Innovations for the next generation Data Center
Cisco Nexus 7000

Maciej Bocian

Consulting System Engineer,
Data Center/Storage Networking

Central & Eastern Europe

mbocian@cisco.com

The New Data Center

Consolidation Needed to Combat Infrastructure Sprawl and its attendant capex/opex impact

Cisco Nexus7000 Delivers Infrastructure Scalability to defer the need to add infrastructure

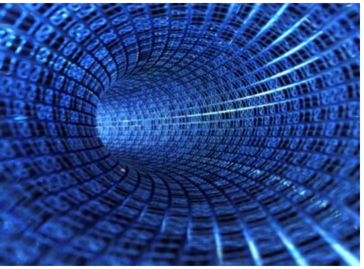

Virtualization of Resources to Easily and Efficiently Adapt to Change

Cisco Nexus7000 Transport Flexibility to meet growing needs and address next-gen protocols

Automation Improves Operations Effectiveness and Infrastructure Availability

Cisco Nexus7000 Operational Continuity through a “Zero Service Loss” system architecture

Nexus 7000 Chassis

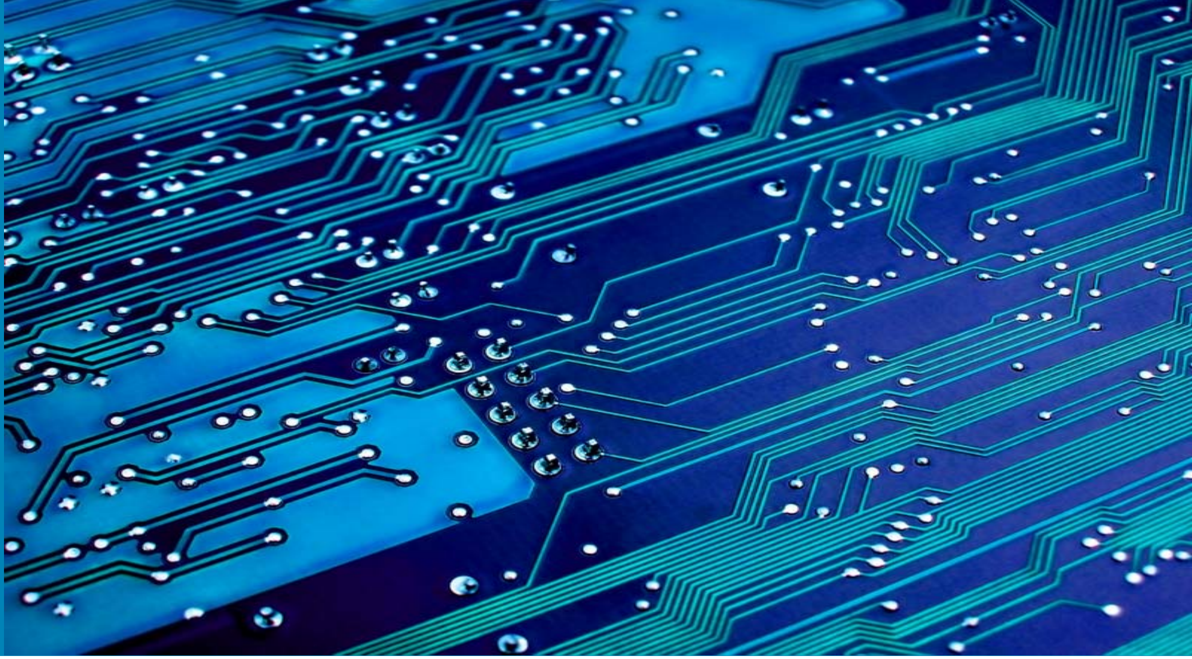

Nexus 7010 10-Slot Chassis

- First chassis in Nexus 7000 product family
- Optimized for data center environments
- High density
  - 256 10G interfaces per system
- High performance
  - 1.2Tbps system bandwidth at initial release
  - Initially 80Gbps per slot
  - 60Mpps per slot
- Future proof
  - Initial fabric provides up to 4.1Tbps
  - Product family scaleable to 15+Tbps
  - 40/100G and Unified Fabric ready

Nexus 7010 Chassis Front

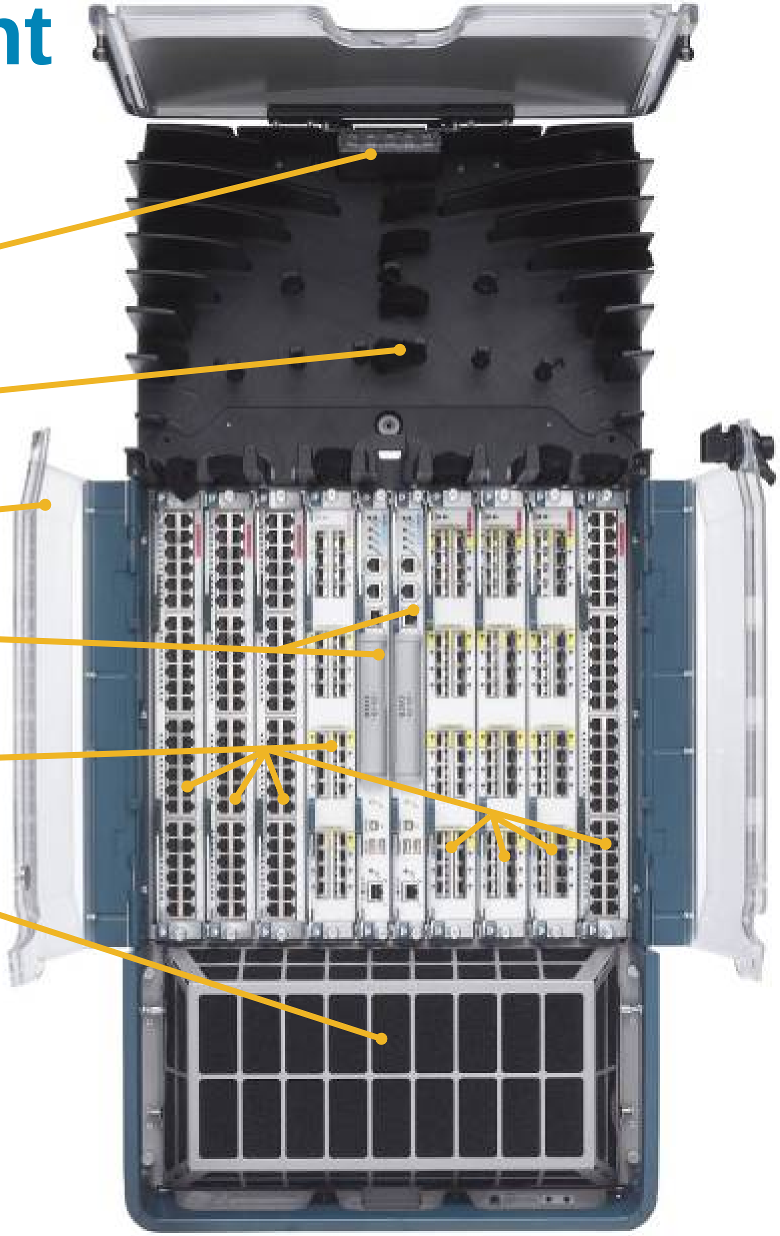

System status LEDs

Integrated cable management with door

Lockable front doors (removable)

2 supervisor slots (5-6)

8 payload slots (1-4, 7-10)

Air intake with optional filter

Nexus 7010 Chassis Back

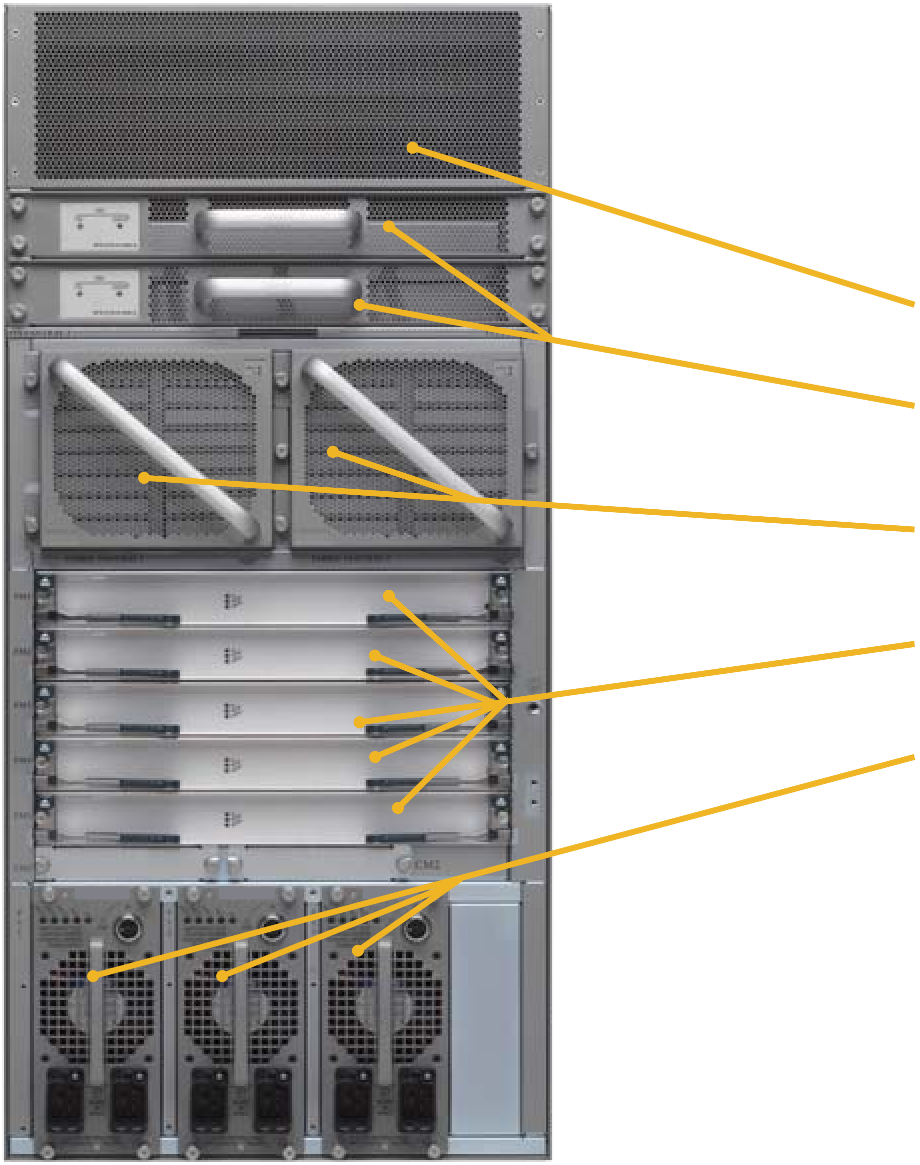

Air exhaust

2 system fan trays

2 fabric fan trays

5 crossbar fabric modules

3 power supplies

System Power

- 6000W AC power supply for Nexus 7000 series chassis
- Dual inputs at 220/240V or 110/120V
- Proportional load-sharing among supplies
- Hot swappable
- Blue beacon LED for easy identification

N7K-AC-6.0KW
Output
Fault
ID
ON
STBY
SWITCH MUST BE IN POSITION STANDBY TO INSTALL/REMOVE POWER SUPPLY

Nexus 7010 Power Redundancy

6 power supplies in 3 physical bays

Power redundancy modes:

- Power Supply Redundancy (default)
- Input Source Redundancy

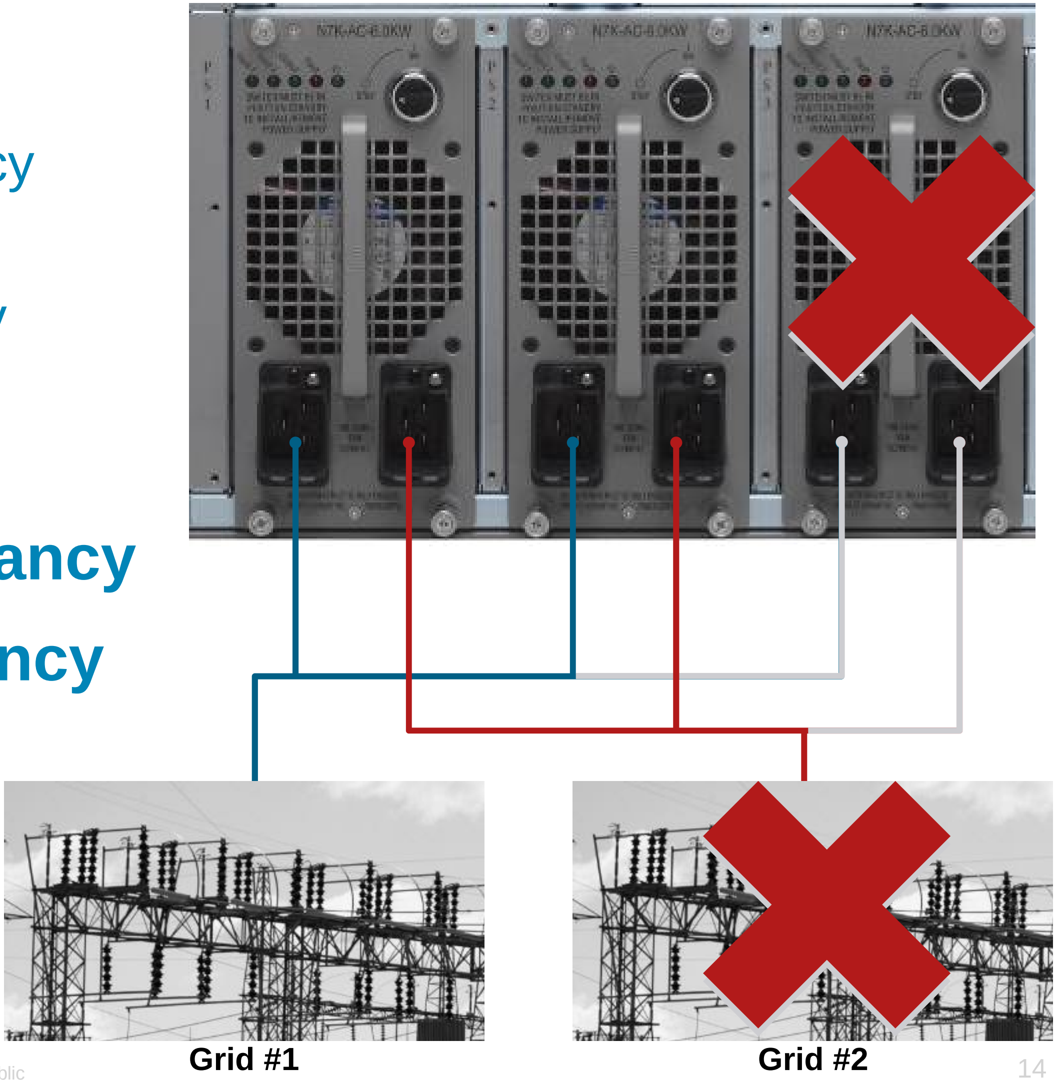

Power Supply Redundancy

Input Source Redundancy

System Cooling

- Variable speed redundant fans provide complete system cooling
- Fans removed from chassis rear – no disruption of cabling
- Hot swappable
- Blue beacon LED for easy identification

- Redundant system fan trays provide cooling of I/O modules and supervisor engines

- Redundant fabric fans provide cooling of crossbar fabric modules

Other Hardware Features

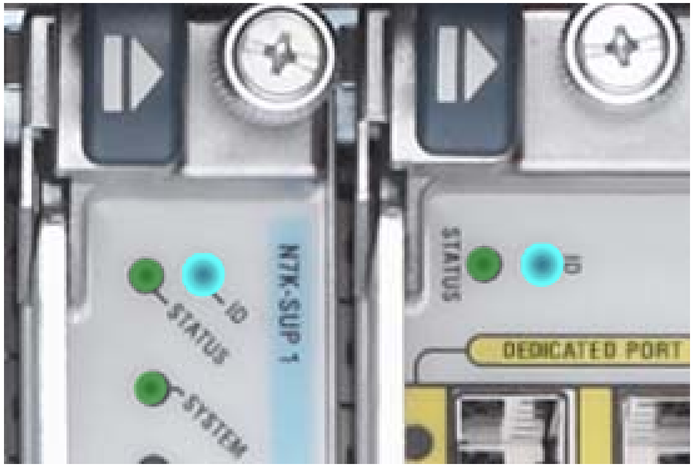

Blue beacon LEDs allow for easy FRU identification for servicing

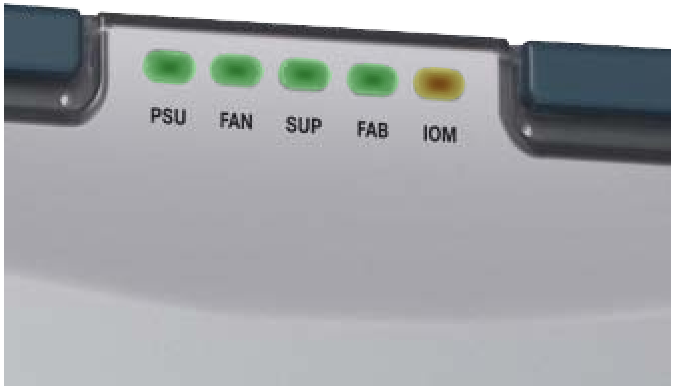

System LEDs provide aggregate view of system status

- Power supplies
- Fan trays
- Supervisor engines
- Fabric modules
- I/O modules

Locking ejector levers ensure proper module seating and prevent accidental disengagement

Nexus 7000 Supervisor Engine

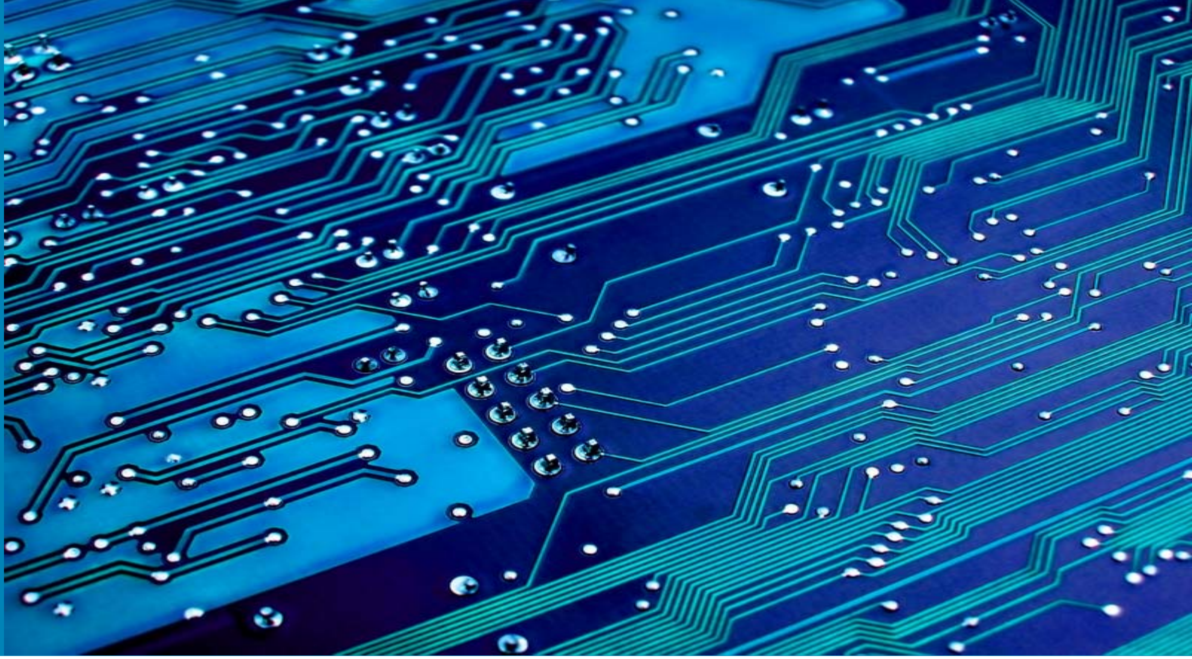

Supervisor Engine

- Dual-core 1.66GHz Intel Xeon processor with 4GB DRAM
- Connectivity Management Processor (CMP) for lights-out management
- 2MB NVRAM, 2GB internal bootdisk, 2 external compact flash slots
- 10/100/1000 management port with 802.1AE LinkSec
- Console & Auxiliary serial ports
- USB ports for file transfer
- Blue beacon LED for easy identification

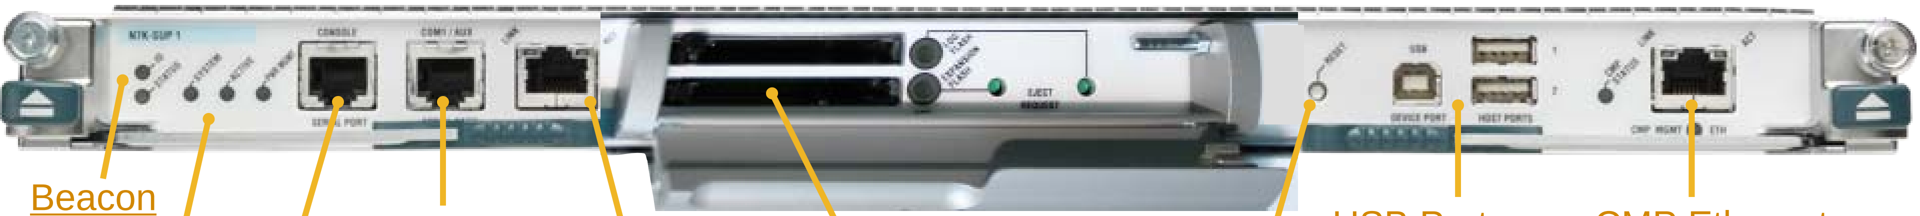

Beacon LED

Status LEDs

Console Port

AUX Port

Management Ethernet

Compact Flash Slots

Reset Button

USB Ports

CMP Ethernet

Connectivity Management Processor (CMP)

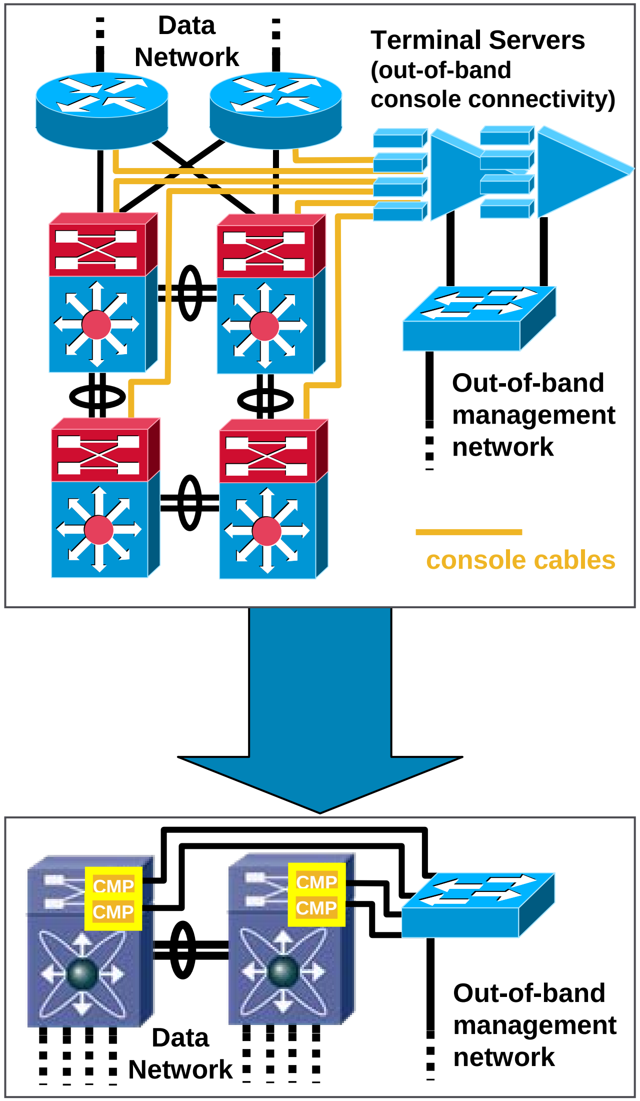

- Standalone, always-on microprocessor on supervisor engine
- Provides 'lights out' remote management and disaster recovery via 10/100/1000 interface
  - Removes need for terminal servers
- Monitor supervisor and modules, access log files, power cycle supervisor, etc.
  - Runs lightweight Linux kernel and network stack
  - Completely independent of DC-OS on main CPU

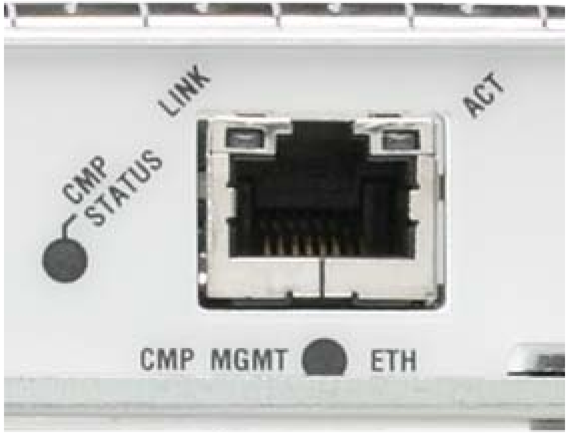

Supervisor Engine Architecture

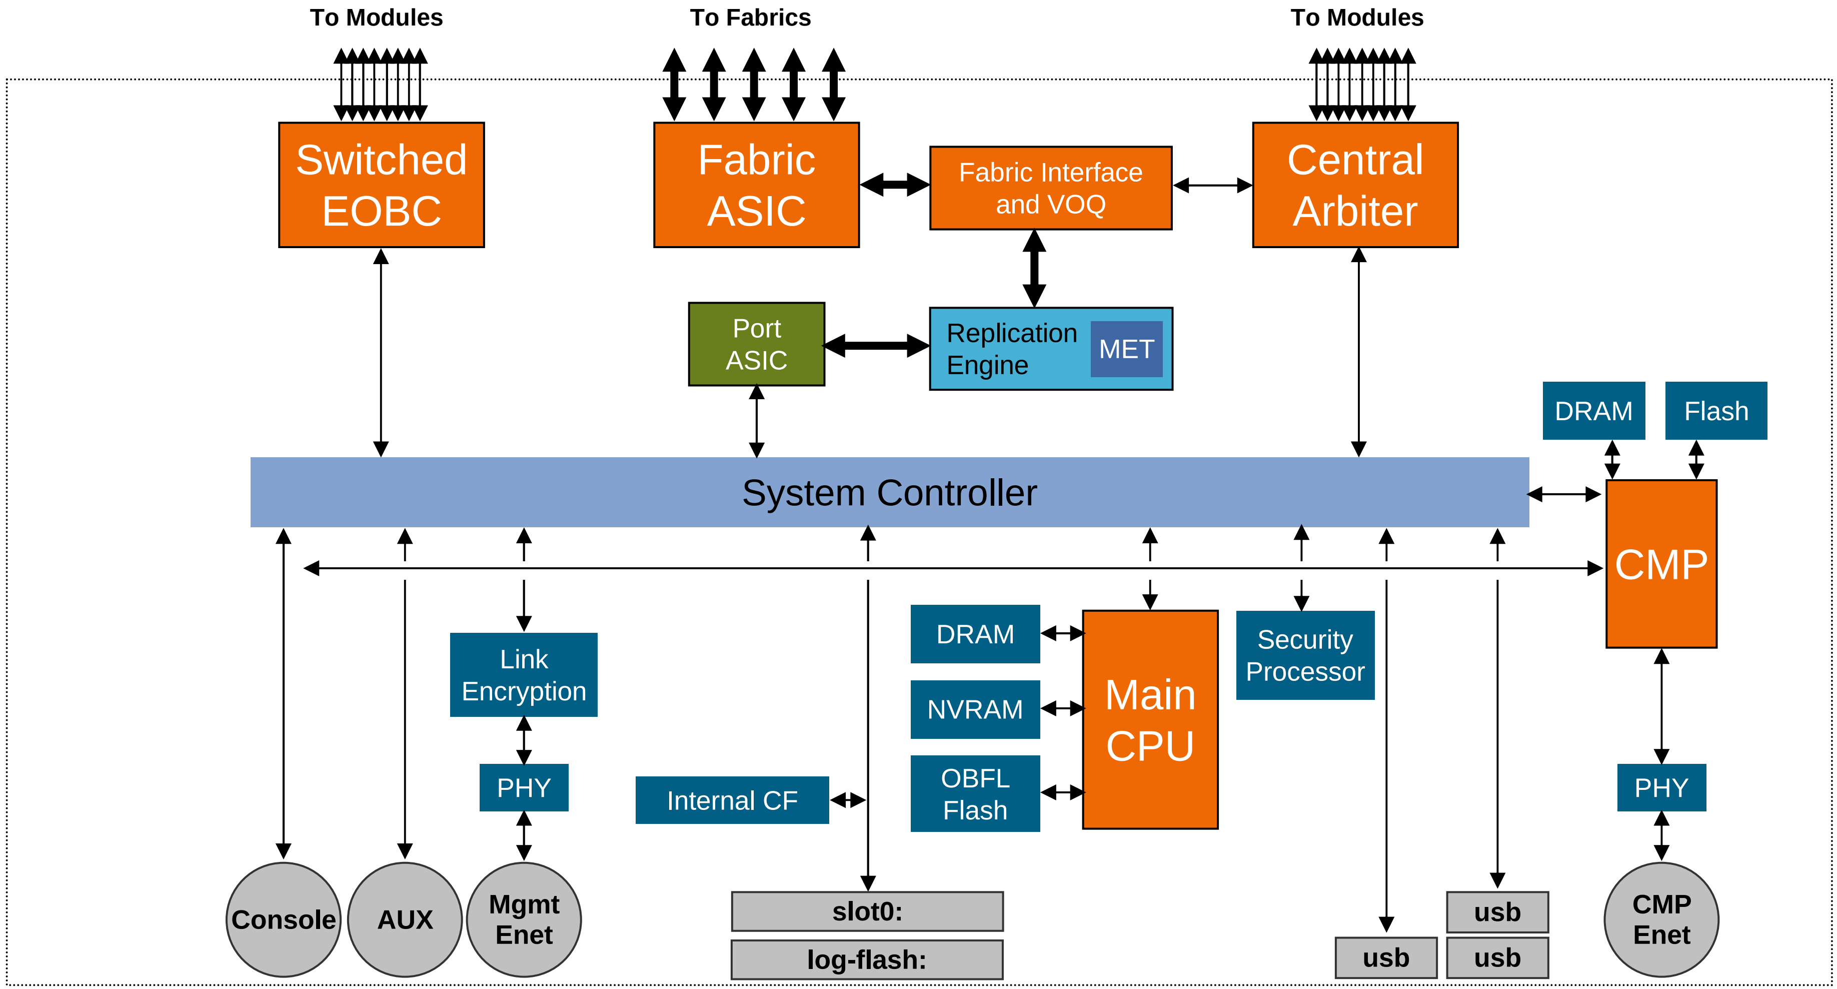

Nexus 7000 I/O Modules

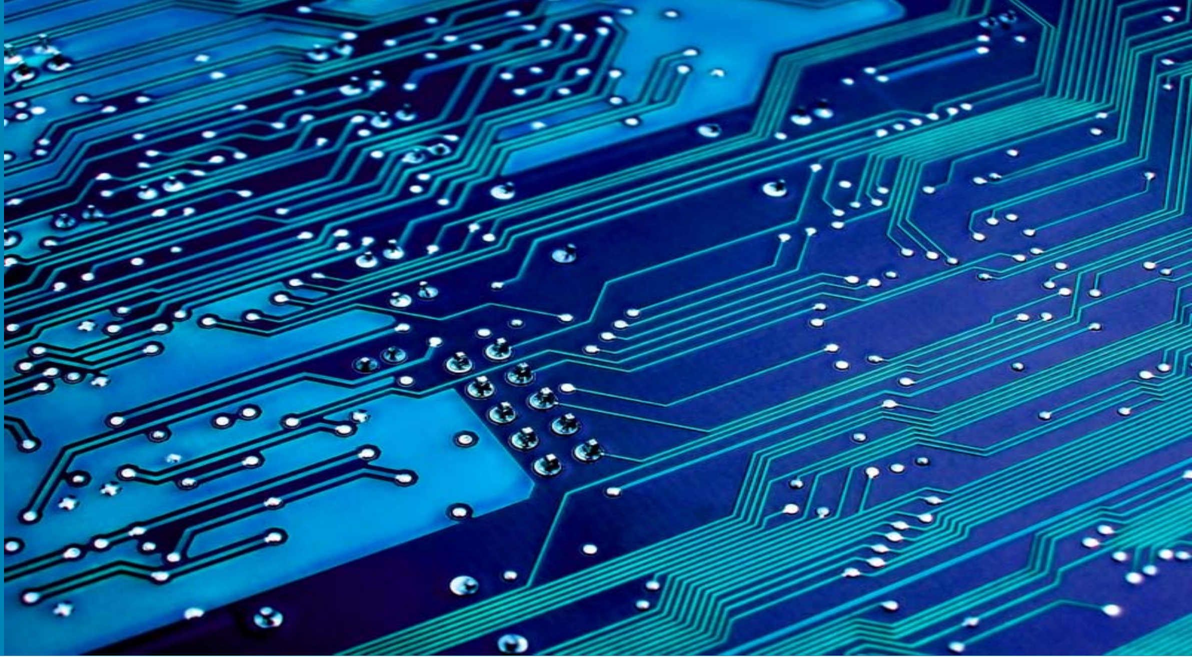

32-Port 10GE I/O Module

- 32 10GE ports with SFP+ transceivers
- 80G full duplex fabric connectivity
- Integrated 60Mpps forwarding engine for fully distributed forwarding
- 4:1 oversubscription at front panel

