Inteligent data center/next generation data center

Internet Users' Conference - CUC 2005
Dubrovnik, November 21.-23., 2005.

Wednesday, 23.11.2005. 11:00-13:00

Josip Zimet

Cisco Systems

List of architectures …

2000/2001
AVVID
- sla

2001/2002
SAFE
- security

2002/2003
V3PN
-sla
-security

2004/2005
SDN/SWAN
-Integrated security
- collaborative security systems
-adaptiive threat defense

2005/2006
IIN: AON/DNA/SONA
-standardization
- virtualization

Architectures …

2000/2001
AVVID
- sla

2001/2002
SAFE
- security

2002/2003
V3PN
-sla
-security

2004/2005
SDN
-Integrated security
- collaborative security systems
-adaptiive threat defense
-

2005/2006
SONA
-standardization
- virtualization

Next generation switches :

1900→ 2900XL → 2950/3550 → 500/2960/3560

Next Generation routers :

2500 → 2600 → 2600XM → 2800 ISR

Next generation PIX :

PIX 5xx → PIX 5xxE → ASA 55xx

Next generation IDS …., Next generation Wireless …

List of acquisitions …

1	Crescendo Communications, Inc.	2	Newport Systems Solutions	3	Kalpana, Inc.	4	LightStream Corporation	5	Combinet, Inc.	6	Internet Junction Networks, Inc.	7	Grand Junction, Inc.	8	Network Translation, Inc.	9	TGV Software, Inc.	10	Stratacom, Inc.
	1993-SEP-21		1994-JUL-12		1994-OCT-24		1994-DEC-08		1995-AUG-10		1995-SEP-06		1995-SEP-27		1995-OCT-27		1996-JAN-23		1996-APR-22
11	Telebit's MICA Technologies		Nashoba Networks, Inc.	13	Granite Systems, Inc.	14	Netsys Technologies, Inc.	15	Metaplex, Inc.	16	Telesend	17	SkyStone Systems Corp	18	Global Internet Software Group	19	Ardent Communications Corp.	20	Dagaz (Integrated Network Corporation)
	1996-JUL-22		1996-AUG-06		1996-SEP-03		1996-OCT-14		1996-DEC		1997-MAR-26		1997-JUN-09		1997-JUN-24		1997-JUN-24		1997-JUL-28
21	LightSpeed International, Inc.	22	WheelGroup Corporation	23	NetSpeed, Inc.	24	Precept Software, Inc.	25	CLASS Data Systems	26	Summa Four, Inc.	27	American Internet Corporation	28	Clarity Wireless Corporation	29	Selsius Systems, Inc.	30	PipeLinks, Inc.
	1997-DEC-22		1998-FEB-18		1998-MAR-10		1998-MAR-11		1998-MAY-04		1998-JUL-28		1998-AUG-21		1998-SEP-15		1998-OCT-14		1998-DEC-02
31	Fibex Systems	32	Sentient Networks, Inc.	33	GeoTel Communications Corp.	34						37	Calista, Inc.	38	MaxComm Technologies, Inc.	39	Monterey Networks, Inc.	40	Cerent Corporation
	1999-APR-08		1999-APR-08		1999-APR-13								1999-AUG-16		1999-AUG-18		1999-AUG-26		1999-AUG-26
41	Cocom A/S	42	WebLine Communications Corp.	43	Tasmania Network Systems, Inc.	44							Internet Engineering Group, LLC		Pirelli Optical Systems S.p.A	49	Compatible Systems Corp.		Altiga Networks
	1999-SEP-15		1999-SEP-22		1999-OCT-26								1999-DEC-17		1999-DEC-20		2000-JAN-19		2000-JAN-19
51	Growth Networks, Inc.	52	Atlantech Technologies Ltd.	53	JetCell, Inc.	54						57	Seagull Semiconductor Ltd.	58	ArrowPoint Communications, Inc.	59	Geyton Systems AB	60	HyNEX, Ltd.
	2000-FEB-16		2000-MAR-01		2000-MAR-16								2000-APR-12		2000-MAY-05		2000-MAY-12		2000-JUN-05
61	Netiverse, Inc.	62	Komodo Technology, Inc.	NuSpeed	NuSpeed Internet Systems, Inc.	64	IPmobile, Inc.	65	PixStream, Inc.	66	Technologies, Inc.	67	Vovida Networks, Inc.	68	CAIS Software Solutions	69	Active Voice Corporation		Radiata, Inc.
	2000-JUL-07		2000-JUL-25		2000-JUL-26		2000-AUG-01		2000-AUG-31		2000-SEP-28		2000-SEP-28		2000-OCT-20		2000-NOV-10		2000-NOV-13
81	ExiO Communications, Inc.	82	AuroraNetics, Inc.	83	Allegro Systems, Inc.		Hammerhead Networks, Inc.		Navarro Networks, Inc.	76	AYR Networks, Inc.	77	Andiamo Systems, Inc.	78	Psionic Software, Inc.		OKENA, Inc.	80	SignalWorks, Inc.
	2000-DEC-14		2001-JUL-11		2001-JUL-27		2002-MAY-01		2002-MAY-01		2002-JUL-25		2002-AUG-20		2002-OCT-22		2003-JAN-24		2003-MAR-19
81	Linksys Group, Inc	82	Latitude Communications	83	Twingo Systems, Inc		Riverhead Networks, Inc		Procket Networks, Inc		Actona Technologies		Parc Technologies		P-Cube, Inc.		NetSolve, Inc.	90	dynamicsoft, Inc.
	2003-MAR-20		2003-NOV-12		2004-MAR-12		2004-MAR-22		2004-JUN-17		2004-JUN-29		2004-JUL-08		2004-AUG-23		2004-SEP-09		2004-SEP-13
91	Perfigo, Inc.	92	Jahi Networks	93	BCN Systems, Inc	94	Protego Networks, Inc	95	Airespace, Inc	96	Topspin	97	Sipura	98	FineGround Networks, Inc.	99	M.I. Secure, Inc.	100	Netsift, Inc.
	2004-OCT-21		2004-NOV-17		2004-DEC-09		2004-DEC-21		2005-JAN-12		2005-APR-14		2005-APR-26		2005-MAY-26		2005-JUN-14		2005-JULY-22
	KiSS Technology		Sheer Networks		Nemo Systems														
	2005-JULY-22		2005-JULY-26		2005-SEP-30														

