TRILL Deployment Nears

August 3rd – 5th 2010

Introduction

An exciting milestone on the road to TRILL deployment occurred from August 3rd to 5th 2010 at the University of New Hampshire InterOperability Laboratory (UNH-IOL), as a group of companies attended the first multi-vendor TRILL plugfest. The plugfest was successful in demonstrating interoperability between TRILL devices. It was a first step towards wide deployment of TRILL and valuable knowledge for increasing TRILL interoperability was gained through the event. The UNH-IOL is excited to be hosting another TRILL plugfest in Q1 of 2011.

This whitepaper discusses the TRILL protocol to help inform potential TRILL equipment customers about the benefits of adoption of TRILL in their networks along with describing what this plugfest means for those looking towards the deployment of TRILL.

The Need for a Replacement for Spanning Tree

In 2002, the Beth Israel Deaconess Medical Center in Boston, MA had a massive system failure caused by a Spanning Tree Protocol loop (1). This led Radia Perlman, the inventor of Spanning Tree Protocol (2), and others to come up with the concept of a “Routing Bridge” (RBridge) (3) and the protocol TRansparent Interconnection of Lots of Links (TRILL). Most of TRILL, the intended replacement for Spanning Tree Protocol, has been approved as a standard and the last part recently went into working group last call, meaning full standardization is nearly complete.

For now, let’s see how Spanning Tree contributed to “one of the worst health-care IT disasters ever.” [i] The first sign that something was wrong was when it started to take several seconds to send and receive e-mail at the medical center. This sluggishness was originally believed to be a normal response to a sudden increase in traffic. Previously, such spikes had worked themselves out, taking from several minutes to a few hours. The IT team at Beth Israel began shutting down switches in an attempt to discover the source of the problem. This exposed the first of Spanning Tree's shortcomings. Each time a switch was shut down, the other switches in the network needed to re-converge. In a network that was already struggling, the recalculations were enough to take the network down completely. Realizing their mistake, they quickly turned all of the switches back on. The network struggled into the evening, and once most users were gone, the IT staff were able to discover the cause: a Spanning Tree Protocol loop. Spanning Tree has a hop limit of 7 switches; if data needs to take more than 7 hops to go from host to host, a loop can be created, flooding the network and causing other resource problems as well. The IT team at Beth Israel cut off redundancy in the network and considered the problem solved.[ii]

Within the first hour of increased network utilization the next morning, it became clear that the problem had not been resolved. The network at Beth Israel supported connections to networks at several other hospitals in the area. Rather than sever those links, the team decided to convert inter-hospital traffic to IP routing instead of bridging. This conversion attempt took all day, and was, in the end, unsuccessful.[iii]

The health-care staff at Beth Israel were facing real problems. They had critical decisions to make regarding their patients health crises, some of which were life-threatening, and had no access to lab reports and other data needed to make the right decisions. That afternoon, Beth Israel closed its ER and called in its equipment vendor to help solve the problem.[iv]

Soon, a second Spanning Tree loop was discovered. The network at Beth Israel was built entirely using Layer 2 switches, with no Layer 3 Routing. Layer 3 Routing is considerably smarter, but more expensive. Most networks at that time incorporated at least some routing capabilities. The team and its vendor installed some routers, removed a few switches, rebooted the system, and believed the problem was solved.[v]

The next morning, the network started experiencing trouble again, and the executive team at the hospital decided to shut down the network altogether, moving back 2 and a half decades and using paper to handle all of its data and communication needs. This move backward created an efficiency and logistics nightmare. It became clear that Beth Israel had neglected the backbone of its information system, dealing more with applications and usage, rather than the network itself, the glue that holds it all together.[vi]

Two more Spanning Tree loops were discovered. More routers were added to the network, and increased redundancy was built in as well. The network was turned back on, and after a few spikes and a few minor fixes, everything was back online and running smoothly.[vii] The specific problem at Beth Israel Deaconess Medical Center was solved, but Spanning Tree's shortcomings were exposed, and would remain.