- Virtual output queueing (VOQ) ensuring fair access to fabric bandwidth
- 802.1AE LinkSec on every port
- Buffering:

  Dedicated mode: 100MB ingress, 80MB egress

  Shared mode: 1MB + 100MB ingress, 80MB egress
- Queues: 8q2t ingress, 1p7q4t egress
- Blue beacon LED for easy identification

SFP+

SR at initial release – 300m over MMF
LR post-release – 10km over SMF

Shared versus Dedicated Mode

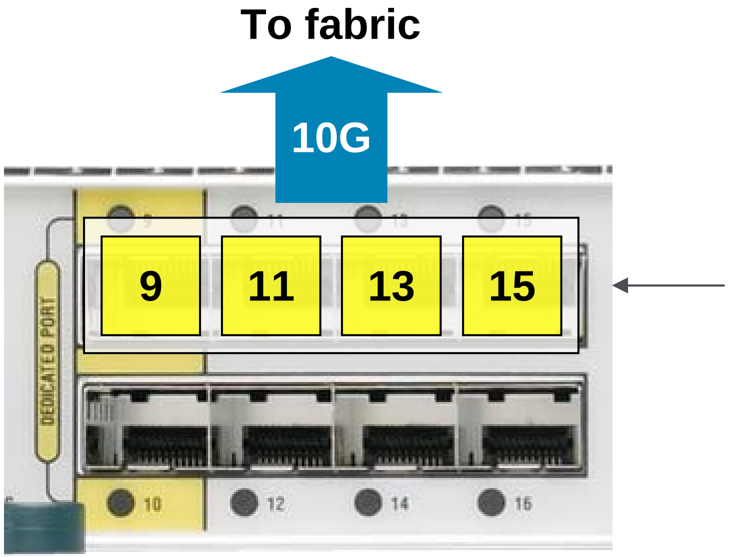

Shared mode

- Four interfaces share 10G bandwidth

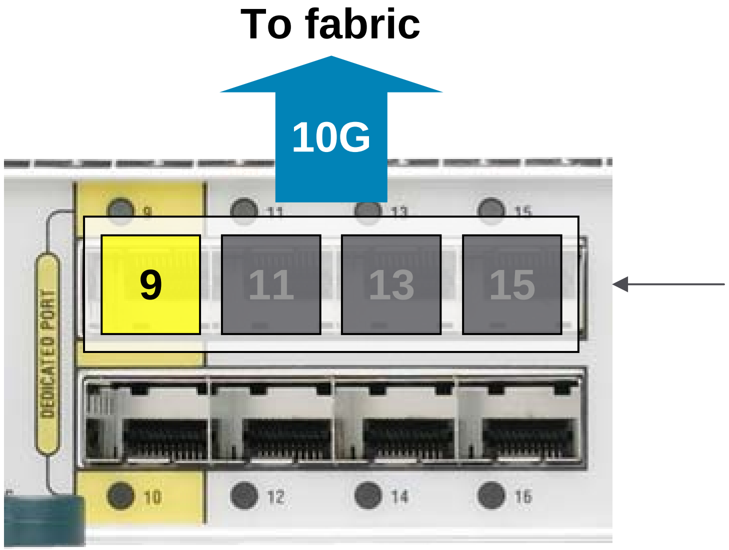

Dedicated mode

- One interface gets 10G bandwidth
- Three interfaces disabled

32-Port 10GE I/O Module Architecture

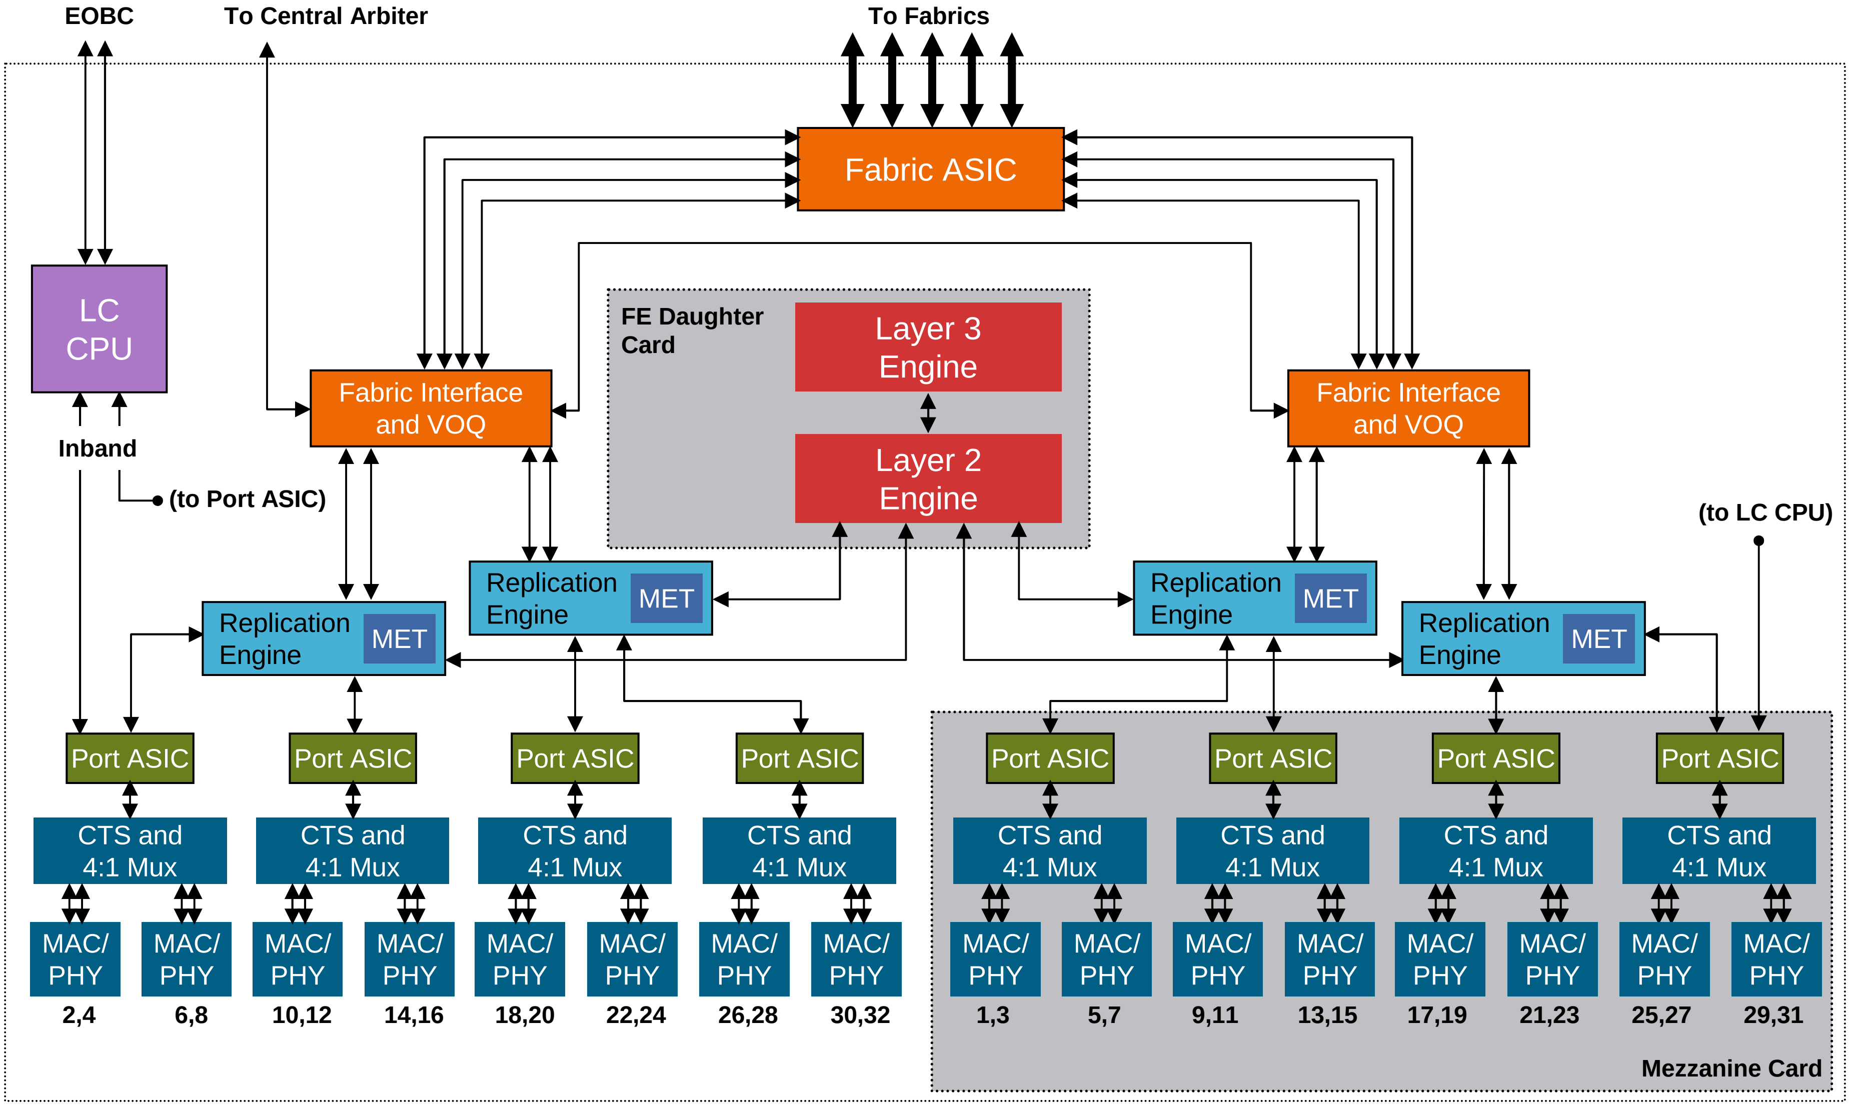

48-Port 1GE I/O Module

- 48 1GE 10/100/1000 RJ-45 ports
- 40G full duplex fabric connectivity
- Integrated 60Mpps forwarding engine for fully distributed forwarding
- Virtual output queueing (VOQ) ensuring fair access to fabric bandwidth
- 802.1AE LinkSec on every port
- Buffer: 7.5MB ingress, 6.2MB egress
- Queues: 2q4t ingress, 1p3q4t egress
- Blue beacon LED for easy identification

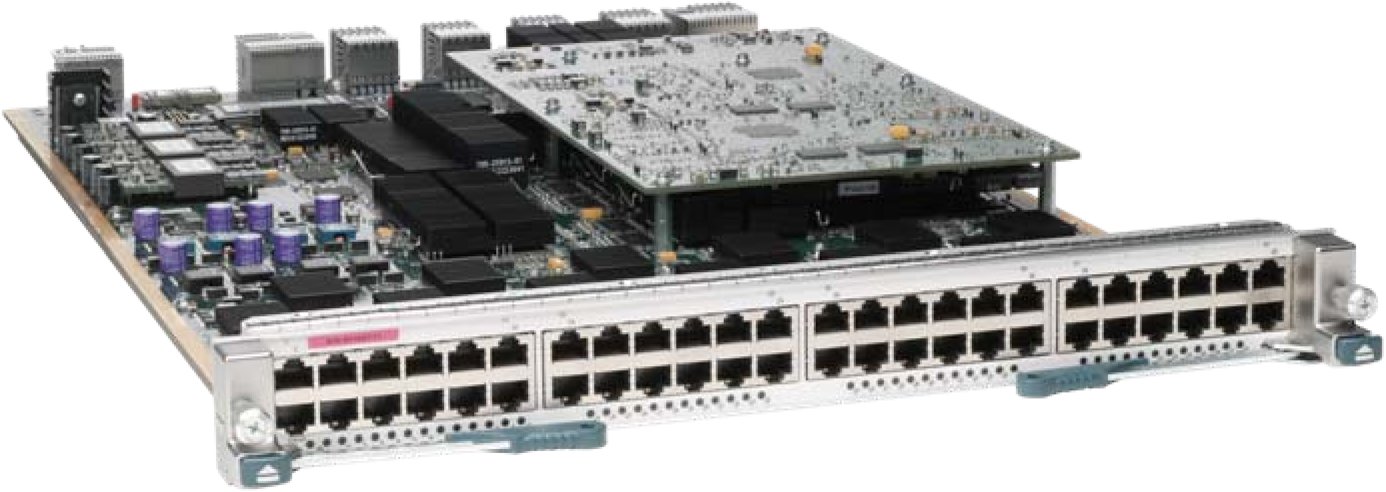

48-Port 1GE I/O Module Architecture

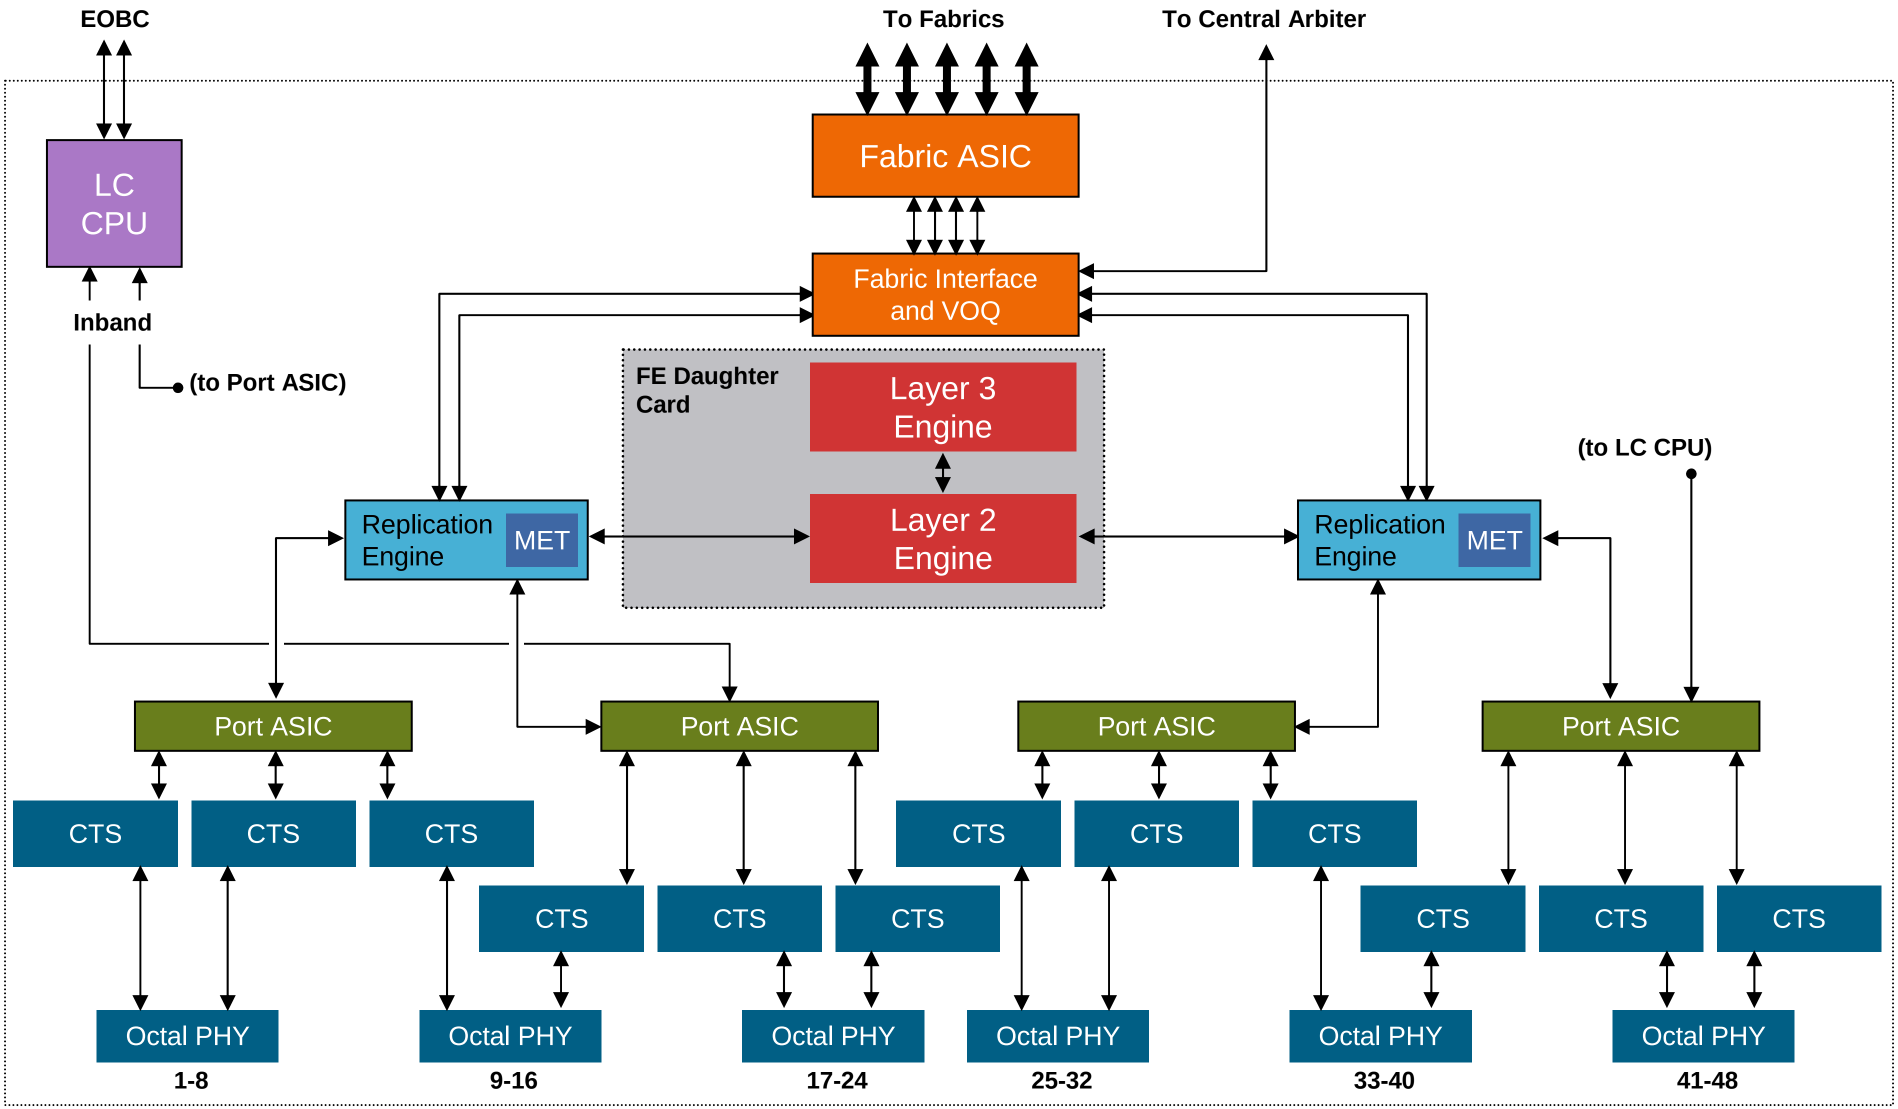

Nexus 7000 Forwarding Engine

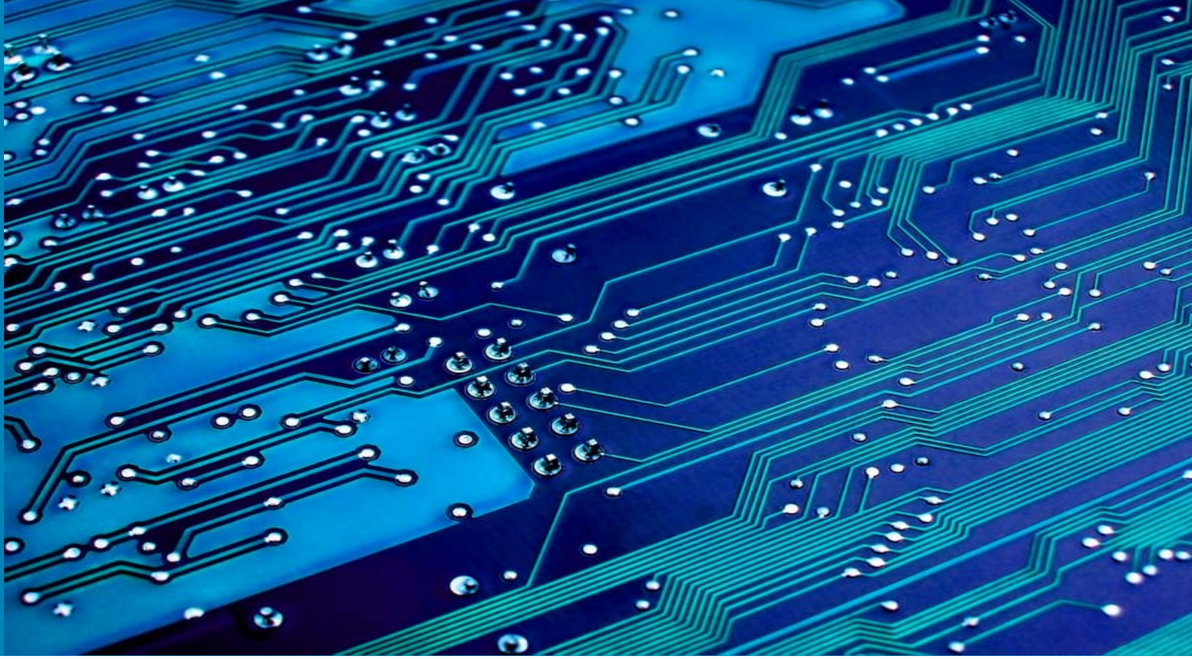

Forwarding Engine Hardware

Advanced hardware forwarding engine integrated on every I/O module

- 60Mpps Layer 2 bridging with hardware MAC learning
- 60Mpps IPv4 and 30Mpps IPv6 unicast
- IPv4 and IPv6 multicast (SM, SSM, bidir)
- IPv4 and IPv6 security ACLs
- Cisco TrustSec security group tag support
- Unicast RPF check and IP source guard
- QoS remarking and policing policies
- Ingress and egress NetFlow (full and sampled)
- GRE tunnels

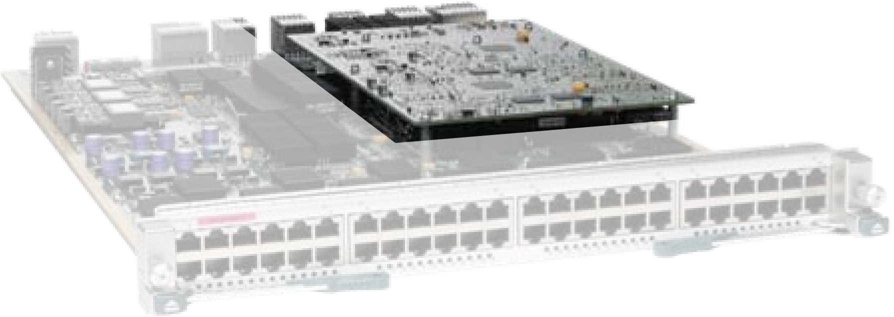

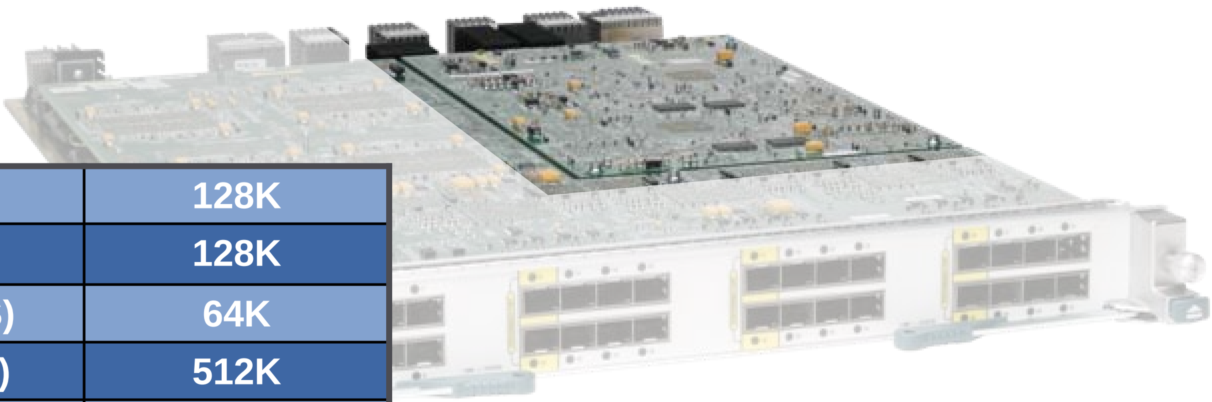

Table sizes optimized for Data Center

FIB TCAM	**128K**
MAC table	**128K**
Classification TCAM (ACL and QoS)	**64K**
NetFlow Table (Ingress and Egress)	**512K**
Policers	**16K**

Forwarding Engine Details

Forwarding engine chipset consists of two ASICs:

- Layer 2 Engine
    - Performs ingress and egress SMAC/DMAC lookups
    - Hardware MAC learning
    - True IP-based Layer 2 multicast constraint
    - Performs lookups on ingress I/O module, and egress I/O module for bridged packets
- Layer 3 Engine
    - 60Mpps IPv4 and 30Mpps IPv6 Layer 3/Layer 4 lookups
    - Performs all FIB, ACL, QoS, NetFlow processing
    - Linear, pipelined architecture – every packet processed in ingress and egress pipe
    - Performs lookups on ingress I/O module, and egress I/O module for multicast replicated packets

Nexus 7000 Fabric and Bandwidth

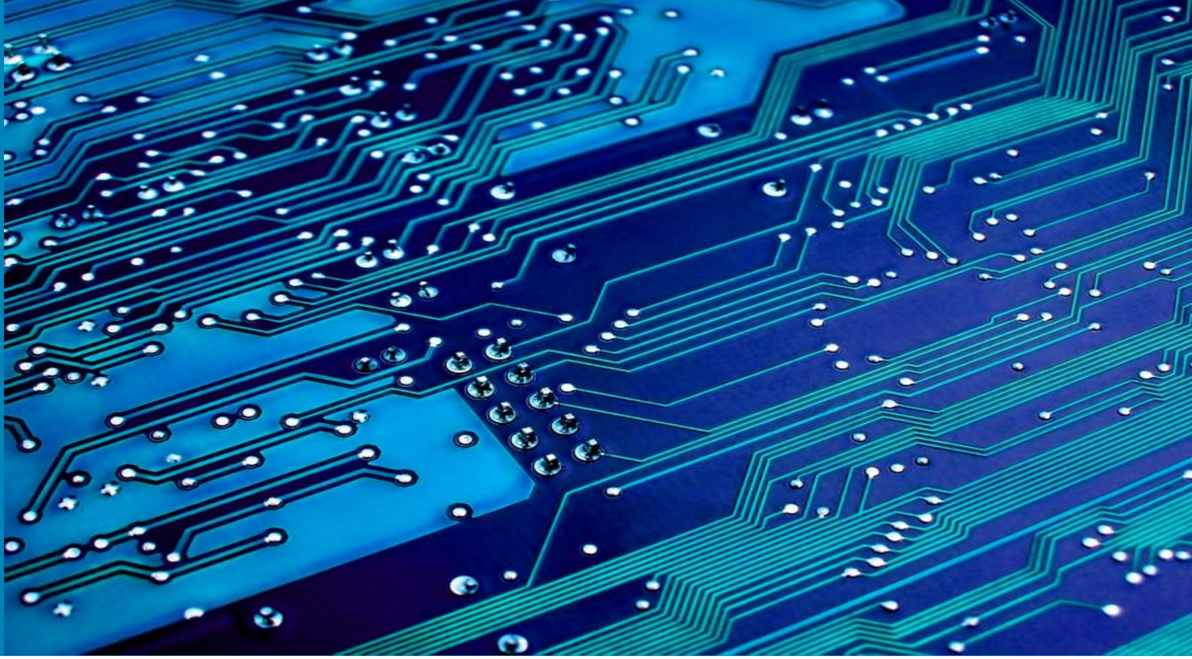

I/O Module Bandwidth Capacity

- Initially shipping I/O module bandwidth: 80Gbps per slot

  Assumes 8 * 10G ports in dedicated mode per module

- In Nexus 7000 10-slot chassis:

