

CISCO VALIDATED PROFILE

# Data Center Switching Multilayer Director Storage Networking Profile

April 2017

---

# Table of Contents

<b>Profile Introduction</b> .....	<b>1</b>
<b>Network Profile</b> .....	<b>3</b>
Topology Diagram .....	3
Hardware Profile .....	3
Test Environment .....	4
<b>Use Case Scenarios</b> .....	<b>5</b>
Test Methodology .....	5
Use Cases .....	5
<b>References</b> .....	<b>8</b>
Data Sheets .....	8
White Papers .....	8

# Profile Introduction

This document focuses on a typical 3-tier design used in storage area network (SAN) environments of banking customers. This is sometimes referred as *edge-core-edge* design. Within the 3-tier design, servers and storage devices are connected to edge switches. Both sets of edge switches connect to core switches through inter-switch links.

Switches used in the profile were Multilayer Director class 9700, as well as the 9500 series in some cases. They provide uncompromised availability, scalability, flexibility, security, ease of management, and transparent integration of new technologies. With data centers continually growing, SAN administrators can now design networks that meet their current needs and can scale for demanding growth. Cisco MDS 9710 and 9706 Multilayer Director class switches, along with MDS 9513, provide embedded features to help SAN administrators in these tasks.

Among the challenges specific to these environments are:

- **High availability**—MDS directors offer redundancy for all major components, with 1+1 redundant supervisors, N:N grid-redundant power supplies, and N+1 redundant backplane fabric modules. The N+1 redundancy on fabric modules provides protection against any loss of bandwidth if one of the fabric modules fails.
- **Performance and scalability**—MDS 9700 directors provide bandwidth of up to 1.5 Tbps per slot. MDS 9700 directors are designed to support line-rate 32 Gbps Fibre Channel on the same platforms in the future, without needing to replace the chassis. The MDS 9700 directors enable scale-out and scale-up architectures by providing the highest scale numbers.
- **Slowdrain detection**—The MDS directors are also architected to detect slow drain problems in the network. Slow drain problem in the fabric can occur due to several reasons such as server performance issues, host bus adapter issues, and speed mismatches between fast and slow devices in a network. Hardware-based slow drain detection in MDS switches happens at a granularity of 1 ms, which is almost instantaneous. MDS switches can also recover based on features such as stuck port recovery, congestion drop, and no-credit drop.
- **Operational simplicity**—Cisco Data Center Network Manager (DCNM) provides end-to-end visibility into the SAN network through a single-pane view, enabling effective health monitoring, diagnostics, and troubleshooting through GUI-based utilities. Cisco MDS also uses single operating system NX-OS across MDS product portfolio, along with Nexus switches. This helps with ease of installation, configuration, monitoring, and troubleshooting using the same set of CLI commands across the Cisco DC product lineup.
- **Interoperability**—MDS directors can interoperate with other SAN vendor switches such as Brocade. The switch needs to operate in interop mode and should be running the recommended code version mentioned here for working with other vendor switches.
- **Replication**—MDS directors can facilitate in remote data replication among the data centers using IP Storage Services Module. The ability to efficiently replicate critical data on a global scale not only ensures a higher level of protection for valuable corporate information but also promotes increased utilization of backup resources, reduces the impact of a catastrophic failure at a single site, and lowers total cost of storage ownership. Inter-VSAN routing (IVR) can help in routing between virtual storage area networks (VSANs) of different data centers during replication.

The following table summarizes key areas on which this profile focuses.

