RBridges/TRILL and IS-IS

Donald Eastlake 3rd
donald.eastlake@stellarswitches.com, +1-508-333-2270
Stellar Switches

CAVEATS

- The base protocol specification may change:
  - It is currently in TRILL WG Last Call.
  - IEEE 802.1 has been asked to comment on the specification from the point of view of its effects on the Ethernet Service Model.
  - It may also be reviewed by the L2VPN WG and non-TRILL-WG IETF experts.
- This presentations is based on "RBridges: Base Protocol Specification" version -10 and "RBridge VLAN Mapping" version -00.

CONTENTS

- Introduction
- TRILL Frames
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL IS-IS Data Requirements
- References

DEFINITIONS

- TRILL –
  TRansparent Interconnection of Lots of Links
  - Being specified by the IETF TRILL Working Group co-chaired by:
    - Erik Nordmark, Sun Microsystems
    - Donald Eastlake 3rd, Stellar Switches
- RBridge – Routing Bridge
  - A device which implements TRILL
- RBridge Campus –
  - A network of RBridges, links, and possibly intervening bridges bounded by end stations.

WHAT/WHO TRILL?

- What is TRILL?
  - TRILL is a new protocol to perform Layer 2 bridging based on IS-IS link state routing technology.
- Who started TRILL?
  - Invented by Radia Perlman of Sun Microsystems, the inventor of the Spanning Tree Protocol.

WHY DO TRILL?

- Provides optimum point-to-point forwarding with zero configuration
- Supports multi-pathing of both unicast and multi-destination traffic
- Reduce the size of forwarding tables in the network core
- Provide a flexible options feature
- Other RBridge features:
  - Compatible with spanning tree bridges; can be incrementally deployed into a customer bridged LAN
  - Scale roughly as well as spanning tree bridging

WHY IS-IS FOR TRILL?

- The IS-IS (Intermediate System to Intermediate System) link state routing protocol was chosen for TRILL 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.

CONTENTS

- Introduction
- TRILL Frames
  - Frame Types
  - The TRILL Encapsulation and Header
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL IS-IS Data Requirements
- References

FRAME TYPES

- Frame Type Terminology in the Protocol Spec.
  - Control Frames – Destination Address is 01-80-C2-00-00-00 to 01-80-C2-00-00-0F or 01-80-C2-00-00-21
    - High Level – 01-80-C2-00-00-00 & 01-80-C2-00-00-21
    - Low Level – all other control frames
  - TRILL Frames – Have either the TRILL Ethertype or a TRILL multi-cast destination address.
    - TRILL core IS-IS frames are sent to All-IS-IS-RBridges
    - TRILL data frames are encapsulated native frames
    - TRILL ESADI (End Station Address Distribution Instance) IS-IS frames look like data frames but the encapsulated destination address is All-ESADI-RBridges
  - Native Frames – all frames that are not TRILL or Control Frames

The TRILL Encapsulation and Header

- TRILL data and ESADI frames sent between RBridges are encapsulated inside a local link header, addressed from the local source RBridge to the local destination RBridge, and a TRILL header.