  (80Gbps/slot) * (8 payload slots) = 640Gbps

  (640Gbps) * (2 for full duplex operation) = 1280Gbps = 1.2Tbps system bandwidth

1.2 Terabits per second initial system bandwidth

Fabric Bandwidth Capacity

- Initially shipping fabric bandwidth: 230Gbps per payload slot, 115Gbps per supervisor slot
  - Initially shipping modules cannot fully leverage fabric bandwidth
  - Assumes future modules that can leverage full bandwidth
- In Nexus 7000 10-slot chassis:
  - (230Gbps/slot) * (8 payload slots) = 1840Gbps
  - (115Gbps/slot) * (2 supervisor slots) = 230Gbps
  - (1840 + 230 = 2070Gbps) * (2 for full duplex operation) = 4140Gbps = 4.1Tbps system bandwidth

4.1 Terabits per second fabric bandwidth capacity

Future Vision for Platform Series

- Future goal to double fabric bandwidth

  500+Gbps bandwidth per slot

  Requires future fabric module

- 10 slot chassis will scale to 9+Tbps system bandwidth
- 18 slot chassis will scale to 15+Tbps system bandwidth

15+ Terabits per second platform bandwidth capacity

Fabric Module

- Provides 46Gbps per I/O module slot

  Also provides 23G per supervisor slot

- Up to 230Gbps per slot with 5 fabric modules

  Initially shipping I/O modules do not leverage full fabric bandwidth

- Load-sharing across all fabric modules in chassis
- Multilevel redundancy with graceful performance degradation
- Non-disruptive OIR
- Blue beacon LED for easy identification

Fabric Capacity and Redundancy

- Per-slot bandwidth capacity increases with each fabric module
- 1G module requires 2 fabrics for N+1 redundancy
- 10G module requires 3 fabrics for N+1 redundancy
- 4th and 5th fabric modules provide additional level of redundancy
- Future modules will leverage additional fabric bandwidth
- Fabric failure results in reduction of overall system bandwidth

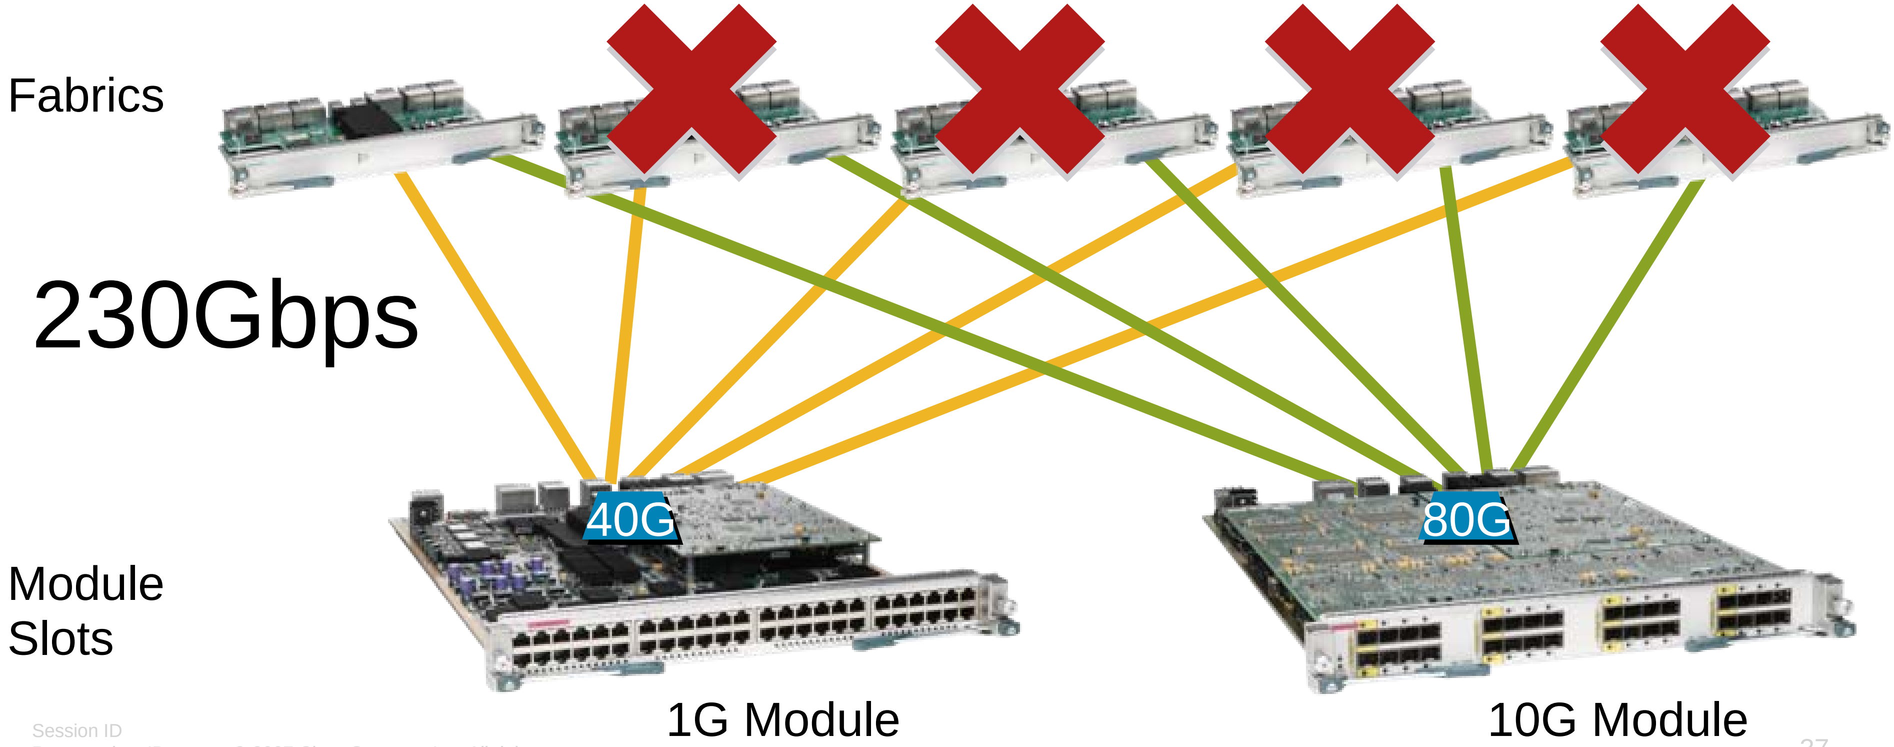

Access to Fabric Bandwidth

- Supervisor engine controls access to fabric bandwidth using central arbitration
- Fabric bandwidth represented by Virtual Output Queues (VOQs)

What Are VOQs?

- Virtual Output Queues (VOQs) on ingress modules represent bandwidth capacity on egress modules
- Guaranteed delivery to egress module for arbitrated packets entering fabric
  - If VOQ available on ingress, capacity exists on egress
- VOQ is NOT equivalent to ingress or egress port buffer or queues
  - Relates ONLY to ASICs at ingress and egress to fabric
- VOQ is “virtual” because it represents EGRESS capacity but resides on INGRESS module
  - It is PHYSICAL buffer where packets are stored

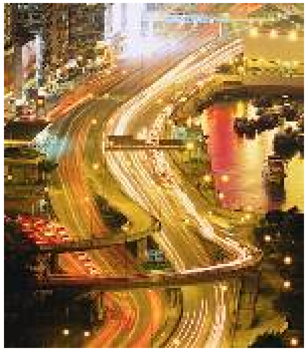

What Is VOQ?

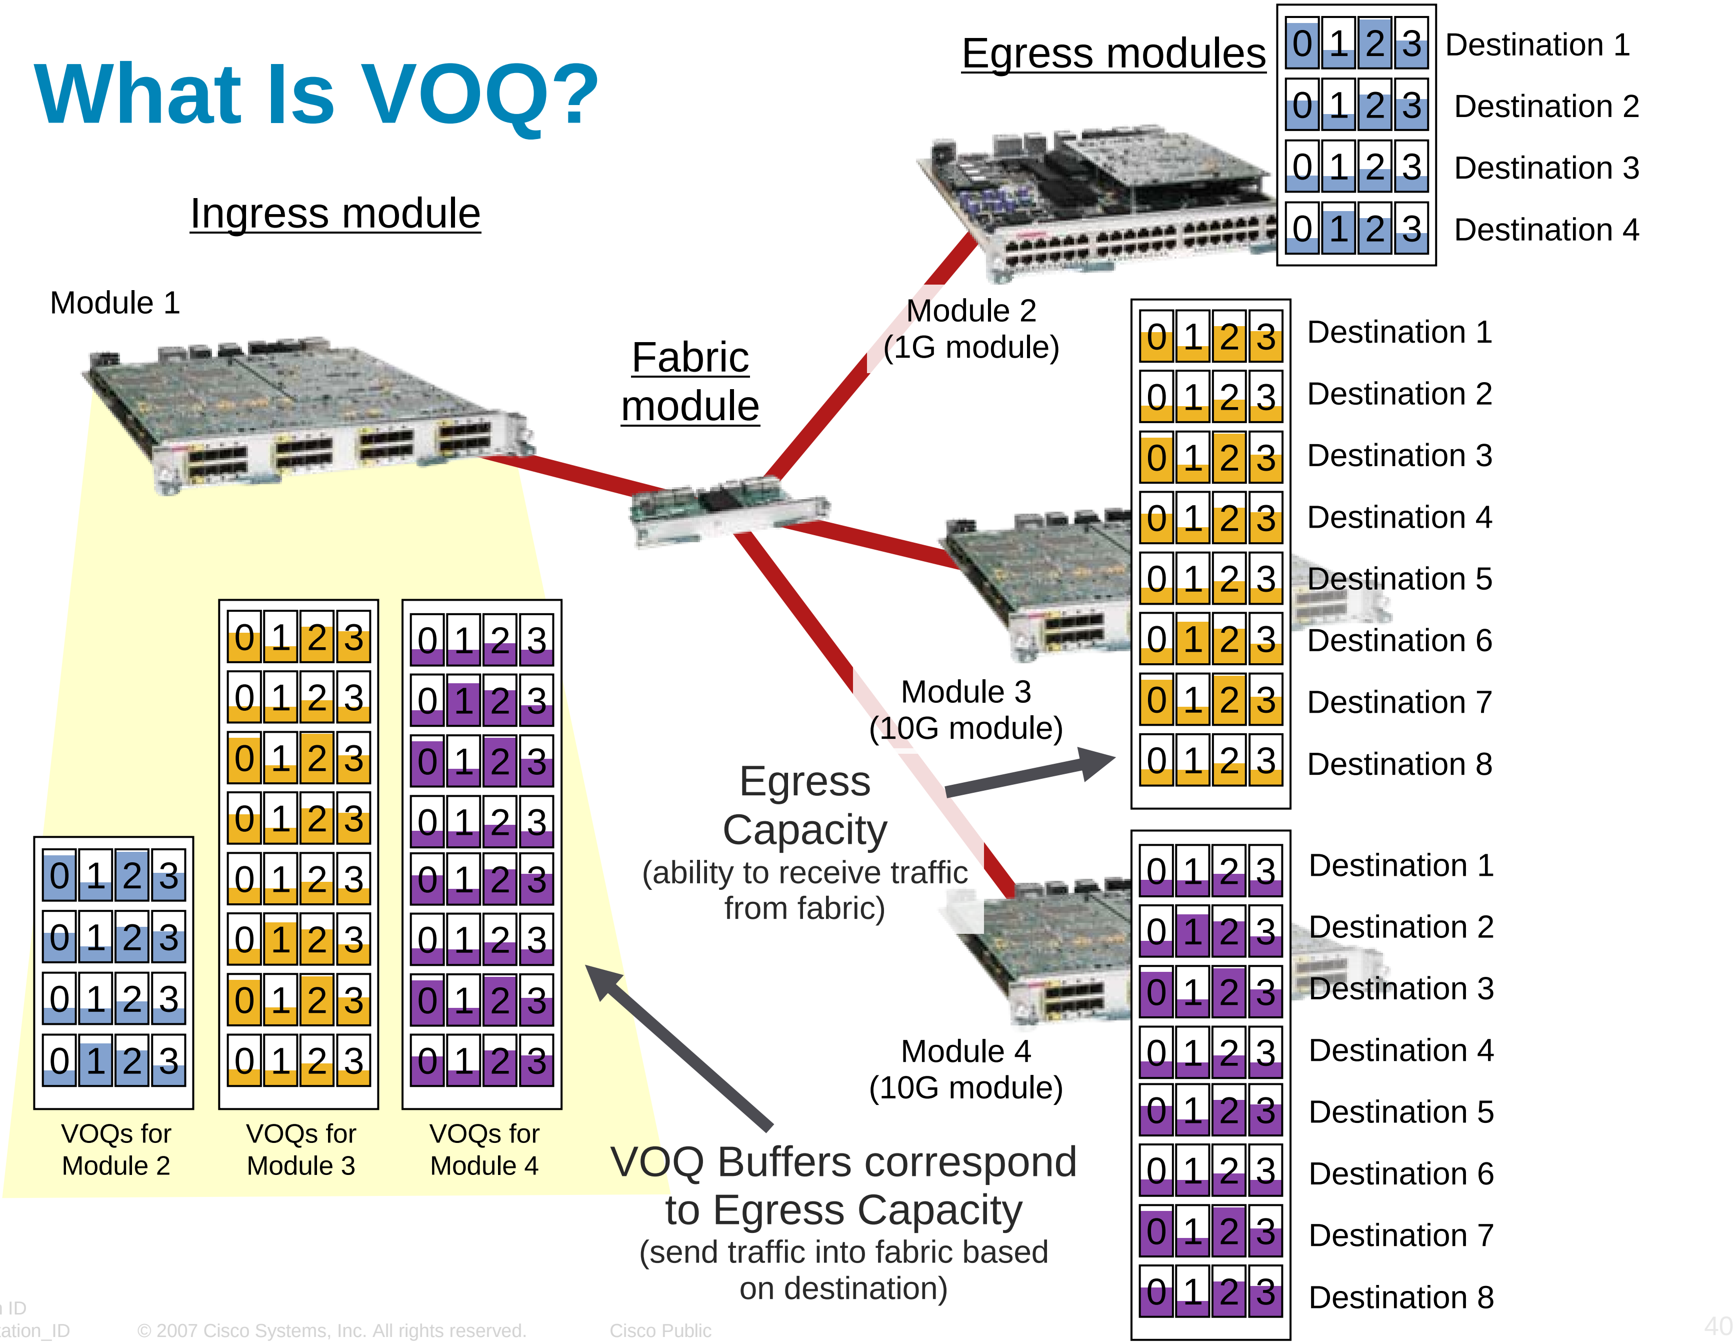

Ingress module

Module 1

Fabric module

Egress modules

Module 2 (1G module)

Destination 1
Destination 2
Destination 3
Destination 4

Module 3 (10G module)

Destination 1
Destination 2
Destination 3
Destination 4
Destination 5
Destination 6
Destination 7
Destination 8

Module 4 (10G module)

Destination 1
Destination 2
Destination 3
Destination 4
Destination 5
Destination 6
Destination 7
Destination 8

VOQs for Module 2

VOQs for Module 3

VOQs for Module 4

Egress Capacity
(ability to receive traffic from fabric)

VOQ Buffers correspond to Egress Capacity
(send traffic into fabric based on destination)

Centralized Fabric Arbitration

- Access to fabric bandwidth on ingress module controlled by central arbiter on supervisor
  - In other words, access to the VOQ for the destination across the fabric
- Arbitration works on credit request/grant basis
  - Modules communicate egress fabric buffer availability to central arbiter
  - Modules request credits from supervisor to place packets in VOQ for transmission to destination over fabric
  - Supervisor grants credits based on egress fabric buffer availability for that destination
- Arbiter discriminates among four classes of service
  - Priority traffic takes precedence over best-effort traffic across fabric

