


IBM System Networking Pureflex and Networking Solutions

Kavish Shah
skavish@au1.ibm.com

GMU/Japan Technical Sales Manager

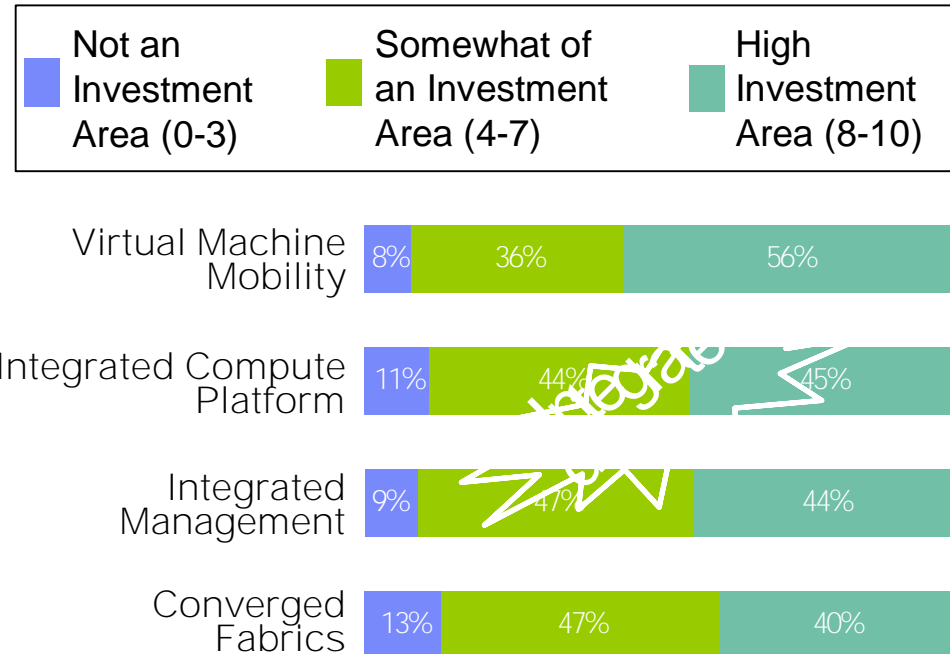



Presentation Includes

- § Customer Requirements
- § Market Trends
- § System Networking Overview
- § Product Overview and Details
- § Competitive Positioning
- § And more....

Customers want larger integrated systems

Priority Investment Trends in next 18 months



Clients are seeking solutions to the complexities associated with inefficient networking, server sprawl and manual virtualization management. Integrated system pre-packages server, storage, network, virtualization and management and provides an automated, converged & virtualized solution with fast time to value & simple maintenance

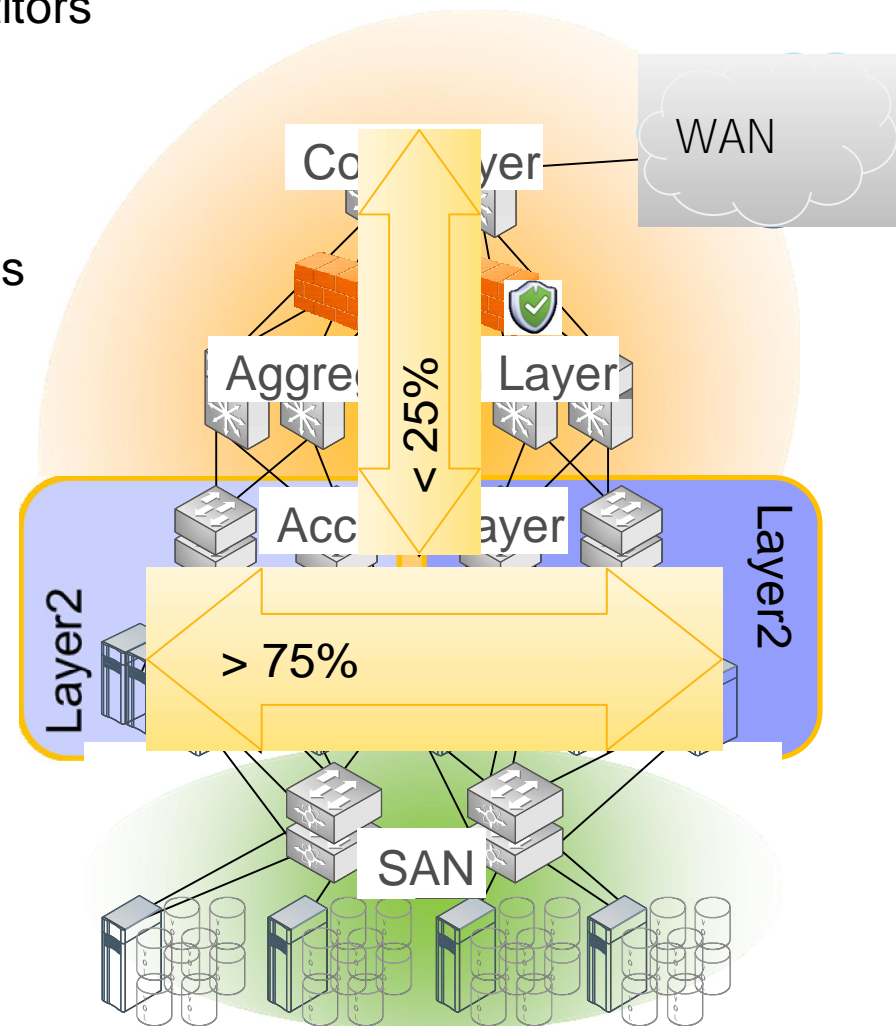

Intelligence is moving to the edge of the network

§ Intelligence moving from Core to the Edge

- Opposite of what some of our competitors want!


§ Why this is happening:

- Performance and latency demand
- Servers and networking optimization is critical
- Convergence of LAN and SAN
- Virtualization more east-west traffic




Product Overview:

Selling the Value and Advantages of 10Gb Scalable Switch and innovation virtualization capabilities



IBM System Networking

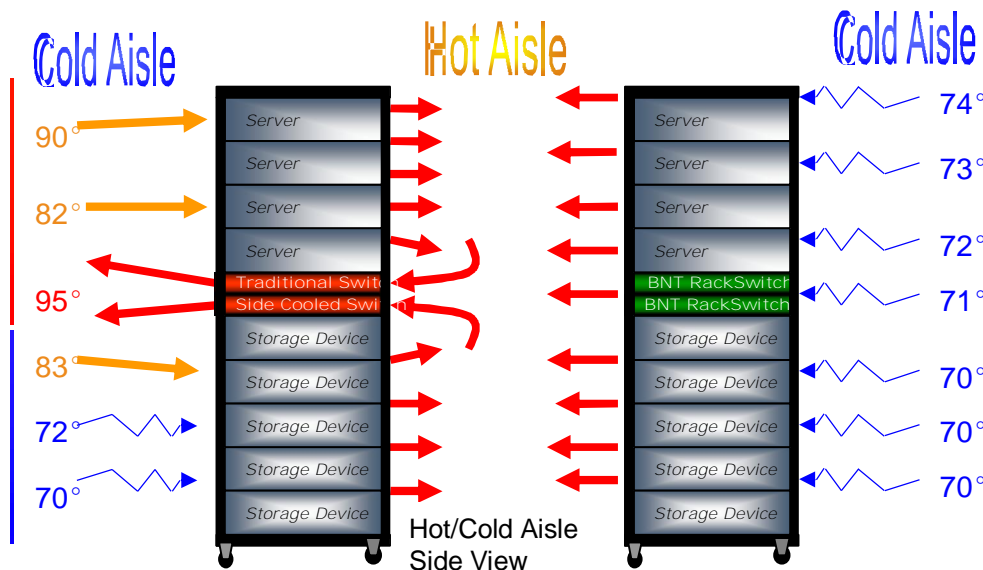

- Formed in October 2010 with IBM acquisition of B Network Technologies (BNT)
- Leadership in Smarter Networking for Smarter Data Centers
- Milestones
 - #1 in 40G Ethernet access switch (Dell'Oro)
 - #2 in data center Ethernet switches (Gartner)
 - Added over 2M data center ports in 2011
 - Over 12M ports deployed
 - Shipped 4th generation of network virtualization
 - Early to market with production-ready OpenFlow



Source: Competitive Landscape: Data Center Ethernet Switches, Worldwide, Gartner, October 2011
<http://www.gartner.com/reprints?id=1-17RXG9D&ct=111025&st=sb>

Best Practices for Data Center switches:

- § Line Rate, non-blocking Ethernet switches = less worry
- § Lower Power Consumption = lower total cost of ownership
- § Server Like Airflow = fewer hot spots, higher allowable temperatures
- § Redundant Power = less downtime
- § Redundant Fans = less downtime
- § Standards-based solutions = Industry-vetted, proven solutions



IBM System Networking RackSwitch 1/10Gb Portfolio



IBM RackSwitch G8000

R, F & DC models

- 44 ports 1G, RJ-45
- 4 ports 1G, SFP
- 4 ports 10G, CX4 or SFP+ uplinks optional
- Redundant fans and power supplies
- Stacking – Today
- 4.6 microsecond Latency
- 3 year warranty & 3 year SW upgrade license

IBM RackSwitch G8052

R & F models

- 48 ports 1G, RJ-45
- 4 ports 10G, SFP+ uplinks standard
- Hot-swap redundant fans & power supplies
- Stacking - Future
- 1.8 microseconds latency
- 3 year warranty & 3 year SW upgrade license


VMready: The industry's 1st automated Virtual Machine-aware networking



Energy Efficient, Choice of Front-to-Rear or Rear-to-Front Airflow



Microburst Performance BLADE vs. Cisco



IBM System Networking RackSwitch 10/40Gb Portfolio

All 10Gb



IBM RackSwitch
G8124E

R, F & DC models

- 24 ports 10G SFP+
- Low Latency – 520ns
- Redundant fans and power supplies
- 3yr warranty & SW upgrade license

10Gb/40Gb



IBM RackSwitch
G8264

R & F models

- 48 ports 10G SFP+
- 4 ports 40G QSFP+ (option: 16x10Gb ports)
- Low Latency - <1ms
- Hot-swap redundant fans & power supplies
- 3yr warranty & SW upgrade license



40Gb/10Gb



IBM RackSwitch
G8316

R & F models

- 16 ports 40G QSFP+
- Up to 64 SFP+ connections break-out cable 1x40G=4x10G
- Low Latency - 880ns
- Hot-swap redundant fans & power supplies
- 1yr warranty & SW upgrade license

Lossless Ethernet (CEE/DCB)



IBM Virtual Fabric




VMready: The industry's 1st automated Virtual Machine-aware networking



Energy Efficient, Choice of Front-to-Rear or Rear-to-Front Airflow



RackSwitch G8316



§ Ports

- 16 QSFP+ 40GbE ports
 - Up to 64 10GbE ports via breakout cables

§ Very similar to RackSwitch G8264

- Same switch ASIC
- Same software
- Same power supplies
- Same Fans

§ Faster MP than G8264

- Quad Core for increased scalability

§ Single ASIC design

- Predictable line rate performance
- Low power consumption

§ Full Layer 2/3

§ Available today



RackSwitch G8264-T

4x QSFP+ ports

48x 10GBase-T ports



§ Ports

- 48 10GBase-T ports (All 10G / 1G)
- 4 QSFP+ ports (All 40G / 10G)

§ Very similar to RackSwitch G8316

- Same switch ASIC
- Same MP & memory
- Same software

§ Needed for 10GBASE-T

- New PHYs
- New power supplies
- New fans


§ Stacking** (up to 8 switches)

§ Full Layer 2/3

§ Available today

****Post GA**

RackSwitch G8264-CS



§ Ports


- 12 Omni Ports (FC or Ethernet)
 - SFP+
 - Fiber Channel Mode: 2G / 4G / 8G
 - Ethernet Mode 1G / 10G
- 36 SFP+ Ports
 - Ethernet only: 1G / 10G
 - FCoE
- 4 QSFP+ ports
 - 40GbE mode or 4x10GbE Mode

§ Fiber Channel Forwarder:

- FCoE FF-BB-5 Compliant
- FCF Gateway
- Full Fabric Services
- NPIV Transparent gateway mode

§ Full Layer 2/3


Flex System / PureFlex Usage Scenario #2: FCoE End-to-End



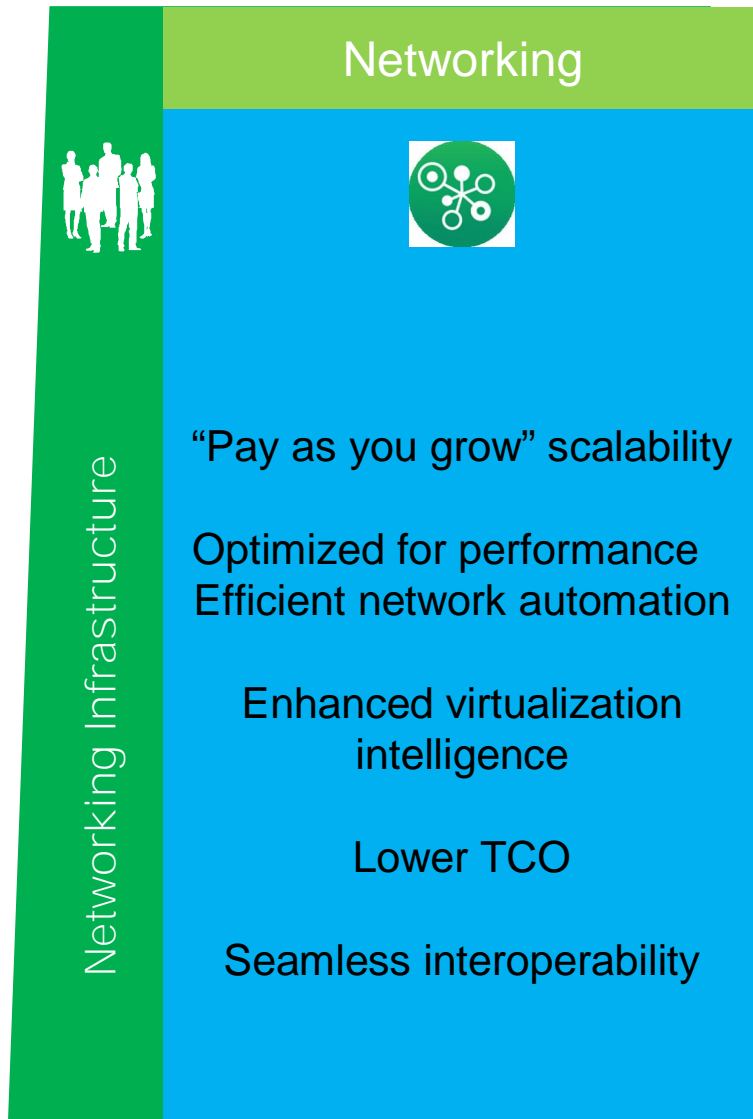
Adapter	Transit	FCoE Switch	SAN Switch	Storage Target	OS levels
LOM & CN4054 4-port adapter (BE3) in pNIC, vNIC2 or UFP Mode	EN4093 / EN4093R	G8264CS	N/A	FC: Storwize V3700 or V7000, SVC	Win2008, WS2012, ESX 4/5, RHEL 5/6, SLES 10/11

Target
Availability
June 2013

New System Networking embedded switches for Flex System



IBM Flex System Fabric EN4093 10Gb Scalable Switch



Scalable 10Gb Ethernet with 10/40Gb Uplinks

Overview


- Layer 2/3 Data Center design – upgrade in order
- Base Switch: 14 x 10Gb server port & 10 x 10Gb uplinks
- Switch upgrade 1: 28 x 10Gb server ports & 10 x 10Gb plus 2 x 40Gb uplinks.
- Switch upgrade 2: 42 x 10Gb server ports* & 14 x 10Gb plus 2 x 40Gb uplinks.*

Leadership

- Proven Operating System
- Exceptional Price/Performance
- Investment Protections – scalable pay-as-you-grow design
- FCoE convergence support planned 4Q 2012
- Partitioning – Physical/Logical 4Q 2012
- VM aware & VM Mobility with VMready
- Virtual Fabric – carve up virtual NIC's and pipes (up to 32)
- Seamless interoperability with other vendors switches
- Warranty: 1 year or warranty of the chassis

Recommended Top-of-Rack switch

- Multiple chassis of 10Gb connection G8264 – PureFlex
- Multiple chassis of 40Gb connection G8316 – Build Own



IBM Flex System EN2092 1Gb Scalable Switch

Networking



Networking Infrastructure

“Pay as you grow” scalability

Optimized for performance
Efficient network automation

Enhanced virtualization
intelligence

Lower TCO

Seamless interoperability

Scalable 1Gb Ethernet with 1/10Gb Uplinks

Overview

- Layer 2/3 Data Center design – upgrade flexibility
- Base Config: 14 x 1Gb server port & 10 x 1Gb uplinks
 - Part Number: 49Y4294
- Upgrade 1: 2x ports - 28x1Gb server ports & 20x1Gb uplinks
 - Part Number: 90Y3562
- Upgrade 2: 10Gb Uplinks Enables the 4 x 10Gb uplinks
 - Part Number: 49Y4298

Leadership

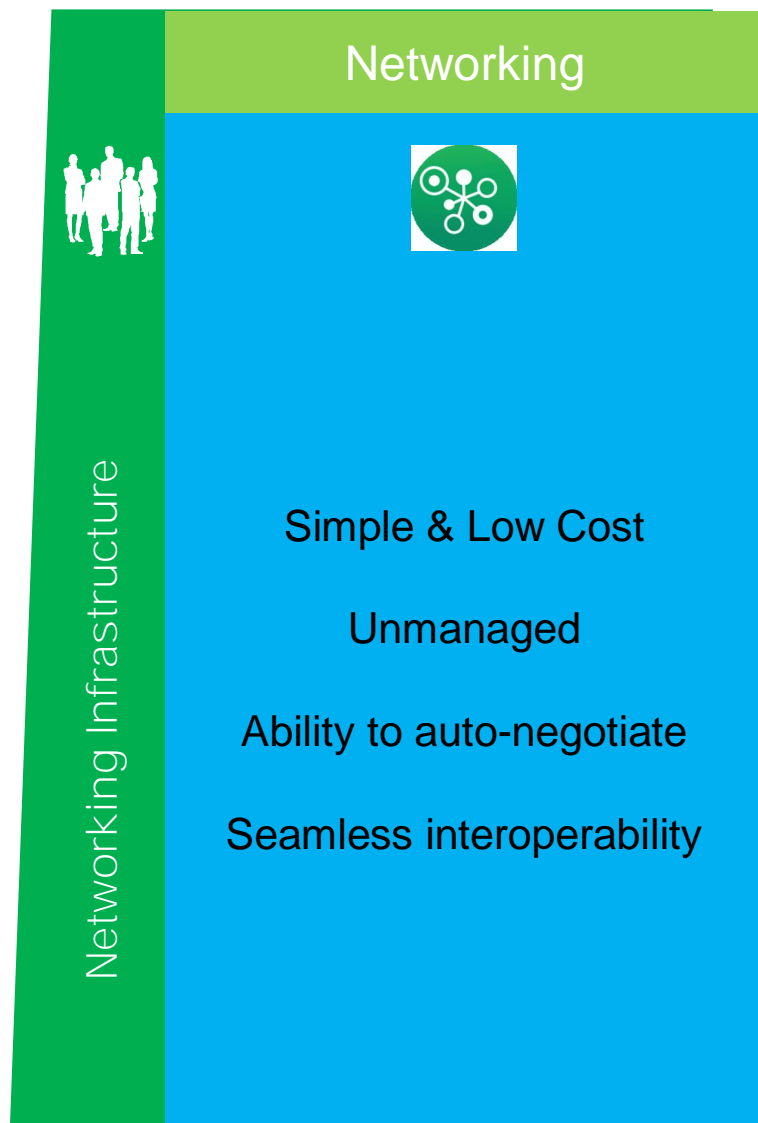
- Proven Operating System
- Exceptional Price/Performance
- Investment Protections – scalable pay-as-you-grow design
- VM aware & VM Mobility with VMready
- Seamless interoperability with other vendors switches
- Warranty: 1 year or warranty of the chassis

Recommended Top-of-Rack switch

- Multiple chassis of 1Gb connection G8052 – Build Own
- Multiple chassis of 10Gb connection G8264 – Build Own



IBM Flex System EN4091 10Gb Ethernet Pass-thru Module



Simple 1/10Gb Pass-through for seamless connectivity to upstream networks

Overview


- 14 x 1/10Gb server port & 14 x 1/10Gb uplink ports

Leadership

- Exceptional Price/Performance
- Seamless interoperability with other vendors switches



Converged Fabric Switch : March 2013 CN4093



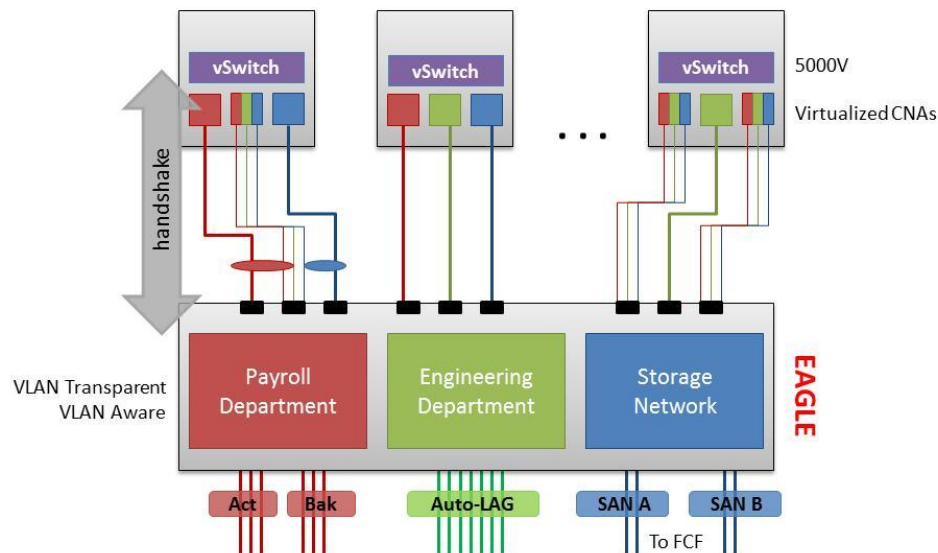
- Downlinks same as 10Gb Switch
- Uplinks
 - 12 Omni Ports : 10Gb ETH or 8Gb FC
 - 2x10Gb ETH
 - 2x40Gb ETH
- 1 to 8 switches stacked as 1
- 1.28 Tbps per switch
 - 5 Tbps per chassis