Acquisitions related to the DNA …

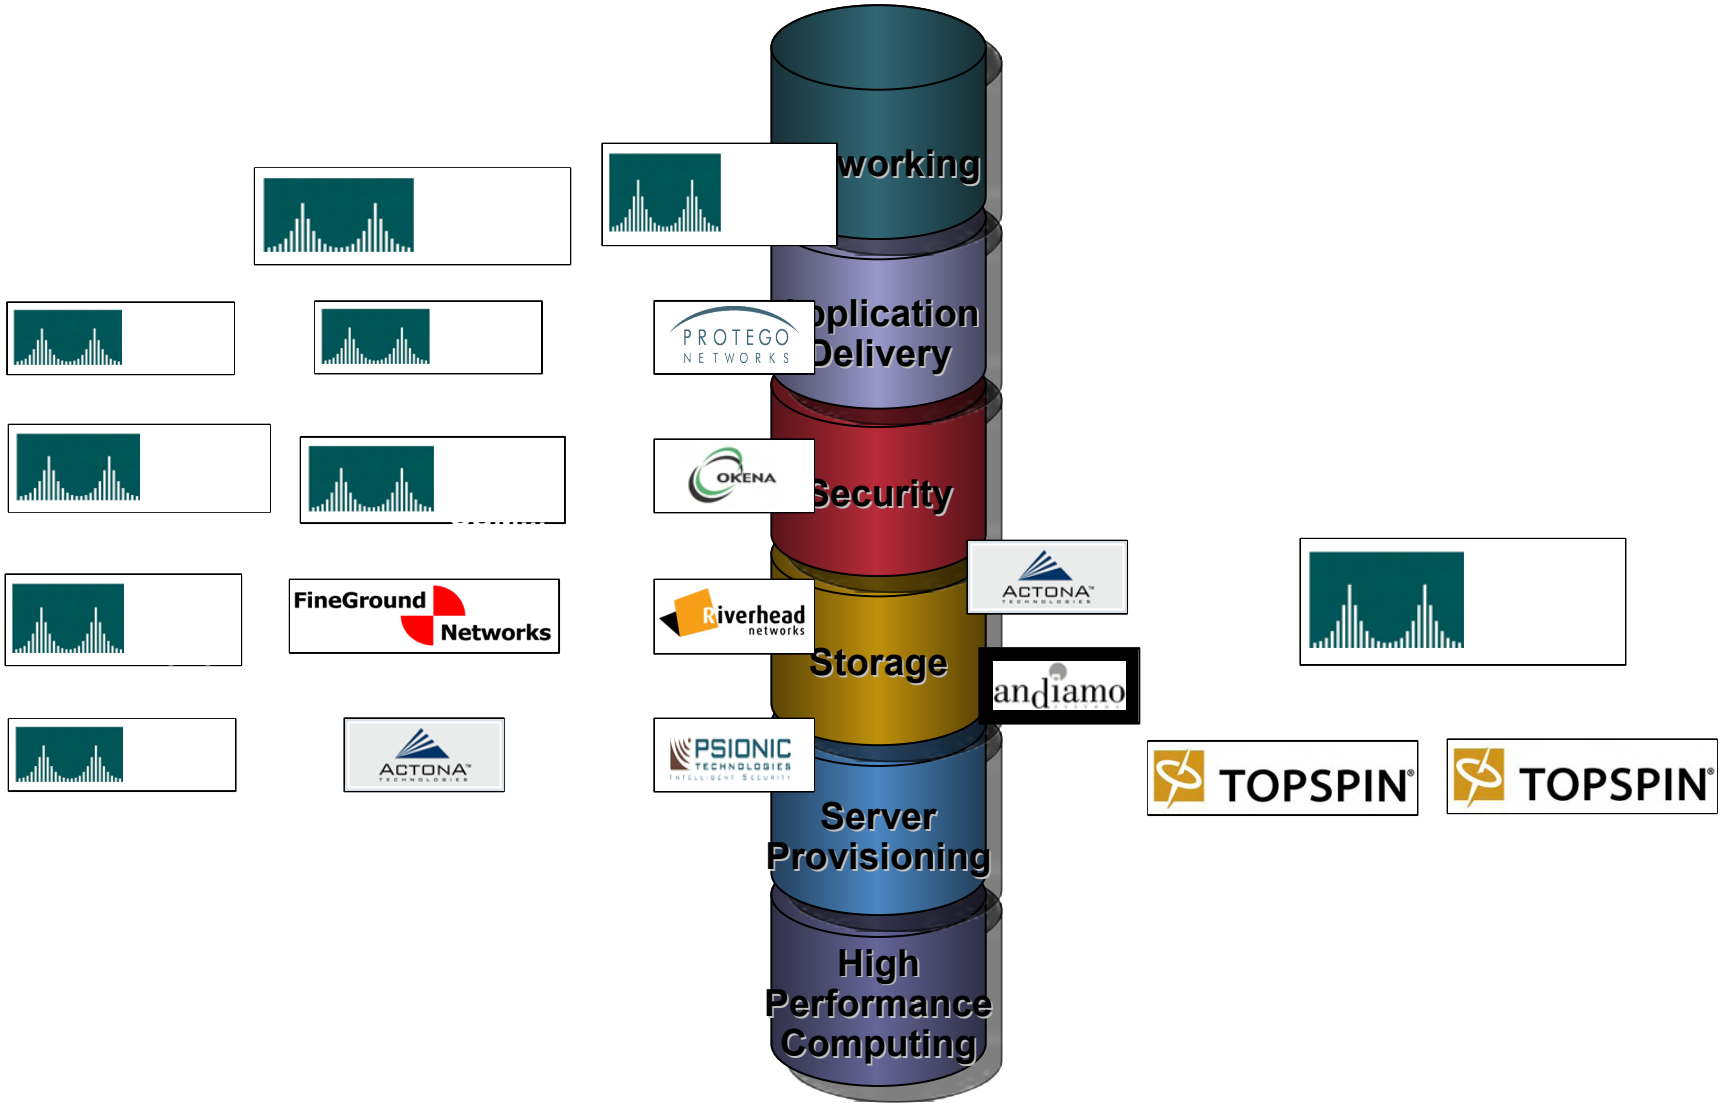

DNA

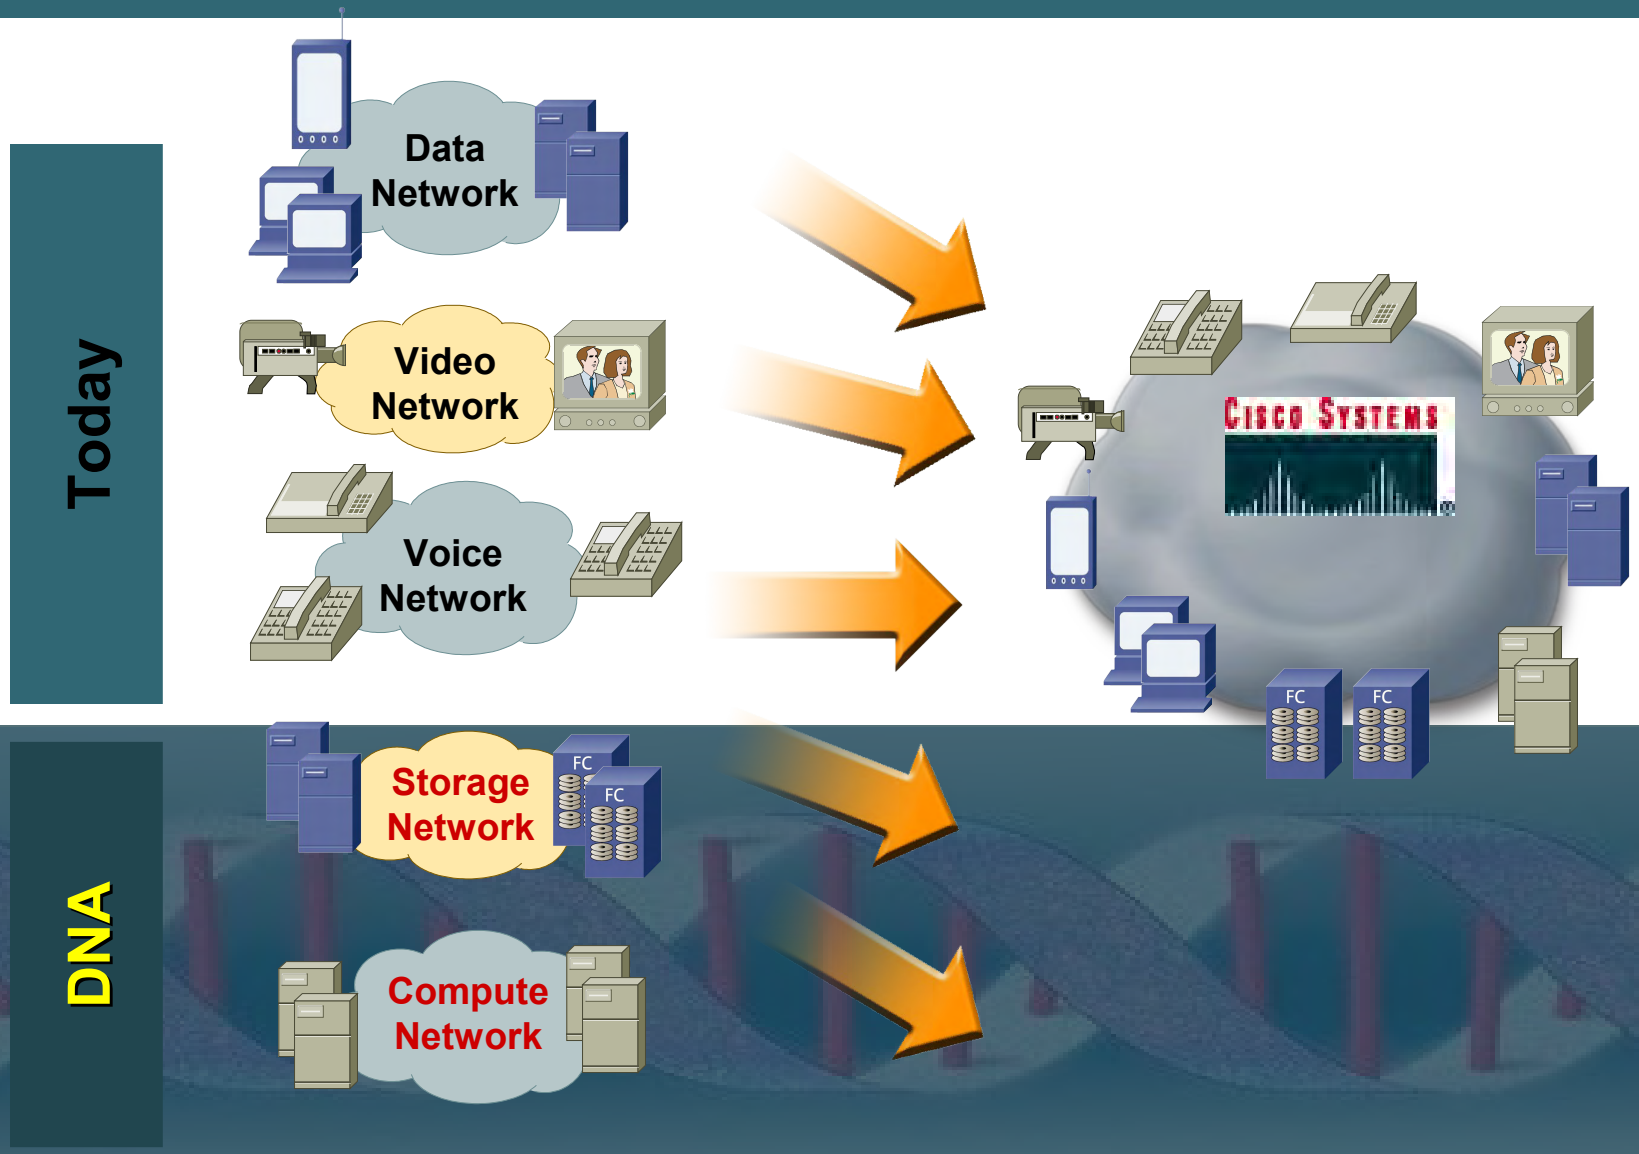

Data Center DNA Products

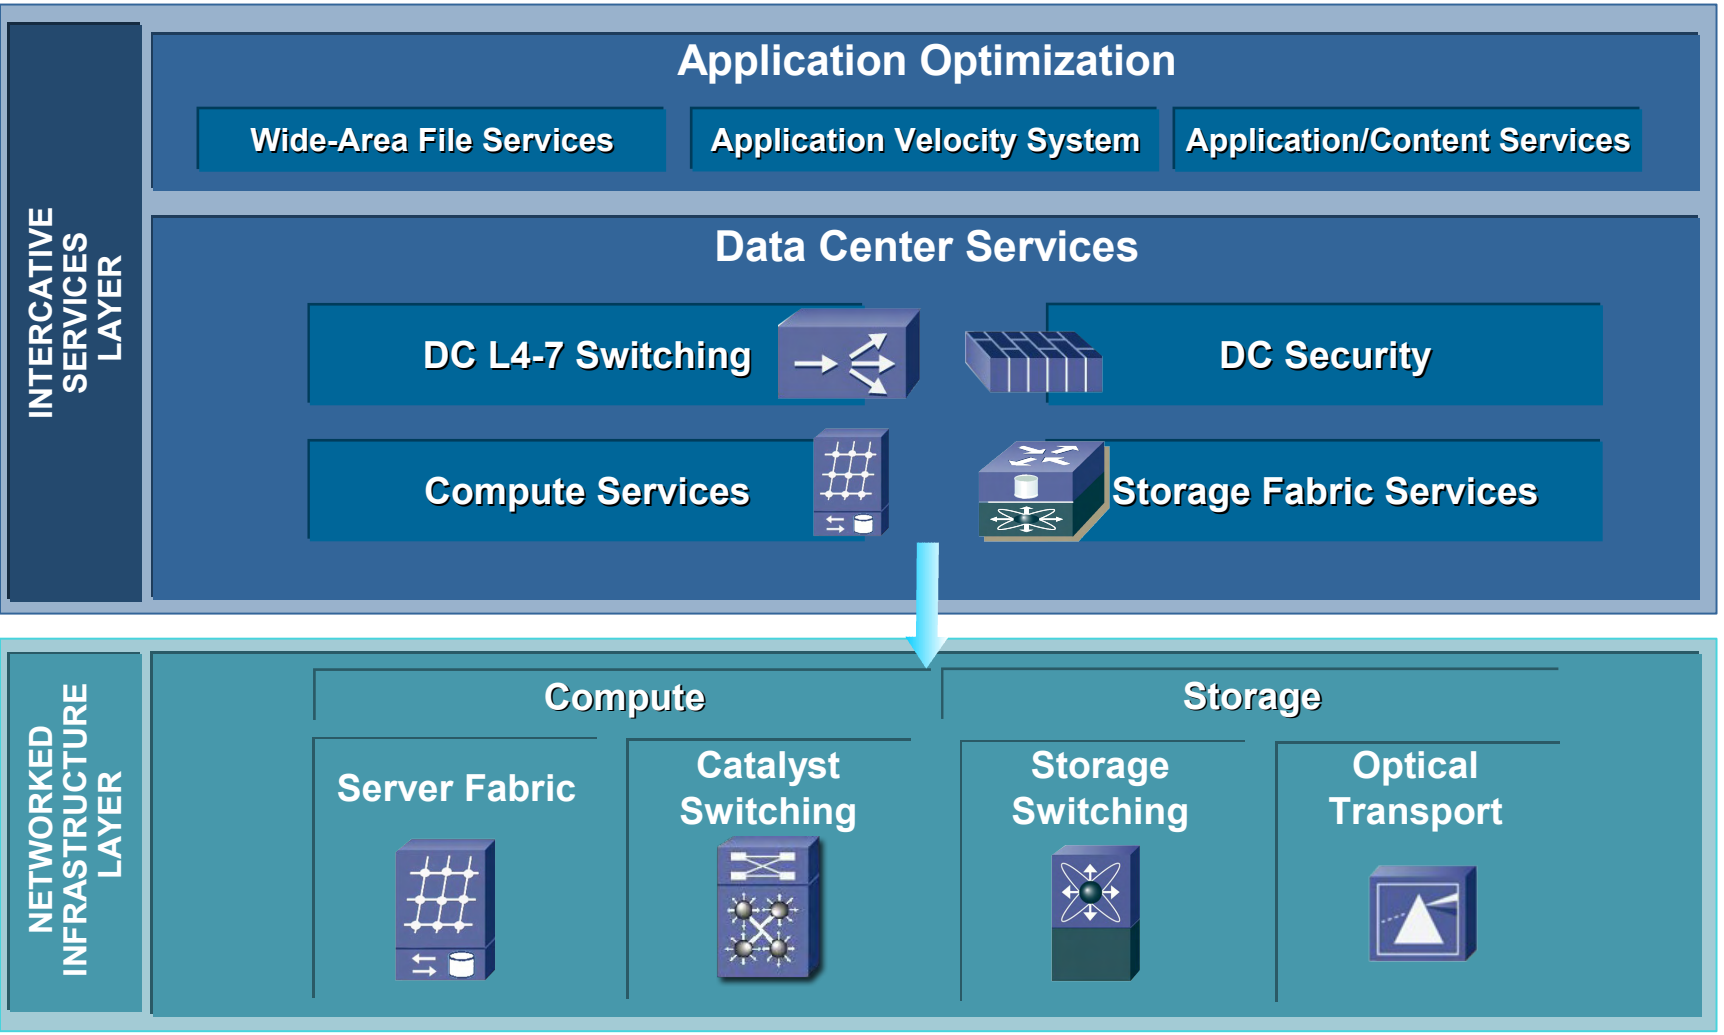

Cisco DNA Infrastructure Evolution/Roadmap

Cat6000 Family

Fabric Enabled Integrated Svcs

Sup720

Continued 6K Evolution

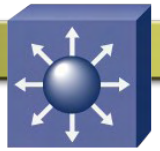

Classic Family
- 32Gb/System
- FE Server Aggregation
- L2/L3 Services

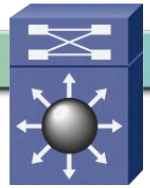

CEF256 Family
- 256Gb/System
- Distributed Forwarding
- L4- L7 Services Integration
- Security Services

CEF720 Family
- 720Gb/System
- Accelerated IP Svcs
- 10/100/1000 and 10GbE high density
- App Compute Engines

CEF1440 Family
- 1.4TB/System
- 8th Gen Fwd Engine
- Dense 10GbE

Storage and Compute Family

MDS9500 Family
- 1TB+ System
- Datacenter Storage
- Virtual SANs
- Highly Available
- Modular Scalable OS

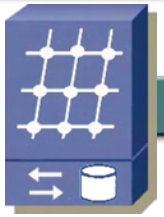

SFS Family
- Compute System
- Low Latency Switching
- High Density IB
- Gateway to FC,GbE
- Virtualization

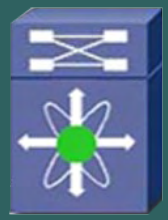

Next Generation Future
- 5Tb+/System
- Datacenter Convergence
- Virtualization
- I/O Consolidation

