




Cisco Expo
2009


Data Center 3.0 Architecture & Evolution (1^{ère} Partie)



Hicham El Alaoui - helalaou@cisco.com

Agenda

- Tendances et Défis du Data Center
- Unification des Réseaux avec FCoE
- Nouvelles Architectures du Data Center avec le Nexus 5000
- Agrégation 10GE avec Nexus 7000
- NX-OS
- Q&A



Vers une Puissance Concentrée et Virtuelle

1. Plus de Densité des Serveurs*


2. Plus de Cœurs dans les Processeurs*

Après 2008 Intel va délivrer **exclusivement** des processeurs à plus de 4 cœurs pour les serveurs


3. Plus de Serveurs Virtuels*

*Source: IDC 2007

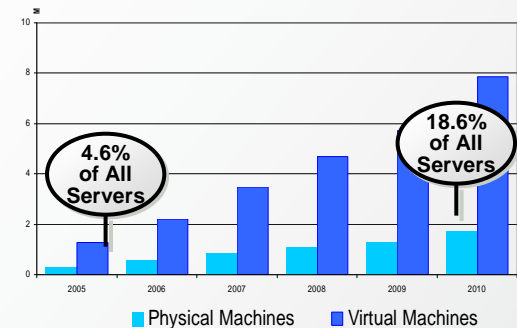
Forme des Serveurs livrés*
(Marché Mondial des Serveurs)



Adoption du Multi-Cœur*
(Marché Mondial des Serveurs X86)



Taux de Virtualisation Des Serveurs*



Les CPUs Multi-Cœur et la Virtualisation des Serveurs Pousse la Demande pour des Connexions Réseau Denses et plus Rapides

Défis de la Puissance Concentrée



Le Nombre Important de Cartes Réseaux **Empêche** l'Adoption des Blades et des Serveurs Rackables



Le **Câblage** dans les Salles Blanches Devient un Vrai Casse-tête




Le Prix Elevé des Cartes HBA FC Freine le **Développement** des SAN



La Duplication des Cartes Réseaux implique Plus de Consommation **Electrique**

Solution : Unification des Réseaux (Unified I/O)

Utilisation des CNAs (Converged Network adapters) au lieu des NICs et HBAs



Avantages de l'Unification des Réseaux




Réduction du Nombre de Cartes Réseaux



Les Serveurs sont Câblés **une Seule fois** puis Connectés à n'importe quel Réseau LAN ou SAN





Tous les Serveurs vont pouvoir Monter des Disques SAN



Consommation Electrique Réduite d'environ **8%**

Fibre Channel over Ethernet (FCoE)

- Méthode pour transport les trames FC sur Ethernet
- Les trames FC sont inchangés
- Pas de translation de Protocole
- FCoE apparaît comme du FC pour les Serveurs et les Baies de Stockage (contrairement à l'iSCSI)
- Préserve l'Infrastructure Existante et les Systèmes d'Administration



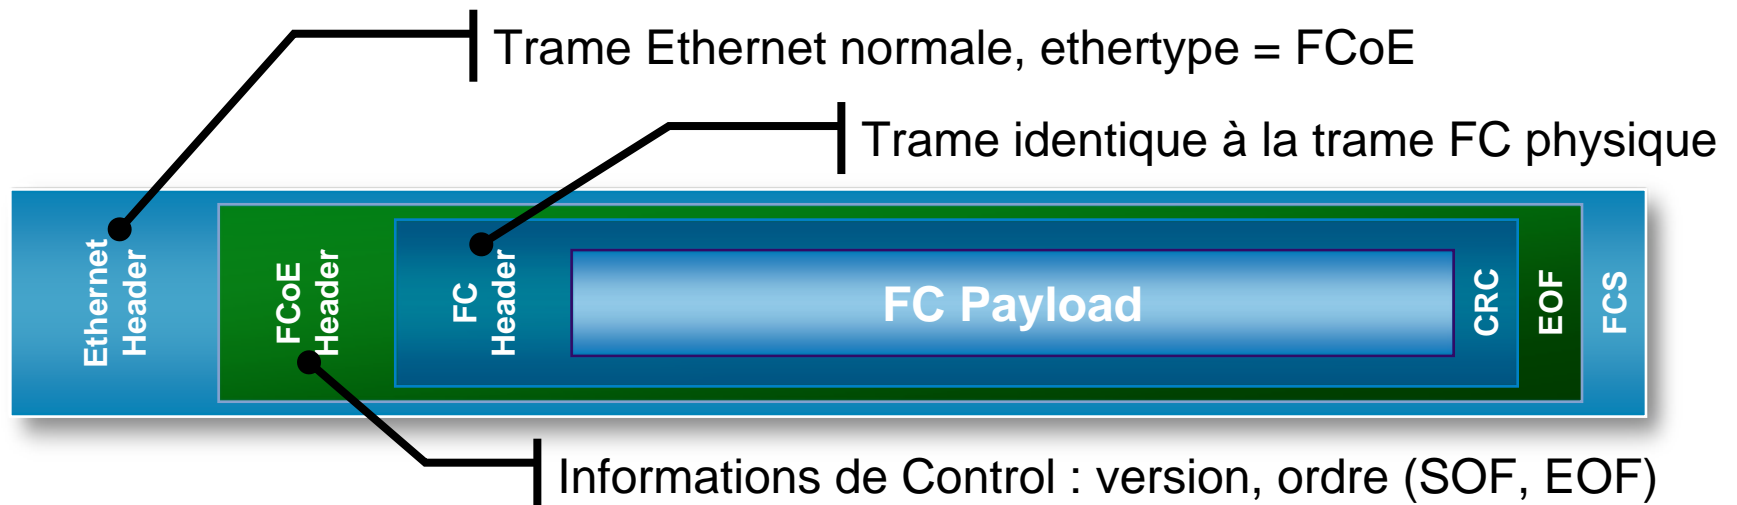
Pourquoi FCoE est Possible Aujourd'hui

1. 10 Gbps Ethernet
2. Lossless Ethernet (*Ethernet Sans Pertes*)

Fournit à Ethernet le même comportement "sans-pertes" qui est garanti par le système "buffer-to-buffer credits" du FC

3. Trames *Jumbo* dans Ethernet

Max FC frame = 2112 bytes



Standard Soutenu par Toute l'Industrie

Progrès de la Spécification FCoE :

Adoptée par le ANSI T11 FC-BB5 en Juin 2007

Format de la Trame adopté en Août 2007


Démontré au Storage Forum en Octobre 2007 (NetApp, QLogic, ...)