- CEE, Direct attach FCoE, FCoE multi-hop, FC Forwarder
- FC Features: NPIV, FC Services, N Port, F port, E port, Native FC connectivity to FC SANs (using NPIV mode)

§ Automated:

- vNIC
- Qbg

Virtual Pass-Thru



Simplified management:

- Profile based setup of NIC bandwidth, VLAN, FCoE across multiple Eagle across multiple chassis.
- Simple way to set some default L2 configurations via FSM

Total 64 port module:

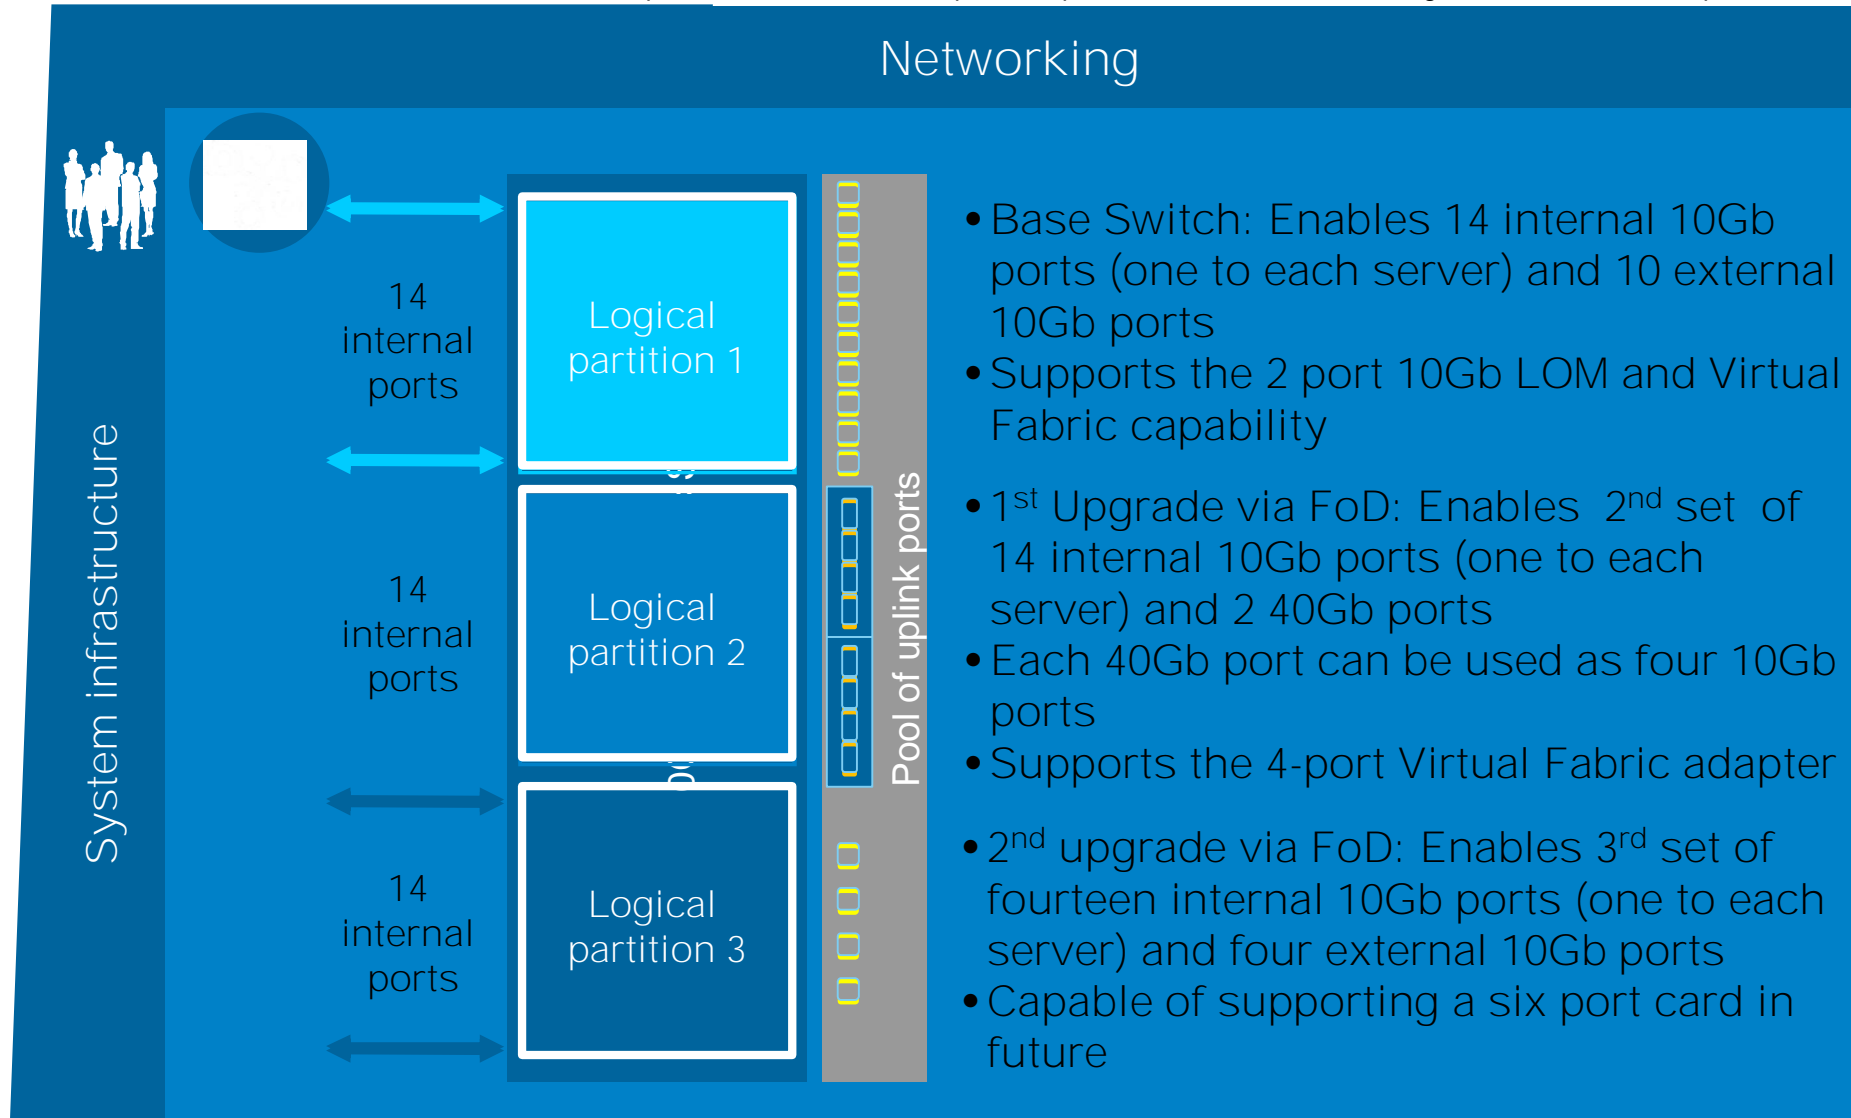
- Base module: 14 internal 10Gb ports & 10 external 10Gb ports
- FOD Upgrade: 14 internal 10Gb ports & 2 external 40Gb ports, or 8 10Gb with breakout cable.
- FOD Upgrade: 14 internal 10Gb ports and 4 external 10Gb ports
- All SFP+ ports can auto negotiate to 1Gb

Function overview:

- Aggregation, basic forwarding, virtualization and QoS
- CEE for connectivity to server CNAs & FCFs upstream
- Simple “NIC” or I/O like management for the device
 - Define your departments or tenants
 - Turn on/off capabilities
- Switch Partitions (SPARs) for traffic and user isolation

Next generation flexibility: Scale for Bandwidth, Ports, or Both

IBM 10Gb Switch: Wired for up to three 10Gb ports per node and twenty two external ports




Smarter Networking for Smarter Data Centers

For Cloud, Big Data (Analytics) & Workload Optimized Systems

Delivering smarter, faster, virtual and flexible system networks that:

- Move to Software Defined Networks
- Bring speed & intelligence to the network edge
- Are standards-based and interoperable
- Flatten and converge
- Reduce cost, power, space and complexity
- Scale easily to meet rising demand
- Simplify management
- Optimize and automate network virtualization
- Provide high availability and enhanced security
- Address today's data center pain points!



Intelligent, Integrated and Flexible network architecture that can fit with your existing or future environment



Networking

- Extreme “Pay as you grow” scalability
- Performance: 1st
 - Support 40Gb
 - 3x10Gb ports
 - < 1ms latency
- Seamless interoperability
- Designed for future:
 - Flat network
 - Convergence
- Physical / Logical Switching*
- Virtual NIC (Virtual Fabric)

Optimized

- Optimize Virtual Machine Mgmt. Virtual Switching (5000V)
- Automate network & server provisioning and failover (Fabric Mgr.)
- VM aware networking for easy management and tracking (VMready)
- Using Standards (Qbg)

Automated

- Manage everything as 1 logical unit. (FSM)
- Stacking & vLAG
- Move from physical network management to workload management (SNEM integration w/ FSM/Director)
- Smart workload provisioning and placement to maximize performance (UFM)

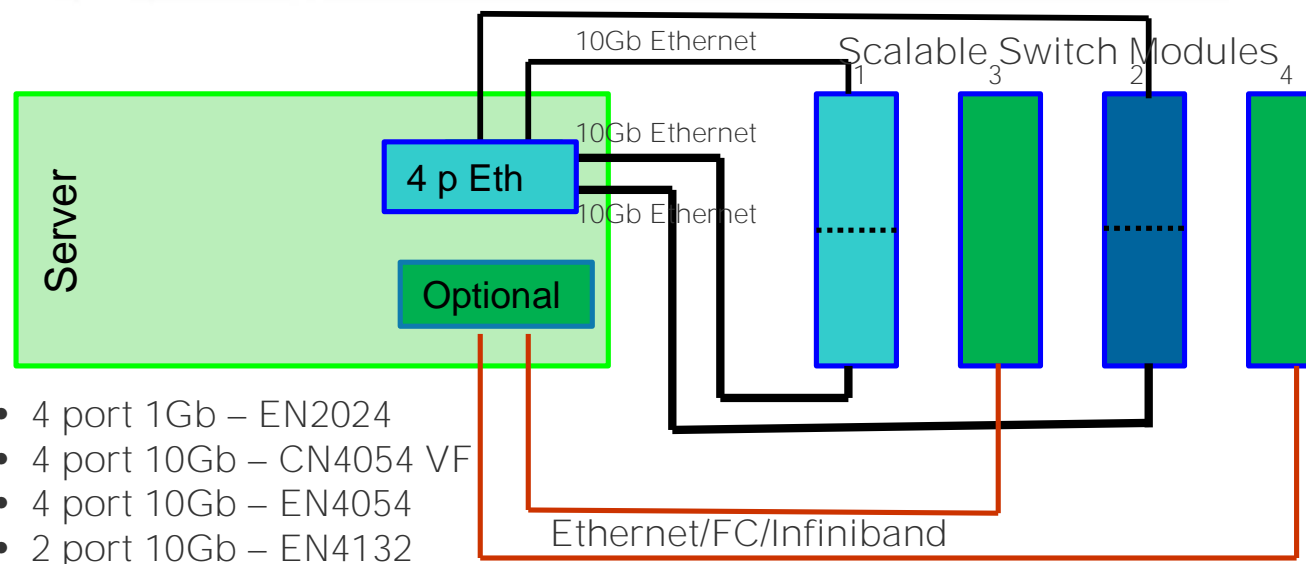

Integrated



Flexible, Integrated, Standard-based No Compromise Networking

Flexible, Integrated, Standard-based No Compromise Networking

IO architecture that reduces complexity and offers future



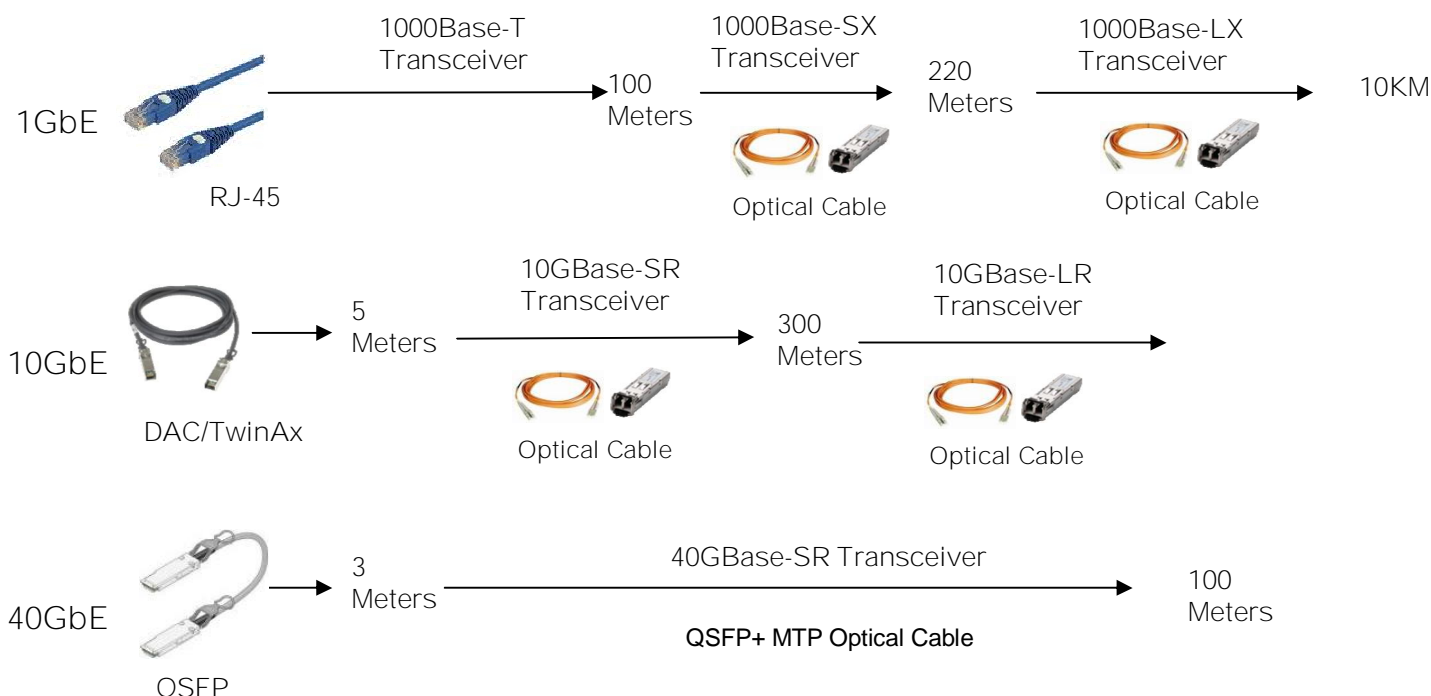
- 4 port 1Gb – EN2024
- 4 port 10Gb – CN4054 VF
- 4 port 10Gb – EN4054
- 2 port 10Gb – EN4132



Scalable Switch
Elements (4X)

Connectivity Options and distances for the EN4093

Networking Cable and Transceiver Questions	
Will your 1Gb connections exceed 100m (less than 550m)?	= 1Gb Short Range Transceivers + Optical Cable
Will your 10Gb connections exceed 5m (less than 300m)?	= 10 Gb Short Range Transceivers + Optical Cable
Will your 40Gb connections exceed 3m?	= 40Gb Short Range QSFP+ Trans+ Optical Cable
Will your 1 or 10Gb connections exceed 300m (less than 10km)?	= Long Range Transceiver + Optical Cable



- DAC cables are ideal within a rack or adjacent racks (lower cost & power)
- Transceivers are best for longer distances

Transceivers and Cables Options

10Gb SFP+ Options

§Direct Attach Copper Cables – Low Cost, Low Power ideal for shorter distances

- IBM PN: 90Y9427 1m IBM Passive DAC SFP+ Cable
- IBM PN: 90Y9430 3m IBM Passive DAC SFP+ Cable
- IBM PN: 90Y9433 5m IBM Passive DAC SFP+ Cable

§Transceivers supported in SFP+ ports

- IBM PN: 81Y1618 IBM (1000BASE-T) SFP RJ45 Transceiver
- IBM PN: 81Y1622 IBM (1000BASE) SFP SX Transceiver
- IBM PN: 90Y9424 IBM (1000BASE) SFP LX Transceiver
- IBM PN: 46C3447 IBM (10GBASE) SFP+ SR Transceiver
- IBM PN: 90Y9412 IBM (10GBASE) SFP+ LR Transceiver

SFP 1000Base-T Copper Transceiver



SFP+ 10GBase-SR SR Transceiver



SFP 1000Base-SX Fiber Transceiver



Direct Attached Copper SFP+ Cable

40Gb QSFP+ Options

§QSFP+ Direct Attach Cables – break out into four 10GbE connections

- IBM PN: 49Y7886 1m IBM QSFP+ DAC Break Out Cable
- IBM PN: 49Y7887 3m IBM QSFP+ DAC Break Out Cable
- IBM PN: 49Y7889 5m IBM QSFP+ DAC Break Out Cable

§QSFP+ Direct Attach Cables – for 40GbE connections

- IBM PN: 49Y7890 1m IBM QSFP+-to-QSFP+ Cable
- IBM PN: 49Y7886 3m IBM QSFP+-to-QSFP+ Cable


§Transceivers supported in SFP+ ports

- IBM PN: 49Y7884 IBM (40GBASE) QSFP+ SR Transceiver



1GbE Mgmt

Next generation IO that Provides IBM Advantages



IBM Flex System	UCS	HP
Extreme Flexibility Designed to meet port and bandwidth requirements for next decade. Pay for what you need today.	None	None
Highest Performance First 40Gb capable Blade Ethernet Switch Up to 220Gb uplink BW and <1microsec latency	None	None
Integrated Management Physical & Logical network management (future) Intelligent network Fabric monitoring Smart network provisioning based on workload (future)	Yes	Partial
Enhanced Virtualization Automate network provisioning and switch via IBM 5000v and IBM Fabric Manager Improve network utilization via Virtual Fabric- Up to 32 vNICs per server	Yes	Partial
Standard Based Seamless Integration with existing Networking Adoptions of standards like Qbg	Partial	Partial

IBM Flex System Fabric E4093 10Gb Scalable Switch

Key Features

Customer Benefits

Cooler

Lower utility costs and more technology per racks power envelop.

- Low power consumption

Easy & Flexible

Reduce IT staff cost to deploy and manage, higher availability/reliability.

- Reduces cable clutter (no cables from server to switch).
- Scalable Ethernet ports up to 42 server ports (6 per server*), plus up to 18 10Gb uplinks and 2x40Gb uplinks
- High availability via enterprise –class features like failover
- Seamless integration with 3rd party switches.

Performance

Better end users/customer experience.

- Low Latency under 1 microsecond
- Industries first 40Gb Ethernet embedded switch

Virtualization

- VMready – helps make the network VM aware for simpler management & VM mobility.
- Virtual Fabric – allows for vNIC's and virtual pipes between adapter and switch

Cost of Ownership

Better and shorter return of investment, savings could be used elsewhere.