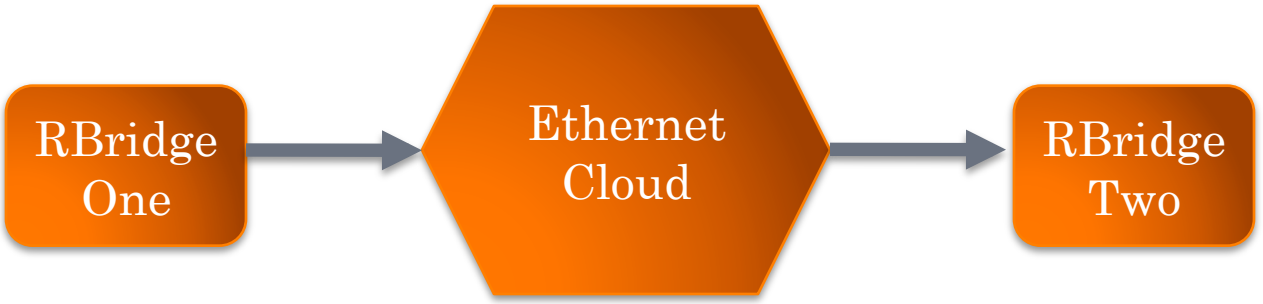

The TRILL Encapsulation and Header

- Some 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 VLAN tag for forwarding traffic between RBridges, independent of the original VLAN of the frame

The 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)
  3. TRILL Header
  4. Inner Ethernet Header
     - Original Source and Destination Addresses
  5. Inner VLAN Tag
  6. Original Payload
  7. Frame Check Sequence (FCS)

The TRILL Encapsulation and Header

- TRILL Header – 64 bits

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)

CONTENTS

- Introduction
- TRILL Frames
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL IS-IS Data Requirements
- References

ARE RBRIDGES BRIDGES OR ROUTERS?

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

Are RBridges Bridges or Routers?

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

Are RBridges Bridges or Routers?

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

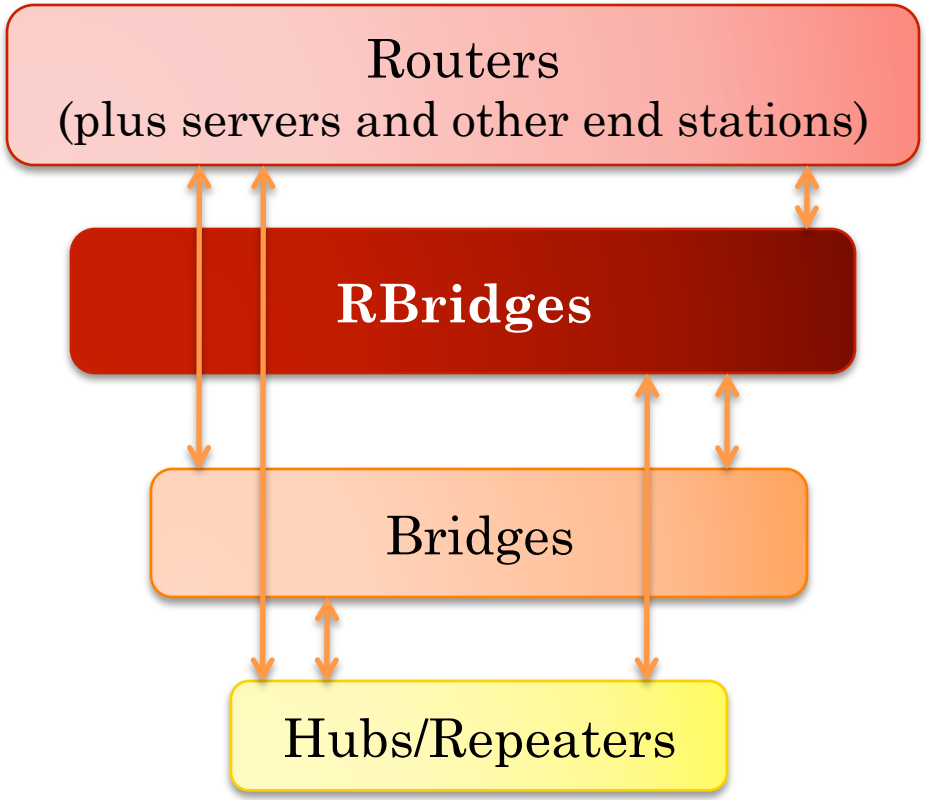

CONTENTS

- Introduction
- TRILL Frames
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL IS-IS Data Requirements
- References

How RBridges Work

- RBridges find each other by exchanging TRILL IS-IS Hello frames
  - Hellos are sent to the All-IS-IS-RBridges multicast address labeled for various VLANs if VLANs have been configured. They are transparently forwarded by bridges, dropped by end stations including routers, and are processed by RBridge ports.
  - The Hellos establish IS-IS connectivity on each RBridge port.
  - Using the information exchanged in the Hellos, the RBridges on each link elect the Designated RBridge for that link.