Ratifiée le 3 juin 2009



Du Côté des Cartes Réseaux

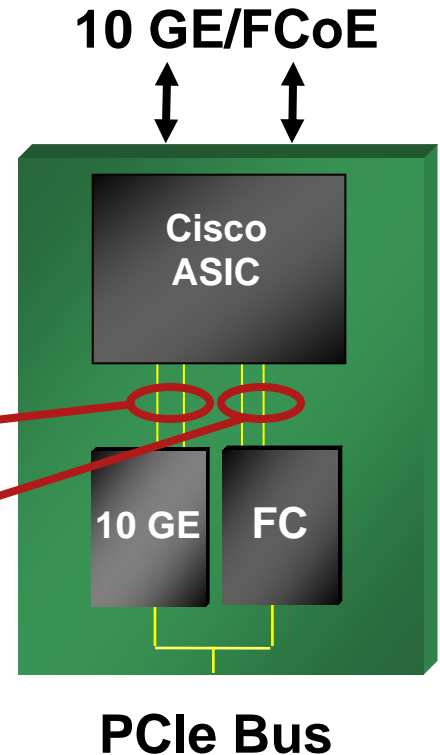
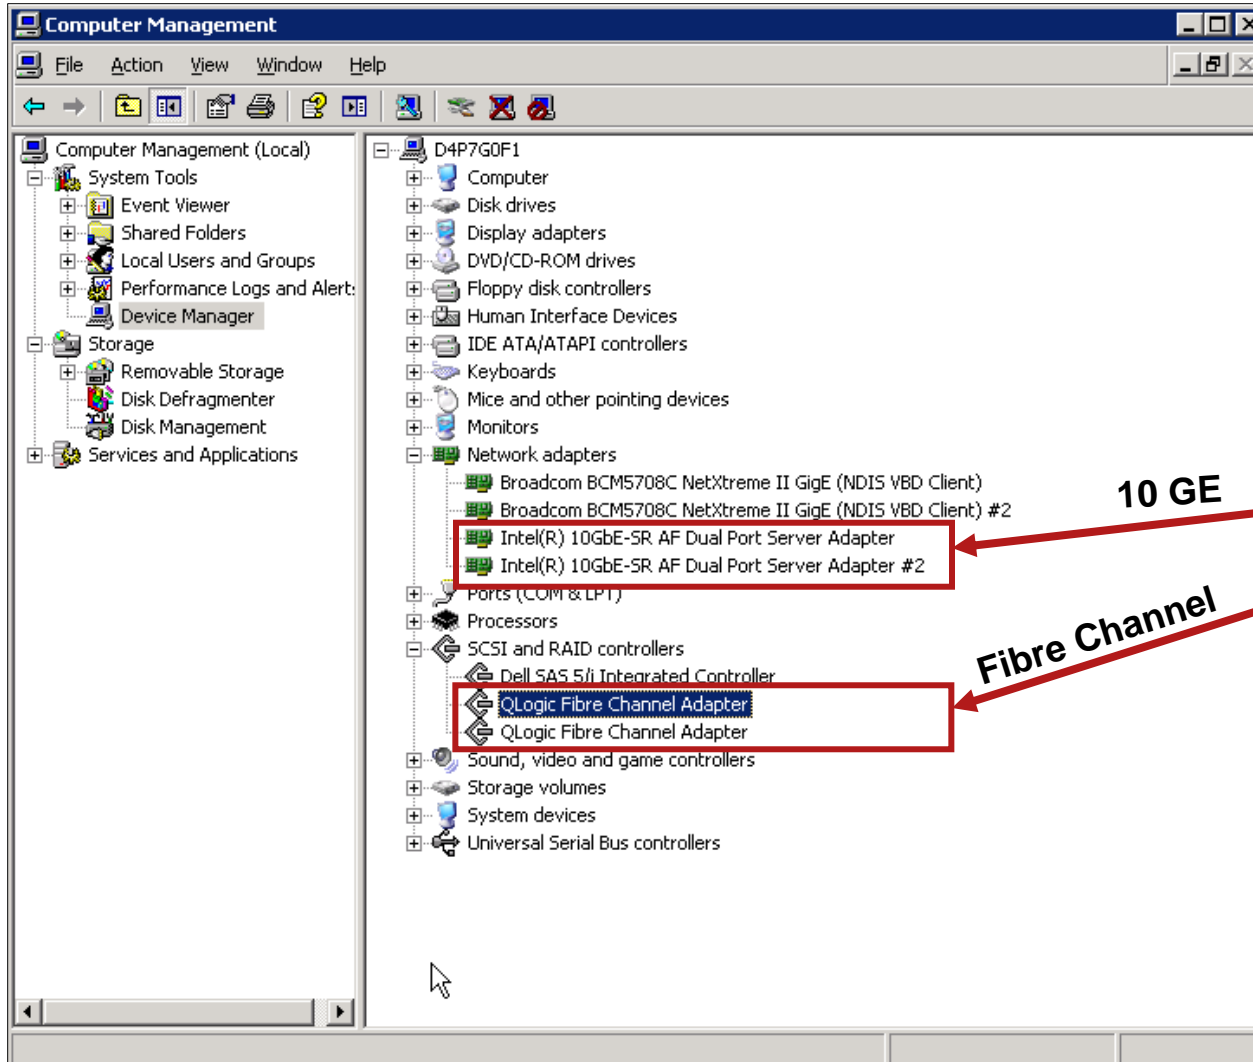
Solution Mixte Carte 10GbE + Pile FCoE



