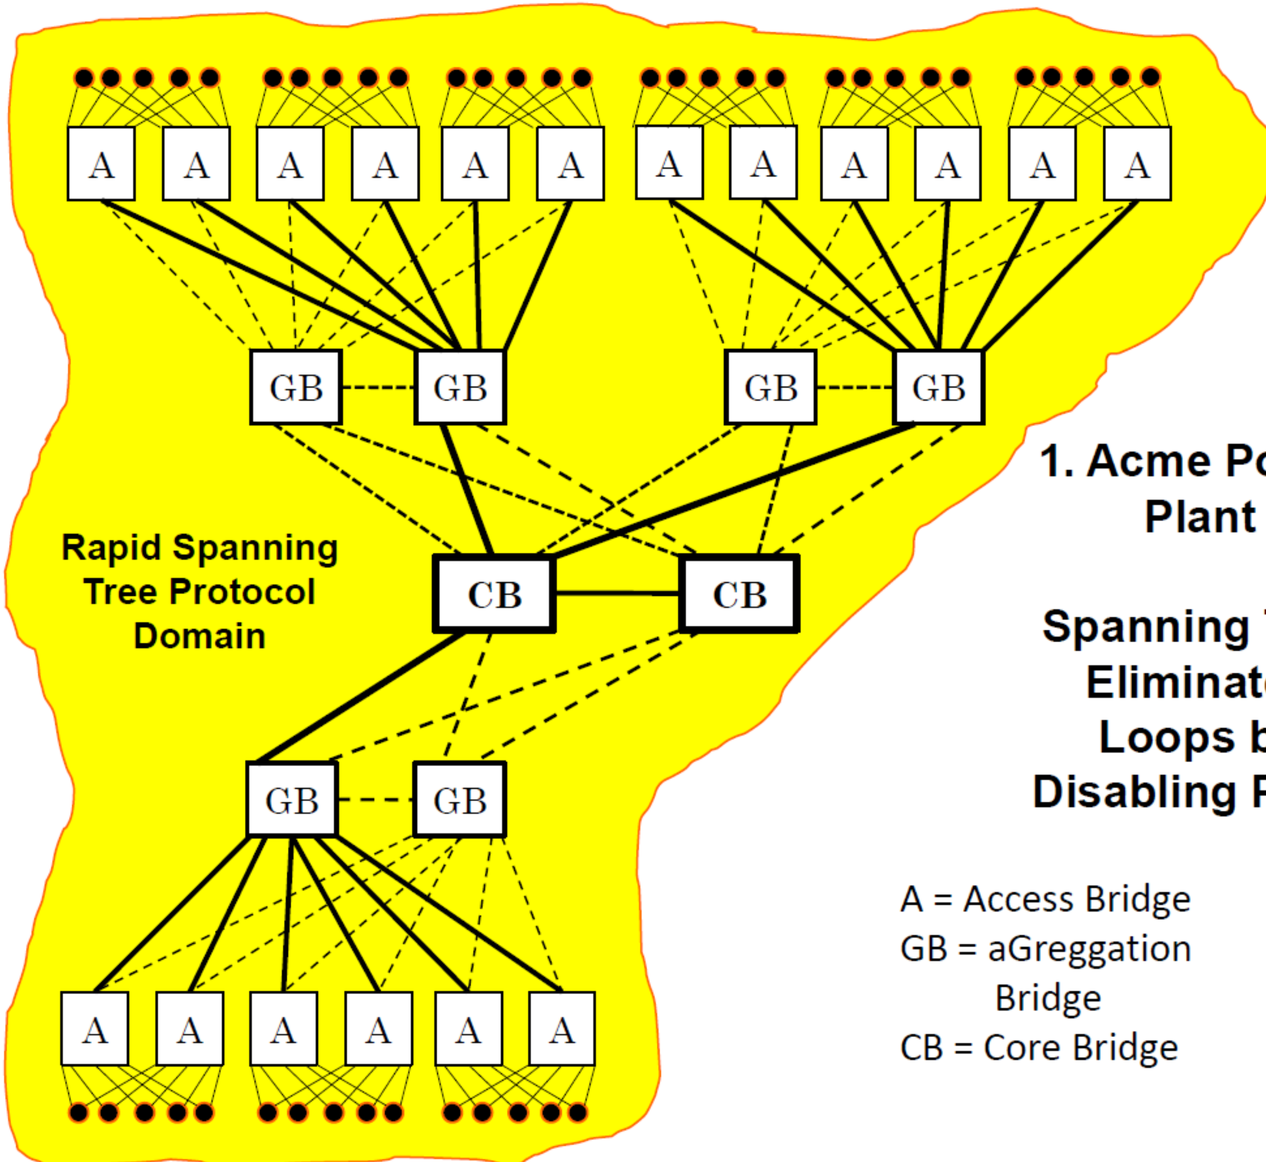

1. Acme Power Plant

Spanning Tree Eliminates Loops by Disabling Ports

A = Access Bridge
GB = aGreggation Bridge
CB = Core Bridge

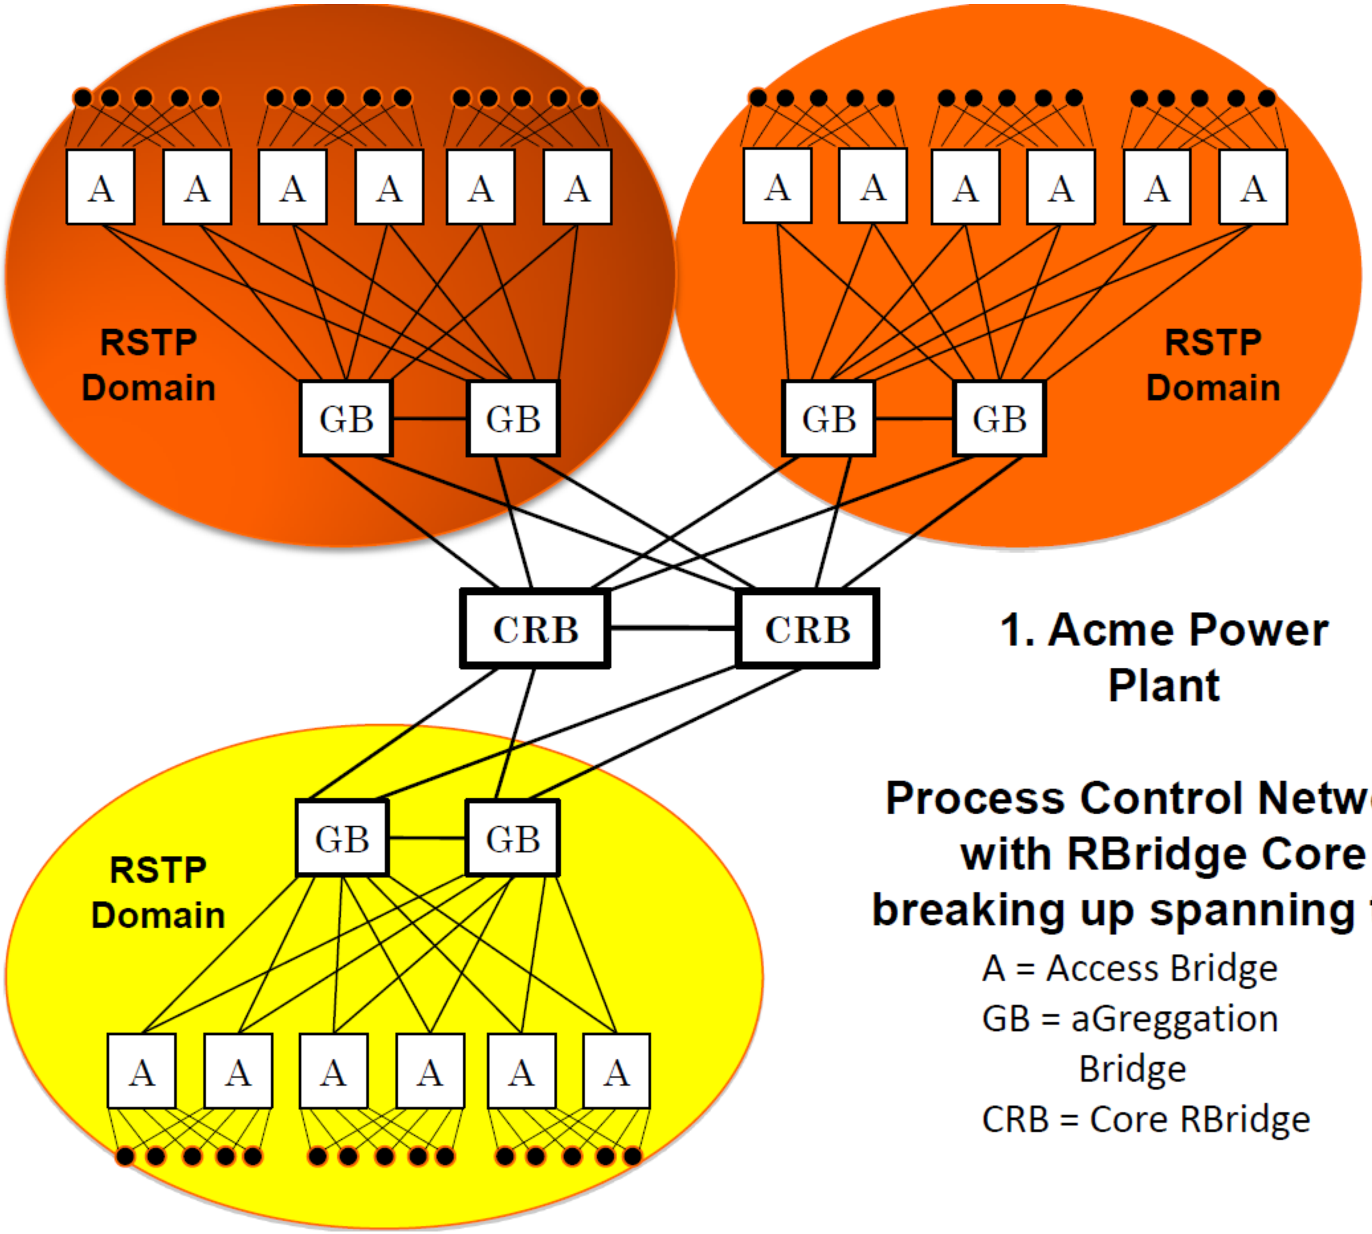

1. Acme Power Plant

Process Control Network with RBridge Core breaking up spanning tree

A = Access Bridge
GB = aGreggation Bridge
CRB = Core RBridge

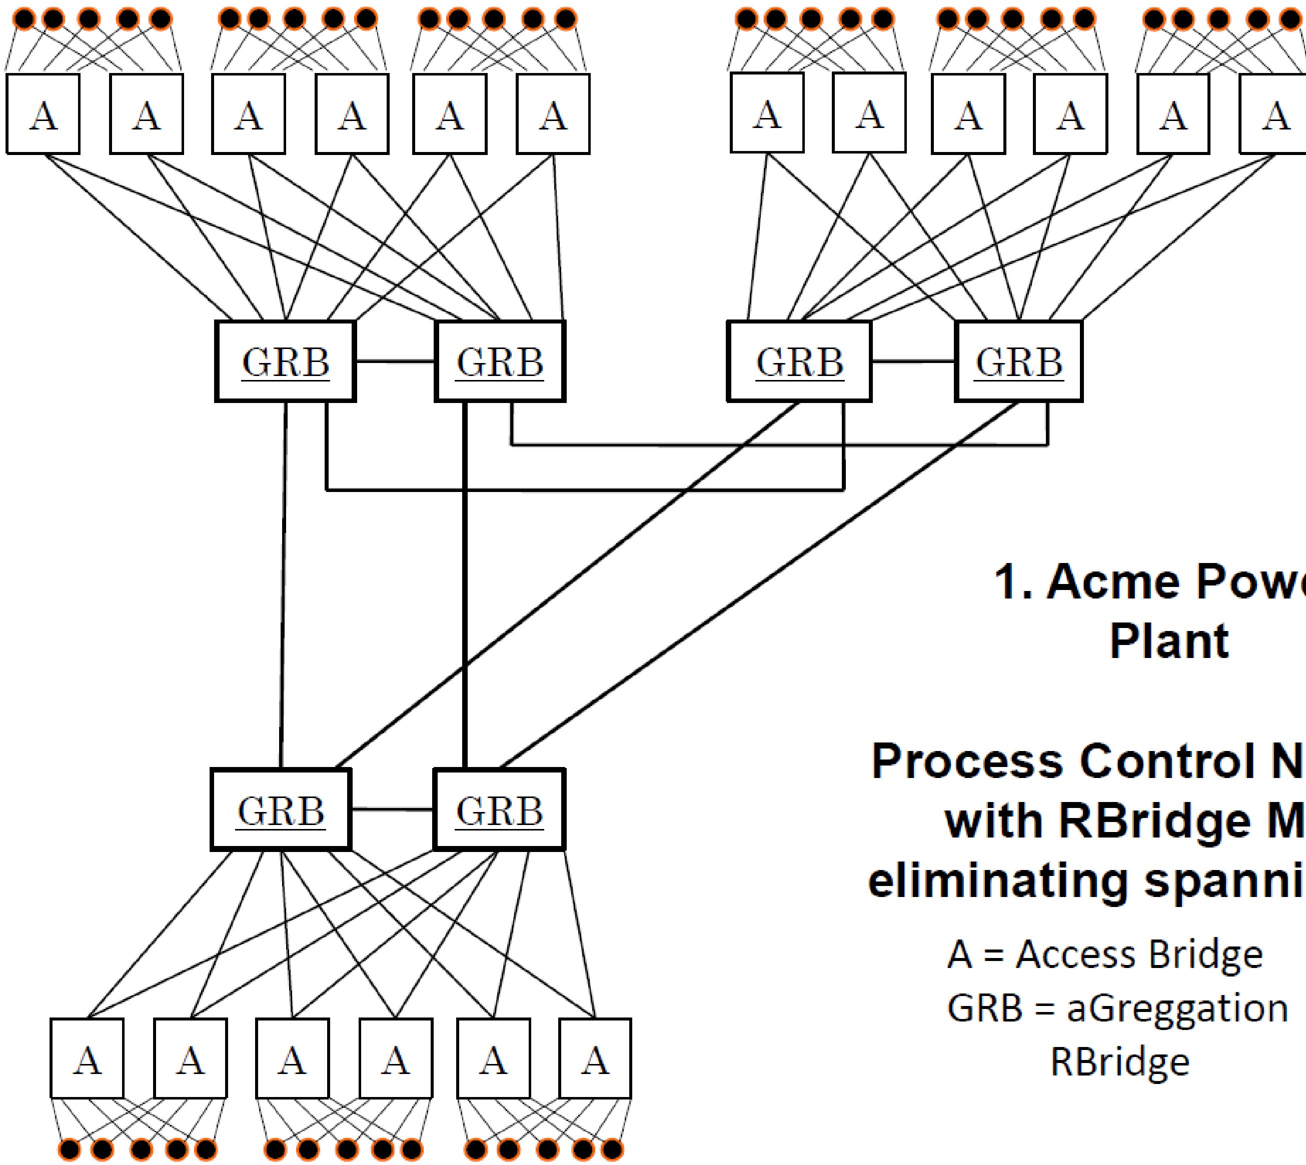

1. Acme Power Plant

Process Control Network with RBridge Mesh eliminating spanning tree

A = Access Bridge
GRB = aGreggation RBridge