How RBridges Work

- RBridges use the IS-IS reliable flooding protocol so that each RBridge in the campus has a copy of the global link state database. Each RBridge campus is a single Level 1 Area.
  - The RBridge link state includes information beyond connectivity and link cost. Information such as VLAN connectivity, multicast listeners and multicast router attachment, claimed nickname, options supported, and the like.
  - The database is sufficient for each RBridge to independently and without further messages calculate optimal point-to-point paths for known unicast frames and the same distribution trees for multi-destination frames.

How RBridges Work

- The Designated RBridge specifies the Appointed Forwarder for each VLAN on the link (which may be itself) and the Designated VLAN for inter-RBridge communication.
- The Appointed Forwarder for VLAN-x on a link handles all native frames to/from that link in that VLAN.
  - It encapsulates frames from the link into a TRILL data frame. This is the ingress RBridge function.
  - It decapsulates native frames destined for the link from TRILL data frames. This is the egress RBridge function.

How RBridges Work

- TRILL data frames with
  - known unicast ultimate destinations are forwarded RBridge hop by RBridge hop toward the egress RBridge.
  - multi-destination frames (broadcast, multicast, and unknown destination unicast) are forwarded on a tree rooted at an RBridge selected by the ingress RBridge from among the distribution trees all RBridges in the campus are computing.

CONTENTS

- Introduction
- TRILL Encapsulation
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
  - Structure of an RBridge
  - Structure of an RBridge Port
- TRILL IS-IS Data Requirements
- References