Solution Matérielle NIC + HBA + MUX

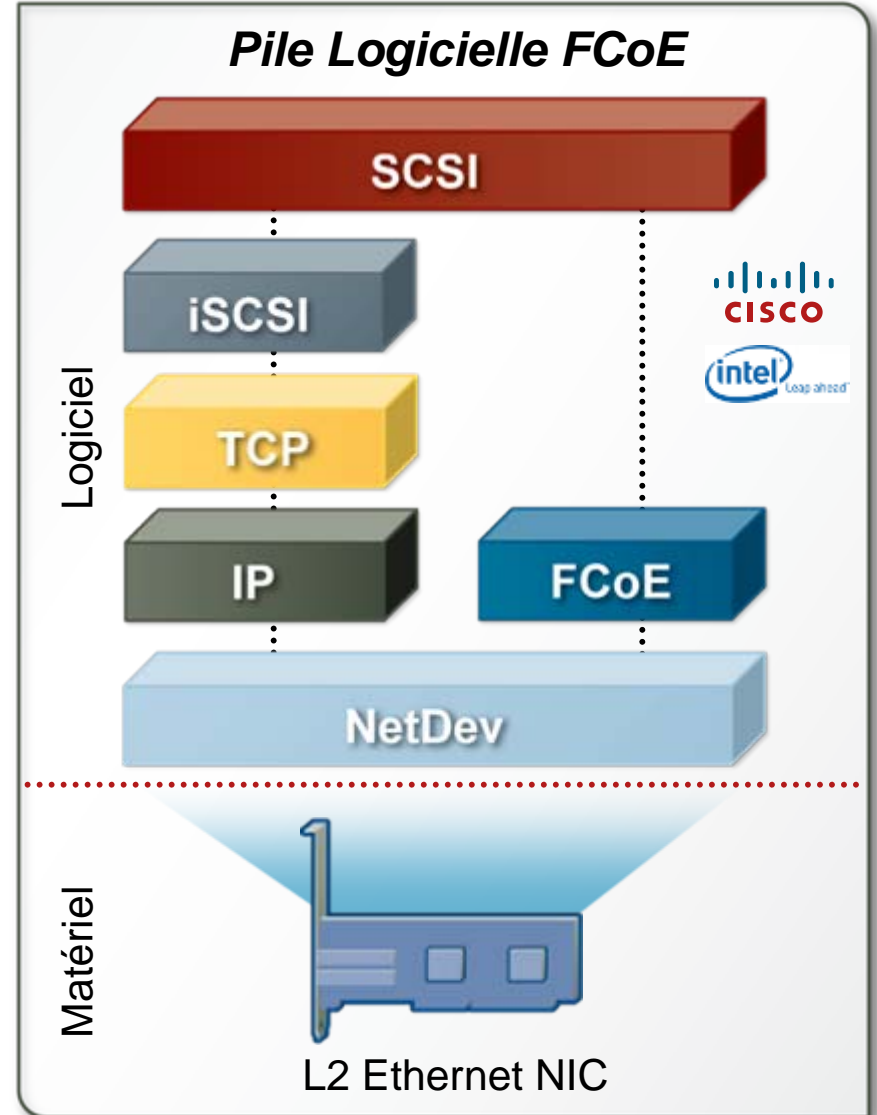


Comment le Serveur voit le CNA ?



Pile Logicielle FCoE

- Supporté sur les cartes 10 GigabitEthernet de Intel
- Implementation Logicielle à 100%
- OS Supportés
 - Linux: Redhat & SLES
 - Windows
- Accès “Gratuit” au SAN



Web site: www.Open-FCoE.org

Evolution de l'Architecture



Phase 0 (Actuellement)



Segregated LAN and SAN


In current architectures, LAN and SAN connectivity is segregated directly from the Servers, where NICs and HBAs connect into Ethernet switches and Fibre Channel Fabrics. This may result in excess of 8+ cables to/from each physical server

In Ethernet, redundancy relies upon technologies such as Spanning Tree Protocol to provide a loop-free topology...



Evolution de l'Architecture

Phase 1



Unified Fabric through FCoE

The Nexus 5000 allows for the consolidation of Ethernet and Fibre Channel to be carried across the same physical piece of cable - Ethernet


Leveraging FCoE technology, the Nexus 5000 is able to provide direct FCoE connectivity from the Server through a Converged Network Adapter (CNA) to the Nexus 5000.

The Nexus 5000 is then able to perform Ethernet switching for regular Ethernet and L3 frames, and Fibre Channel forwarding for FC frames...



Evolution de l'Architecture

Phase 2




FCoE & DCE

Once DCE-enabled modules become available on the Nexus 7000 or the MDS 9500 series platforms, multi-hop DCE topologies may be possible by keeping FCoE encapsulated within an Ethernet Header

