CISCO SYSTEMS

Cisco Data Center Infrastructure Design Guide 2.1 Release Notes

OL-11569-01

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.

CONTENTS

iv

Cisco Data Center Infrastructure Design Guide 2.1 Release Notes

This Release Note highlights the changes in Versions 1, 2 and 2.1 of this guide, and describes the hardware and software components that have been validated for each version.

Note   Version 1.1 and 2.0 of this design guide is referred to as SRND and 2.1 and future releases will be referred to as design guide.

Topics Covered in Version DG 1.1

The Data Center Infrastructure DG version 1.1 was first released in March, 2004. Note that DGv2 builds on DGv1.1 design recommendations and is not intended to be a replacement. A more detailed table listing items covered in DGv1.1 is provided in the Readme file. Topics covered in v1.1 include the following:

- Layer 3 design with OSPF and EIGRP
- Layer 3 security
- Layer 2 design with STP, UDLD, LoopGuard, VTP, and BPDU Guard
- Layer 2 security—Private VLANs, Port Security
- NIC teaming
- Mainframe connectivity and OSA/OSPF design
- Spanning tree design with Rapid-PVST+ and MST

version 1.1 hardware and software were based on the use of Sup2 and the Cisco IOS 12.1 release train on the Cisco Catalyst 6500.

What's New in the Version DG 2.0?

This design guide (DG) provides guidelines for designing and building the data center switching infrastructure. The major differences between DG 2.0 and version 1.1 are as follows:

- Next generation hardware and software—The Sup720 and 10GE line modules now dominate all data center-related designs as the need for higher performance and scalability are even more critical today in the data center.

- Data center design models—The multi-tier model that dominated most data center designs is now complimented with the server cluster model as large clustering environments begin to take a pivotal role in large enterprise data centers.

- Next generation 1RU switches—Next generation high performance 1RU switches are examined in the access layer.

- Scalability focus—Data center consolidation is a growing trend, while centralized data center requirements are growing in terms of number of servers, processing power, and bandwidth consumption. This is placing more strain on some critical technology areas in the data center such as spanning tree and HSRP.

- Access layer switch design—The access layer requirements are more dynamic, and now consist of a hybrid of modular chasses as well as rack-based switches (1/2RU). The various approaches taken in building an access layer design have implications that affect scalability in other areas such as spanning tree, port density, and service module interoperability.

Hardware and Software used in this DG

Since the first (version 1.1) of this DG, many new hardware modules, platforms, and related software images have been released and used in the data center architecture. This section identifies the hardware and software used in testing for this version of the DG.

Although many different hardware solutions may be appropriate for a particular data center design, this document and the related test lab activities that support it are focused on a subset of hardware and software that has been selected based on many requirements (for example, technical features, scalability, best practices, and customer acceptance). Other hardware platforms not mentioned in this DG may also be perfectly suitable for data center deployments.

Platforms

The following table lists the platforms and software versions used in support of this DG.

Layer	Platform and CPU/Supervisor	Software Version
Core layer	Catalyst 6500/Sup720/PFC3A	12.2(18)SXD3
Aggregation layer	Catalyst 6500/Sup720/PFC3A	12.2(18)SXD3
Access layer	Catalyst 6500/Sup720/PFC3A	12.2 (18) SXD3, 12.2.(18) SXF for FlexLinks
Access layer	Catalyst 4948-10GE	12.2(25)EWA1
Access layer	Catalyst 4948-10GE	12.2(25)EWA1
Services layer	Catalyst 6500/Sup2/PFC	12.2(18) SXD

Modules

The following table lists the modules used in support of this DG.

Platform	Module	Software Version
Catalyst 6500	6704- 4 port 10GE	
Catalyst 6500	6748- 48 port SFP (sx and tx) and copper-only modules	
Catalyst 6500	6724- 24 port SFP	
Catalyst 6500	CSM	4.2(2)
Catalyst 6500	FWSM	2.3(2)

Note The software version used for the solution verification does not constitute a software recommendation. See Limitations and Restrictions for Version 2, page 3 for caveats and in which release they are addressed.

Limitations and Restrictions for Version 2

The following are limitations and restrictions for version 2 of this document:

- Multicast sources and SPAN reflector—When using Sup720 with an FWSM in the chassis running Cisco Native IOS, by default a SPAN session is used. If you check for unused sessions with **show monitor**, you see that "session 1" is in use:

  ```
  agg#show monitor
  Session 1
  ---------
  Type : Service Module Session
  ```

  This session is automatically installed for the support of hardware multicast replication when a firewall blade is in the Catalyst 6500 chassis. This is because an FWSM cannot replicate multicast streams, so if multicast streams sourced behind the FWSM must be replicated at Layer 3 to multiple line cards, the automatic session copies the traffic to the supervisor through a fabric channel.

  If you have a multicast source that generates a multicast stream from behind the FWSM, you need the SPAN reflector. If you place the multicast source on the outside VLAN, the SPAN reflector is not necessary. The SPAN reflector is incompatible with bridging BPDUs through the FWSM. You can disable the SPAN reflector by using the **no monitor session service module** command.