1999 2004 2007+

The Big Picture - The Cisco Data Center

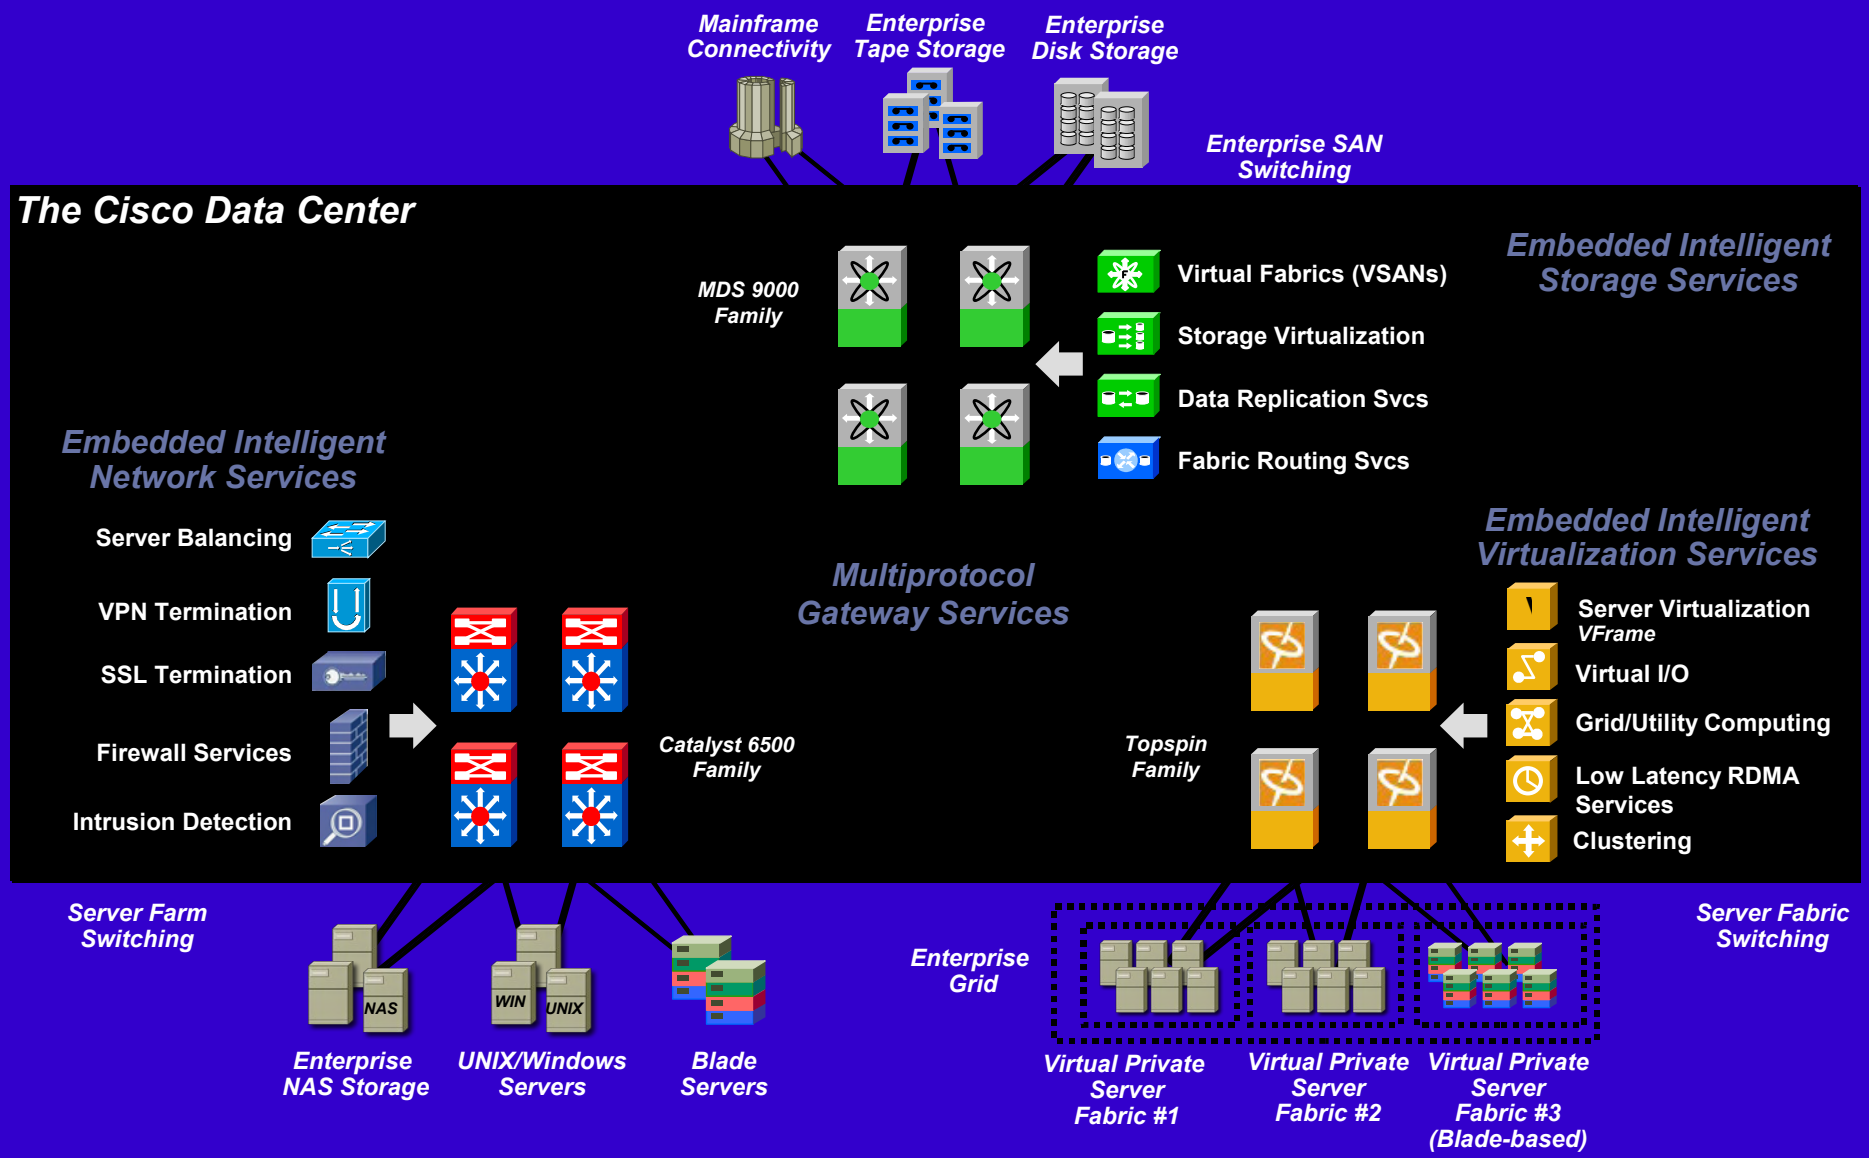

Business Ready Data Center Architecture to Topology

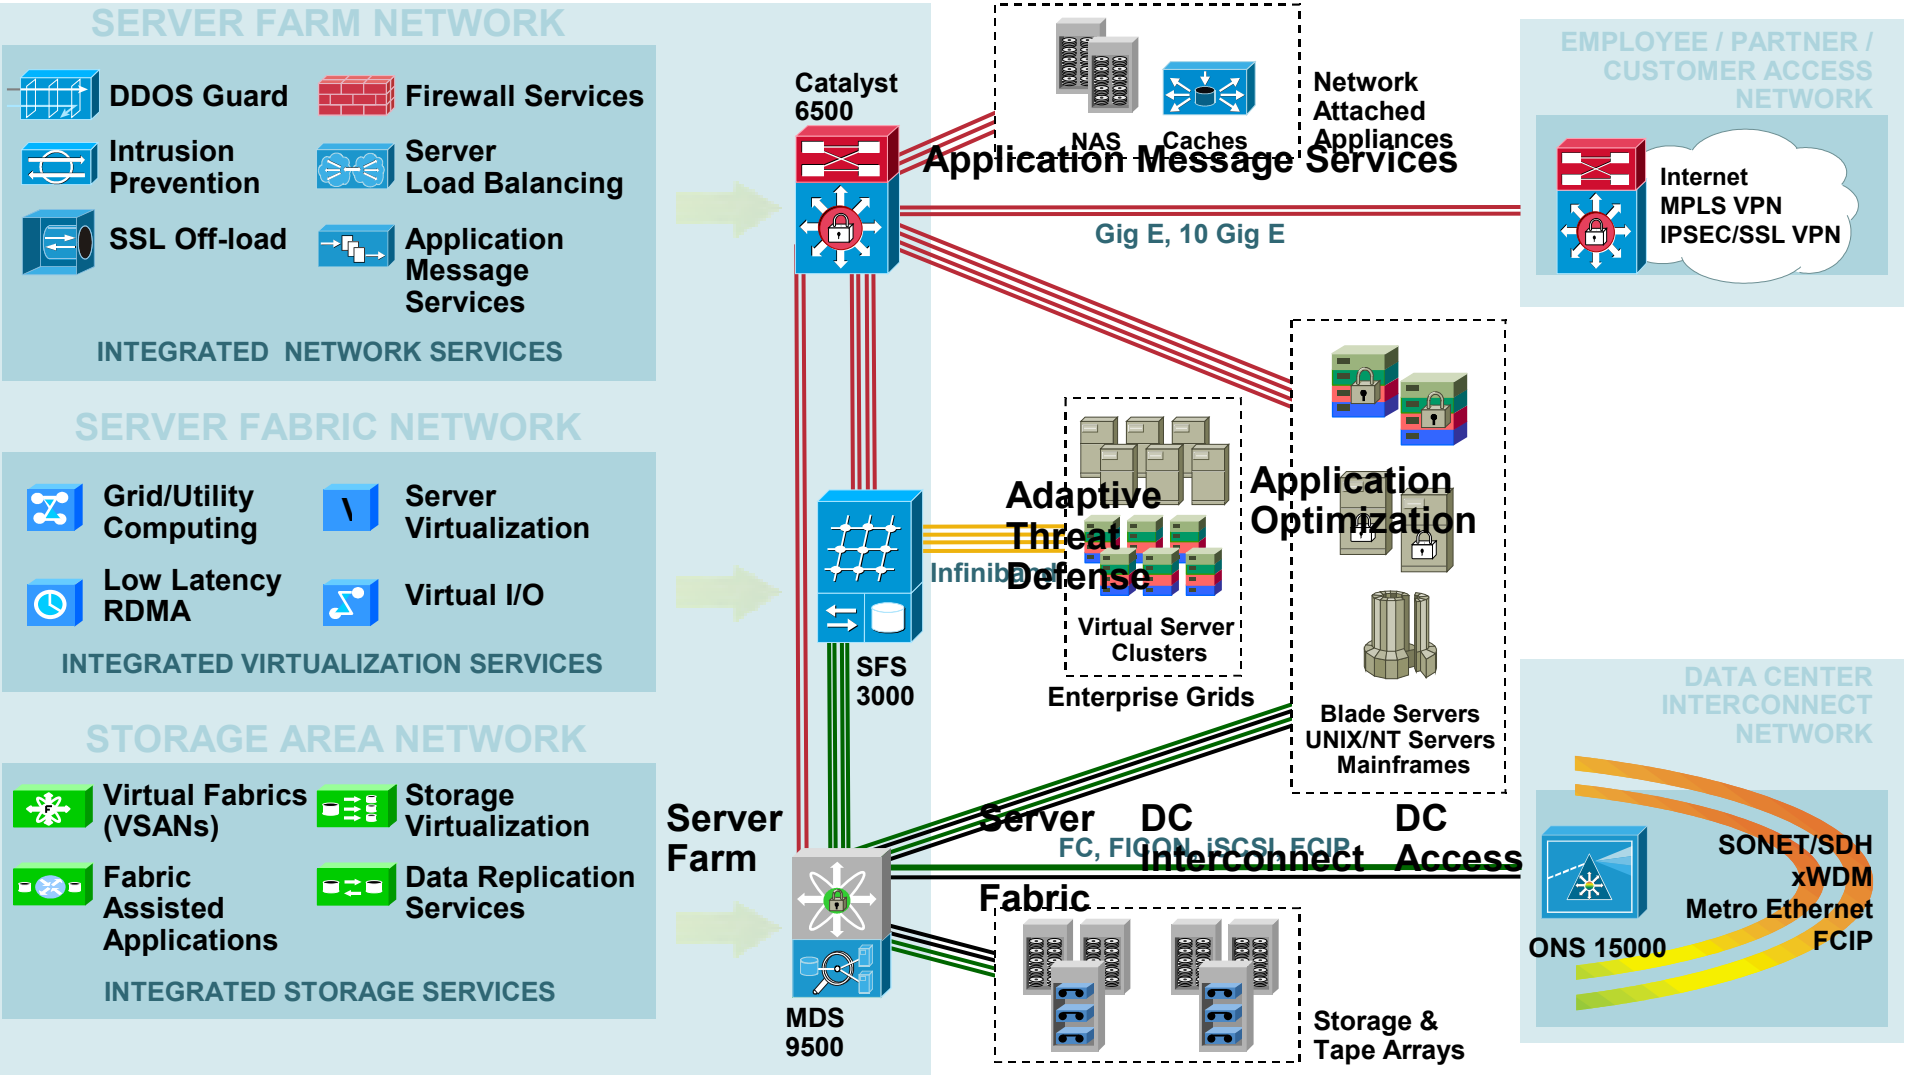

Enterprise Data Center Network Topology

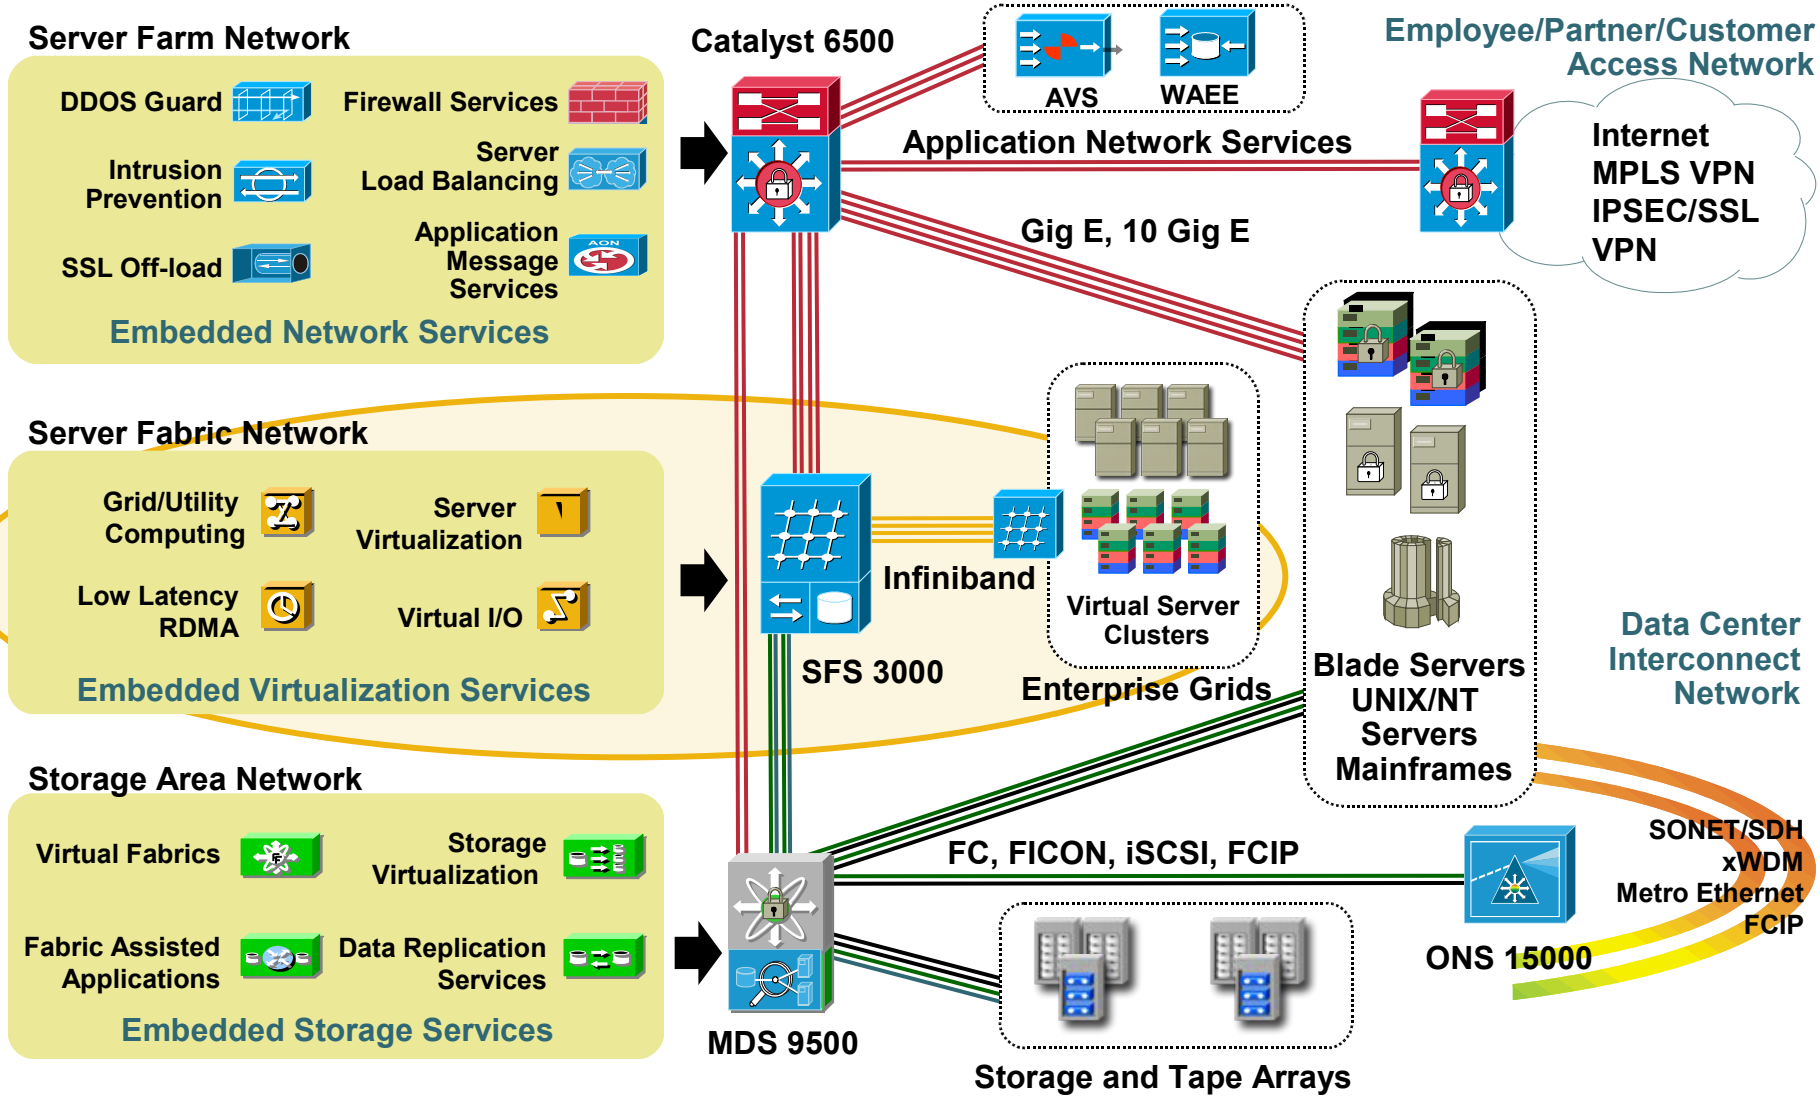

The Evolution of the Data Center

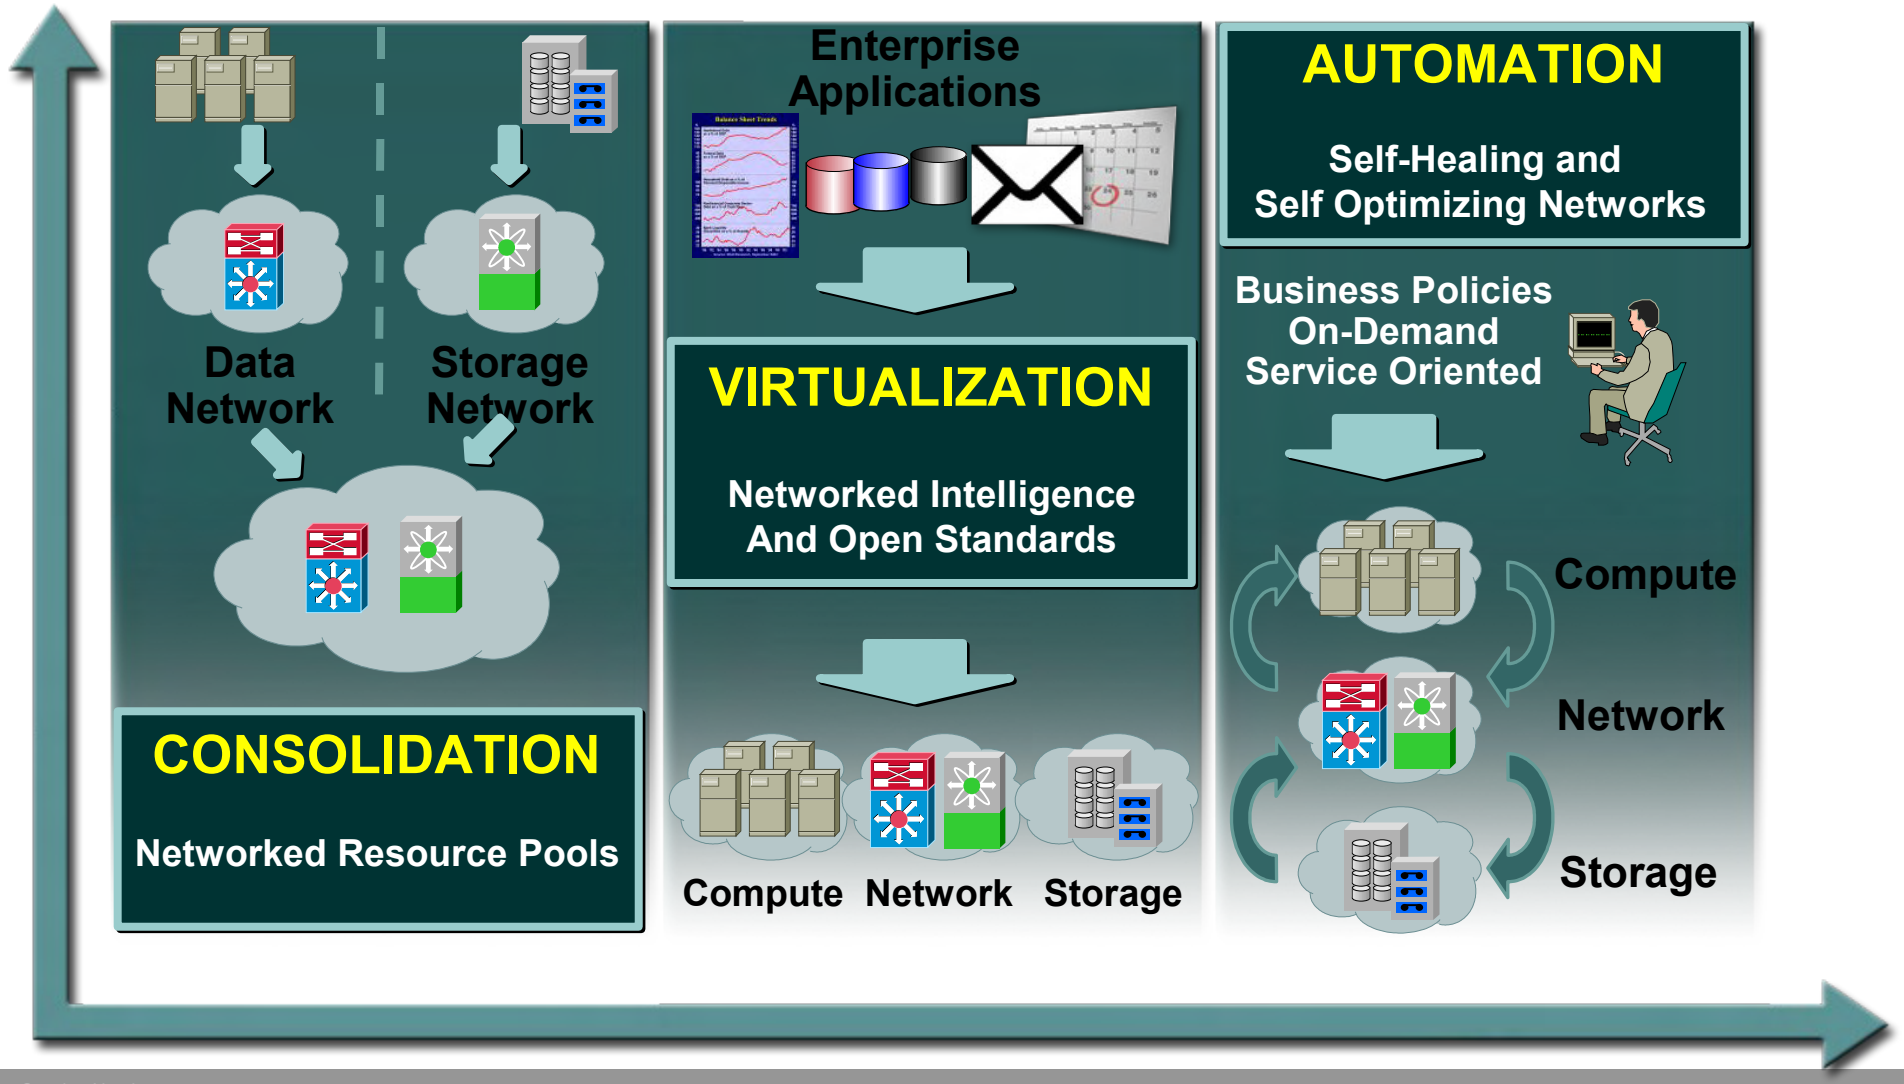

Server Switching Architectural Evolution

Server/Compute Consolidation → **Server and I/O Virtualization** → **Automated Application Services**

