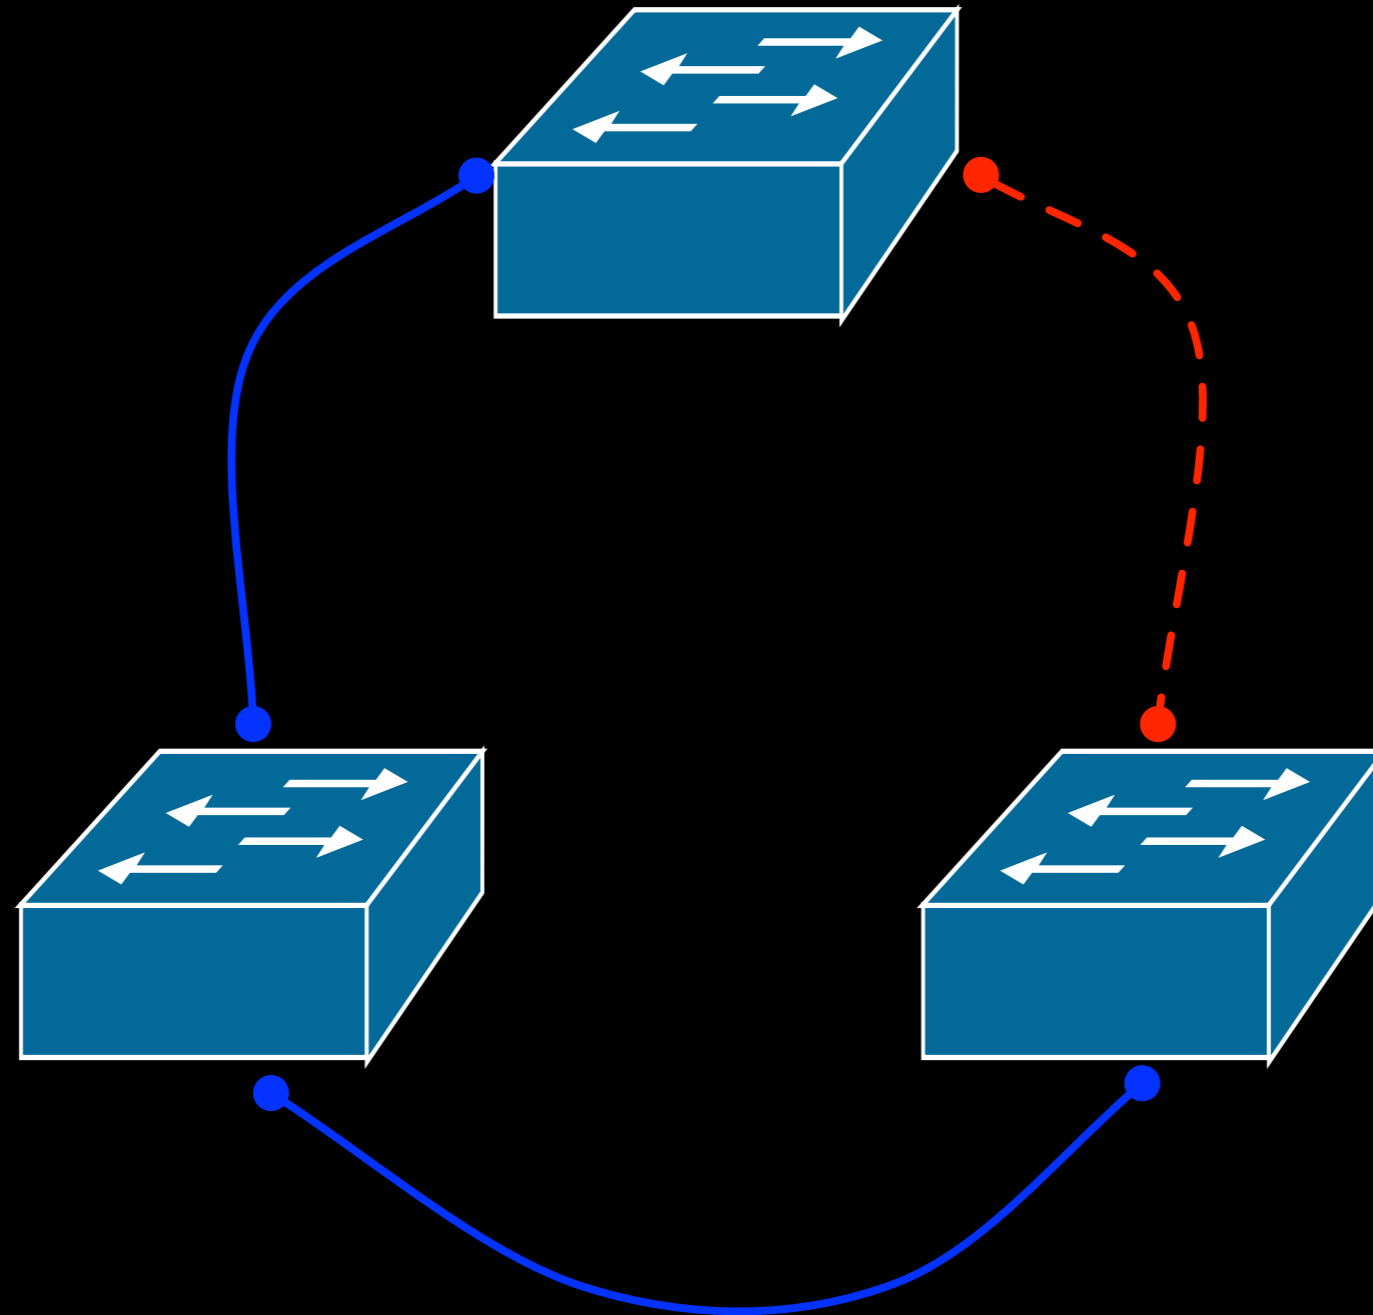



TRILLing new features for your IXP