Internet Exchange Point (IXP)

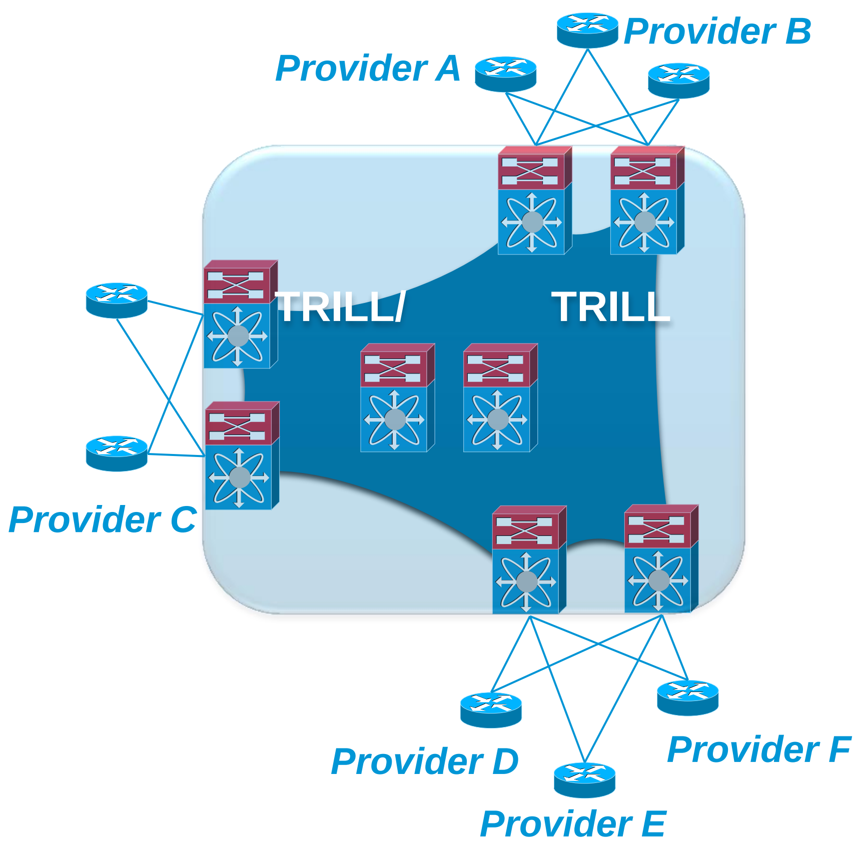

IXP Requirements

- Layer 2 Peering
- 10GE non-blocking Fabric
- Scale to thousands of ports

TRILL Benefits for IXP

- Layer 2 Fabric
- Non-blocking up to thousands 10GE ports
- Simple to manage
- No design constraint, easy to grow

TRILL Operation

New Control Plane

Plug-n-Play L2 IS-IS Manages Forwarding Topology

- IS-IS assigns addresses to all TRILL switches automatically
- Compute shortest, pair-wise paths
- Support equal-cost paths between any TRILL switch pairs

