Connectrix DS-16M, DS-16M2

The Connectrix DS-16M (1 Gb/s) and DS-16M2 (2 Gb/s) are Fibre Channel Departmental Switches. The DS-16M/DS-16M2 is installed in a Connectrix EC-1200 cabinet, which also houses a Connectrix service processor.

E-Lab Navigator lists specific versions of supported firmware, as well as fabric topology constraints associated with the DS-16M/DS-16M2.

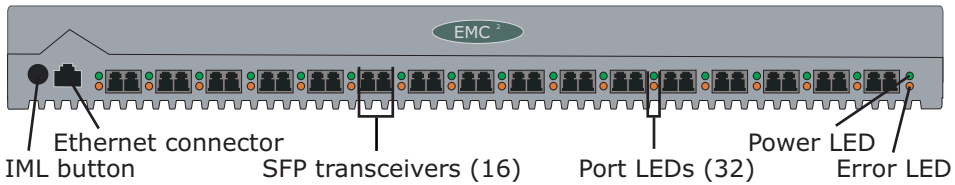

Figure 67 Connectrix DS-16M, DS-16M2

Key features

Connectivity features include:

- 16 ports available.
- Single fiber-optic SFP on each port enables connectivity.
- Multimode fiber (shortwave) support.
- 60 buffer-to-buffer credits (BB_Credits) available per port.
- Supports Class 2 and Class 3 Fibre Channel protocols.
- All ports auto-negotiate to provide either switched N_Port or E_Port connections.
- Full duplex 200 MB/s data rate per Fibre Channel port in DS-16M2.
- LC (Small Form Factor) switch port connectors.

Supported features

For the most up-to-date information on supported features, consult the *EMC Support Matrix* located on the E-Lab Interoperability Navigator.

Unsupported features

For the most up-to-date information on unsupported features, consult the most current version of the appropriate Connectrix M Series release notes and product documentation on Powerlink.

Management

Management features include:

- Centralized configuration and management of products and fabrics through Connectrix Manager, which is a client/server architecture. The Connectrix service processor laptop management station provides server functionality, and can also act as a client. Clients can be installed onto remote workstations running Windows NT, Windows 2000, Solaris, Linux, HP-UX, or AIX.
- Centralized management of up to 48 units.
- Support for online, non-disruptive code upgrades.
- 10/100 Mb Ethernet connections to CTP (control processor) for out-of-band management.
- Extensive centralized logging: event, audit, hardware, and session logs.
- Embedded Web Server and SNMP support.

Reliability, availability, and serviceability

Reliability, availability, and serviceability features include:

- Error notification:
  - Automatic communication to central support facility through modem connection
  - Automatic error notification sent through email
- Preemptive hardware health checking and redundant hardware switchover.

Further reading

Several related EMC Connectrix M manuals and release notes are available on Powerlink, under Support > Technical Documentation and Advisories > Hardware/Platforms Documentation > Connectrix Directors and Switches.

Connectrix M Series (Brocade M Series) Open Trunking

The Open Trunking feature monitors the average data rates of all traffic flow on ISLs (from a receive port to a target domain), and periodically adjusts routing tables to reroute data flows from congested links to lightly loaded links and optimize bandwidth use.

The objective of Open Trunking is to make the most efficient possible use of redundant ISLs between neighboring switches, even if these ISLs have different bandwidths. Load balancing across the ISLs does not require user configuration, other than enabling Open Trunking. However, you can modify or tweak default settings for congestion thresholds (per port) and low BB_Credit threshold if desired. In particular, you do not need to manually configure ISLs into trunk groups of redundant links where data can be off-loaded. Candidate links for rerouting flow are identified and maintained automatically. This means that flow may be rerouted onto a link that goes to a different adjacent switch, as long as that link is on the least-cost/shortest path to the destination domain ID.

The Open Trunking feature is installed and enabled on each switch/director. The algorithm controls the exits the frames will take.

To fully balance an ISL, Open Trunking must be enabled on both switches connected to the ISL.

Using Open Trunking in interop mode

If a switch with Open Trunking has an ISL to a Connectrix ED-1032, Brocade, or Cisco switch, the Connectrix M Series switch uses the Open Trunking algorithms to balance the transmit flows to the interop switch.

Potential impact on environment

Significant improvements to the Open Trunking algorithm have been made with Connectrix E/OS 6.01 to lower the likelihood of a reroute causing frames to arrive out of order at the N_Ports. Out-of-Order frame handling by N_Ports typically requires a retry of the FCP Exchange. In some cases this retry may depend on upper-level protocol timers in the host device driver, which may be as long as 60 seconds. On the other hand, overly congested links can lead to discarded frames. Error recovery for discarded frames often depends

on the upper-level protocol timers. Open Trunking helps minimize the likelihood of congested links.

Out-of-order and discarded frames caused by Open Trunking reroutes can cause the following:

- Occasional BF2D, xx1F, AB3E errors could occur on Symmetrix Fibre Channel adapters.
- Some Open Systems hosts might log temporary disk access types of events. (For example, Solaris hosts might log SCSI transport errors in the file **/var/adm/messages**.)
- Windows hosts attached to CLARiiON arrays might see Event 11s.
- In FICON environments, an IFCC error can result from an out-of-order frame.

The storage environment can be monitored for a change in the rate of these errors and events to determine whether Open Trunking improves system throughput.

Supported topologies

Open Trunking is fully supported in all fabric topologies currently supported by EMC. However, it is strongly recommended that the number of hops between server and storage pairs be kept to a minimum (ideally one hop maximum). The total hop count is the number of ISLs that a frame will have to cross when routed between the server and its storage. Not all switches in the fabric need to have Open Trunking enabled.

Example — Single fabric

In Figure 68, no limitations are placed on which host can access which storage array. For all possible combinations, Open Trunking is fully supported.

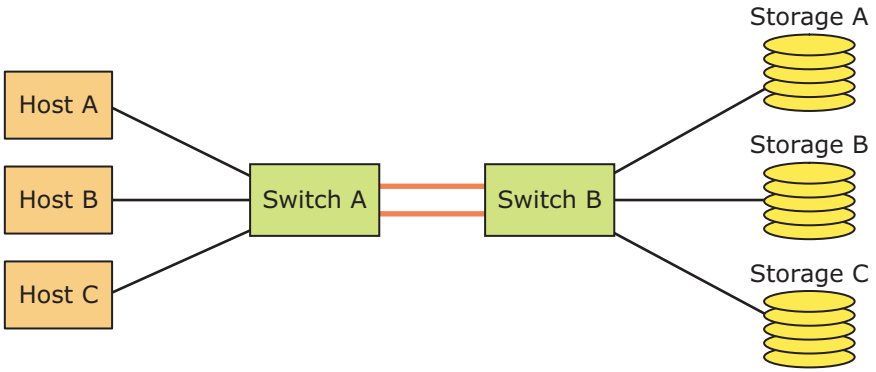

Figure 68 Open Trunking on single fabric

Example — Balanced fabrics

In Figure 69, no limitations are placed on which host can access which storage array. For all possible combinations, Open Trunking is fully supported.

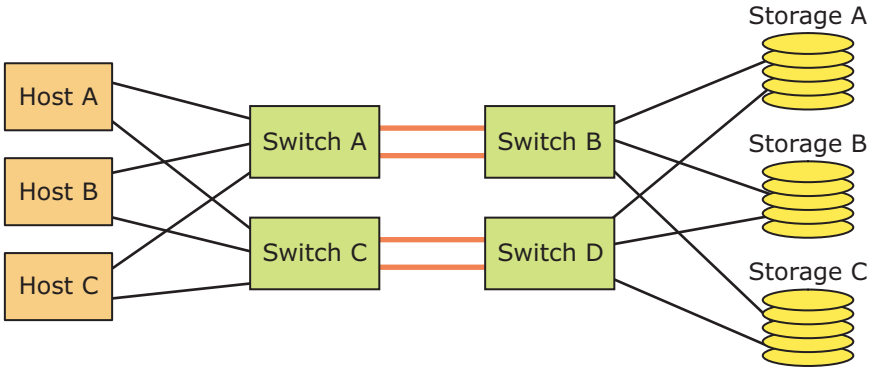

Figure 69 Open Trunking on balanced fabrics

Example — Four-corner cross-connect (full mesh)

In Figure 70, no limitations are placed on which host can access which storage array. For all possible combinations, Open Trunking is fully supported.

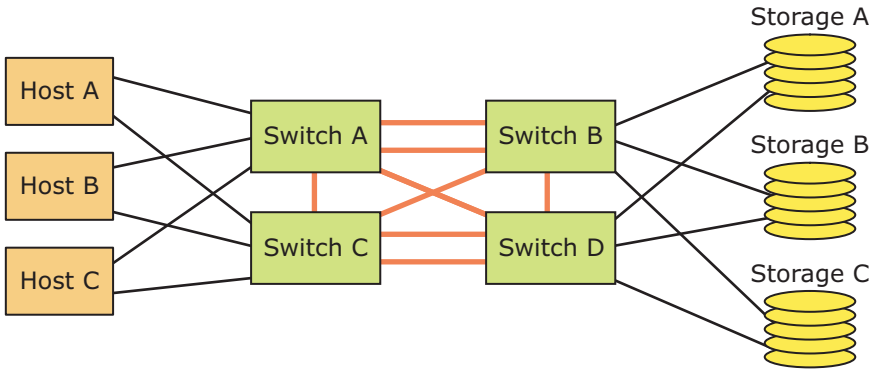

Figure 70 **Open Trunking on full mesh**

Example — Three-hop core/edge; one hop logical

In Figure 71 on page 268, although there are three hops in the fabric, all hosts logged in to the core switches (A, B, C or D) can access each storage array also logged in to the core switches without crossing more than one hop. Therefore, the hosts are considered to be one logical hop away from each other, and can access each other without restriction.

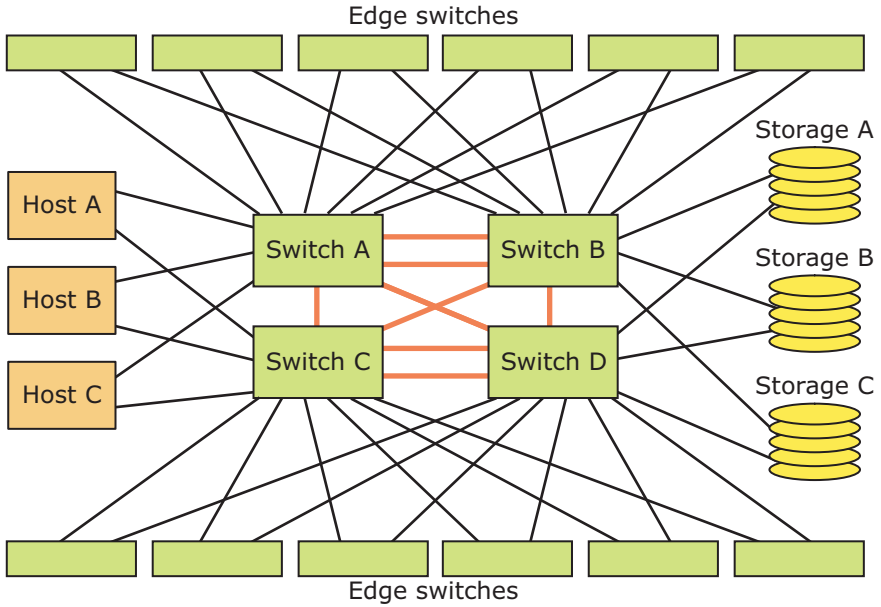

Figure 71 Open Trunking on three-hop core/edge; one hop logical (Example 1)

In Figure 72, although Host D is logged into an edge switch, it is still only one hop away from Storage D, even though there are three hops in the environment. Host D and Storage D are still considered to be one logical hop away from each other, and can access each other without restriction.

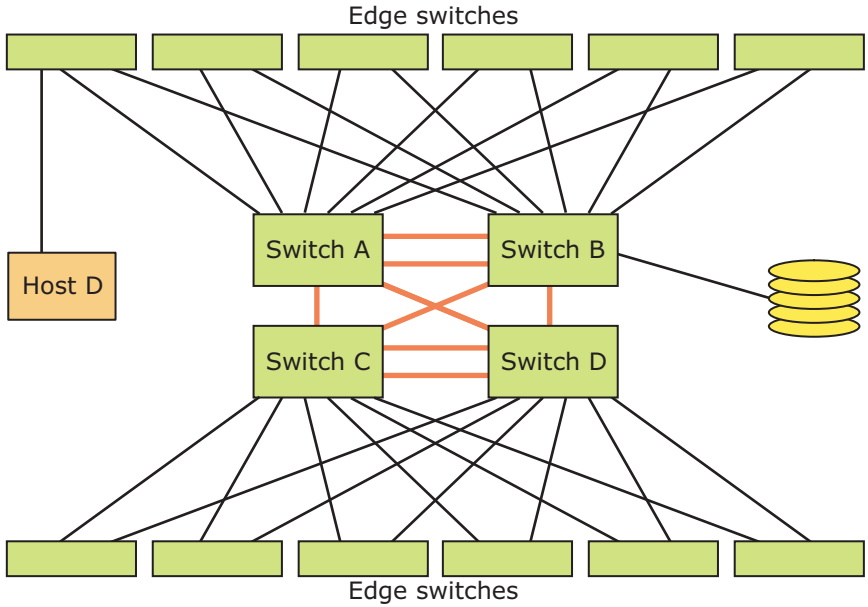

Figure 72 **Open Trunking on three-hop core/edge; one hop logical (Example 2)**

Example — Three-hop core/edge; three-hop logical

In Figure 73 on page 270, not only are there three hops in the fabric, but Host E is three hops away from Storage E. This configuration is not recommended, because as the number of hops increases, so does the chance of a frame being delivered out of order when an Open Trunking reroute occurs. Again, it is important to note that if Open Trunking did not cause a reroute, frames might be discarded anyway, due to ISL congestion.

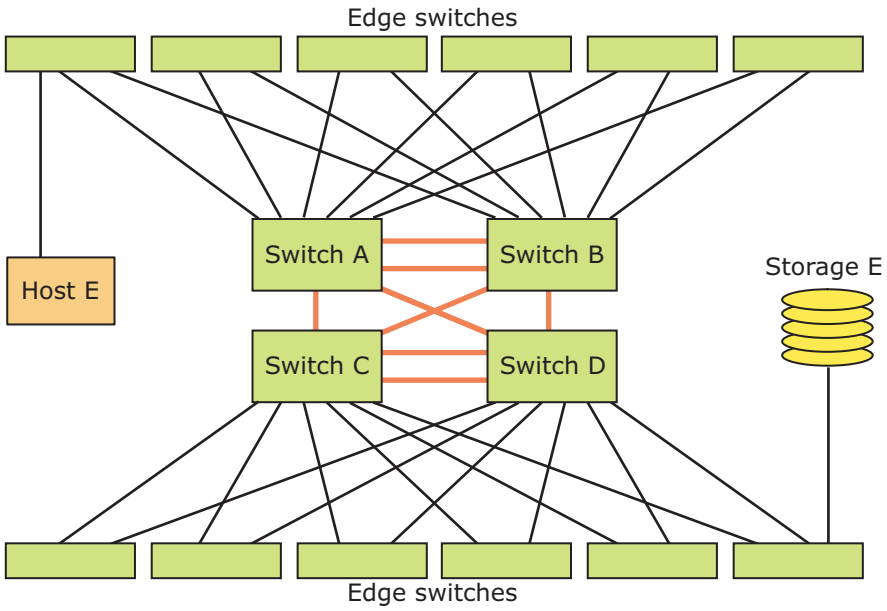

Figure 73 Open Trunking on three-hop core/edge; three-hop logical