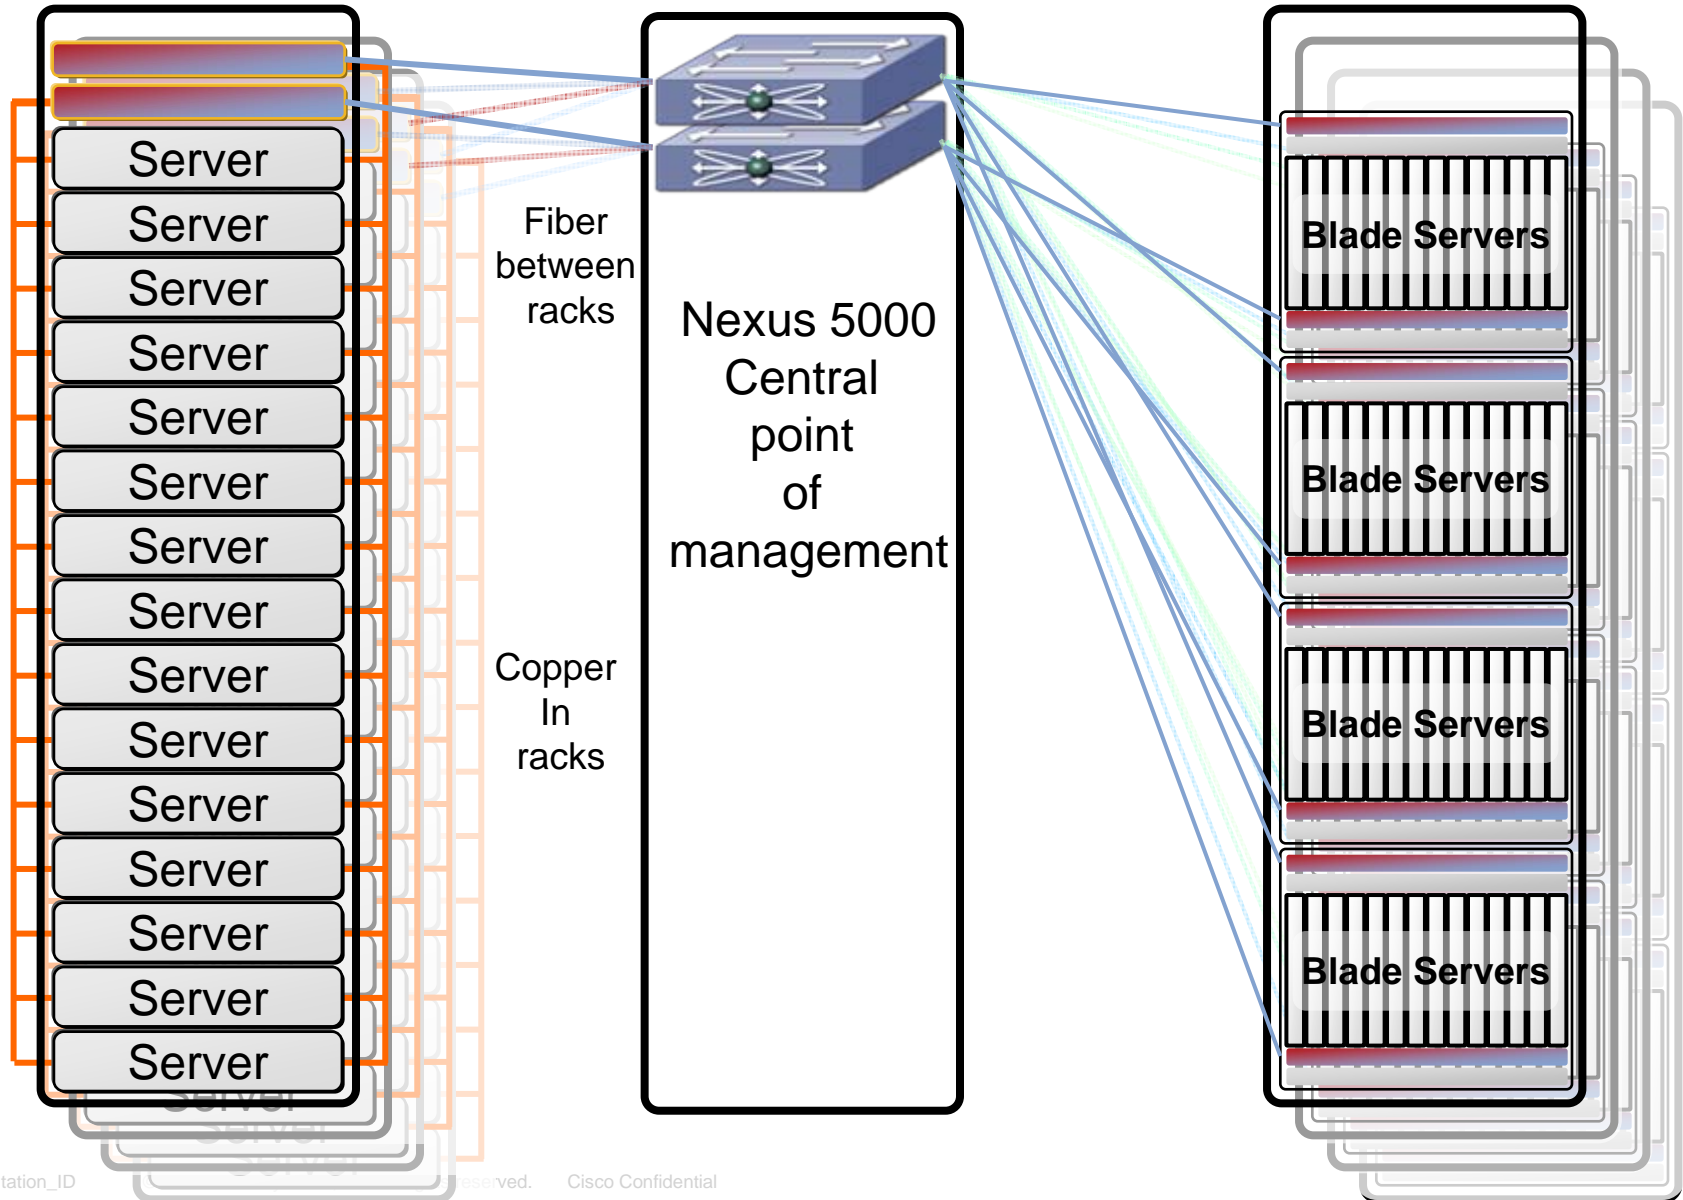
Additionally, with the introduction of direct FCoE attached targets, these may also be directly connected to any of these FCoE-enabled devices...

— Fibre Channel
— FCoE or DCE

Cablage Typique dans un Data Center



Cablage avec l'Unified IO



Famille de Switches Cisco Nexus 5000

Premier Switch dans l'Industrie pour la Consolidation et Unification des I/O pour le Datacenter

Switch Nexus 5000



56-Port L2 Switch (Nexus 5020)

- 40 Ports 10GE/FCoE/DCE, fixed
- 2 Expansion Modules



28-Port L2 Switch (Nexus 5010)

- 20 Ports 10GE/FCoE/DCE, fixed
- 1 Expansion Module

Modules d'Extension



FC + Ethernet

- 4 Ports 10GE/FCoE/DCE
- 4 Ports 1/2/4G FC



Ethernet

- 6 Ports
10GE/FCoE/DCE

Partenaires



2x10GE/ FCoE/DCE



SW FCoE, DCE

OS


Cisco NX-OS

Management

Cisco Fabric Manager et Cisco Data Center Network Manager

Nexus 2000 Fabric Extender (FEX)