TRILL Routing Table

Switch	IF
S10	L1
S20	L2
S30	L3
S40	L4
S200	L1, L2, L3, L4
…	…
S400	L1, L2, L3, L4

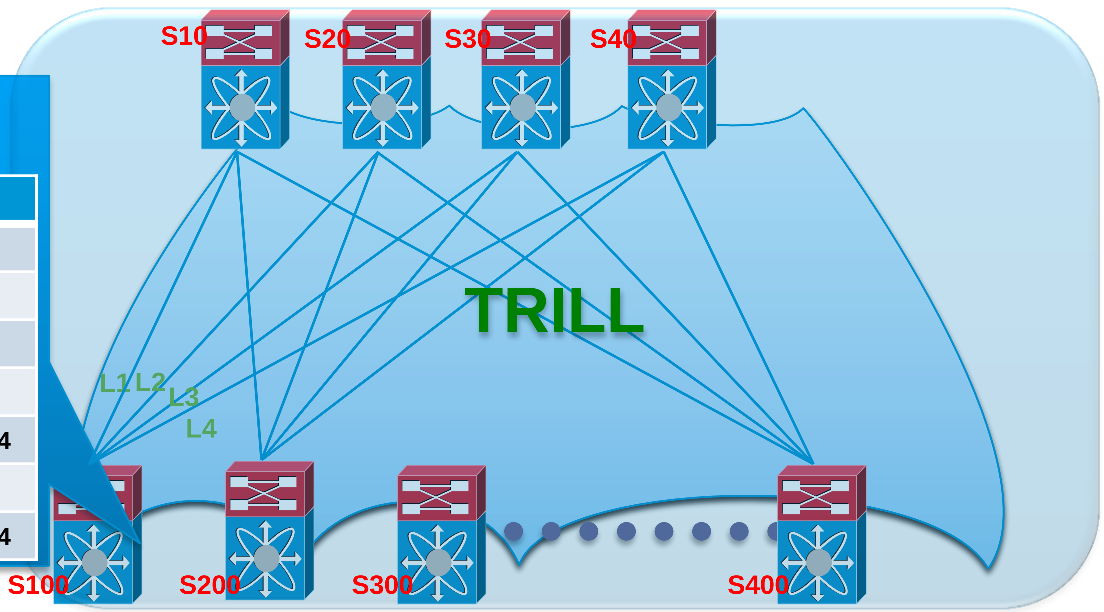

TRILL Control Plane Operation: ISIS

Plug-N-Play L2 IS-IS is used to manage forwarding topology

- Assigned switch addresses to all TRILL enabled switches automatically (no user configuration required)
- Compute shortest, pair-wise paths
- Support equal-cost paths between any TRILL switch pairs

TRILL Routing Table

Switch	IF
S1	L1
S2	L2
S3	L3
S4	L4
S12	L1, L2, L3, L4
…	…
S42	L1, L2, L3, L4

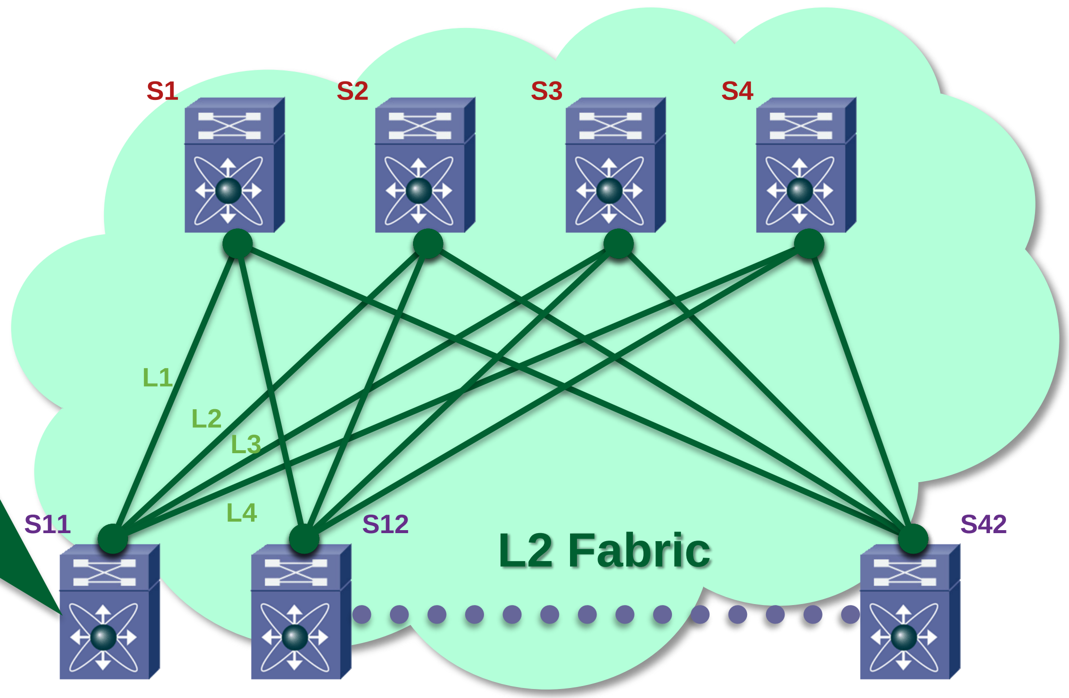

New Data Plane

- The association MAC address/Nickname is maintained at the edge

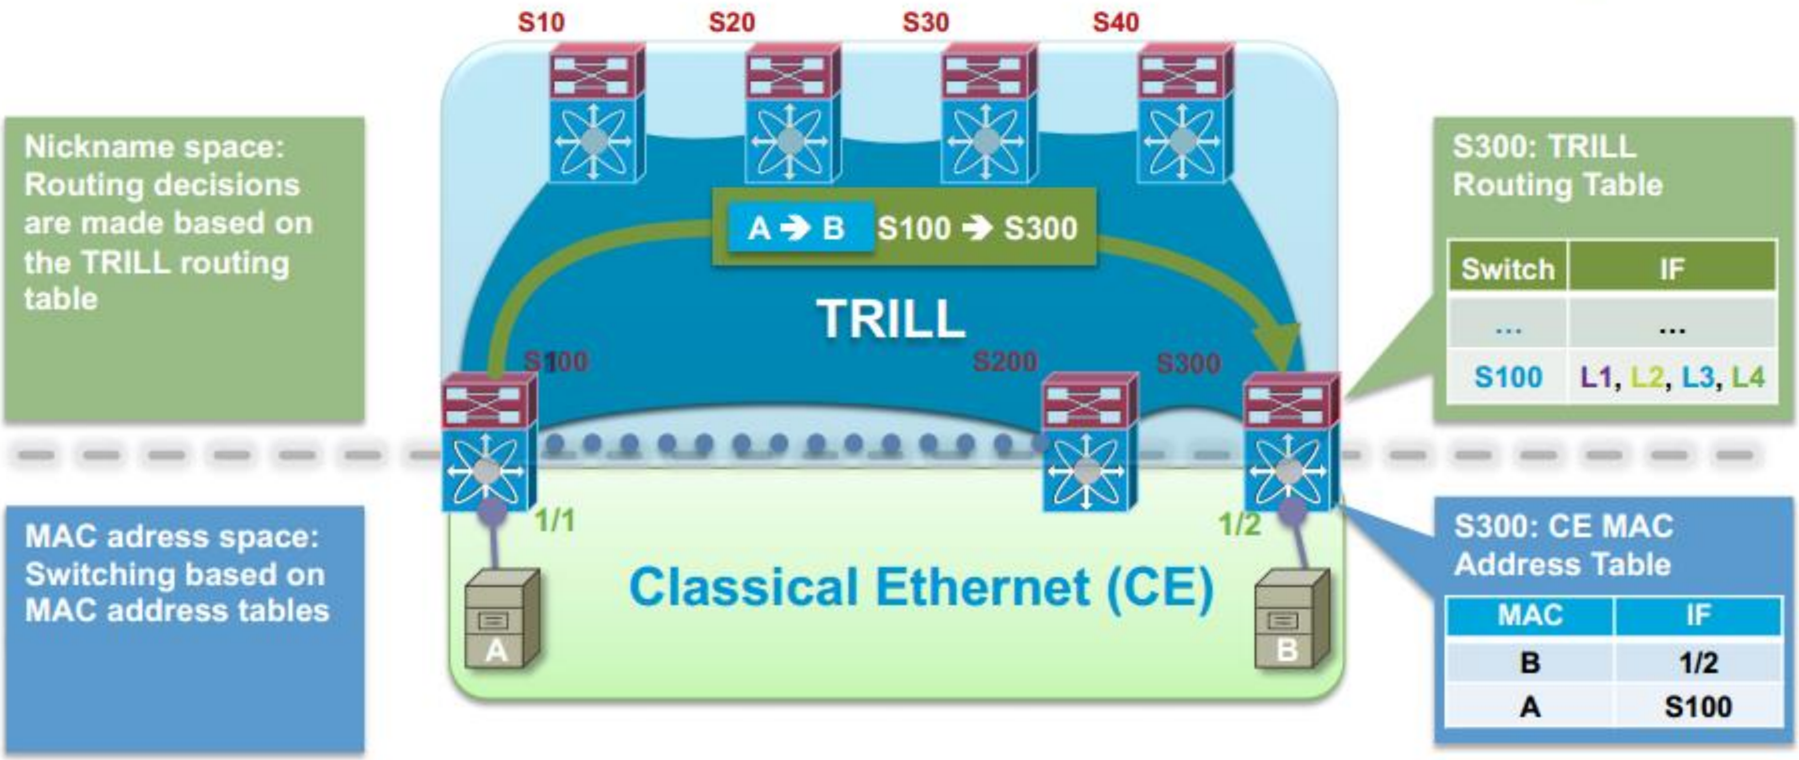

- Traffic is encapsulated across the Fabric

Unknown Unicast in TRILL

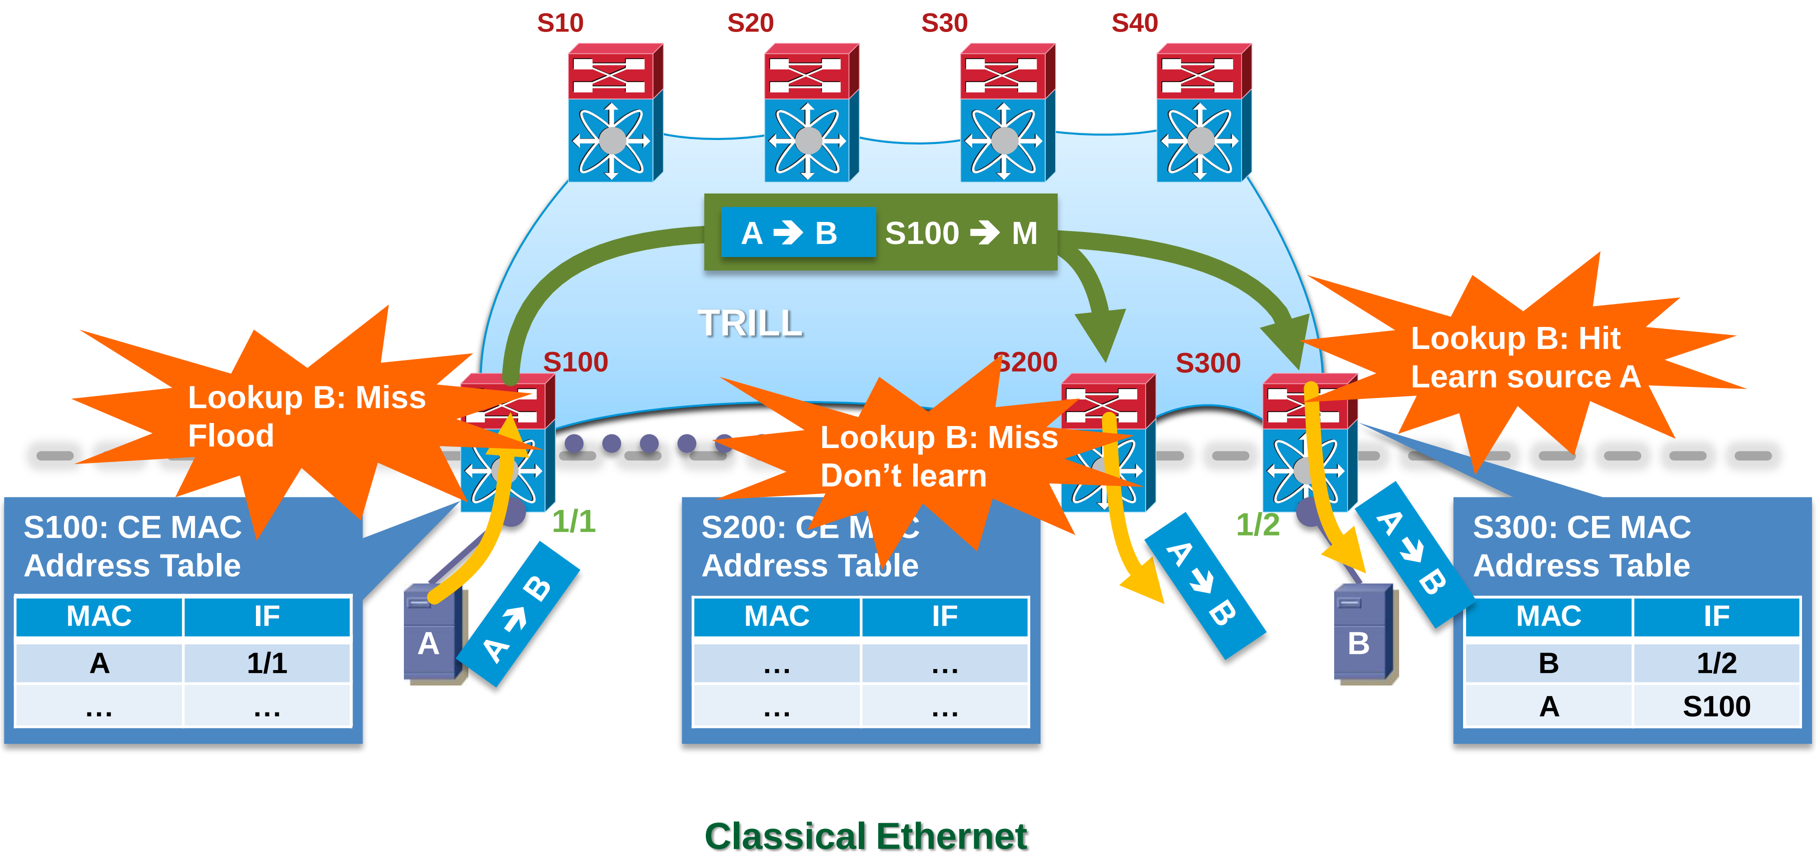

Known Unicast, Conversational Learning in TRILL

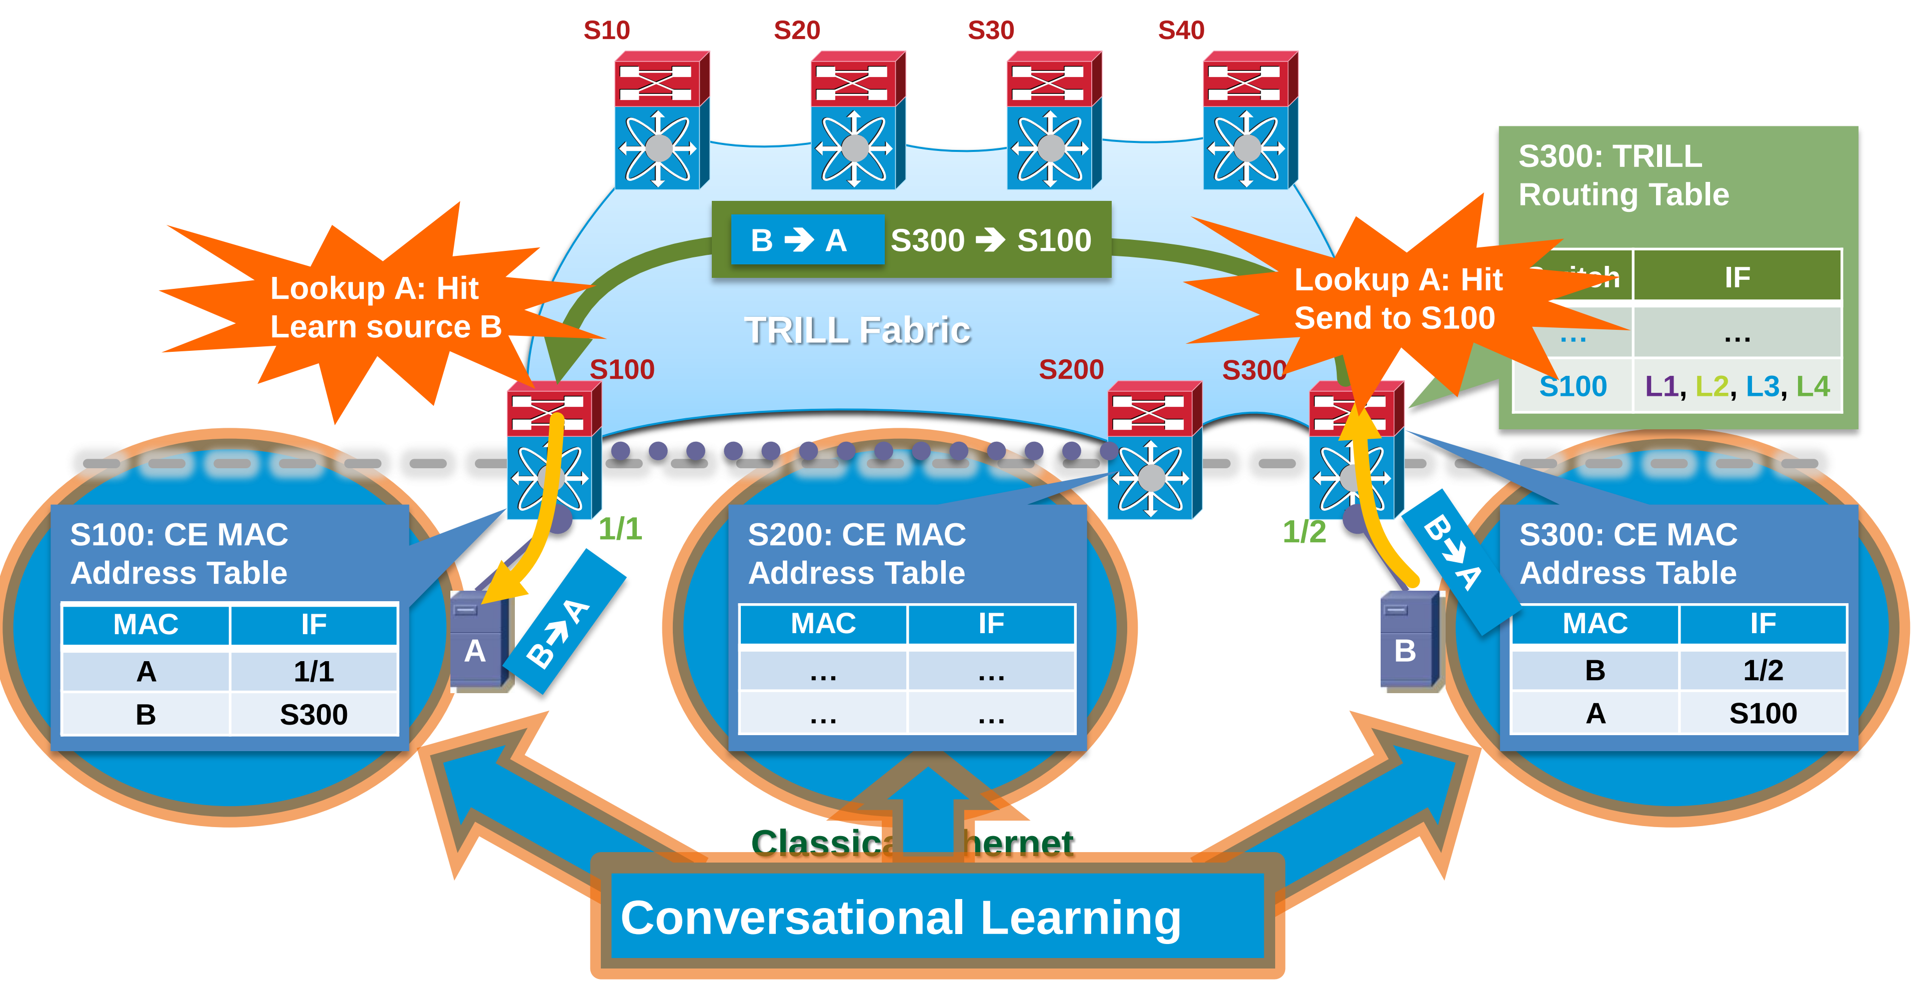

TRILL Frames Overview

Frame Types

- Frame Type Names Used in TRILL
  - TRILL Frames
    - TRILL IS-IS Frames – Used for routing control between RBridges.
    - TRILL Data Frames – Used for encapsulated native frames.
  - Layer 2 Control Frames – Bridging control, LLDP, LACP, etc. Never forwarded by RBridges.
  - Native Frames – All frames that are not TRILL or Layer 2 Control Frames.

TRILL
Encapsulation and Header

- Summary of reasons for encapsulation:
  - Provides a hop count to mitigate loop issues
  - To hide the original source address to avoid confusing any bridges present as might happen if multi-pathing were in use
  - To direct unicast frames toward the egress RBridge so that forwarding tables in transit RBridges need only be sized with the number of RBridges in the campus, not the number of end stations