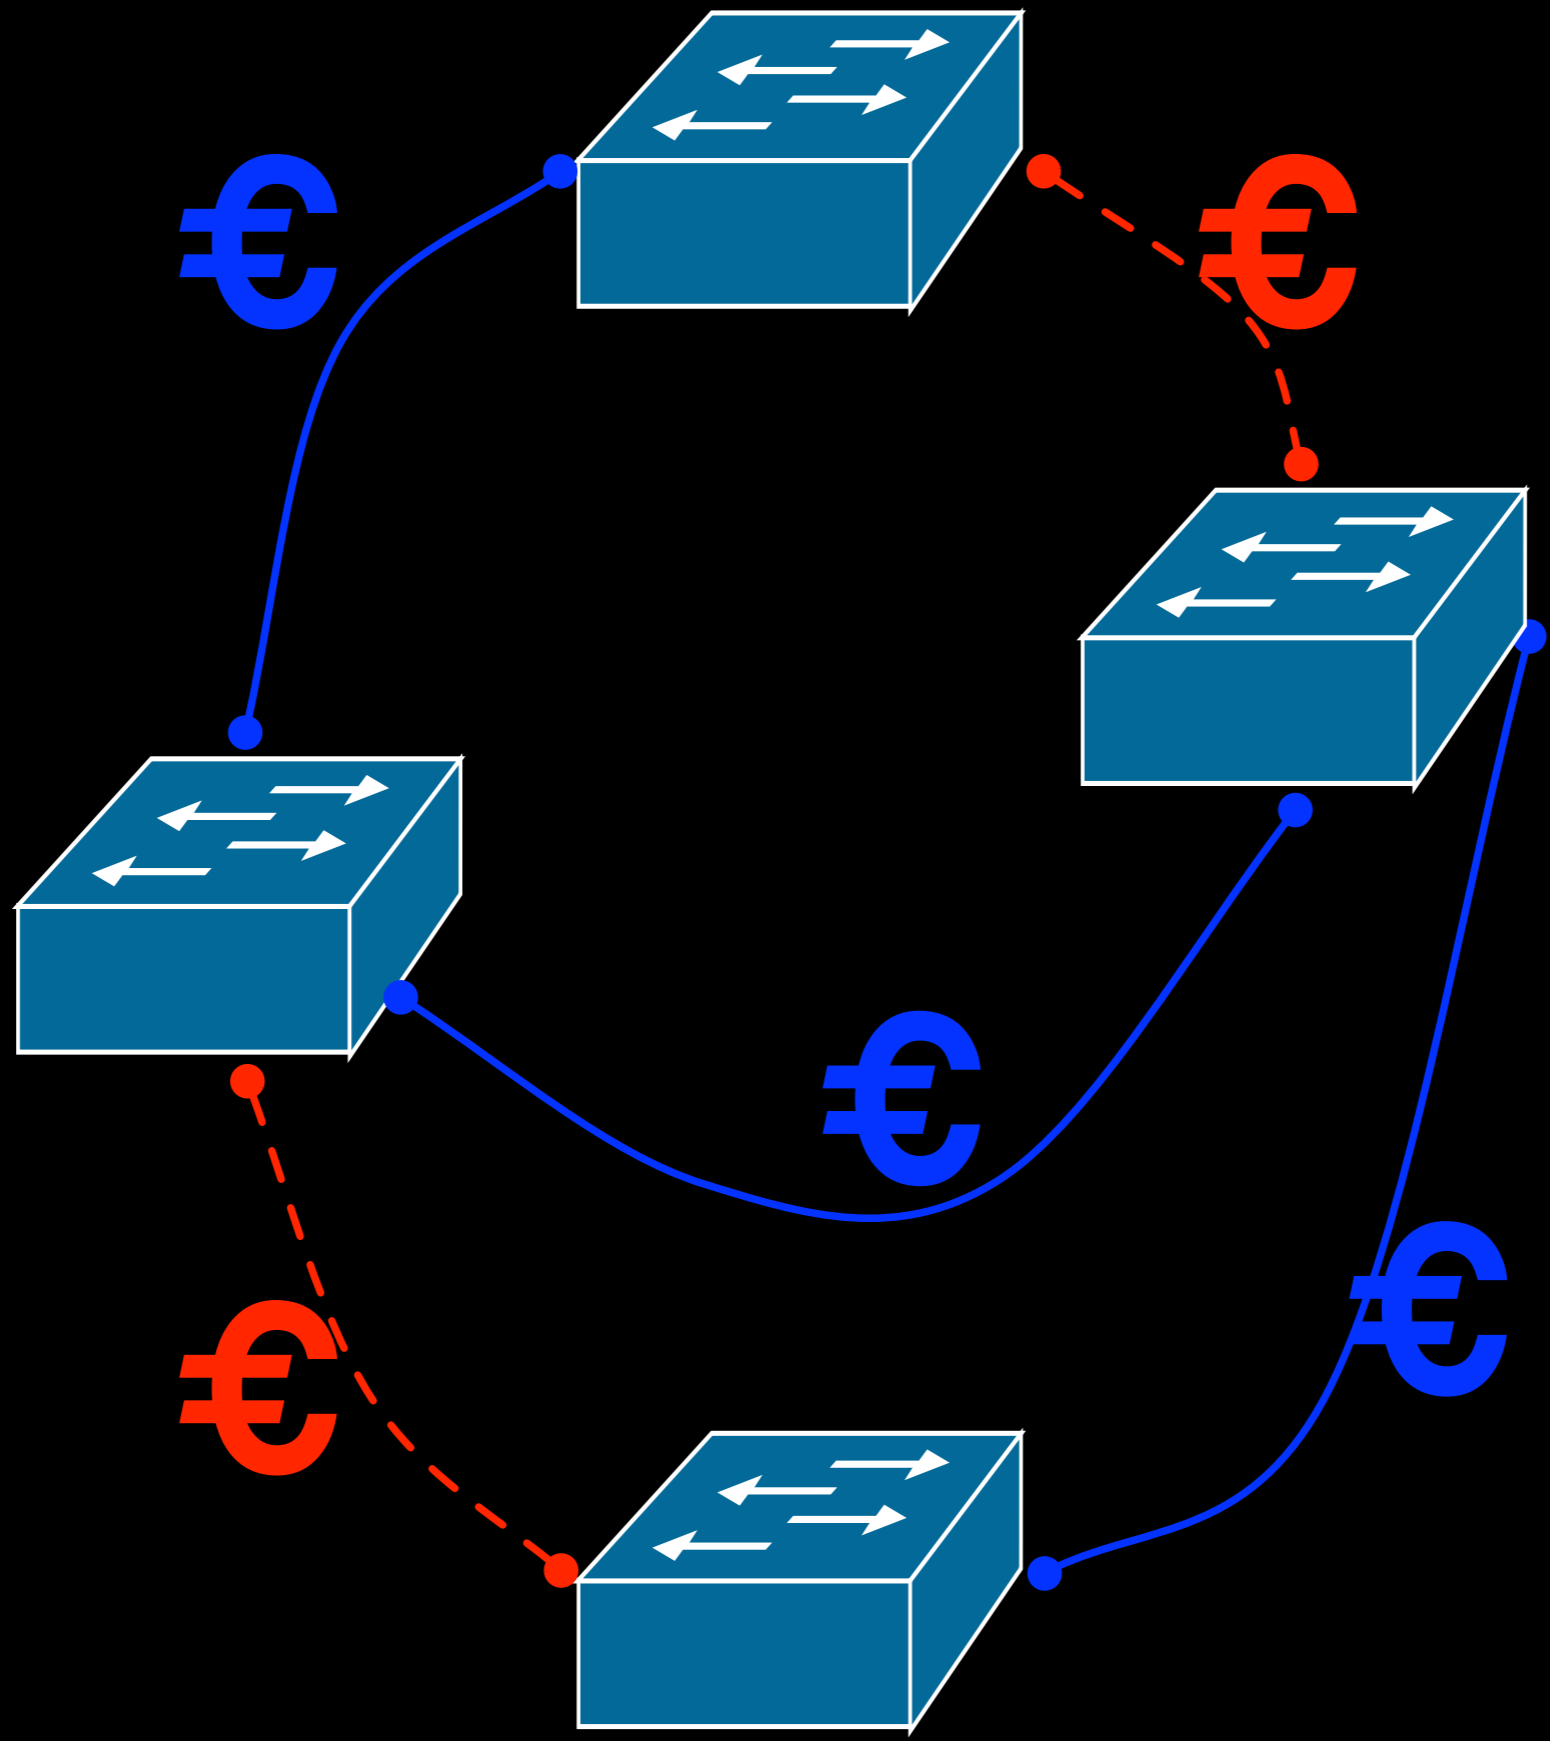
Andy Davidson
RIPE60, Prague, CZ