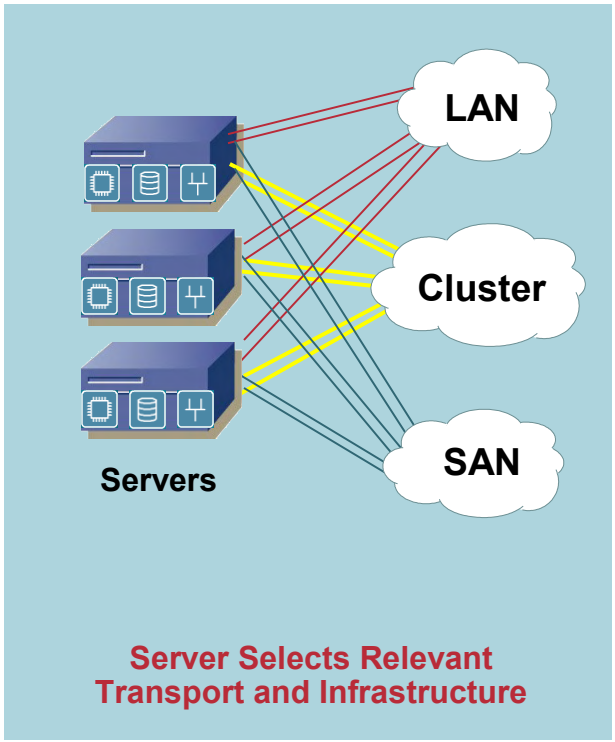

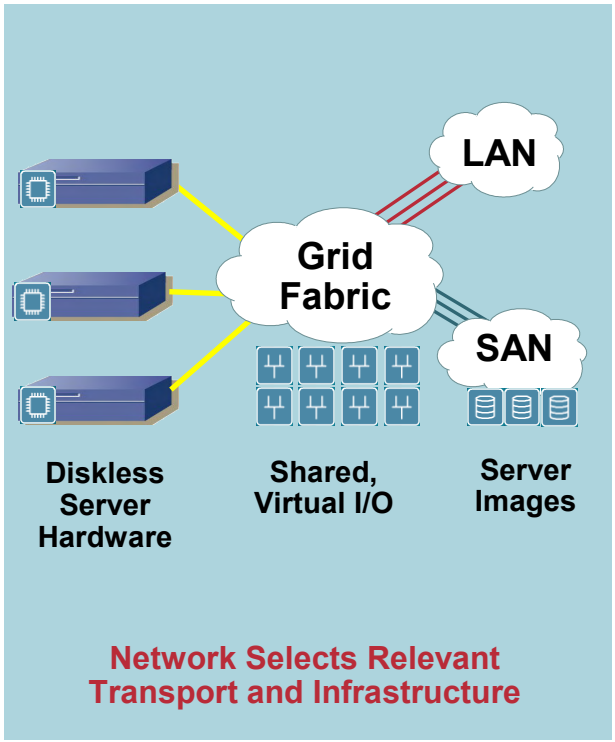

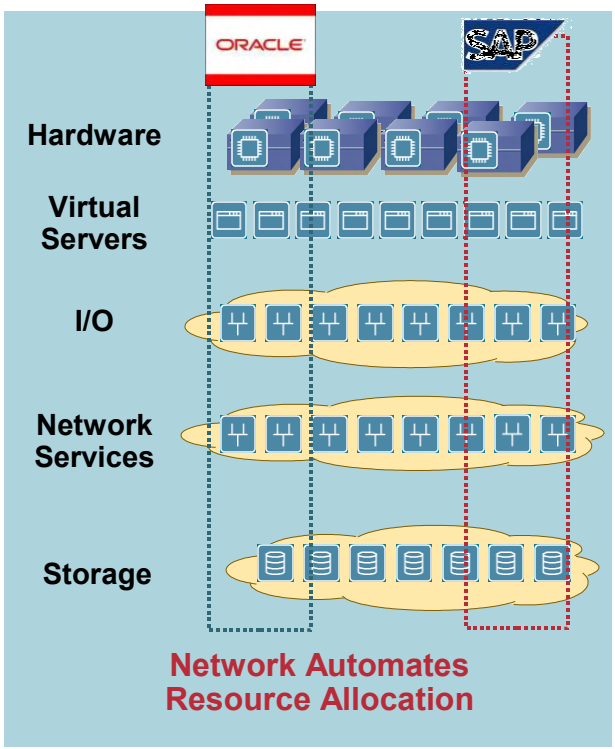

Today: Complex servers with parallel networks for LAN, SAN, and Clustering

Next: Unified fabric enables shared pools of stateless commodity servers

Future: Virtual Resource Networks dynamically map apps to resources

The Evolution of the Data Center

SMP

Pro: Single managed entity, fast backplane

Con: Expensive, Proprietary server + backplane

Dis-aggregation

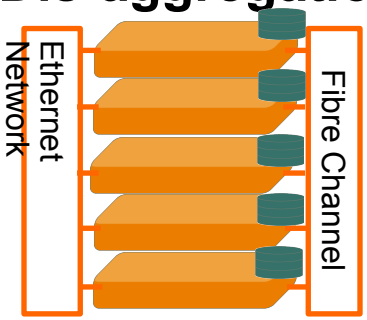

Pro: Standard servers, inexpensive

Con: Lots of managed components, low-performing interconnect

Virtualization

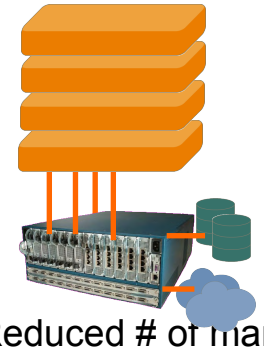

