BROCADE DATA CENTER TECHNOLOGY LEADERSHIP

DATA CENTER

ISL Trunking for Brocade VCS

HIGHLIGHTS

- Combines up to eight Inter-Switch Links (ISLs) into a single logical trunk that provides up to 80 Gigabits per second (Gbps) data transfers (with 10 Gbps solutions)

- Optimizes link usage by evenly distributing traffic across all ISLs at the frame level in hardware

- Maintains in-order delivery to ensure data reliability

- Helps ensure Ethernet fabric reliability and availability even if a link in the trunk fails

- Simplifies Ethernet fabric management by reducing the number of ISLs required

INTRODUCTION

The Brocade® VCS™ Ethernet fabrics utilize hardware-based ISL Trunking to provide high-performance, high-throughput Inter-Switch Links (ISLs), to simplify interconnect configuration, to provide automatic link failover with no interruption of traffic on unaffected links, and to provide plug-and-play fabric scalability. Instead of manually configuring Link Aggregation Groups (LAGs) on individual ports on multiple switches, Brocade VCS Ethernet fabrics automatically form trunks when multiple ISL connections are added between switches.

ISL TRUNKING AND ETHERNET FABRIC NETWORKS

Classic Ethernet networks are hierarchical with three or more tiers creating congestion on ISLs. Spanning Tree Protocol (STP) allows only one ISL between any two switches. This means that ISL bandwidth is limited to a single connection, since multiple paths between switches are prohibited.

LAGs enable multiple links between switches to be treated as a single connection without forming loops. But a LAG must be manually configured on each port in the LAG, is not very flexible, and does not provide guarantees that bandwidth will be handled evenly across ISL links, shown in Figure 1.

Ethernet fabrics prevent loops without using STP. Flatter networks built using Ethernet fabric technologies require self-forming and multipathing connections between end points, eliminating manual configuration of LAG ports while providing non-disruptive, scalable bandwidth within the fabric. Ethernet fabrics support any network topology (tree, ring, mesh, or core/edge) and avoid bottlenecks on paths as traffic volume grows, since traffic flows are load balanced between multiple shortest paths.

BROCADE ISL TRUNKING

Brocade ISL Trunking is a built-in capability available for all Brocade VDX™ Data Center Switches. This ASIC-based technology is ideal for optimizing performance and simplifying the management of multi-switch Brocade VCS-based Ethernet fabrics.

ISL Trunking is built into Brocade VDX switches to automatically aggregate multiple eligible ISL links without configuration into one physical layer trunk. This link aggregation method is very different from standard IEEE 802.3adLAGs, which can ride on top of Brocade ISL trunks. Brocade ISL Trunking traffic is load balanced among the trunk members by distributing frames in round-robin fashion.

BROCADE

When a Brocade VDX switch joins an Ethernet fabric, ISLs automatically form between directly connected switches within the fabric. If more than one ISL exists between two switches, then Brocade ISL trunks will automatically form. All ISLs connected to the same neighboring Brocade VDX switch attempt to form a trunk. No user intervention is necessary to form these trunks, and the trunks are formed only when the ports belong to the same port group. For successful trunk formation, all ports on the local switch must be part of the same port group and must be configured at the same speed. Rules for these trunks are similar to ISL trunks on proven Brocade Fibre Channel switches, that is, eight ports are allowed per trunk group.

INCREASED ETHERNET FABRIC PERFORMANCE

ISL Trunking is designed to significantly reduce traffic congestion in Ethernet fabrics. Up to eight ISLs can be combined into a single logical ISL with a total bandwidth of 80 Gbps, which can support any number of devices. To balance workload across all of the ISLs in the trunk, each incoming frame is sent across the first available physical ISL in the trunk, shown in Figure 2.

As a result, transient workload peaks are much less likely to impact the performance of other parts of the Ethernet fabric and bandwidth is not wasted by inefficient traffic routing. ISL Trunking can help simplify fabric design, lower provisioning time, and limit the need for additional ISLs or switches. Brocade ISL Trunking technology also accounts for frame sizes going on the member links of the trunk, which further results in a very even flow traffic across these ISLs in the trunk group.