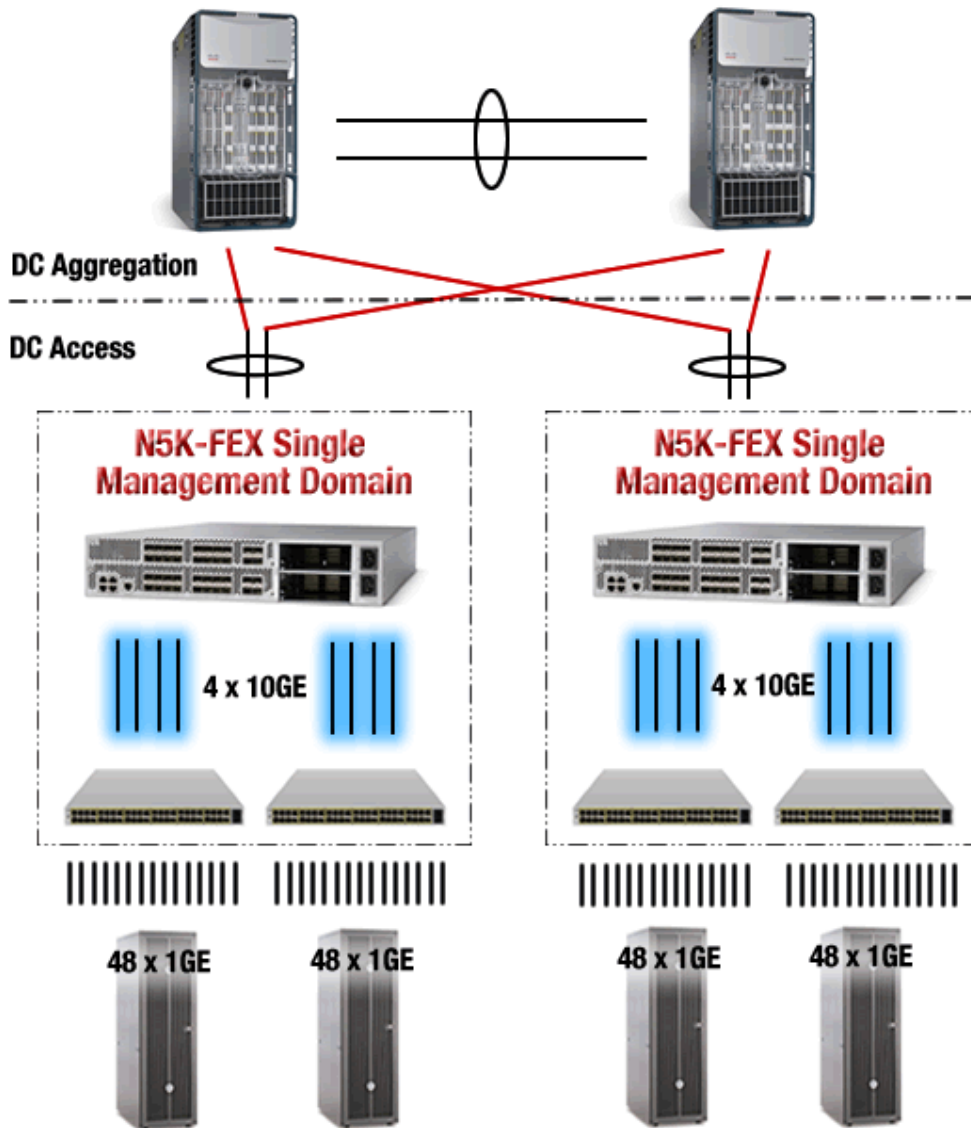
Pour une Meilleure Transition



Nexus 5000 Fabric Extender (FEX)

The Nexus 5000 Fabric Extender (FEX) acts as a remote line card (module) for the Nexus 5000, retaining all centralized management and configuration on the Nexus 5000, transforming it to a Virtualized Chassis

Nexus 2000 Fabric Extender (FEX)