Pro: Reduced # of managed components, virtual I/O, *fast* standards backplane

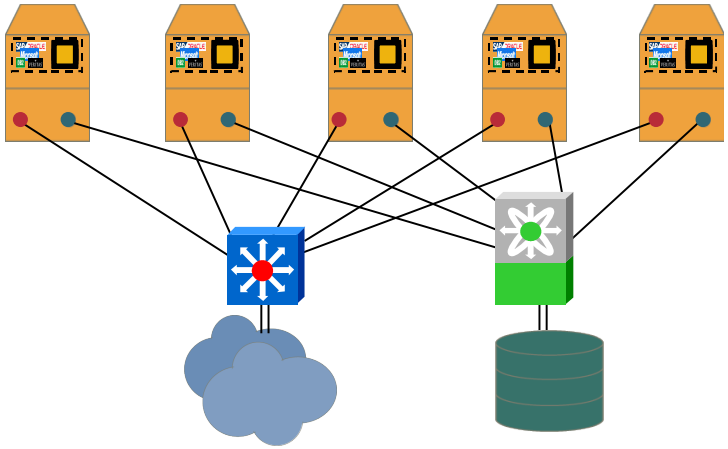

Network and Storage Virtualization

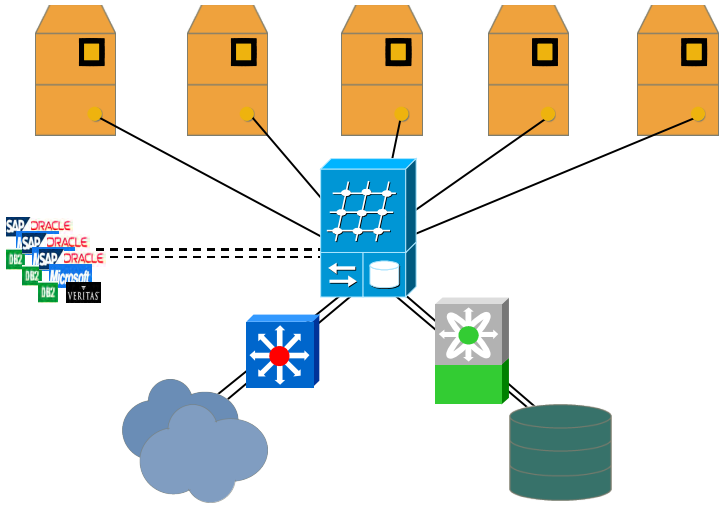

Server Virtualization - The Server Switch

The Evolution of the Data Center

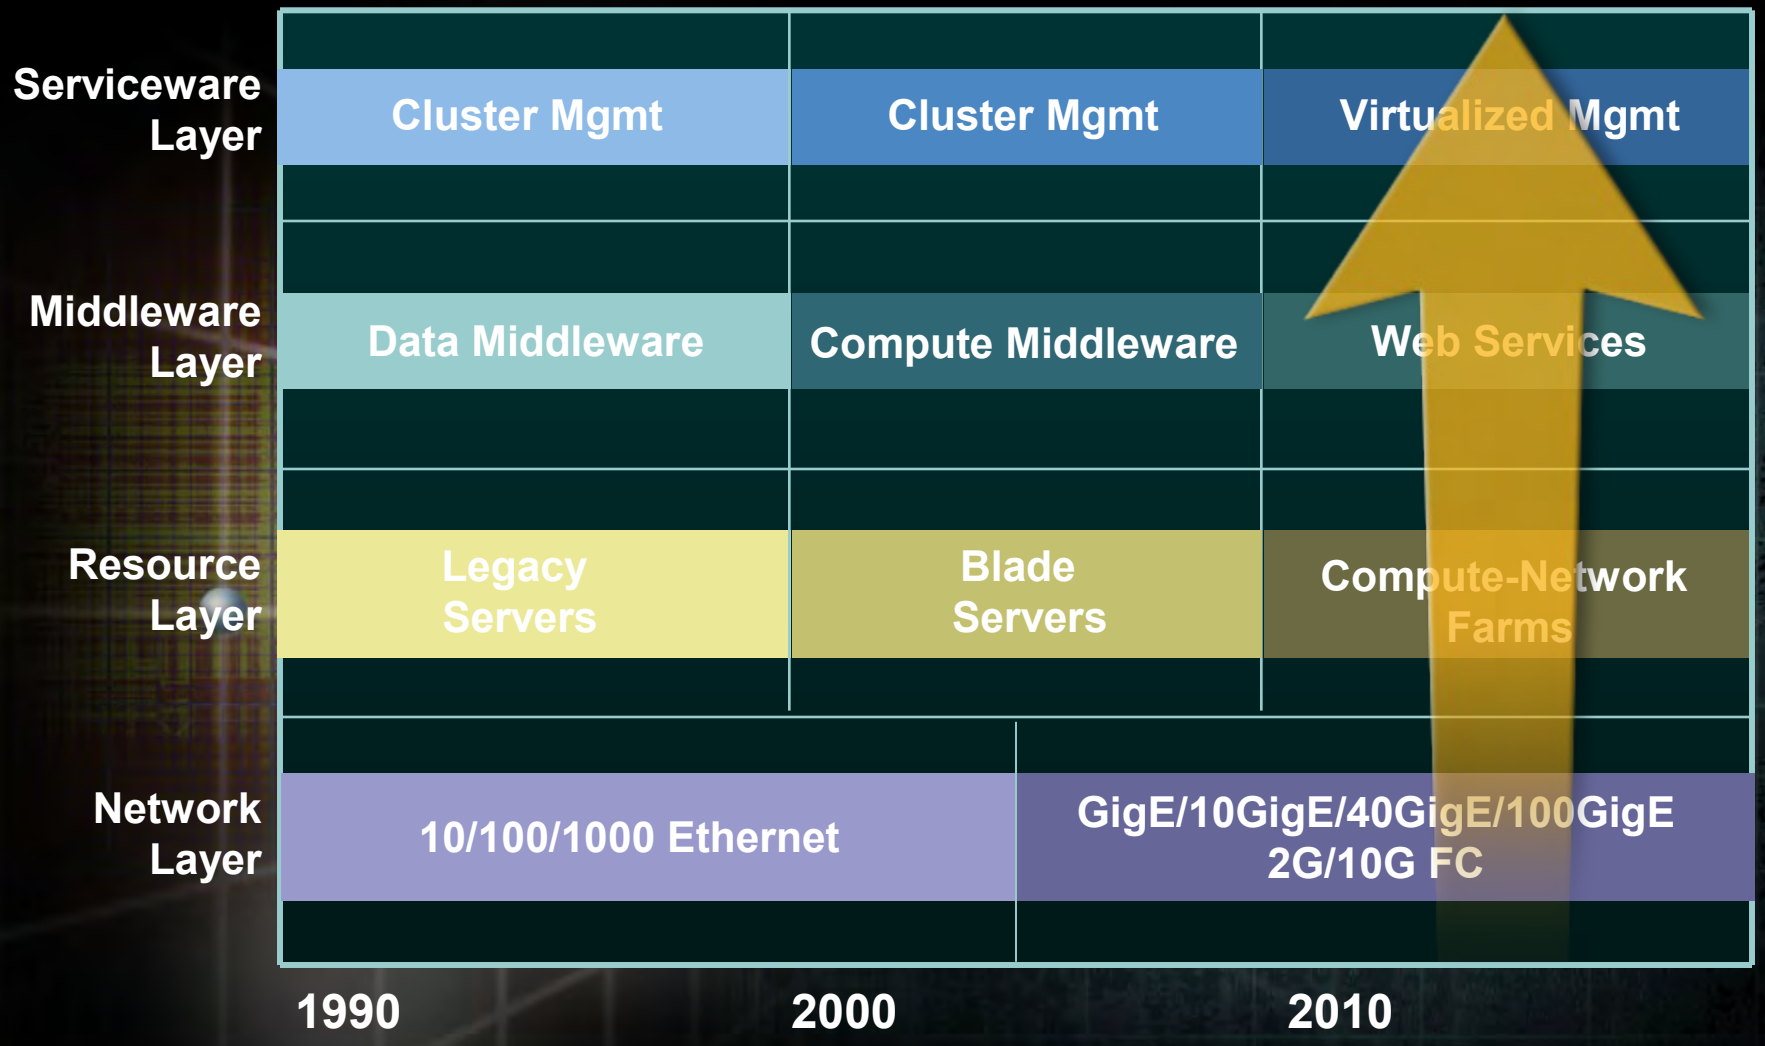

Cisco DNA Architecture

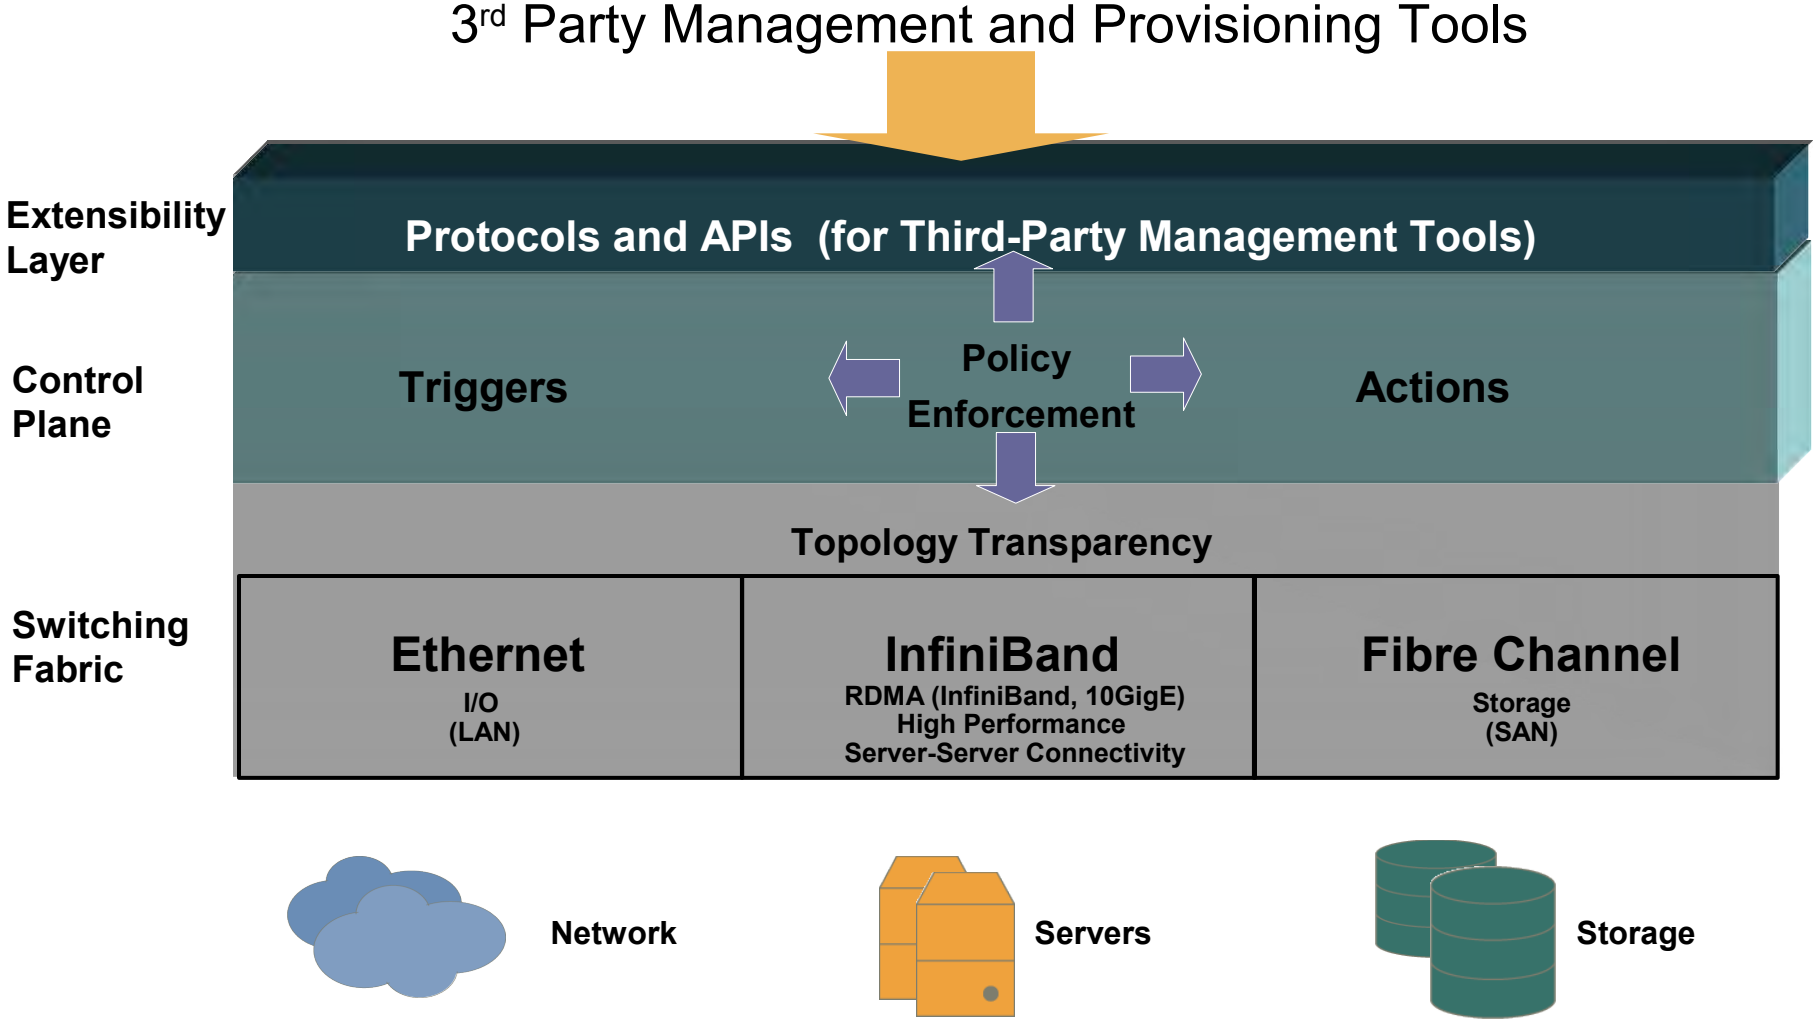

A New Category of Data Center Infrastructure-
The Server Fabric Switch

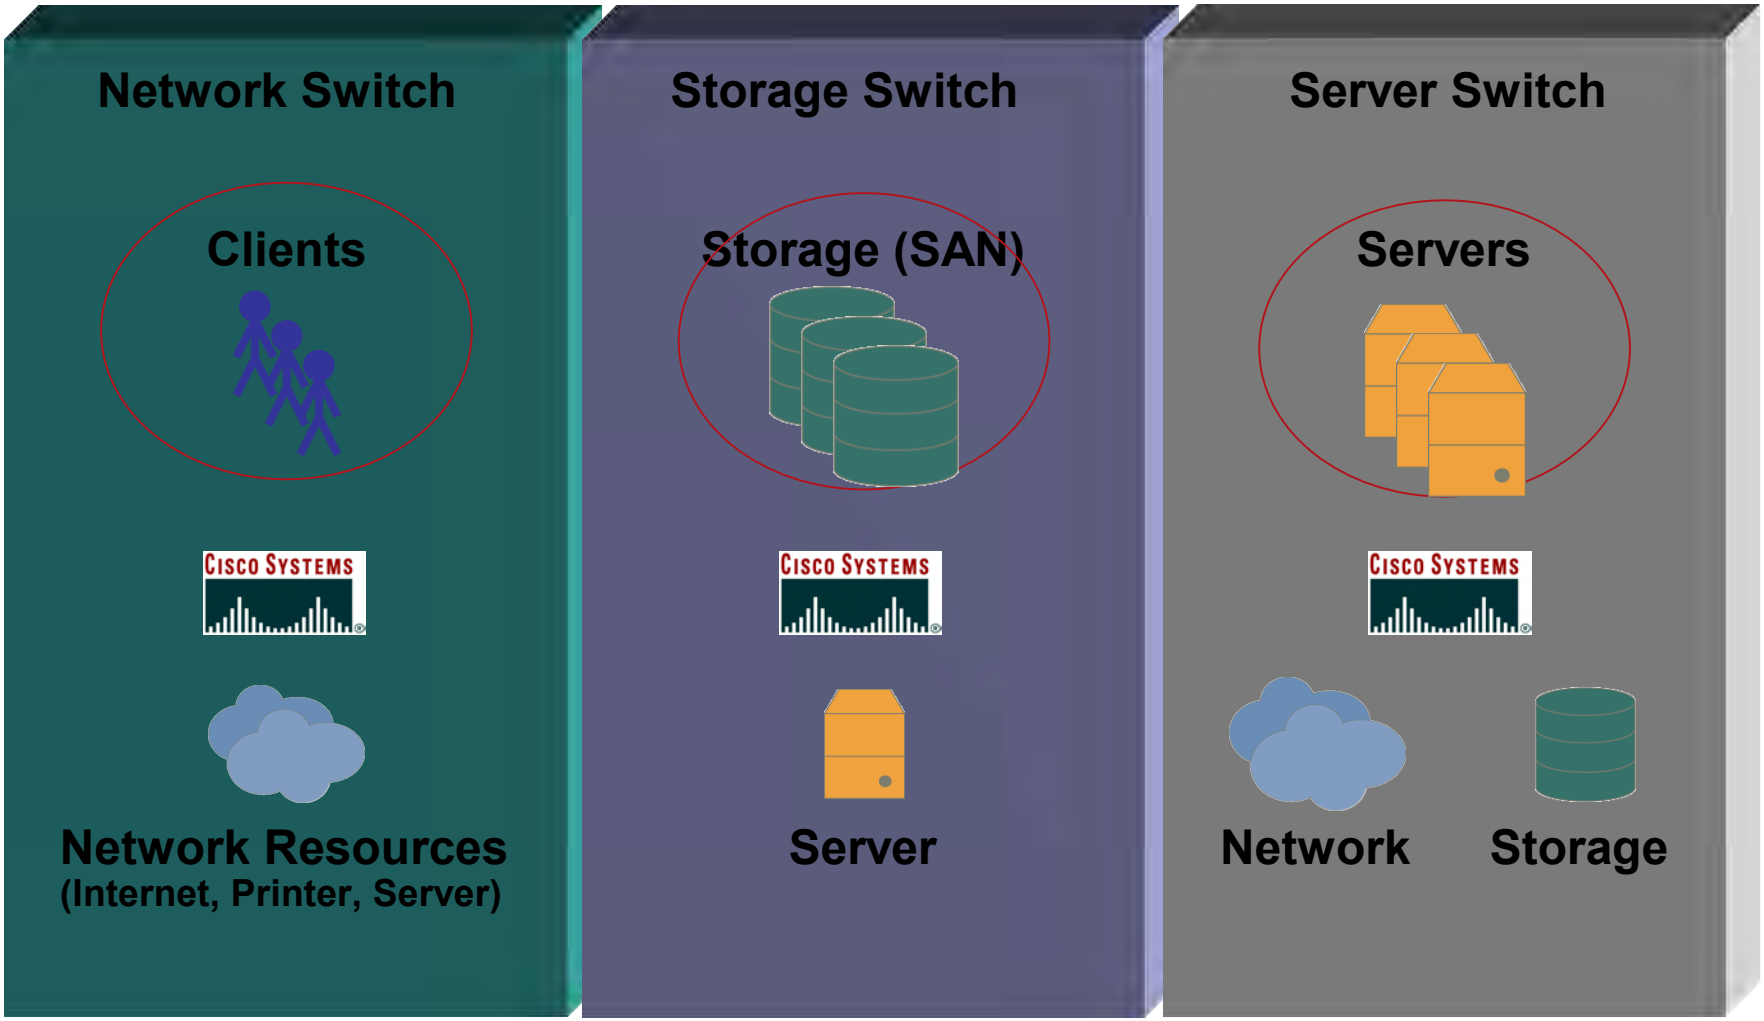

DNA Virtualization Vision

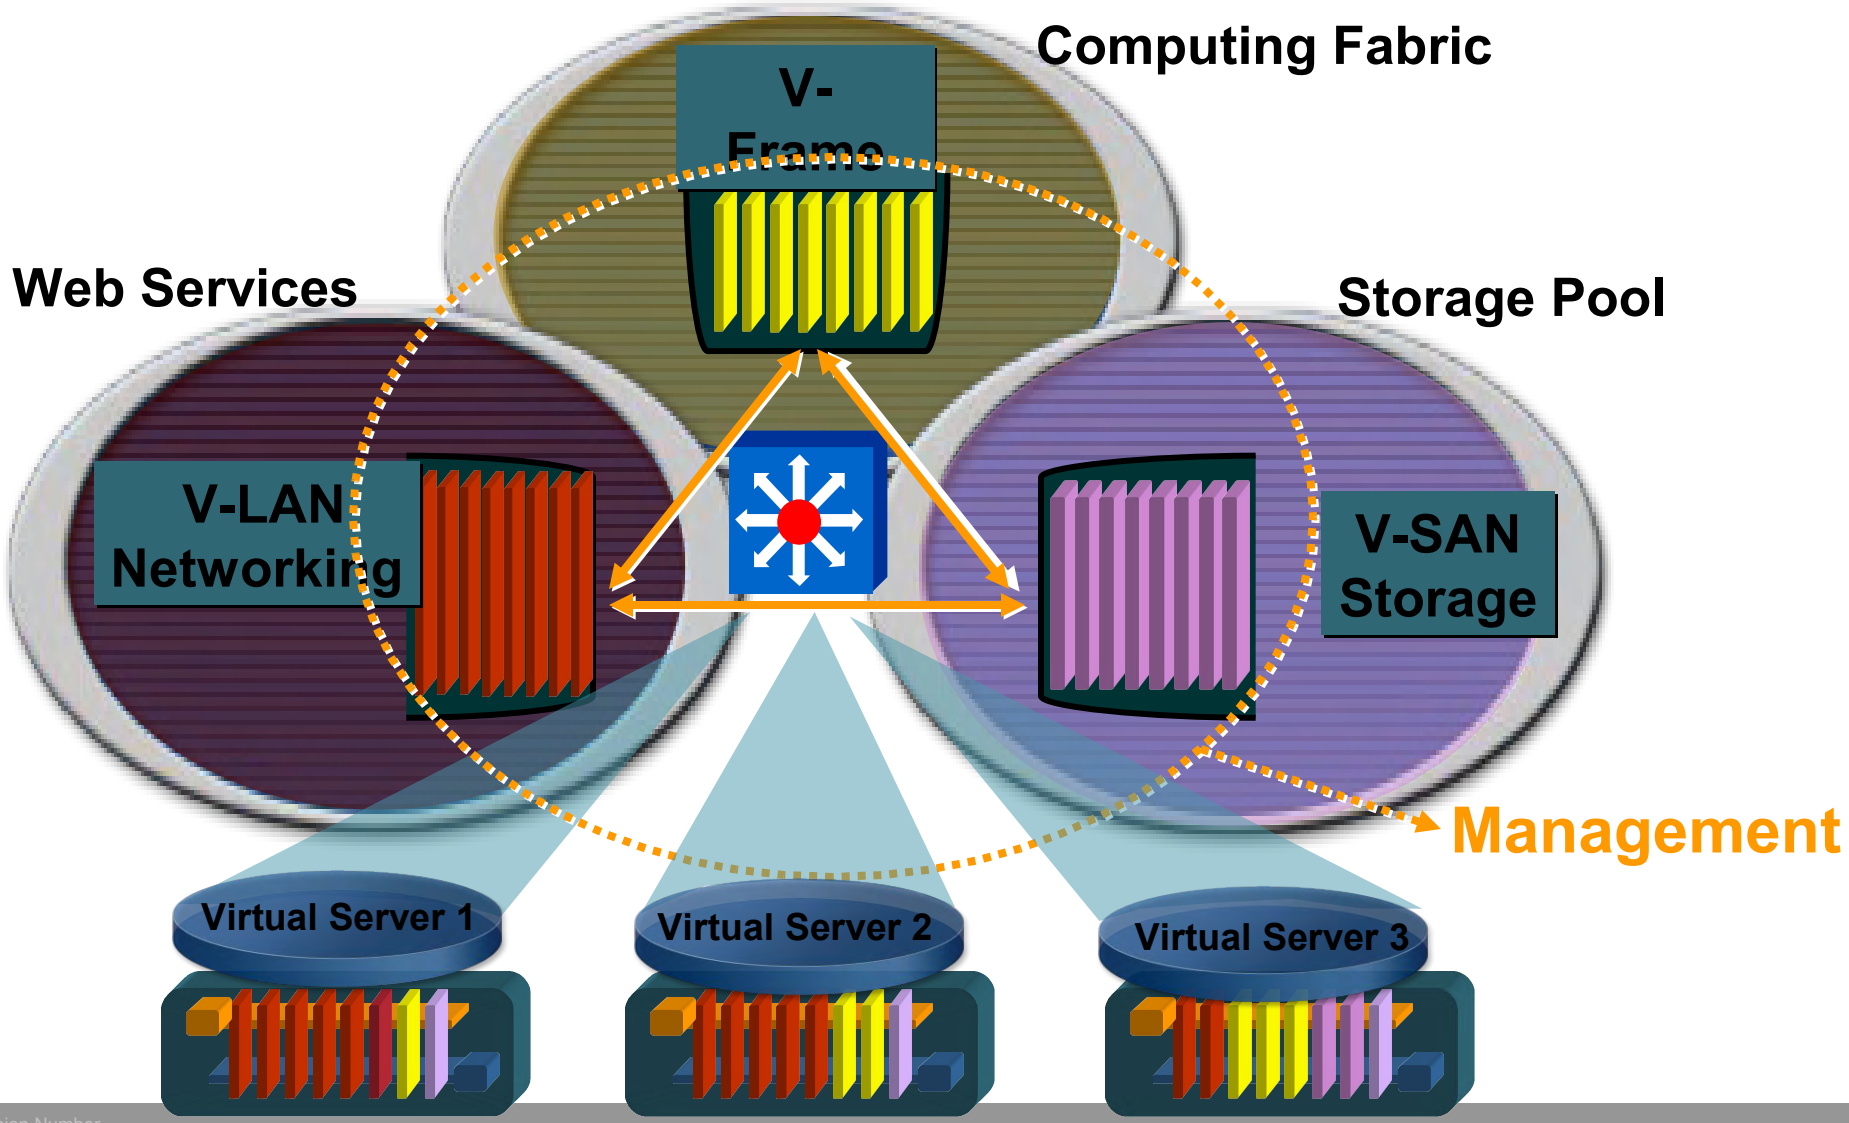

What Makes The Server Fabric Switch Different?

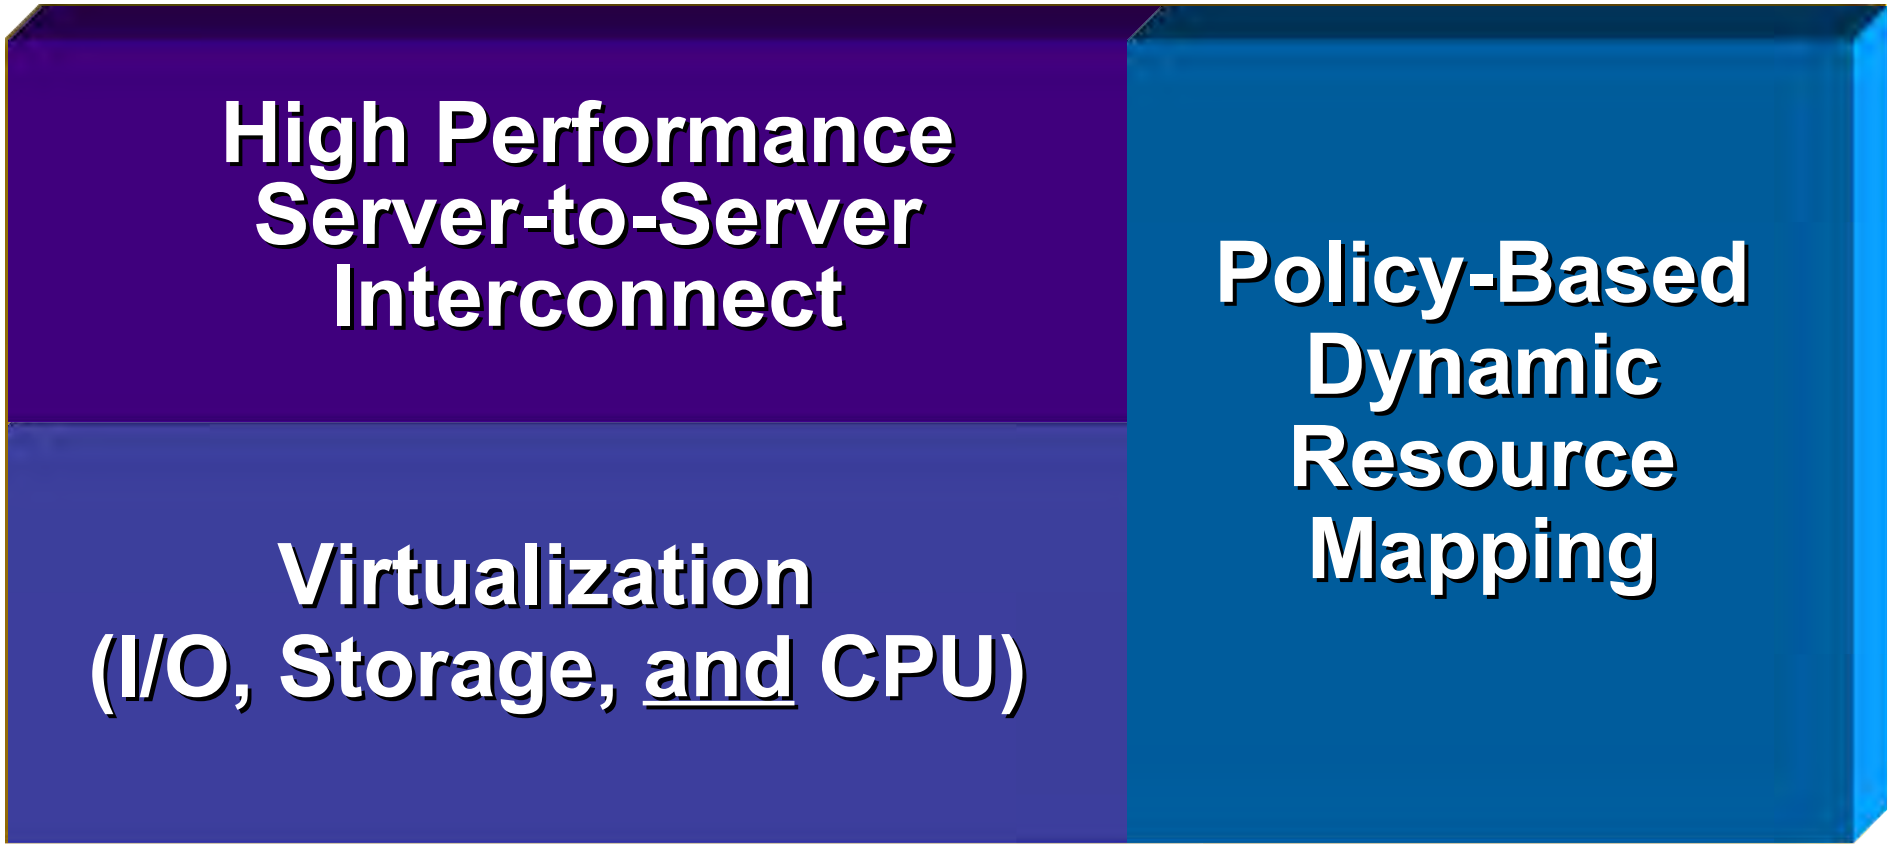

Performance and Control

Server Fabric Switch Applications

Why Performance and Control?