  - To provide a separate outer VLAN tag, when necessary, for forwarding traffic between RBridges, independent of the original VLAN of the frame

TRILL Encapsulation and Header

- TRILL Data frames between RBridges are encapsulated in a local link header and TRILL Header.
  - The local link header is addressed from the local source RBridge to the next hop RBridge for known unicast frames or to the All-RBridges multicast address for multi-destination frames.
  - The TRILL header specifies the first/ingress RBridge and either the last/egress RBridge for known unicast frames or the distribution tree for multi-destination frames.

TRILL Data
Ethernet Encapsulation

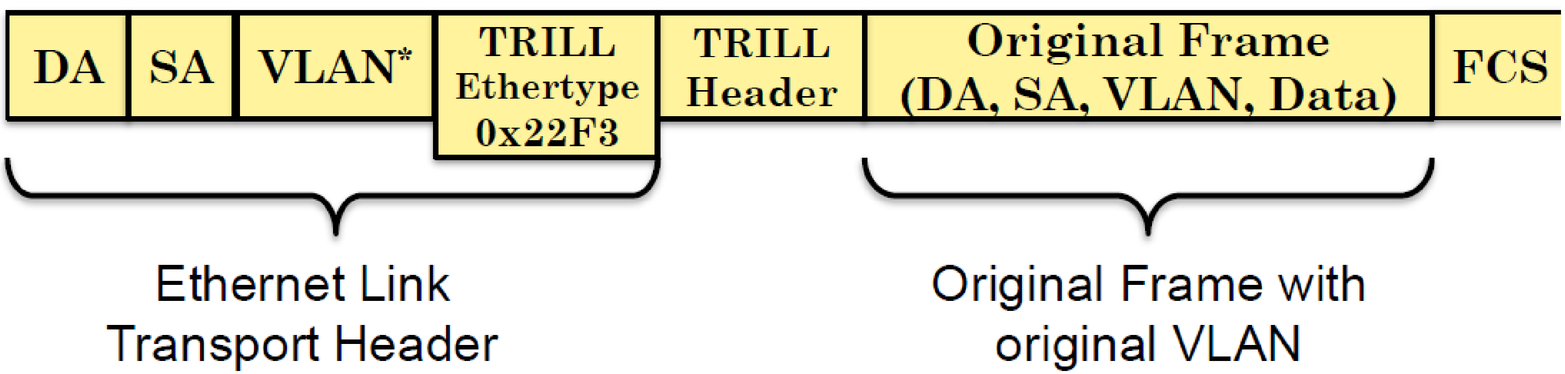

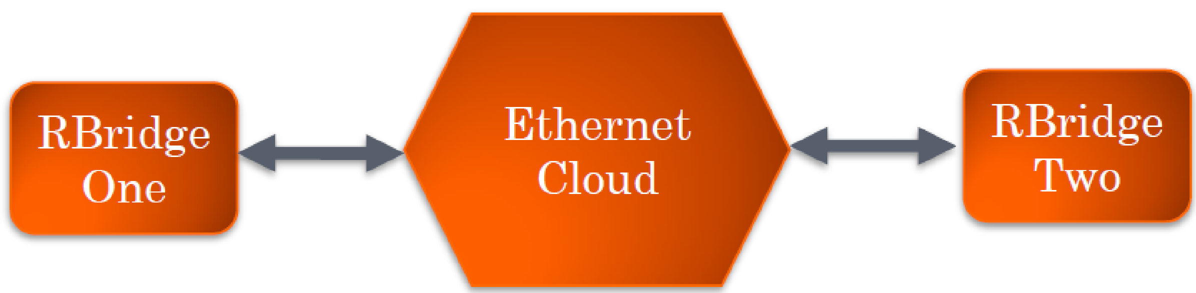

* Link Transport VLAN need only be present if RBridges are connected by a bridged LAN or carrier Ethernet requiring a VLAN tag or the like.

TRILL Encapsulation and Header

- Assuming the link is Ethernet (IEEE 802.3) the encapsulation looks like:
  1. Outer Ethernet Header
     - Source RBridge One, Destination RBridge Two
  2. (Outer VLAN Tag if Necessary)
  3. TRILL Header
  4. Inner Ethernet Header
     - Original Source and Destination Addresses
  5. Mandatory Inner VLAN Tag
  6. Original Payload
  7. Frame Check Sequence (FCS)