VOQ Operation

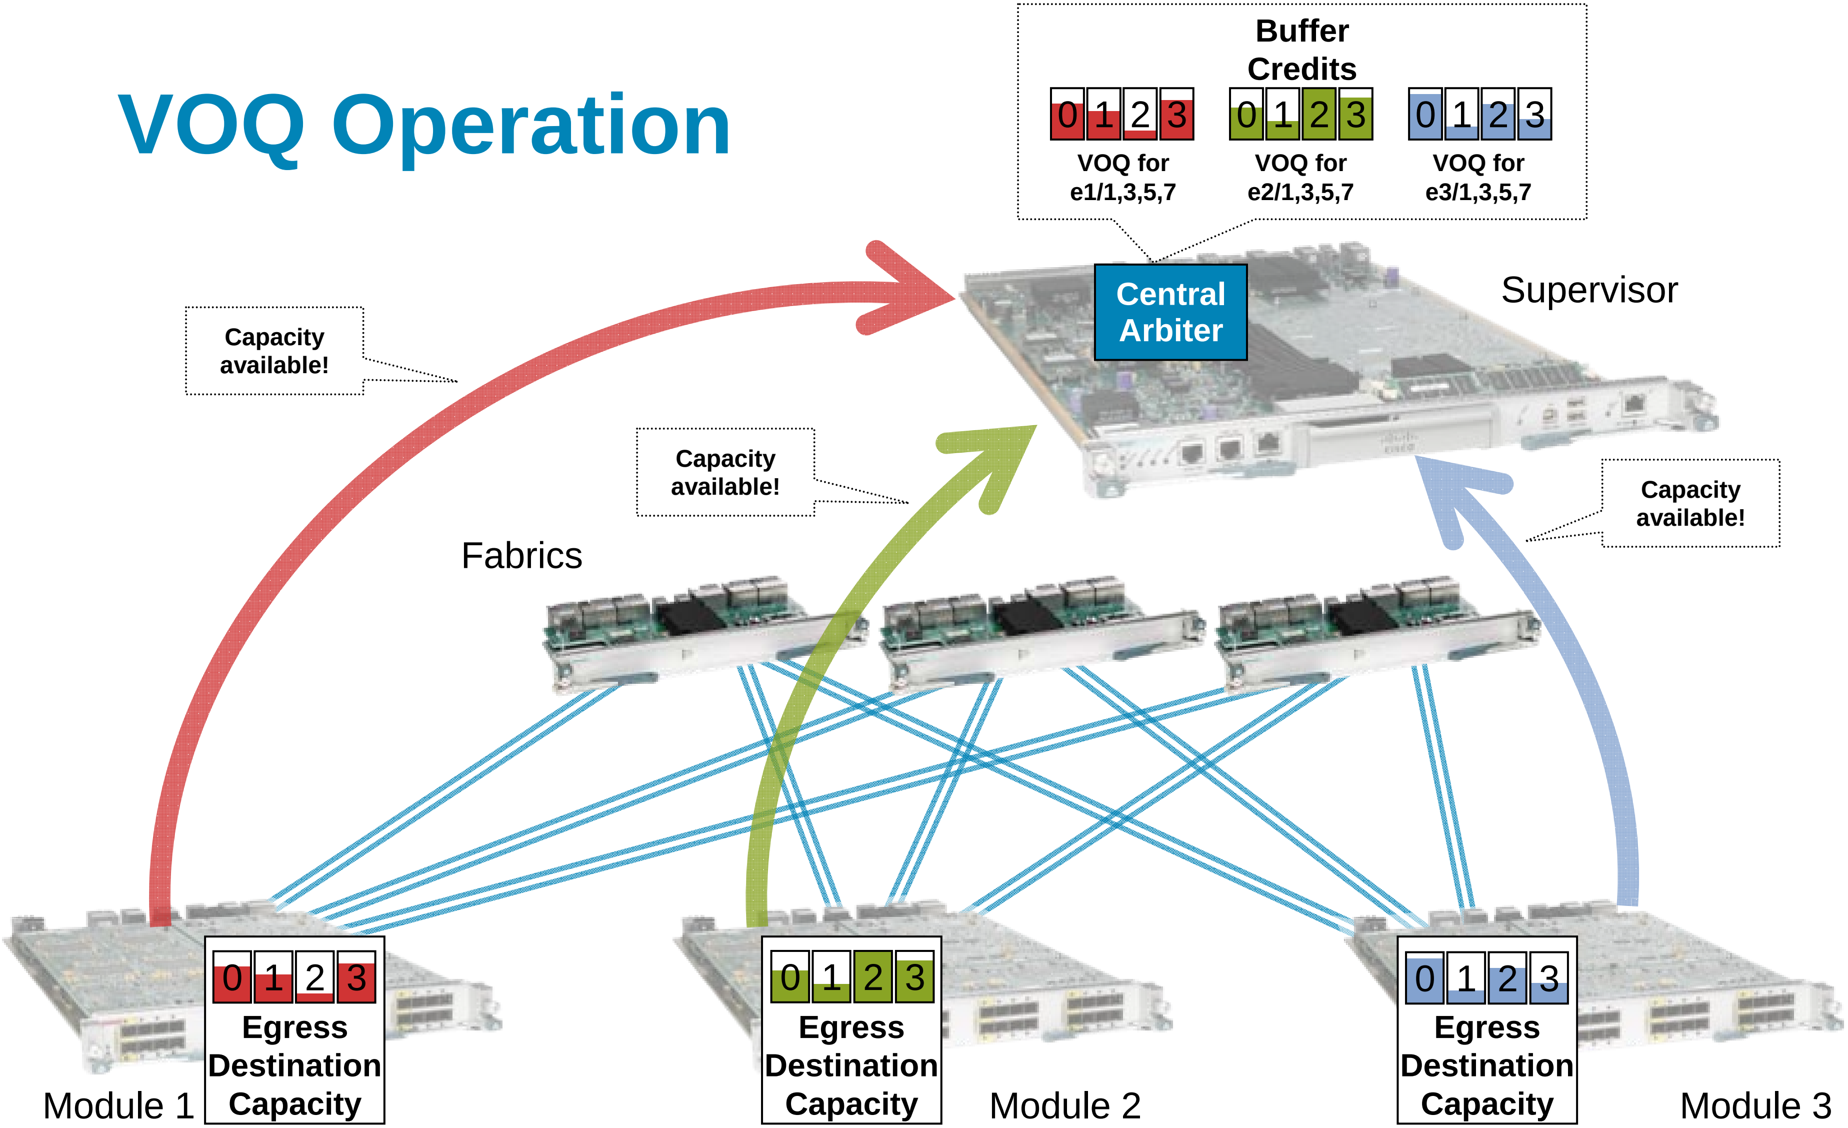

Egress modules have capacity to receive traffic from fabric

VOQ Operation

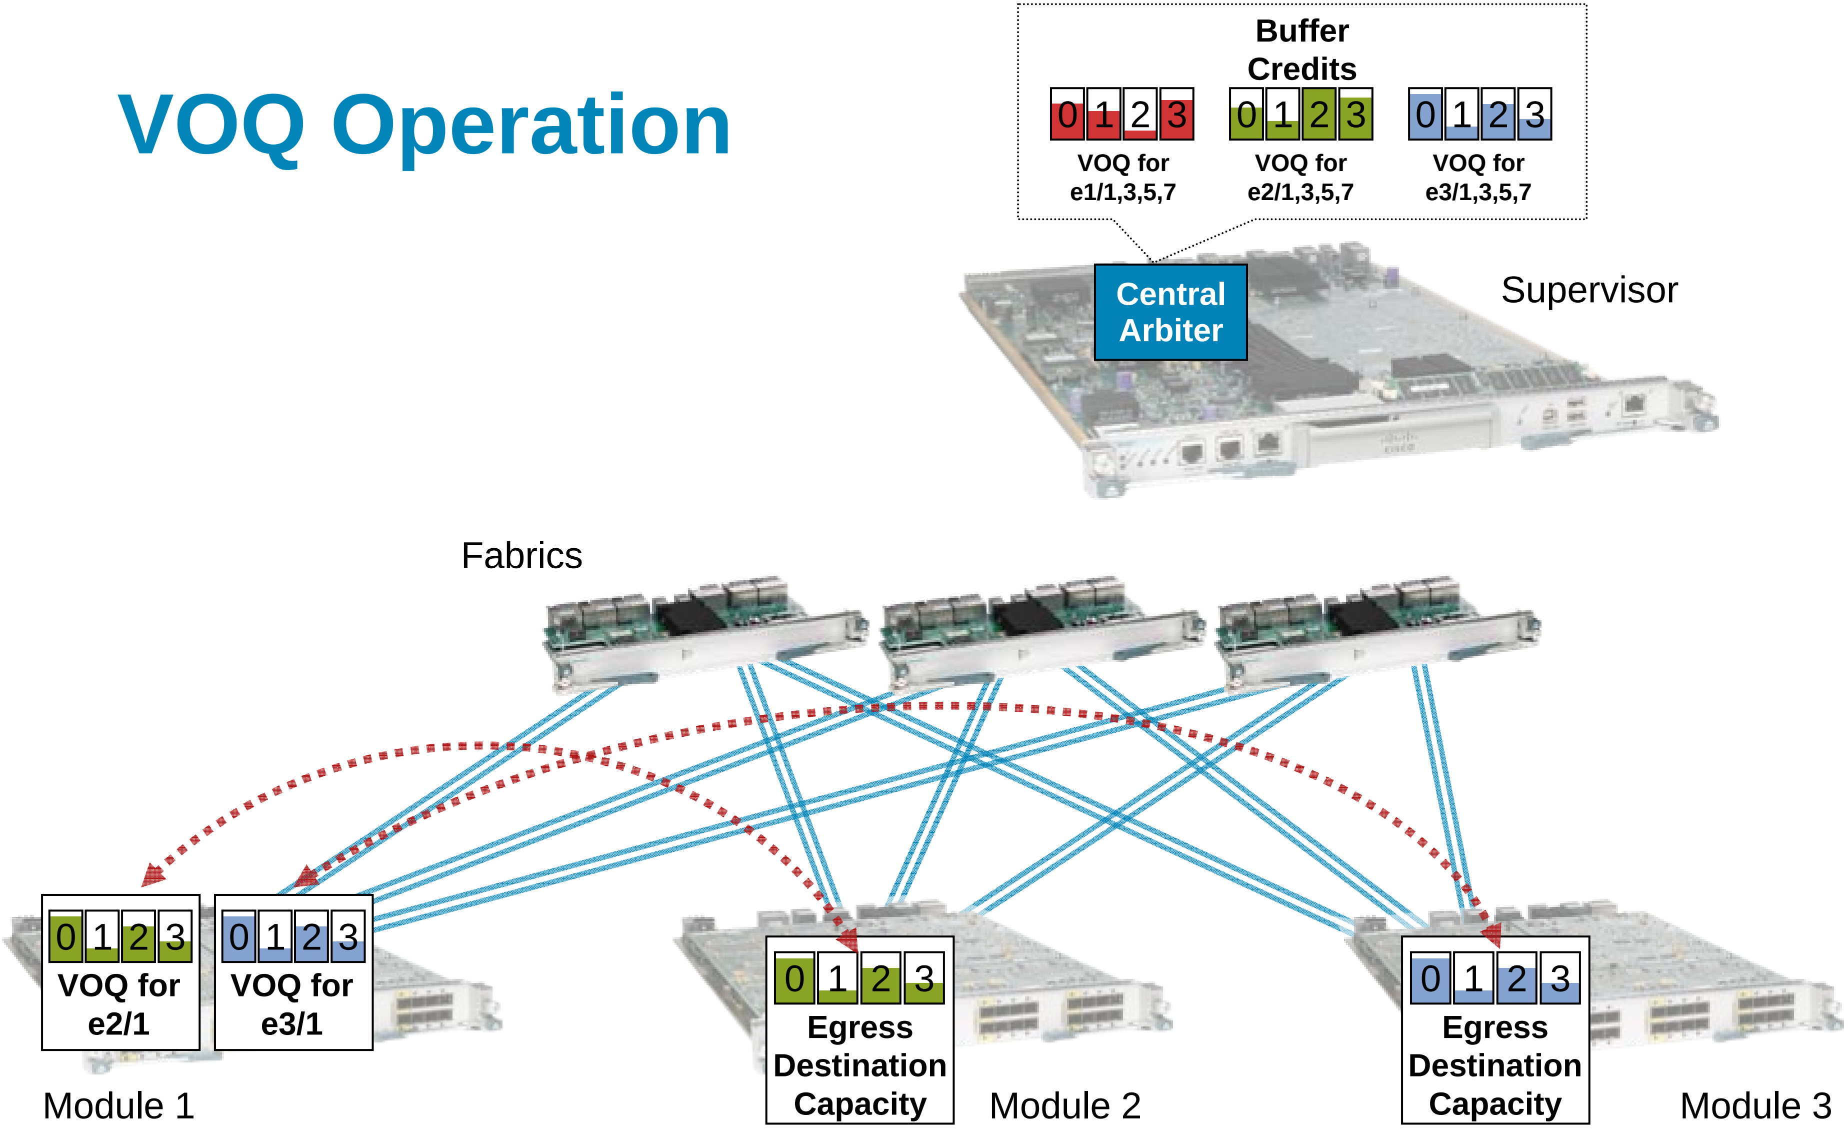

Buffer Credits

0	1	2	3

VOQ for e1/1,3,5,7

0	1	2	3

VOQ for e2/1,3,5,7

0	1	2	3

VOQ for e3/1,3,5,7

Central Arbiter

Supervisor

Fabrics

0	1	2	3

VOQ for e2/1

0	1	2	3

VOQ for e3/1

Module 1

0	1	2	3

Egress Destination Capacity

Module 2

0	1	2	3

Egress Destination Capacity

Module 3

VOQs on ingress module correspond to capacity on egress modules

INGRESS MODULE

EGRESS MODULES

VOQ Operation

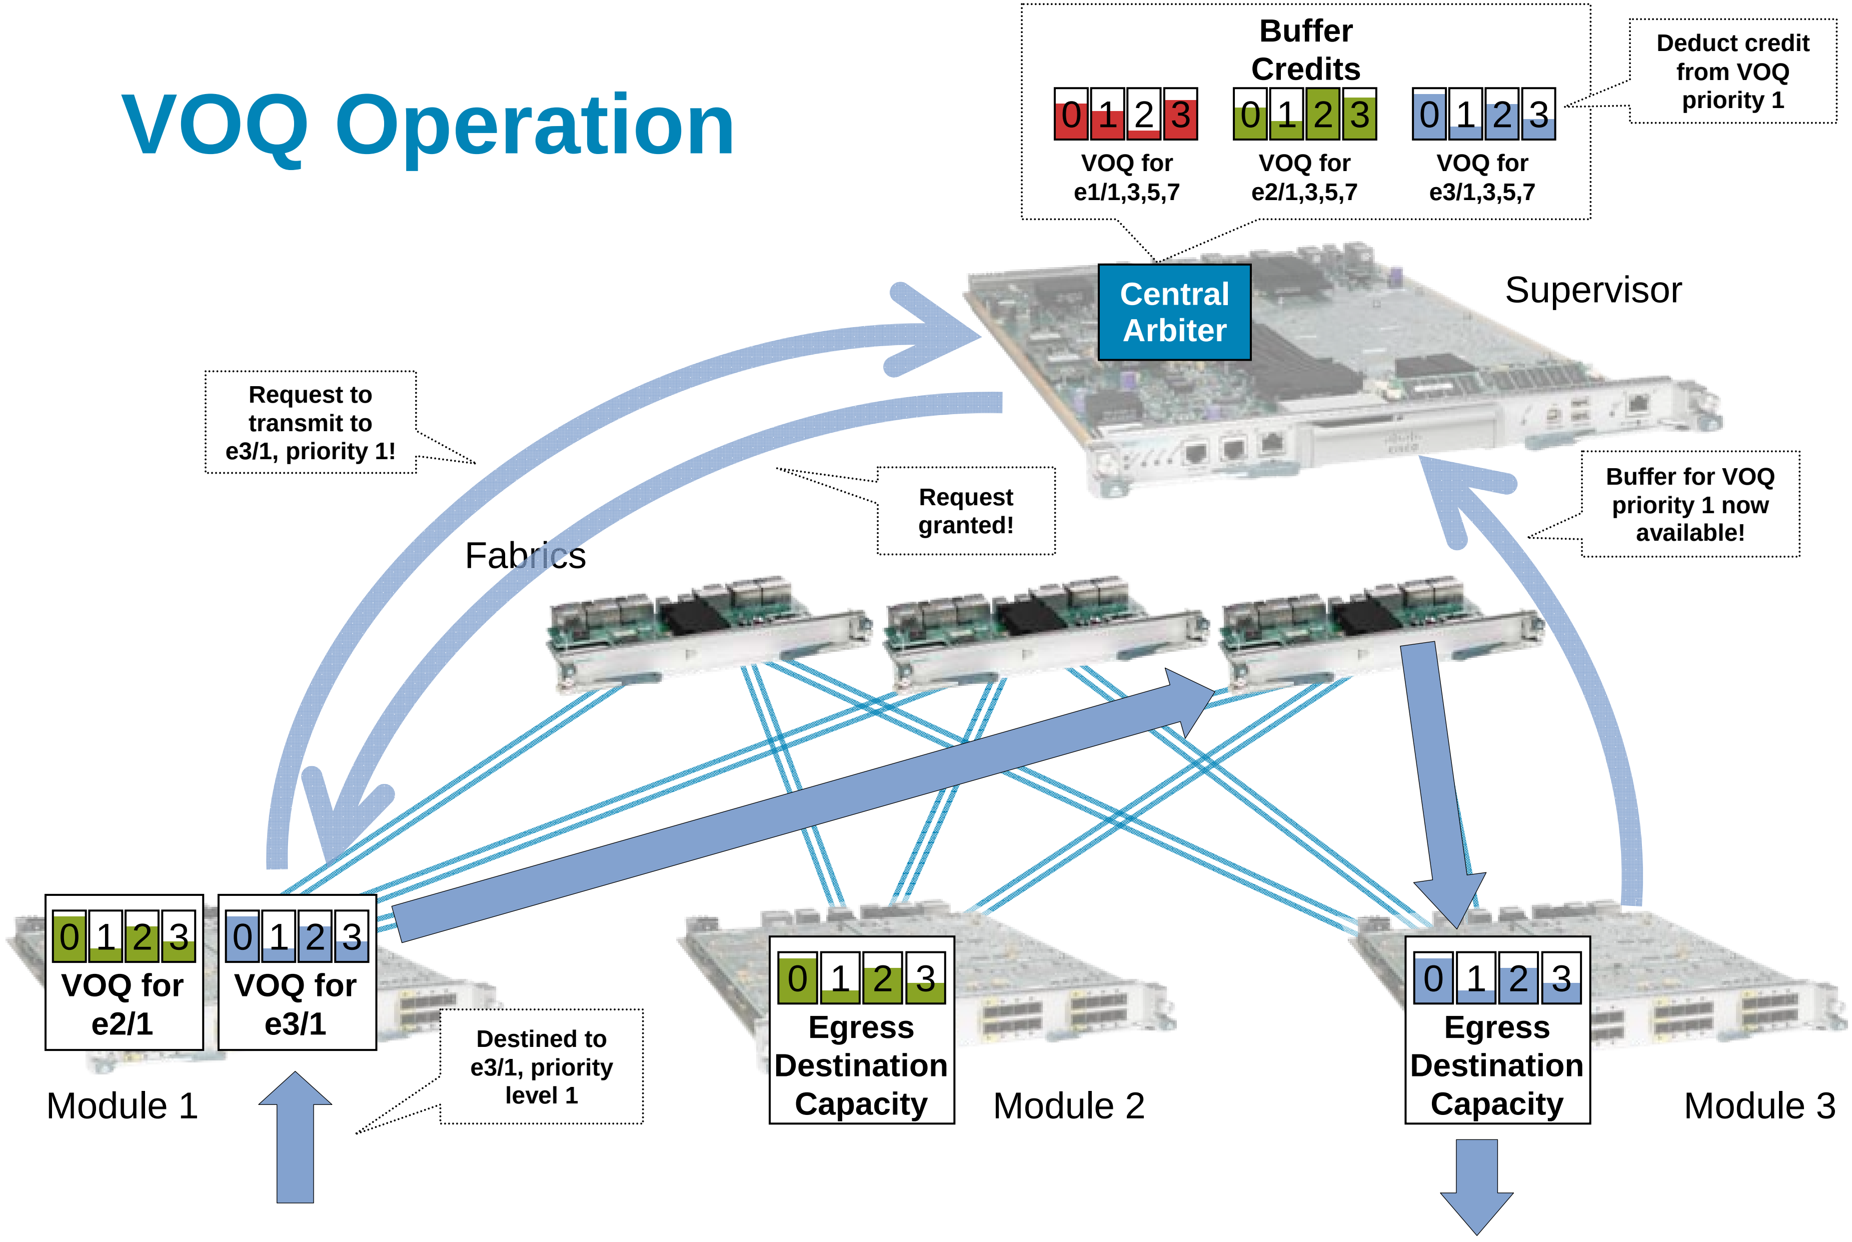

INGRESS MODULE

EGRESS MODULES

Benefits of Central Arbitration and VOQ

- Ensures fair access to bandwidth for multiple ingress ports transmitting to one egress port
- Prevents congested egress ports from blocking ingress traffic destined to other ports
- Priority traffic takes precedence over best-effort traffic across fabric
- Engineered to support Unified I/O

  Can provide no-drop service across fabric for future FCoE interfaces

Fabric Superframing

- Fabric interface ASIC performs superframing for fabric-bound unicast packets
- When packet hits VOQ, arbitration request generated immediately
- When grant returned, packet immediately transmitted into fabric
- When packet transmission complete, check if other packets are enqueued in the same VOQ – if so, transmit them
- Transmit until superframe limit reached, or queue empty
- Superframe up to 2560 bytes, or up to 32 packets, whichever comes first
- If additional packets need transmission, new arbitration request generated and new superframe begins
- Superframe disassembled on egress fabric interface ASIC

Fabric Superframing

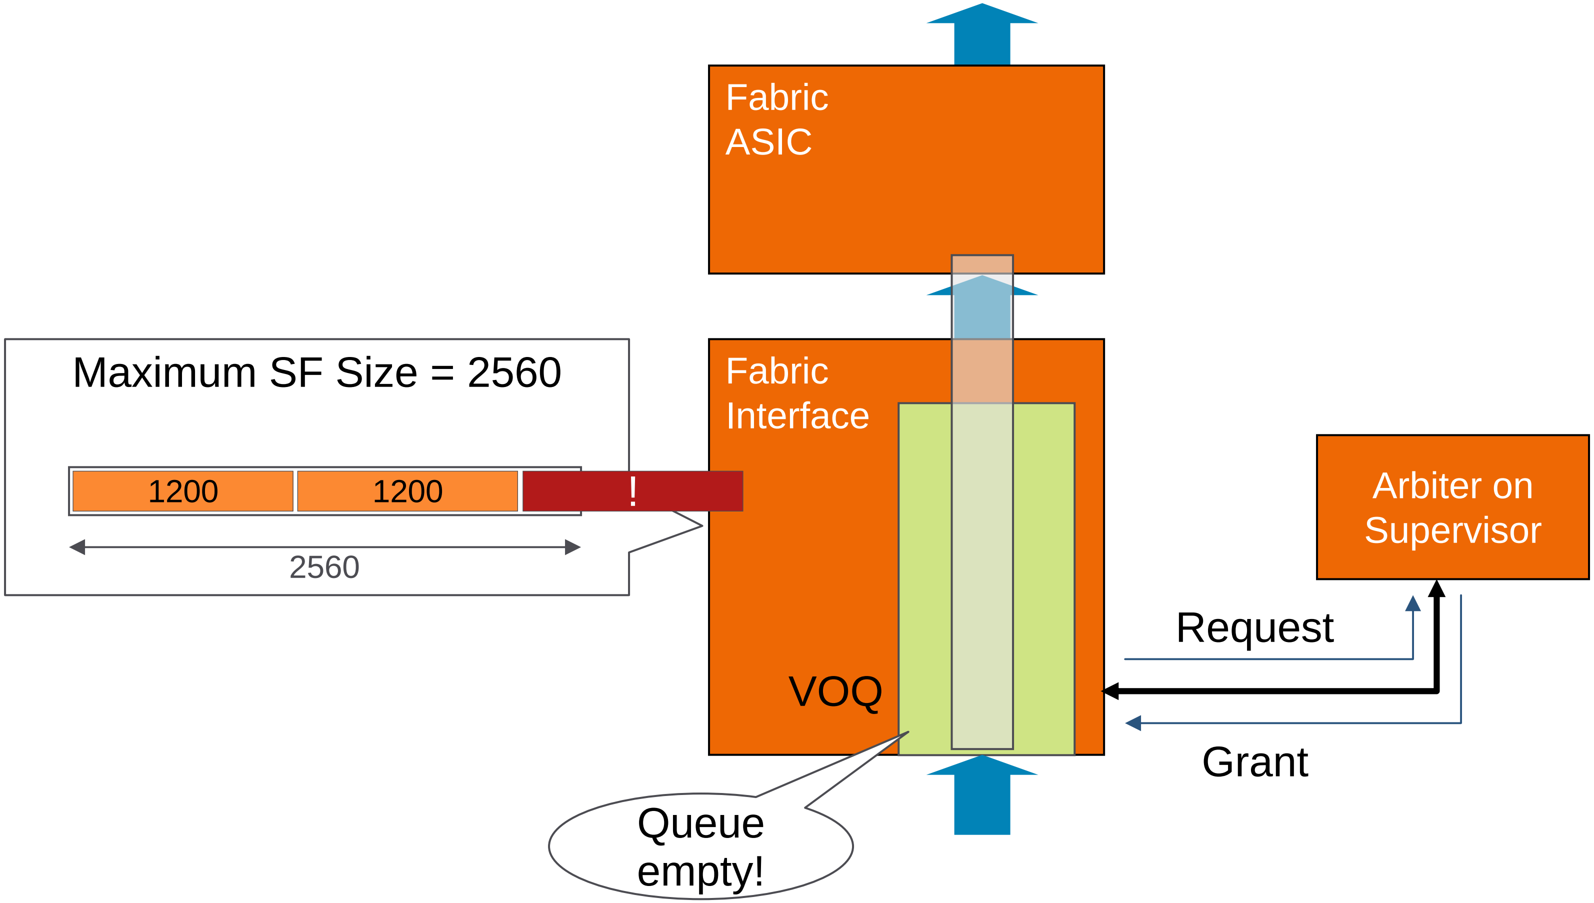

Nexus 7000 Pack

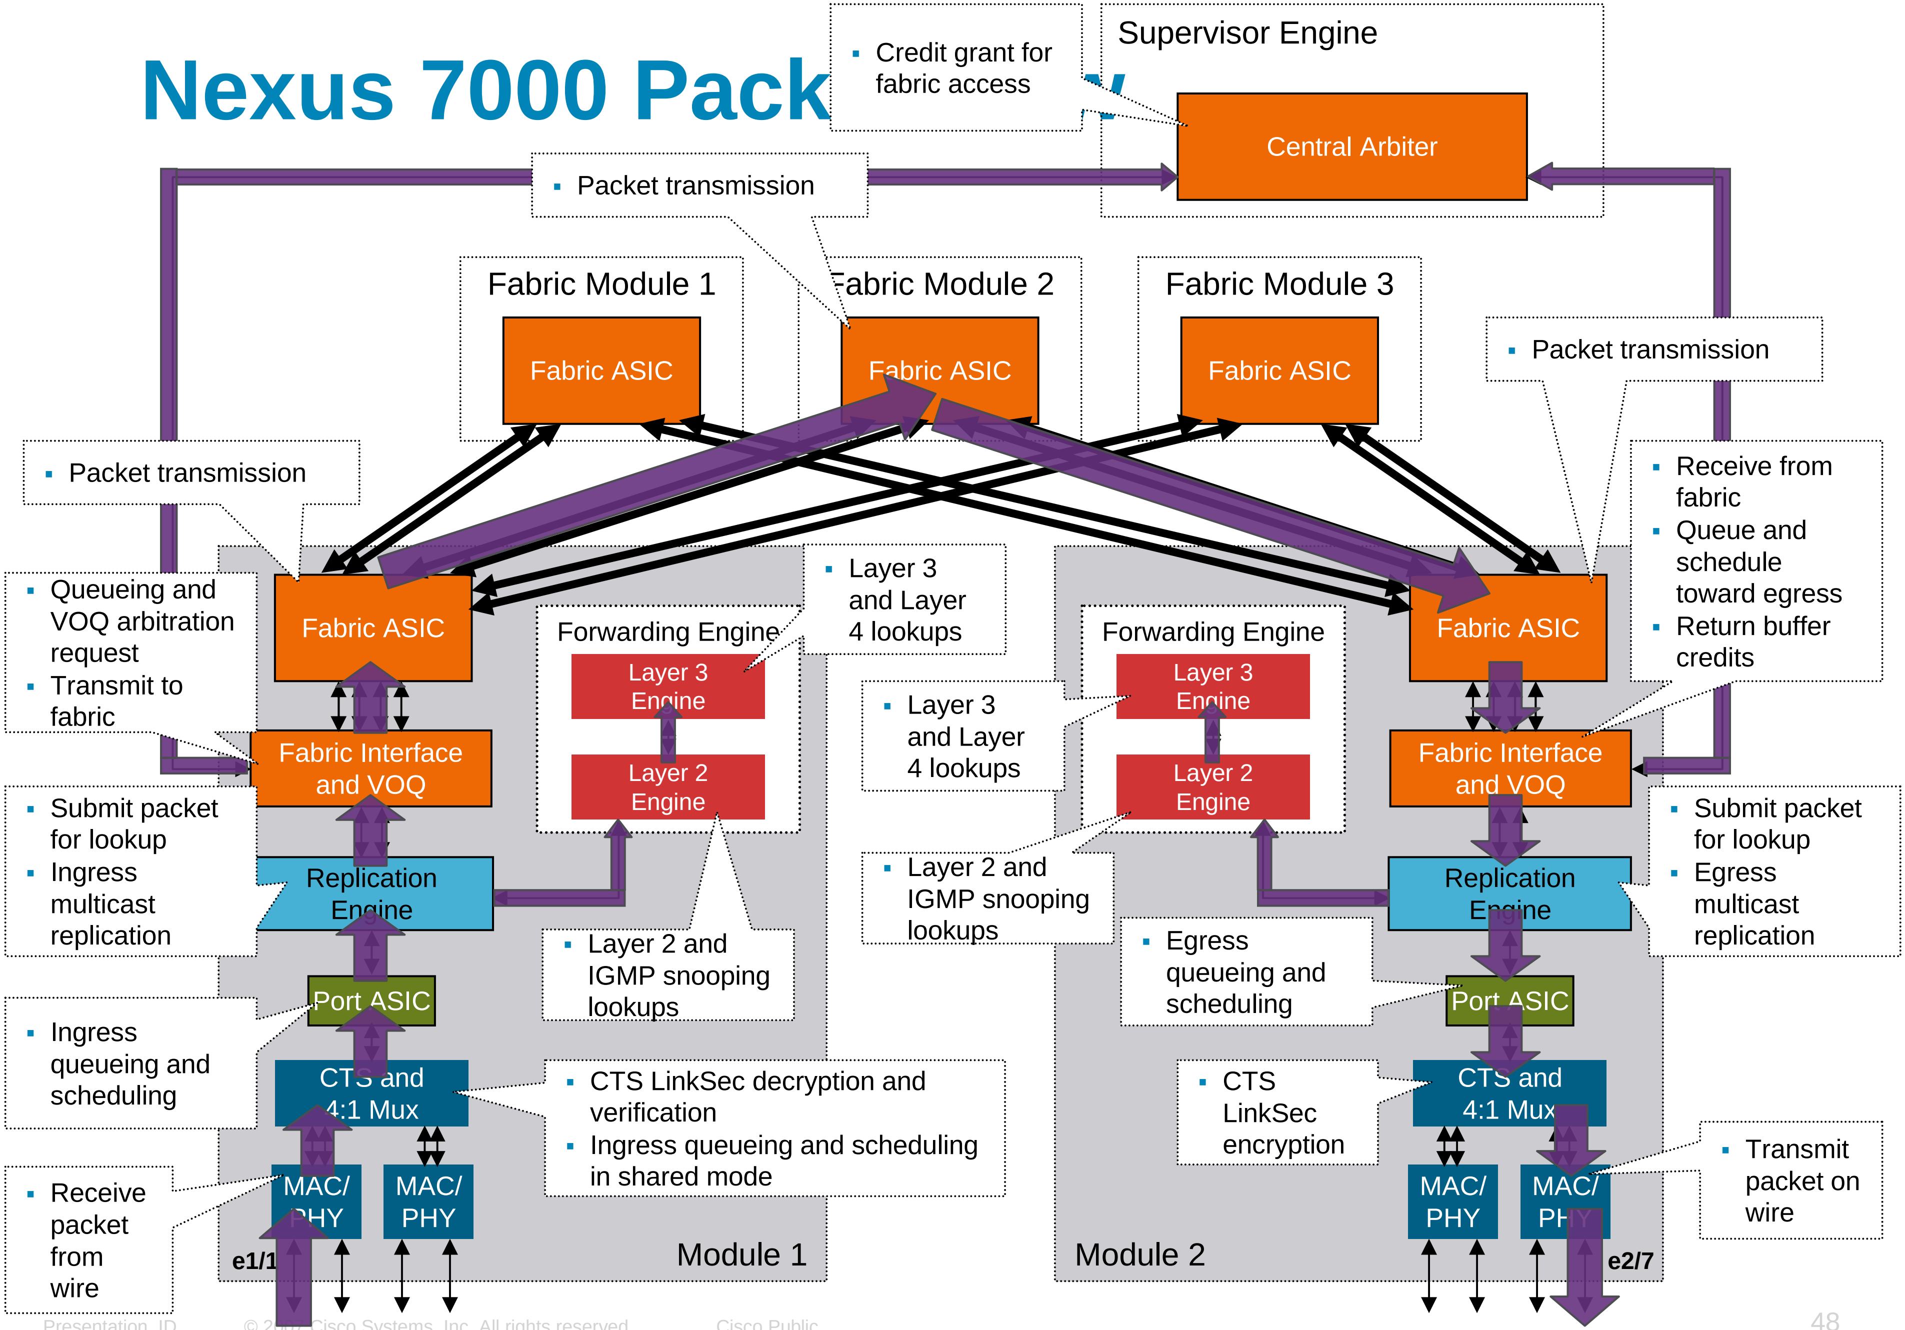

Supervisor Engine

Central Arbiter

- Credit grant for fabric access

- Packet transmission

Fabric Module 1

Fabric ASIC

Fabric Module 2

Fabric ASIC

Fabric Module 3

Fabric ASIC

- Packet transmission

- Packet transmission

- Queueing and VOQ arbitration request
- Transmit to fabric

Fabric ASIC

Fabric Interface and VOQ

- Submit packet for lookup
- Ingress multicast replication

Replication Engine

Port ASIC

- Ingress queueing and scheduling

CTS and 4:1 Mux

- Receive packet from wire

MAC/ PHY

MAC/ PHY

e1/1

Module 1

Forwarding Engine

Layer 3 Engine

Layer 2 Engine

- Layer 3 and Layer 4 lookups

- Layer 2 and IGMP snooping lookups

- CTS LinkSec decryption and verification
- Ingress queueing and scheduling in shared mode

- Layer 3 and Layer 4 lookups

- Layer 2 and IGMP snooping lookups

Forwarding Engine

Layer 3 Engine

Layer 2 Engine

- Egress queueing and scheduling

- CTS LinkSec encryption

Module 2

- Receive from fabric
- Queue and schedule toward egress
- Return buffer credits

Fabric ASIC

Fabric Interface and VOQ

- Submit packet for lookup
- Egress multicast replication

Replication Engine

Port ASIC

CTS and 4:1 Mux

MAC/ PHY

MAC/ PHY

- Transmit packet on wire

e2/7

DC-OS
Data Center class operating system

DC-OS: Delivering DC Class Attributes

Operational Continuity

Granular stateful process restart provides increased uptime and improved network stability

Integrated Manageability toolset improves troubleshooting and reduces time-to-resolution

ISSU+ allows software upgrades without service interruption

Transport Flexibility

Virtual Device Contexts allow the switch to be split into multiple logical switches for better utilization and isolation

Multi-transport control plane natively supports unified fabric without external gateways

Infrastructure Scalability

Multi-core/multi-thread architecture means turning on features will not impact performance

DC-OS Software Architecture

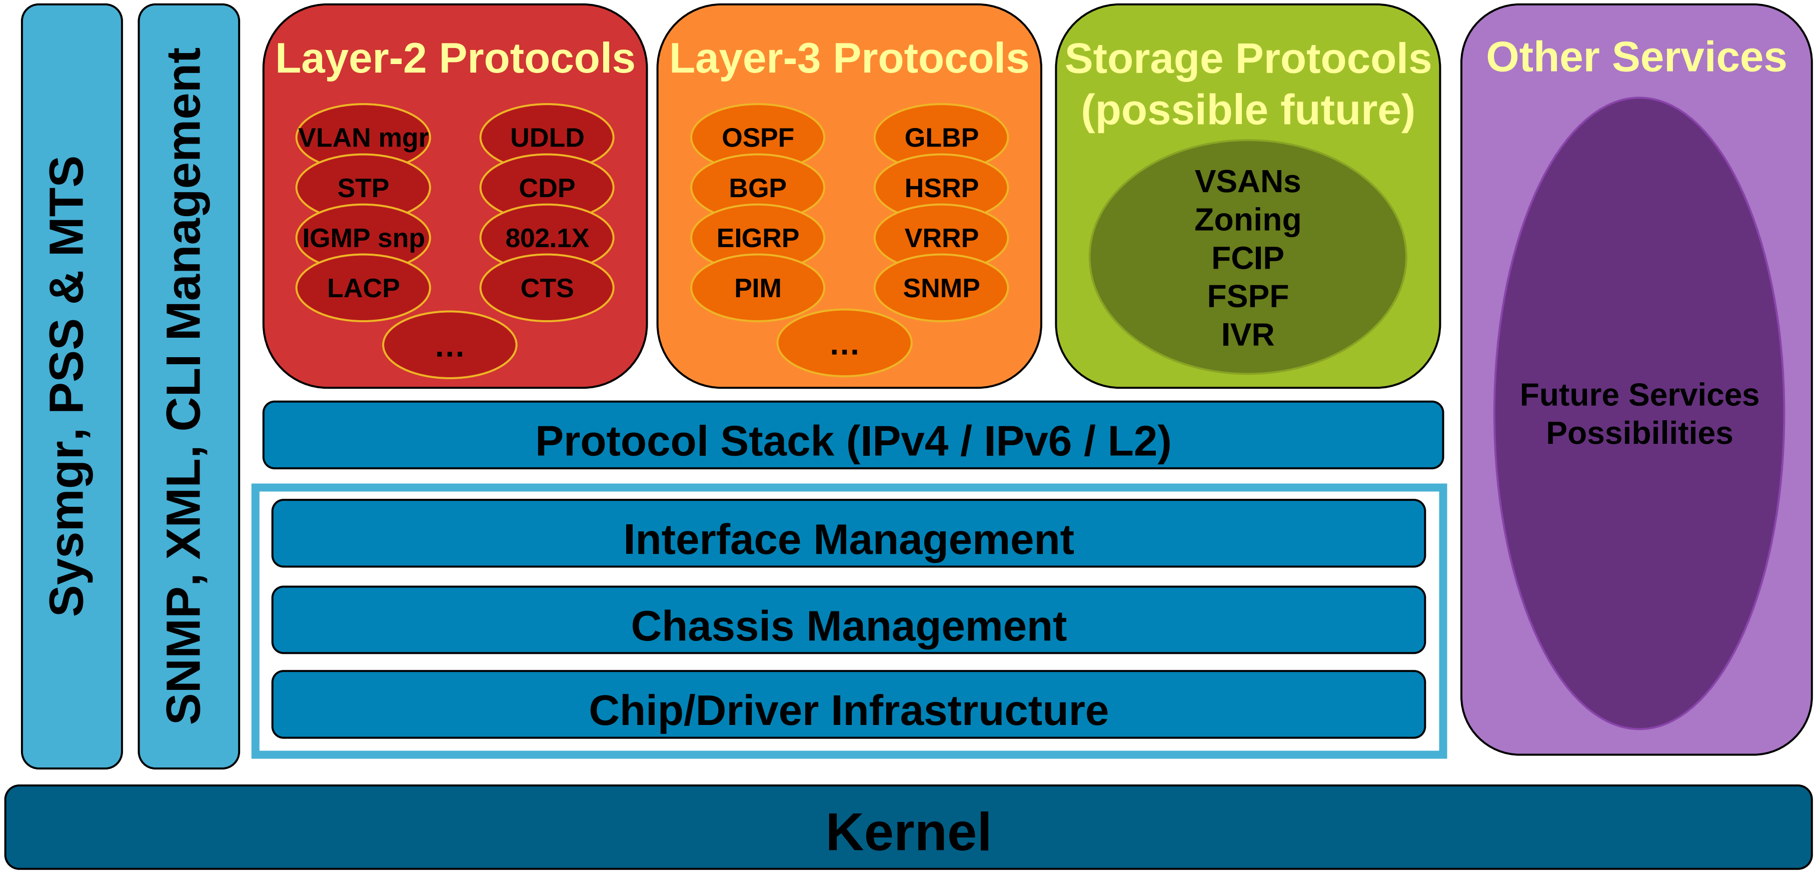

Licensing

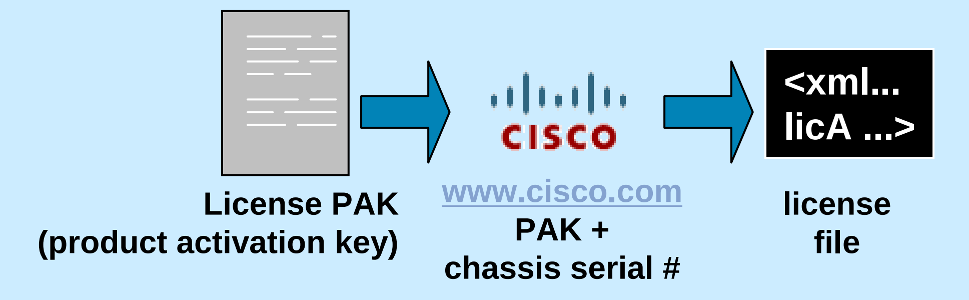

- Licenses are enforced on the switch

```
# show license host-id
```

License tied chassis serial # stored in dual redundant NVRAM modules on backplane

- Licenses are issued in the form of a digitally signed text file

```
# install license bootflash:DC3-1234.lic
```

Grace Period

- Enables features to be run for a certain period without installing a license
- Allows feature testing/trials without buying a license (e.g. 120 days)
- Periodic syslog, callhome and SNMP traps warning when grace period nears expiry

Time-bound licenses

- License with expiry date
- Currently used in SAN-OS as an emergency when grace period is over and need time to buy license
- Expiry date is absolute (expires at midnight UTC on expiry date)
- Periodic syslog, callhome and SNMP traps warning when time bound license nears expiry
- After expiry date feature will continue to run if grace period has not been exhausted

NX-OS Licensing

Simple, Flexible Licensing Model

- There are three levels of enforced licensing: Base, Enterprise Services, and Advanced Services
- Grace periods facilitate feature testing and trials without buying a license (for example, 120 days), with some restrictions. The Cisco Trusted Security does not have a grace period because of export restrictions on strong cryptography