Server Clustering	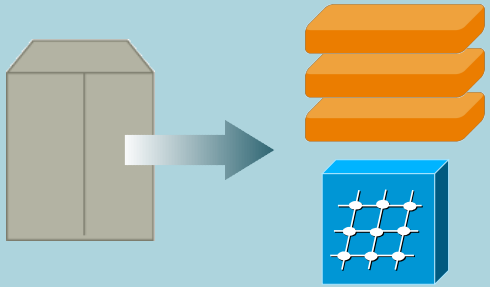	• High Performance Computing (HPC) • "Enterprise-class" HPC • Database scalability
I/O Virtualization	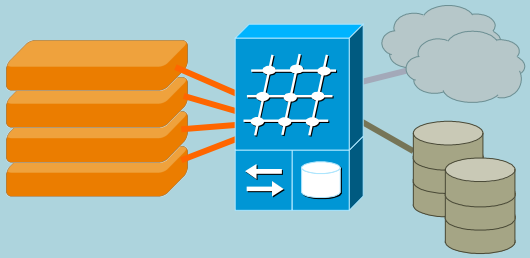	• I/O consolidation • I/O aggregation • Server consolidation
Utility or Grid Computing	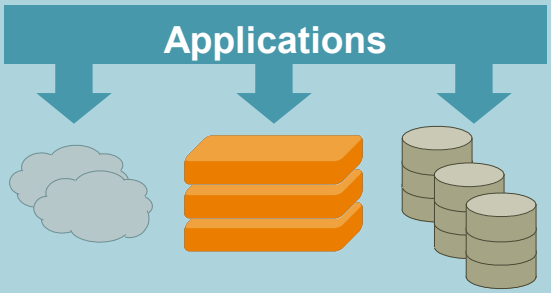	• Application provisioning • Server re-purposing • Server migration

RDMA & OS Bypass, Kernel Bypass

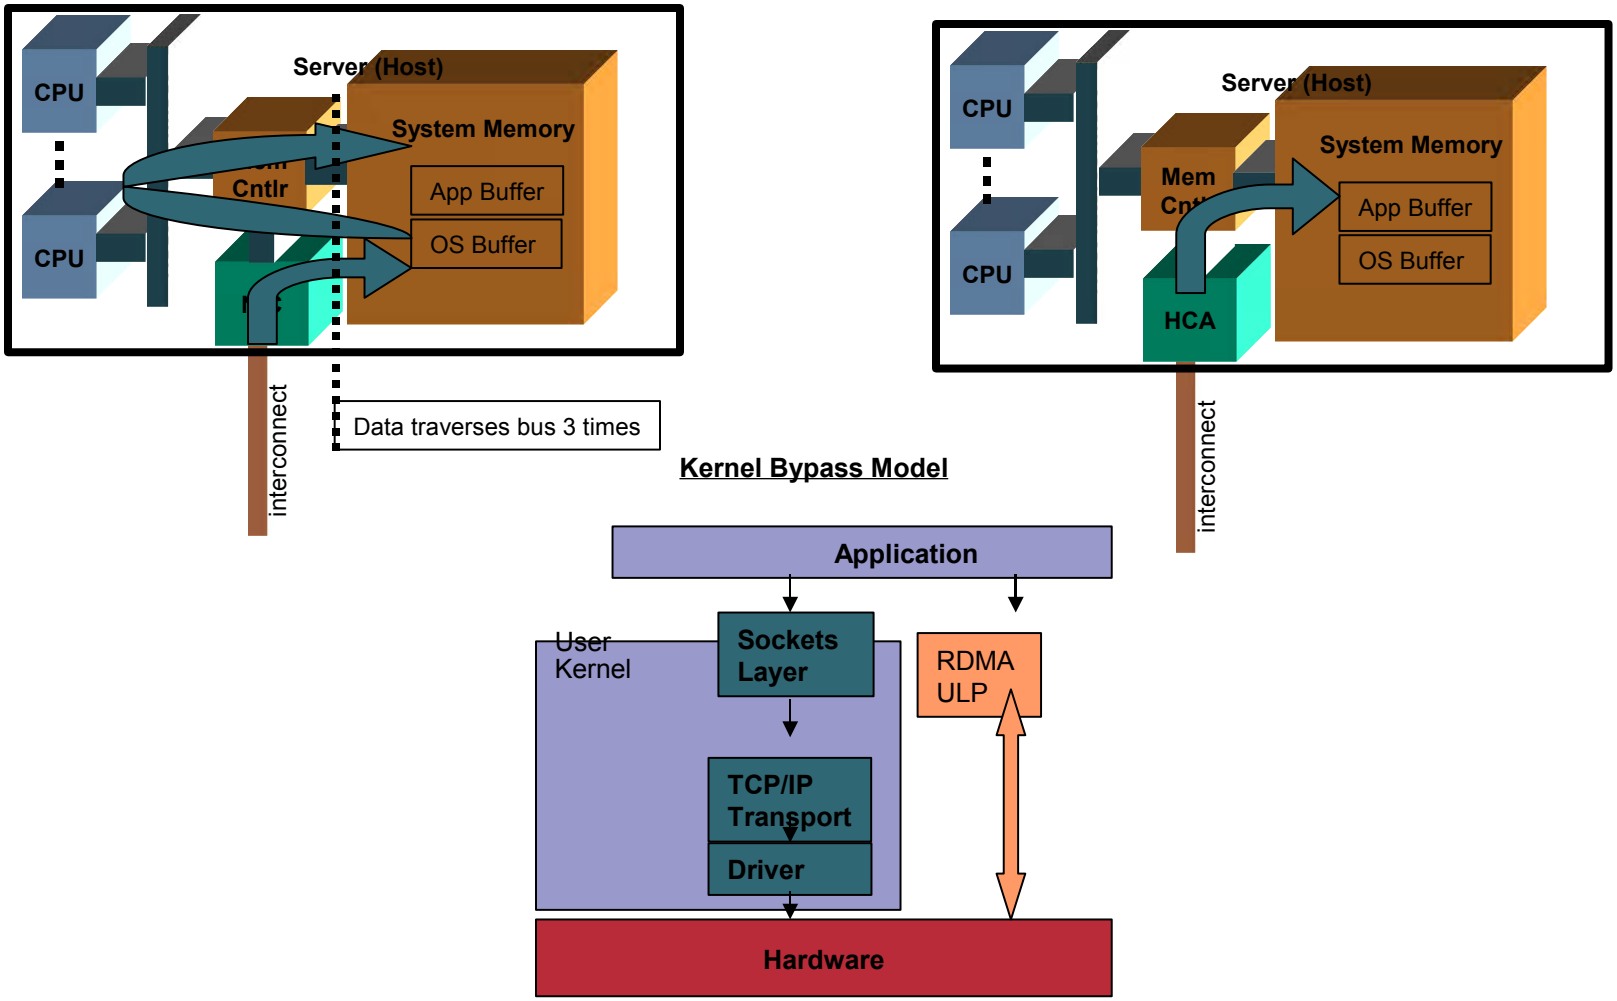

Cluster Application Interconnect

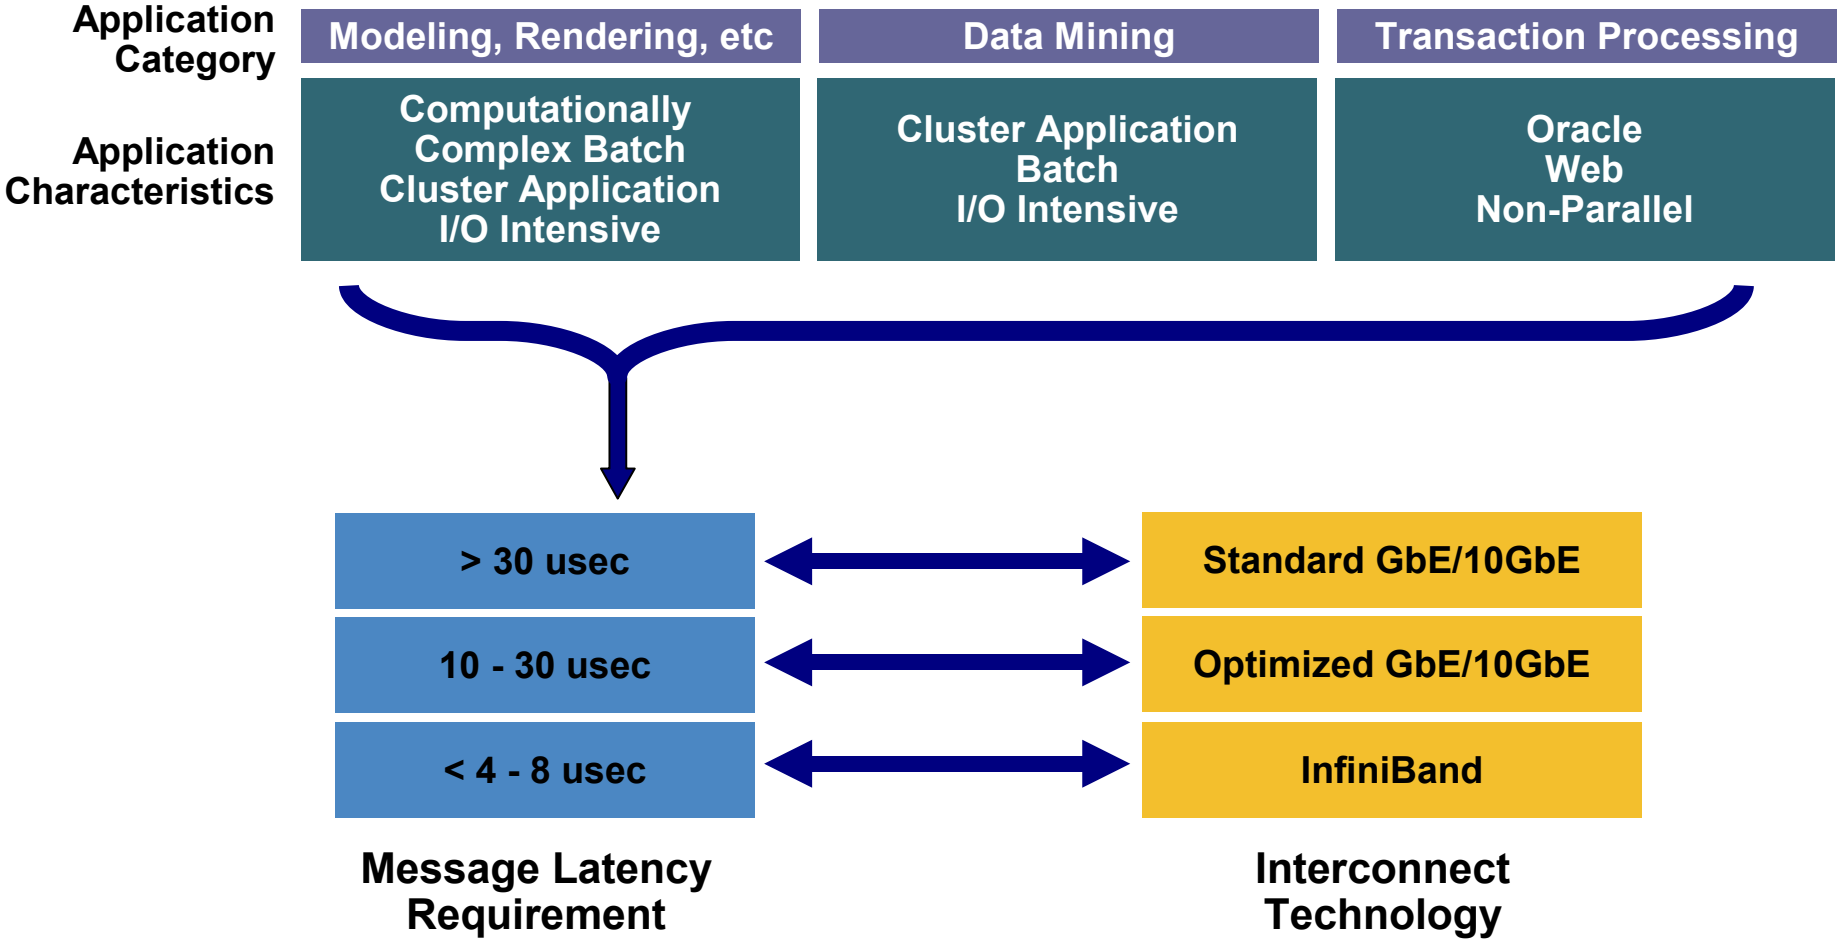

InfiniBand Performance
Measured Results

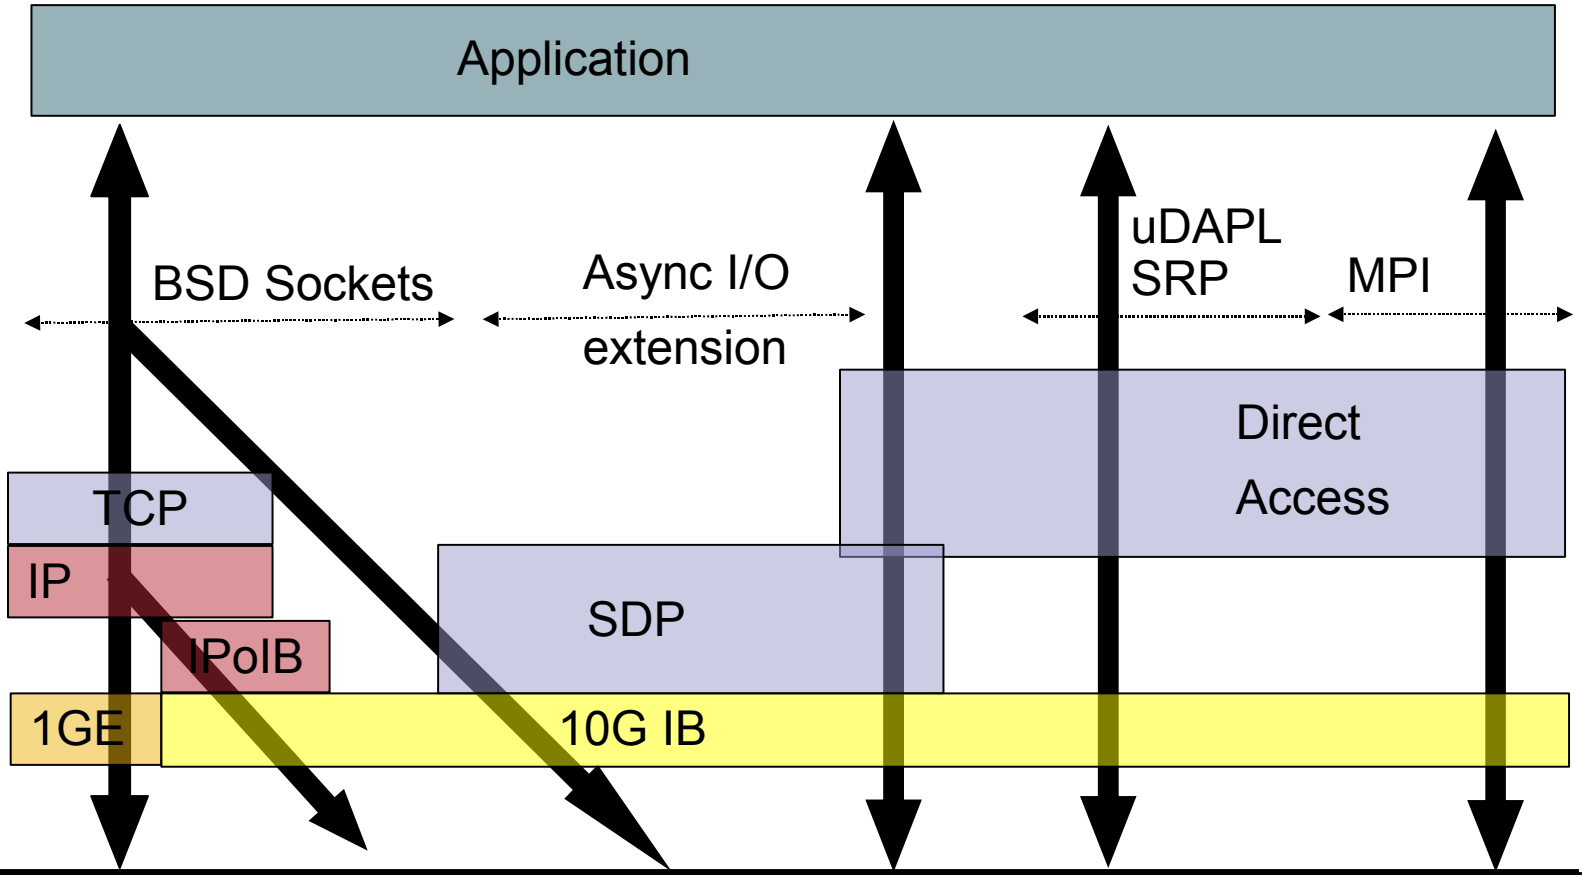

Throughput	1 Gb/s	4.1Gb/s	4.5 Gb/s	7.9Gb/s	8Gb/s	8Gb/s
Latency	40-60 usec	30 usec	18 usec	18 usec	8 usec	3.5 usec