TRILL Unicast Ingress

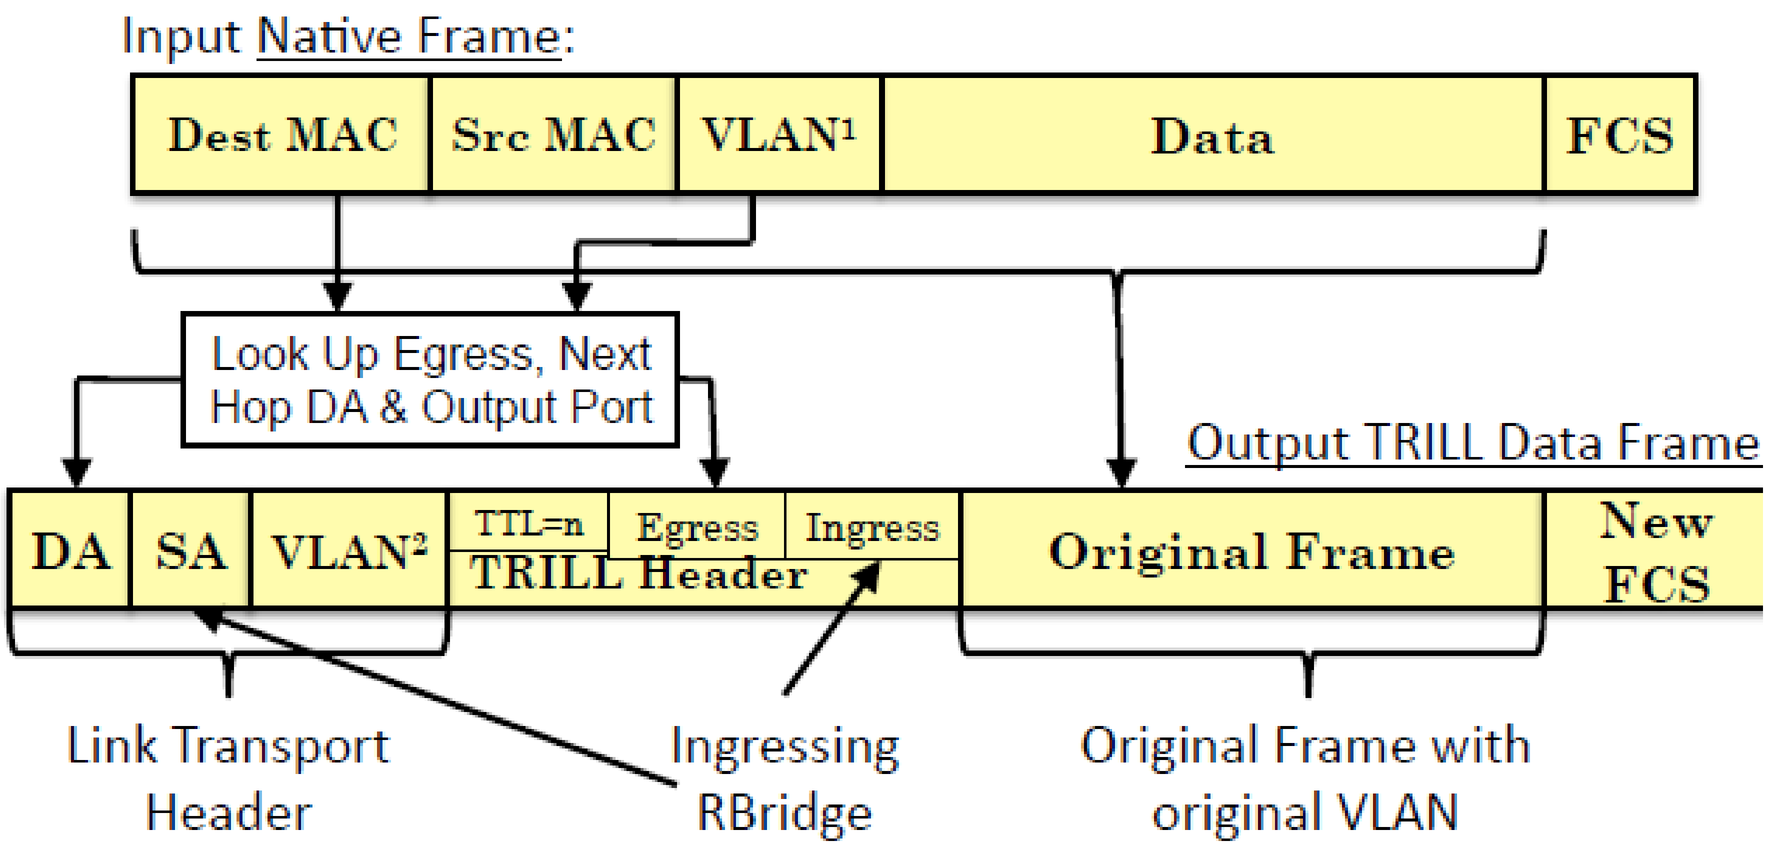

[1] If initial VLAN tag not present, it will be added in encapsulation.
[2] Outer VLAN tag (Designated VLAN) is a transport artifact and is only needed if RBridges are connected by a bridged LAN or carrier Ethernet requiring a VLAN tag or the like.

TRILL Unicast Transit

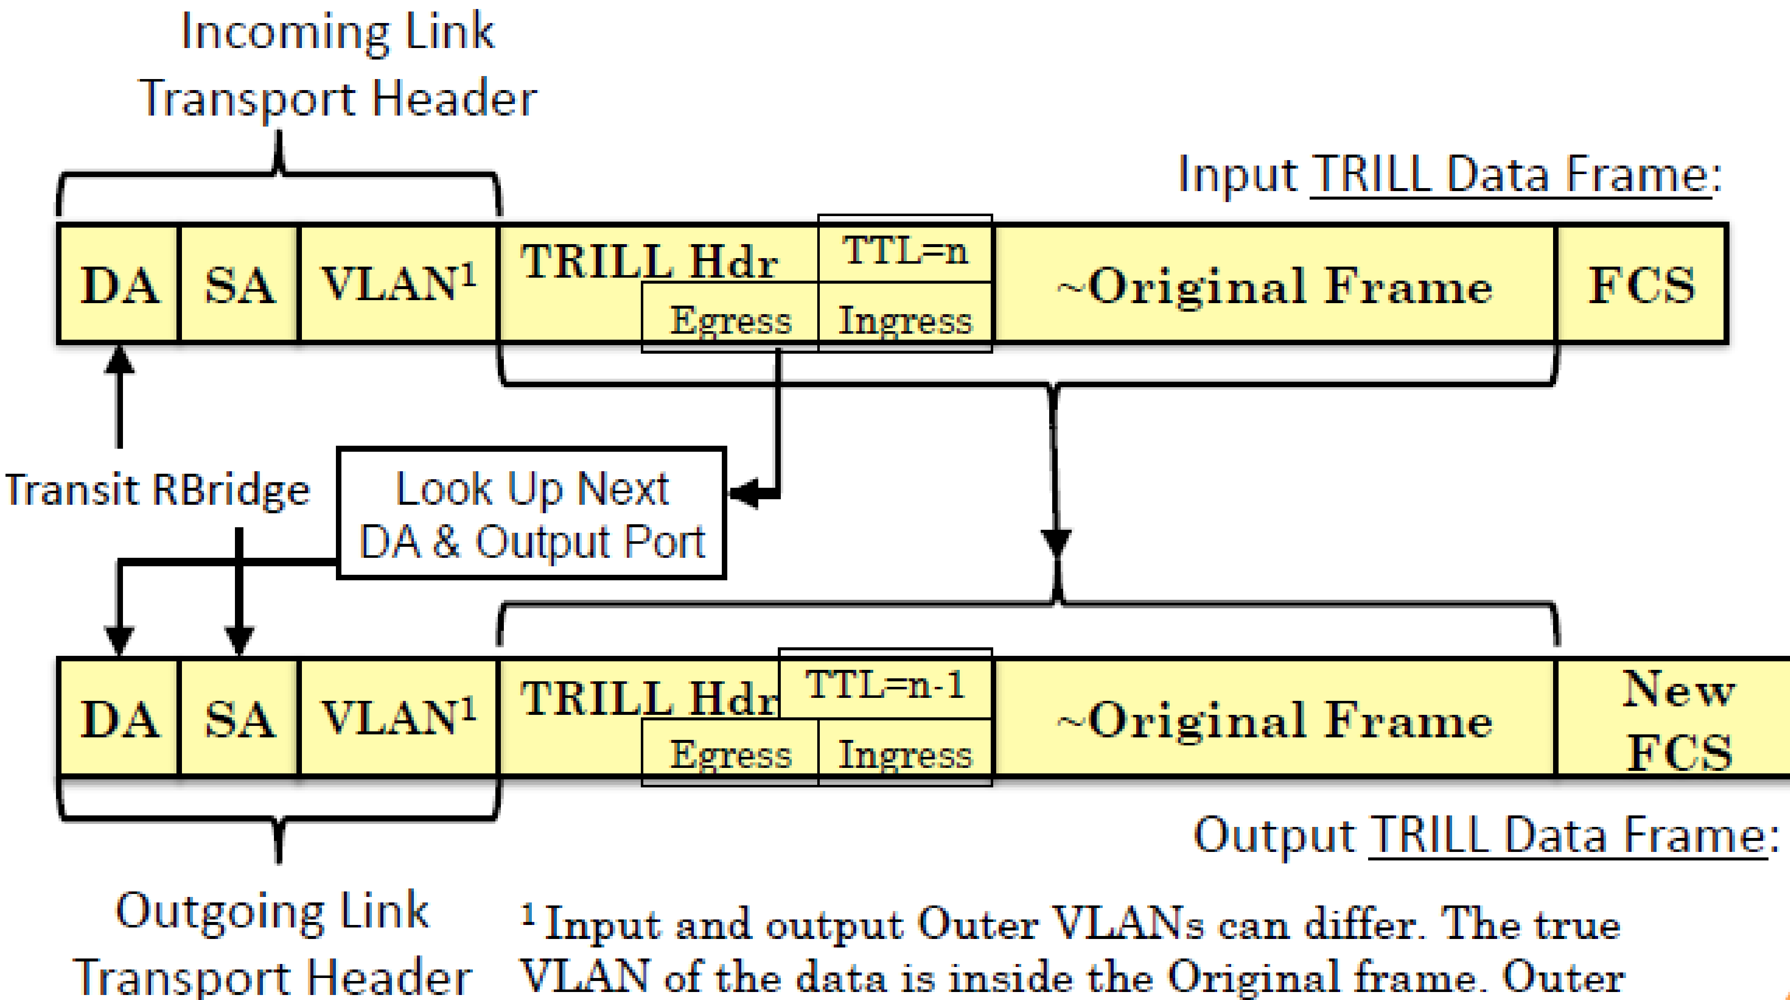

[1] Input and output Outer VLANs can differ. The true VLAN of the data is inside the Original frame. Outer VLAN only needed if RBridges are connected by a bridged LAN or carrier Ethernet requiring a VLAN tag or the like.

TRILL Unicast Egress

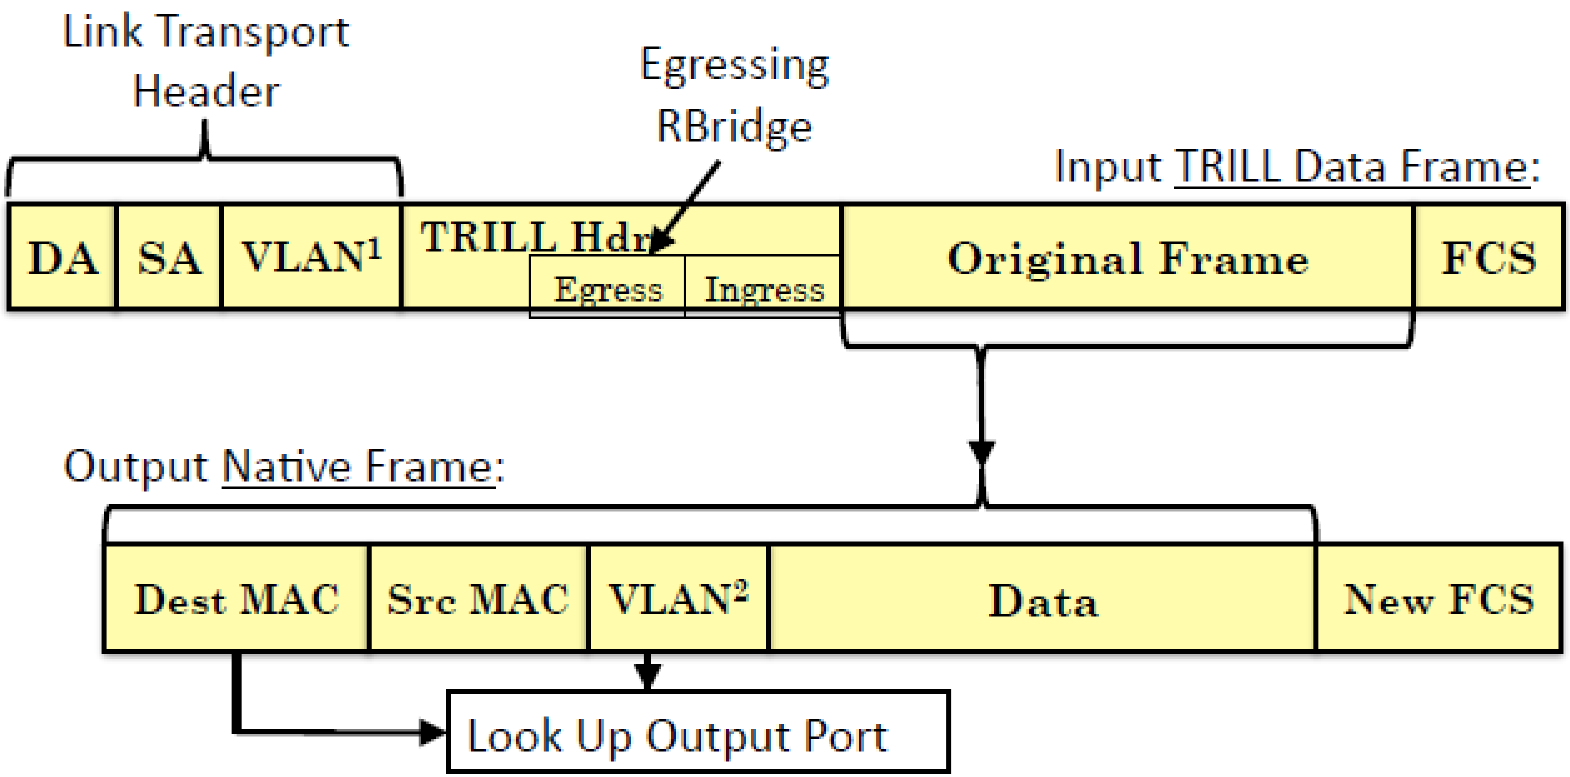

[1] Outer VLAN only needed if RBridges are connected by a bridged LAN or carrier Ethernet requiring a VLAN tag or the like

[2] Final native frame VLAN tag may be omitted depending on RBridge output port configuration.

TRILL Encapsulation and Header

- TRILL Header – 8 bytes

TRILL Ethertype	**V**	**R**	**M**	**OpLng**	**Hop**
Egress RBridge Nickname	**Ingress RBridge Nickname**				

  - Nicknames – auto-configured 16-bit campus local names for RBridges
  - V = Version (2 bits)
  - R = Reserved (2 bits)
  - M = Multi-Destination (1 bit)
  - OpLng = Length of TRILL Options
  - Hop = Hop Limit (6 bits)

Multicast and Multi-destination Handling in TRILL

Packet Flow — Multicast/Broadcast/Unknown Unicast

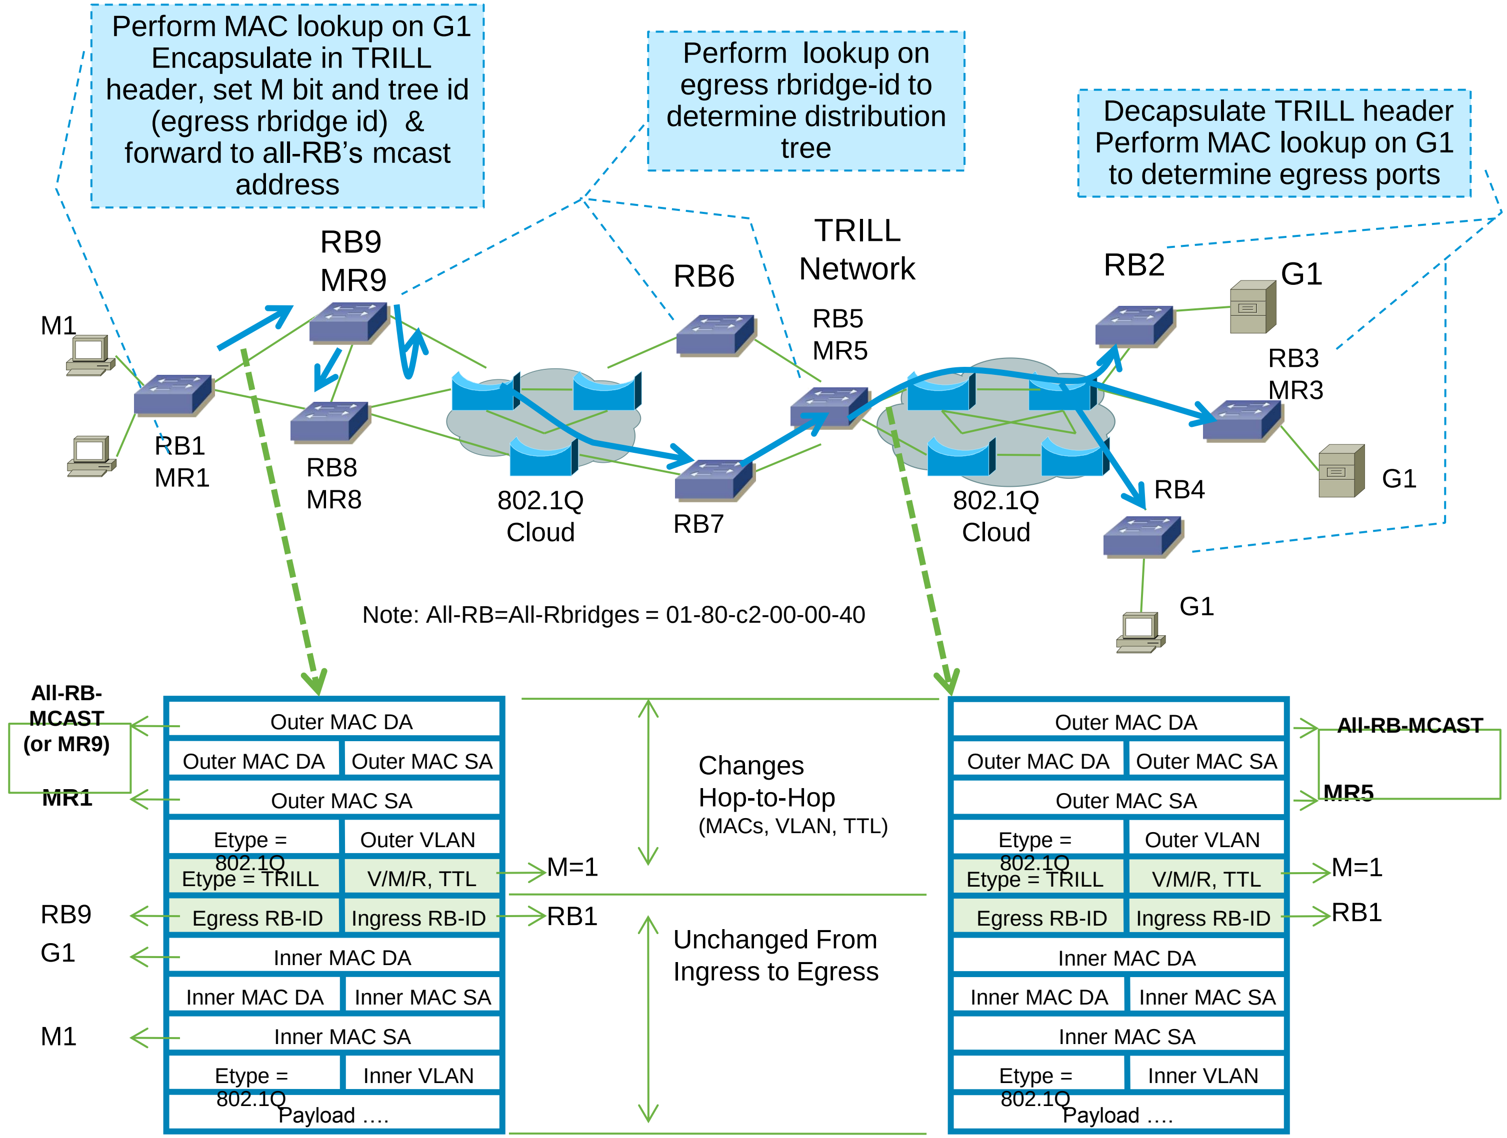

TRILL IP Multicast

- **Control plane:**
  - IGMP snooping operates as usual in TRILL edge switches
  - TRILL IS-IS learns multicast group membership from IGMP snooping on edge switch
  - TRILL edge switch announces group interest by using GM-LSPs, creating “pruned trees” for each group on each multidestination tree
- **Data plane:**
  - Hardware selects which multidestination tree to use for each flow based on hash function
  - Once tree selected, traffic constrained to pruned tree for that IP multicast group, based on MAC table lookup