Structure Of An RBridge

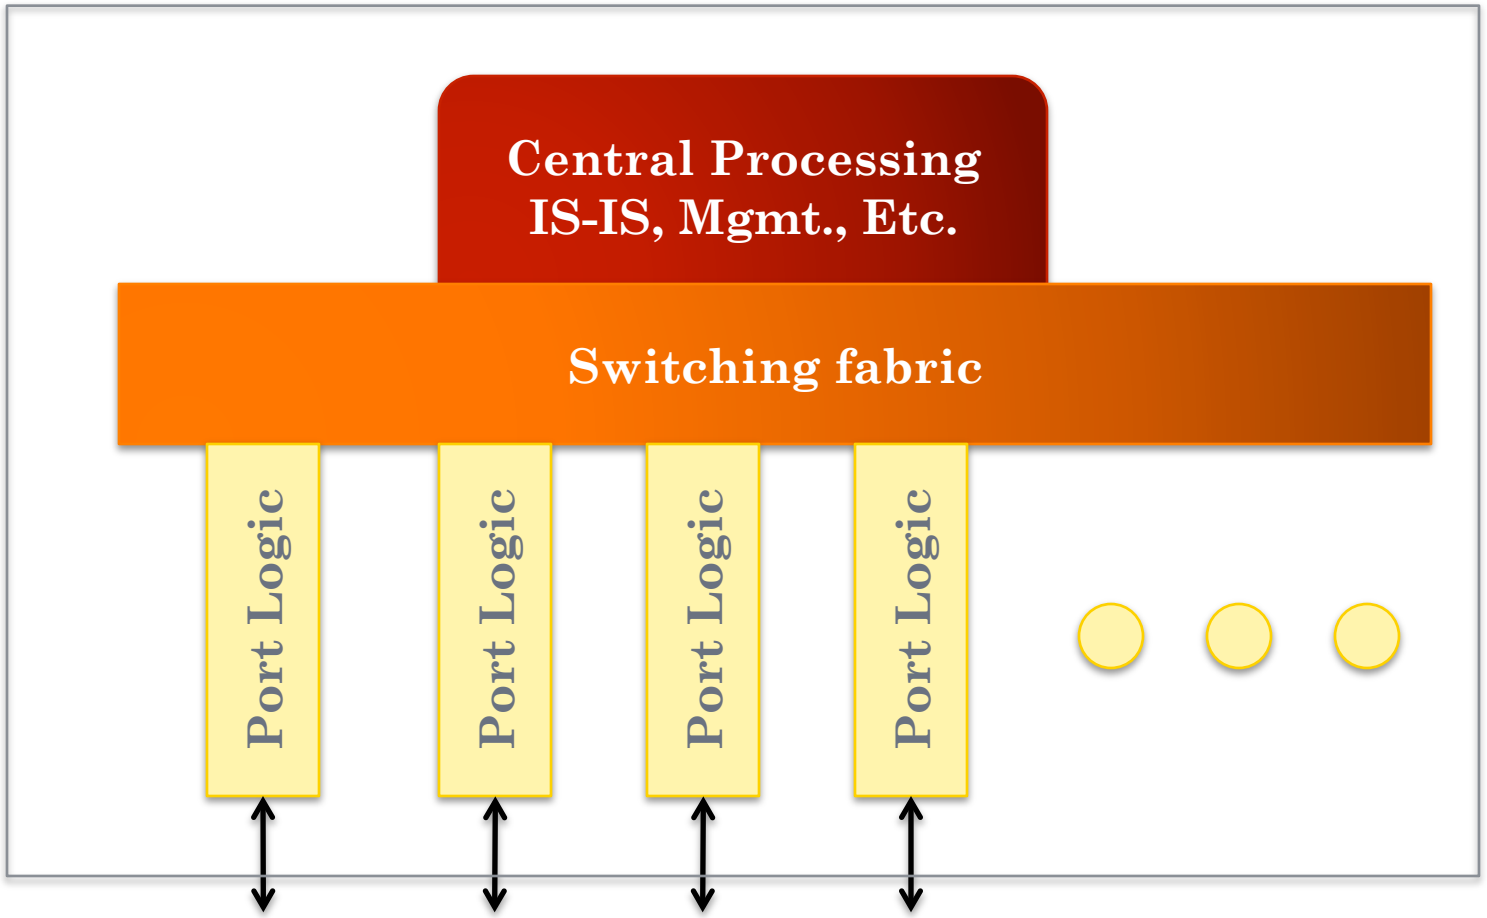

Links to other devices. Could be 802.3 (Ethernet), 802.11 (Wi-Fi), PPP, …

Structure Of An RBridge Port

Assumes an Ethernet (802.3) link.

