

CISCO VALIDATED PROFILE

# Data Center Switching Unified Fabric Converged Network Profile

April 2017

---

# Table of Contents

Profile Introduction .....	1
Network Profile.....	3
Topology Diagram .....	3
Hardware Profile .....	4
Features and Scale .....	4
Use-Case Scenarios.....	5
Test Methodology .....	5
Use Cases .....	5
References .....	8

# Profile Introduction

This document focuses on the deployment of Unified Fabric Converged Network Solution in a data center. Cisco Unified Fabric provides the networking foundation for the Cisco Unified Data Center, on which you can build the data center architecture, whether you run a traditional data center or are on the journey to full private cloud computing or hybrid private and public cloud computing.

Cisco Unified Fabric is built on three main pillars: convergence, reduced operational cost, and scalability. Cisco Unified Fabric can help you reduce costs, migrate to the next generation data center, and bring value to your business.

## Convergence

Convergence of the data center network is the melding of the storage network (SAN) with the general data network (the local area network). Cisco Unified Fabric supports data center convergence by consolidating multi-protocol storage and local area network (LAN) traffic onto a single, scalable, and intelligent network. This cuts costs and increases efficiency. Companies need to keep using their current SAN infrastructure while extending it gradually, transparently, and non-disruptively into the Ethernet network. The traditionally separate LAN and SAN fabrics evolve into a converged, unified storage network through normal refresh cycles that replace old servers containing host bus adapters with new servers containing converged network adapters, and storage devices undergo a similar refresh process.

Cisco customers can deploy an Ethernet network for the data center that conforms to the needs of storage traffic, with a lossless, in-order, highly reliable network for the data center by using Fiber Channel over Ethernet (FCoE).

## Reduced Operational Cost

Consolidation of the general data and storage network can save customers a lot of money. For example, customers can significantly decrease the number of physical cables and ports by moving to a converged 10 Gigabit Ethernet network because the number of cables required for reliability and application bandwidth is significantly reduced.

A standard server requires at least four networking cables: two for the SAN and two for the LAN with current 1 Gigabit Ethernet and Fiber Channel technology. Often, more 1 Gigabit Ethernet ports are needed to meet bandwidth requirements and to provide additional connections for server management and for a private connection for server clusters. Two 10 Gigabit Ethernet converged ports can replace all these ports, providing a cable savings of at least 2:1. From a larger data center perspective, this cable reduction means fewer ports and the capability to decrease the number of switches and layers in the data center, correspondingly reducing the amount of network oversubscription. Reducing cabling saves both acquisition cost and the cost of running the cables, and it reduces cooling costs by improving airflow.

Also, by eliminating or reducing the second network, customers end up with less equipment in the data center, saving on costly rack space, power, and cooling and making the overall data center much more efficient. However, the biggest cost savings is the capability for administrators to shift their time from maintenance of two separate networks and their associated cables and hardware to working on projects that directly benefit the business.

## Scalability

A simple definition of scalability is the ability to grow as needs change, often described by the number of nodes that a given architecture can ultimately support. Cisco Unified Fabric delivers true scalability: not just enabling increased growing port count as needed, but doing so without compromising on performance, manageability, or cost.

Scalability begins with 10 Gigabit Ethernet. 10 Gigabit Ethernet allows customers to consolidate their networks, which means fewer tiers to the network and fewer overall ports while providing exponentially more usable bandwidth for servers and storage. By moving to 10, 40, and 100 Gigabit Ethernet technologies, customers will be able to consolidate the number of ports and cables dedicated to servers as well as the overall number of switches under management in the data center. The reduction of devices reduces management overhead and comes with a concomitant reduction in rack space use, power, and cooling.

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

**Table 1** *Unified Fabric Converged Network profile summary*

Deployment areas	Features
Converged network	FCoE, Fex, zone, device alias
High availability	Redundant supervisor Redundant fabric module Redundant power supply
Efficient network management	DCNM
Performance and scalability	10G to 40G FCoE
Interoperability	Nexus 7706, Nexus 5696 and Nexus 2248P

# Network Profile

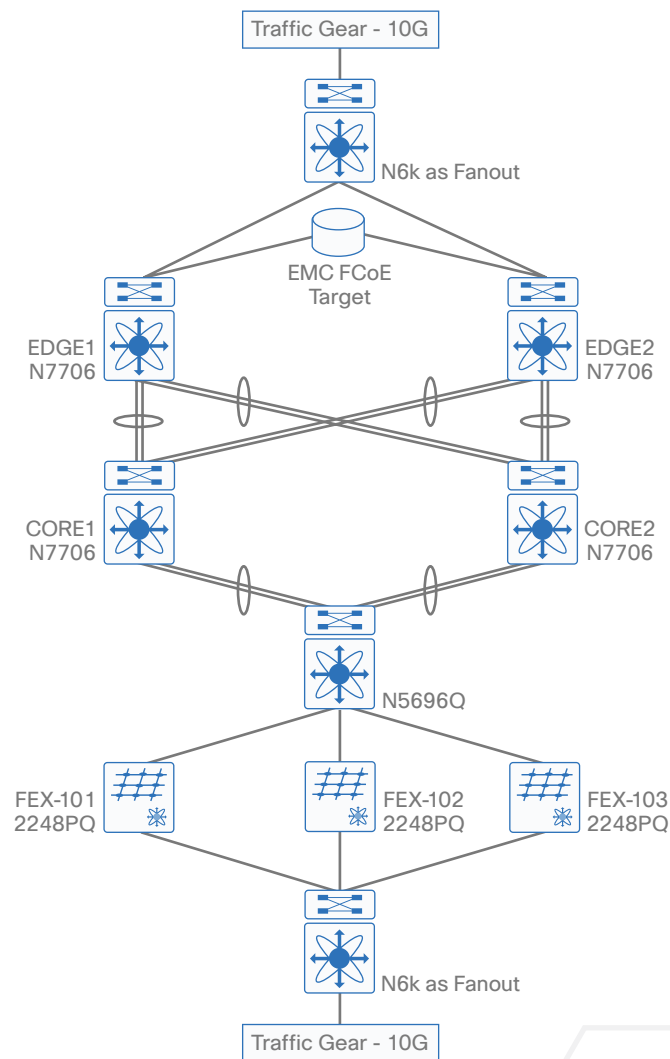
Based on the research, customer feedback, and configuration samples, this profile is designed with a deployment topology that is generic and can easily be modified to fit any specific deployment scenario. Please note that in this profile we have included only the SAN part of the data center network.