Benefits of TRILL

It is outside the scope of this paper to discuss in detail the differences between using Layer 2 switching and Layer 3 routing in an enterprise network. However, switching has some well-known benefits, arguably, when compared to routing, which include cost, efficiency, and simplicity. For those networks, and network implementers, that decide to use Layer 2 switching, TRILL offers many improvements over Spanning Tree, making the continued use of switching more feasible and less likely to experience the types of failures encountered with Spanning Tree.

Here are some benefits that users of TRILL will enjoy over users of Spanning Tree:

- The introduction of shortest path routing to Layer 2 to eliminate inefficient paths. This is in contrast to STP whose active topology disables forwarding on many redundant links by reducing the physical topology to a single spanning tree (or a few spanning trees).
- The ability to use multipath forwarding to spread traffic out among the available paths and decrease traffic congestion. STP in contrast sends all traffic along the spanning tree meaning that all traffic from one device to another will take the same path.
- Decreased convergence time by use of a link state protocol versus spanning tree's distance vector algorithm.
- Safety in forwarding. STP fails in the un-safe direction: defaulting to flooding traffic out on all ports. In contrast, when TRILL fails in the safe direction by defaulting to dropping traffic. This means it only temporarily blocks network connectivity and does not cause a loop in network traffic. TRILL furthermore has a hop count that provides an additional guarantee of safety.

For companies looking into adopting TRILL, these points mean TRILL provides greater bisectional bandwidth in layer 2 networks. This in turns mean more efficient utilization of network infrastructure and a decreased cost to benefit ratio. These benefits are particularly important for companies deploying TRILL in data centers that in turn support cloud computing, a rapidly growing market.

TRILL is more stable than Spanning Tree because it provides greater safety against loops and decreased recovery time in a network when an RBridge fails. This benefit would have prevented the 2002 Beth Israel Deaconess Medical Center failure. TRILL provides protection from just such an event.

Standardization

Since 2005 the IETF's TRILL Working Group has been working on the solution known as TRILL. Earlier this year the IESG approved the TRILL working group's standardization of the RBridge concept as a proposed standard. The draft (4)is currently in the RFC editor's queue waiting for the IESG to receive another draft (5)that specifies some Intermediate System To Intermediate System (IS-IS) frame formats used by TRILL.

Plugfest

The TRILL Plugfest was hosted at the University of New Hampshire's InterOperability Laboratory in Durham NH and was attended by Broadcom, JDSU, and Oracle. For companies looking to adopt TRILL, this is an exciting sign saying there may be deployments of TRILL in the near future. However, as with any new technology, early adopters may have some concerns. Will TRILL truly provide the benefits it advertises? Will the implementations be mature enough for a high availability environment? And will buying an early implementation from one company lock one into buying all future equipment from that same company or will implementations between companies be interoperable?

Conclusion

A plugfest is an event that answers these questions. The companies that attended the event in August brought their implementations of TRILL and ran them against each other using the UNH-IOL's TRILL Interoperability Test Suite (6) as the basis for measuring interoperability. By TRILL implementations demonstrating interoperability at a plugfest, early adopters of TRILL can be confident TRILL is ready for deployment. Plans are under way for another plugfest in Q1 of 2011.

References

1. **Berinato, Scott.** All Systems Down. *CIO Magazine.* April 11, 2003.

2. 802.1Q-2005 - Virtual LANs. s.l. : IEEE, December 5, 2005.

3. *RBridges: Transparent Routing.* **Perlman, Radia.** March 2004. Proc. Infocom 2005.

4. **Perlman, Radia, et al.** Rbridges: Base Protocol Specification, Work In Progress. s.l. : IETF, March 3, 2010.

5. **Eastlake, Donald, et al.** TRILL Use of IS-IS. s.l. : IETF, August 23, 2010.

6. **Bond, David.** TRILL Interoperability Test Suite. Durham NH : University of New Hampshire InterOperability Laboratory, March 16, 2010.

7. **Touch, J. and Perlman, Radia.** Transparent Interconnection of Lots of Links (TRILL): Problem and Applicability Statement. *RFC 5556.* s.l. : IETF, May 2009.

[i] **Berinato, Scott**. All Systems Down. *CIO Magazine*. April 11, 2003
[ii] **ibid**
[iii] Ibid

[iv] ibid
[v] ibid
[vi] ibid
[vii] ibid