If throughput greater than 80 Gbps is required between adjacent switches in the fabric, several Brocade ISL Trunks can be further aggregates using standard IEEE 802.3ad-based LAGs. This balances traffic flows between adjacent switches using hashing algorithms that take into account "6-tuples" in the frame headers.

DYNAMIC PATH SELECTION FOR OPTIMIZED ROUTING

To further optimize network performance, Brocade VCS technology supports Dynamic Path Selection (DPS). Available as a standard feature in Brocade Network OS (NOS), flow-based DPS optimizes Ethernet

fabric-wide performance by automatically routing data to the most efficient available shortest path in the fabric and load balancing across multiple shortest paths if available. DPS augments ISL Trunking to provide more effective load balancing in certain configurations, such as routing data between multiple trunk groups.

SIMPLIFIED MANAGEMENT AND DESIGN

In almost any network, management costs increase with complexity—rising with the number of elements managed. With ISL Trunking, Brocade VCS views the group of physical ISLs as a single logical ISL, a design that:

- Reduces the number of lines on a logical topology map
- Improves traffic and capacity provisioning to keep systems and applications running at full speed
- Simplifies network design, capacity planning, and fabric administration
- Trunks are self provisioning and auto-forking, leading to fewer management touch points in VCS fabrics
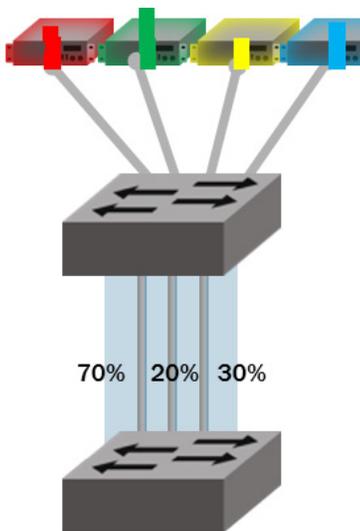
Figure 1.
Hashing algorithms in traditional switch-to-switch link aggregation can cause imbalances.

70% 20% 30%
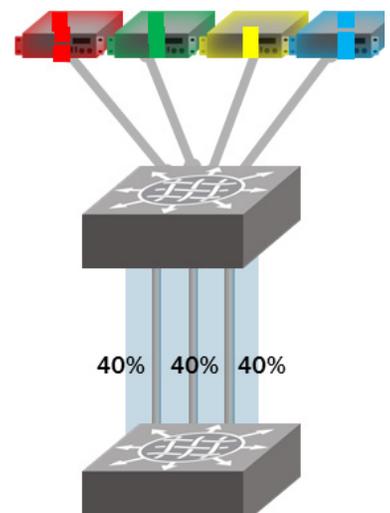
Figure 2.
Frame-based ISL Trunking in Brocade VCS ensures even link utilization.

40% 40% 40%

HIGHER AVAILABILITY

The failure of a link in a route causes the network to reroute traffic that was using that particular link—as long as an alternate path is available. Brocade Fabric Shortest Path First (FSPF) is a highly efficient routing algorithm that reroutes around failed links in less than a second.

AUTOMATIC CONFIGURATION

ISL Trunking in a built-in capability of Brocade VDX Data Center switches—no additional license is required. As a result, Brocade VDX users get optimal connectivity within Ethernet fabrics automatically when ISLs are added between any two Brocade VDX switches.

SUMMARY

Brocade ISL Trunking is a built-in capability available for all Brocade VDX Data Center Switches. This ASIC-based technology is ideal for optimizing performance and simplifying the management of multi-switch Brocade VCS-based Ethernet fabrics. In summary, ISL Trunking for Brocade VDX offers:

- Optimal Ethernet fabric performance and throughput while enabling simplified network design

- Simplified management of Ethernet fabrics by reducing the number of logical entities to manage

- Improved Ethernet resiliency if individual ISL links fail

LEARN MORE

Brocade networking solutions help the world's leading organizations transition smoothly to a virtualized world where applications and information reside anywhere. This approach is based on the Brocade One™ unified network strategy, which enables a wide range of consolidation, convergence, virtualization, and cloud computing initiatives.

Learn more at www.brocade.com.

PTS
DATA CENTER SOLUTIONS

For More Information:
(866) 787-3271
Sales@PTSdcs.com

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE