



Cisco Expo  
2008




Modern Data  
Center Design  
**Cisco DC 3.0**

Dré Van Brussel

Enable Your Network  
Empower Your Business

# Business Challenges for Connected Industries



Collaboration



Empowered User



Global visibility



Provisioning of new services



Reg. Compliance

## New Business Pressures



## Operational Limitations

Power & Cooling



Asset Utilization



Provisioning



Security Threats



Bus. Continuance



# Data Center and Network Evolution



# New Business-led Data Center design goals

## Traditional DC

Silo'd, fragmented

Box-based, manual provisioning

Optimised for expediency

IT as a cost centre

“Accidental Architectures”

## New DC3.0


Virtualised, collaborative

Service-based, orchestrated

Optimised for scale, agility, resilience

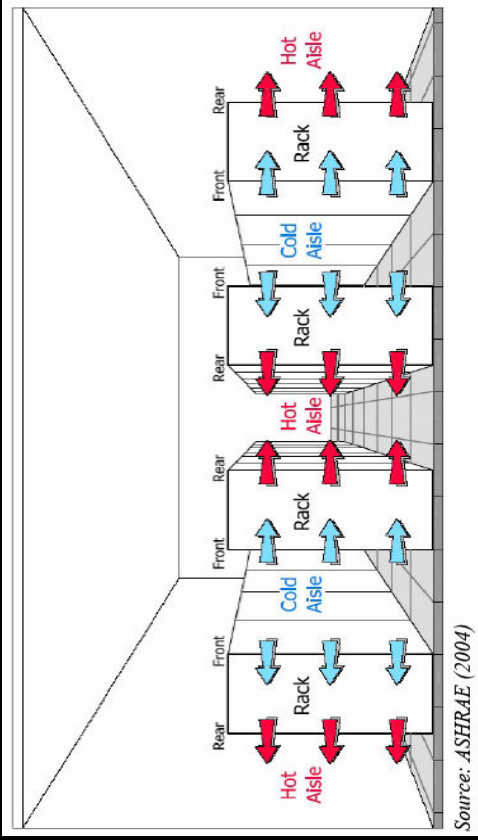

IT as a Service

Architected for the future





# Modern DC Design - Airflow

## Cold Aisle / Hot Aisle



### Cold Aisle/Hot Aisle Concept

### Air Circulation Components

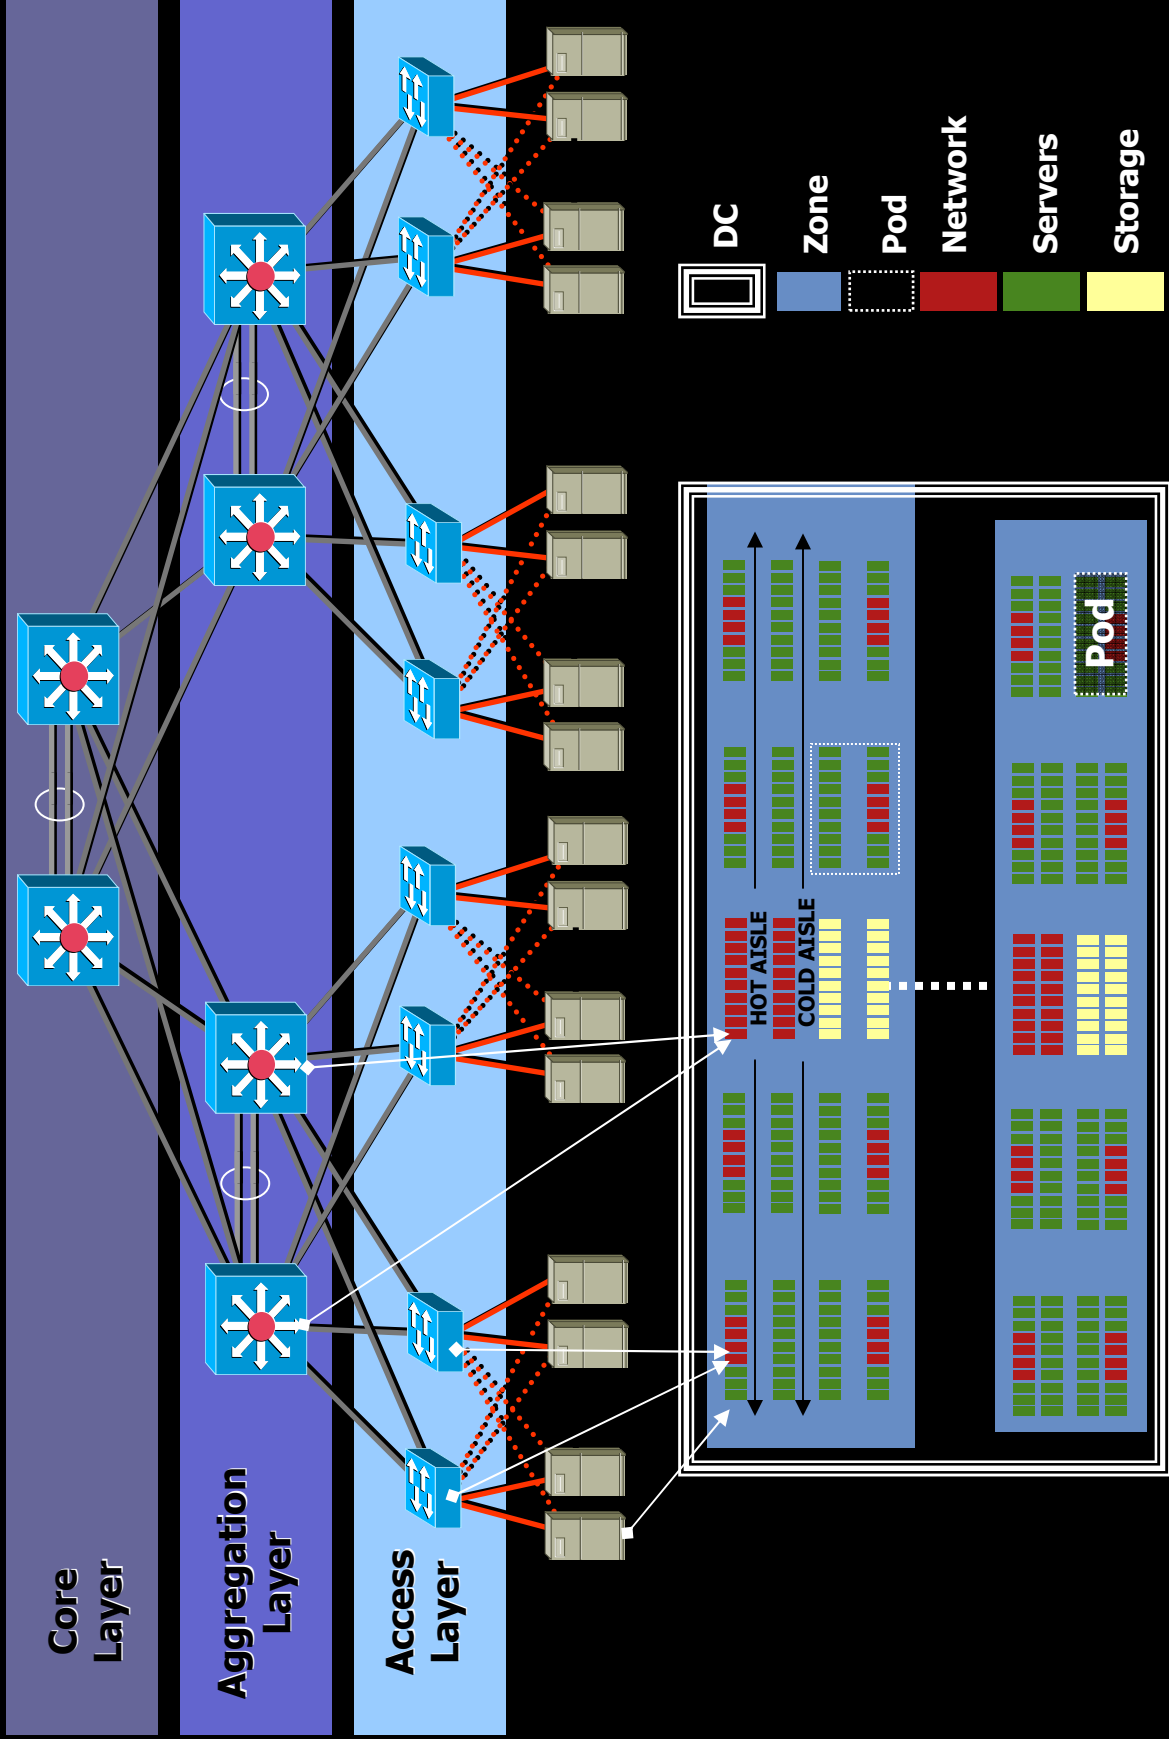


### Cable Placement

### Additional Cooling Capacity

# Data Center Topology

## Network Equipment and Zones



# Network Equipment Distribution

## End of Row and Middle of Row

### End of Row

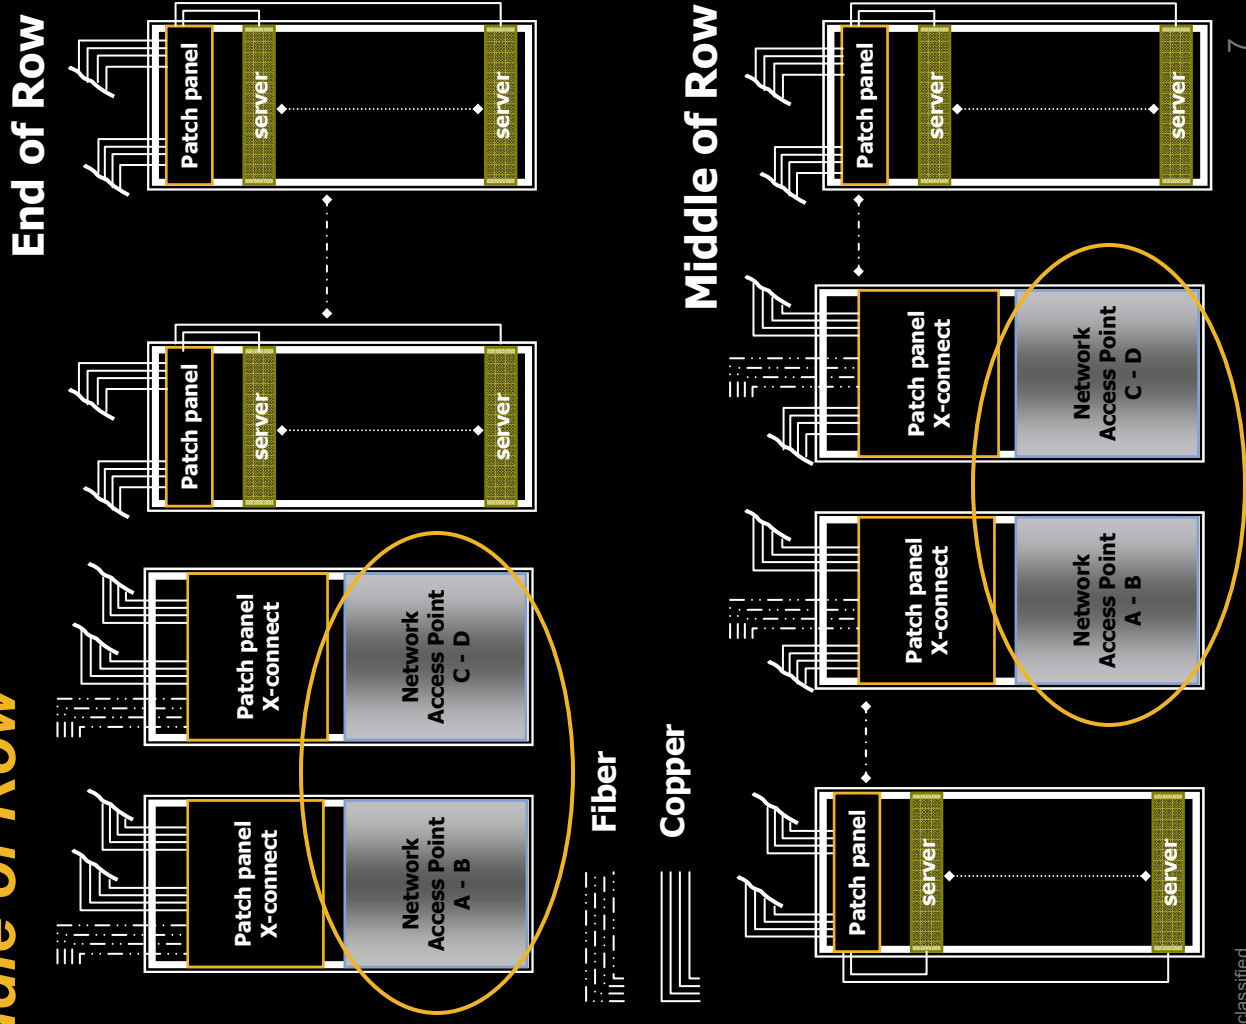
- Traditionally used
- Copper from server to access switches
- Poses **challenges on highly dense server farms**
  - Distance from farthest rack to access point
  - Row length may not lend itself well to switch port density
  - Compromised flexibility compared to MoR

### Common Characteristics

- Typically used for **modular access**
  - Cabling is done at DC build-out
  - Model evolving from EoR to MoR**
    - Shorter cabling distances (lower cost)
    - Allows denser access (better flexibility)
    - 6-12 multi-RU servers per Rack
    - 4-6 Kw per server rack, **10Kw-20Kw per network rack**
    - Subnets and VLANs: one or many per switch.
- Subnets tend to be medium and large

### Middle of Row

- Use is starting to increase given EoR challenges
- Copper from servers to access switches
- Fiber may be used to aggregate ToR
- It addresses aggregation requirements for ToR access environments



# Network Equipment Distribution

## Top of Rack

### Tor

#### Used in conjunction with dense access racks (1U servers)


- Typically one access switch per rack
- Sometimes two switches per rack + cluster
- Typically:
  - ~10-15 server per rack (enterprises)
  - ~15-30 server per rack (SP)
- Use of either side of rack is gaining traction

#### Cabling:

- Within rack:
  - Copper for server to access switch
- Outside rack (uplink):
  - Copper (GE): MoR for fiber aggregation
  - Fiber (GE or 10GE): more flexible, also MoR for aggregation

#### -Subnets and VLANs:

- one or many subnets per access switch
- Subnets tent to be small



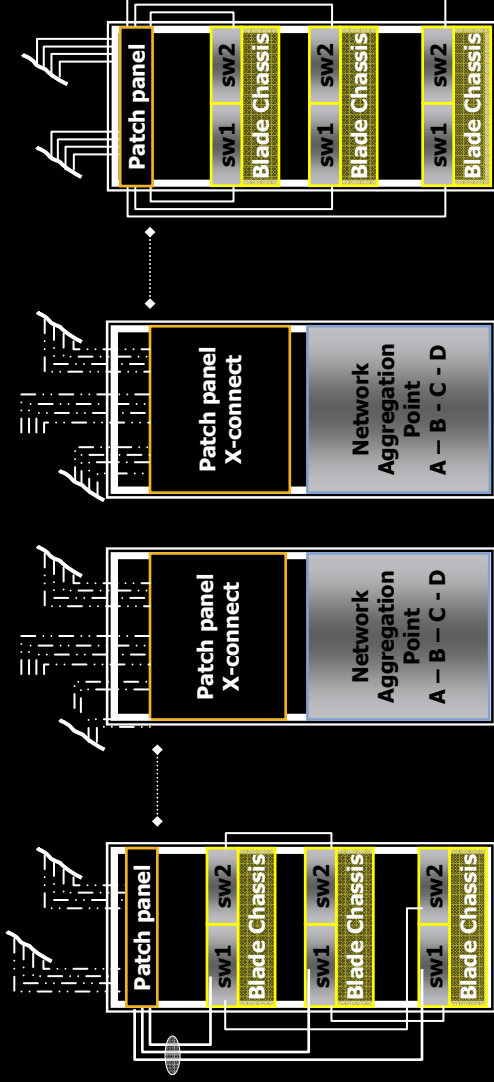


# Network Equipment Distribution

## Blade Chassis


### End of Row (Switch to Switch)

- Scales well for blade server racks (~3 blade chassis per rack)
- Most current uplinks are copper but the NG switches will offer fiber



### ToR

- No real use in conjunction with imbedded blade switches
- May be a viable option for **pass-through** environments if the access port count is high



### Middle of Row (Pass-through)

- Scales well for pass-through blade racks
- Copper from servers to access switches

# Next Generation DC drivers - Green ICT

Corporations are adopting standards to improve the **environmental performance** of their operations:

- Green building standards
- Environmental management systems and processes
- Citizenship and corporate responsibility reporting
- Network as the Platform

## Organization

**Green ICT =**

**80% Process, 20% Product**



# Next Generation DC drivers - more **bandwidth** to the Servers



Multi-Core CPU architectures allowing bigger and multiple workloads on the same machine



Server virtualization driving the need for more I/O bandwidth per server



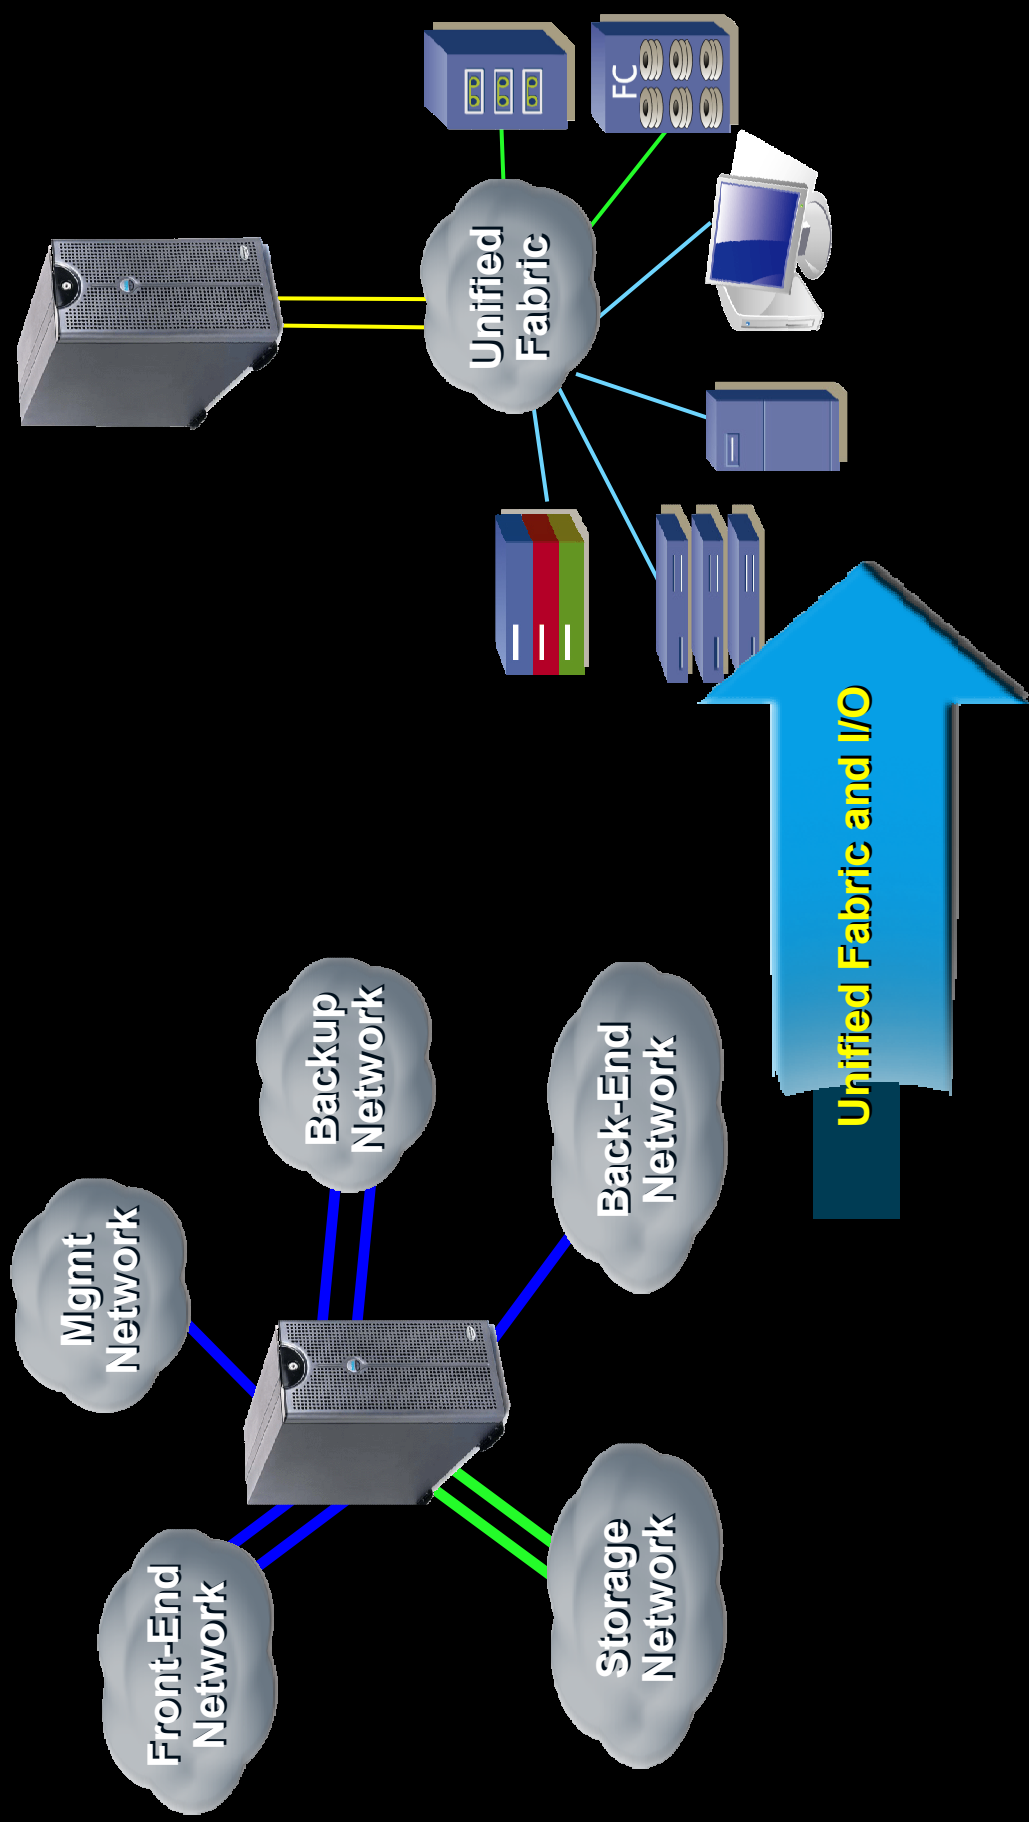
Growing need for network storage driving the demand for higher network bandwidth to the server



10GE LAN on server Motherboards (LoM) beginning mid-2008

**10GE LoMs, Multi-Core CPUs and Server Virtualization driving the adoption of 10GE network connections**

# Next Generation DC drivers - Unified Fabric and I/O for increased efficiency, simpler operations



# Critical core infrastructure components for Data Center 3.0



## Unified Fabric and I/O Interfaces

- Enables storage, Ethernet, IP, and HPC traffic to converge to a single network
- Lowers overall data center power draw
- Enables unified fabric management



## Cisco® Nexus Switching Platforms

- Multi-Terabit platform with ops-centric design
- Delivers unified fabric and I/O deployments
- Designed for the most stringent availability needs



## NX-OS Operating System

- First multi-protocol DC-class operating system
- Virtualized control plane and scalable design sets new standard for usability




## Data Center Network Manager


- Builds on Cisco Fabric Manager to extend L2, L3, Fibre Channel, and unified fabric management
- Provides end-to-end systems visibility

# Cisco Data Center 3.0 Strategy & Technologies

## Next Generation Data Center



Data Center  
Class OS




Performance  
& Density



Unified MGMT  
Architecture



### Modular Switching



### Blade Switching



### ToR Switching



Nexus  
7000


# Requirements Definition (1) Defining Data Center Class



## Continuous Operations

Hitless Software Upgrades, Zero Packet Loss  
Lossless System Backplane and Forwarding Architecture

Fault Tolerant and Loosely Coupled Architecture



## Virtualized Resources

Virtualized Control Plane

Virtualized Data Plane



## Operational Efficiency

Lights-Out Management Processors  
Graceful System Operations

Programmatic API for all functions

Integrated Diagnostics and Packet Analysis



## Power and Cooling Efficiency

Front to Back Airflow and High-Density Systems


Efficient Power Designs and Handling

# Requirements Definition (2) Supporting Data Center 3.0 Architecture



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-generation protocols**



Automation Improves Operations Effectiveness and Infrastructure Availability

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



# Nexus 7000 - First In Class Data Center Switch



## Data Center Class Platform

- Multi-Terabit system
- 550Gb/slot capable
- Optimized for 10 / 40 / 100 Gbps interfaces
- Extreme availability
- Multi-protocol (Ethernet, Storage and Unified I/O)

## Data Center Class Operating System

- Self Healing Operating system
- Graceful system operation
- Virtualized Control Plane and Data Plane
- Fully Modular
- Security

## Data Center Network Manager (DCNM)

- Unified Data Center Manager
- Configuration / Provisioning / Service Enablement / Network Ops / Status / Statistics / Event Management
- Powerful feature rich web services API (XML)

# Nexus 7010 10-Slot Chassis Optimized for the Data Center

## High density

256 10G interfaces/system

## High performance

1.2Tbps I/O system bandwidth initially  
80Gbps / slot, 60Mpps / slot

## Future proof

Initial fabric bandwidth capacity 4.1Tbps  
Product family scalable to 15+Tbps  
40/100G and Unified Fabric ready

## 2 Supervisor slots

independent Connectivity Management Processor (CMP) for  
lights-out management via 10/100/1000 port

## 8 Payload slots w/ Virtual Output Queuing (VOQ) \*

32-Port 10GE I/O Module w/ SFP+, 80Gbps fdx to Fabric  
48-Port 1GE I/O Module w/ RJ-45, 40Gbps fdx to Fabric  
Advanced forwarding engine h/w – 60Mpps, QoS, NetFlow,  
GRE tunnels..

## 5 Fabric Module slots


230Gbps / Payload slot  
115Gbps / Supervisor slot  
Load-sharing across all fabric modules in chassis  
Multilevel redundancy with graceful performance degradation \*  
Non-disruptive OIR



# Nexus 7000 – 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
- 4<sup>th</sup> and 5<sup>th</sup> fabric modules provide additional level of redundancy
- Future modules will leverage additional fabric bandwidth
- Fabric failure results in reduction of overall system bandwidth

Fabrics



230Gbps

Module  
Slots

1G Module

10G Module


# What Is VOQ?

Virtual Output Queues (VOQs) on ingress modules represent bandwidth capacity on egress modules

Module 1



## Ingress module




VOQ ensures fair access to fabric bandwidth

## Fabric module




Egress Capacity  
(ability to receive traffic from fabric)

Module 2  
(1G module)



## Egress modules



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



Module 3  
(10G module)



Module 4  
(10G module)



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

# Cisco High End Switching Portfolio



## Catalyst 6500 and Nexus 7000



### Nexus

230 GbE / Slot  
10G optimized

15T Switching; 500G+ / slot  
10G/40G/100G optimized  
Unified Fabric (DCE)



### Catalyst 6500

720G switching; 40Gb /slot  
1Gb/10Gb optimized  
Service modules

2T Switching; 80Gb+ / slot  
1Gb/10Gb/40Gb optimized  
Service Modules



# What others say about the Cisco Nexus 7000

## Nexus 7000: The best-ever enterprise product from Cisco

*The Nexus 7000 is a multipurpose data center product that integrates storage and Ethernet features. I believe that the Nexus 7000 is the best enterprise-focused product ever to emerge from Cisco in terms of software architecture, focus on high availability, capacity, performance, and security.*

*While network operator and enterprise features remain different in some areas, the Nexus 7000 demonstrates the commonality that now exists in the areas of high performance and high availability.*

**Mark Seery, Vice President of switching and routing at Ovum RHK**

# What NX-OS brings to the table...

## **Availability**

ISSU+ (In Service Software Upgrades)  
Process Modularity  
Optimized architecture  
(Reliable IPC, Dual EOBC)

## **Serviceability**

Call Home  
GOLDS (Generic OnLine Diagnostics)  
EEM (Embedded Event Manager)

## **Manageability**

Programmatic XML Interface  
Data Center Network Manager  
RBAC (Roles Based Access Control)  
Configuration Verification and Rollback

## **Virtualization**

VDC (Virtual Device Contexts)  
Built-in Feature Virtualization  
(VDC/VRF awareness)

## **Security**

Security Group ACLs (User communities)  
Link-layer Cryptography (TrustSec)  
CoPP (Control Plane Policing)  
CISF (Cisco Integrated Security Features)  
DHCP Snooping / Dynamic ARP inspection / IP source Guard