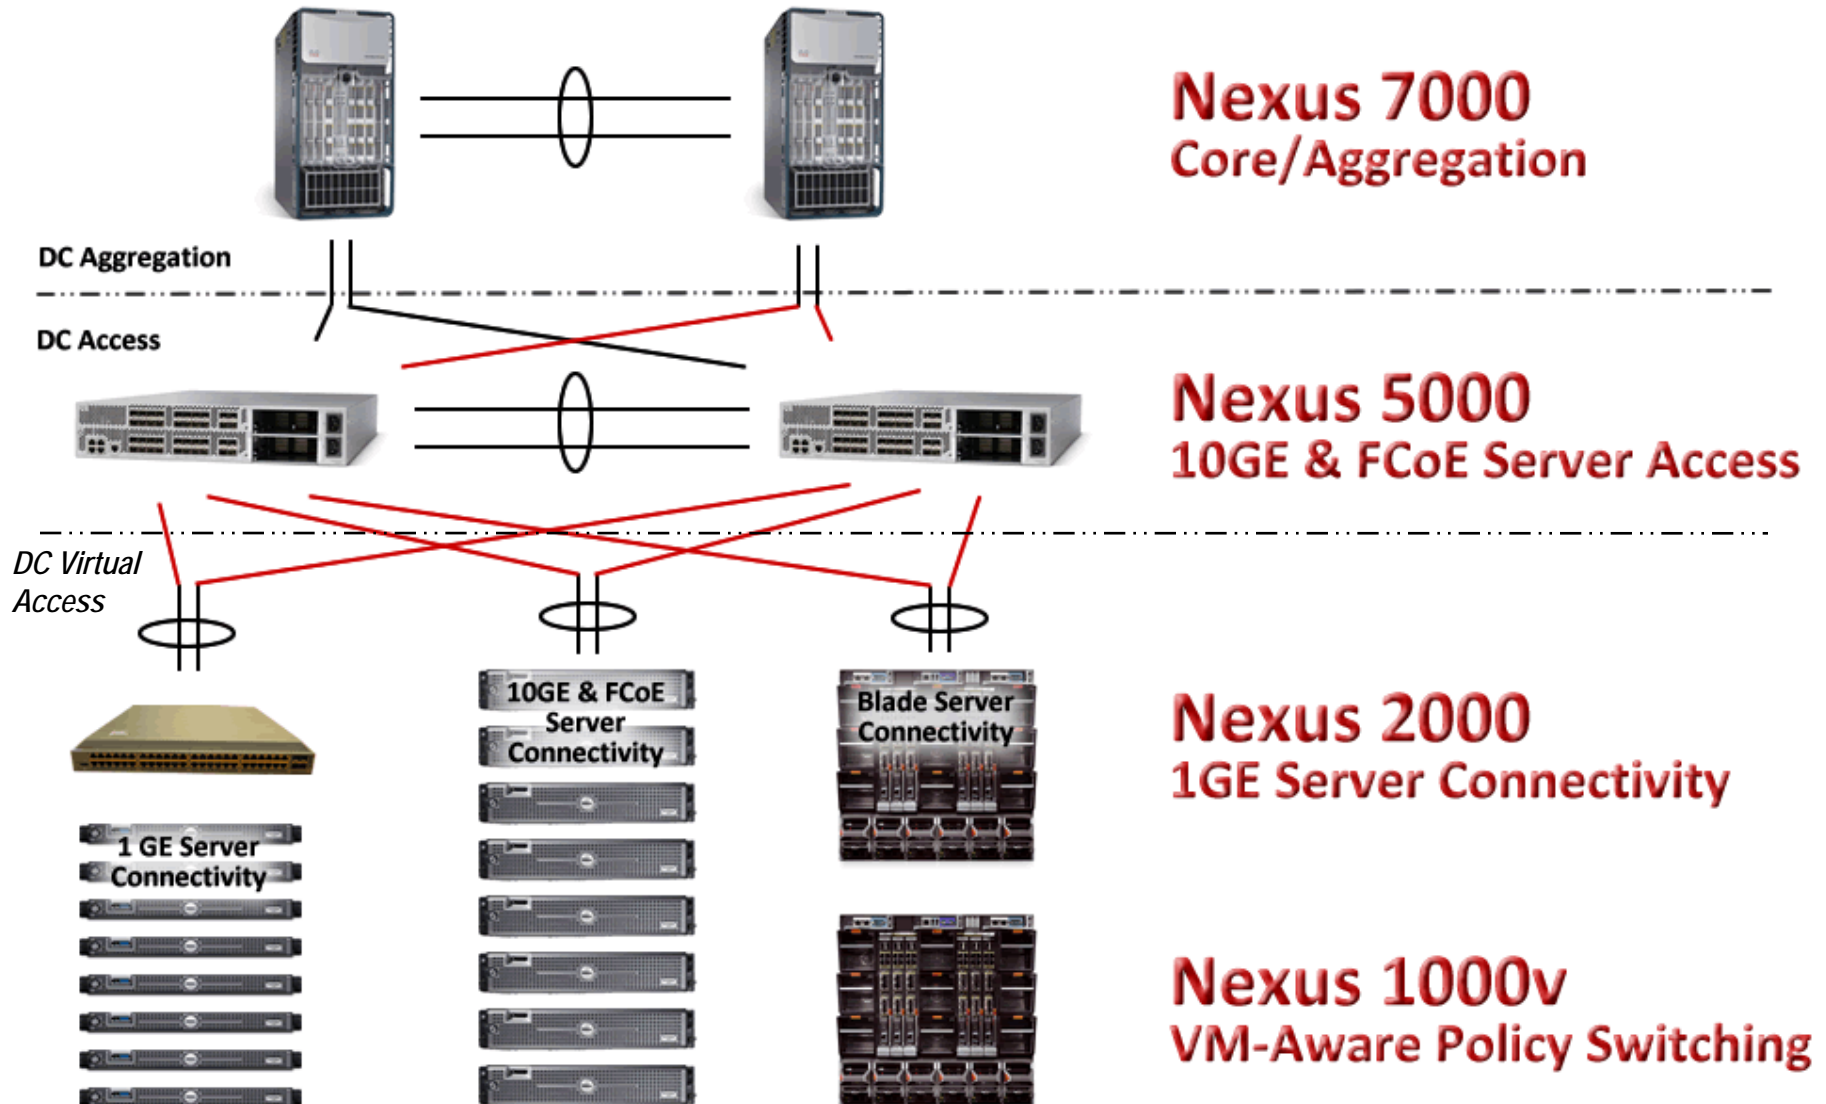


Nexus 5000 Fabric Extender

- ..Facilitates transition from 1G to 10G
- ..48 x 1G Downlinks, 4 x 10G Uplinks provide 1.2 : 1 oversubscription
- ..Centralized Management through Nexus 5000
- ..Virtualized Data Center Access
- ..Integrated Hardware-based VN-Link for Server Virtualization
- ..Physical Top-of-Rack, Virtual End-of-Row
- ..Virtual Line Card on Nexus 5000

Evolution de l'Architecture du Data Center

Famille Nexus – Optimisée pour le Data Center



Catalyst et Nexus : Positionnement Complémentaire



Cisco® Nexus 7000

**15 Terabit Scalability
Unified Fabric**

100GbE

40GbE

Transport Flexibility

Operational Continuity

10GbE

1GbE

Cisco Catalyst® 6500

**2 Terabit Scalability
Integrated Services**



15T


7.5T

3.7T

720G

2T

Nexus 7018 Chassis



Data Center Class Availability



Redundancy at Every Component Level

24 x 7 x 365

Fault Tolerant Hardware



No Single Point of Failure

- ✓ Dual supervisors offer active/standby Control Plane
- ✓ Fabric redundancy
- ✓ Dual control channels to ensure control packet delivery
- ✓ Power supply designed for grid redundancy



Even the Internal Out of Band Channel is redundant!