**Table 1** *Multilayer Director storage profile summary*

Deployment areas	Features/hardware
High availability	Redundant supervisor Redundant fabric module Redundant power supply and fans
Performance and scalability	Maximum supported zones, per the scalability guide Maximum supported device-alias, per the scalability guide Maximum supported zone db size, per the scalability guide
Slowdrain detection	Port Monitor Port Guard (using Port Monitor) DCNM monitoring
Operational simplicity	DCNM Device Manager Same NX-OS In-Service Software Upgrade (ISSU)
Interoperability	VSAN interop mode
Replication	FCIP IVR

# Network Profile

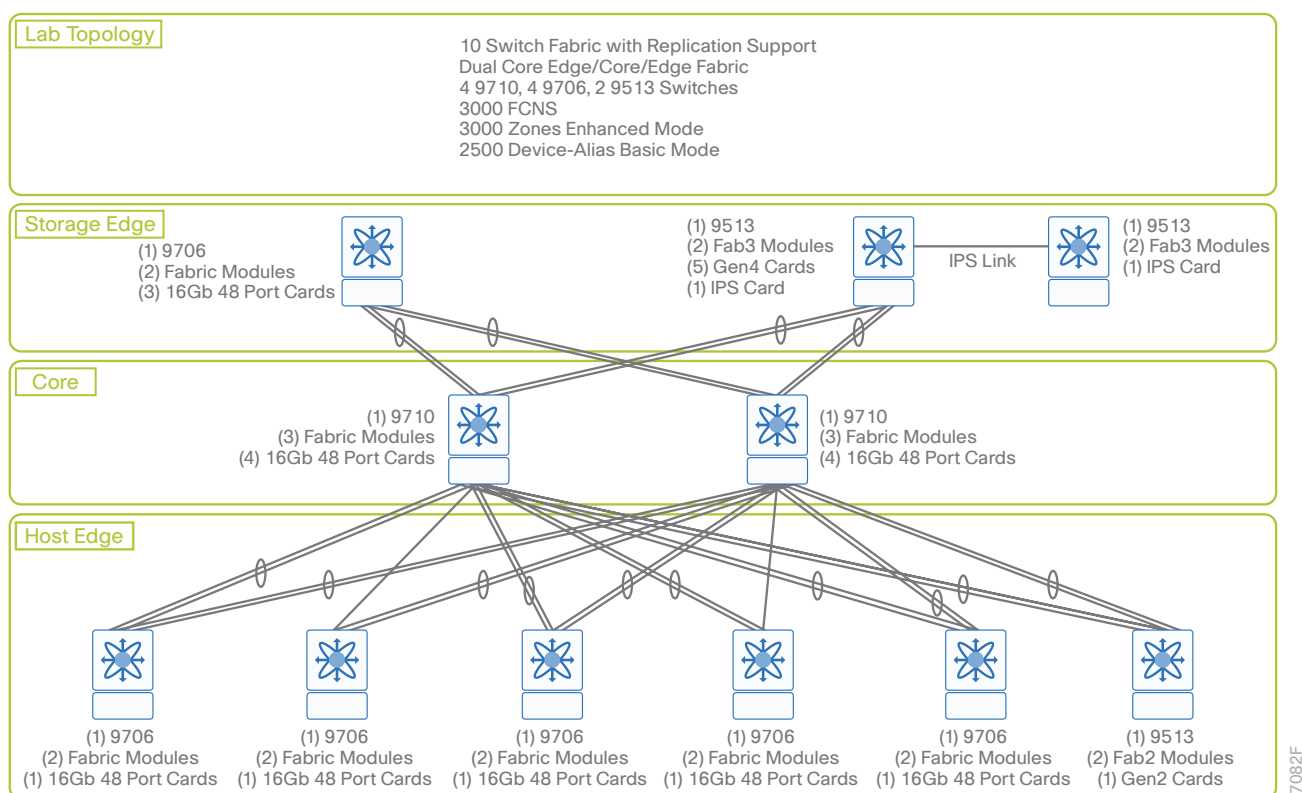
Based on research, customer feedback, and configuration samples, this profile is designed with a deployment topology that is generic and can easily be modified to fit any specific deployment scenario.

## TOPOLOGY DIAGRAM

Figure 1 shows the 3-tier design used for validation of this profile.

The topology represents a typical edge-core-edge deployment with 9710s in the core layer and 9706/9513 as edge switches. There is also an FCIP link between 9513 switches in the edge.

**Figure 1** Topology overview



## HARDWARE PROFILE

Table 2 defines the set of relevant hardware, servers, test equipment, and endpoints that are used to complete the end-to-end profile deployment.

This list of hardware, along with the relevant software versions and the role of these devices, complements the actual physical topology that is defined in Figure 1.

**Table 2** Hardware profile

VM and HW	Description
MDS 9710	Core switches
MDS 9706	Edge switches for server and storage
MDS 9513	Edge switches for server and storage+replication
JDSU	Traffic generation and congestion for slow drain scenarios
UCS servers	TACACS server, DCMN server, and SNMP polling
Laptop	DCNM client

## TEST ENVIRONMENT

This section describes the features and the relevant scale (wherever applicable) at which the features are deployed across the physical topology. Based on the use cases, which are mentioned in a later section, scale is increased to the maximum supported number, per the scalability guide. Table 3 lists the scale for each feature.

**Table 3** Base feature scale

Feature	Scale
Fabric size	10 switch fabric
Platform	Four 9710, four 9706, two 9513 switches
VSANs	2 including default
FCNS	3000
Zoning (enhanced mode)	3000
Device-alias (basic mode)	2500
NPIV	Enabled in edge hosts
RBAC	Customized roles defined
AAA	TACACS +
Port Monitor	Configured on trunk and access ports
CallHome	Enabled
CFS	CFS region used for device-alias
Slow drain	Alerts are configured
DCNM	Integrated with external oracle database
SNMP	SNMP MIB walks

# Use Case Scenarios

## TEST METHODOLOGY

The use cases shown in Table 4 will be executed using the topology shown in Figure 1, along with the test environment already explained in this document. With respect to the longevity for this profile setup, the CPU and memory usage would be monitored overnight as well as during the weekends, along with any mem-leak checks. In order to test robustness, certain negative events would be triggered during the use case testing.

## USE CASES

Table 4 describes the use cases that were executed on this profile. These use cases are divided into buckets of technology areas to show the complete coverage of the deployment scenarios.

The customer use case is composed of system upgrade/bring-up, operational triggers/configuration changes, steady state/usability, network events/link flaps, and resiliency/error recovery.

**Table 4** List of use case scenarios

No.	Focus area	Use cases
System upgrade/bring-up		
1	Hardware migration	<ul style="list-style-type: none"> <li>▪ Brownfield (mixed fabric) migration from 9500 to 9700</li> <li>▪ Pure 9500 fabric -&gt; mixed fabric -&gt; pure 9700 fabric</li> <li>▪ Greenfield (pure 9700 fabric) upgrade in 9700 environment</li> <li>▪ Interop migration from brocade to 9700</li> </ul>
2	Software upgrade—ISSU	<ul style="list-style-type: none"> <li>▪ Different software version combinations in Brownfield (mixed fabric) and Greenfield (pure 9700 fabric) deployments</li> <li>▪ Use of CLI for upgrade</li> <li>▪ Use of DCNM for upgrade (batch upgrade of multiple switches)</li> <li>▪ Scale configuration</li> </ul>
3	Control plane sync	<ul style="list-style-type: none"> <li>▪ Synchronization monitored across fabric for different applications like zone, device alias, fcns, fcdomain, cfs, etc. using CLI and DCNM</li> </ul>
4	Traffic flow	<ul style="list-style-type: none"> <li>▪ Traffic flows monitored across fabric using CLI and DCNM</li> <li>▪ Flows are mainly FC (host to storage) and FCIP (storage to storage) flows</li> </ul>

Table 4 continued

Steady state/usability		
5	Multiple instances of DCNM	<ul style="list-style-type: none"> <li>Managing different software versions of switches</li> </ul>
6	Longevity and soak	<ul style="list-style-type: none"> <li>Memory leak checks</li> <li>Traffic flow verification</li> <li>Control plane application counter verification</li> </ul>
Operational triggers/configuration changes		
7	Zoneset activation/merge with maximum supported scale numbers	<ul style="list-style-type: none"> <li>Maximum supported zone db size, per the scalability guide in Brownfield (mixed fabric) and Greenfield (pure 9700 fabric) deployments</li> <li>Maximum supported zones/zonesets, per the scalability guide in Brownfield (mixed fabric) and Greenfield (pure 9700 fabric) deployments</li> </ul>
8	Device-alias activation/merge with max supported scale numbers	<ul style="list-style-type: none"> <li>Maximum supported device-alias, per the scalability guide in Brownfield (mixed fabric) and Greenfield (pure 9700 fabric) deployments</li> </ul>
9	Addition of IVR in replication environment	<ul style="list-style-type: none"> <li>IVR activation in replication environment</li> <li>IVR merger in replication environment</li> </ul>
10	SNMP polling	<ul style="list-style-type: none"> <li>Several instances of SNMP walk running.</li> <li>Specific SNMP OID polling</li> </ul>
Network events/link flaps		
11	Slowport detection/diagnostics	<ul style="list-style-type: none"> <li>Slow port scenario including sub 100 ms detection</li> <li>Stuck port scenario</li> <li>TxWait scenario for Gen4 and Luke LC</li> <li>Relevant Port Monitor/Port Guard configuration for all slow drain scenarios</li> <li>Diagnostics in DCNM</li> <li>Logs in form of syslog, rmon and OBFL</li> </ul>
12	DCNM	<ul style="list-style-type: none"> <li>Migration from post gre to oracle database</li> <li>OHMS/GOLD diagnostics access via Device Manager and CLI</li> <li>Loopback tests using DM/CLI</li> </ul>
13	Convergence/redundancy	<ul style="list-style-type: none"> <li>Measurement/comparison of traffic convergence during module reload, switch reload, link flap, and fabric module reload triggers</li> </ul>
14	Switch reload	<ul style="list-style-type: none"> <li>Fabric convergence during switch reload</li> </ul>



Table 4 continued

Resiliency/error recovery		
15	Process crash	<ul style="list-style-type: none"> <li>▪ Stateless and stateful process restart after crash</li> <li>▪ Process killing on every LC and SUP</li> <li>▪ Process crashes leading to system switchover</li> <li>▪ Process crashes when flash partition is full</li> </ul>
16	Hardware failure	<ul style="list-style-type: none"> <li>▪ System recovery after ASIC, module, Xbar, SUP failure</li> <li>▪ Any collaterals due to ASIC, module, Xbar, sup failure</li> </ul>
17	Online Insertion and Removal (OIR)	<ul style="list-style-type: none"> <li>▪ System recovery after OIR of module, Xbar, power supply, fan, sup</li> <li>▪ Any collaterals due to OIR of module, Xbar, power supply, fan, sup</li> </ul>
18	ISSU with scale configuration	<ul style="list-style-type: none"> <li>▪ Any failures during ISSU with multi D scale configuration, per the scalability guide</li> </ul>
19	Software and hardware counters	<ul style="list-style-type: none"> <li>▪ Clearing different software and hardware counters</li> <li>▪ Sync between software and hardware counters</li> </ul>

# References

## DATA SHEETS

[Cisco MDS 9710 Multilayer Director Data Sheet](#)

[Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module Data Sheet](#)

## WHITE PAPERS

[Cisco MDS 9000 Family Diagnostics, Error Recovery, Troubleshooting, and Serviceability Features White Paper](#)

[Monitoring and Alerting in Cisco MDS Fabric White Paper](#)

[Five Reasons Why You Should Choose the Cisco MDS Directors Cheat Sheet](#)

[Large SAN Design Best Practices Using Cisco MDS 9700 and MDS 9500 Multilayer Directors White Paper](#)

[Slow-Drain Device Detection, Troubleshooting, and Automatic Recovery](#)



Please use the [feedback form](#) to send comments and suggestions about this guide.



**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2017 Cisco Systems, Inc. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com/go/trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)