

# Fibre Channel Trunking

## Emulex<sup>®</sup> Gen 7 HBAs and Brocade<sup>®</sup> Gen 6 Switches

---

### Overview

Broadcom trunking combines Brocade<sup>®</sup> switch and Emulex<sup>®</sup> host bus adapter (HBA) technologies to provide end-to-end connectivity solutions for mission critical storage solutions. Brocade Inter-Switch Link (ISL) trunking aggregates multiple ISLs into a single, logical high-bandwidth connection that is efficient, resilient, and simple. Emulex HBA trunking extends this technology between endpoints and switches to provide end-to-end trunking benefits. Broadcom trunking technology has been data center-proven over two decades in the world's most demanding storage environments.

### Highlights

Broadcom trunking is implemented in Brocade Gen 6 switching ASICs<sup>1</sup> and Emulex Gen 7 LPe35000-series Fibre Channel HBAs to provide:

- Efficient hardware-based link aggregation
- End-to-end network resilience
- Simplified management
- Integrated Broadcom trunking technology

### Key Features

- Efficient, hardware-based link aggregation:
  - Balances traffic at a frame level
  - Aggregates bandwidth
  - Reduces response time
  - Enhances performance
- End-to-end network resilience:
  - Maintains in-order delivery
  - Minimizes disruptions and eliminates outages when links are added or removed
  - Eliminates fabric events due to addition or loss of links
  - Complimentary to multipath software and requires no software changes to the failover system
- Simplified management:
  - Enables flexible assignment and grouping of trunks to fit bandwidth needs
  - Provides a single point of management
- Integrated Broadcom trunking technology:
  - Supports Slow Drain Device Quarantine (SDDQ), encryption, and compression in sensitive Inter-Switch Links
  - Leverages frame-based granularity
  - Enables up to 8 links per trunk for fabric ISLs
  - Enables up to 4 links per trunk for end points

---

1. May require an optional license on some platforms.

## Efficient, Hardware-Based Link Aggregation

Broadcom trunking is not hash-based routing like other in-fabric and operating system-based link aggregation solutions. Such hash-based algorithms rely on characteristics of the I/O stream to equitably distribute bandwidth among the available ports. However, real-life application I/O consists of different operations and block sizes, yielding inefficient use of the available bandwidth resources.

Broadcom trunking avoids these restrictions by operating at a more granular level. By making routing decisions in hardware, at a frame-by-frame level, there is no dependence on I/O characteristics to optimize the use of the trunk’s bandwidth resources.

More specifically, Broadcom trunking optimizes throughput performance and response times by expediting transmission via selection of the shortest egress queue (see Figure 1). Trunking fills the first link available on a frame-by-frame basis, providing optimal link efficiency regardless of the I/O pattern. Further, these decisions are independently made at each “hop” in the path from initiator to target, automatically adjusting to differing conditions.

Example: Figure 1 has two links (0 and 1) and three FC flows/exchanges (A, B, and C). With Broadcom trunking, a flow or exchange can span multiple links, and bandwidth is truly aggregated. Conversely, hash-based algorithms do not actually aggregate bandwidth because if a particular flow requires additional bandwidth, it would not be permitted to access an additional link.