LONAP
6th May 2010

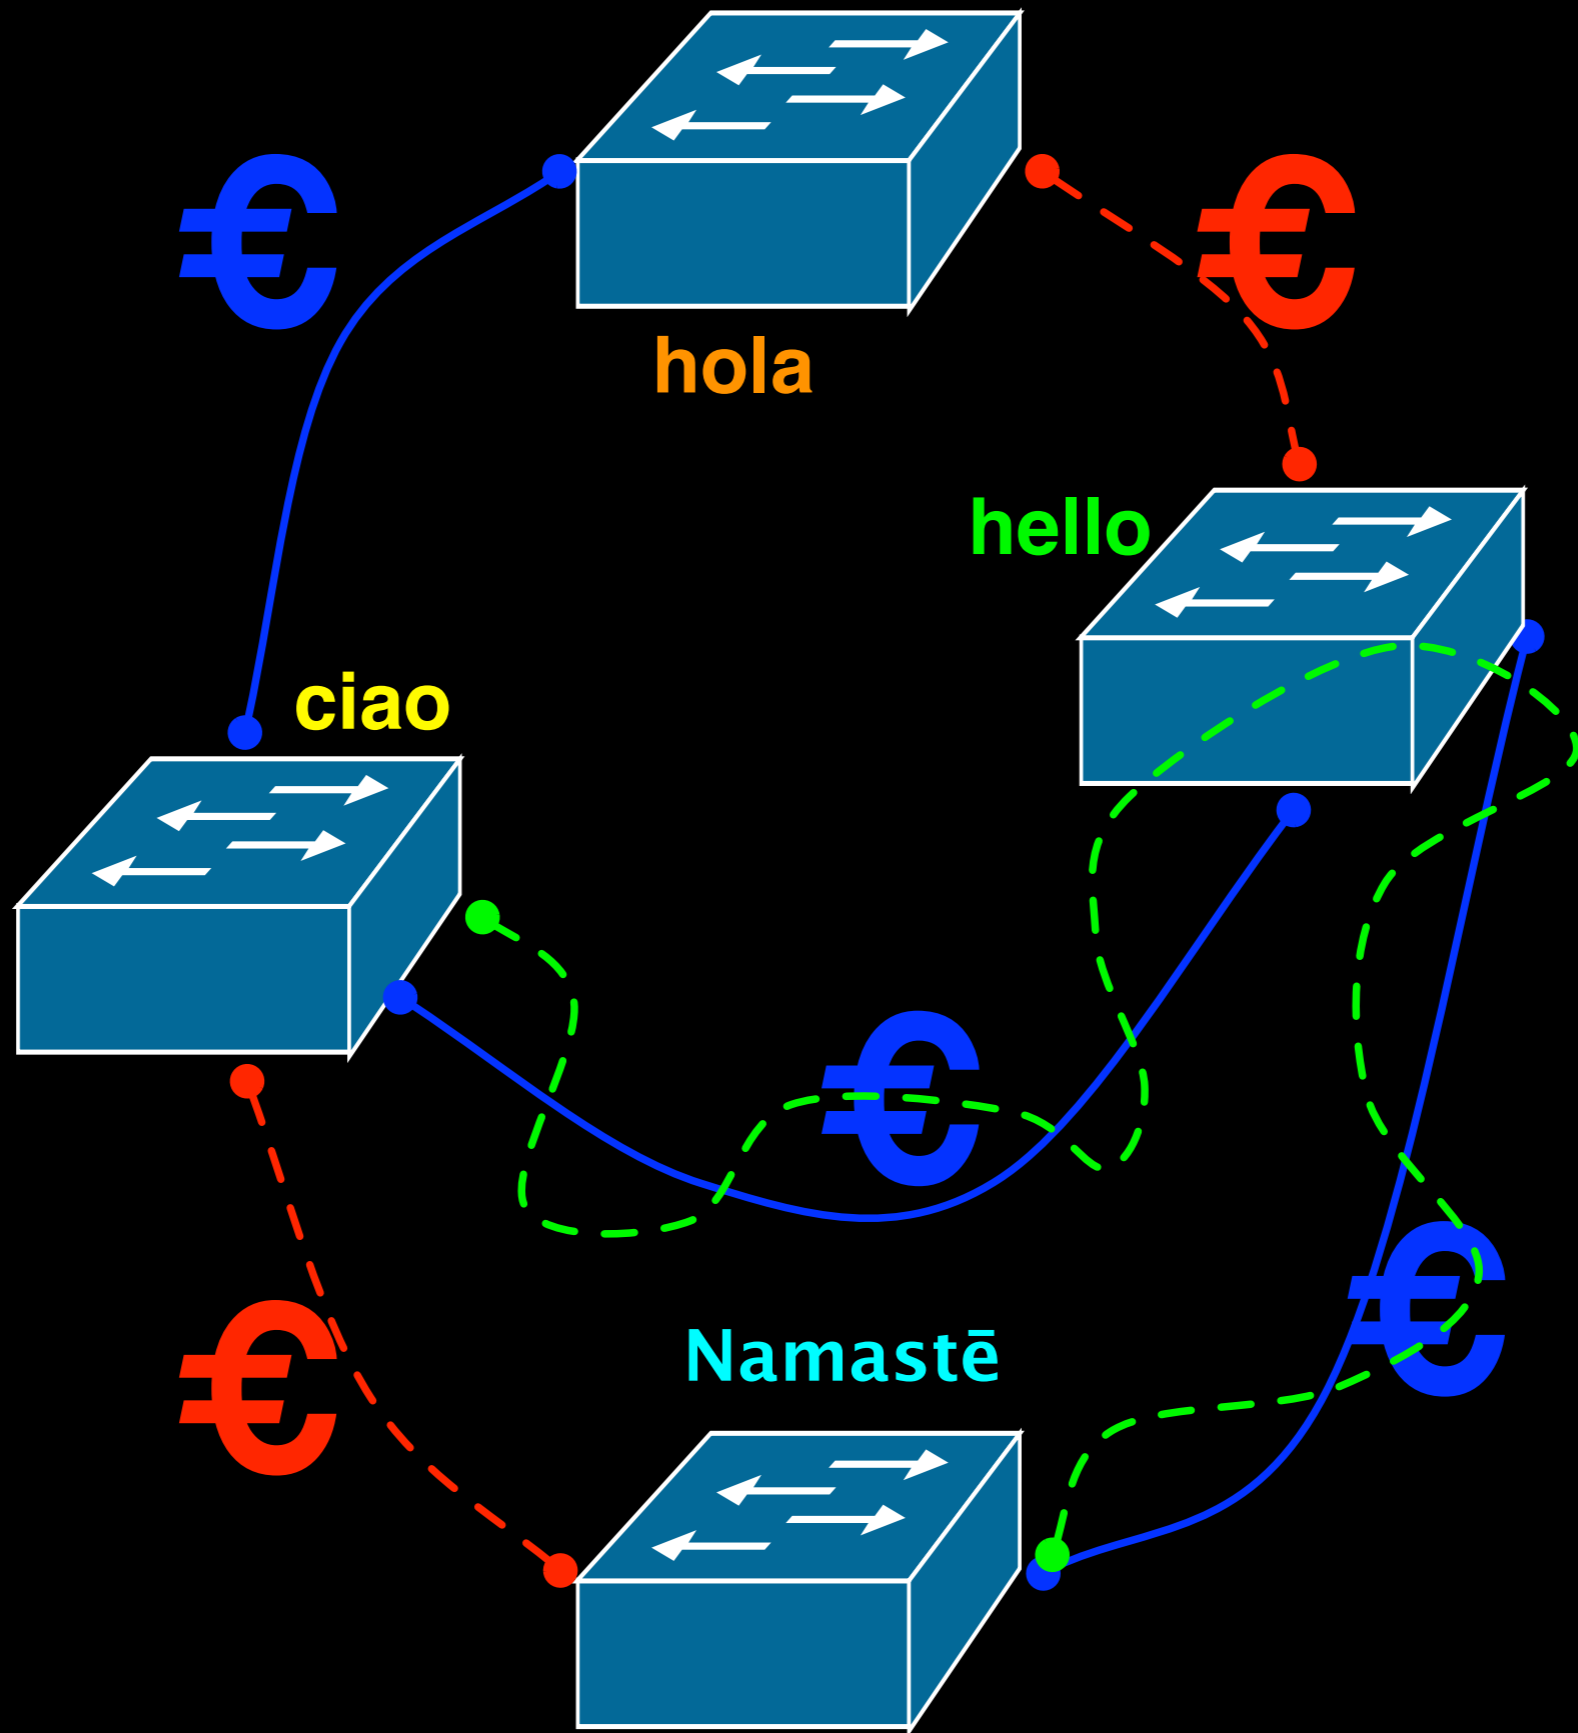
Typical layer 2 segment

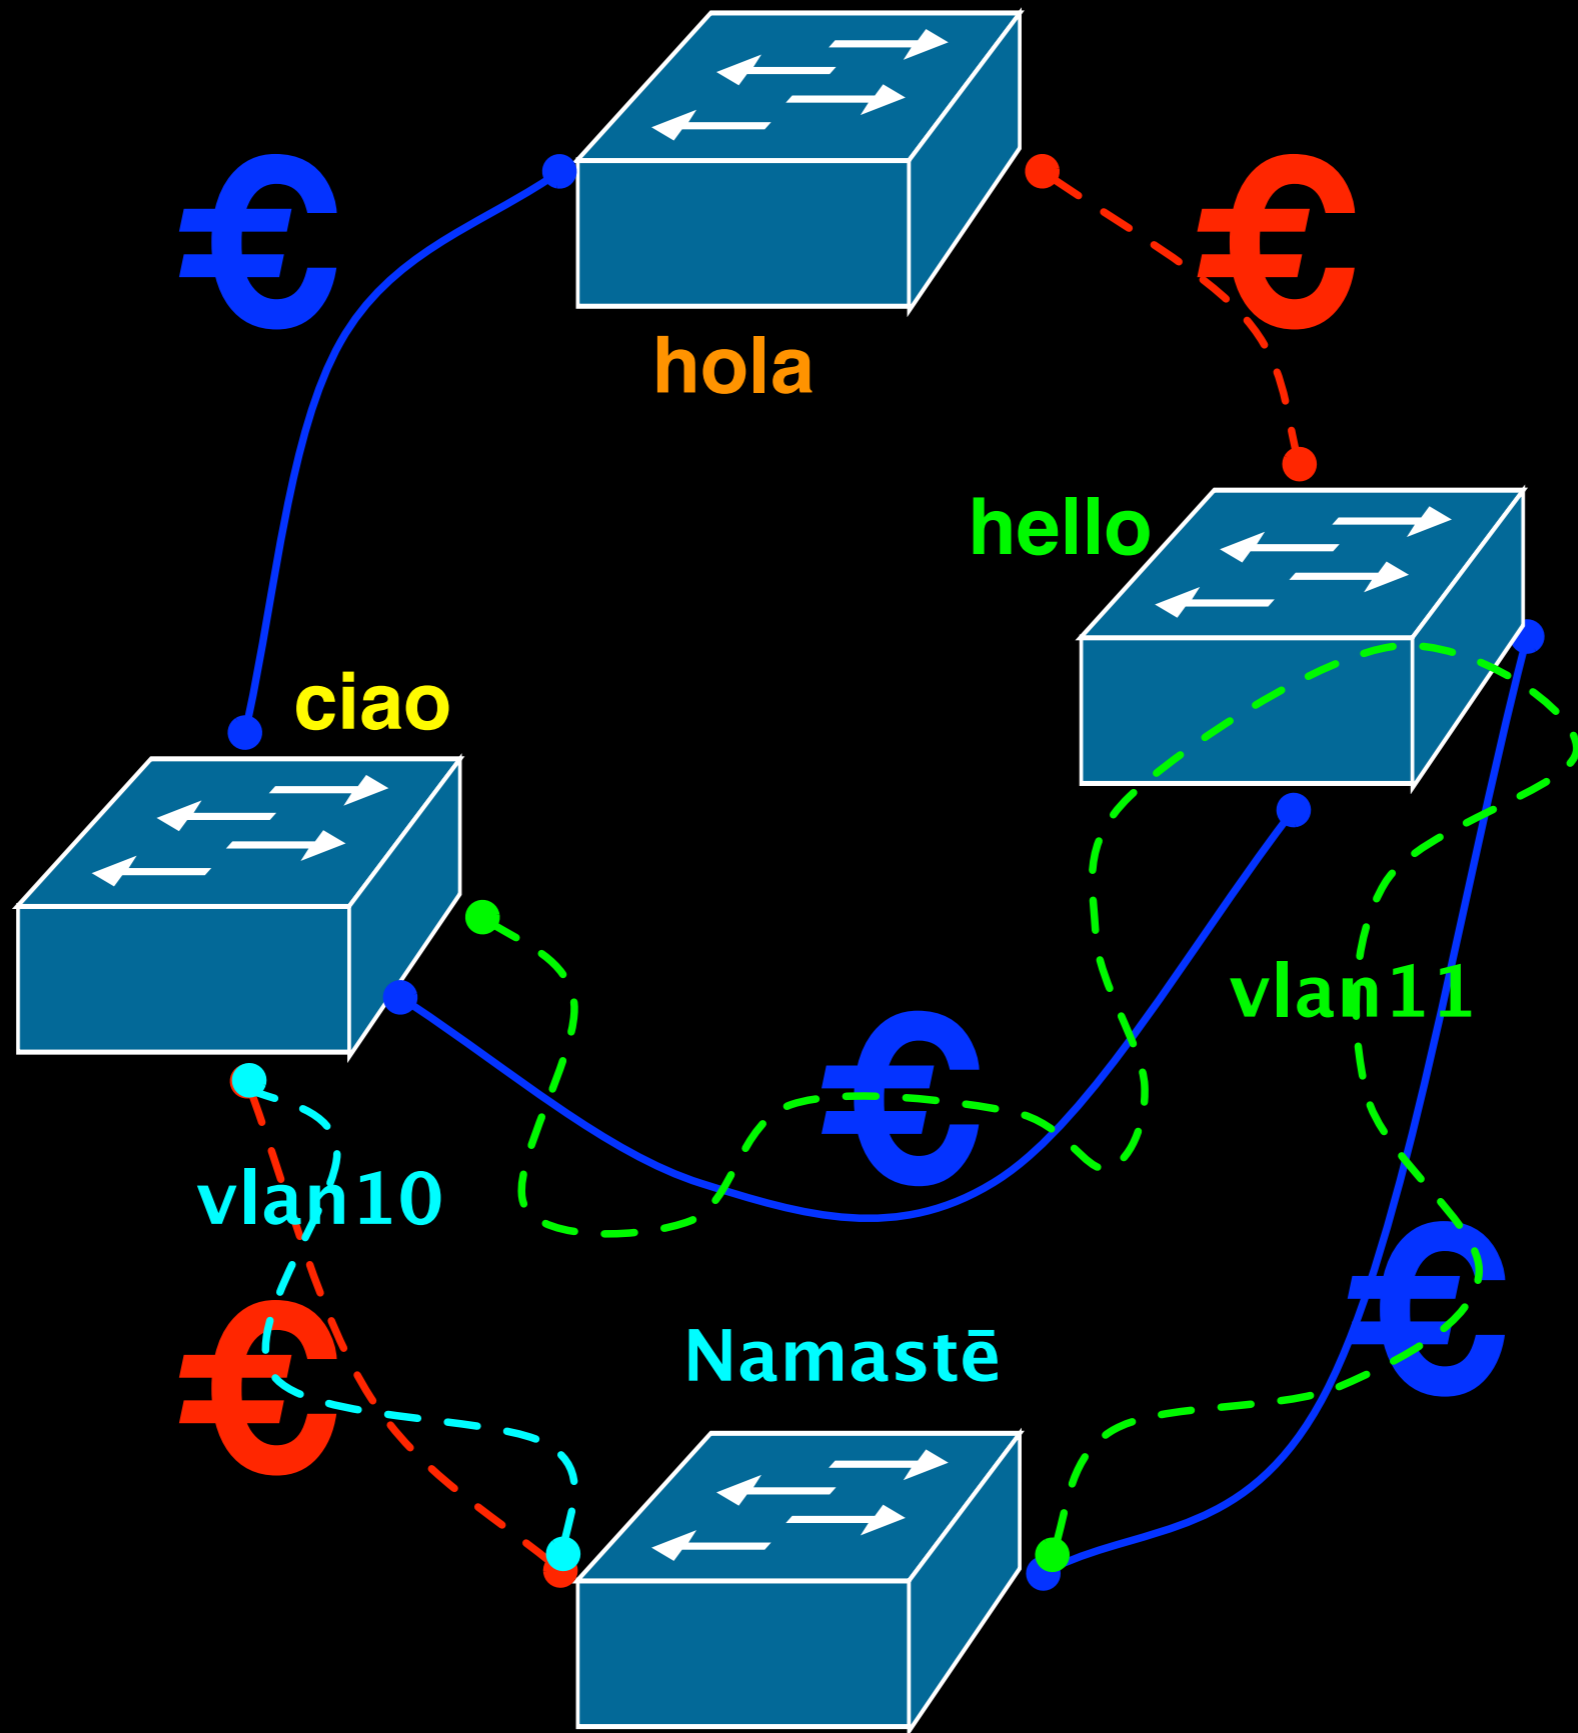