I/O Gateways for Network and Storage
Eliminating Technology Islands

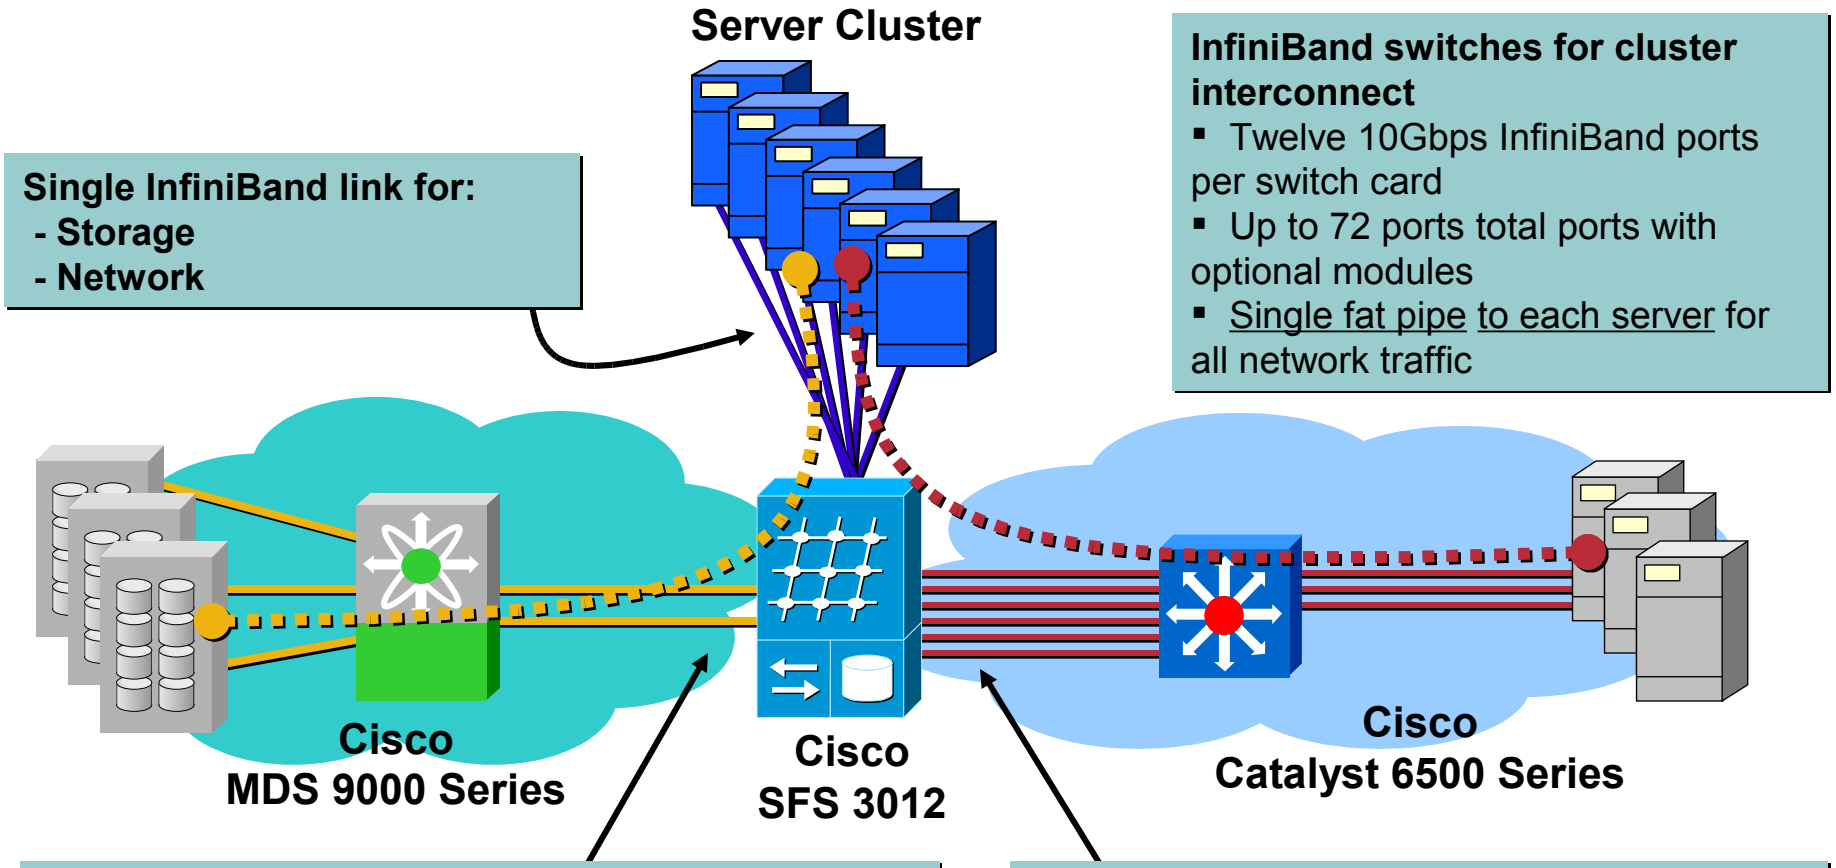

Programmability
VFrame™

1) Server Switch receives policy from VFrame™ Director or 3rd party software.
2) Based on policy, Server Switch assembles the virtual server
   - Selects server(s) that meet minimum criteria (e.g. CPU, memory)
   - Boot server(s) over the network with appropriate app/os image
   - Creates virtual IPs in servers and maps to VLANs for client access.
   - Creates virtual HBAs in servers and maps to Zones, LUNs, and WWNNs for storage access

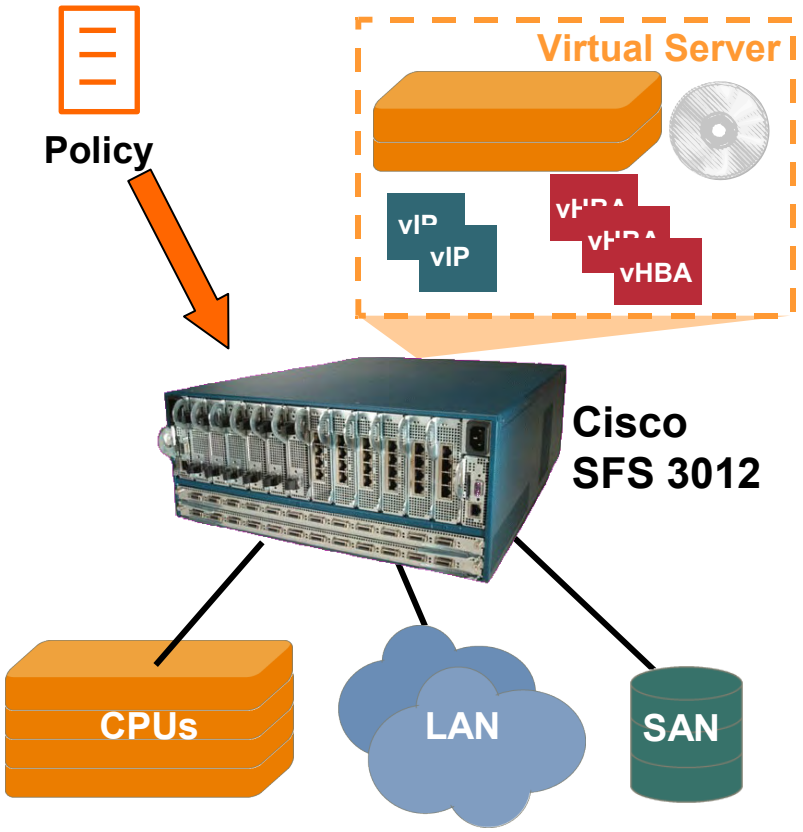

Grid or Utility Computing

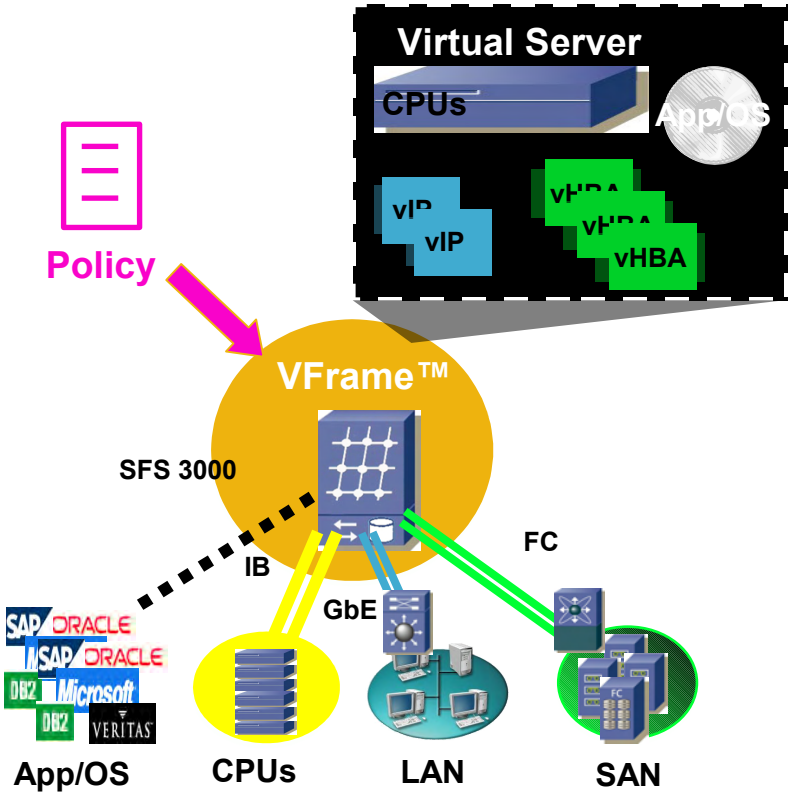

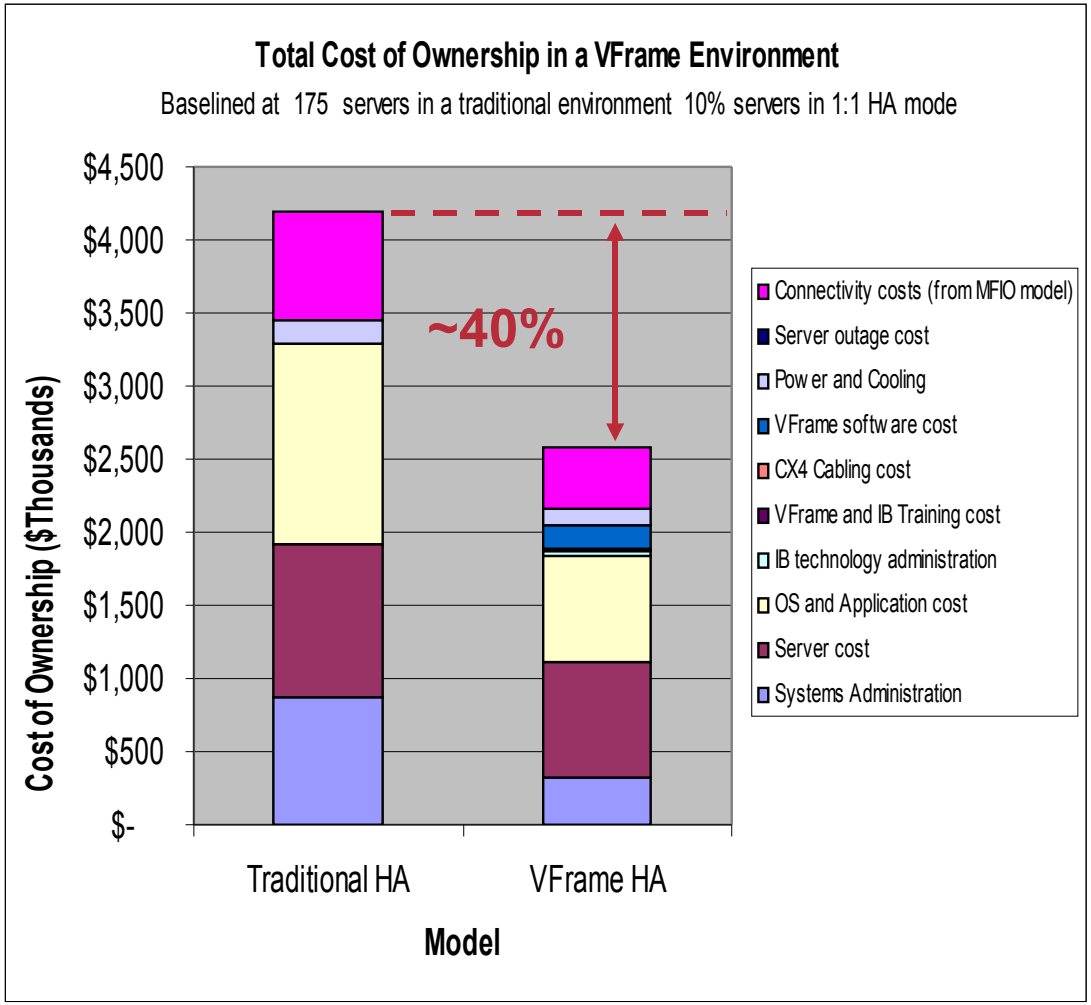

VFrame

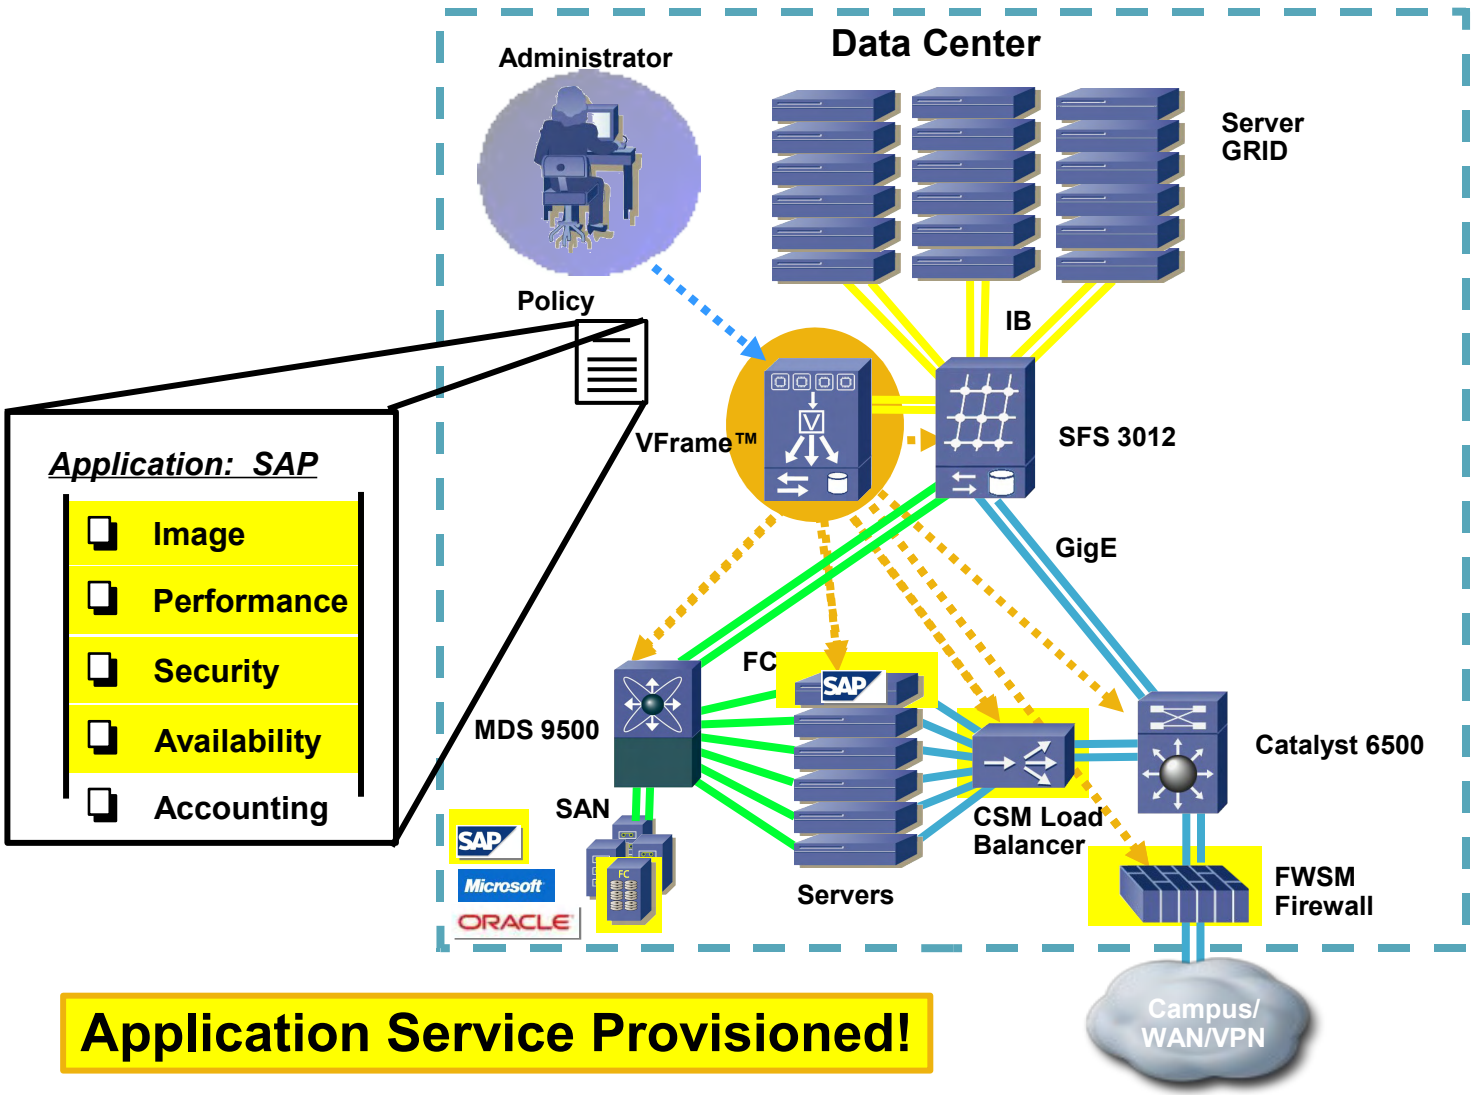

Application Service Provisioned!