- Spanning Tree and FWSM—When providing transparent firewall services with the FWSM, it is important to bridge BPDUs through the transparent context to prevent a loop condition in the event that redundant FWSM pairs incorrectly become active/active. The global command **no monitor session service module** must be enabled to prevent a spanning tree loop occurring on DEC trunks passing the context VLANs.

- Transparent services while using 802.1s—When using the FWSM transparent mode in conjunction with the 802.1S (MISTP) spanning tree protocol, an active/active misconfiguration on the FWSM can cause a spanning tree loop. The reason is that with 802.1s, you cannot bridge two VLANs together. This is not an FWSM limitation but is imposed by the way 802.1s operates according to the standard. If transparent mode contexts are used, Cisco recommends using 802.1w Rapid PVST+.

- Distributed EtherChannel (DEC) and service modules—It may be necessary to implement the global command **fabric switching-mode force bus-mode** when using DEC in the presence of service modules. This command forces the service modules to operate in bus mode, which can have performance implications because it forces all traffic that goes through the FWSM to use the local bus through the supervisor (CSCee10005).

- Distributed EtherChannel and DFCs—There are possible implications when using DEC across DFC-enabled line cards. The workaround is to remove DFCs or to use an EtherChannel that does not span across line cards (CSCee10005).

- HSRP CPU activity spikes—Once ~90 or more HSRP instances are configured and enabled on a Catalyst 6500 with a Sup720, a CPU spike appears in a 10–12 minute cycle. This is a cosmetic spike only and does not adversely affect operation.

- LoopGuard global and FWSM errdisable issue—You should not enable LoopGuard globally on the aggregation switches if an FWSM transparent mode is present, because LoopGuard would be automatically applied to the internal EtherChannel between the switch and the FWSM. After a failover and a failback, this configuration would cause the secondary firewall to be disconnected because the EtherChannel would go into *err-disable* state.

- RootGuard fallback—When using Rapid PVST+, an interface that has been placed into a Root_Inc_State by RootGuard does not automatically recover. The current workaround is to shut-noshut the interface (CSCsc95631).

- RHI routes during SSO switchover—When using NSF/SSO in a switch that has a CSM configured for RHI with advertise active, the flows that are going through the CSM stop flowing for a short duration after the SSO switchover. The workaround is to extend the RHI failover and retry timers.

- HSRP state during SSO switchover—HSRP state between aggregation switches is not maintained during an SSO switchover. Ideally, the HSRP state stays in standby on the adjacent aggregation switch during an SSO switchover. It currently moves to an active state based on the hello timer value expiration. More detail is provided in Chapter 7 of this guide(CSCeg33278, CSCec27709).

- FlexLinks and uplink failover convergence time—The aggregation layer switch may take too long to converge/resume after an uplink FlexLink failover because the MAC address table is not refreshed quickly enough. This would be noticeable if a large number of MAC addresses are present on the access layer switch (CSCsd69806).

What's New in the Version 2.1

The guide was updated to reflect the new WS-X6708-10G-3C availability on the Catalyst 6500 Series switch. This new line card increases the overall system density at 10GigE rates. Testing is still forthcoming, but the module is included to reflect the overall port density increase.

Related Documents

- Catalyst 6500 Series WS-X6708-10G-3C Data Sheet—

  http://www.cisco.com/en/US/products/hw/switches/ps708/products_data_sheet09186a00801dce34.html

- Data Center Infrastructure Architecture Version 1.1—
  http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/cdccont_0900aecd800e4d2e.pdf

- Server Farm Security in the Business Ready Data Center Architecture v2.0—
  http://www.cisco.com/warp/public/732/systems/docs/dcsrndbk.pdf

- Integrating Security, Load Balancing, and SSL Services—

  – http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/cdccont_0900aecd800f252b.pdf

  – http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidance0900aecd800f252b.html

- Internet Edge Design Architectures Design Guide—

  – http://www.cisco.com/application/pdf/en/us/guest/netsol/ns304/c649/ccmigration_091 86a008014ee4e.pdf

  – http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidance09186a008014ee4e.html

- Securing Server Farms Design Guide—

  – http://www.cisco.comv/application/pdf/en/us/guest/netsol/ns304/c649/ccmigration_09186a008014edf3.pdf

  – http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidance09186a008014edf3.html

- Load Balancing Data Center Services Design Guide—

  – http://www.cisco.com/application/pdf/en/us/guest/netsol/ns377/c649/cdccont_0900aecd800eb95a.pdf

  – http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidance0900aecd800eb95a.html

- SAN Extension for Business Continuance Design Guide—

  – http://www.cisco.com/application/pdf/en/us/guest/netsol/ns378/c649/cdccont_0900aecd800ed146.pdf

  – http://www.cisco.com/en/US/netsol/ns656/networking_solutions_design_guidance0900aecd800ed146.html