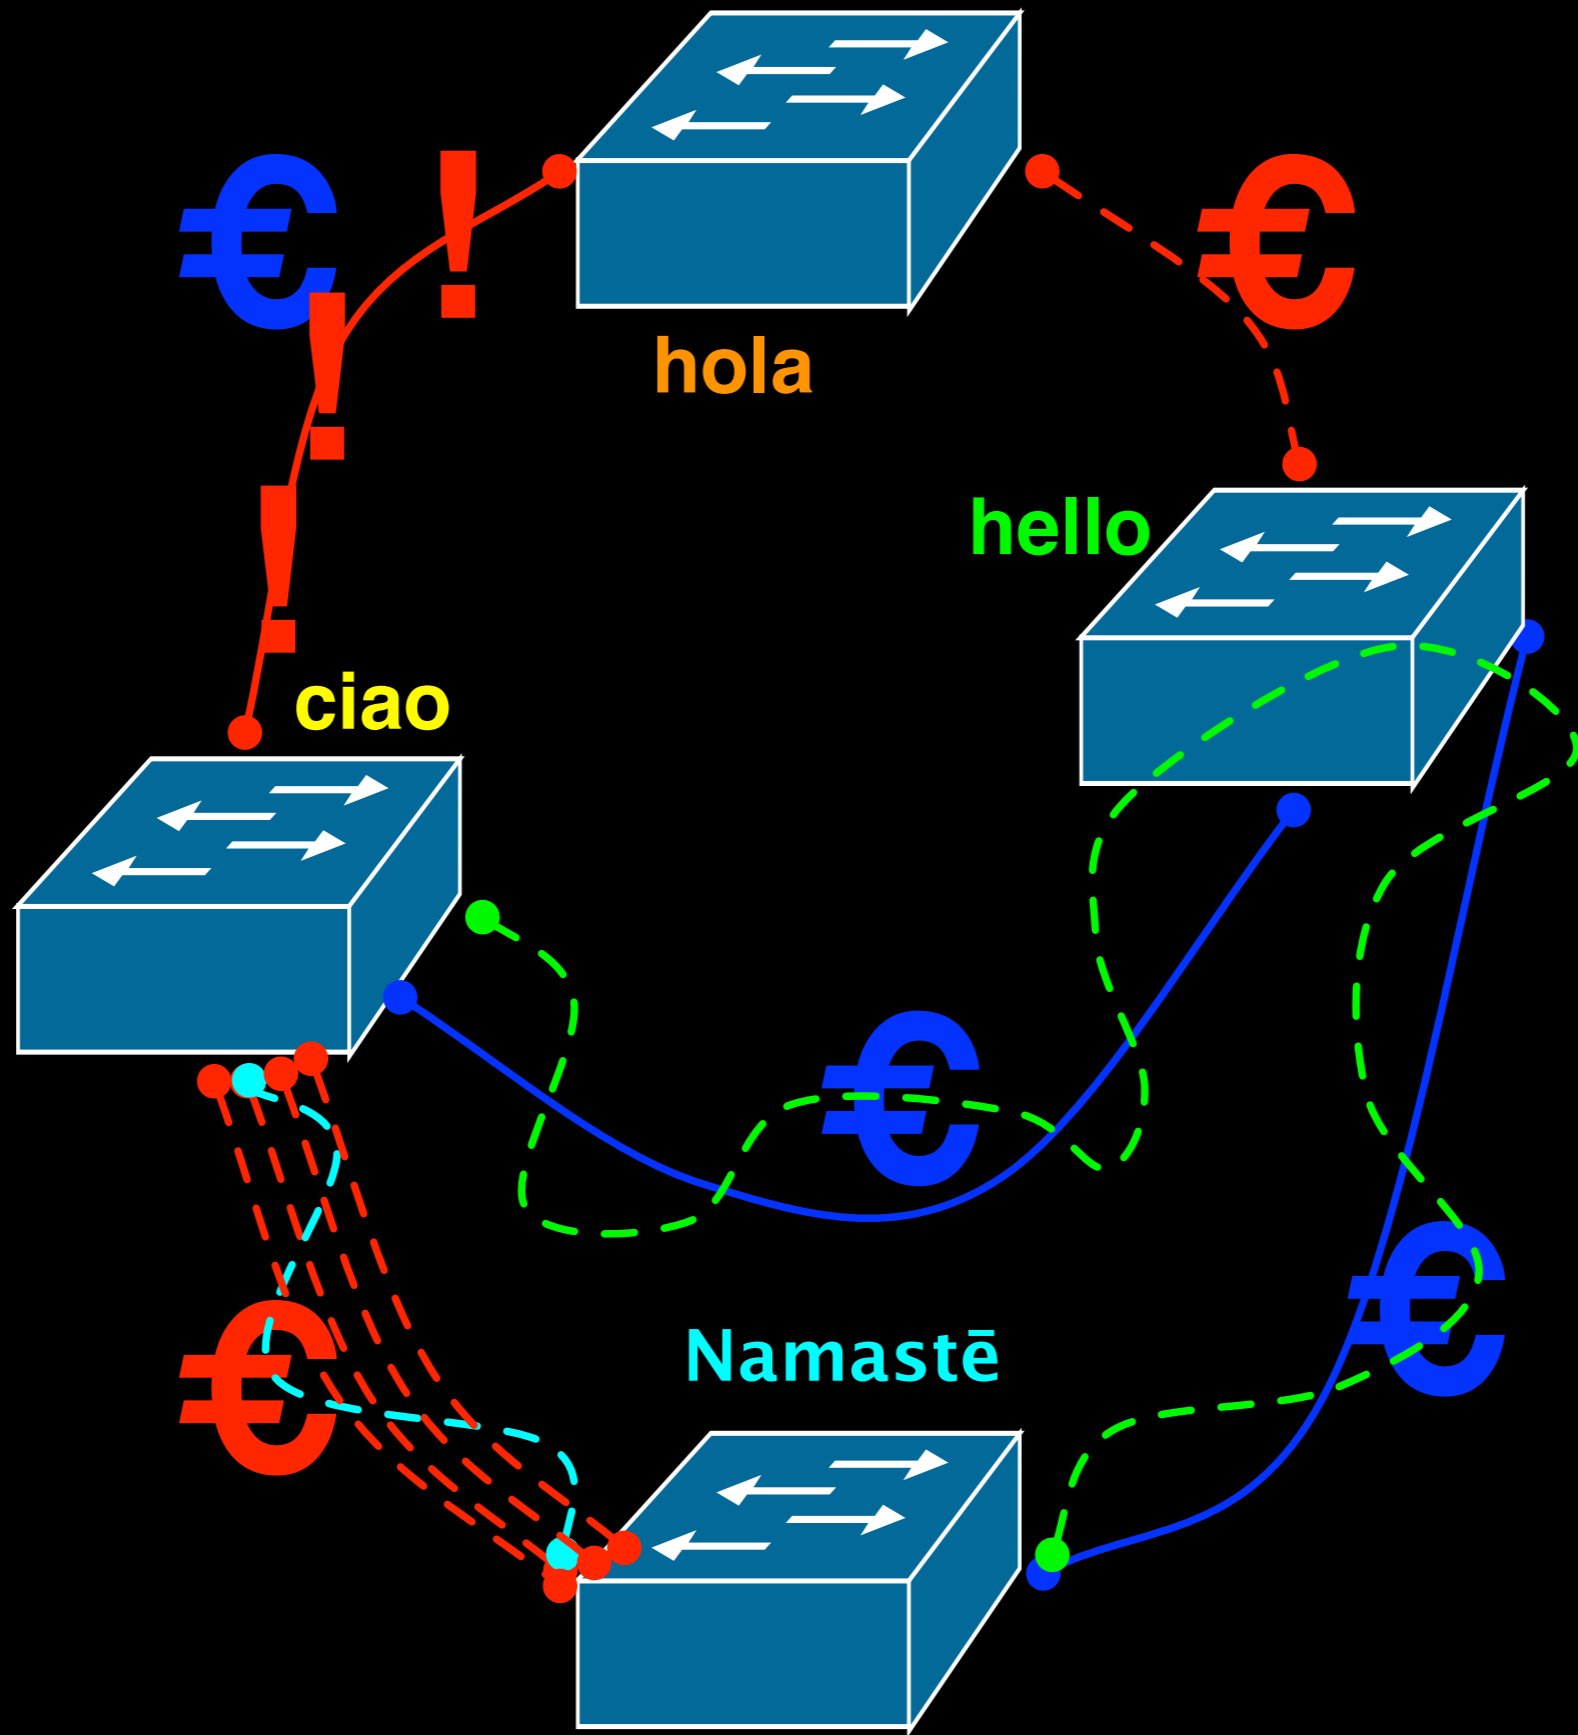












There are a number of features of modern layer 3 routing protocols which would be beneficial if available at layer 2

RFC5556

J Touch (USC), R Pearlman (Sun)

With packet hop counts we now see,
The network need not be loop-free!

draft-ietf-trill-rbridge-protocol-16

Ray Perlner, Algorhyme v2

Magic here:

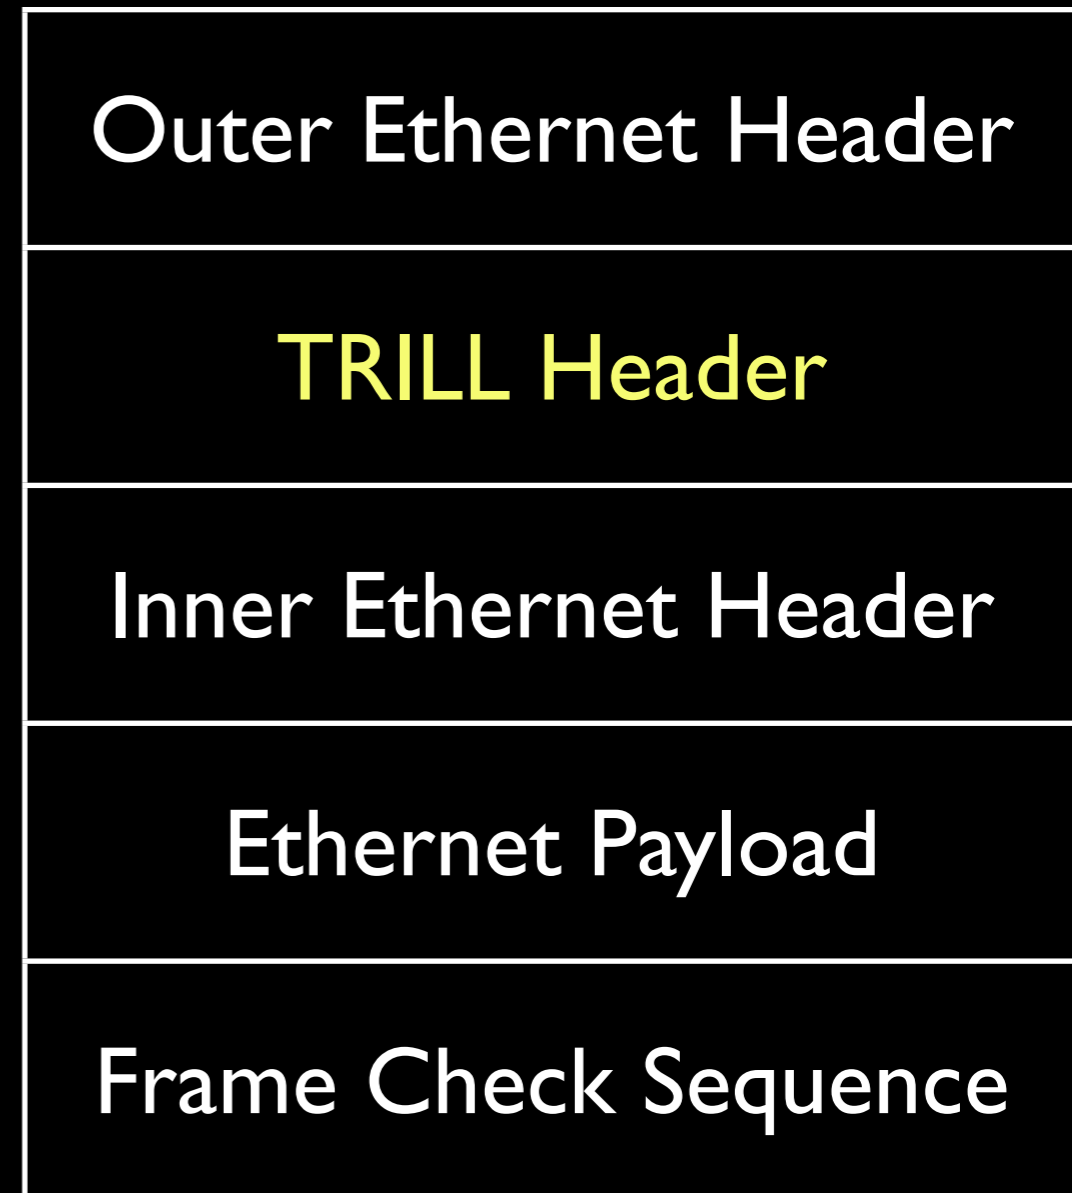
- Link State protocols run between switches (Routing Bridges)
- All bridges know all other bridges
- **Optimal paths** converged for unicast destinations
- **Loopless distribution** trees configured for unknown or multicast destinations

Link State Protocol


- TRILL-IS-IS
- Election for Designated RBridge per LINK
- DR assigns an appointed forwarder role to one RBridge for each VLAN on the link
- The forwarder listens and learns all end-node MAC addresses, ports and attached RBridge names on Layer 2 link

Frame Encapsulation

- Frames on Inter-switch links encapsulated with a header containing **hop count** and RBridge **nickname** (to mitigate reconvergence loops)
- Also details **exit RBridge** name to calculate frame forwarding path