Figure 1: Flow Transmission Diagram

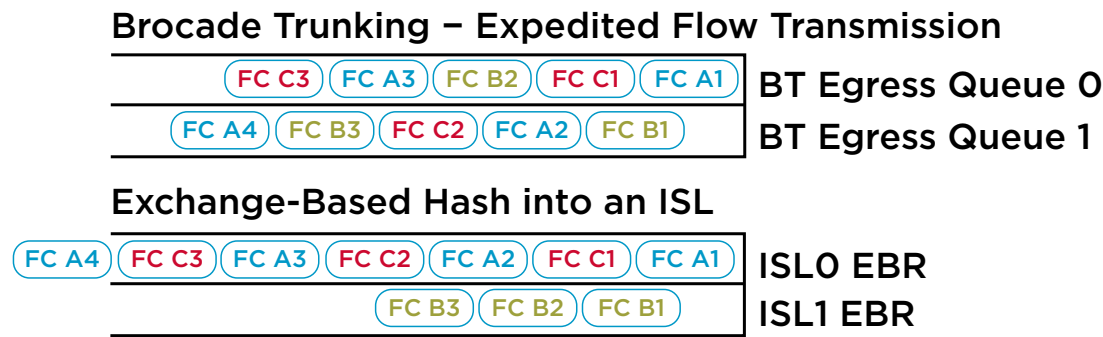
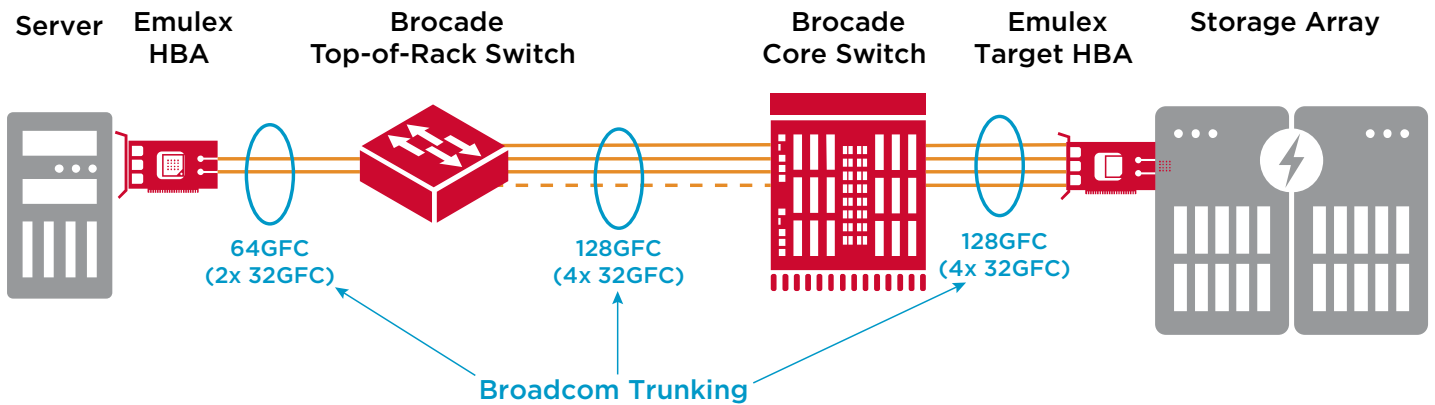


Figure 2: Broadcom Trunking Diagram



## End-to-End Network Resilience

Broadcom trunking provides additional distinct advantages in terms of resilience against link failures. If any trunked link from initiator to target goes down, the Broadcom trunking hardware automatically diverts frames to the remaining links, allowing continuous operation without disruptive network or port events that would otherwise invariably affect application performance.

Figure 2 shows a series of trunked connections—from the server to a top-of-rack switch, to the director-class core switch, and then to storage—each with two to four links.

In a production environment, any link in a Broadcom trunk may go offline (for example, the SFP+ optic may fail)—see the dotted line in the ISL link. The expected behavior is as follows:

- The maximum bandwidth of 128GFC (4x 32GFC) becomes 96GFC (3x 32GFC) for trunked links alone. Similarly, 2x 32GFC becomes 1x 32GFC at the server.
- No FSPF routing changes occur because it is a single logical link that is still online.
- No fabric reconfiguration occurs; as a result, there is no disruption to the running application.
- If the link comes back on line, the link is dynamically added back to the trunk, resuming full bandwidth operation.

## Simplified Management

Existing solutions require software-based port aggregation techniques (either at server or storage). Setting aside the latency introduced by such software layers, software solutions can be difficult to manage when one considers the differences between various OS and storage arrays. Being integral to the SAN fabric, Broadcom trunking offers a complete port aggregation solution.

## Integrated Broadcom Trunking Technology

Customers have enjoyed the benefits of Brocade trunking for years. Foundational to SAN environments, Broadcom's trunking improves performance and reduces I/O response times. Multiple links create a single logical connection of true aggregated bandwidth up to 256GFC between two domains or devices. Advanced ISL functions like Slow Drain Device Quarantine (SDDQ), encryption, and compression are supported on Gen 5 and Gen 6 platforms.

Emulex HBA trunking extends this functionality to the ends of the link, using the same fundamental signaling. Applications see a single (fat) pipe connection with bandwidth based on the sum of the port bandwidth. Emulex HBA trunking also supports Brocade's Advanced ISL functions.

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, Emulex, Brocade, and the A logo are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2018 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit [www.broadcom.com](http://www.broadcom.com).

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.