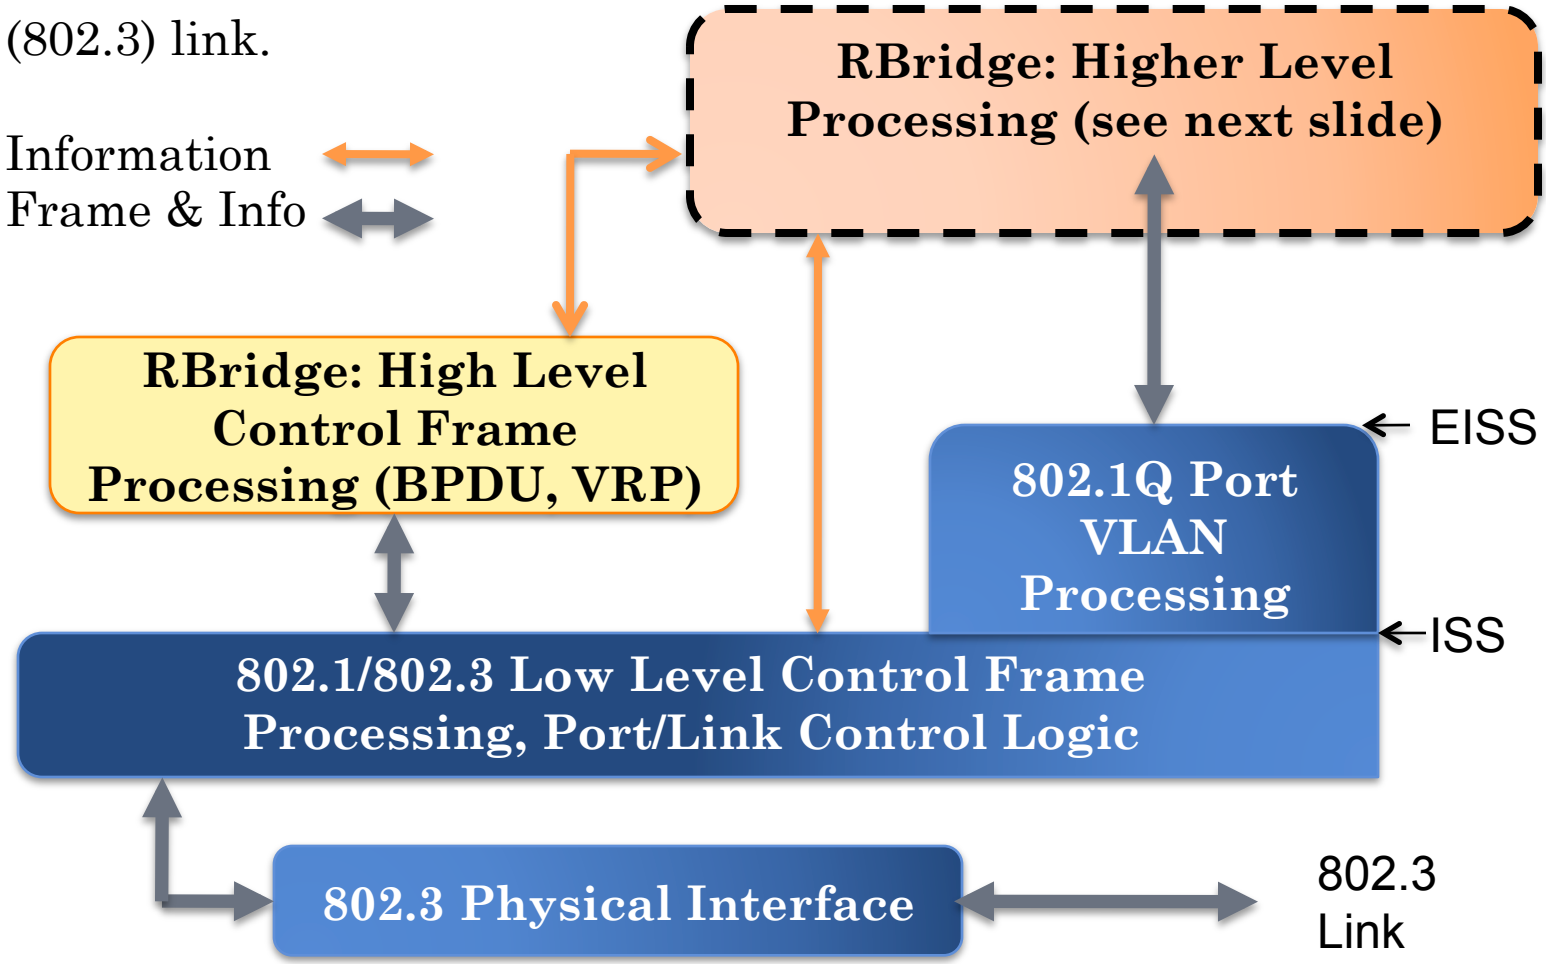

Structure Of An RBridge Port

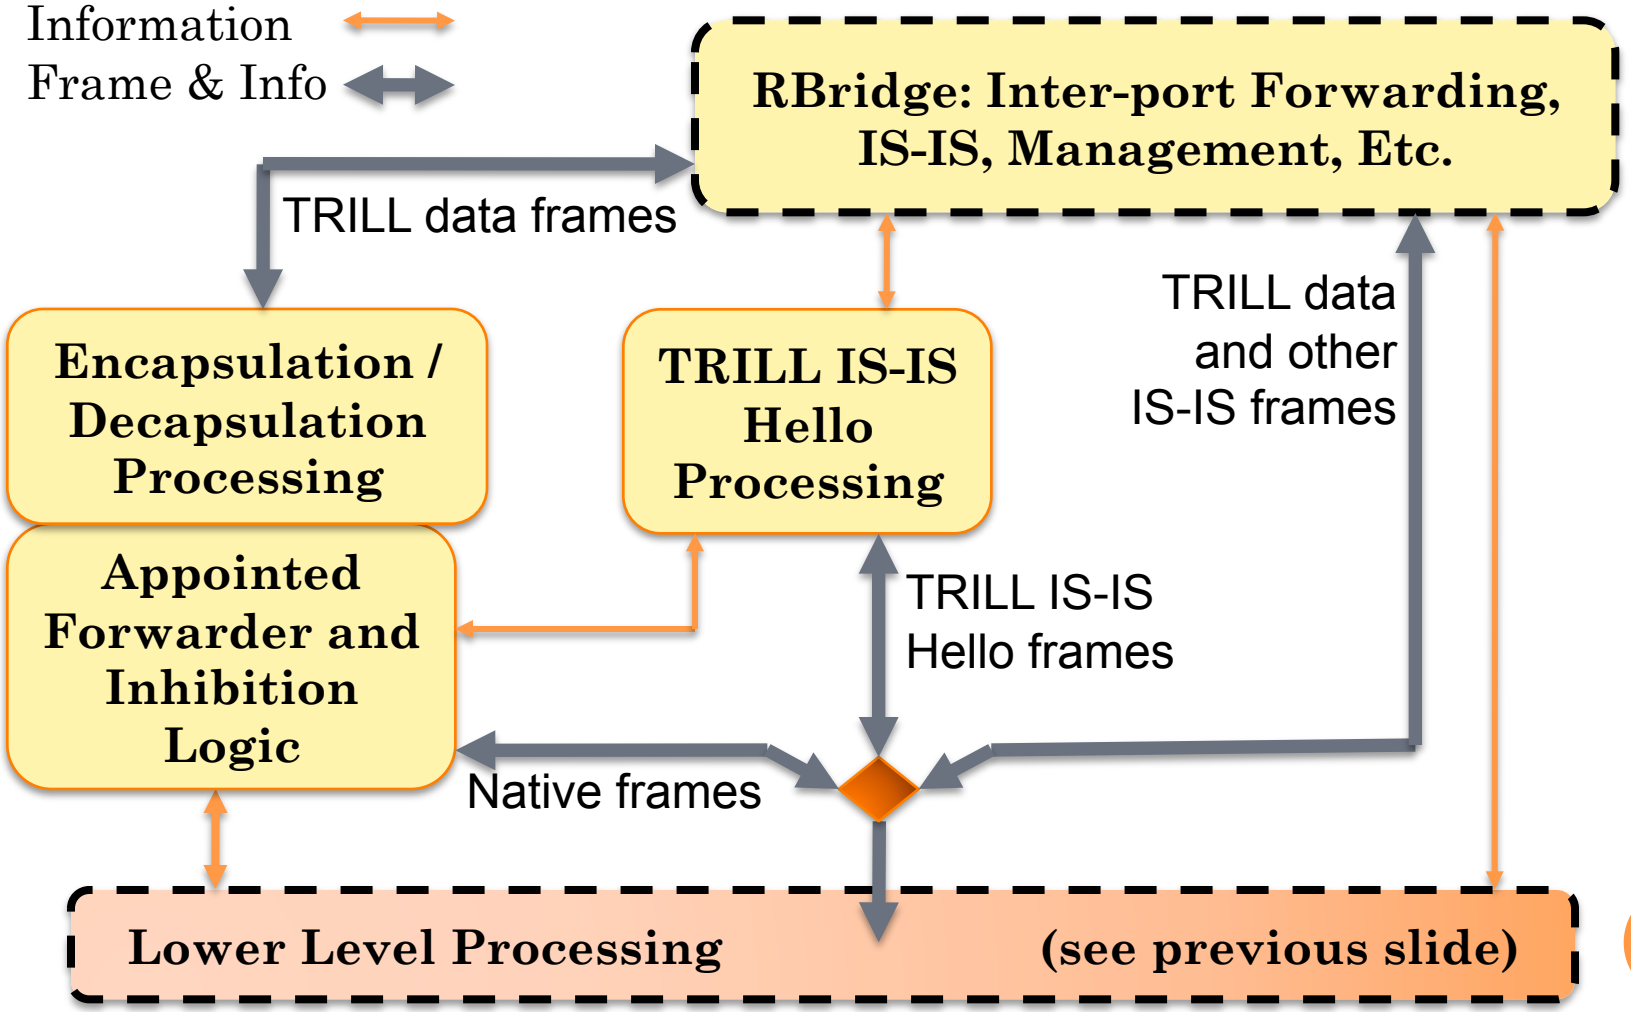

CONTENTS

- Introduction
- TRILL Frames
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL Data Requirements
  - TRILL Hello Data
  - TRILL Core LSP Data
  - TRILL ESADI Data
- References

TRILL HELLO DATA

- Additional data for TRILL is all VLAN Stuff:
  - The Designated VLAN
  - Copy of Outer VLAN ID (to detect VLAN mapping in the link)
  - Flag indicating VLAN mapping detected
  - Set of VLANs with native frame service enabled on the port
  - Flag indicating the sender thinks it is Appointed Forwarder
  - Flag indicating an access port
  - Appointed Forwarder appointments by the Designated RBridge

TRILL CORE LSP DATA (SLIDE 1 OF 2)