Known Unicast Forwarding




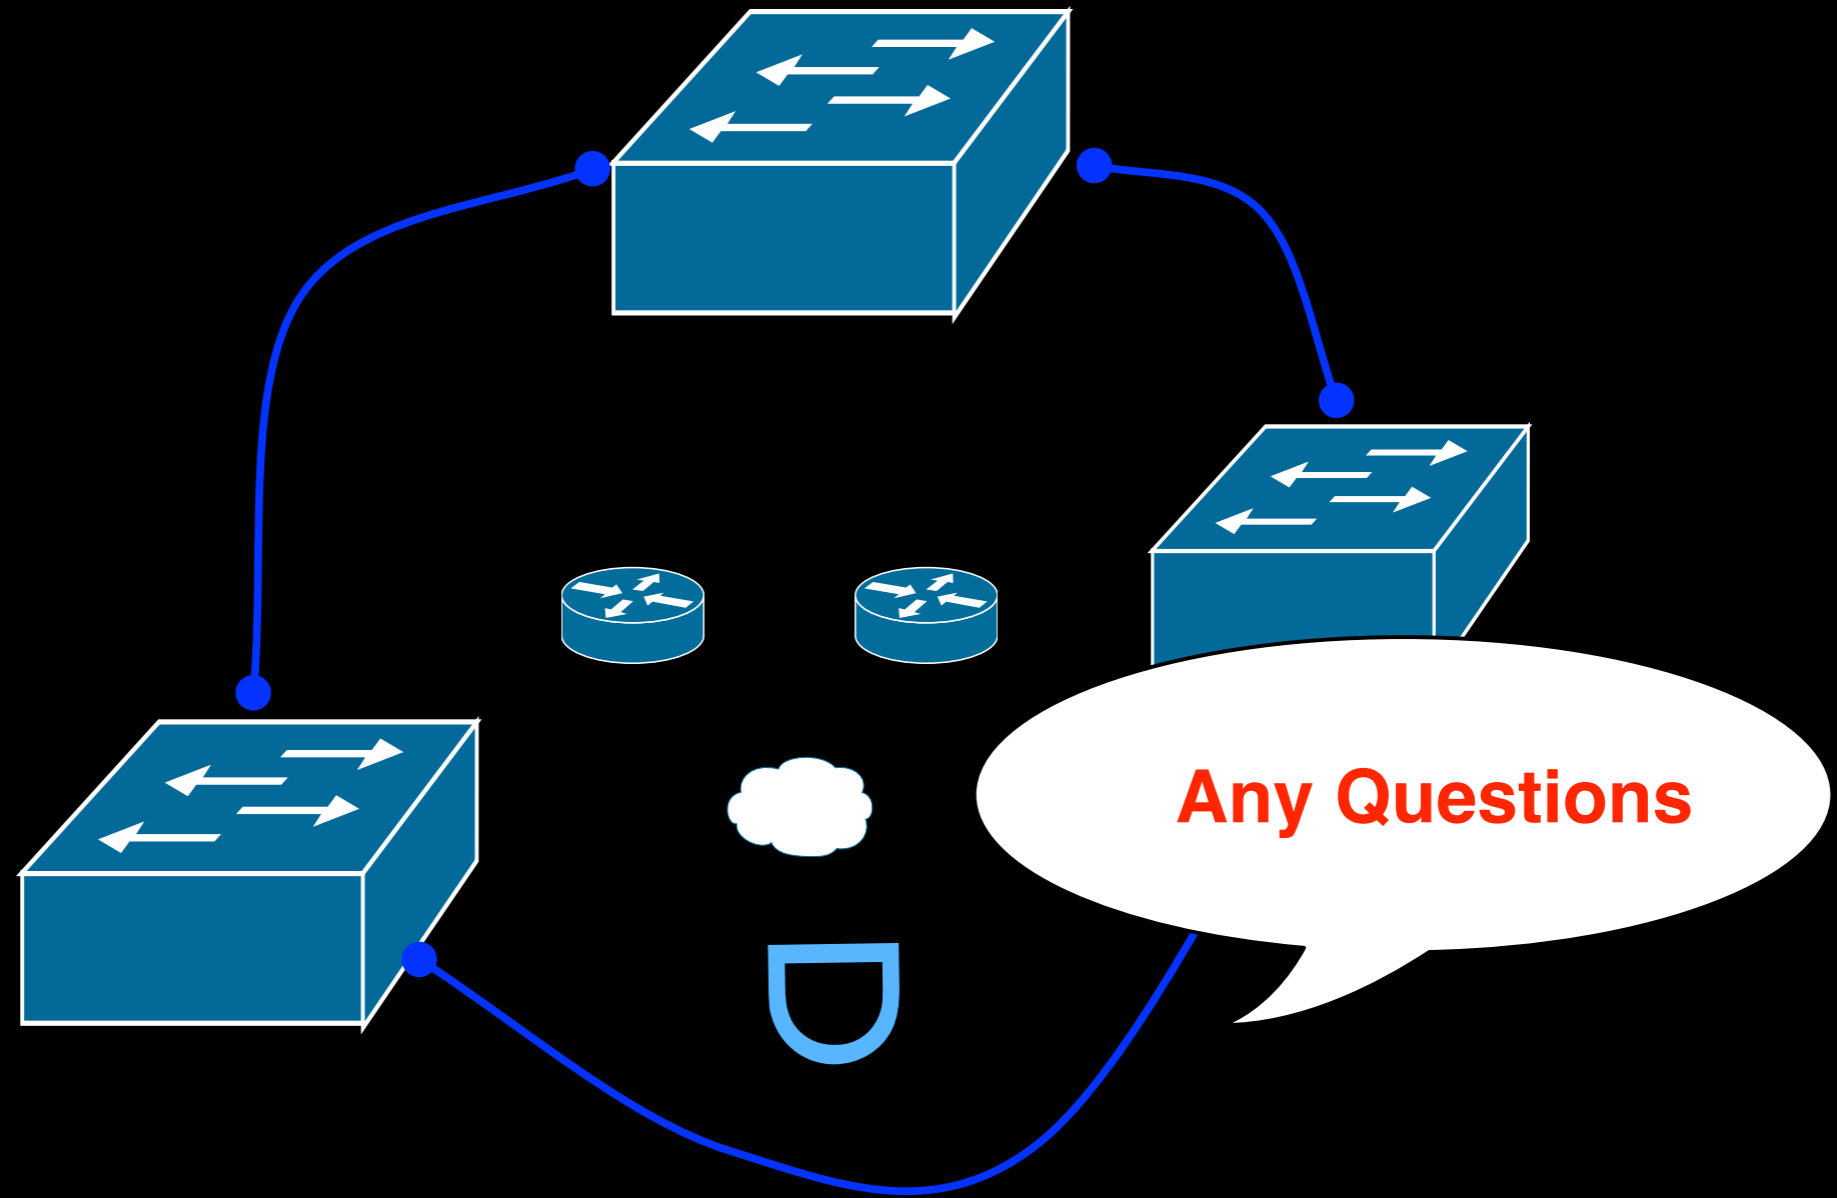
- Defined by the presence of an encapsulated frame with known unicast destination MAC
- First RBridge encapsulates the frame with a TRILL header, that identifies the exit RBridge
- Forwarded “hop by hop”

Multi-destination

- Broadcast, multicast, or unknown unicast MACs.
- One or more **bidirectional trees** calculated and nicknamed
- **SPF** calculation, not Spanning Tree
- Links with no downstream nodes are **pruned**.
- Forwarding, generally, is handled by delivering frames to adjacent, downstream RBridges, according to the tree nickname specified in the TRILL header.

Benefits

- Shorter Layer 2 paths with meshing
- Therefore improved latency
- **RoI** from resilient links
- No full  loop prevention protocol faults
- Multipath forwarding to handle increased unexpected traffic



Join in

- <http://datatracker.ietf.org/wg/trill/>
- <http://mailman.postel.org/mailman/listinfo/rbridge>