## **Scalability**

Scale

Distributed, symmetric-multiprocessor, multi-threaded

Density

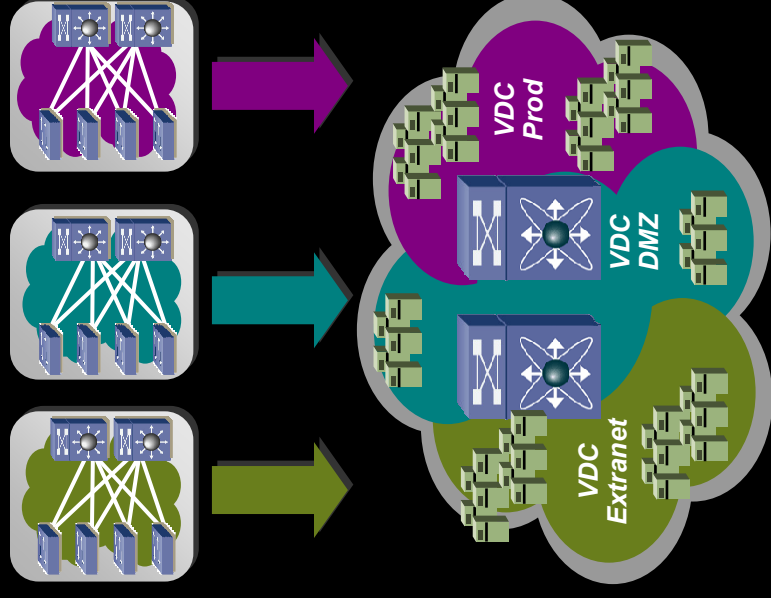
128K logical interfaces (LIFs), 10K+ ACLs, 16K VLANs

## **General**

Comprehensive L2 & L3 features  
Industry leading and standards based

# Virtual Device Contexts (VDCs) benefits

- **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




Network “islands” are virtualized onto common datacenter networking infrastructure



# Cisco TrustSec on Nexus Network Device Admission Control

Assumes Network Devices (Routers, Switches, etc.) are not trusted by default and will need to authenticate themselves to a trusted node via 802.1X with EAP-FAST infrastructure prior to any network connectivity becoming available. Once Authenticated, polices are then acquired as well as any applicable session keys or cipher suites...

## Network Device Admission Control



# 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

# DCNM Server Network Model



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

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

# NX-OS/Nexus High Availability

- Hardware provides redundancy at every component level:
  - Supervisors**
  - Power**
  - Fabrics**
  - Fans**
- Software offers multi-layered, multi-faceted resiliency:
  - Stateful process restarts**
  - Graceful restart for routing protocols**
  - Stateful supervisor engine switchovers**
  - True in-service software upgrades**

**Nexus 7000 + NX-OS =**

**Hardware and software combined to deliver data-center class high availability – zero service disruption**



# Data Center 3.0 and Power Efficiency

## Investment in the Network drives Greater Utilization

### Network Enabled Storage Consolidation

1. Consolidate SAN's and Slow Growth
2. Implement Virtual Fabrics
3. Implement Storage Virtualization


IVR

Storage Services

Load Balancing

### Data Center 3.0 Allows for:

1. "Silo-busting"
2. Homogenous services
3. Better power efficiencies
4. Power management
5. Less disruptive scalability



# Improving IT Responsiveness

Adapting to the changing requirements of business



No Virtualization



Data Center  
Virtualization




Static Service  
Orchestration



Service  
Orchestration


# VFrame Data Center

## Graphical user interface



2

## Server Agent



4

## VFrame appliances

VFrame Data Center  
(Primary)




Active synchronization  
Automated Failover

VFrame Data Center  
(Secondary)




1

## SOAP/XML API



3





## Macros



5

# VFrame's Workflow for Service Oriented Data Center

## Logical, Structured for Ease of Use





## Slide 32

---

### GC1

Good if in the "deploy" prt of the script for this slide, you define the following: DHCP, VIP, LB, LUN, NFS

Geta Carlsson; 6.7.2007

### GC2

Image mgmt  
Remote Boot  
VM mappings

Commas missing here--should this be Image mgmt remote boot, VM mappings

or

is each of these a separate item?

Geta Carlsson; 6.7.2007

# Storage Networking Consolidation & Virtualization Case Study: Petrobras

## Challenge

- Implement a core edge architecture to facilitate future growth and avoid disruptive interventions
- Virtualize storage to reduce TCO

## Solution

- Joint Cisco/EMC solution incorporating:

### Core-Edge architecture

**Consolidation:** VSAN technology to consolidate multiple SANs into a single network

**Backup:** Backup VSAN with IVR and FCIP technologies to allow inter-site backups

**Virtualization and NASB:** SAN prepared for future implementations

## Benefit

- Reduced TCO through total manageability and fabric virtualization
- Improved performance, troubleshooting and provisioning of new services



# High Performance Computing Case Study: ENI

## Challenge


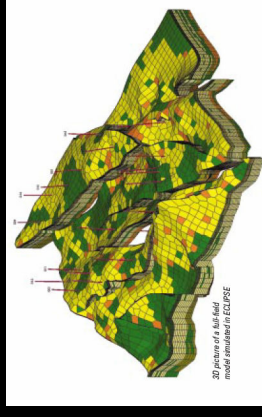

- Find and access new oil & gas reserves while reservoirs get more complex
- Finish simulation jobs faster – hours compared to days
- Procure more cost-effective hardware and reduce cost of ownership

## Solution

- High Performance Computing System formed by 8 clusters, 6x256 nodes and 2x512 nodes, for a total count of 2560 server nodes and 10.240 CPUs.
- Joint Cisco/IBM solution incorporating SFS and Catalyst switches and IBM blade servers

## Benefit

- The computational power will be mainly used for seismic data processing
- With this state of the art solution Eni will be more competitive in international tenders for gaining exploration rights