IGMP Snooping in TRILL

- IGMP snooping learns of interested receivers on TRILL edge switches
- Membership tracked on CE ports based on receiving IGMP reports/leaves
  - Only locally connected receivers tracked on a given edge switch
- Group membership advertised in TRILL IS-IS using GM-LSPs

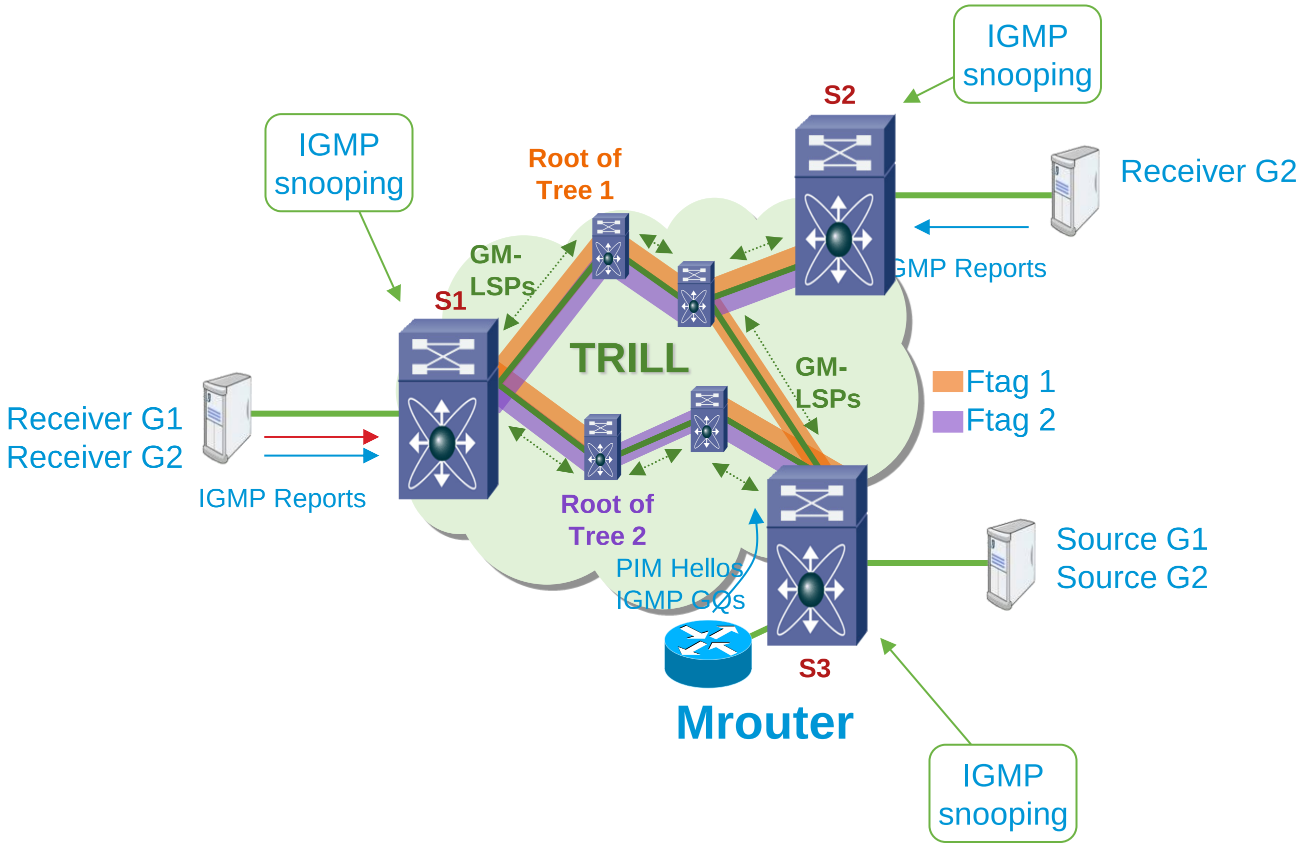

Standardization Status

- The -16 version of the TRILL base protocol specification was approved as an IETF standard on 15 March 2010. The protocol is being implemented by many companies.

TRILL in 2009/2011

- 2009: RFC 5556 “TRILL: Problem and Applicability Statement”
- 2009: TRILL Protocol passed up to IESG for Approval.
- **2010: TRILL approved IETF Standard (2010-03-15)**
  - Ethertypes, Multicast addresses & NLPID assigned
- 2010: Successful TRILL control plane interop at UNH IOL
- 2011: TRILL Protocol base document set:
  - RFC 6325: “RBridges: TRILL Base Protocol Specification”
  - RFC 6326: “TRILL Use of IS-IS”
  - RFC 6327: “RBridges: Adjacency”
  - RFC 6361: “TRILL over PPP”
  - RFC 6439: “RBridges: Appointed Forwarders”
- 2011: TRILL Working Group Re-Chartered to do further development of the TRILL protocol

STANDARDIZATION STATUS

- The TRILL protocol RFCs
  - RFC 5556, "TRILL Problem and Applicability"
  - **RFC 6325, "RBridges: TRILL Base Protocol Specification"**
  - RFC 6326, "TRILL Use of IS-IS"
  - RFC 6327, "RBridges: Adjacency"
  - RFC 6361, "TRILL over PPP"
  - RFC 6439, "RBridges: Appointed Forwarders"
  - RFC 6847, "FCoE over TRILL"
  - RFC 6850, "Definitions of Managed Objects for RBridges" (MIB)

TRILL in 2012/2013

2012: Second Successful TRILL control plane interop at UNH IOL

2013: Additional TRILL documents published:

- RFC 6447: FCoE (Fibre Channel over Ethernet) over TRILL
- RFC 6850: RBridge MIB
- RFC 6905: TRILL OAM Requirements

2013: Third TRILL interop for control and data plane at UNH IOL week of May 20th

2013: TRILL Working Group Re-Chartered to do further development of the TRILL protocol

Standardization Status

- The TRILL protocol RFCs (**bold** = stds track)
  - RFC 5556, “TRILL Problem and Applicability”
  - **RFC 6325, “RBridges: TRILL Base Protocol Specification”**
  - **RFC 6326, “TRILL Use of IS-IS”**
  - **RFC 6327, “RBridges: Adjacency”**
  - **RFC 6361, “TRILL over PPP”**
  - **RFC 6439, “RBridges: Appointed Forwarders”**
  - RFC 6847, “FCoE over TRILL”
  - **RFC 6850, “Definitions of Managed Objects for RBridges” (MIB)**
  - RFC 6905, “TRILL OAM Requirements”

Standardization Status

Document that are fully approved and in the RFC Editor's Queue. These are expected to issue as standards track RFCs soon:

- "TRILL: Fine Grained Labeling:
  - https://datatracker.ietf.org/doc/draft-ietf-trill-fine-labeling/
- "TRILL: BFD Support"
  - https://datatracker.ietf.org/doc/draft-ietf-trill-rbridge-bfd/
- "TRILL: RBridge Channel Support"
  - https://datatracker.ietf.org/doc/draft-ietf-trill-rbridge-channel/
- "TRILL: Edge Directory Assistance Framework"
  - https://datatracker.ietf.org/doc/draft-ietf-trill-directory-framework/
- "TRILL: Clarifications, Corrections, and Updates"
  - https://datatracker.ietf.org/doc/draft-ietf-trill-clear-correct/
- "TRILL: Header Extension"
  - https://datatracker.ietf.org/doc/draft-ietf-trill-rbridge-extension/

Standardization Status

- Non-IETF Assignments:
  - Ethertypes assigned by IEEE:
    - TRILL Data: 0x22F3
    - TRILL IS-IS: 0x22F4
    - TRILL Fine Grained Labeling: 0x893B
    - RBridge Channel: 0x8946
  - Block of multicast addresses assigned to TRILL by IEEE:
    - 01-80-C2-00-00-40 to 01-80-C2-00-00-4F
  - TRILL NLPID (Network Layer Protocol ID) assigned from ISO/IEC: 0xC0

More TRILL References

- TRILL Introductory Internet Protocol Journal Article:
  http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_14-3/143_trill.html

- The first paper: Perlman, Radia. “Rbridges: Transparent Routing”, Proceeding Infocom 2004, March 2004.
  http://www.ieee-infocom.org/2004/Papers/26_1.PDF

Comparison with MPLS

- TRILL versus MPLS
  - MPLS is an older, more mature technology with better Quality of Service features, etc.
  - MPLS is more configuration intensive. TRILL can be auto-configuring.
  - TRILL provides easier support of multicast
  - TRILL can scale better because
    - MPLS requires a label entry at each LSR (Label Switched Router) for each MPLS path through that LSR
    - TRILL requires a nickname entry at each RBridge for each TRILL switch in the campus

Comparison with IP

- TRILL versus IP
  - IP is an older, more mature technology
  - TRILL supports VM mobility. Changing subnets changes IP Address, breaking TCP connections
  - TRILL is better at multicast because
    - IP requires a complex protocols like PIM to do multicast
    - TRILL has simple multicast distribution, with pruning for optimization, designed in from the start

Thank you.