NX-OS: Purpose Built for the Data Center



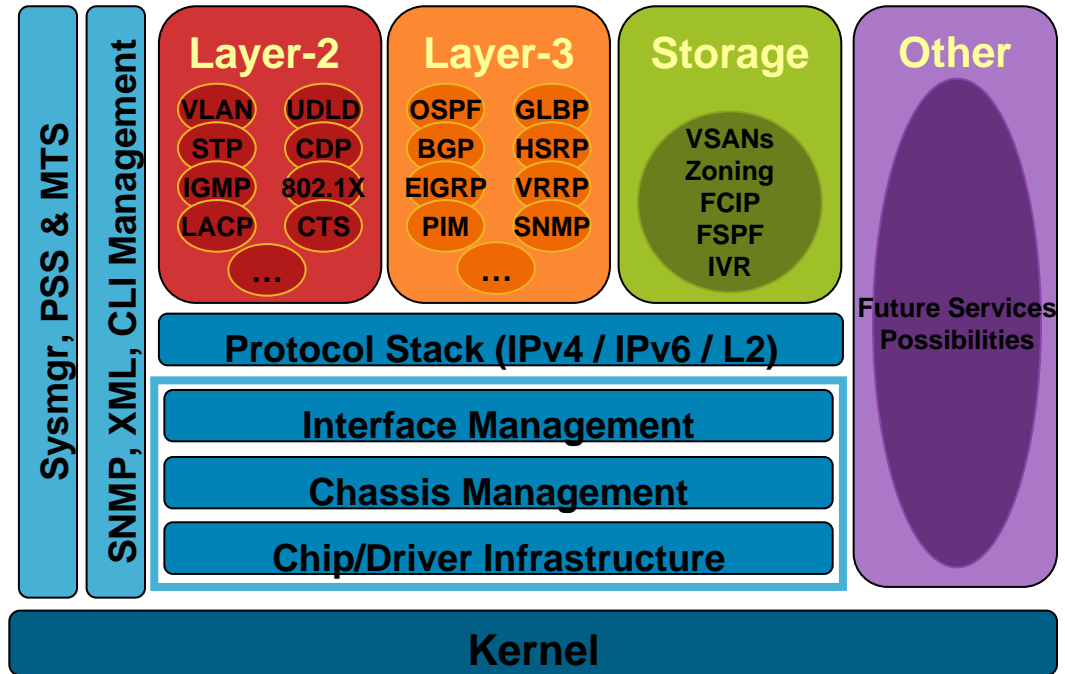
Data Center Class Availability

Redundancy at Every Component Level

NX-OS Modular Architecture

Zero Service Disruption Operating system

- ✓ Granular Modularity – every service exists in protected memory
- ✓ Process restart provides fault containment
- ✓ Conditional Services – process and CLI are not allocated until a service is enabled



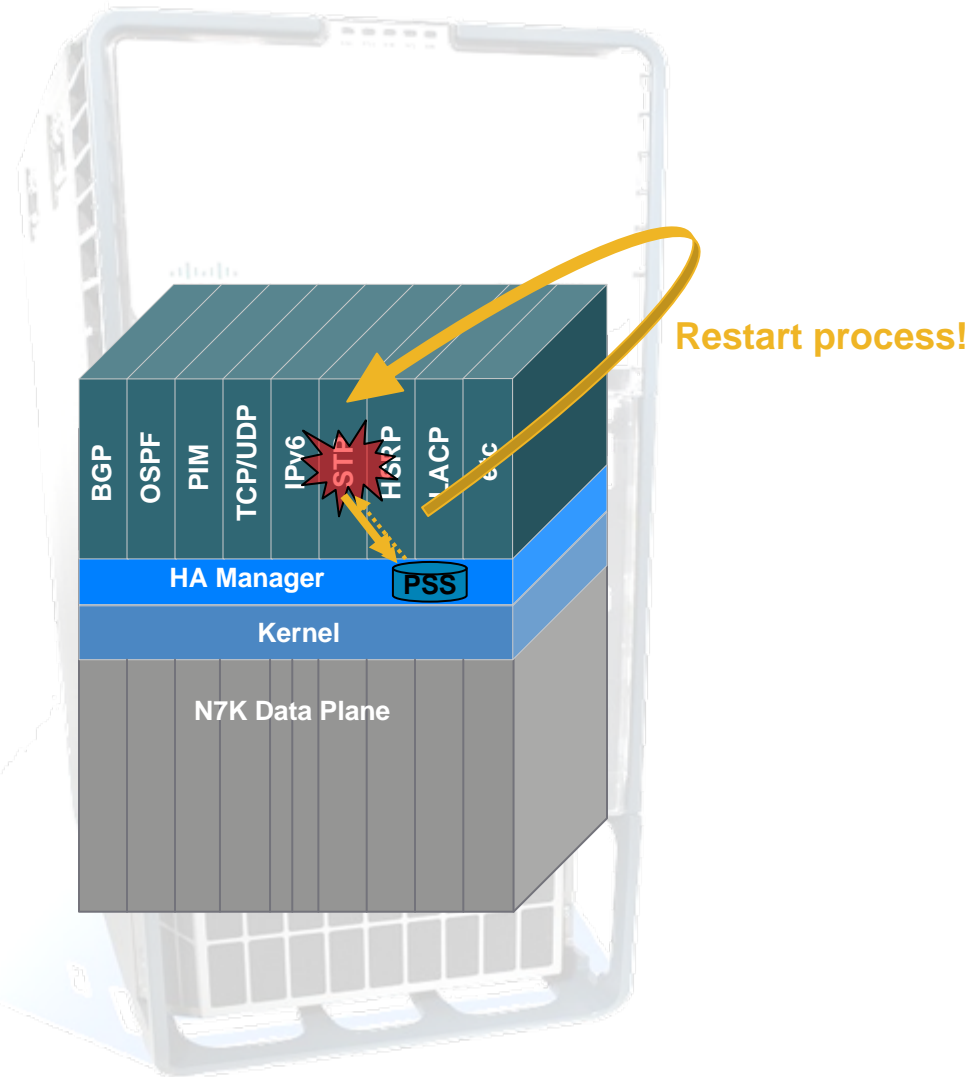
Data Center Class Availability

Stateful Process Restart

Stateful Process Restart

Avoid Network Reconvergences

- ✓ Processes can restart in milliseconds and maintain state from state database (PSS)
- ✓ Net effect is zero impact to neighbor relationships
- ✓ Supported for all L2 protocols as well as OSPFv2