- Acquisition cost is competitive & improves as you scale up 10Gb connections



§ Base Configuration –

- 14 x 10Gb server port, 10 x 10Gb uplinks
- Part Number: 49Y4270

§ Switch upgrade 1: 2x the ports

- 28 x 10Gb server ports, 10 x 10Gb and 2 x 40Gb uplinks
- Part Number: 49Y4798

§ Switch upgrade 2*:

- 42 x 10Gb server ports, 14 x 10Gb and 2 x 40Gb uplinks
- Part Number: 88Y6037

Competes Against:


- HP Virtual Connect Flex10
- HP ProCurve 6120XG
- Dell PowerConnect M8024
- Cisco UCS Fabric Extender and Fabric Interconnect

Where to win

- Target Customers: Clients only wanting 10G Ethernet and consider HP or Dell PowerConnect switches
- Clients looking for a low TCO but also wanting features that help improve performance & available, lower power & complexity
- Clients looking for 10Gb Ethernet or Clients deploying virtualization and looking to reduce cost and complexity
- Clients considering 40Gb Ethernet aggregation in the future

Delivering High Availability

HA

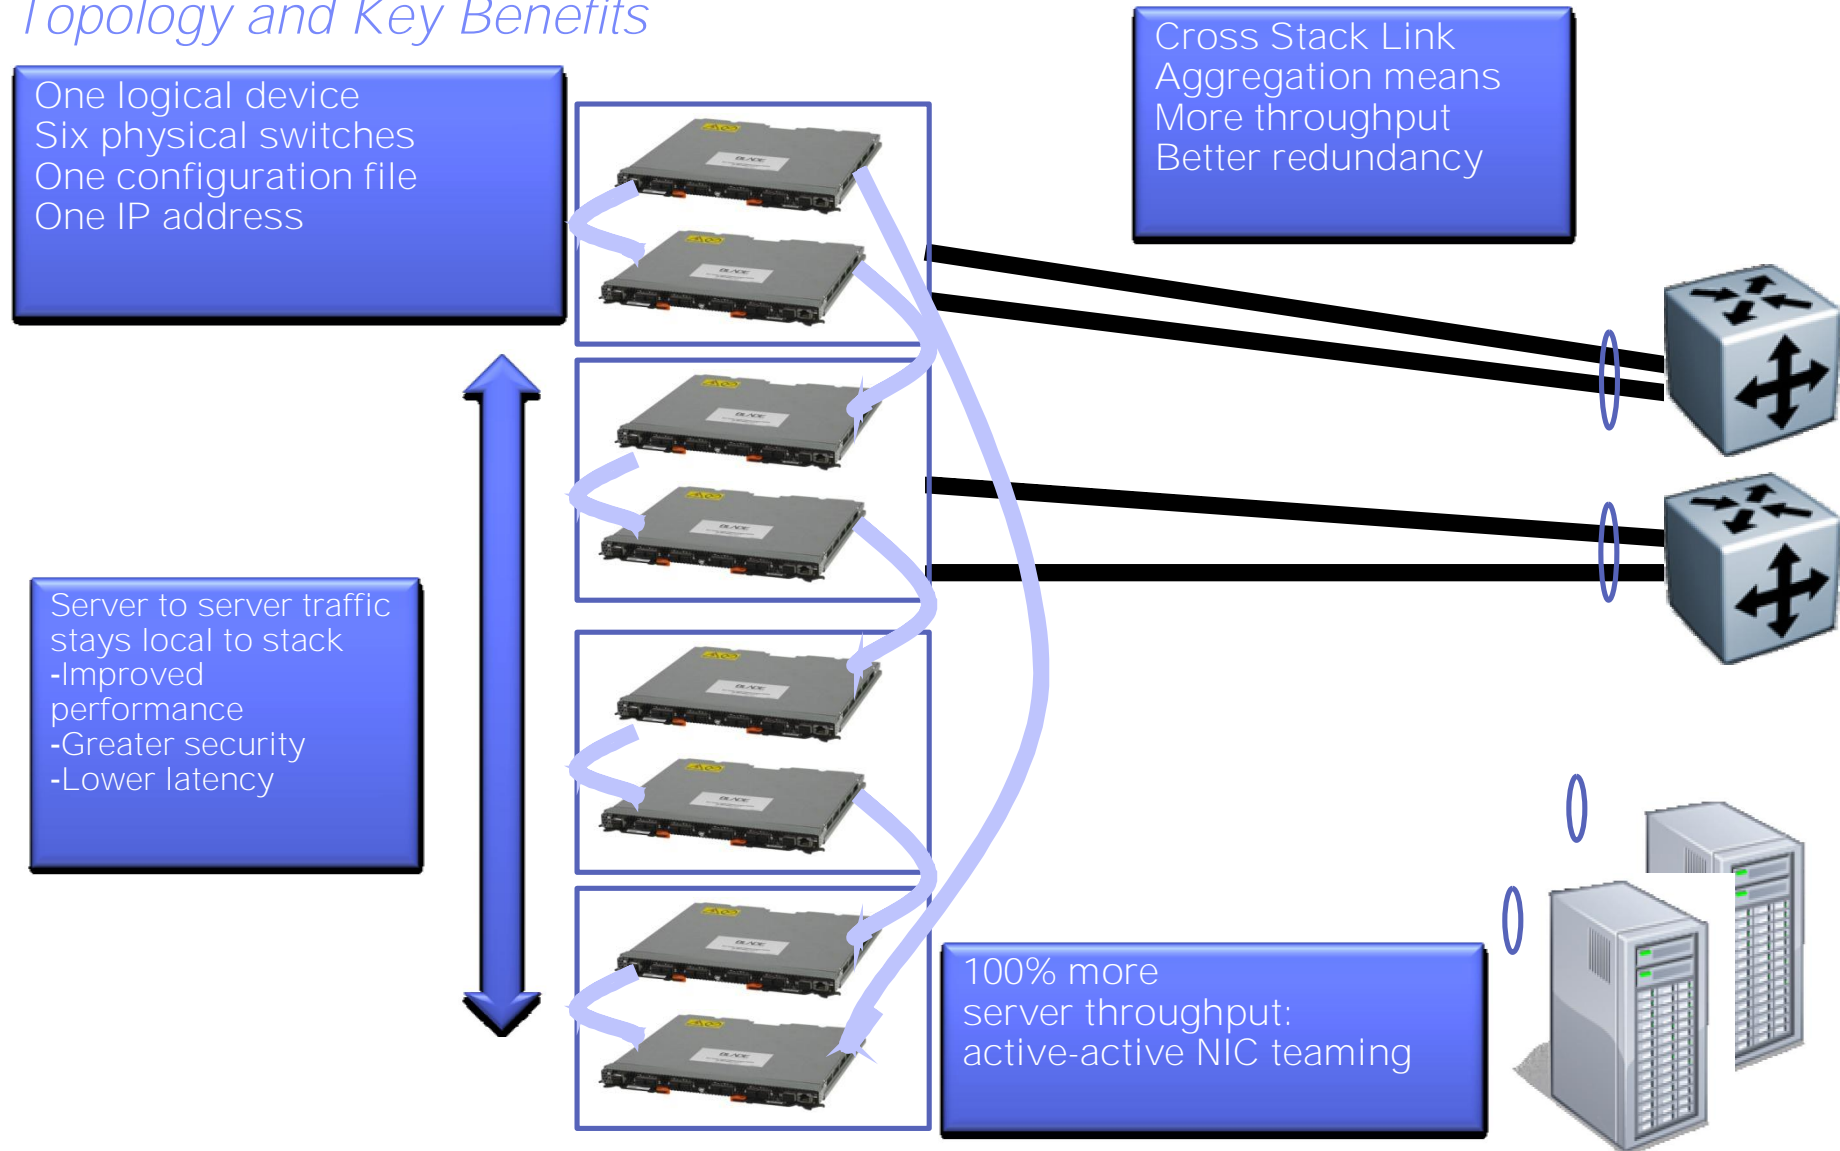


What can we do?

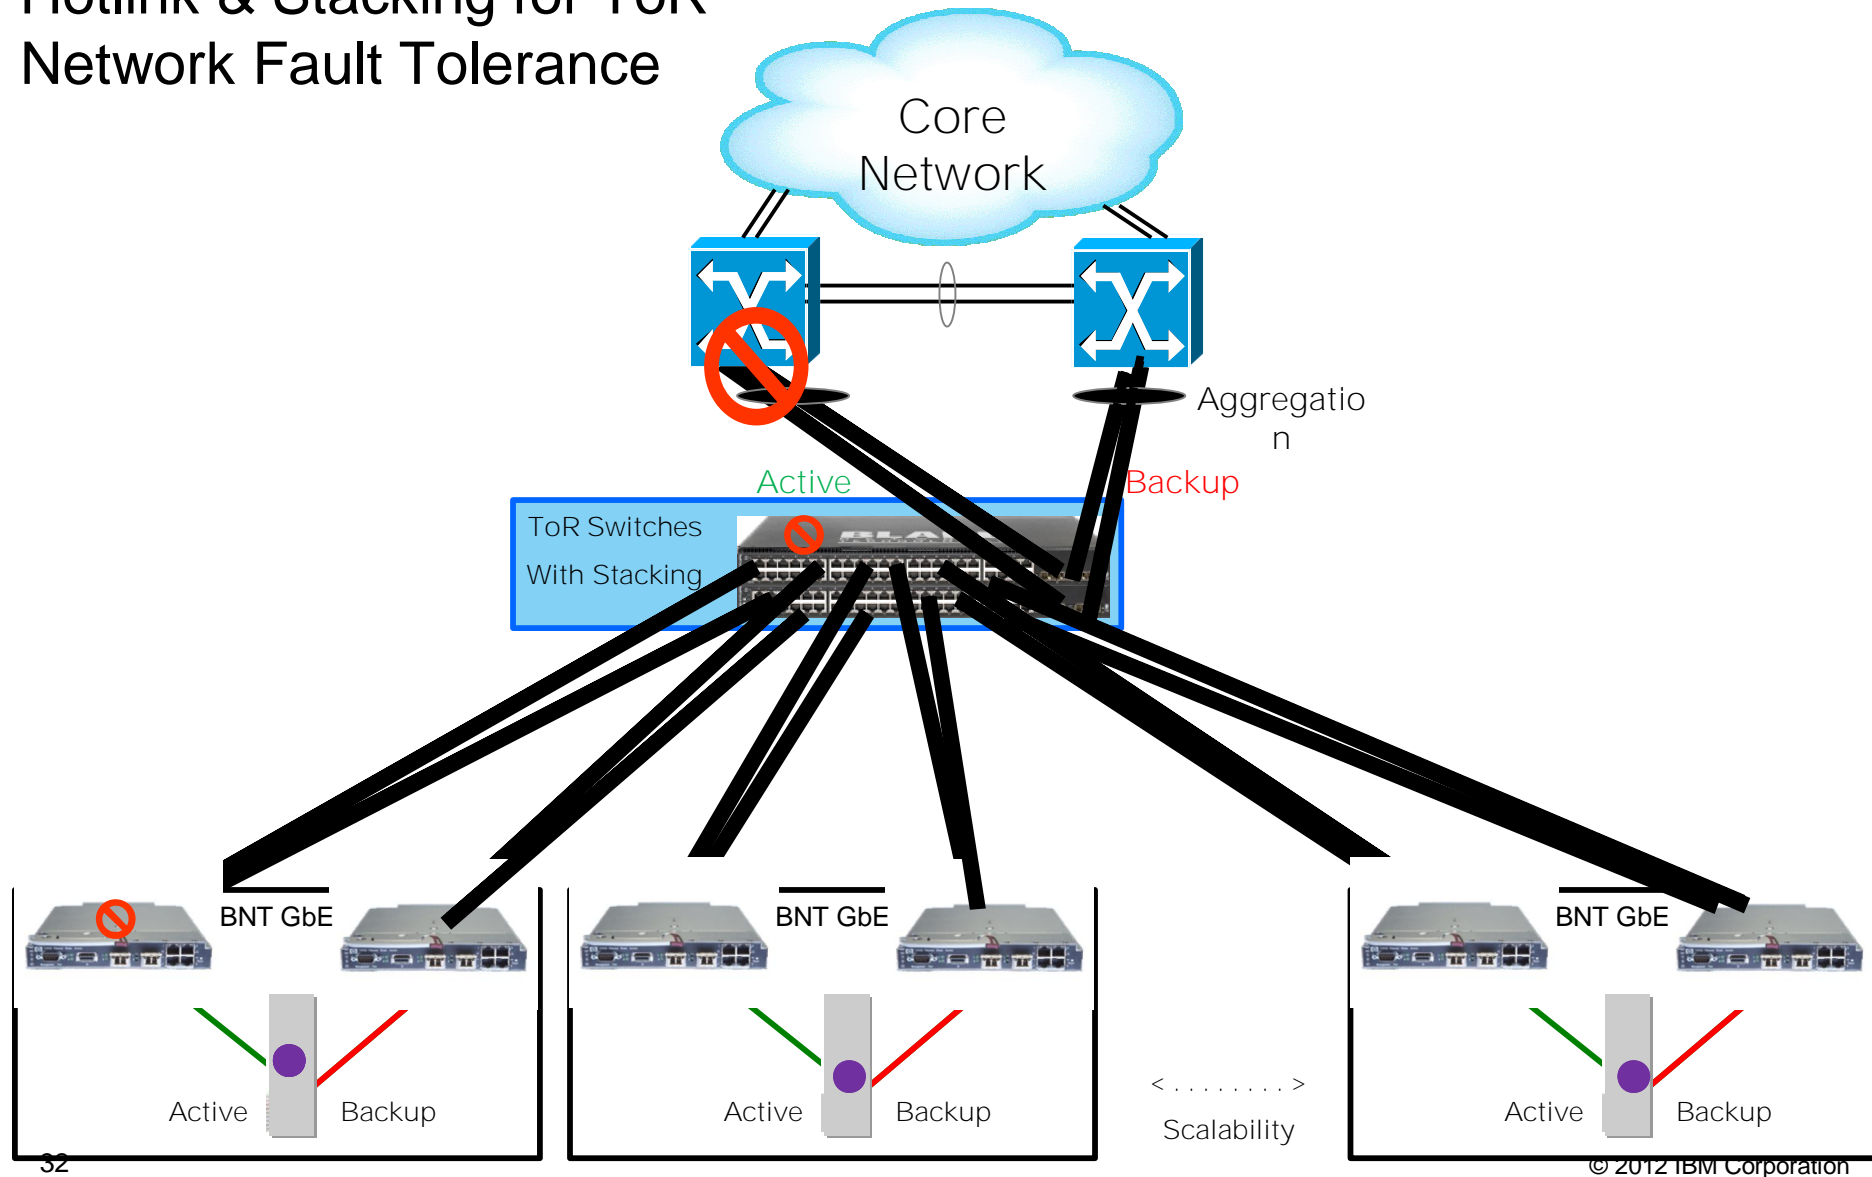


Stacking


Topology and Key Benefits



Hotlink & Stacking for ToR Network Fault Tolerance




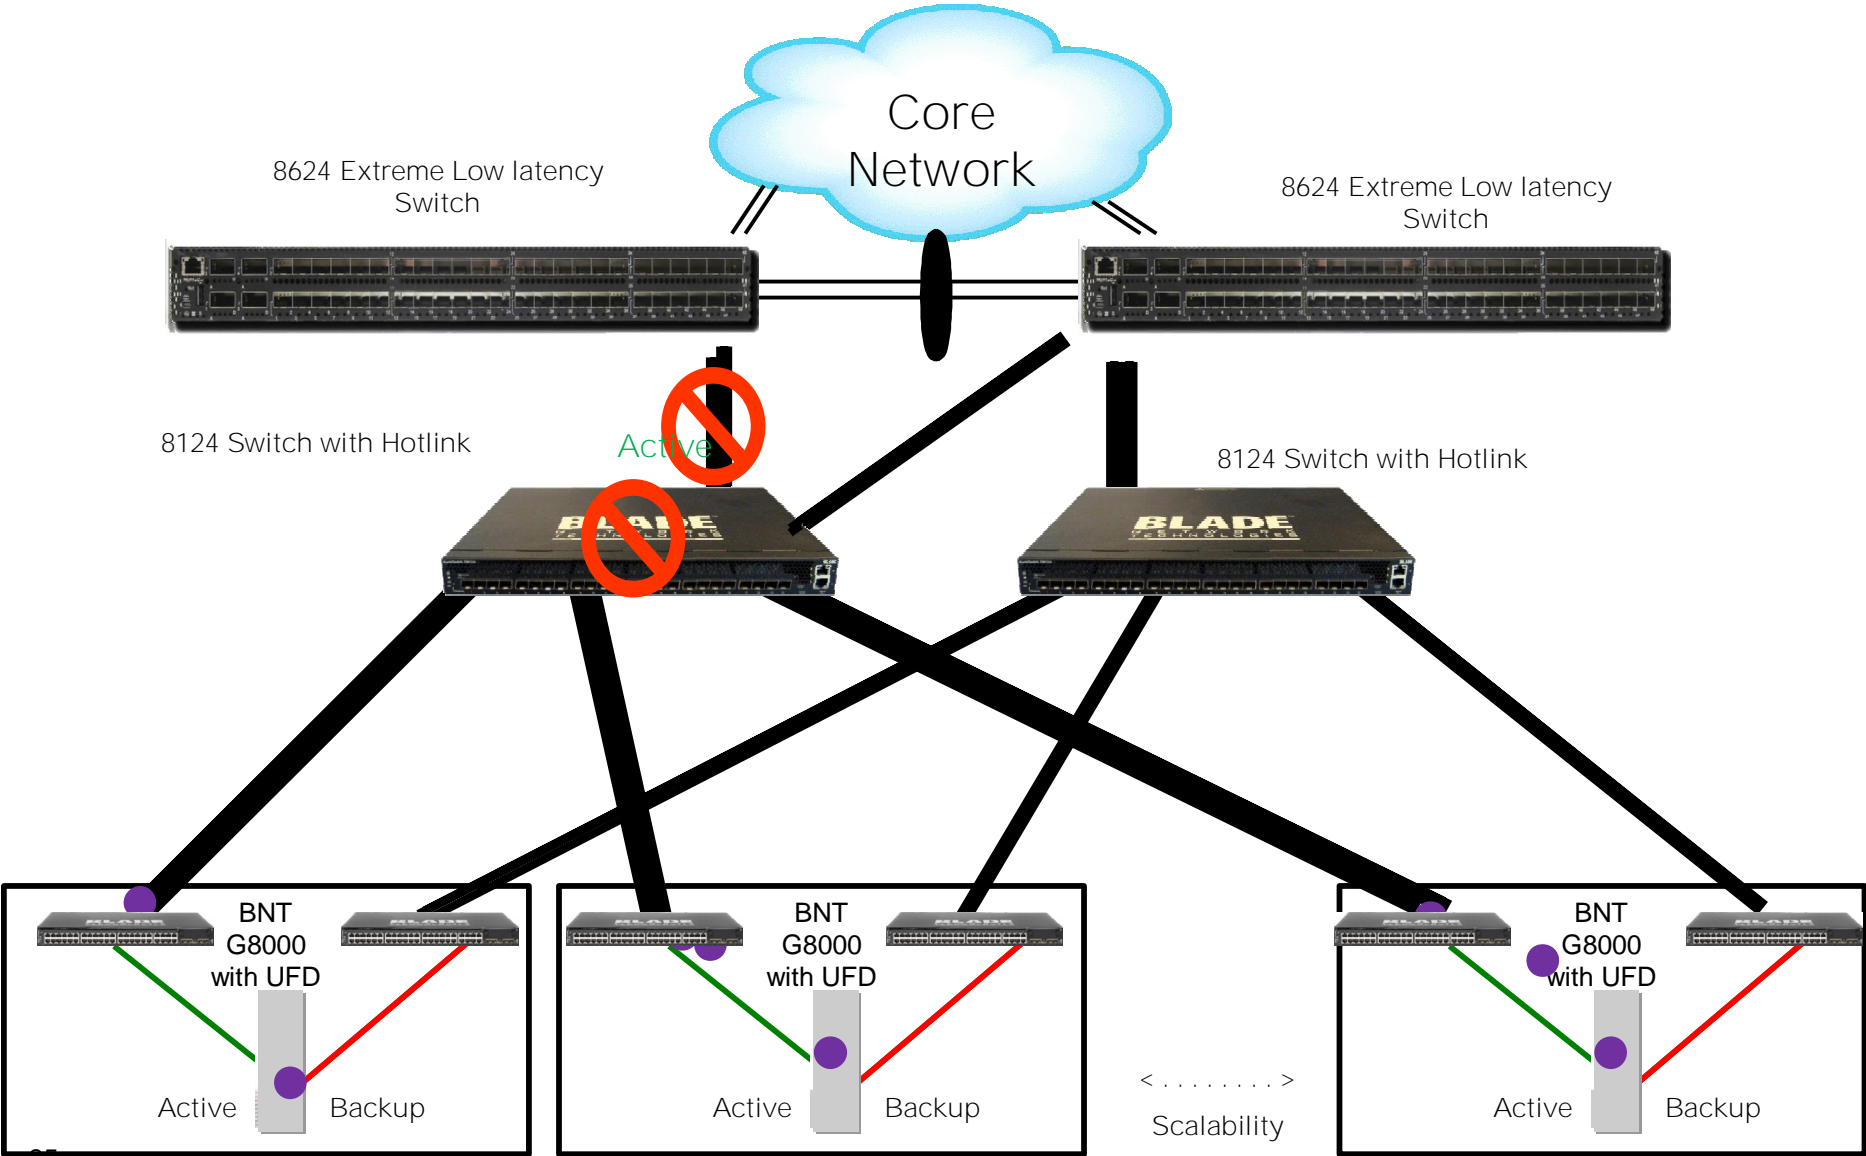
What can we do?



Hotlinks


- § Designed to eliminate spanning tree in simple configurations
- § Two ports are configured as Master and Backup ports
 - When the Master port is Link up the Backup port is held down
 - When the Master port is Link down the Backup port is automatically enabled



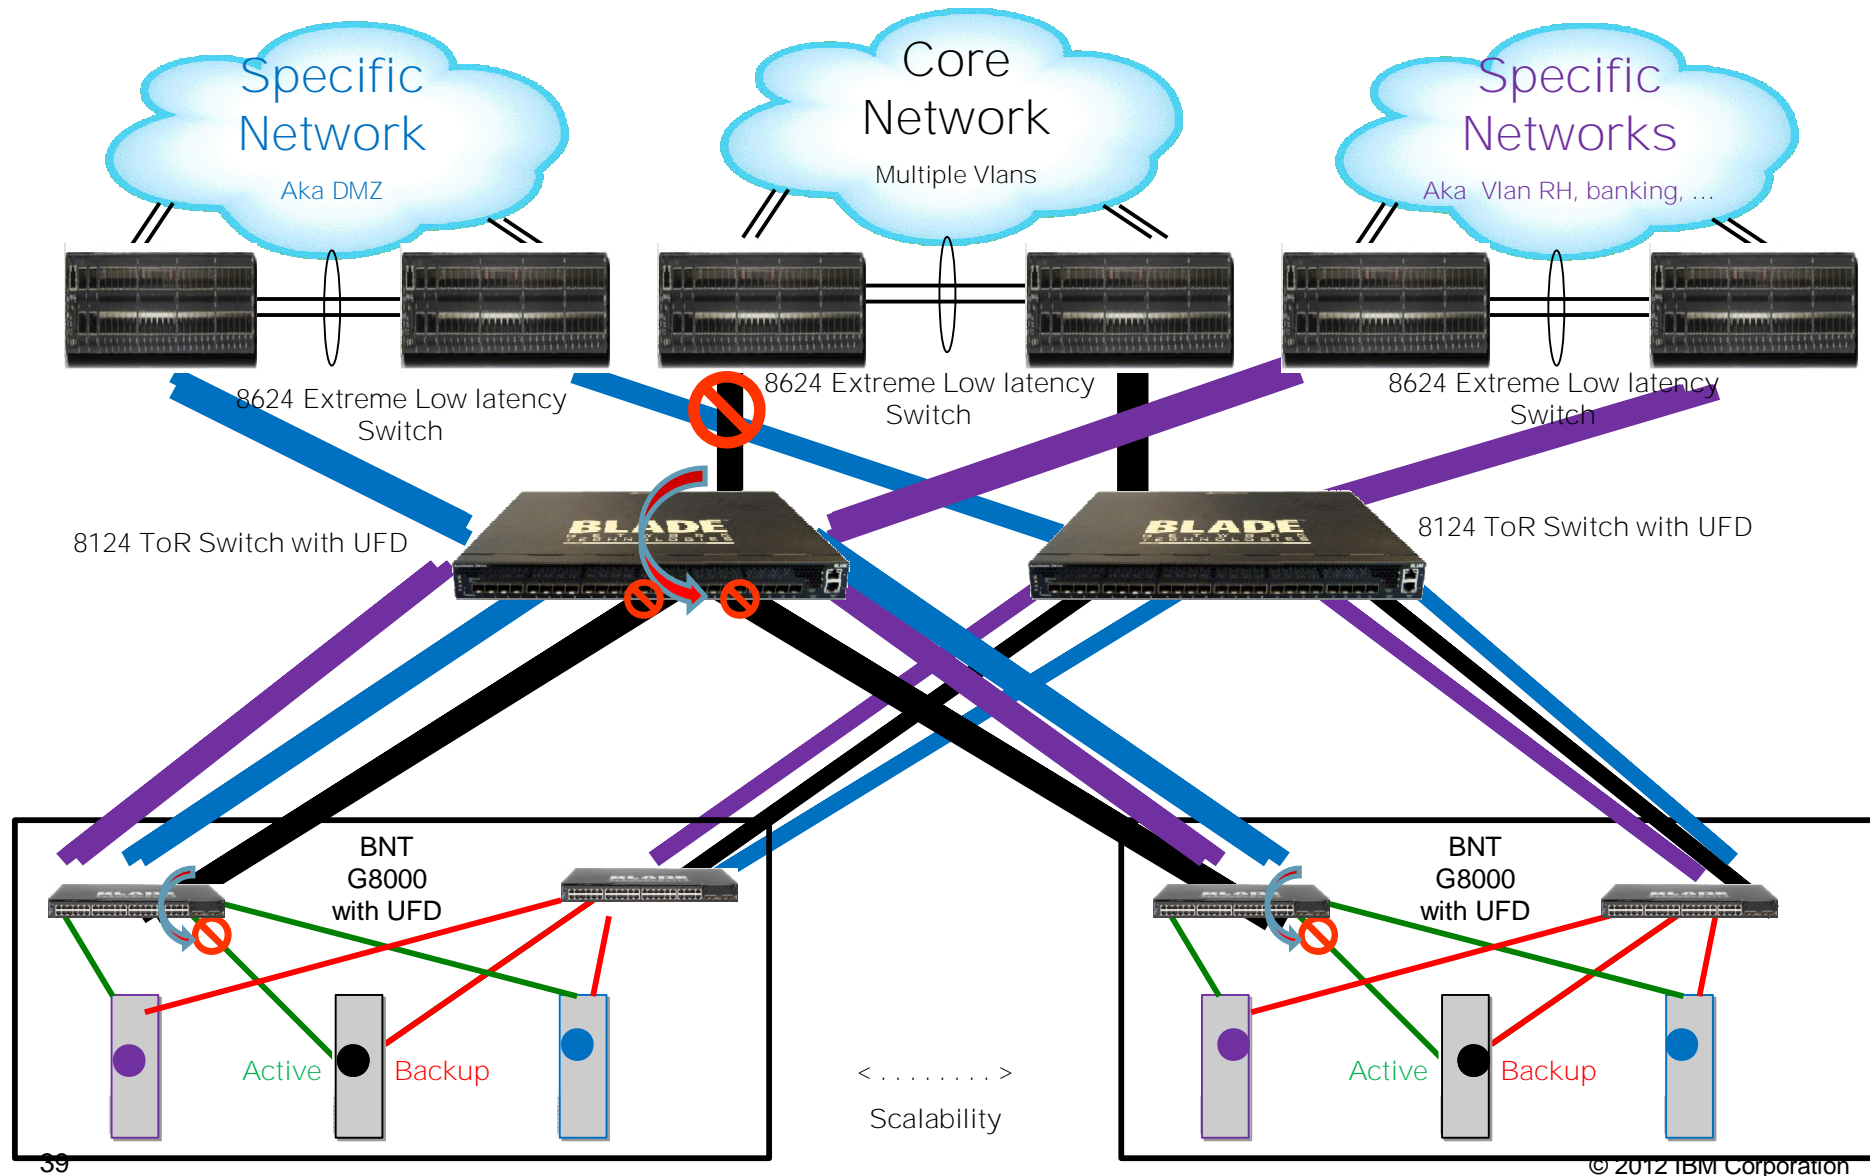


The diagram illustrates a network architecture for BNT GbE with UFD units. At the top, three clouds represent different network types: 'Specific Network Management Vlan', 'Core Network Productions Vlan', and 'Specific Networks Vlan RH, banking, ...'. These are connected to a row of six blue square switches. Below these switches are two black 'BLADE' switches labeled 'ToR Switch With Hotlink'. Thick blue lines represent the primary 'Active' paths from the switches to the ToR switches. A green line labeled 'Active' points to the left ToR switch. A red 'X' with a slash is placed over the connection between the second switch and the left ToR switch. Purple dashed lines represent backup paths. At the bottom, three server racks labeled 'BNT GbE with UFD' are shown. Each rack has a traffic light indicator with green, yellow, and red lights. The first two racks show the green light lit, indicating 'Active' status. The third rack is preceded by an ellipsis '< >' and labeled 'Scalability', indicating the architecture can be expanded. The bottom right corner contains the copyright notice '© 2012 IBM Corporation'.


What can we do?







What can we do?




„ What is vLAG

- ü vLAG: Virtual Link Aggregation Group
- ü IBM solution for Multi-Chassis Link Aggregation
- ü Allows a single device to use a port channel across two upstream devices


- ü Eliminates Spanning Tree Protocol (STP) blocked ports
- ü Uses all available uplink bandwidth
- ü Provides link-level resiliency
- ü Provides fast convergence if either the link or a device fails
- ü Assures high availability
- ü It is not stacking!!

ü PHYSICAL REQUIREMENTS

- § vLAG can only be configured between same type of switches (G8264-G8264, G8052-G8052, etc.)
- § vLAG peers require a dedicated inter-switch link (ISL) for synchronization
- § A minimum of two ports on each switch are recommended for ISL use
- § A management or data port is required for Health Check




vLAG v/s STP



- **STP Topology**

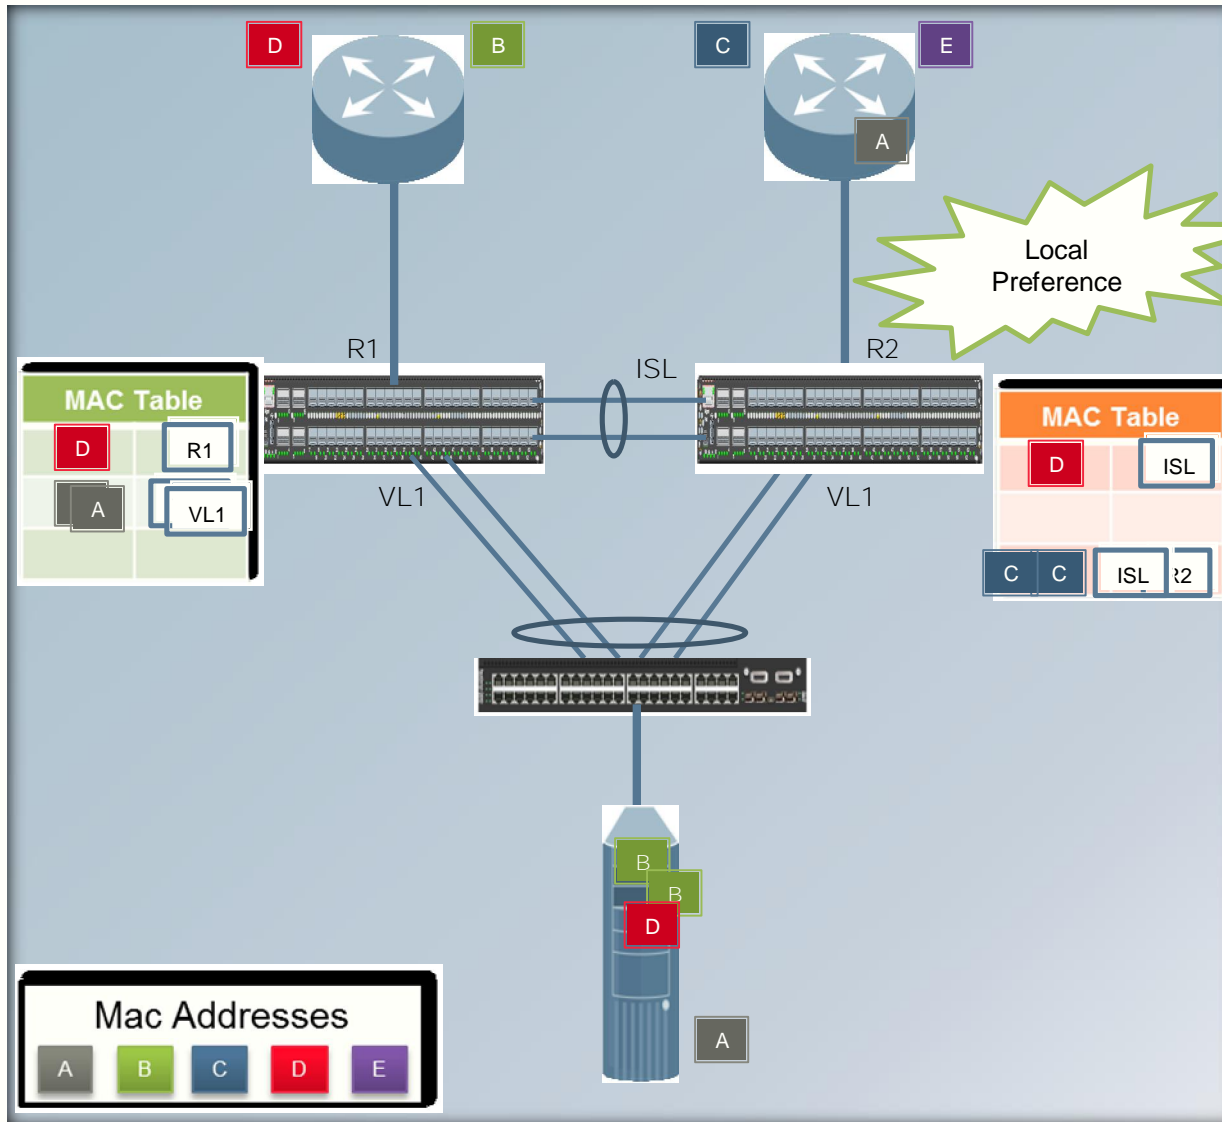
- Blocked ports reduce network bandwidth
- Slow convergence on link failure with STP
- Topology changes result in Mac flooding due to Mac flush
- MSTP solves problem partially
 - Some links are still blocked
 - Not suitable for deployments with few Vlans



- **vLAG Topology**

- No blocked ports
- Faster and predictable convergence during a link failure
- No unnecessary Mac flooding
- Standards based – Uses LACP or Static trunking
- STP can be used with vLAG, if required

FDB Synchronization in vLAG




FDB synchronization

- ∅ Avoids Mac address flooding
- ∅ Both static and dynamic addresses synchronized
- ∅ "Local Preference" used to reduce ISL traffic
- ∅ Fast synchronization methods used to reduce unnecessary flooding

Synchronization for both

1. Macs learned on vLAG ports
2. Macs learned on non-vLAG ports

Failure Handling in vLAG



Primary and Secondary communicate over management link upon ISL link failure

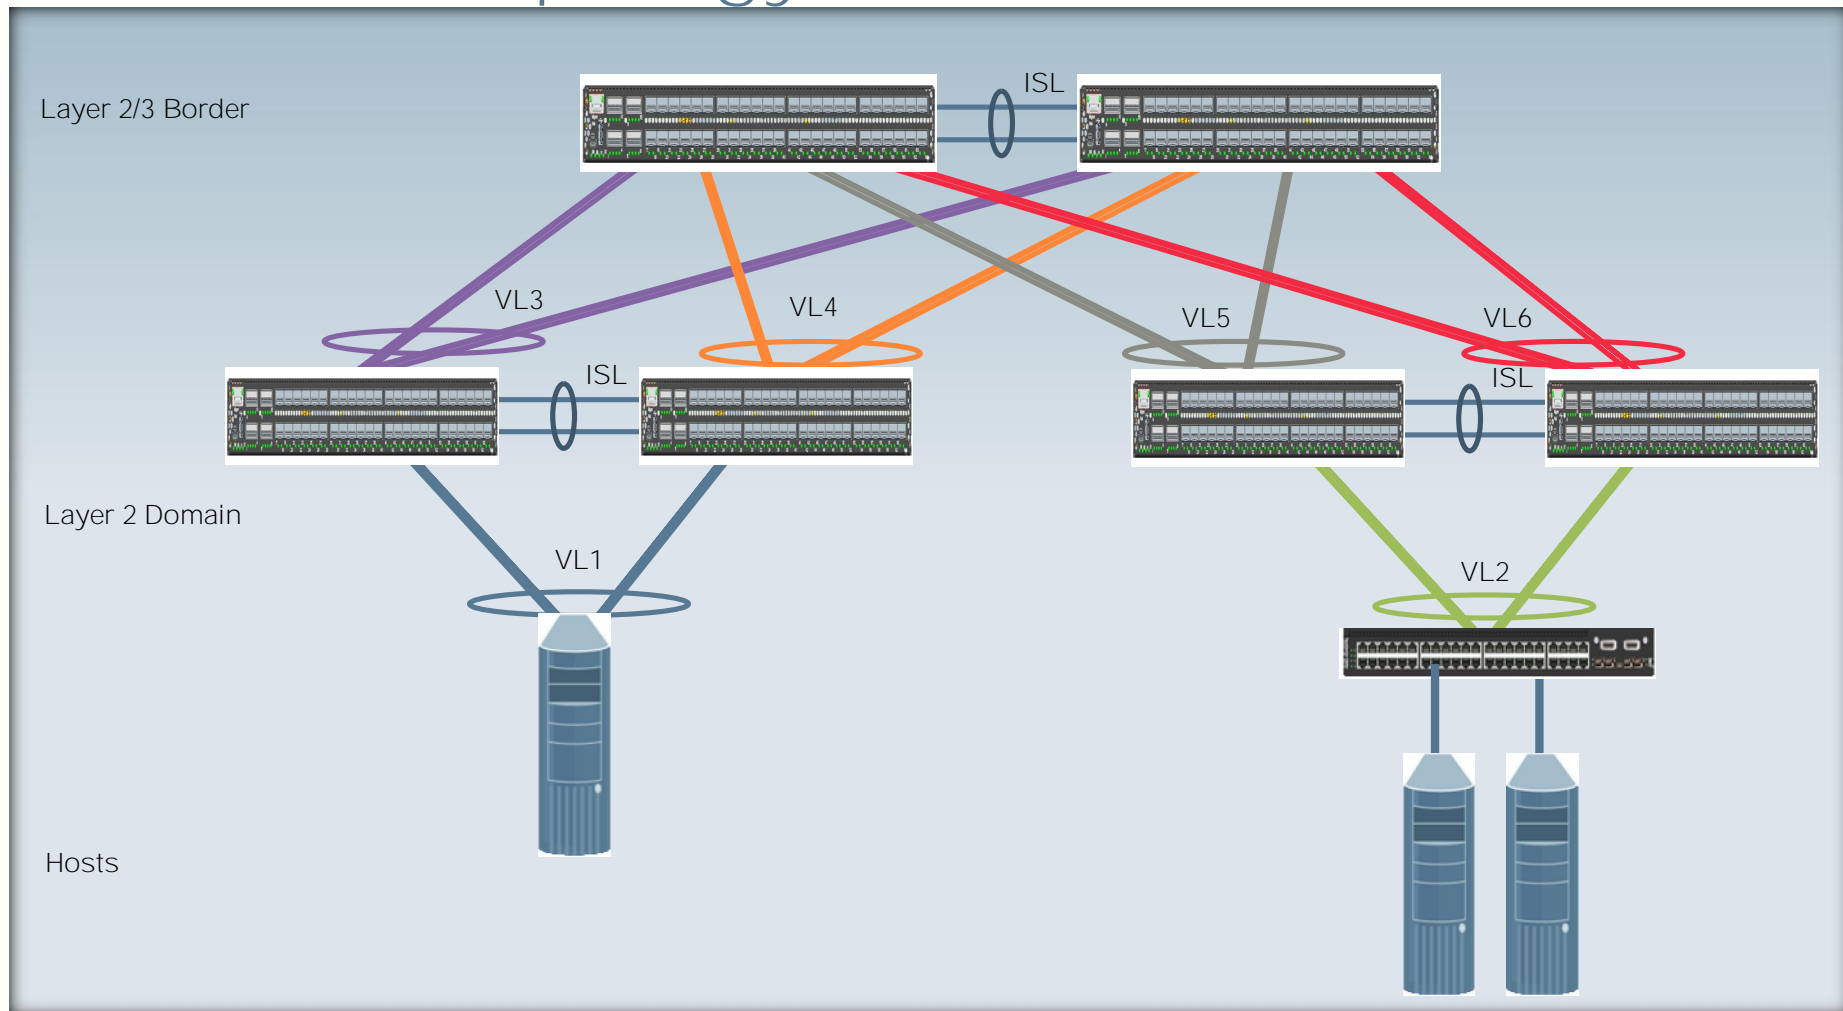
Scenario: ISL link failure

§The vLAG link on the secondary is shutdown to prevent any potential loops

Scenario: Switch goes down

§Traffic goes through the other switch and connectivity is maintained


Use Case Multi-tier Topology



vLAG scales seamlessly, in a multi-tier network topology

Delivering Optimization

Optimized



Optimized for current and future workloads


Optimized

- “Pay as you grow” scalability
 - Can help clients save up to 34% (10Gb)*
 - Helps reduce OPEX - Avoids retesting, now new switch HW to install
- Exceptional Performance
 - First embedded switch with latency less than 1 microsecond
 - First embedded switch with support for 40Gb connections
 - Extreme performance with for up to 3 x 10Gb server ports per switch
- Seamless interoperability
- Designed for future
 - Flat network – leverage an IBM RackSwitch Upstream
 - Convergence – leverage EN4093 with FCoE Gateway upstream
- Virtual NIC (Virtual Fabric)
 - Up to 32 Virtual NICs


Performance vs. Dell & HP

Optimized


Unmatched Latency



Exception Throughput



Oversubscription



Cisco UCS does not offer a true integrated switch!

Performance against Cisco UCS – (Part 1)

Optimized

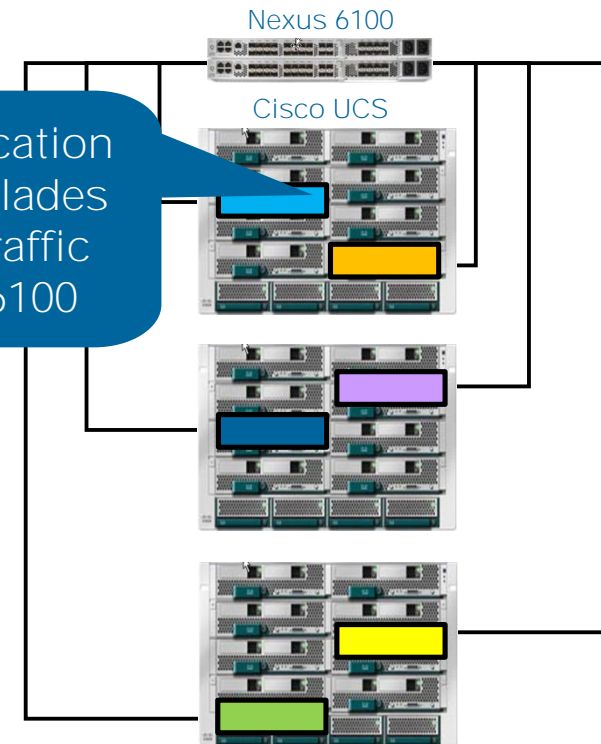
§ IBM delivers Lower Latency up to 70%

- IBM communications stay within the chassis at less than 1 microsecond (us)
 - Node to EN4093 to Node
- Cisco communications must exit the chassis to go to the Fabric Interconnect adding an extra step meaning you have latency of the Fabric Extender (0.65us) and Fabric Interconnect (2us) to consider.
 - Node to FEX to UCS6k to FEX to Node
 - Latency $3.3\mu s = 0.65\mu s + 2\mu s = 0.65\mu s$

§ Management Performance

- IBM has separate management ports
- Cisco UCS has management traffic going over the same networking ports as their data!

Communication between blades requires traffic to Nexus 6100



Blade to blade communication flows north-south through the external Nexus, causing latency from request / response traffic. Added network latency will impact the overall workload the servers can support.

Designed for Energy Efficiency

Optimized


§ IBM delivers Energy Advantages

- Help clients reduce Utility Expenses
- Help clients get more per limited power envelope per rack
- Help reduce heats related expenses like AC requirements

§ Competitive Bullets

- Slight savings for clients using 2x10Gb per server
- 50-75% power savings for clients using 4x10Gb per server
- Savings against Cisco UCS can be significant when you factor in the power consumption of the Fabric Extender and the upstream Fabric Interconnect required to compete with the IBM switch

Power Consumption (Watts)




Simplify Network Infrastructure

Key Benefits:


- Lower Cost and Power
- Better Performance
- Lower latency
- Easier to manage
- Less space
- Critical for Virtualization

Flatten network topology via enhanced stacking that reduces number of network devices needed and allows users to manage all devices as a Single Logical network




One Logical Switch

Traditional setup (Today)



IBM vision to flatten network topology



Converged Enhance Ethernet & Data Center Bridging

Optimized

§ IBM Flex System Fabric E4093 10Gb supports

- CEE/DCB Protocols implementation to enable lossless converged Ethernet Fabric - Ideal for NAS, iSCSI and FCoE environments
- FCoE Transit switch capability to deploy FCoE in a multi-hop

§ IBM Flex System Fabric E4093 10Gb enables FCoE implem

- By providing lossless Ethernet fabric
- By providing fair bandwidth usage in converged Ethernet fabri
- By providing FC equivalent security in Ethernet fabric


**Planned Support
Mid 2012**

Priority-based Flow Control




- § Supports zero loss under congestion
- § Pauses traffic based on classes of service
- § Manages a high-volume of data from a single protocol source
- § Guarantees delivery via on-off flow control

Enhanced Transmission Selection



- § Enables optimal bandwidth management of multiple protocols on single virtual link
- § Assigns each class of traffic a specified bandwidth percentage
- § Allows other traffic to use the spare bandwidth as available

Congestion Notification



- § Signals data transmission systems to indicate the status of network congestion
- § Allows network devices to determine when to send packets
- § Ensures end-to-end notification of network congestion

IBM Virtual Fabric

Optimized

Fast, flexible, reliable I/O that's ready for the future and fits into your existing data center



Reduce Cost

- Reduce CAPEX up to 54%
- Reduce energy costs up to 52%
- Protect Investment with Convergence Ready*



Reduce Complexity


- Reduce cables up to 86%
- Reduce Switches, Adapters up to 75%
- More VMs per systems (Up to 32*)
- Higher availability & security – vNIC isolation
- Dynamic bandwidth allocation - No downtime
- Ability to operate in transparent mode



Improve Performance

- Up to 10x I/O bandwidth
- Lower Latency
- Maximum bandwidth control in both directions

IBM Virtual Fabric



Reduce Costs

- Reduce CAPEX up to 50%
- Reduce energy up to 50%
- Investment Protection with Convergence Ready Solution

Reduce Complexity

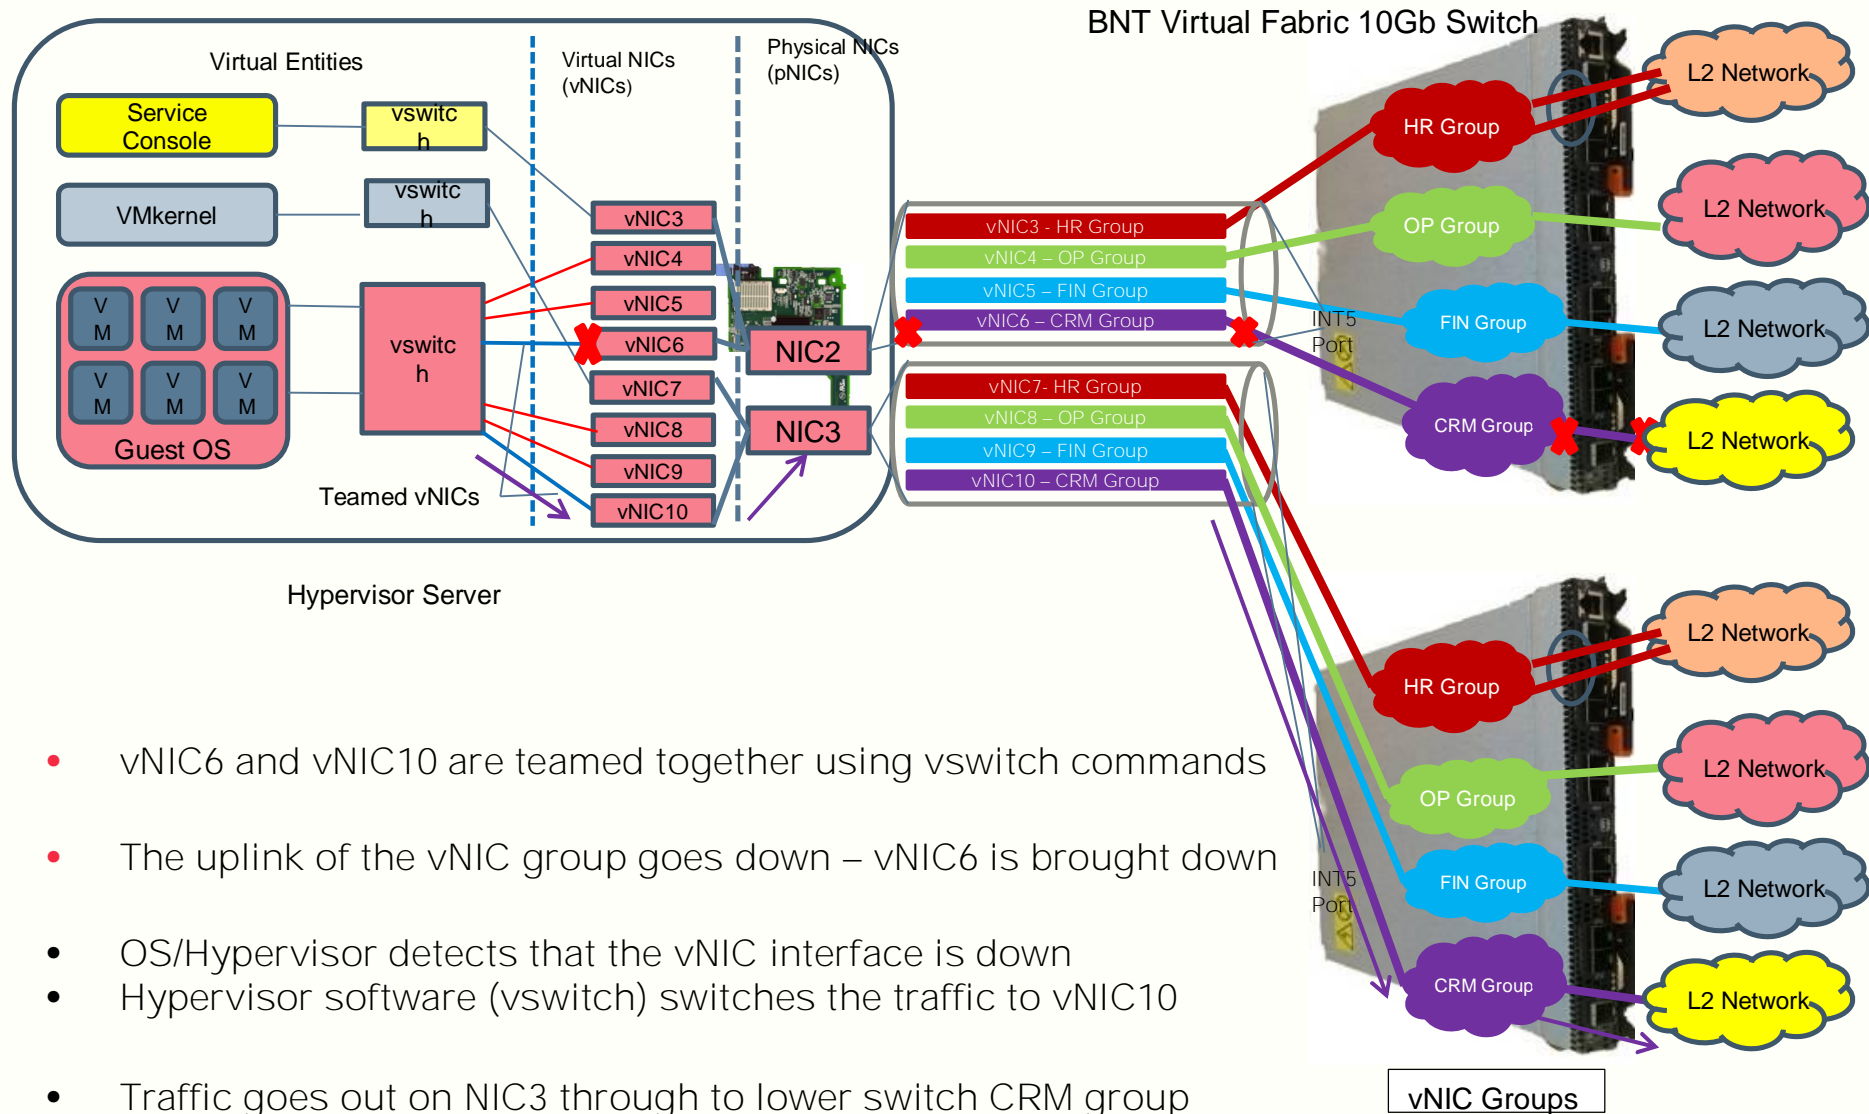
- Reduce hardware up to 75%
- Reduce cables up to 86%
- More VMs onto fewer systems
- Higher availability & security

Improve Performance

- Up to 10x I/O bandwidth
- Dynamically allocate bandwidth with no downtime
- Lower Latency
- Max performance per vNIC

BUILT ON INDUSTRY STANDARDS


Virtual Fabric Solution – vNIC Teaming



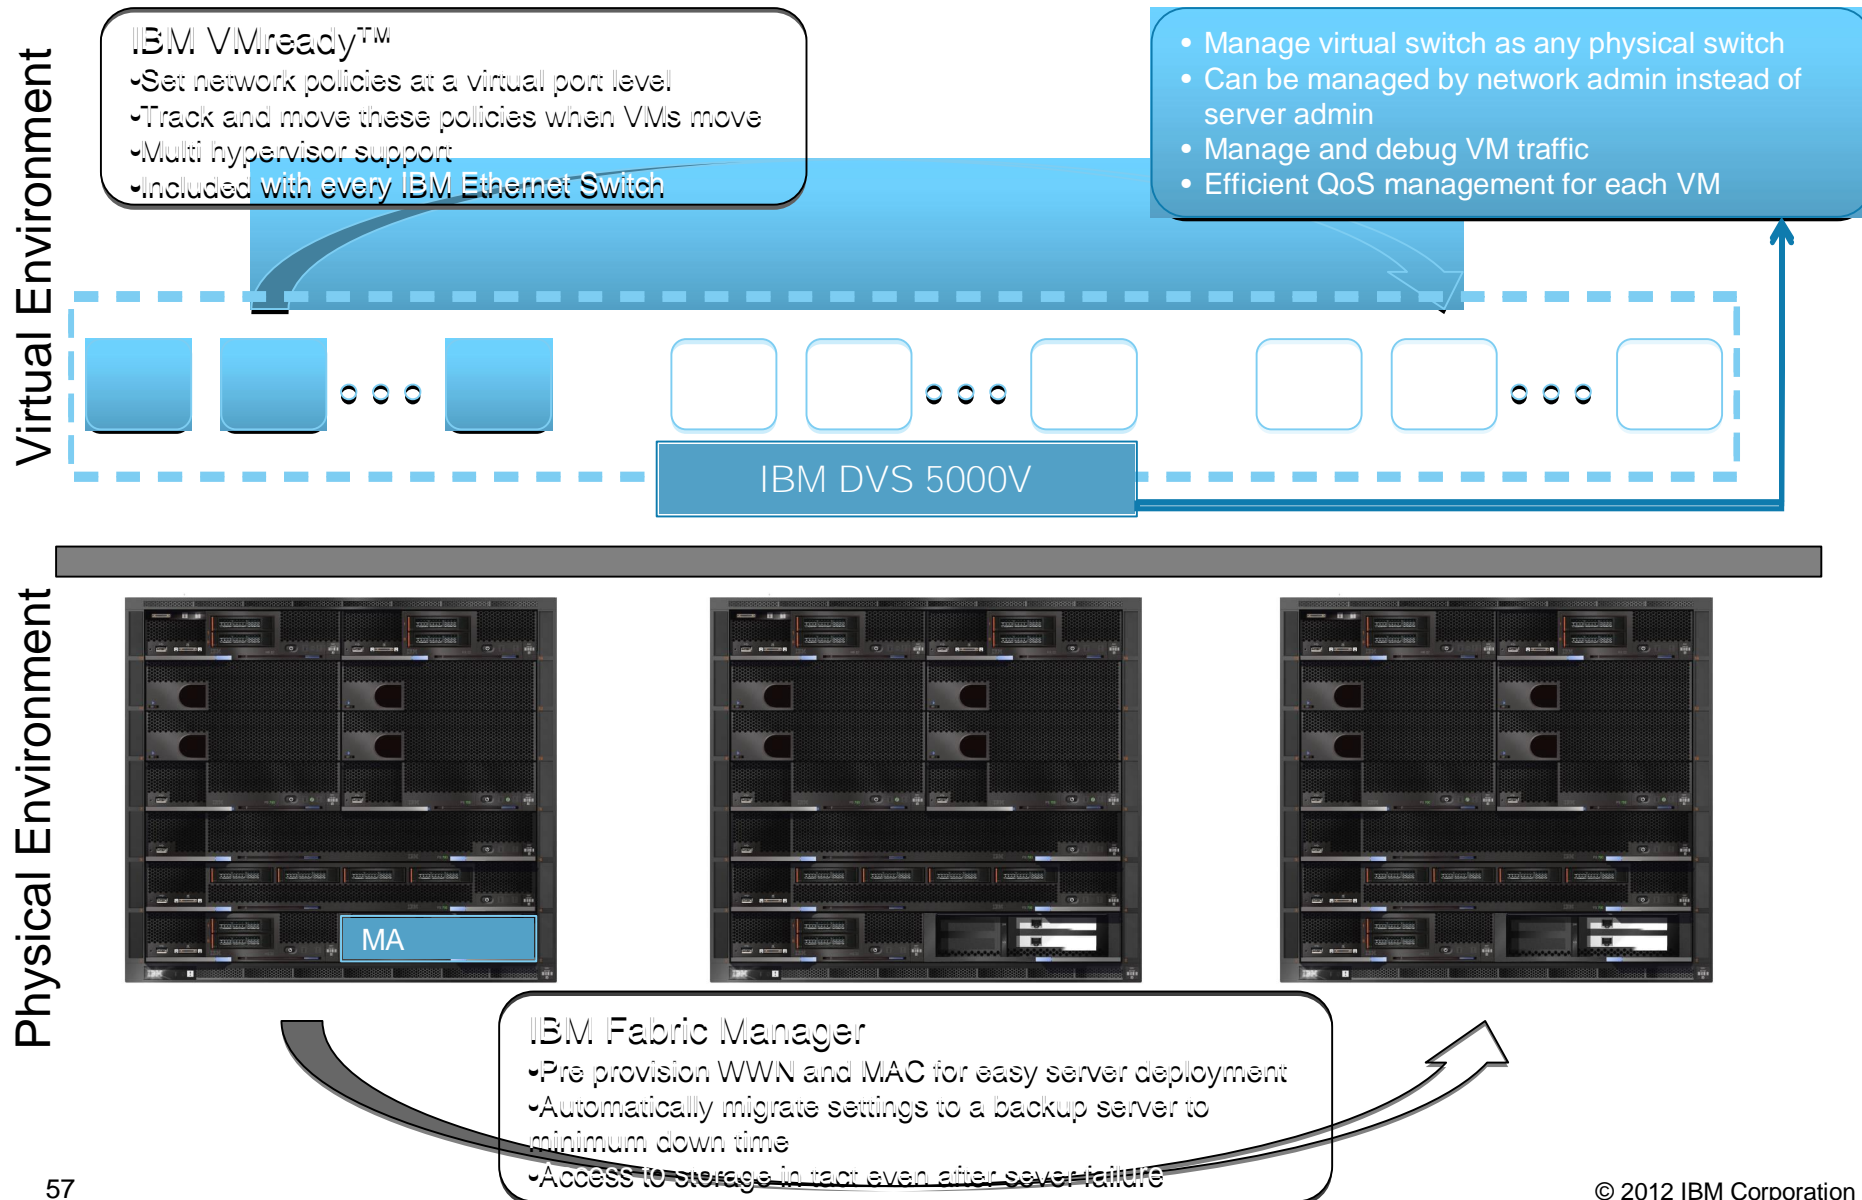
- vNIC6 and vNIC10 are teamed together using vswitch commands
- The uplink of the vNIC group goes down – vNIC6 is brought down
- OS/Hypervisor detects that the vNIC interface is down
- Hypervisor software (vswitch) switches the traffic to vNIC10
- Traffic goes out on NIC3 through to lower switch CRM group and out on the same L2 network

Designed for Optimization

Automated



Automated to manage both Physical and Virtual environment



VMready™ - making the network VM aware

Automated

§ Enterprise Data Center Network Virtualization

- VM Detection: automatically discovers and identifies the Virtual Machines
- VM groups – Grouping of similar Virtual Machines to simplify management tasks
- NMotion™ - Automatically track migrating VMs to maintain network configuration

§ Vendor Neutral

- VMready works with all major virtualization offerings



vmware™



§ Integration with VMware vCenter

- Single pane of management for both VMready switches and ESX vSwitches
- Rich display of VM info such as IP addresses, VM name and ESX server location

§ Low cost – Software license included with the switch


§ Software runs on switch, not the servers

§ Accounting and Auditing

- VM Migration is detected by VMready and reported
- Traffic per VM can be audited

How VMready works

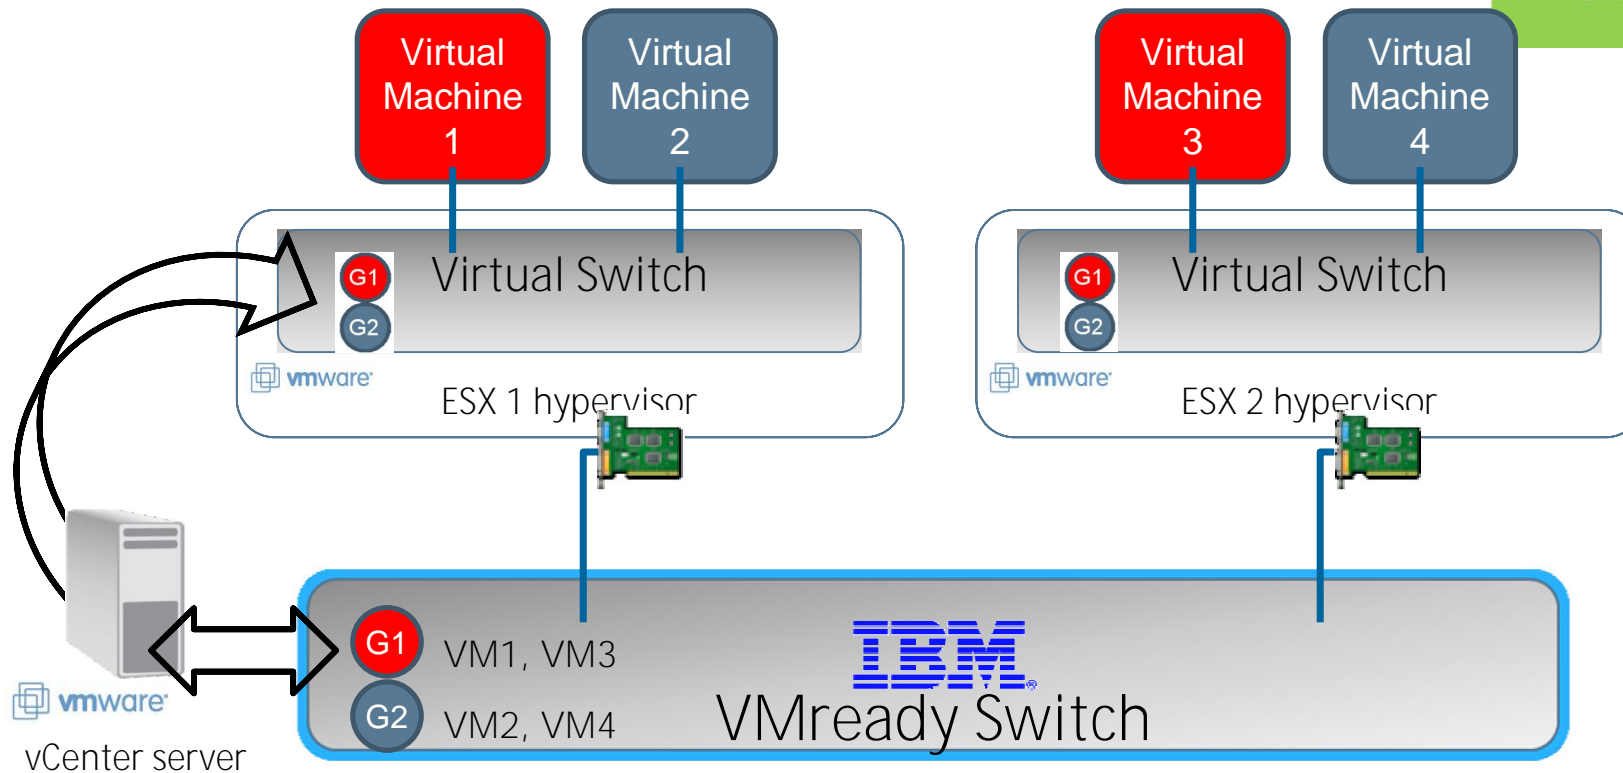
Automated



1. VMready creates a virtual port for each VM can be configured for VLANs, ACLs, QoS etc.
2. Virtual port configurations can be synchronized with vSwitches via APIs
3. VMready see the packets sent from VMs as they migrate and moves the virtual ports and policies in real time with Nmotion™
 - Virtual Machine stay attached and secure


Integration with VMware vCenter

Automated



- § VMready integrates with VMware vCenter
- § Can automate the creation of port groups on ESX vSwitches
- § Ensures ensure network settings are consistent
- § VMready also polls vCenter for VM info such as VM names, location

Software Defined Networking



What are networks

Networks are a mechanism to move traffic from one node to another node, based on networking rules and definitions.

These networking rules and definitions are created using complex algorithms where each individual network devices needs to understand what it's neighbours are doing.

Defining these rules take time and resources, so networks like to remain static to provide optimal connections.

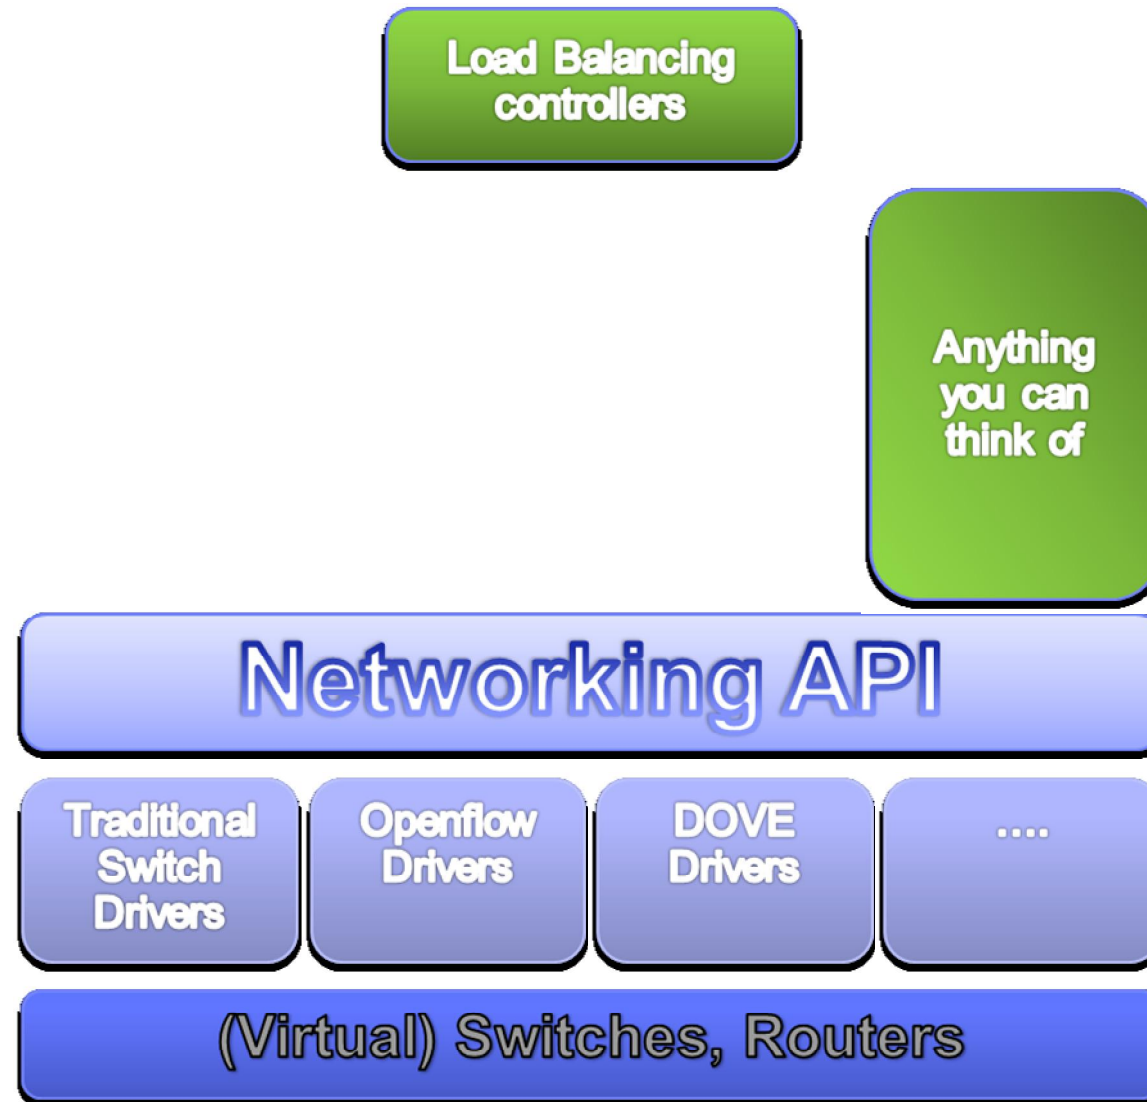
What is Software defined networking

Networks that are controlled by your application needs and not your physical infrastructure

SDN networks are networks that can change when needed without the penalty of resources and convergence time.

SDN is whatever you want your network to be

SDN stack



Integrated System Fabric Trends

What scaling technologies are we pursuing?



Standalone

- *Optimal local switching*
- *Low latency*
- *Converged*
- *Full L2/3*



Stacking
Single stacked switch

- *Active-active multi-link redundancies*
- *Single logical switch*

Technology Trend




Fabric
Switch cluster

+


- *Arbitrary topologies*
- *Scale-out cross-sectional bandwidth*
- *Numerous redundant paths*
- *Scales to large number of ports*
- *Unification of physical & virtual*

Layer-3




- Established technology
- Standards based
- Distributed control plane
- Large scalability
- HA with fast convergence
- Small layer 2 (limits VMs)
- Many devices to manage

vLAG

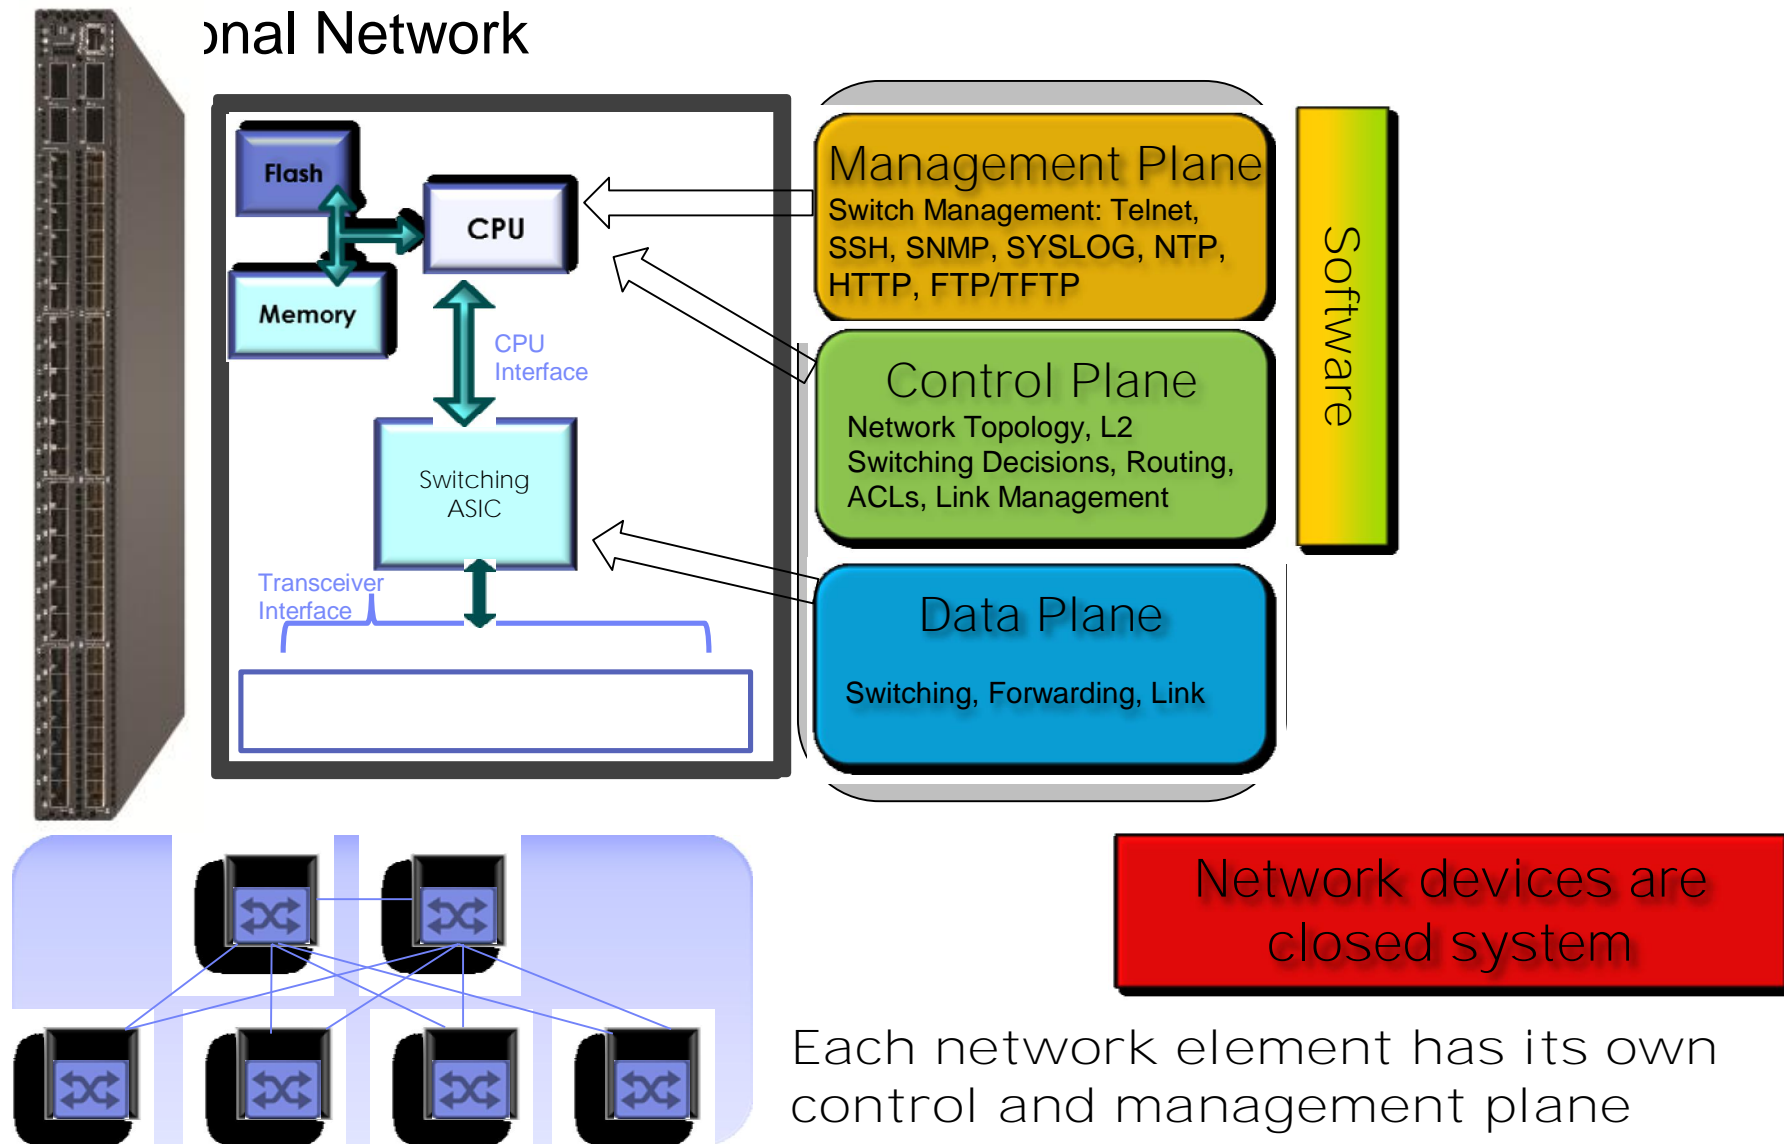


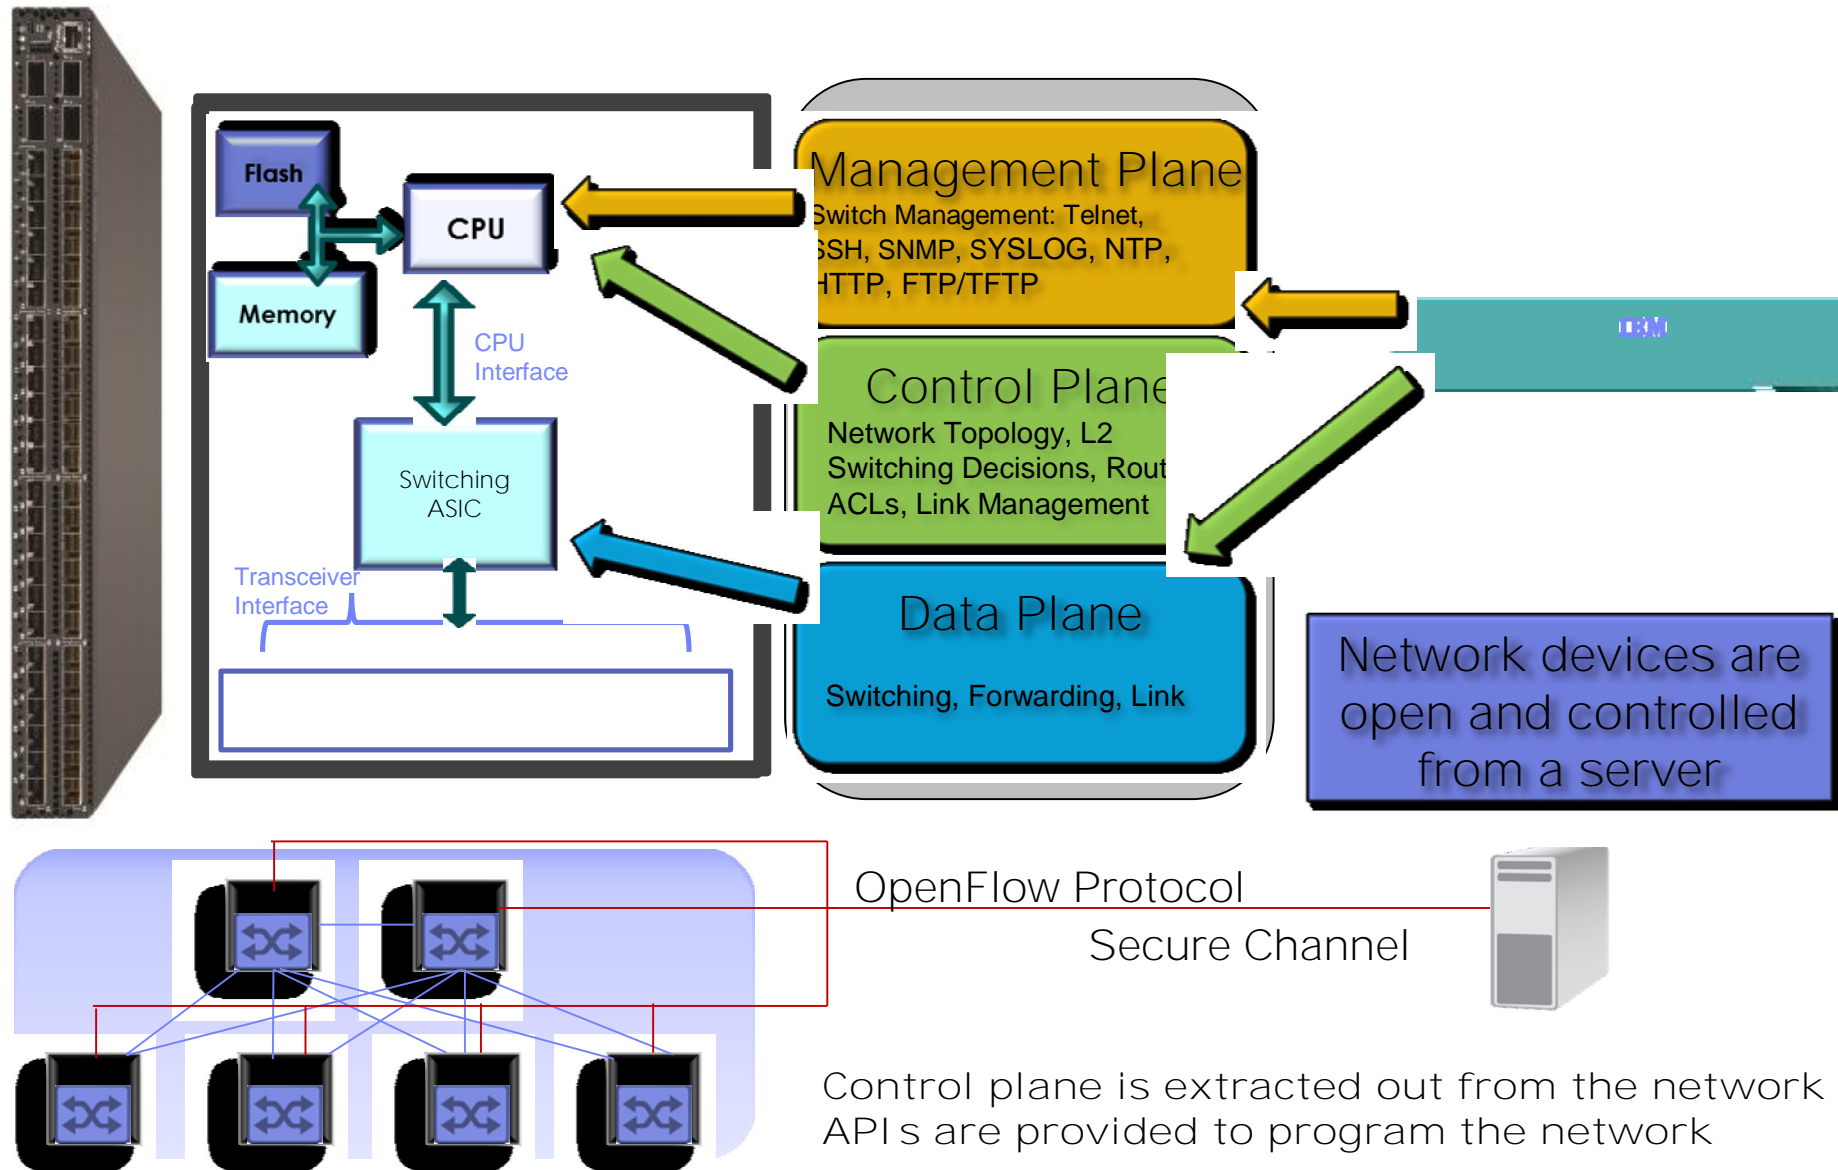
- Large layer-2
- Distributed control plane
- Large scalability
- HA with fast convergence
- Works with FCoE and RoCE
- Proposed standards (currently vendor-specific)

OpenFlow




- Large layer-2
- Distributed control plane
- Large scalability
- HA with fast convergence
- Enables disjoint multi-pathing
- Emerging technology
- Client acceptance barrier





OpenFlow can also be used to create a disjoint multi-pathing fabric.



- § Each switch has Layer-2 forwarding turned off
- § Each switch connects to OpenFlow Controller
- § OFC discovers switches and switch adjacencies
- § OFC computes disjoint physical paths and configures switch forwarding tables
- § All ARPs & FCoE FIPs go to OFC, with request for disjoint or not. If disjoint, OFC selects disjoint path

IBM OpenFlow : ONLY Supported OF Switch

§ IBM first OpenFlow 1.0 with advanced functionalities on IBM RackSwitch G8264

§ IBM first vendor to ship OpenFlow functionality on 10G switches

§ Works with ALL Openflow controllers – being DEPLOYED with:

- Beacon
- Big Switch
- NEC pFlow



IBM RackSwitch G8264


§ IBM is an original member of ONF (Open Networking Forum)

§ IBM first to demonstrate 10G-based OpenFlow solution



- Interop in May 2011; Open Networking Summit in October 2011
- Now GA level

§ Multiple connections to controllers for high availability

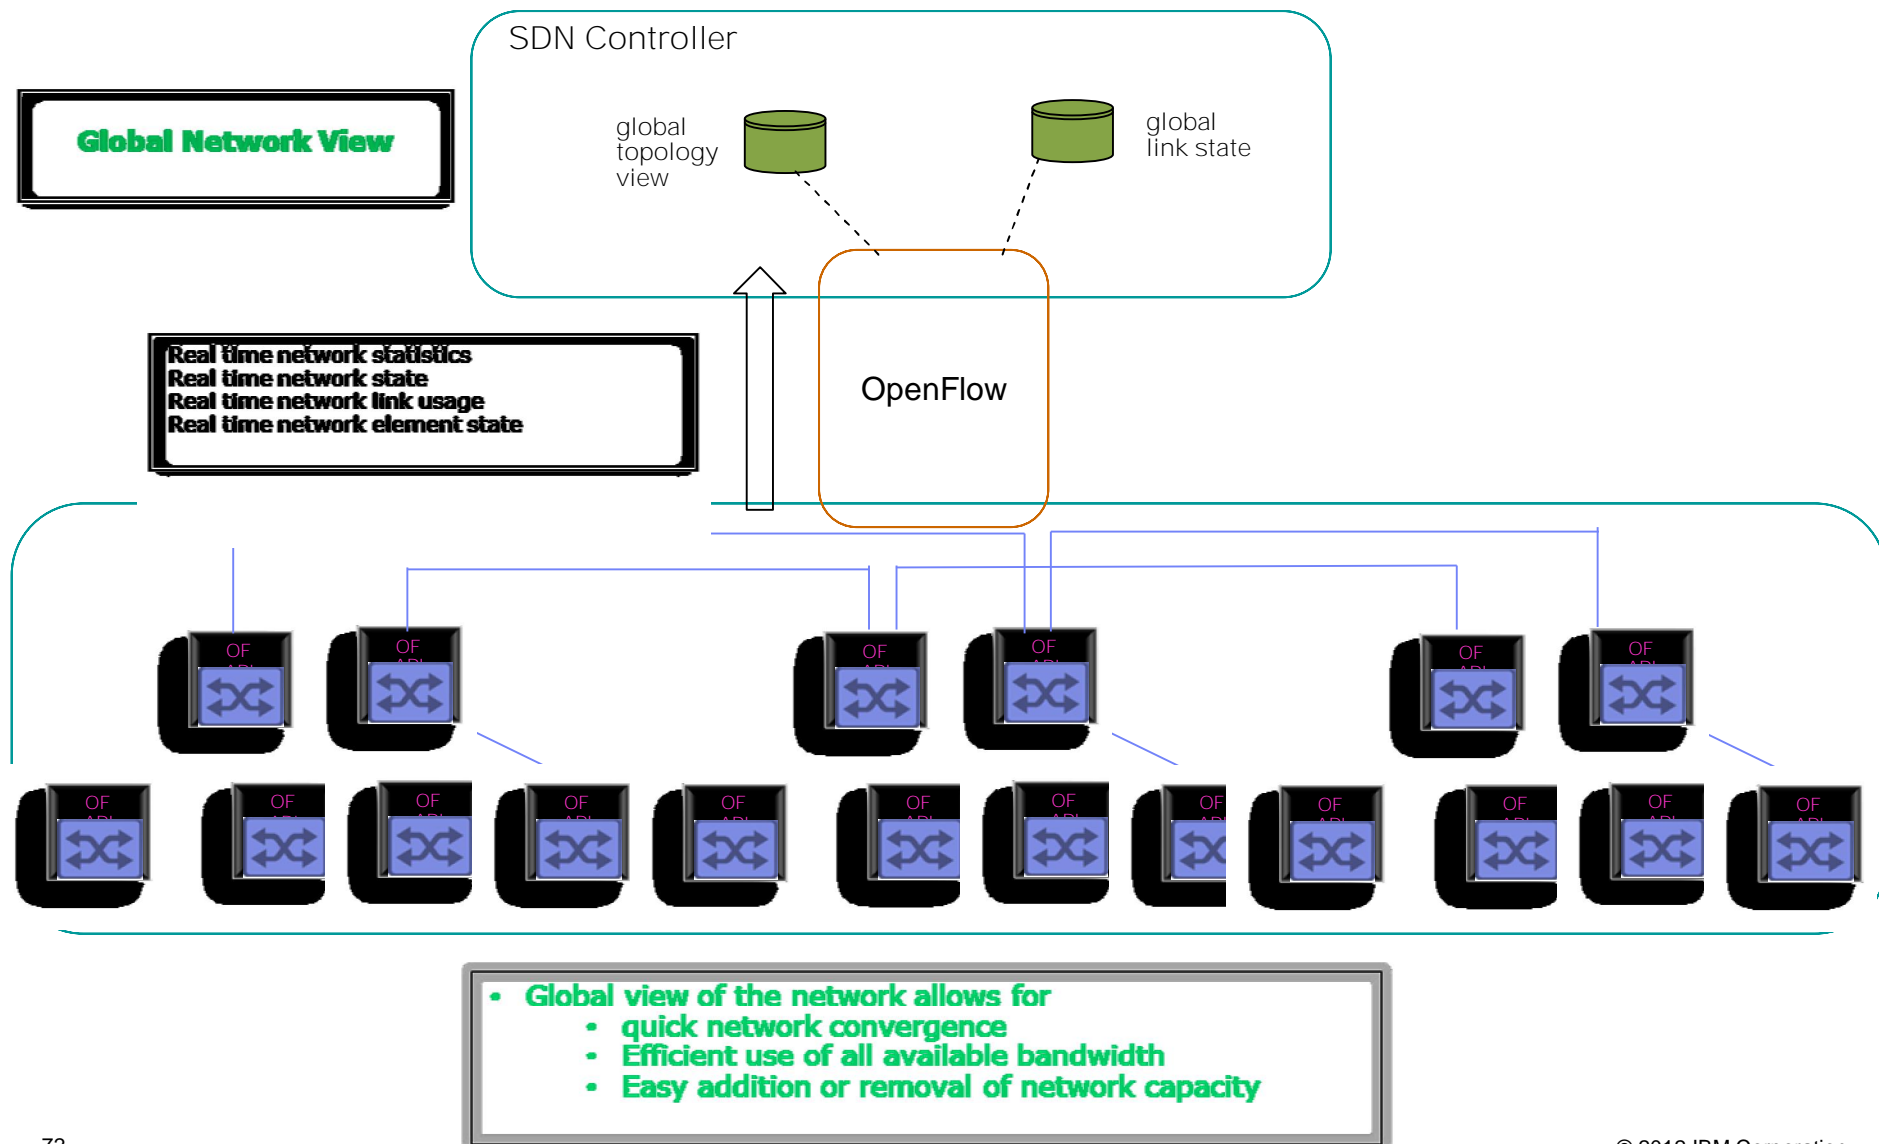
- Number of FDB based flow entries : 128000
- Number of ACL based flow entries : 1000
- L3 connectivity to the controller
- Openflow ports
 - Flow based switching
 - No L3 based switching
 - No STP and its variants



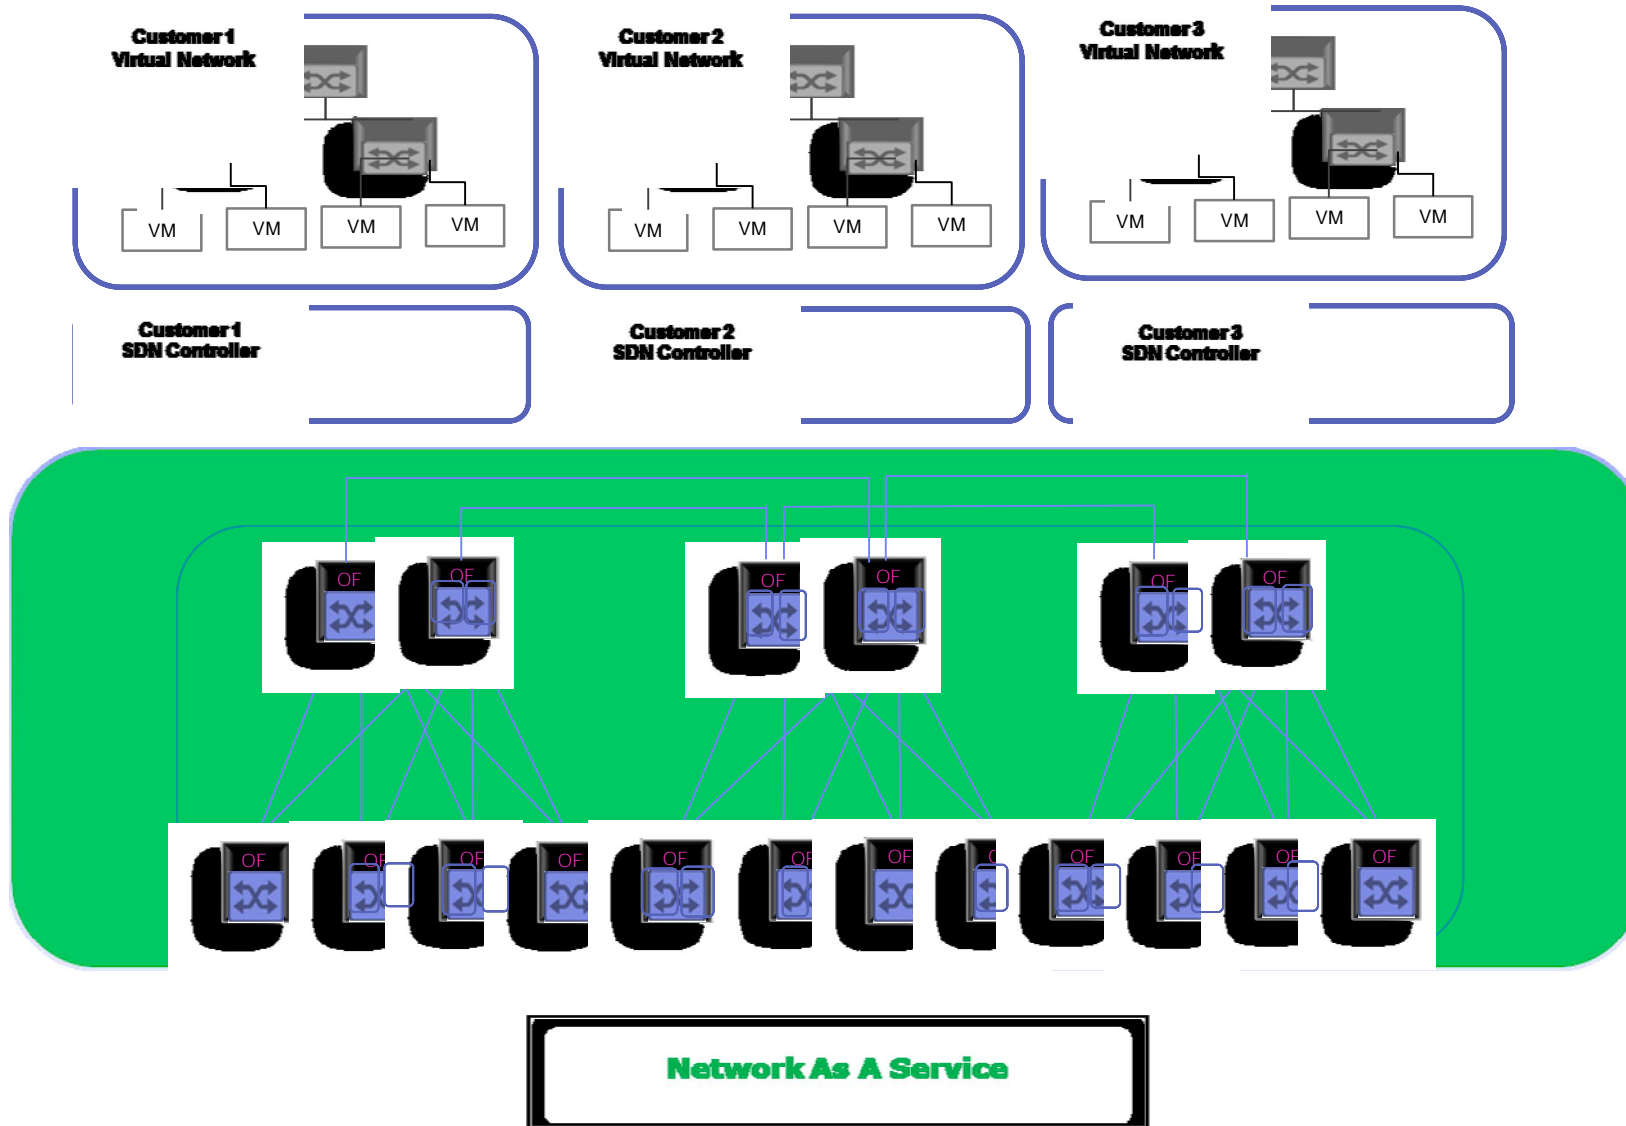
Software Defined Network: Let's build a smarter network



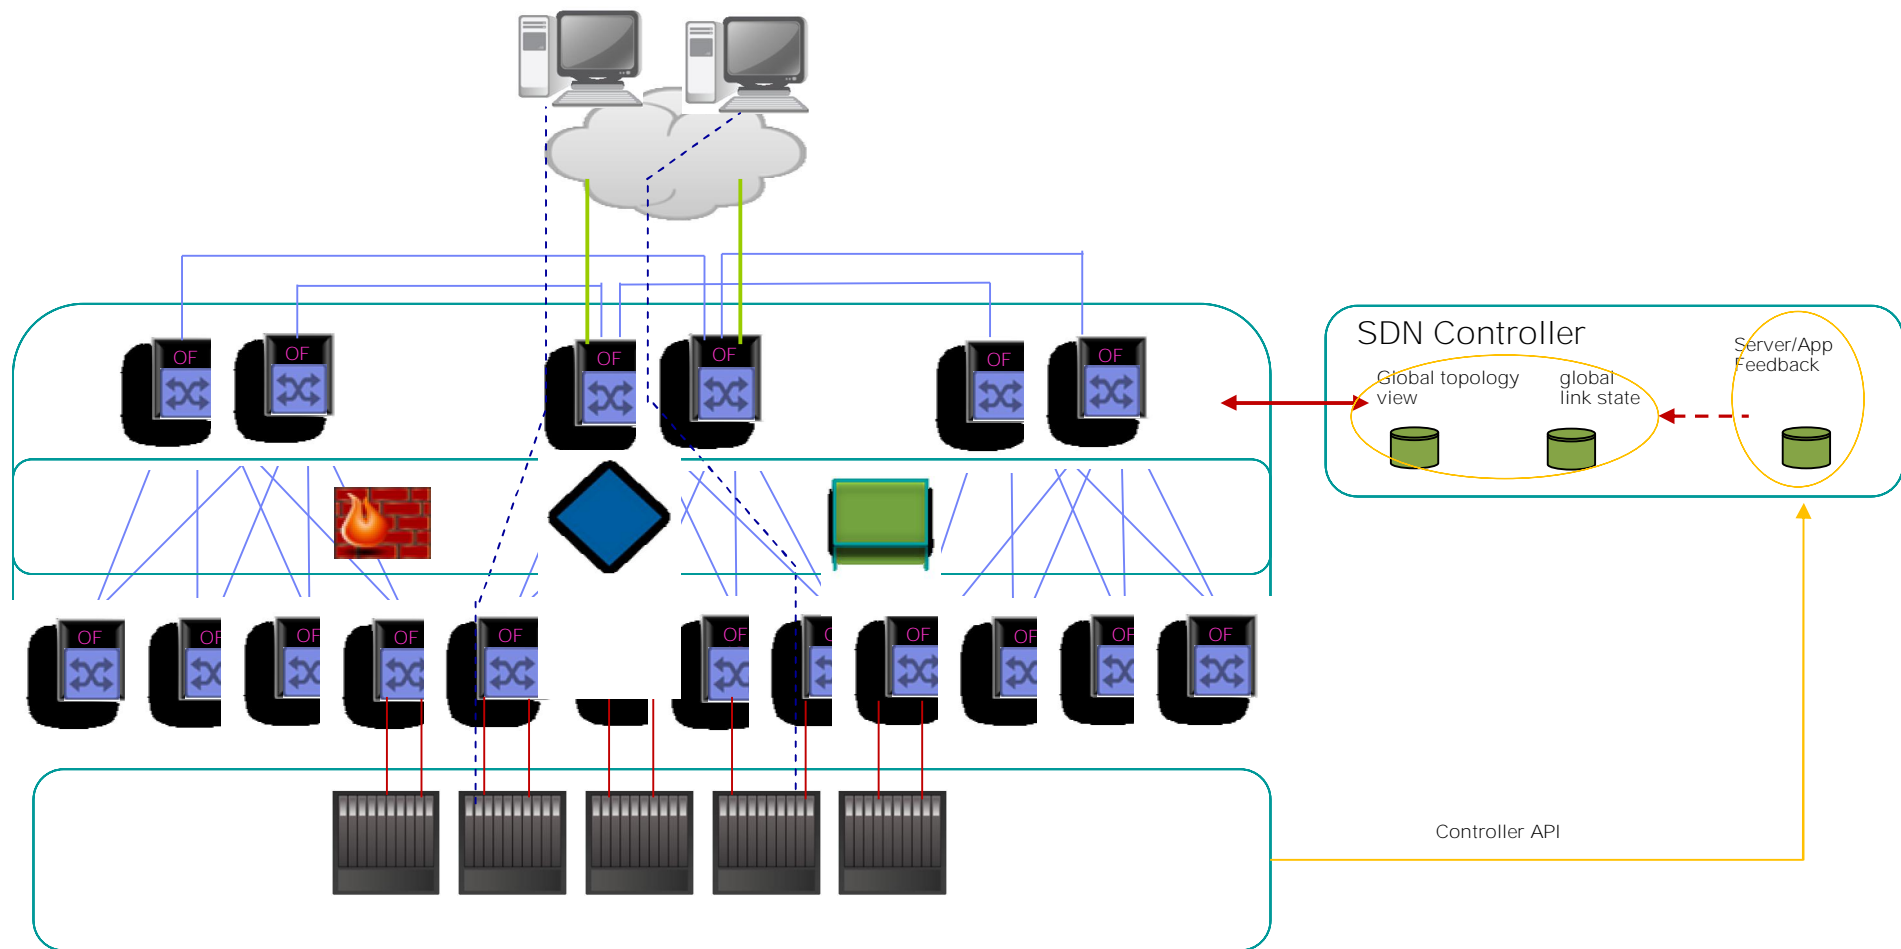
SDN: Smarter Network for a Smarter Planet



SDN: Cloud Multi Tenancy




SDN: Smarter use of Network and Appliances




Appliance sharing let's network provider use the resources more effectively

OpenFlow Based SDN for Cloud Infrastructure



Are You an Innovator?




http://en.wikipedia.org/wiki/Crossing_the_Chasm
<http://en.wikipedia.org/wiki/File:Technology-Adoption-Lifecycle.png>

IBM is a leader in SDN/OpenFlow space

- § First vendor to provide a 10Gb/40Gb OpenFlow switch as a GA product – IBM G8264
- § Initial member of ONF (Open Network Foundation)
- § Active participation in ONF standards process and market education effort
- § Live demonstration of IBM G8264 OpenFlow switch at
 - May 2011 : Interop at Las Vegas, NV
 - Oct 2011 : ONS (Open Network Summit) at Stanford, CA
 - Nov 2011 : SC11 (Supercomputing 11) at Seattle, WA)
- § OpenFlow Controller Independent standards based OpenFlow 1.0 implementation
- § Customers deploying and doing PoC with IBM G8264 OpenFlow switch
- § IBM continuing to enhance OpenFlow offering on IBM switches


Designed for Integration

Integration



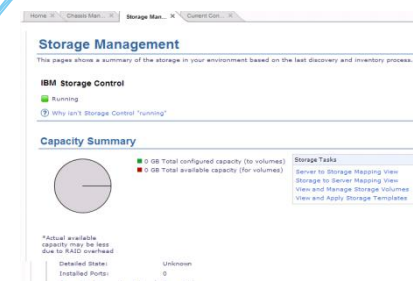
Integrated: Tightly couple for Easier management

Integration



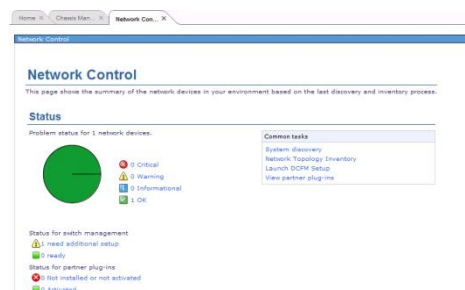
- Easy step by step server and chassis provisioning
- Intelligent virtual machine placement for optimum performance
- Image management and deployment
- Individual server control via remote presence

Server



- Manage internal and external storage as one virtualized pool
- Dynamic Storage provisioning
- Dynamic zoning and masking
- Upward integration with IBM TPC

Storage




- Logical network management – allows management of port profiles, VLANs, ACLs and QoS
- Optional Fabric Management extends QoS Management providing advanced monitoring, VM priority and rate limiting
- Network monitoring at a glance via network topology perspectives with the ability to see the components affected by network outages

Network

Integrated Management Appliance

Integration



*Management ITE / Software
Manager*

- § Integrated Management Appliance
- § Basic care and feeding for the hardware
- § Advanced virtualization management
- § Server, storage and network unification
- § Flexibility and choice for virtualization

Offering Includes Two Software Editions to Choose From

IBM Flex System Manager

- § Hardware Management
- § Energy Management (Active Energy Mgr)
- § Storage Management (Storage Manager, Storage Control)
- § Network Management (Network Control)
- § VS Management (VMControl Express)
- § Fabric Management

IBM Flex System Manager Adv. Manager

- § IBM Flex System Manager
- § Image Management (VMControl Standard)
- § Systems Pools (VMControl Enterprise)
- § Workload Management (VMControl Enterprise)
- § Service Fabric Provisioning