## TOPOLOGY DIAGRAM

Figure 1 shows the topology that is used for the validation of the profile.

This topology represents a typical Edge-Core-Edge SAN network. Nexus7K devices are positioned in the storage core and storage edge with N6K at the access layer as the converged network, with Fex facing the hosts

*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, complement the actual physical topology defined in Figure 2.

**Table 2** *Hardware profile of servers and endpoints*

VM and HW	Description
Nexus 7706	As Storage Core and Edge running 6.2(12) which was upgraded to 7.2(1)D1(1) to support 40G FCoE
N77-F248XP-23E	FCoE 10G line card
N77-F324FQ-25	FCoE 40G line card
Nexus 5696Q Fex 2248PQ	Converged Access running 7.0(6)N1(1)
UCS Server	To manage and host the virtual machines
NetApp target	Storage device
Nexus 56128 Nexus 5648	As fanout device
Spirent	Generate traffic streams

## FEATURES AND SCALE

This section contains the description of the features covered in Unified Fabric Converged Network profile deployment. The relevant scale numbers at which the features are deployed across the physical topology are also captured. Table 3 lists the scale for the profile (only FCoE). For the FCoE scale limits supported by Nexus 7k, see [Configuration Limits for FCoE](#).

**Table 3** *Unified Fabric Converged Network profile*

Feature	N7k scale	N6k scale
Zone	1396	1396
Device alias	974	974
Host	0	104
Target	88	0
FCoE VLAN	1	1
VSAN	1	1
AAA	Enabled	Enabled
TACACS+	Enabled	Enabled

# Use-Case Scenarios

## TEST METHODOLOGY

The use cases shown in Table 4 will be executed using the topology defined 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 the robustness, certain negative events would be triggered during the use case testing.

## USE CASES

Table 4 describes the use cases that were executed on the Unified Fabric Converged Network Profile. The test coverage is divided into buckets of customer use-case 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	ASCII replay ISSU Copy run start, reload	Network administrator should be able to perform upgrade and downgrade between releases seamlessly.  Validate image copy to the system and check for any incompatibility  Validate that upgrade will be non-disruptive.  Validate that after upgrade none of the configurations are lost.  Validate that none of the devices have logged out during upgrade.
2	Platform additions	Addition of Nexus platforms with existing network topology should be possible.  For example, upgrading the network from 10G to 40G FCoE by inserting new 40G FCoE LC in the N7ks.  Validate that on bringing up the 40G FCoE links, traffic should seamlessly move over to the new 40G links.
3	Reload	All of the configuration should be migrated seamlessly during the reload operation.  Verify that all of the devices have logged in back after reload  Verify that the FCNS database is in sync after the switch comes up
4	Traffic forwarding	No loss of FCoE traffic across the network topology.

Table 4 continued

Operational triggers/configuration changes		
5	Config changes and modifications	Verification of adding/deleting/modifying zone members Interface configurations—create port-channels
6	SNMP	Verification of SNMP traps for: Link up/down Module insertion/remove
Steady state/usability		
7	Soak	Verification of system stability. Soak the network for at least 48hrs. System is stable and no loss of traffic is seen. No core dump or module reloads or memory leaks seen.
Network events/link flaps		
8	Link flaps	Verify that the system holds well and recovers to working condition after the following events are triggered: Link flaps Port flap
9	Power OFF/ON	Verification of FCoE ability to re-login to fabric and resume traffic post outage. Zone, device alias, and FCNS are in sync after the system is up.
10	Module reload	Verification of the feature after system goes for a reload All devices login back once module is back online. FCNS database is in sync Traffic resumes
11	Traffic based triggers	Verification of different traffic flows and streams with Spirent. Streams of different frame size Streams with multiple oxid
12	Congestion	Verify that when there is a congestion, PAUSE is generated and traffic is flow controlled No drop in FCoE traffic <i>Note: Validated between N77 for 10G FCoE</i>
13	Slow drain	Validate slow drain traffic with default timeout values. The default congestion timeout is 500ms Validate the drop counters for slow drain <i>Note: Validated between N77 for 10G FCoE</i>



Table 4 continued

Resiliency/error recovery		
14	System switchover	Verify that system switchover has no impact on FCoE traffic. All features like zone and device alias continue to work No cores dumps are seen
15	ISSU	Verify that ISSU is not disruptive and all the features continue to work after upgrade

# References

[Cisco Unified Fabric: Enable the Data Center Network white paper](#)

[Cisco Data Center Convergence web site](#)





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)