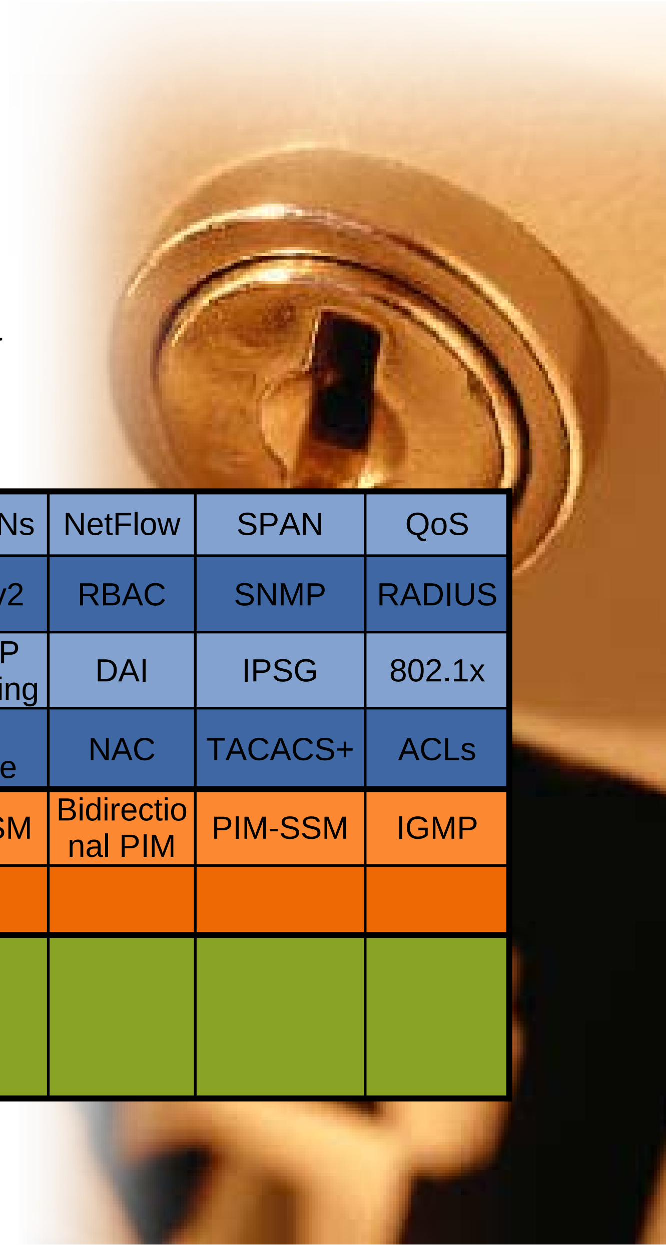

Base	ISSU	PVRST+	MSTP+	802.1Q	LACP	PVLANs	NetFlow	SPAN	QoS
	RIP/RIPng	IGMP snooping	DHCP helper	uRPF check	Port Security	SSHv2	RBAC	SNMP	RADIUS
	HSRP	GLBP	VRRP	VRF lite	CoPP	DHCP snooping	DAI	IPSG	802.1x
	Jumbo Frames	UDLD	Storm control	EEM	Cisco GOLD	Call Home	NAC	TACACS+	ACLs
Enterprise Services	OSPF	EIGRP	IS-IS	BGP	Graceful Restart	PIM-SM	Bidirectional PIM	PIM-SSM	IGMP
	MSDP	PBR	GRE						
Advanced Services	VDCs	Cisco Trusted Security							

Note: Enterprise Services is NOT included with Advanced Services license

New NX-OS Feature Navigator

Available NOW

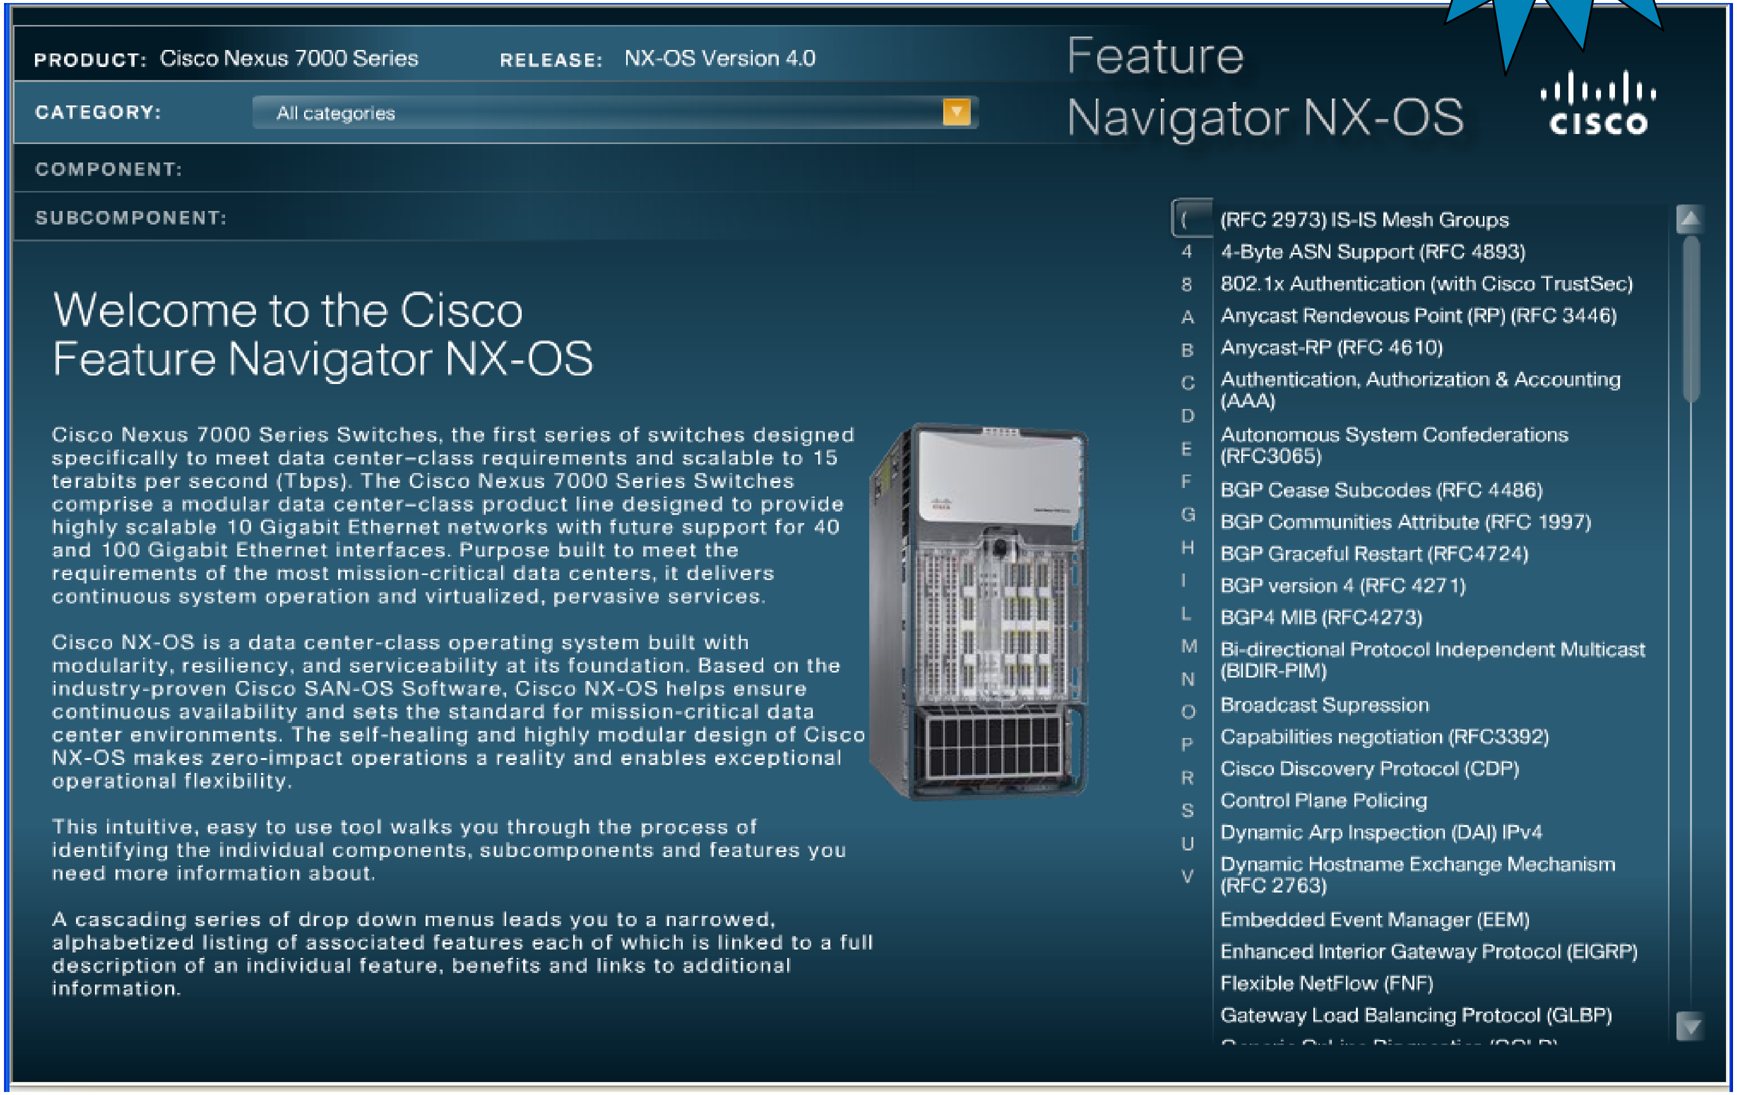

PRODUCT: Cisco Nexus 7000 Series
RELEASE: NX-OS Version 4.0
CATEGORY: All categories
COMPONENT:
SUBCOMPONENT:
Feature Navigator NX-OS
cisco
(4 8 A B C D E F G H I L M N O P R S U V

Welcome to the Cisco Feature Navigator NX-OS

Cisco Nexus 7000 Series Switches, the first series of switches designed specifically to meet data center–class requirements and scalable to 15 terabits per second (Tbps). The Cisco Nexus 7000 Series Switches comprise a modular data center–class product line designed to provide highly scalable 10 Gigabit Ethernet networks with future support for 40 and 100 Gigabit Ethernet interfaces. Purpose built to meet the requirements of the most mission-critical data centers, it delivers continuous system operation and virtualized, pervasive services.

Cisco NX-OS is a data center-class operating system built with modularity, resiliency, and serviceability at its foundation. Based on the industry-proven Cisco SAN-OS Software, Cisco NX-OS helps ensure continuous availability and sets the standard for mission-critical data center environments. The self-healing and highly modular design of Cisco NX-OS makes zero-impact operations a reality and enables exceptional operational flexibility.

This intuitive, easy to use tool walks you through the process of identifying the individual components, subcomponents and features you need more information about.

A cascading series of drop down menus leads you to a narrowed, alphabetized listing of associated features each of which is linked to a full description of an individual feature, benefits and links to additional information.

- (RFC 2973) IS-IS Mesh Groups
- 4-Byte ASN Support (RFC 4893)
- 802.1x Authentication (with Cisco TrustSec)
- Anycast Rendevous Point (RP) (RFC 3446)
- Anycast-RP (RFC 4610)
- Authentication, Authorization & Accounting (AAA)
- Autonomous System Confederations (RFC3065)
- BGP Cease Subcodes (RFC 4486)
- BGP Communities Attribute (RFC 1997)
- BGP Graceful Restart (RFC4724)
- BGP version 4 (RFC 4271)
- BGP4 MIB (RFC4273)
- Bi-directional Protocol Independent Multicast (BIDIR-PIM)
- Broadcast Supression
- Capabilities negotiation (RFC3392)
- Cisco Discovery Protocol (CDP)
- Control Plane Policing
- Dynamic Arp Inspection (DAI) IPv4
- Dynamic Hostname Exchange Mechanism (RFC 2763)
- Embedded Event Manager (EEM)
- Enhanced Interior Gateway Protocol (EIGRP)
- Flexible NetFlow (FNF)
- Gateway Load Balancing Protocol (GLBP)
- Generic Online Diagnostics (GOLD)

http://www.cisco.com/cdc_content_elements/flash/dataCenter/ciscofeaturenavigator/index.html

Stateful Fault Recovery

- DCOS services checkpoint their runtime state to the PSS for recovery in the event of a failure

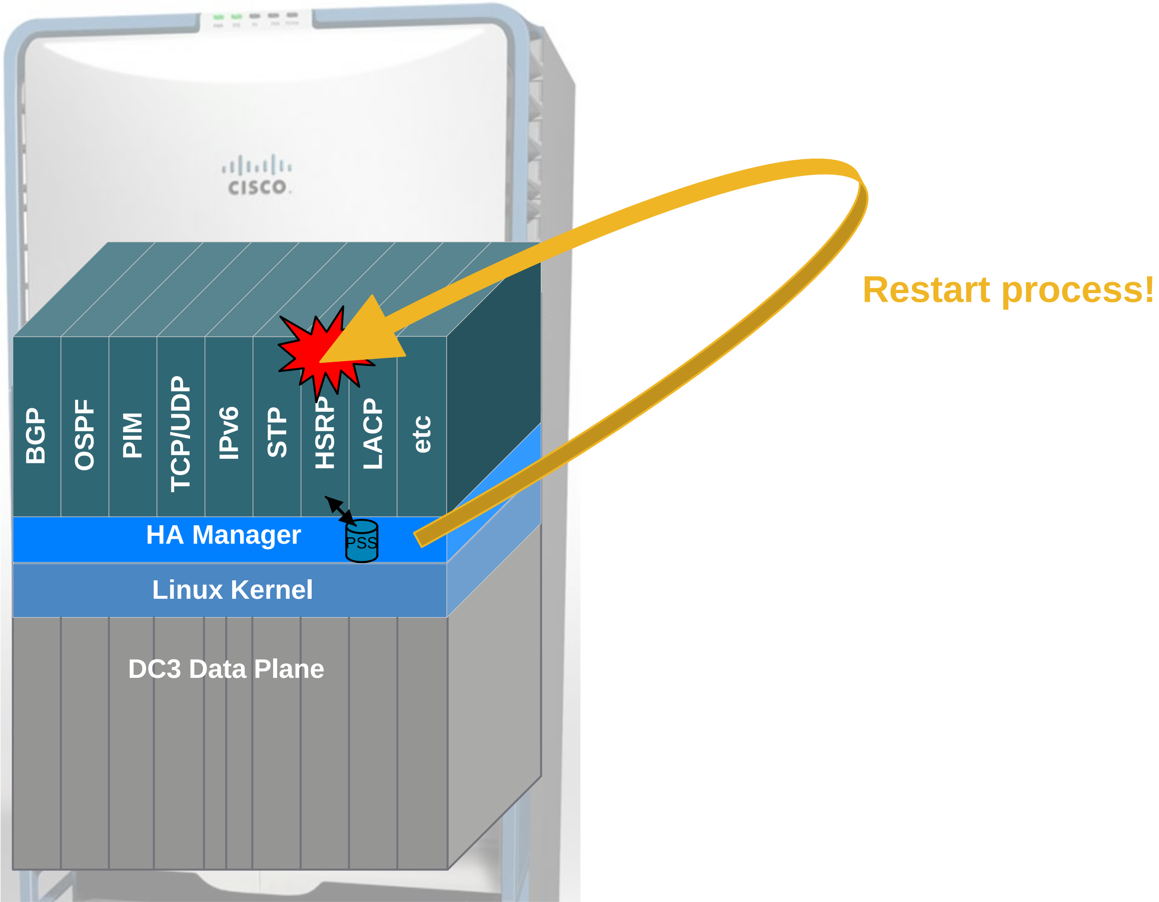

If a fault occurs in a process…

- HA manager determines best recovery action (restart process, switchover to redundant supervisor)
- Process restarts with no impact on data plane
    - State checkpointing (PSS) allows instant, stateful process recovery
    - Software utilizes Graceful Restart where appropriate

Stateful Fault Recovery

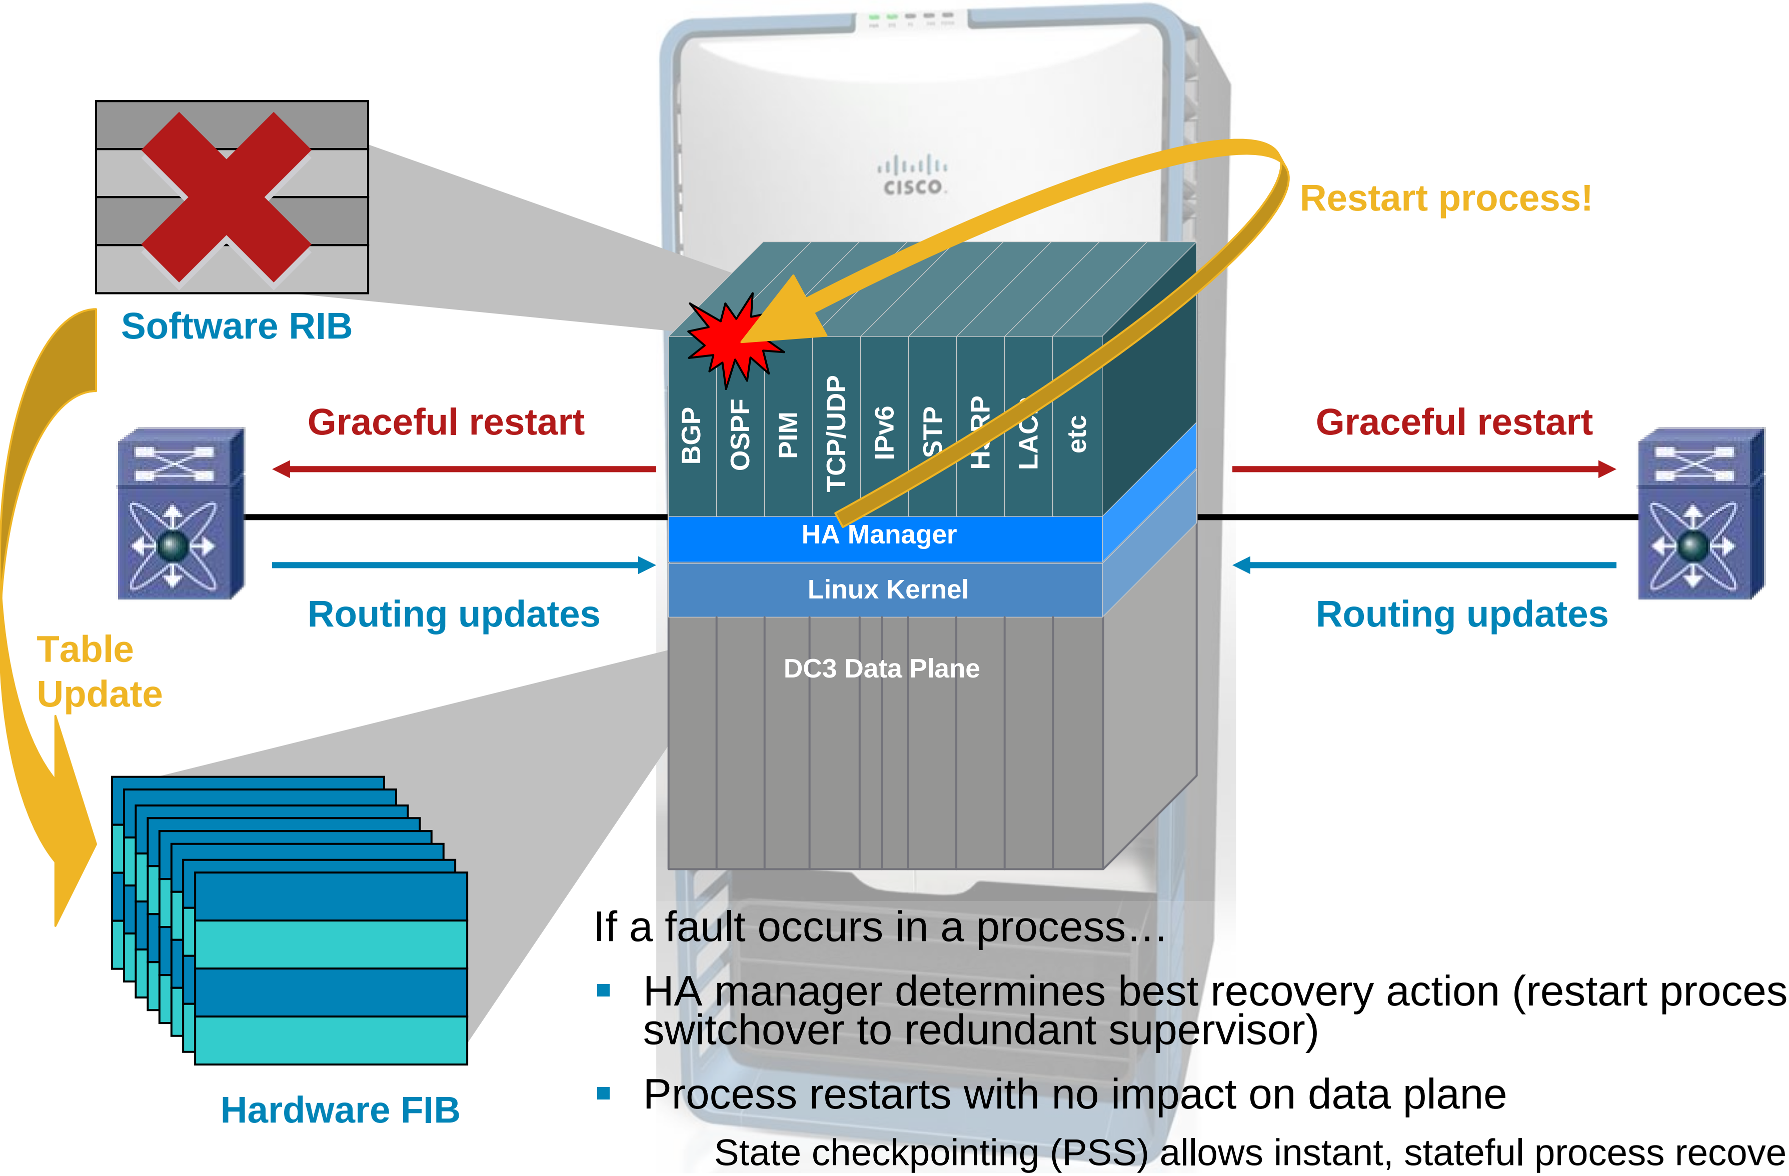

If a fault occurs in a process…

- HA manager determines best recovery action (restart process, switchover to redundant supervisor)
- Process restarts with no impact on data plane
  - State checkpointing (PSS) allows instant, stateful process recovery
  - Software utilizes Graceful Restart where appropriate

In-Service Software Upgrade

```
dc3# install all kickstart bootdisk:4.1-kickstart system bootdisk:4.1-system
dc3#
```

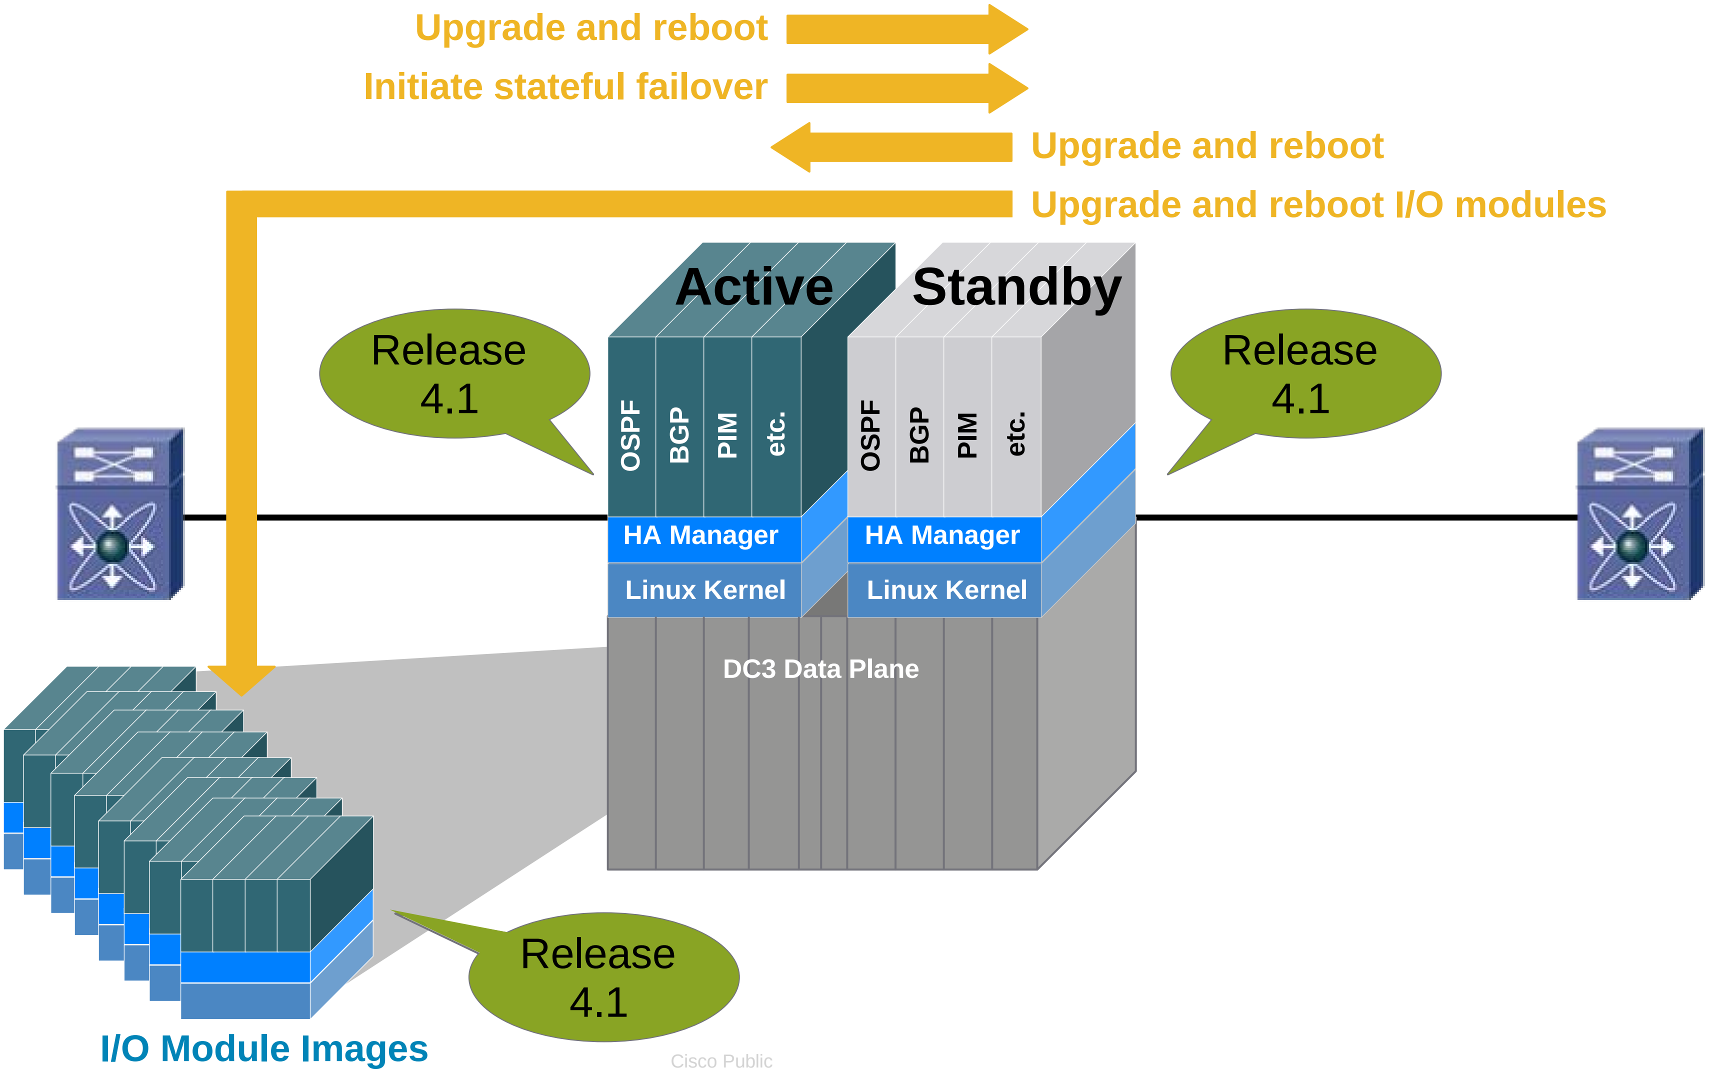

Virtual Device Contexts (VDCs)

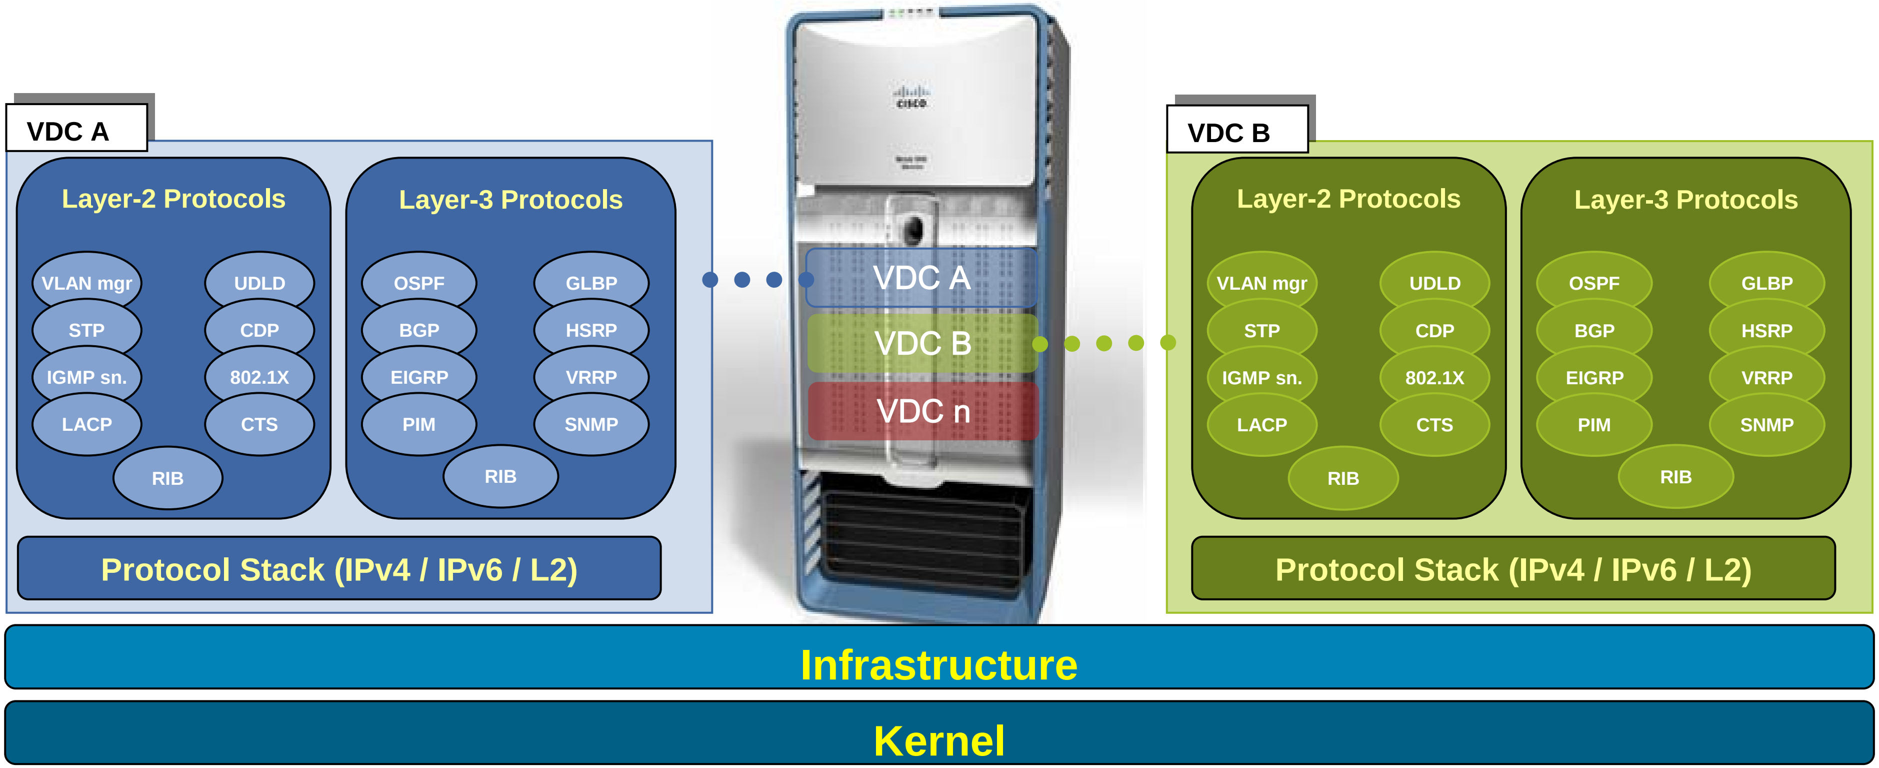

VDC A

Layer-2 Protocols

VLAN mgr, UDLD, STP, CDP, IGMP sn., 802.1X, LACP, CTS, RIB

Layer-3 Protocols

OSPF, GLBP, BGP, HSRP, EIGRP, VRRP, PIM, SNMP, RIB

Protocol Stack (IPv4 / IPv6 / L2)

VDC A

VDC B

VDC n

VDC B

Layer-2 Protocols

VLAN mgr, UDLD, STP, CDP, IGMP sn., 802.1X, LACP, CTS, RIB

Layer-3 Protocols

OSPF, GLBP, BGP, HSRP, EIGRP, VRRP, PIM, SNMP, RIB

Protocol Stack (IPv4 / IPv6 / L2)

Infrastructure

Kernel

- **VDC – Virtual Device Context**

  Flexible separation/distribution of **Software Components**

  Flexible separation/distribution of **Hardware Resources**

  Securely delineated **Administrative Contexts**

- **VDCs are not…**

  The ability to run different OS levels on the same box at the same time

  based on a ***hypervisor*** model; there is a single ‘infrastructure’ layer that handles h/w programming…

Virtual Device Contexts
An Introduction to the VDC Architecture

Virtual Device Contexts provides virtualization at the device level allowing multiple instances of the device to operate on the same physical switch at the same time…

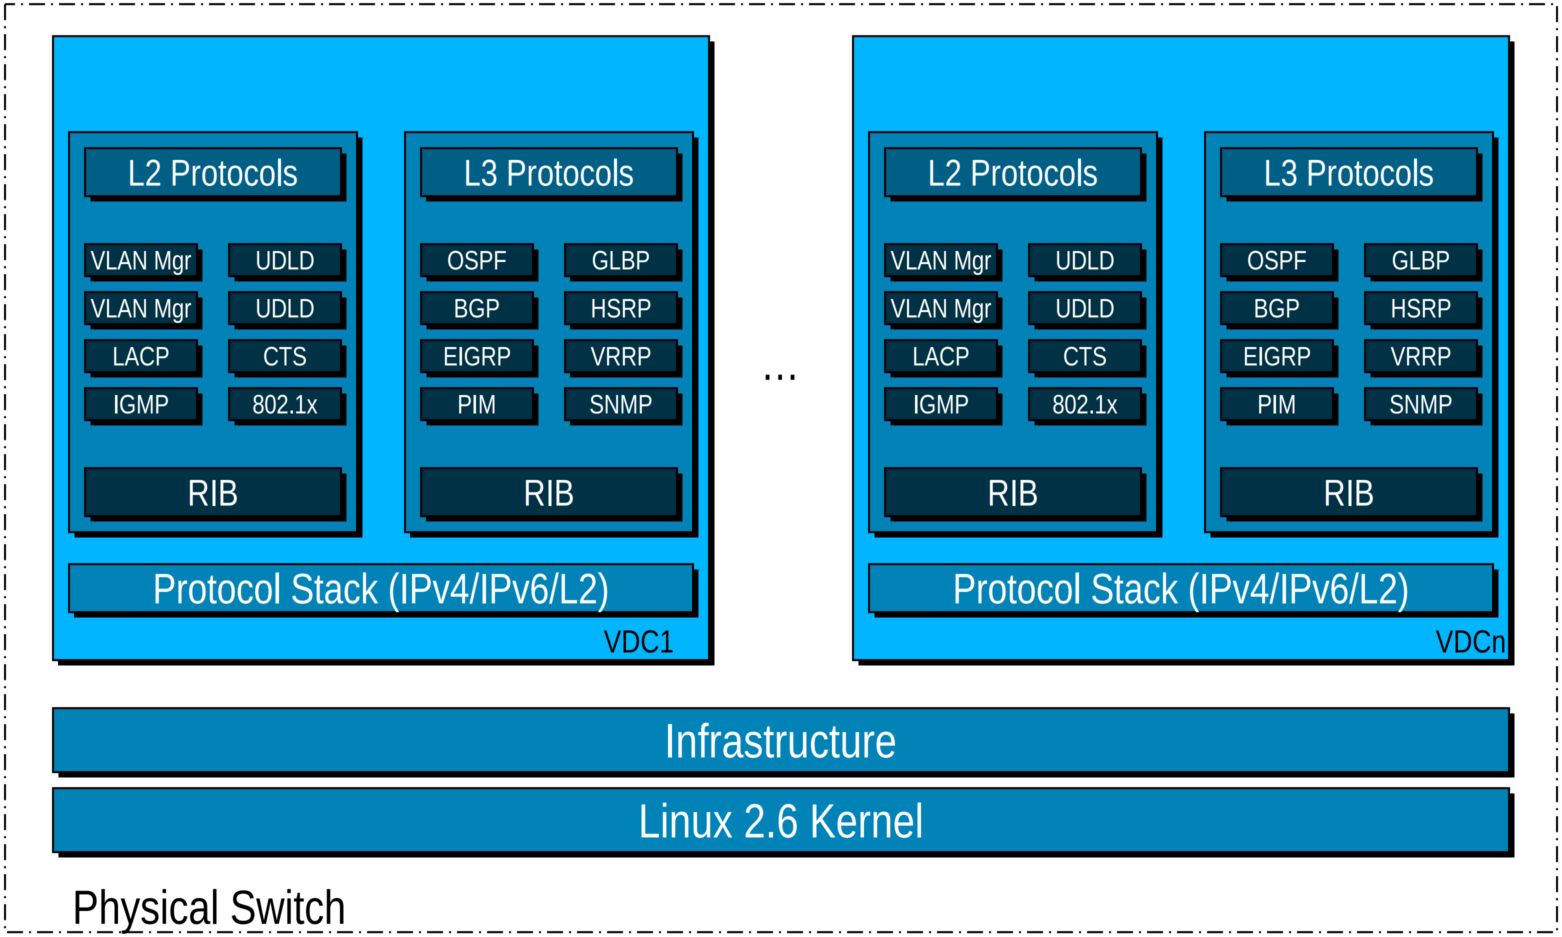

Virtual Device Contexts
VDC Fault Domain

A VDC builds a fault domain around all running processes within that VDC - should a fault occur in a running process, it is truly isolated from other running processes and they will not be impacted…

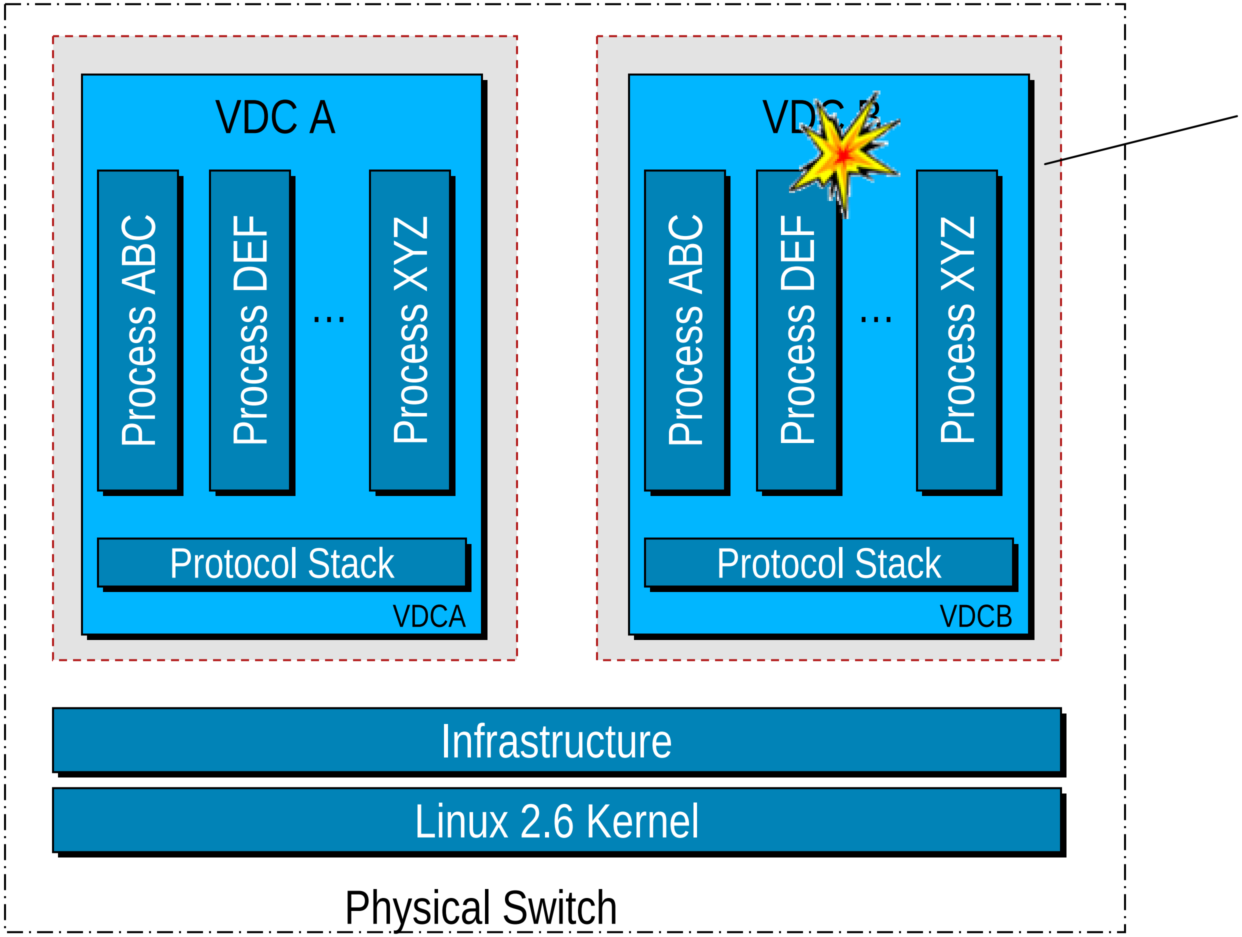

Fault Domain

Process “DEF” in VDC B crashes

Processes in VDC A are not affected and will continue to run unimpeded

This is a function of the process modularity of the OS and a VDC specific IPC context

Virtual Device Contexts (VDCs)

- **Network Consolidation:**

  Multiple logical nets/single physical net

  Maintain clear delineation between nets

    Independent Topologies

    Clear Management Boundaries

    Fault Containment

- **Service Velocity:**

  In-line tests

  Rapid deployment and rollback

  e.g. Enable Utility Computing

- **Device Consolidation:**

  Logical Appliances

  Multi-switch emulation

  Pwr, Cooling & Real-Estate efficiencies

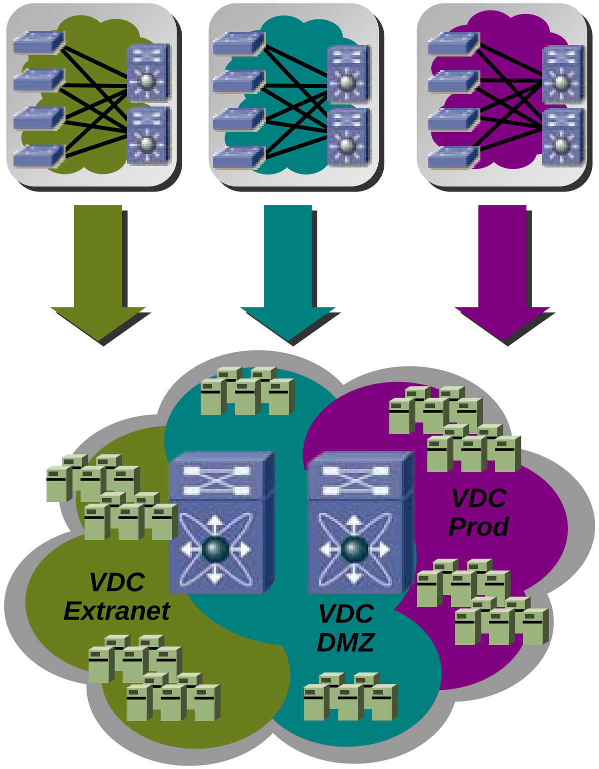

Physical network islands are *virtualized* onto common datacenter networking infrastructure

Command Line Interface

CLI

IOS look-and-feel CLI with enhancements...

- Show commands can be executed identically from exec mode and configuration mode
- Show commands have parser help even in configuration mode

```
tstevens-dc3-10# sh ip ospf nei
 OSPF Process ID 10 context default
 Total number of neighbors: 1
 Neighbor ID     Pri State            Up Time  Address         Interface
 10.255.255.2      1 INIT/DROTHER     00:00:04 10.1.2.2        Eth1/2
tstevens-dc3-10# config t
tstevens-dc3-10(config)# router ospf 10
tstevens-dc3-10(config-router)# sh ip ospf nei
 OSPF Process ID 10 context default
 Total number of neighbors: 1
 Neighbor ID     Pri State            Up Time  Address         Interface
 10.255.255.2      1 FULL/BDR         00:00:03 10.1.2.2        Eth1/2
tstevens-dc3-10(config-router)#
```

CLI
Routing Configuration

Two configuration models for routing protocols

- BGP follows neighbor-centric model

```
tstevens-dc3-10(config)# router bgp 100
tstevens-dc3-10(config-router)# address-family ipv4 unicast
tstevens-dc3-10(config-router-af)# network 10.0.0.0/8
tstevens-dc3-10(config-router-af)# neighbor 10.1.2.2 remote-as 200
tstevens-dc3-10(config-router-neighbor)# address-family ipv4 unicast
tstevens-dc3-10(config-router-neighbor-af)# soft-reconfiguration inbound
```

- IGPs follow interface-centric model

```
tstevens-dc3-10(config)# router ospf 10
tstevens-dc3-10(config-router)# int e2/22
tstevens-dc3-10(config-if)# ip router ospf 10 area 0
tstevens-dc3-10(config-if)# ip ospf hello-interval 1
tstevens-dc3-10(config-if)#
```

CLI
Slash Notation

- “Slash” notation supported for all IPv4/IPv6 masks

```
tstevens-dc3-10(config)# int e2/23
tstevens-dc3-10(config-if)# ip add 10.2.23.1/24
tstevens-dc3-10(config-if)# ipv6 add ::abcd:223/120
tstevens-dc3-10(config-if)# ip access-list test
tstevens-dc3-10(config-acl)# permit ip 10.1.1.0/24 any
tstevens-dc3-10(config-acl)#
```

CLI
Interface Ranges

- Same configuration used for interface ranges as for single interfaces

```
tstevens-dc3-10(config)# int e1/1-3
tstevens-dc3-10(config-if-range)# no sh
tstevens-dc3-10(config-if-range)# int e2/3
tstevens-dc3-10(config-if)# ip add 10.2.3.1/24
tstevens-dc3-10(config-if)# int e2/1-4,e1/1-2,e1/15
tstevens-dc3-10(config-if-range)# mtu 9216
tstevens-dc3-10(config-if-range)#
```

CLI
Parser Help

- <TAB> key displays brief list of all available options at current branch
- ? key displays full parser help strings

```
tstevens-dc3-10(config-if)# <TAB>
bandwidth        description      exit             mac              rate-mode        storm-control
beacon           dot1x            flowcontrol      mdix             service-policy   switchport
cdp              duplex           ip               mtu              shutdown         vrrp
channel-group    eou              ipv6             nac              spanning-tree
delay            errdisable       link             no               speed
tstevens-dc3-10(config-if)# ?
  bandwidth       Set bandwidth informational parameter
  beacon          Disable/enable the beacon for an interface
  cdp             CDP Interface Configuration parameters
  channel-group   Add to/remove from a port-channel
  delay           Specify interface throughput delay
  <etc>
```

CLI
Piping Terminal Output

- Variety of advanced pipe options for CLI output , including egrep, less, no-more, wc
- Multiple levels of pipe

```
tstevens-dc3-10# sh run | ?
  egrep    Egrep
  grep     Grep
  less     Stream Editor
  no-more  Turn-off pagination for command output
  wc       Count words, lines, characters
  begin    Begin with the line that matches
  count    Count number of lines
  exclude  Exclude lines that match
  include  Include lines that match

tstevens-dc3-10# sh run | egrep ?
  -A    Print <num> lines of context after every matching line
  -B    Print <num> lines of context before every matching line
  -c    Print a total count of matching lines only
  -i    Ignore case difference when comparing strings
  -n    Print each match preceded by its line number
  -v    Print only lines that contain no matches for <expr>
  -w    Print only lines where the match is a complete word
  -x    Print only lines where the match is a whole line
  WORD  Search for the expression
```

```
tstevens-dc3-10# sh run | egrep -A 2 -B 2 ospf
interface Ethernet2/22
  ip address 10.2.22.1/24
  ip router ospf 10 area 0

interface Ethernet2/23
  ip address 10.2.23.1/24
  ip router ospf 10 area 0

interface Ethernet2/24
--
interface loopback0
  ip address 10.255.255.1/32
  ip router ospf 10 area 0
router ospf 10
hostname tstevens-dc3-10
tstevens-dc3-10# sh run | in ospf | wc -l
4
tstevens-dc3-10#
```

CLI
Configuration Rollback

- Provides checkpointing and rollback facility to return configuration to any previous state
- Options to name checkpoints, view contents of checkpointed configuration, diff checkpoints versus each other or running/startup configuration, etc.

```
tstevens-dc3-10# sh checkpoint
---------------------------------------------------------------------
Checkpoint_id   Label           UserName                TimeStamp
---------------------------------------------------------------------

16777476        10-8            tstevens                Mon Oct  8 21:55:45 2007

tstevens-dc3-10# rollback destination label 10-8
Note: Processing the Request... Please Wait
Note: Generating the Rollbackpatch... Please Wait
Note: Executing the patch... Please Wait
`conf t`
`interface Ethernet1/1`
`no service-policy type qos input foo stats-enable`
`no ip access-group test in`
tstevens-dc3-10#
```

CLI
running-config permutations

- ‘show running-config’ (“show run”) works as expected, but there are many other enhancements over IOS:

```
dc3# show running-config ?
  <CR>
  >                Redirect it to a file
  aaa              Display aaa configuration
  all              Current operating configuration with defaults
  am               Display am information
  arp              Display arp information
  bgp              Display bgp information
  callhome         Display callhome configuration
  cdp              Display cdp configuration
  cmp              Display CMP information
  copp             show running config for copp
  dhcp             Display dhcp snoop configurations
  diagnostic       Display diagnostic information
  diff             Show the difference between running and startup configuration
  dot1x            Display dot1x configuration
  eem              Show the event manager running configuration
  eigrp            Display eigrp information
  icmpv6           Display icmpv6 information
  igmp             Display igmp information
  interface        Interface configuration
  ip               Display ip information
  ipqos            show running config for ipqosmgr
...
```

DCNM
Data Center Network Manager

DCNM Solution Components

DCNM is a Client Server Solution

- DCNM Server communicates with the DC-OS devices
- DCNM Client communicates with the DCNM Server

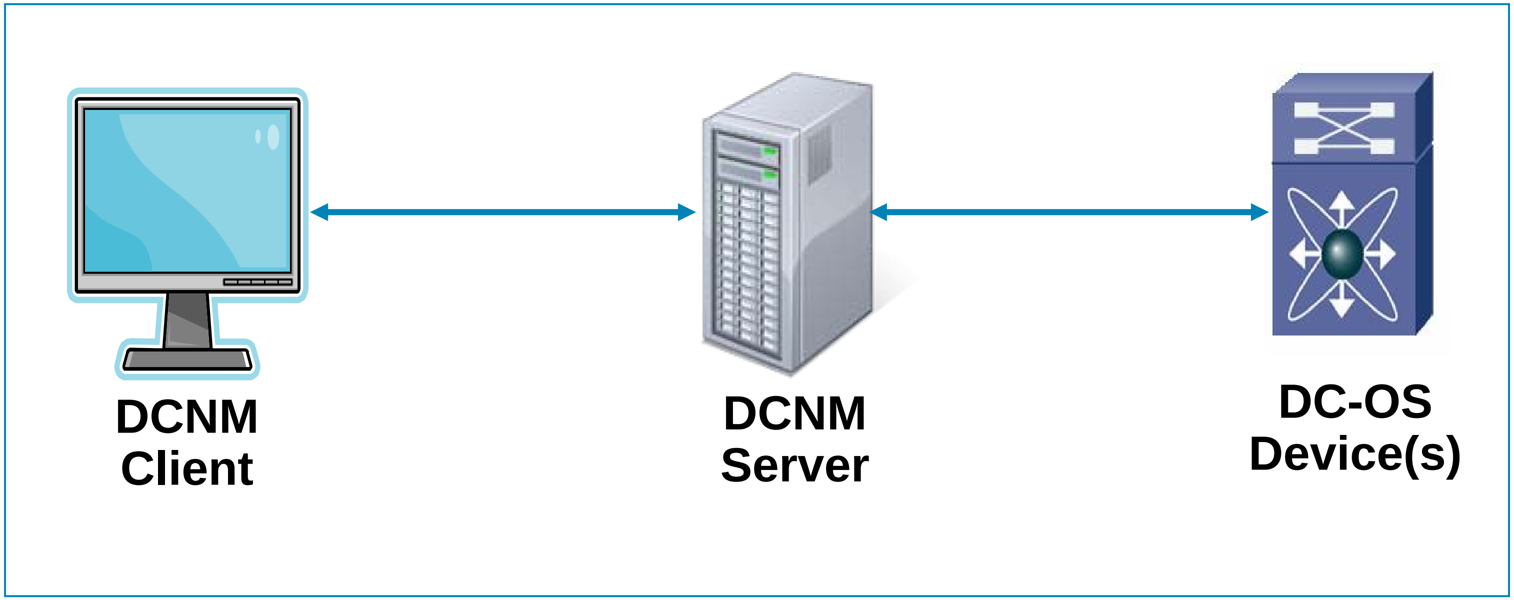

Server Hardware Specifications

System Requirements

- CPU Speed: 3+GHz dual-core processor (32 bit)
- RAM: Minimum 4GB
- 100GB High Performance Hard disk

Operating Systems Supported

- Windows Server 2003 Standard Edition Service Pack 2
- Red Hat Enterprise Linux AS release 4.

DCNM Discovery

- Discovers DC-OS and Cisco IOS devices
- Discovers adjacent devices if CDP enabled
- Server collects extensive switch inventory and configuration details. Based on the collected information, DCNM Server builds a virtual network model.
- As part of discovery process, DCNM establishes an SSH session with each DC-OS device managed by DCNM and each Cisco IOS device discovered
- SSH session is left in place after discovery. DCNM relies on the SSH session to gather information at regular intervals.

DCNM Server Network Model

DCNM Server builds an intelligent Network data model that enables the server to intelligently serve user requests.

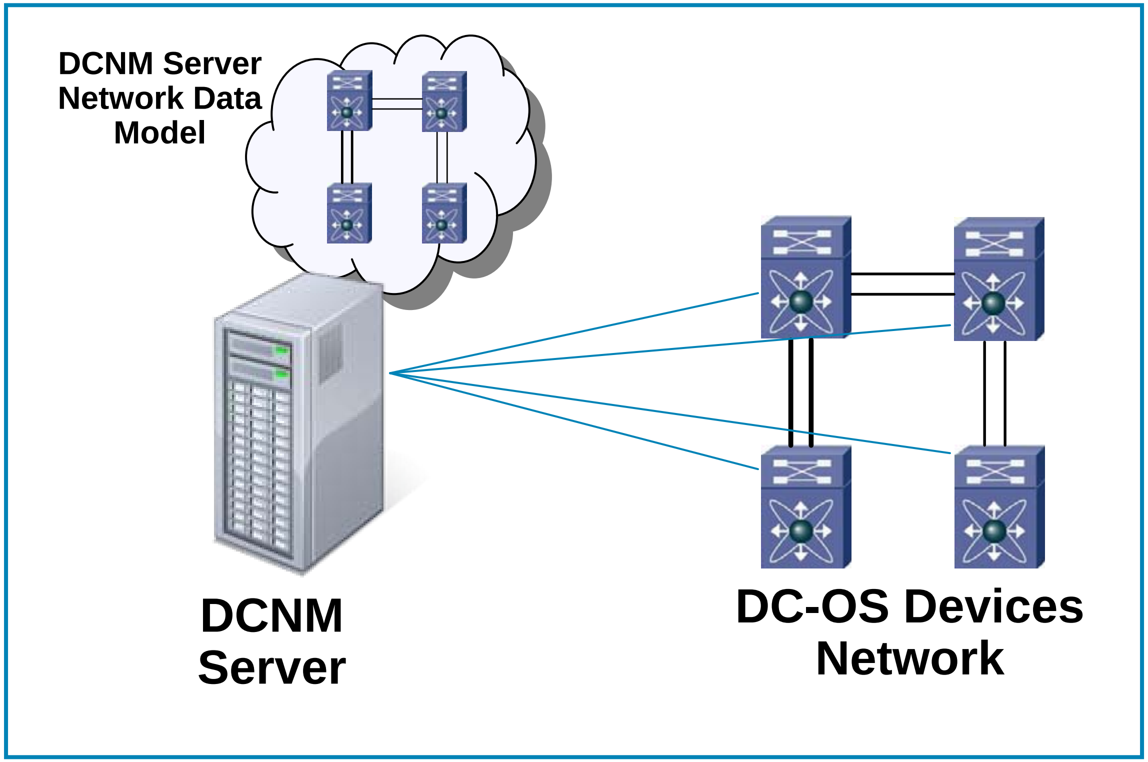

DCNM Client

- DCNM Client is a Java Application
- DCNM Client is downloaded from the DCNM Server using Java Web Start technology.

  Java Web Start technology enables Java software applications to be deployed with a single click over the network.

  Java Web Start ensures that the most current version of the application is deployed, as well as the current version of the Java Runtime Environment (JRE).

- DCNM Client is a thin client – all business logic on the DCNM Server.

Communications

- DCNM Server connects to the DC-OS devices over SSH.
- DCNM Client communicates to the DCNM server over Java RMI. No direct communication between DCNM Client and the DC3 devices.
- DCNM Server notifies DCNM Client of asynchronous events as JMS messages.

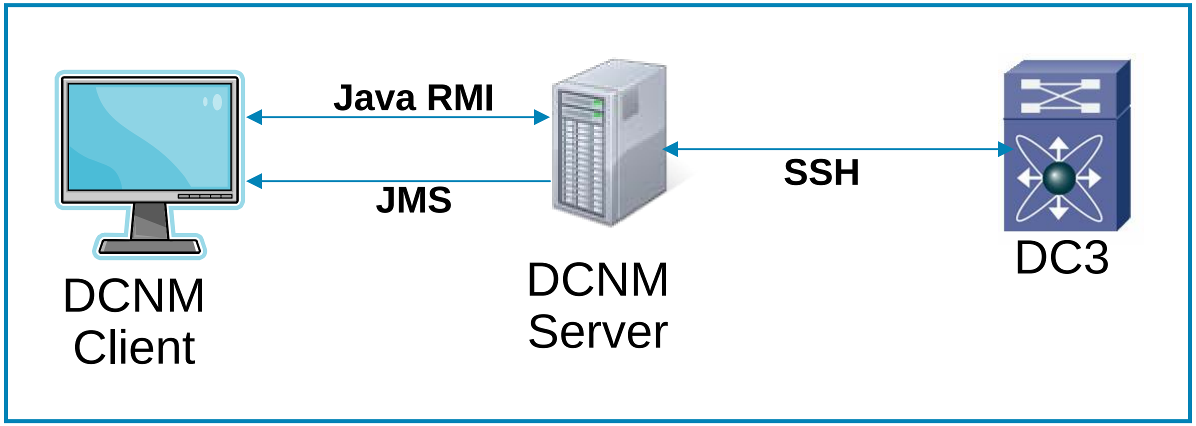

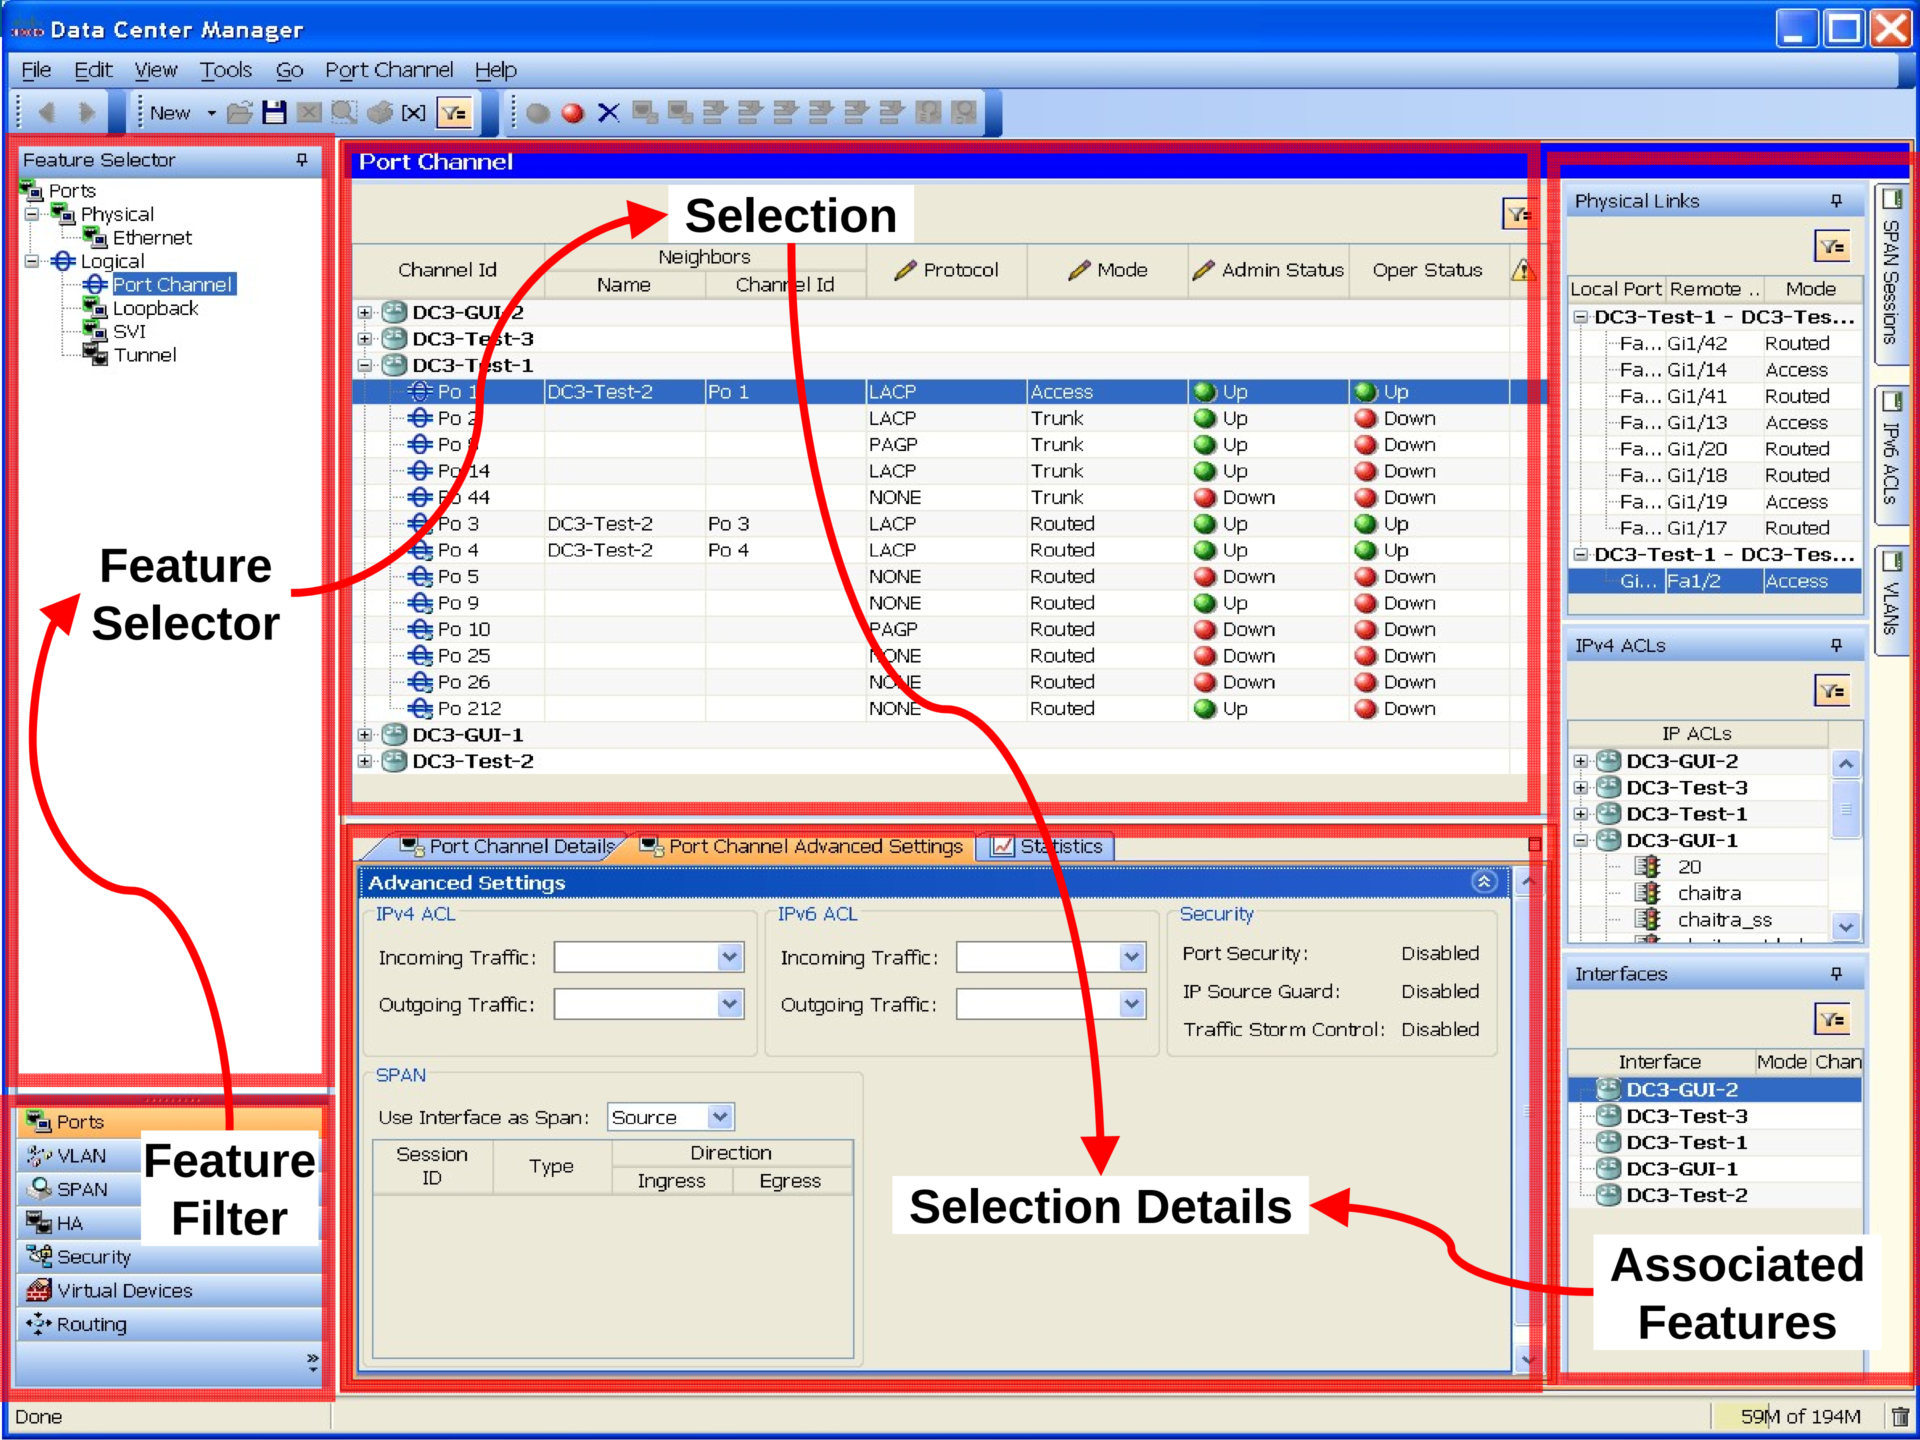

Port Channel

Channel Id	Neighbors Name	Neighbors Channel Id	Protocol	Mode	Admin Status	Oper Status
DC3-GUI-2						
DC3-Test-3						
DC3-Test-1						
Po 1	DC3-Test-2	Po 1	LACP	Access	Up	Up
Po 2			LACP	Trunk	Up	Down
Po 8			PAGP	Trunk	Up	Down
Po 14			LACP	Trunk	Up	Down
Po 44			NONE	Trunk	Down	Down
Po 3	DC3-Test-2	Po 3	LACP	Routed	Up	Up
Po 4	DC3-Test-2	Po 4	LACP	Routed	Up	Up
Po 5			NONE	Routed	Down	Down
Po 9			NONE	Routed	Up	Down
Po 10			PAGP	Routed	Down	Down
Po 25			NONE	Routed	Down	Down
Po 26			NONE	Routed	Down	Down
Po 212			NONE	Routed	Up	Down
DC3-GUI-1						
DC3-Test-2						

Advanced Settings

IPv4 ACL
Incoming Traffic:
Outgoing Traffic:

IPv6 ACL
Incoming Traffic:
Outgoing Traffic:

Security
Port Security: Disabled
IP Source Guard: Disabled
Traffic Storm Control: Disabled

SPAN
Use Interface as Span: Source

Session ID	Type	Direction Ingress	Direction Egress

Physical Links

Local Port	Remote ..	Mode
DC3-Test-1 - DC3-Tes...		
Fa...	Gi1/42	Routed
Fa...	Gi1/14	Access
Fa...	Gi1/41	Routed
Fa...	Gi1/13	Access
Fa...	Gi1/20	Routed
Fa...	Gi1/18	Routed
Fa...	Gi1/19	Access
Fa...	Gi1/17	Routed
DC3-Test-1 - DC3-Tes...		
Gi...	Fa1/2	Access

IPv4 ACLs

IP ACLs
DC3-GUI-2
DC3-Test-3
DC3-Test-1
DC3-GUI-1
20
chaitra
chaitra_ss

Interfaces

Interface	Mode	Chan
DC3-GUI-2		
DC3-Test-3		
DC3-Test-1		
DC3-GUI-1		
DC3-Test-2		

Data Center network design with Nexus

VSS Design with Service Modules and Nexus 7k in the core

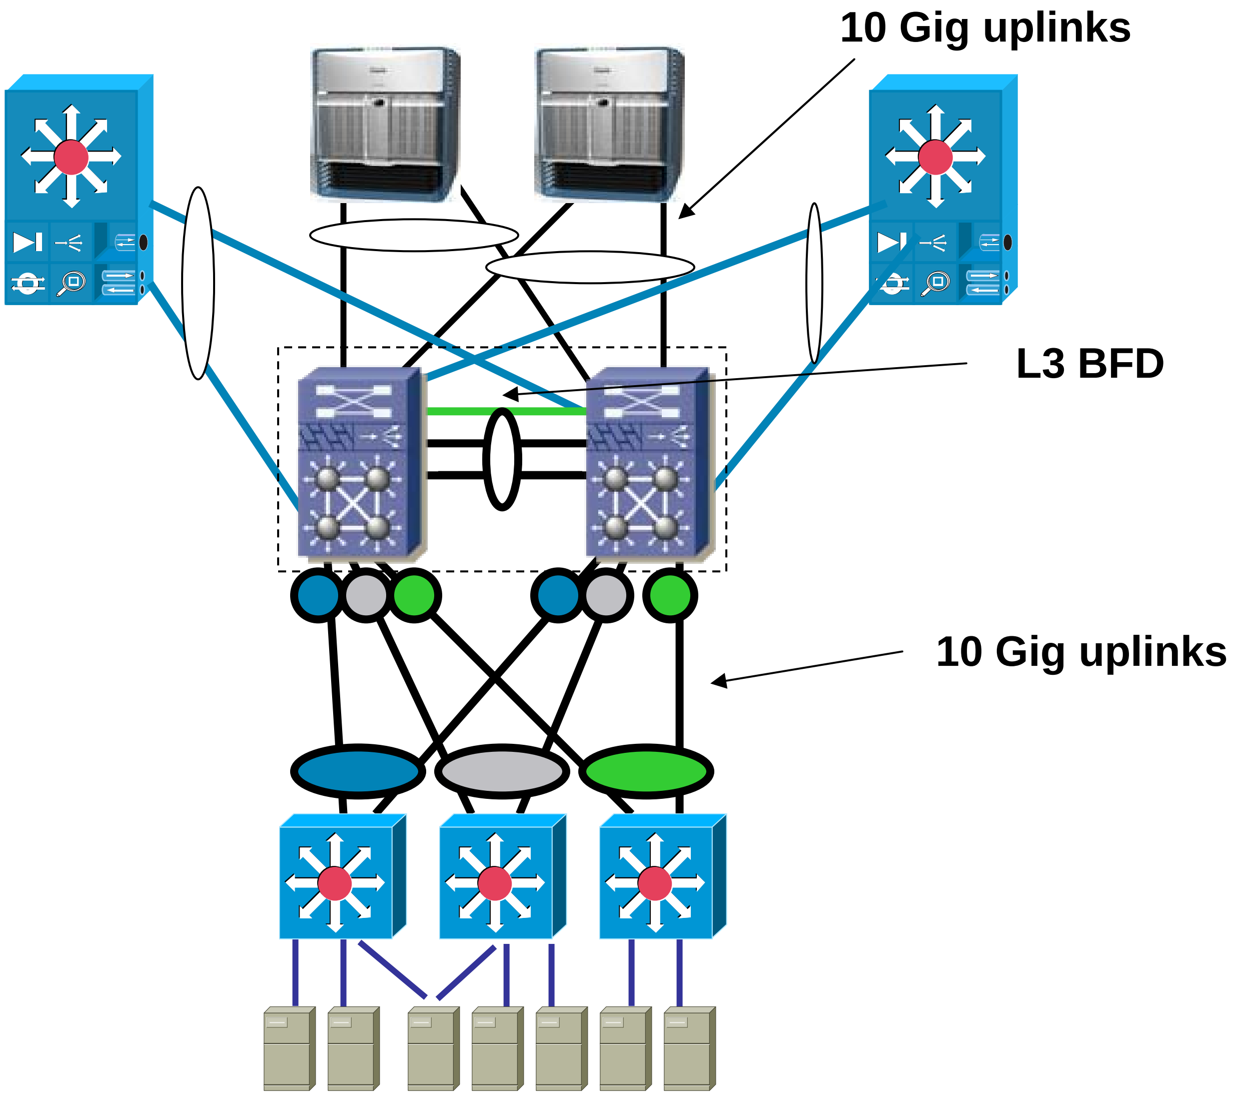

Nexus 7k for 10 Gig Aggregation

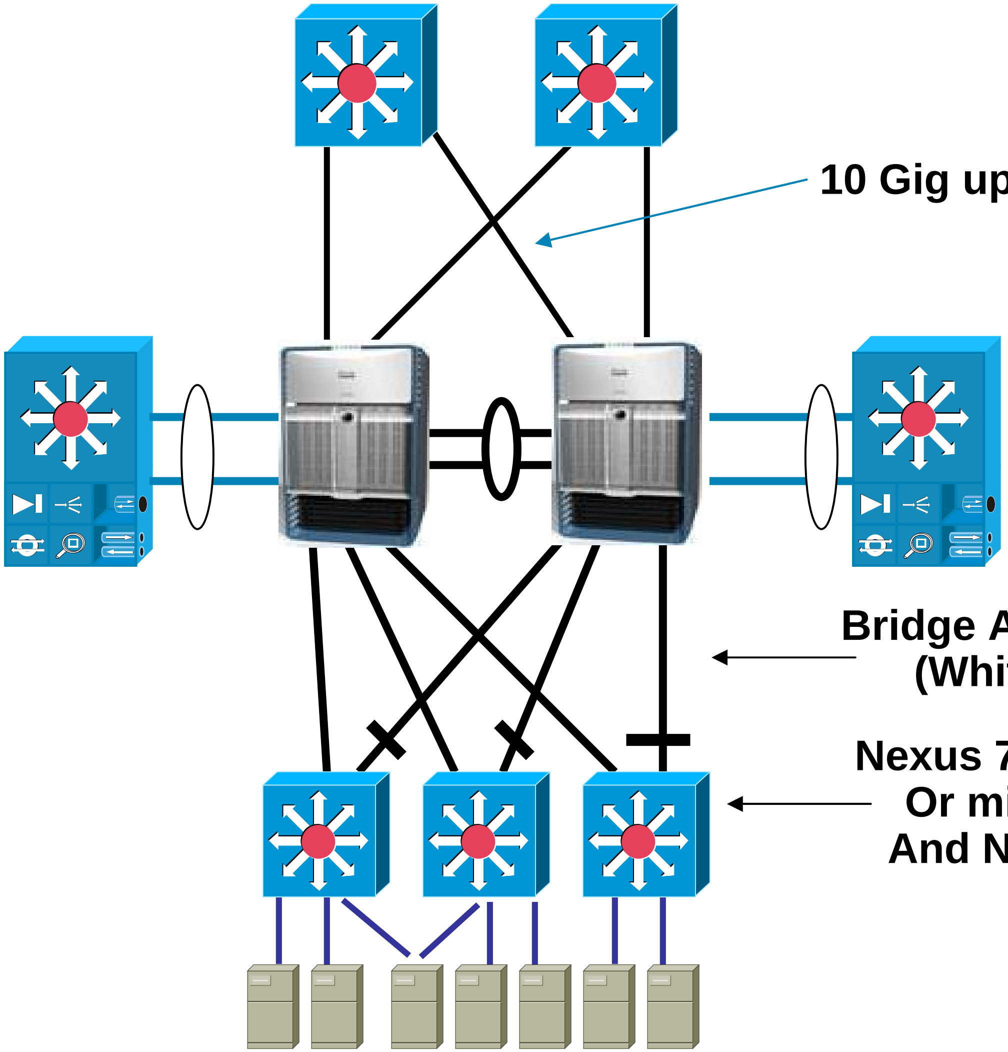

10 Gig uplinks

Bridge Assurance (Whitney2)

Nexus 7k Access
Or mix of 6k
And Nexus 7k

Collapsed Aggregation/Access

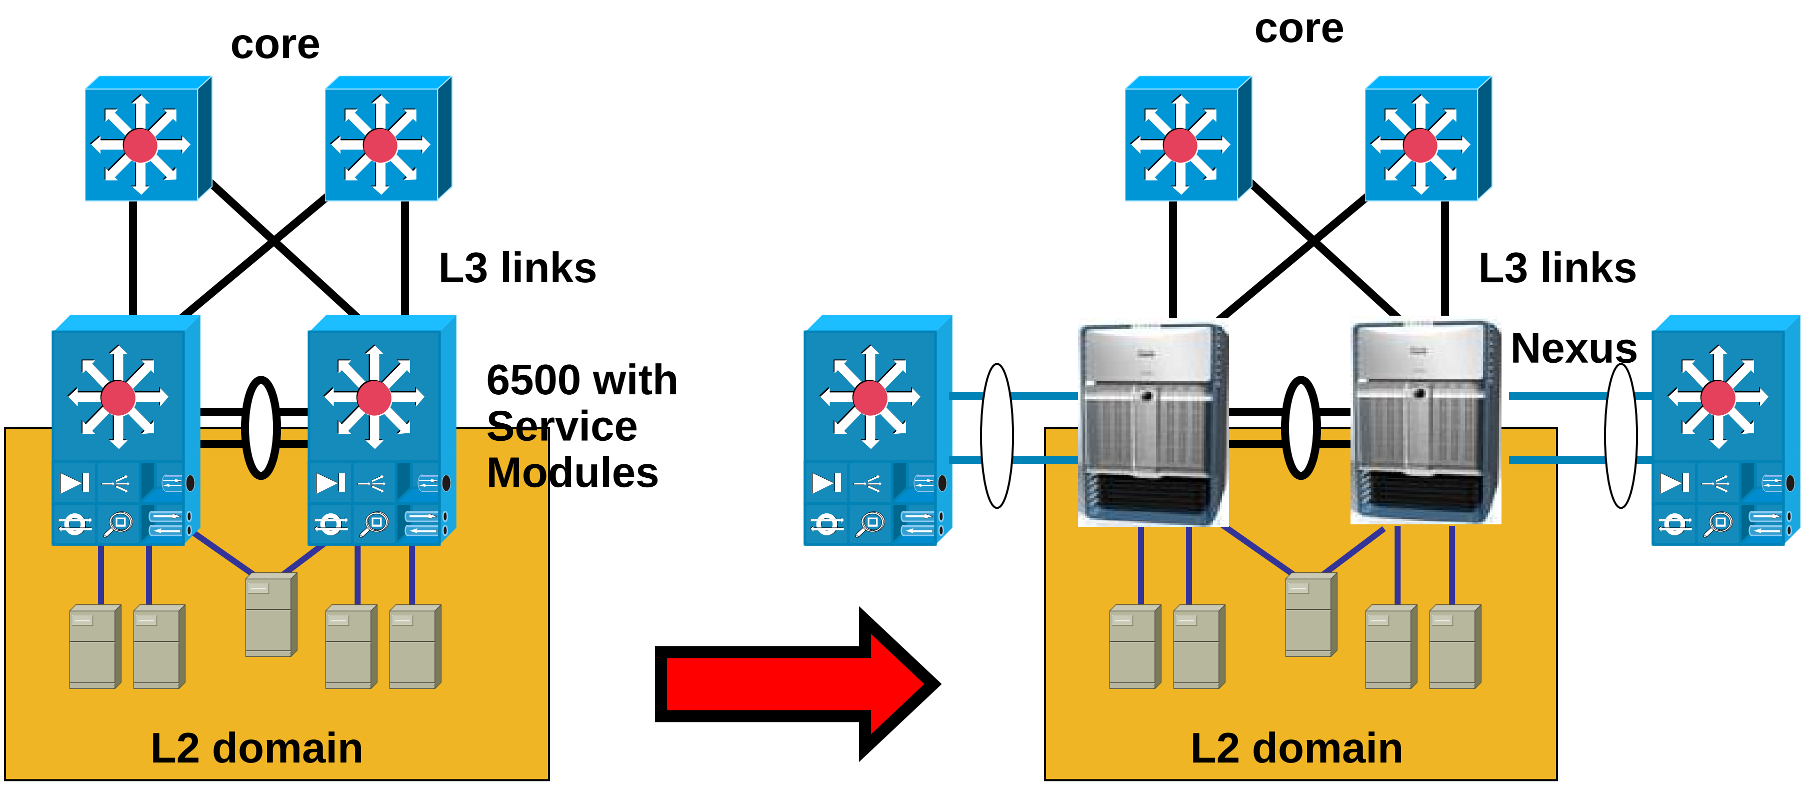

Data Center topologies with Nexus and virtualization

Reference Network Topology

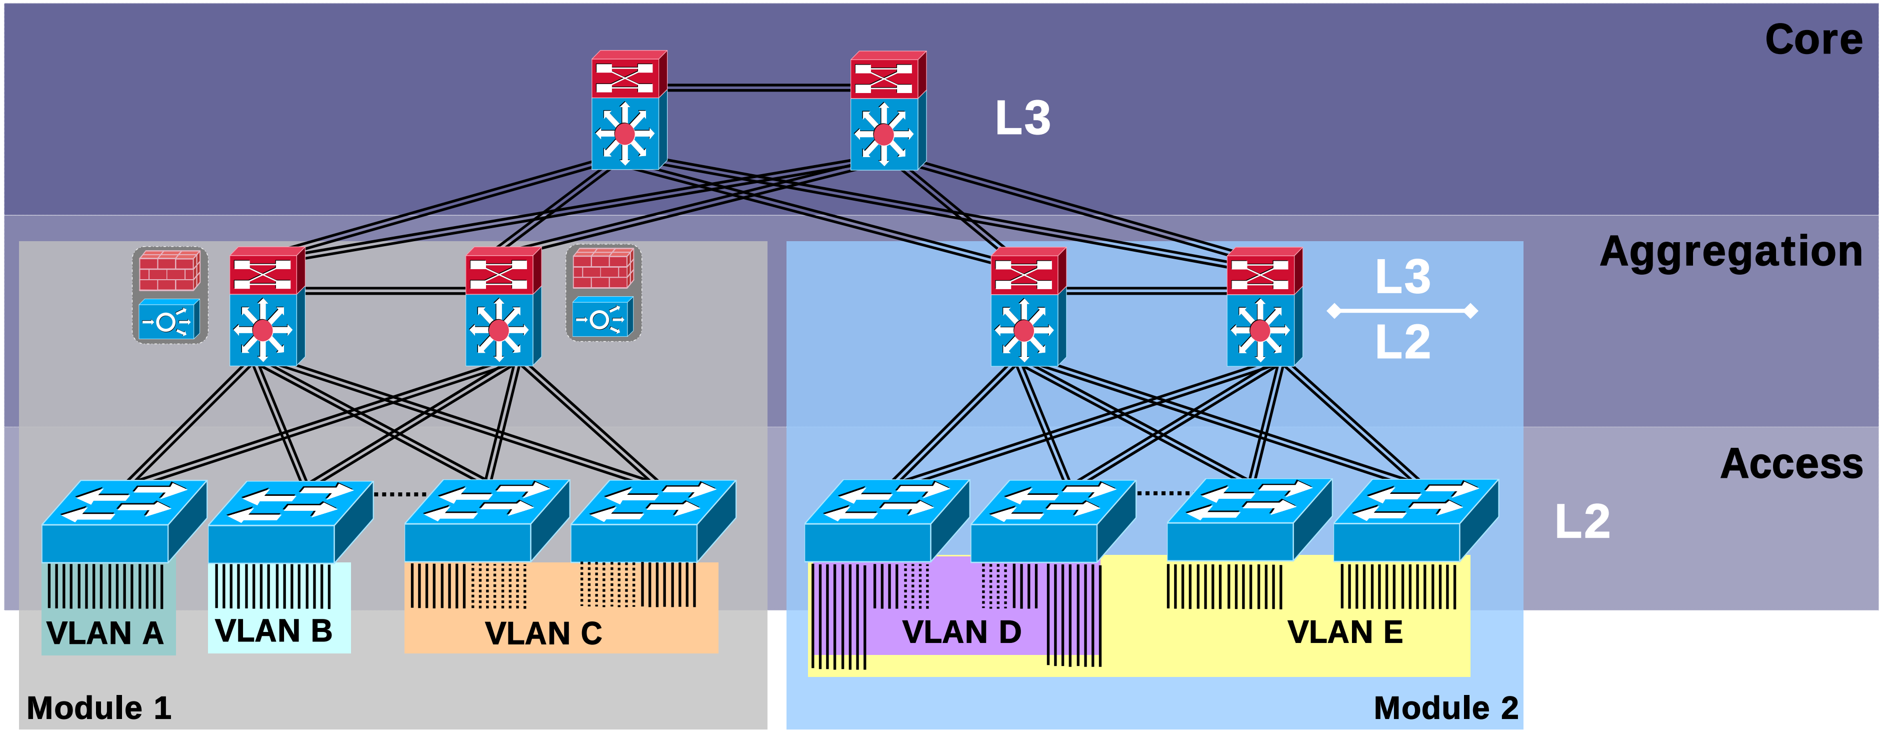

Hierarchical Design

Triangle and Square Topologies

Multiple Access Models: Modular, Blade Switches and ToR

Multiple Oversubscription Targets (Per Application Characteristics)

2000 – 10000 Servers

10,000 to 50,000 ports

New Topology

Classic Design

1. Common Topology – Starting Point

   Nexus at Core and Aggregation Layers

   2-Tier L2 topology

   VLANs contained within Agg Module

2. Topology Highlights

   Lower Oversubscription - if Needed

   Higher Density 10 GE at Core and Agg Layers

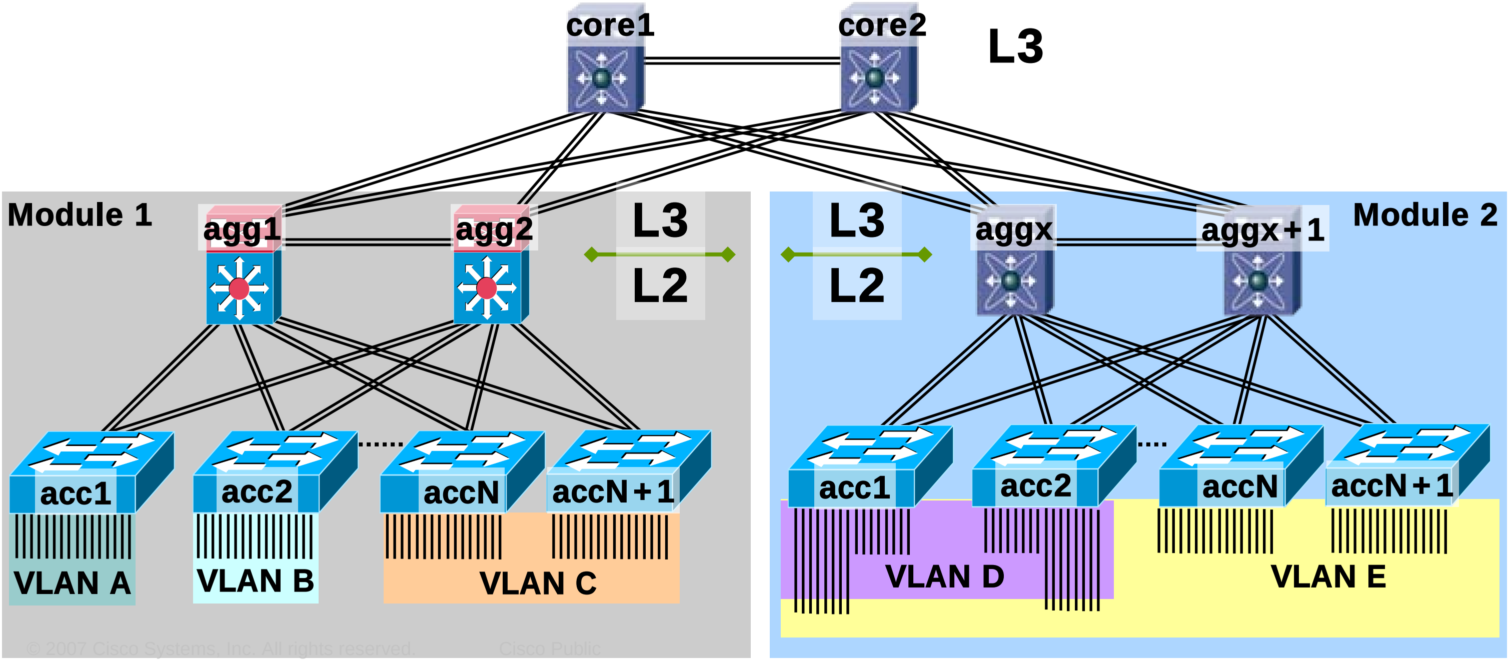

High Density GE Server Farms

10GE Aggregation and Server Farm Capacity

High Density Optimization Areas

- More Ports per Access Switch
  - i. New Supervisor with 10GE Uplinks
  - ii. From 240 to 336 Access Ports
- More Access Switches per Aggregation Module
  - i. 16-port on Catalyst 6500 and 32-port on Nexus 7000
  - ii. From 30 to 60 or 120 Access Switches
- More Aggregation Modules per Core Module
  - i. New I/O modules: 8 wire rate 10GE ports
  - ii. From 32 to 64 Wire rate 10GE port per core switch

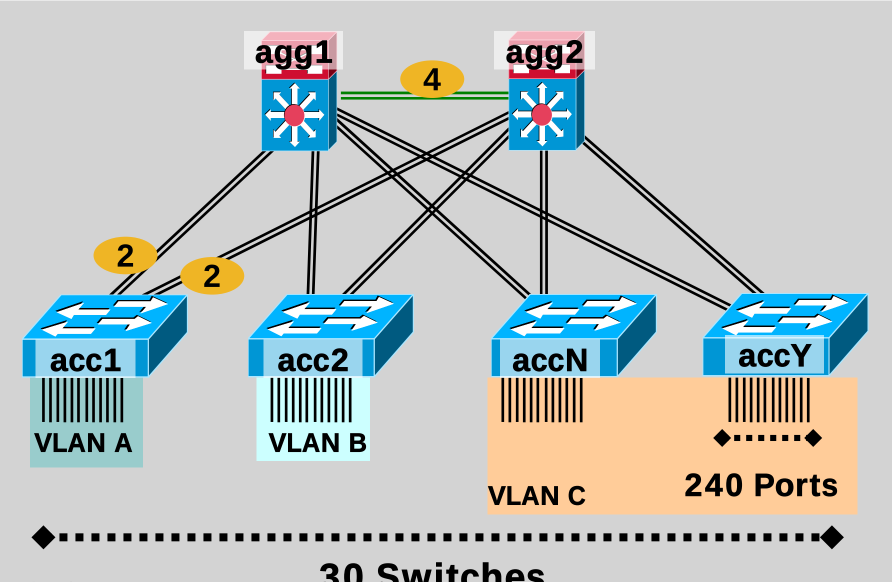

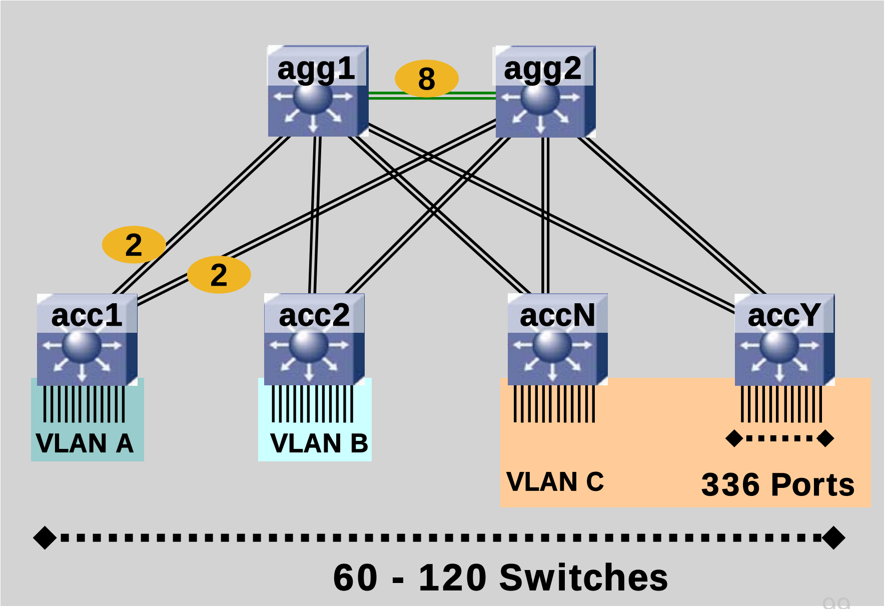

New Topology…

Enhanced L2 Design

- Enhanced L2 Topology
  - 3-tier L2 Topology
  - Nexus at Core and Aggregation Layers
  - 6500 at Aggregation and Services Layers
- Topology Highlights
  - DC-Wide VLANs
  - Higher Stability of STP environment – New STP Features
  - Lower Oversubscription - if Needed
  - Higher Density 10 GE at Core and Agg Layers

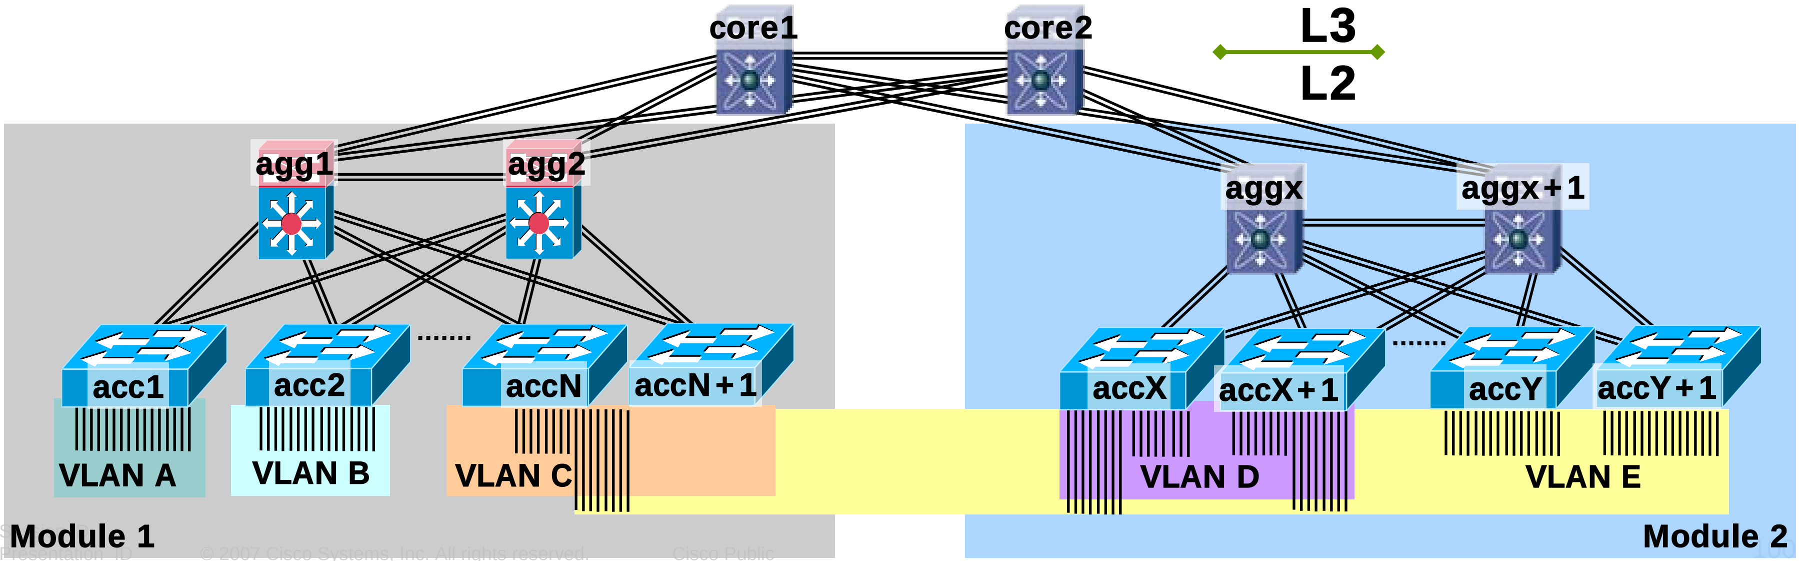

Enhance L2 Topology…

End to end Virtual Switching

- Enhanced L2 Topology
  - 3-tier L2 Topology
  - Nexus at Core and Aggregation Layers
  - 6500 at Aggregation and Services Layers
- Topology Highlights
  - DC-Wide VLANs
  - Higher Stability of STP environment – New STP Features
  - Lower Oversubscription - if Needed
  - Higher Density 10 GE at Core and Agg Layers

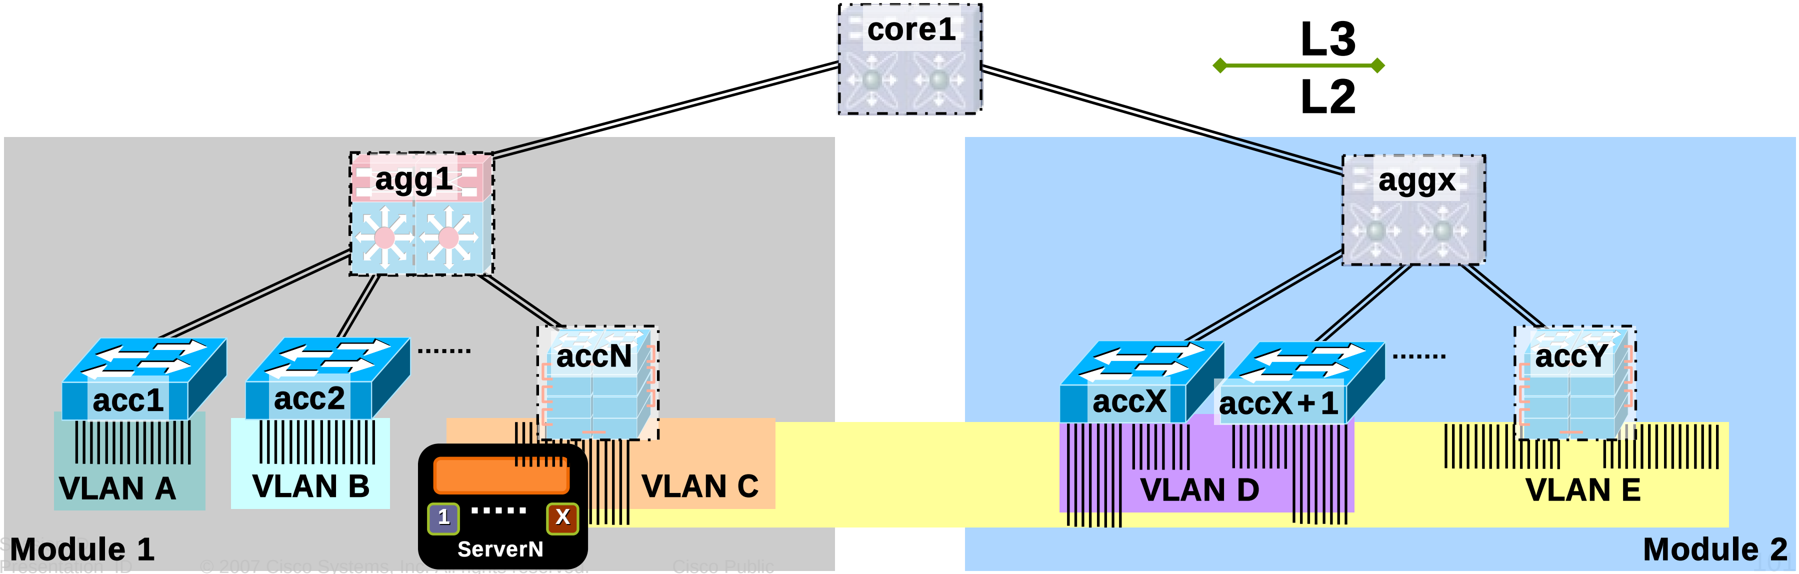

New Topology…

Classic Design + Integrated Services

Common Topology

- Nexus at Core and Aggregation Layers
- 6500 at Aggregation and Services Layers
- 2-Tier L2 topology
- VLANs contained within Agg Modules

Topology Highlights

- Lower Oversubscription - if Needed
- Higher Density 10 GE at Core and Agg Layers
- Services Integrated through Service Chassis
- Service Chassis & Virtual PortChannels through VSS

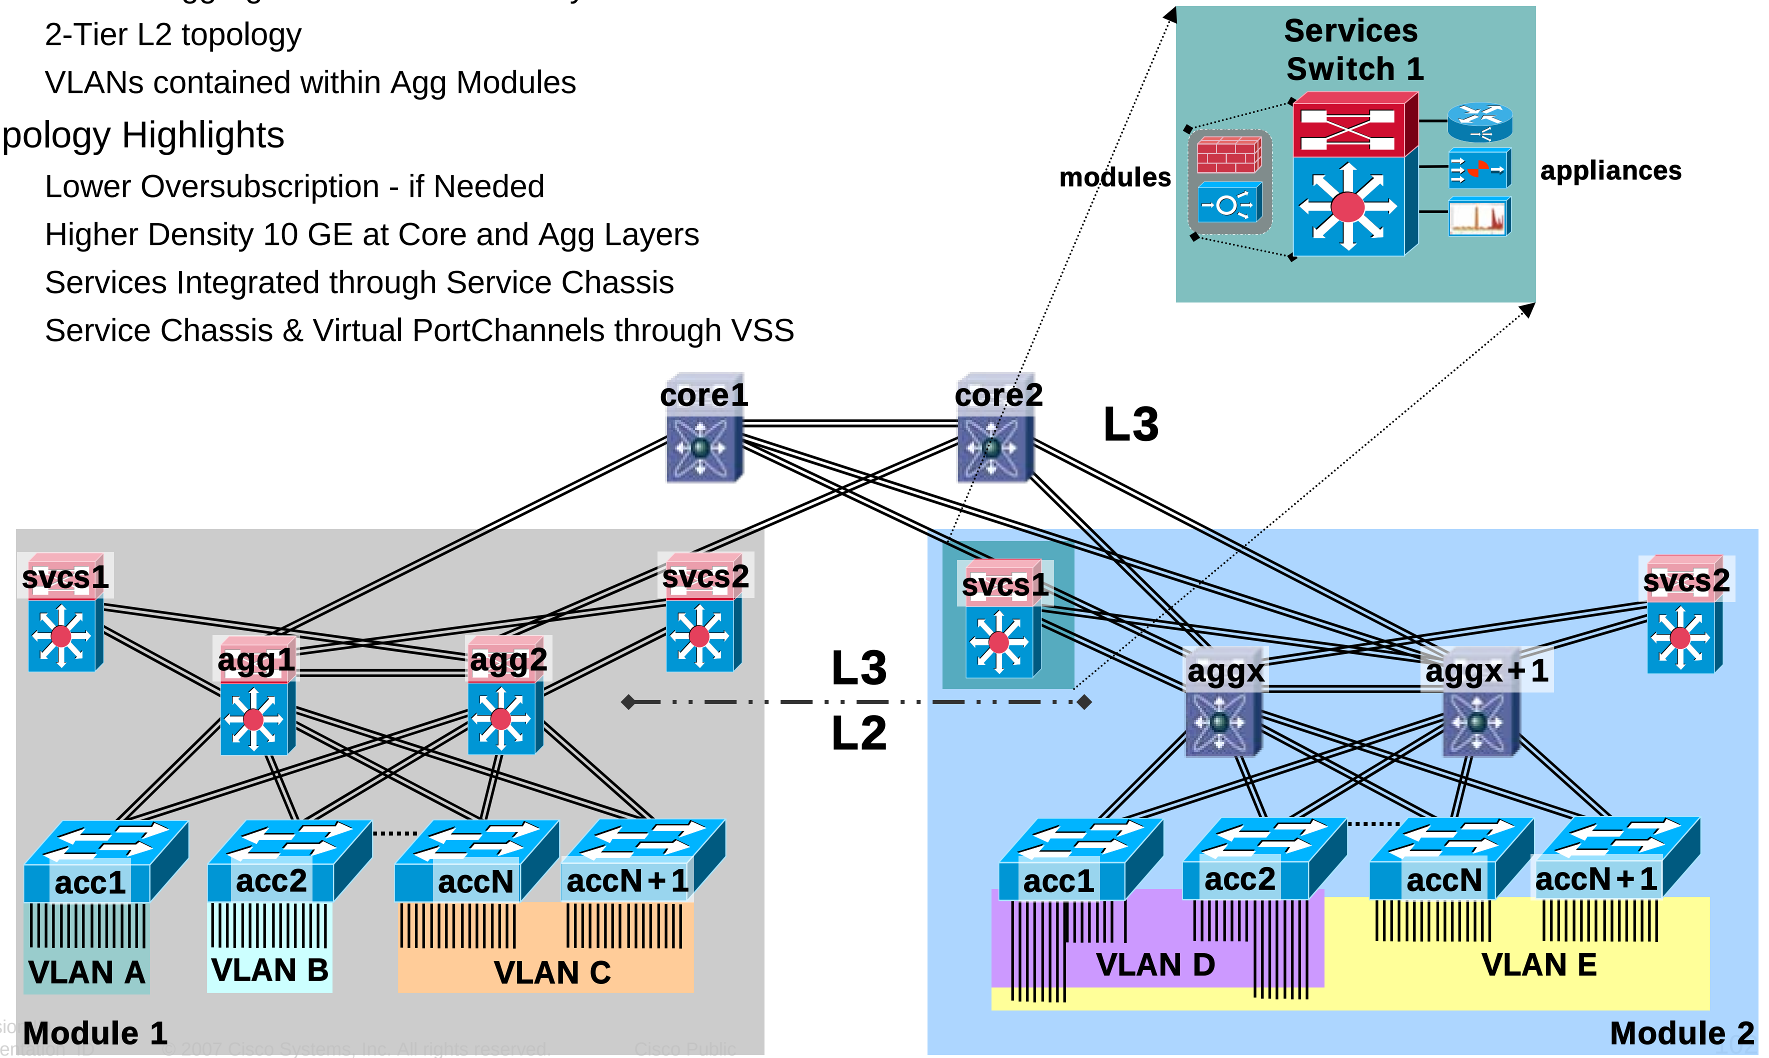

New Topology…

Classic Design, Integrated Services + Virtual Switching

- Service Appliances of Service Switches
  - Leverage Virtual Port Channels
  - Non-blocking path to STP root/HSRP primary
- Topology Highlights
  - Simplifies topology
  - Applies equally to
    - service appliances
    - Service chassis

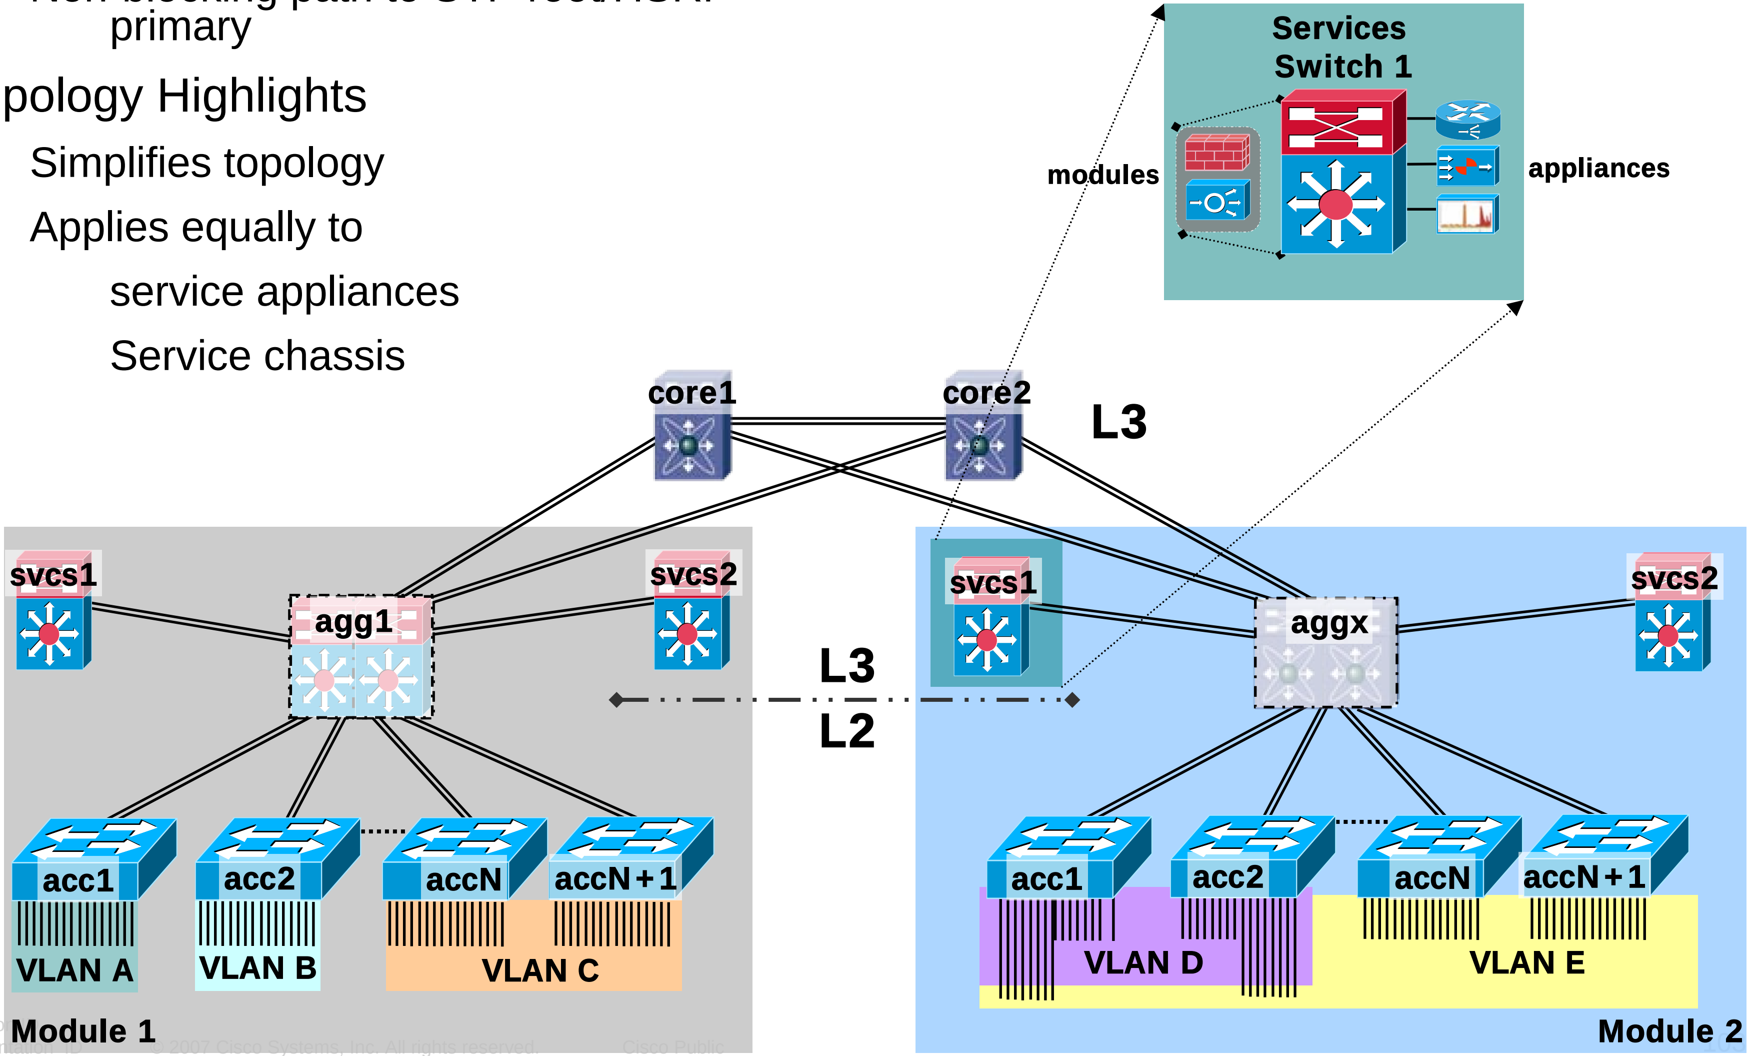

New Topology – Isolating Collapsed L2 Domains

Virtual Device Contexts @ Agg Layer

Pods are isolated at aggregation layer

- Each Pod runs its own STP instance (instance per VDC)
- Multiple pods could exist in a single VDC
- VLANs contained within Agg Module per VDC

Pods are logically isolated – two topologies

- Each Pod belong to multiple VDCs
- Each VDC topology requires dedicated Ports
- VLANs contained within Agg Module per VDC

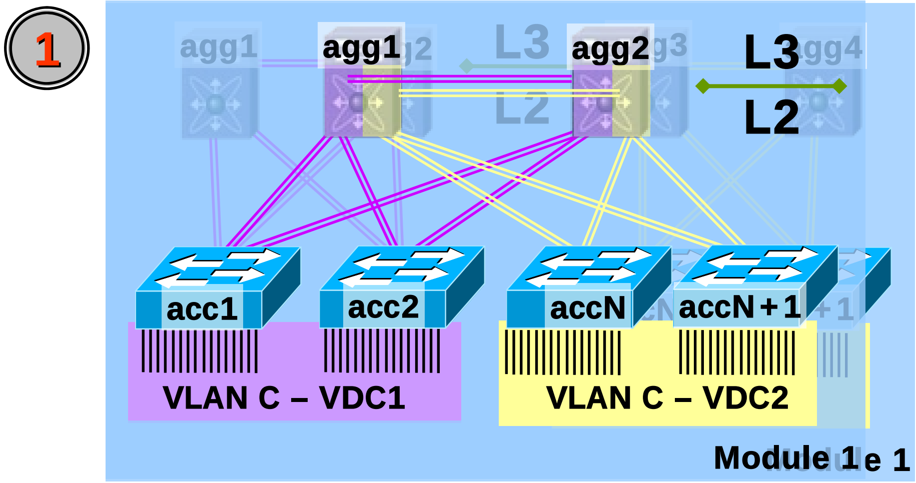

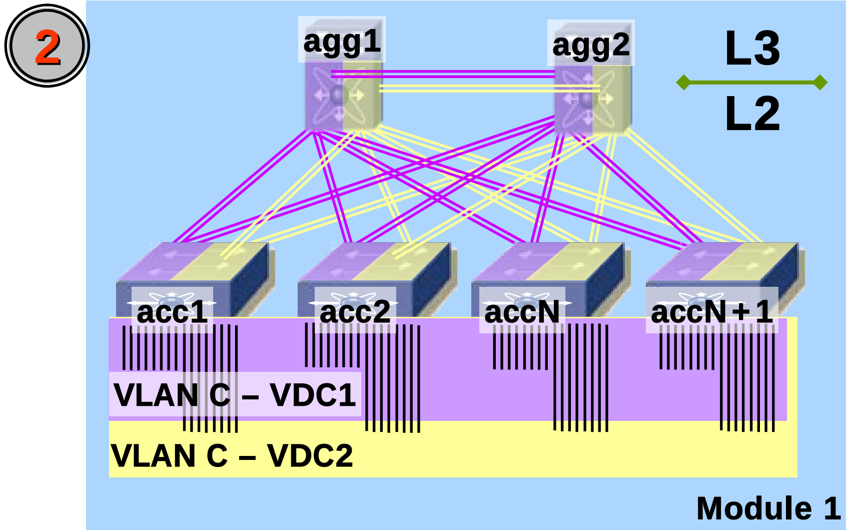

Higher 10GE Port Density Allows multiple Agg Pairs to be collapsed

Collapsed Agg Pair could still be L2 isolated (different STP instances)

VLAN IDs could be replicated on different VDC – shared infrastucture

New Topology – Enhanced L2 Collapsed Core

Virtual Device Contexts @ Core Layer

- Enhanced L2 Topology with Collapsed Core

  Benefits of 3-tier L2 Topology

  Zone are still isolated (An STP instance per zone)

  Core switches are managed independently by VDC

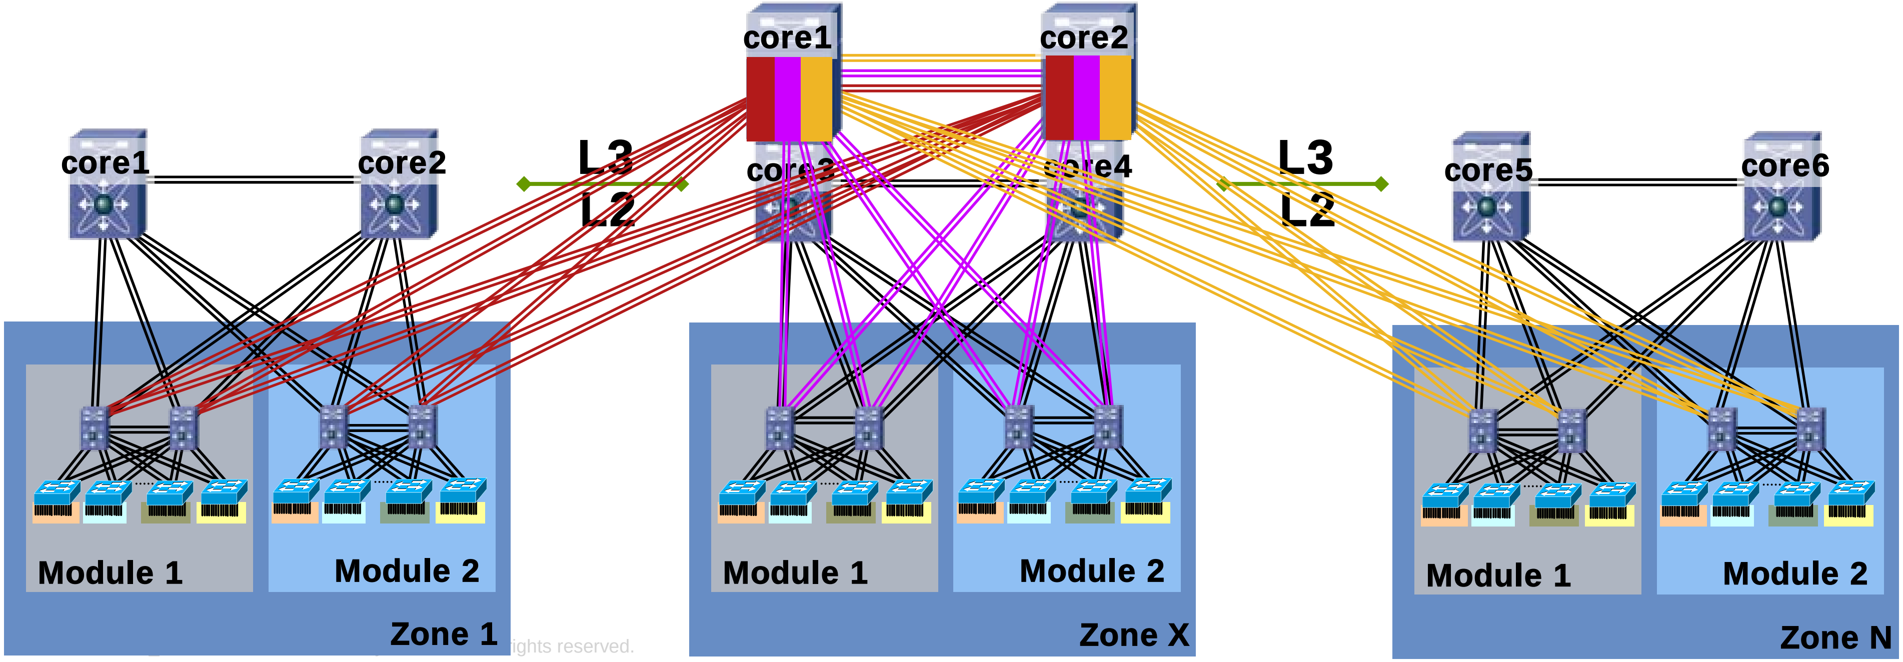
core1
core2
core1
core2
core3
core4
core5
core6
L3
L2
L3
L2
Module 1
Module 2
Zone 1
Module 1
Module 2
Zone X
Module 1
Module 2
Zone N