Landmark Graphics VIP: 3D View of Reservoir Simulation – Coarse Models for Risk Assessment

# Application Networking Services

## Case Study: Holcim Croatia

### Challenge:

- Many remote offices across the country with limited bandwidth
- Poor application performance from data center to branches

### Solution:

- Cisco WAAS optimizes bandwidth from data center to branches

### Benefit:


- Significant cost reduction in bandwidth
- User experience improved significantly
- Faster roll out of new centralized applications to branches



# Cisco DC IT: Storage Utilization

## Improving Efficiency through SAN Consolidation


2005



~2 PB's Storage

Multiple SAN Islands  
Isolated Fabrics  
Over-provisioned Power & Cooling



Today



~5 PB's Storage

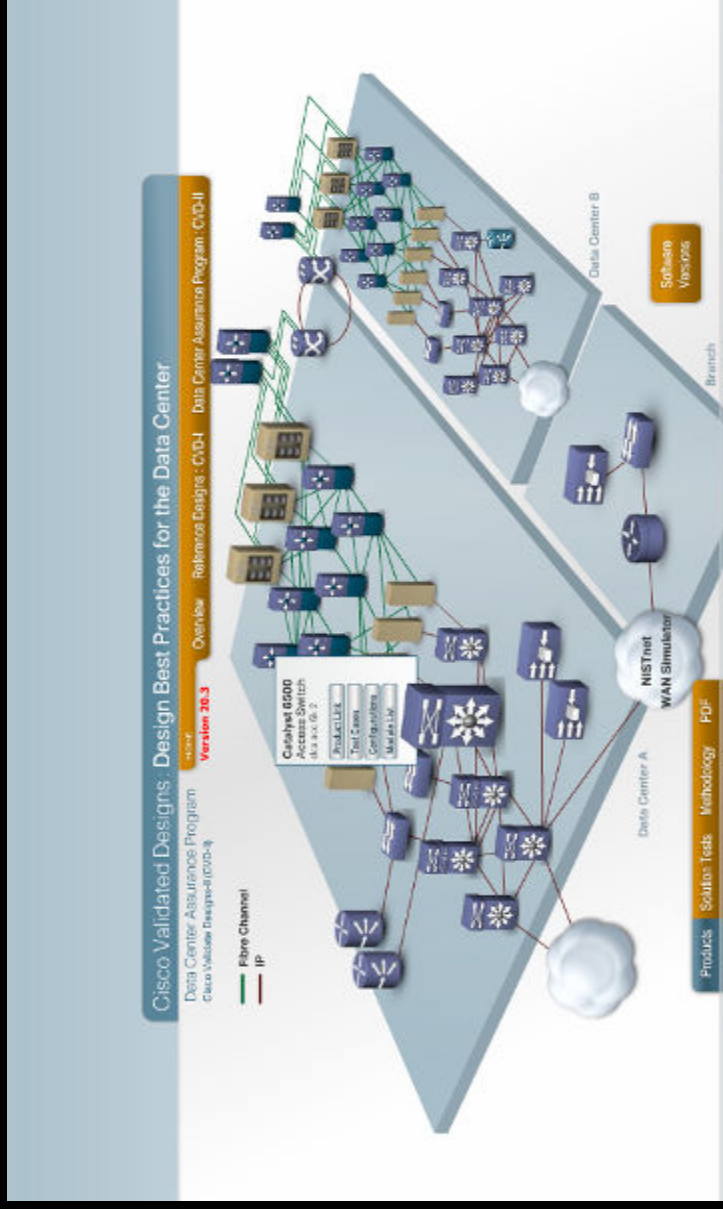
Inter-vSAN Routing  
Virtual Fabrics  
Targeted Power & Cooling

Common Physical Fabric



# Data Center Assurance Program


- Design Best Practices
- Real-world Tested Configurations
- Downloadable Spec's and Results
- Full Test Plan and Documentation Kit
- Intuitive 3D graphical interface
- Testing updated Quarterly



<http://www.cisco.com/go/datacenter/dcap>

- Best Practice Design Zone
- Integrated Discussion Forum
- Operational Best Practices'


# Cisco Data Center Technology Enablers



# Cisco Data Center Technology Enablers

- DCNM
- VFrame
- ISC
- Virtual Machine Integration

- MDS 9500
- MDS 9100
- Storage Services



- MDS 9500
- SSM
- Cisco Routers (network services transport)
- Cisco Switches (network services)
- ONS 15454
- FCiP

- WAAS
- ACE
- ACE Appliance
- GSS

- Catalyst 6500
- Catalyst 4900
- Nexus Family
- HPC SFS Infiniband

- Nexus Family



# Cisco: integral Partner in Your evolving Data Center

- leading provider of **network-based services for holistic data center operations**
- **Strong partnerships** with other important data center vendors
- Fully tested end-to-end solutions : Cisco **Data Center Assurance Program (DCAP)**
- Consistent record of **innovation**, eg **Nexus 7000**
- **Financial stability**
- **Award-winning customer support**

More information

<http://www.cisco.com/go/datacenter>