- TRILL versions supported
- Nickname and priority
- Distribution Tree Stuff:
  - Priority to be a distribution tree root
  - Number of trees desired
  - List of RBridges to be used as roots by this RBridge

TRILL CORE LSP DATA (SLIDE 2 OF 2)

- List of VLANs for which the RBridge is an Appointed Forwarder and, for each VLAN:
  - IPv4 and IPv6 multicast router connectivity
  - Non-IP derived multicast request flag
  - Optional: Listeners for IP-derived multicast addresses
  - Appointed Forwarder status lost counter
  - Root bridges of attached bridged LANs
  - Optional: VLAN groups
  - Optional: VLAN mappings being performed by this RBridge

TRILL ESADI DATA

- End Station Address Distribution Information: End station MAC address information and confidence

CONTENTS

- Introduction
- TRILL Frames
- Are RBridges Bridges or Routers?
- How RBridges Work
- Structure of RBridge Ports
- TRILL IS-IS Data Requirements
- References

REFERENCES

- Current Specification Draft:
  "Rbridges: Base Protocol Specification"
  - http://www.ietf.org/internet-drafts/draft-ietf-trill-rbridge-protocol-10.txt
- "RBridge VLAN Mapping"
  - http://www.ietf.org/internet-drafts/draft-perlman-trill-rbridge-vlan-mapping-00.txt
- "RBridges: Use of IS-IS"
  - http://www.ietf.org/internet-drafts/draft-eastlake-trill-rbridge-isis-02.txt
- Original Paper by Radia Perlman:
  "Rbridges: Transparent Routing"
  - http://www.postel.org/rbridge/infocom04-paper.pdf

END

Donald Eastlake 3rd

donald.eastlake@stellarswitches.com, +1-508-333-2270

Stellar Switches