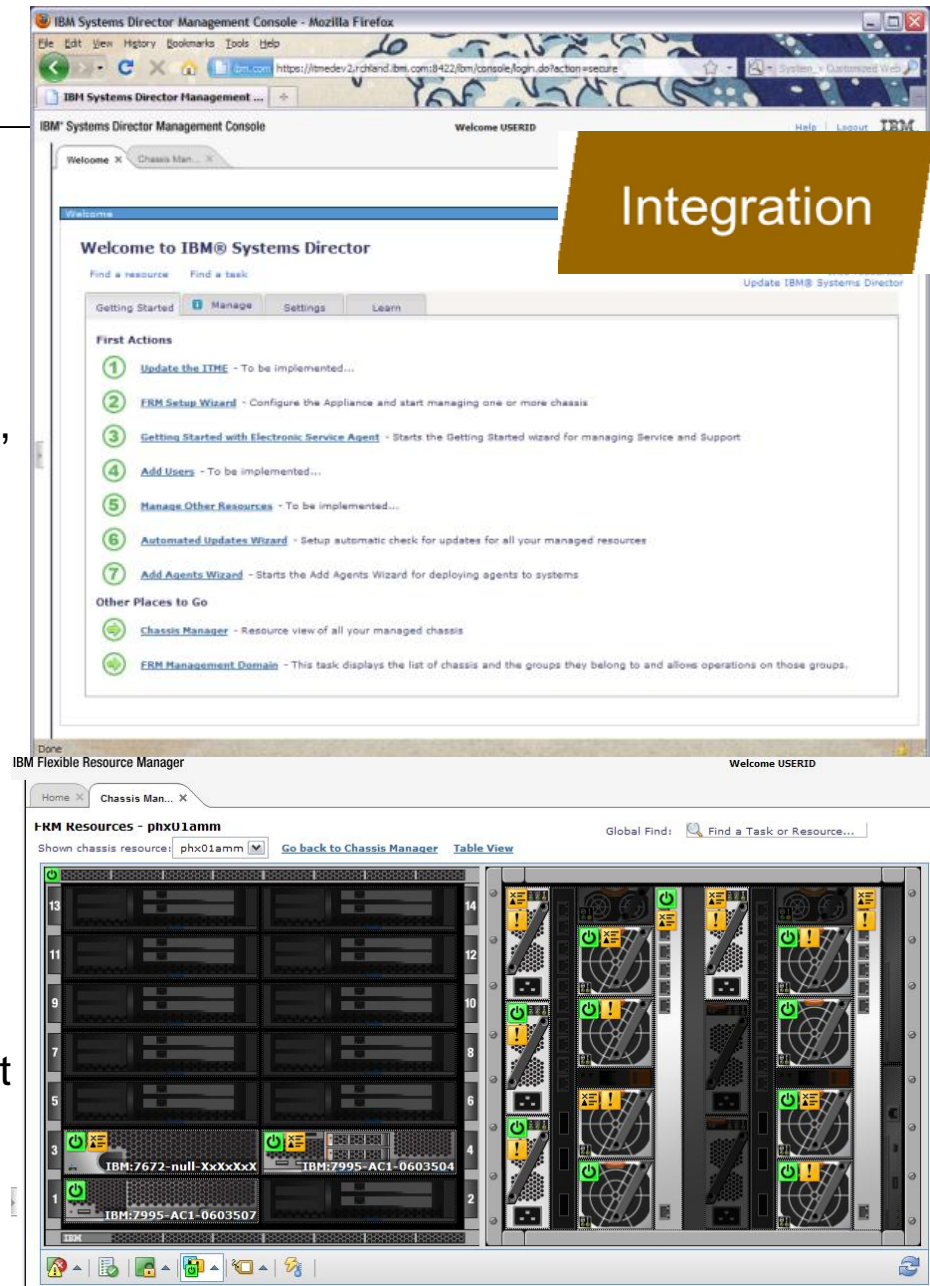
High Availability

Sold as stand alone product

* All software products licensed per chassis

Hardware Management

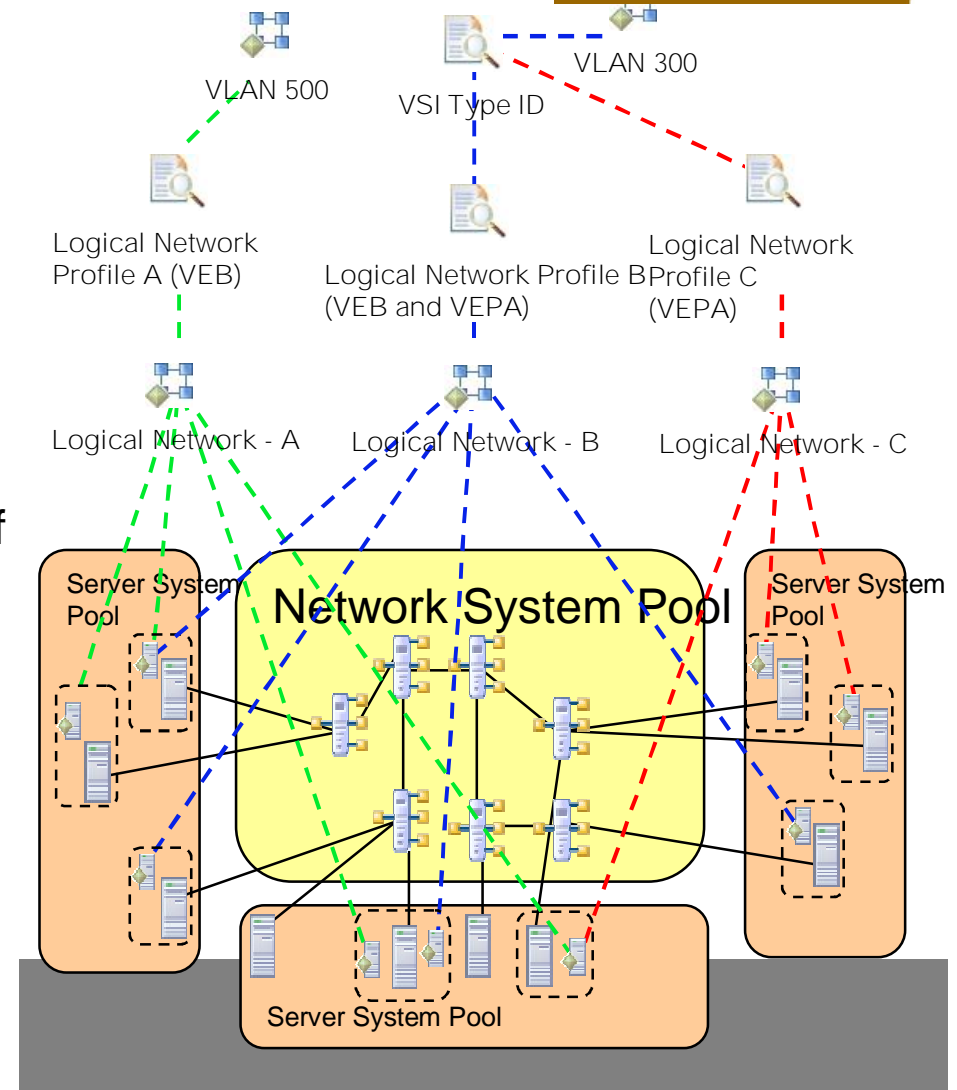
- § Simplified Setup: Start, manage, learn
- § Monitor Resources: Health Summary, Alerts, Thresholds, Updates, Service / Support
- § Visualize / Navigate Physical and Virtual Relationships: Intuitive drilldown / Views, Topology map, finger tip trouble shooting
- § Automate Responses: Custom actions / filters, configure, edit, relocate, automation plans
- § Manage firmware and Software updates: Set policies to track and automate firmware and software compliance.



Network management and network control

- § Provides management of network switches from a variety of vendors
- § Discovery, inventory, and status monitoring of switches
- § Graphical network topology views
- § Support for KVM, pHyp & VMware virtual switches as well as physical switches
- § Basic protocol and VLAN configuration of switches
- § Integration with server management
- § Provides per-VM network usage and performance statistics to VMControl
- § Logical views of servers and network devices grouped by subnet and VLAN
- § Provides the ability to automatically provision and move VLANs for VMs

Integration




System Networking Element Manager

Integration


A comprehensive tool to automate Virtualized
Data Center Network workflows



Engineer
§Physical
§Firmware
§Configuration
§Advanced PD




Operator
§Monitoring
§Initial PD
§Automation



Planner (or Engineer)
Investment = business direction + application
requirements + utilization trends

Today this is not officially integrated into the overall management solution!

Why an Integrated switch vs. a Pass-through




Why consider an Integrated switch vs. pass-through

- § Reduced Complexity – Fewer Cables
- § Simpler management – through automation capabilities with IBM FSM and advanced capabilities
- § Lower Latency – VM and server traffic can stay within the chassis
- § Better Security – less exposure by keeping traffic in the chassis
- § Higher Reliability – failover detection inside chassis and upstream
- § Lower TCO
 - CAPEX (fewer cables and upstream ports)
 - OPEX (power savings, less cabling = fewer points of failure)

IBM PureFlex System Networking



EN/CN4093 Easy-Connect




What is Easy-Connect


- The outcome is a clever pass-through. In this mode, the EN/CN4093 and if applicable the ToR G8264 are simple I/O module connecting servers, storage and the Core network
- Server ports are aggregated together towards the core network
- From the core network perspective, this is not a switch, there is no network integration. It is seen as one big pipe with server traffic coming in/out.
- No loops possible, no Spanning-Tree, no network integration
- This is very similar to Cisco FEX module with one important exception, server traffic within the same group can communicate with each other without having to go to the core.
- Specific Networking features (Virtual Fabric Mode, VMready etc...) can be activated on selected ports when/if needed.
- In a nutshell: We create groups inside the switch. We select internal server & external ports and assigned them to a group. Traffic within the same group can communicate with each other. External ports are connected to an external device and inter-group traffic will communicate via this external device. Servers can have overlapping VLAN id or IP address, as long as they are in different groups.

Easy-Connect

- Groups inter-connecting different ports
- Traffic “stays” within the group
- Complete isolations between groups
- Traffic is presented to the Core Network as server traffic. – No networking protocols
- Additional features can be activated on specific ports if desired.
- Multi-Chassis concept is also possible



Transparent Mode



Easy-Connect EN/CN4093



- Easy-Connect Transparent Mode

EN/CN4093 Easy-Connect Transparent Mode

- Allows EN/CN4093 to ACT as a Fabric Extension Module off a Cisco Network
- Looks like a “dumb” device to Nexus 5K/Edge Switch
- No Spanning Tree Protocol (STP) – eliminates Network admin loop concerns
- Provides traffic consolidation in the chassis to minimize ToR port utilization
- Provides intra-chassis switching, even in Transparent Mode

How to set the EN/CN4093 into Transparent Mode

- Set the EN/CN4093 uplinks as one logical unit via Static Port-Channel
- Turn on vNIC and configure all internet “ports” and external “port-channel” into vNIC “Group 1”
- Turn on Spanning Tree BPDU Guard and Edge on the Nexus 5K/Edge Switch
- If connecting to a Nexus 2k both BPDU Guard and Edge are already enabled by default and cannot be disabled.

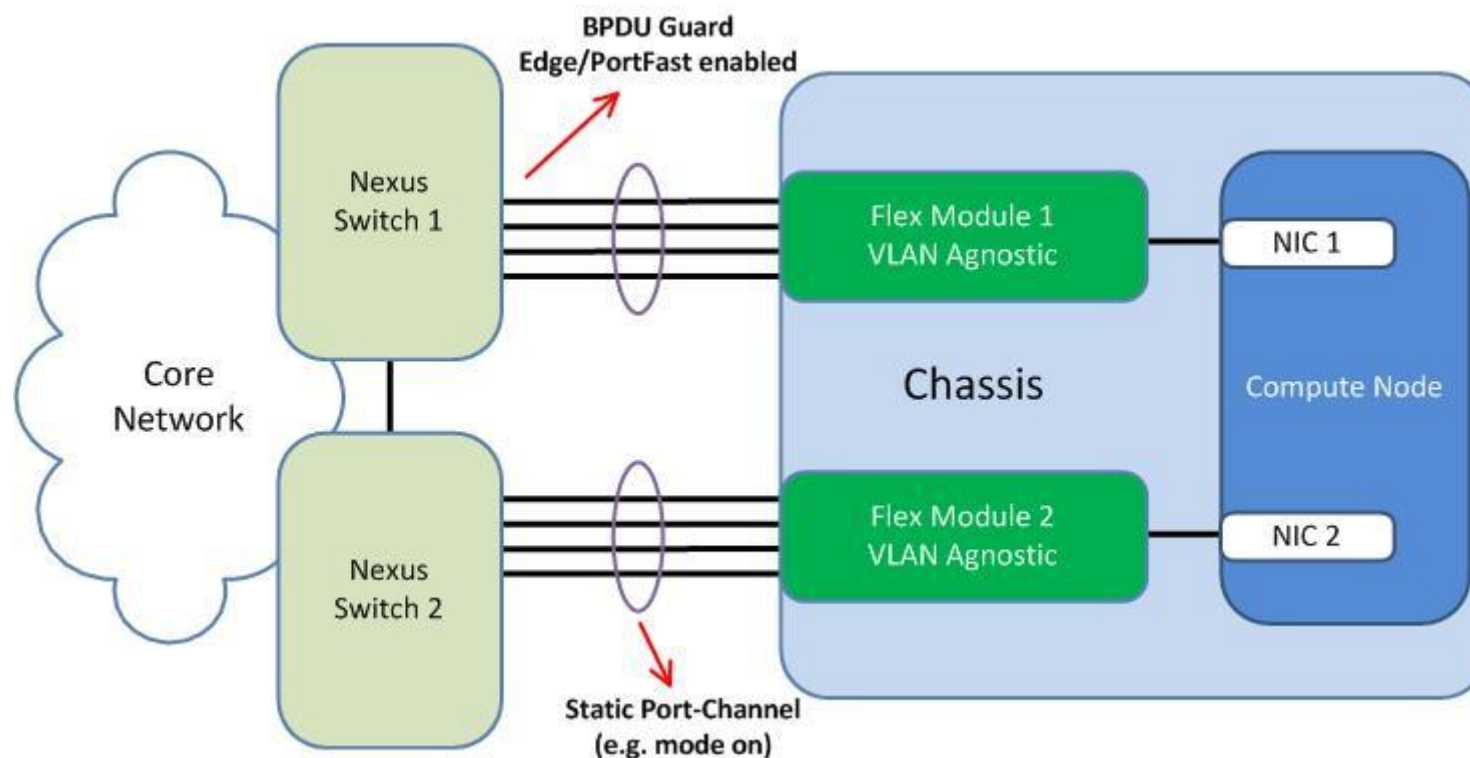
Client Benefits

- Allows client to manage networking from the Nexus management system
- Allows for intra-chassis switching for optimal performance (i.e. vMotion)
- Allows for pay as you grow expandability in the chassis, both to the network and the compute nodes.
- VLAN's are transparent to the EN/CN4093 when in “Transparent Mode” between the NIC and Nexus/Edge Switch
- Allows for pNIC or Switch Independent vNIC to be utilized
- Allows for growth into “Virtual Fabric Mode” using multiple groups



Easy-Connect EN/CN4093

What you end up with in Transparent Mode connecting to Nexus



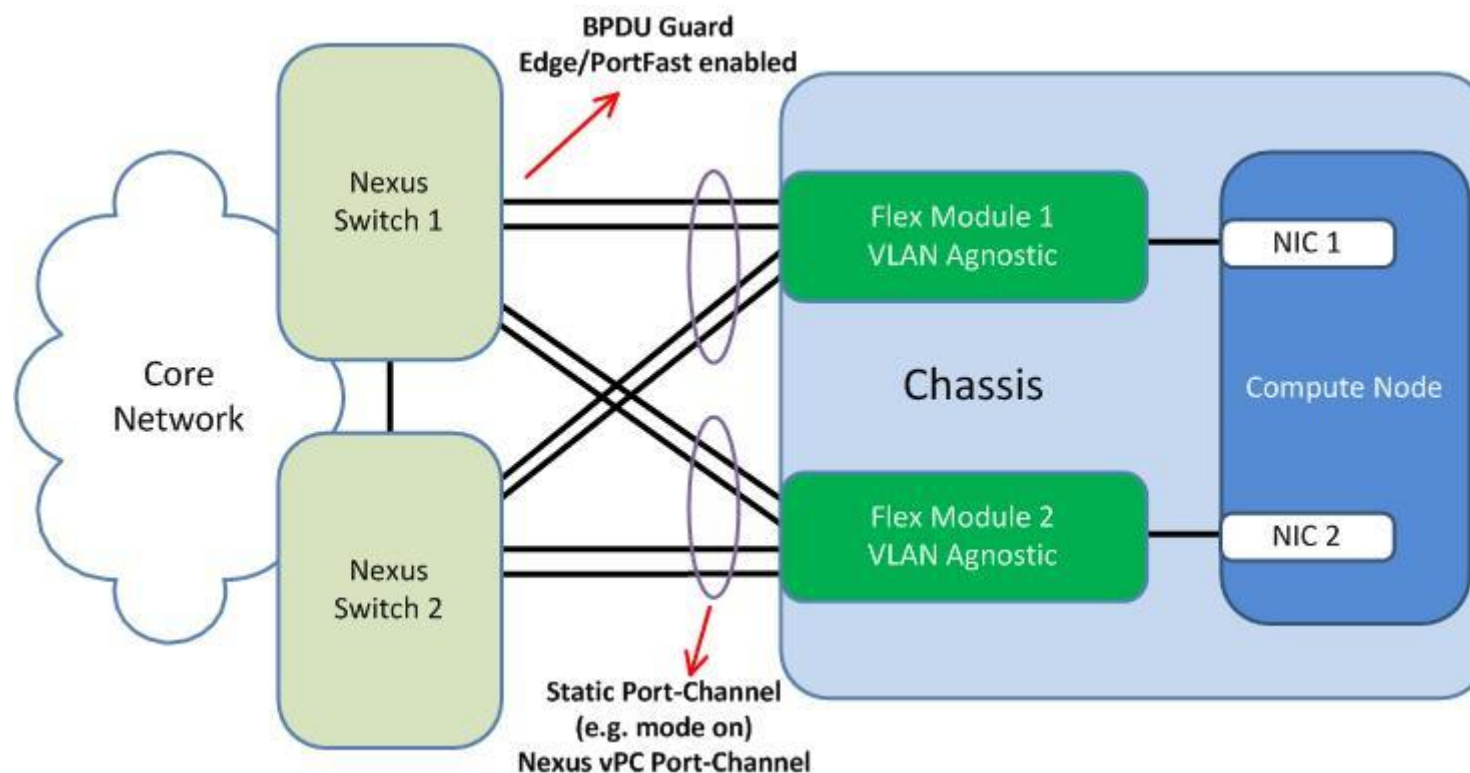
Ports EXT1-EXT10 in one Static Port-Channel
No Native VLAN / No Spanning-Tree enabled

All other ports disabled or non working



Easy-Connect EN/CN4093

What you end up with in Transparent Mode with Nexus vPC



Ports EXT1-EXT10 in one Static Port-Channel
No Native VLAN / No Spanning-Tree enabled

All other ports disabled or non working



Easy-Connect EN/CN4093



Configuration to convert from Switch Mode to Transparent Mode

Configuration for the EN/CN4093 from A to Z

- ssh to each EN/CN4093
- boot the EN/CN4093 to use isCLI (Cisco-Like CLI) for a more familiar look
 - /boot/mode iscli
 - /boot/reset
- Set the following commands to convert the EN/CN4093 into “Transparent Mode”
 - ssh to each EN/CN4093
 - enter into enable mode (e.g. enable)
 - enter into configuration terminal mode (e.g. config t)
 - type the following;

=====Begin Script =====

```
spanning-tree mode disable
portchannel 1 port ext1-ext10 enable
vnic enable
vnic vnicgroup 1
vlan 4091
port inta1-inta14
trunk 1
enable
failover
exit
write memory
```

=====End Script =====



Performance – Advantage of Layer 2/3 - Integrated


Switching (L2) allows servers in same VLANs (Virtual LANs) to communicate within the chassis

Routing (L3) allows various back-end services, on different VLANs, to “talk” within the chassis

Routing Advantages


- Reduces latency (response times)
- Passes data traffic faster and makes decisions
- Conserves bandwidth to Core or Internet – more bandwidth for other traffic
- Secures data inside the chassis
- Saves firewall CPU

SSL Security





Competing Against HP and Dell Offerings

Thank You



Thank You!





Legal Disclaimer

- © IBM Corporation 2011. All Rights Reserved.
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- If the text contains performance statistics or references to benchmarks, insert the following language; otherwise delete:
Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.
- If the text includes any customer examples, please confirm we have prior written approval from such customer and insert the following language; otherwise delete:
All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.
- Please review text for proper trademark attribution of IBM products. At first use, each product name must be the full name and include appropriate trademark symbols (e.g., IBM Lotus® Sametime® Unyte™). Subsequent references can drop "IBM" but should include the proper branding (e.g., Lotus Sametime Gateway, or WebSphere Application Server). Please refer to <http://www.ibm.com/legal/copytrade.shtml> for guidance on which trademarks require the ® or ™ symbol. Do not use abbreviations for IBM product names in your presentation. All product names must be used as adjectives rather than nouns. Please list all of the trademarks that you use in your presentation as follows; delete any not included in your presentation. IBM, the IBM logo, Lotus, Lotus Notes, Notes, Domino, Quickr, Sametime, WebSphere, UC2, PartnerWorld and Lotusphere are trademarks of International Business Machines Corporation in the United States, other countries, or both. Unyte is a trademark of WebDialogs, Inc., in the United States, other countries, or both.
- If you reference Adobe® in the text, please mark the first use and include the following; otherwise delete:
Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- If you reference Java™ in the text, please mark the first use and include the following; otherwise delete:
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- If you reference Microsoft® and/or Windows® in the text, please mark the first use and include the following, as applicable; otherwise delete:
Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.
- If you reference Intel® and/or any of the following Intel products in the text, please mark the first use and include those that you use as follows; otherwise delete:
Intel, Intel Centrino, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- If you reference UNIX® in the text, please mark the first use and include the following; otherwise delete:
UNIX is a registered trademark of The Open Group in the United States and other countries.
- If you reference Linux® in your presentation, please mark the first use and include the following; otherwise delete:
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.
- If the text/graphics include screenshots, no actual IBM employee names may be used (even your own), if your screenshots include fictitious company names (e.g., Renovations, Zeta Bank, Acme) please update and insert the following; otherwise delete: All references to [insert fictitious company name] refer to a fictitious company and are used for illustration purposes only.