Evolution of Ethernet Physical Media

Role of Transport in Enabling 10GE Technology

Mid 1980's

10Mb

UTP Cat 3

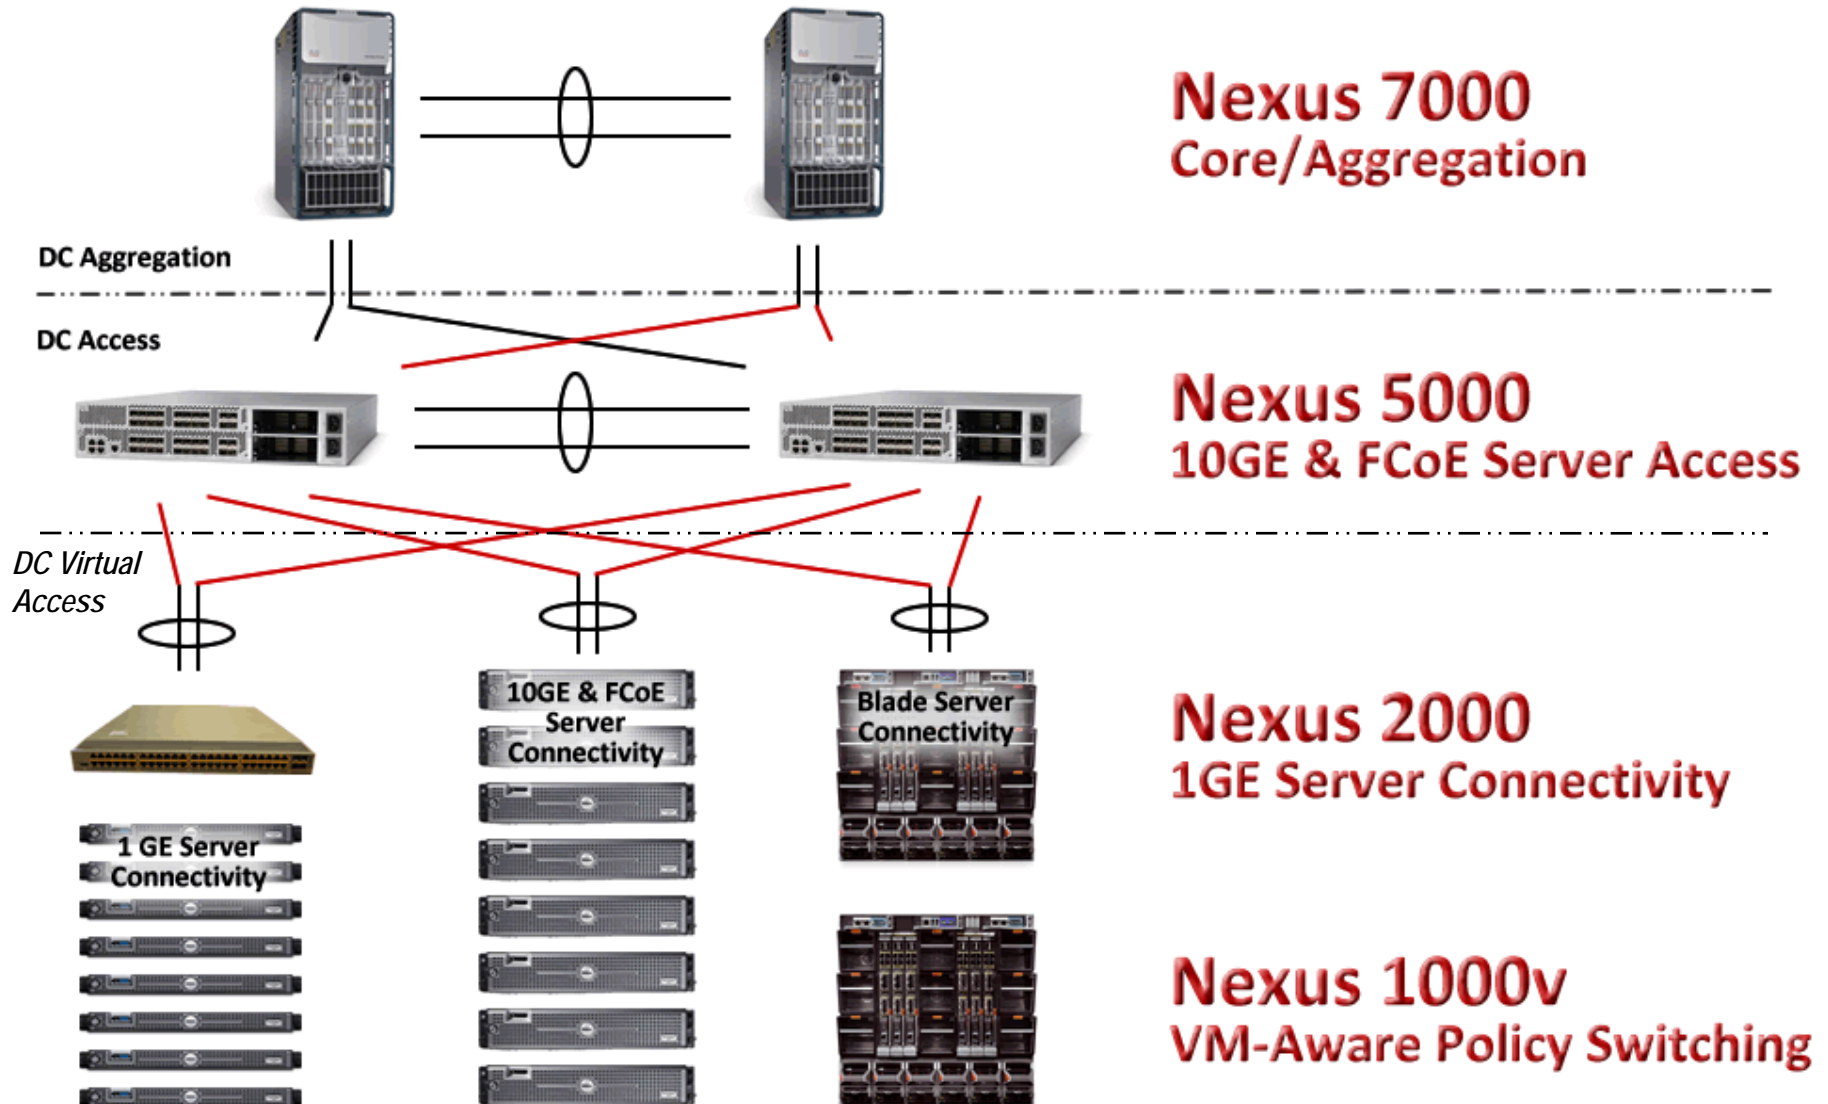


10GE Copper Solution

- Low cost
- Low power and latency
- Up to 10 meters (in-rack and adjacent rack cabling)

Technology	Cable	Distance	Power (each side)	Transceiver Latency
SFP+ CU Copper	Twinax	10m	~0.1W	~0.25 μ s
SFP+ USR ultra short reach	MM OM2 MM OM3	10m 100m	1W	~0.1 μ s
SFP+ SR short reach	MM OM1 MM OM3	33m 300m	1W	~0.1 μ s
10GBASE-T	Cat6 Cat6a/7 Cat6a/7	55m 100m 30m	~8W ~8W ~4W	2.5 μ s 2.5 μ s 1.5 μ s

Résumé et Questions/Réponses



MERCI



=



Contact : Hicham El Alaoui - helalaou@cisco.com