Define application services and pass policy to VFrame

VFrame translates policies to actions and passes to infrastructure

VFrame identifies right App / OS Image From storage

VFrame picks server with right criteria to run application and boots server

VFrame gives new server right VLAN and LUN info so it can find/be found by right clients and storage

VFrame provisions security policies to FWSM

VFrame provisions CSM to add new server to load balancing pool

The Next Frontier: Datacenter Management

Device and Service Provisioning + Virtualization

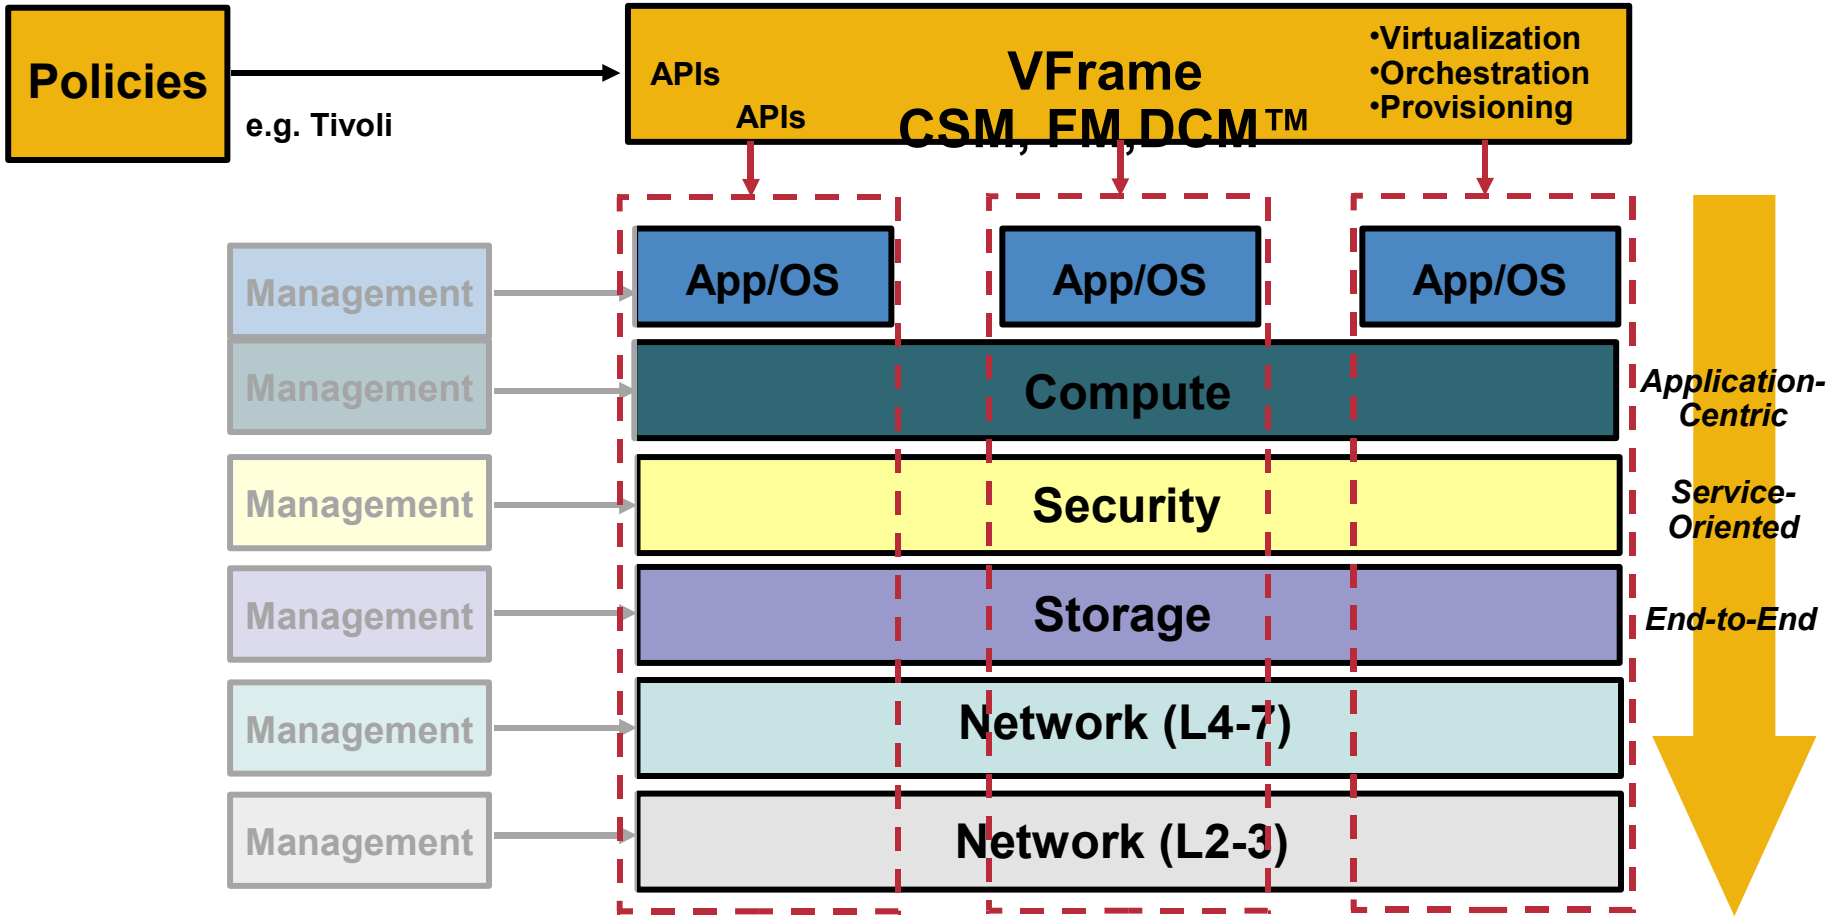

Horizontal versus Vertical Provisioning

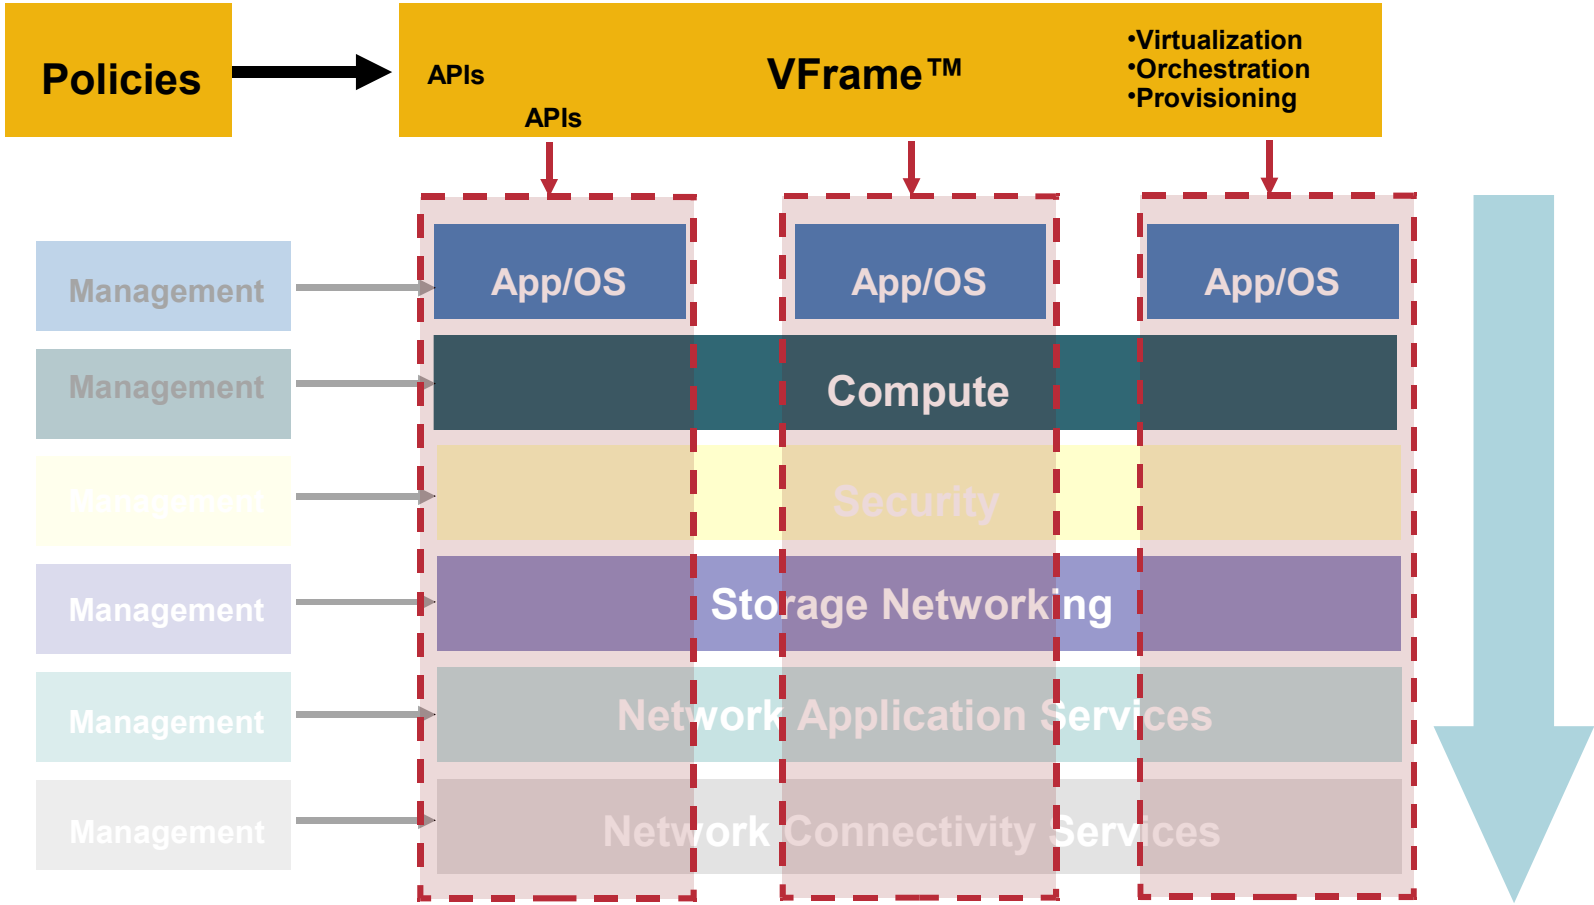

Datacenter Ecosystem

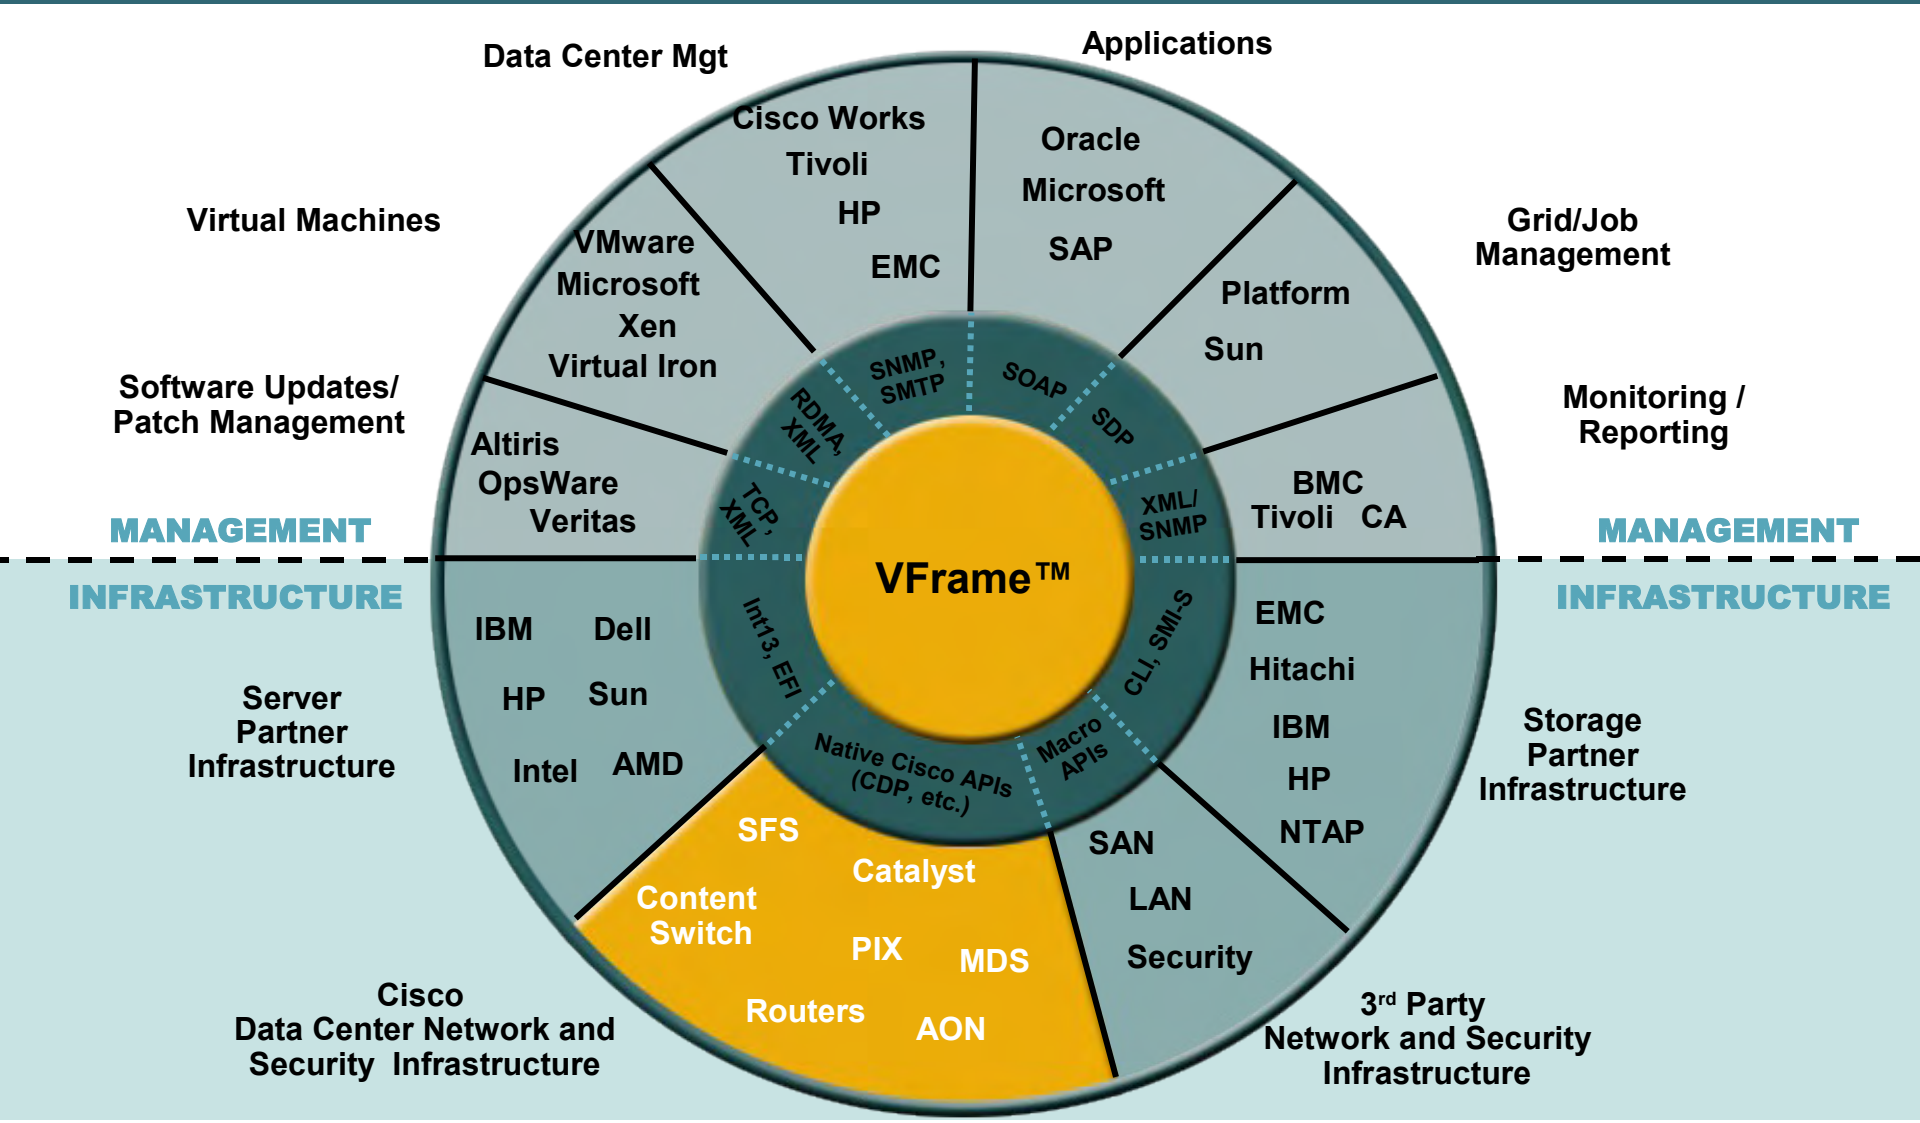

What is InfiniBand?

- **InfiniBand is a high speed – low latency technology used to interconnect servers, storage and networks within the datacenter**
- **Standards Based – InfiniBand Trade Association**

  http://www.infinibandta.org
- **Scalable Interconnect:**

  1X = 2.5Gb/s (2Gb/s data)

  4X = 10Gb/s (8Gb/s data)

  12X = 30Gb/s (24Gb/s data)

High Performance Server Interconnect
InfiniBand

- Industry Standard
- RDMA for Ultra-Low Latency
- 10Gbps Bandwidth (moving to 30Gbps)
- Connection Oriented Control

Manageability

- Standard
- Connection Oriented
- Built-in Control
- Partitionable
- Boot Over IB
- Interconnect Agnostic Storage and I/O
  - –GigE, 10GigE, FC, iSCSI, etc.

Economics

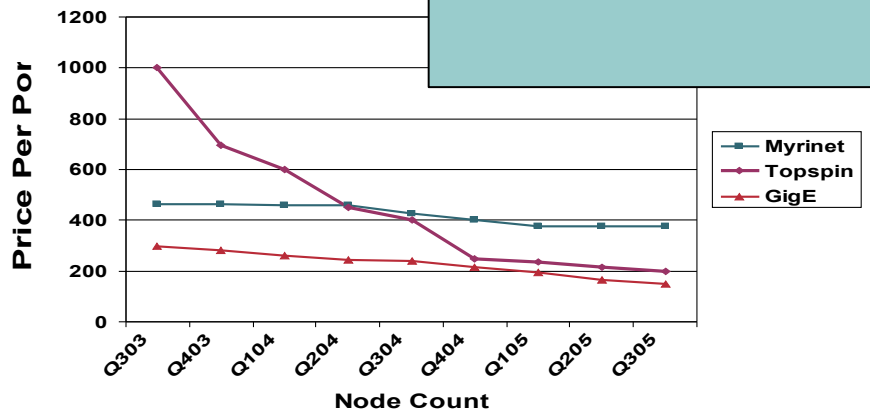

Performance

Throughput		
Latency	70 usec	5 usec
CPU Utilization	40-50%	1-3%

Cluster Application Interconnect

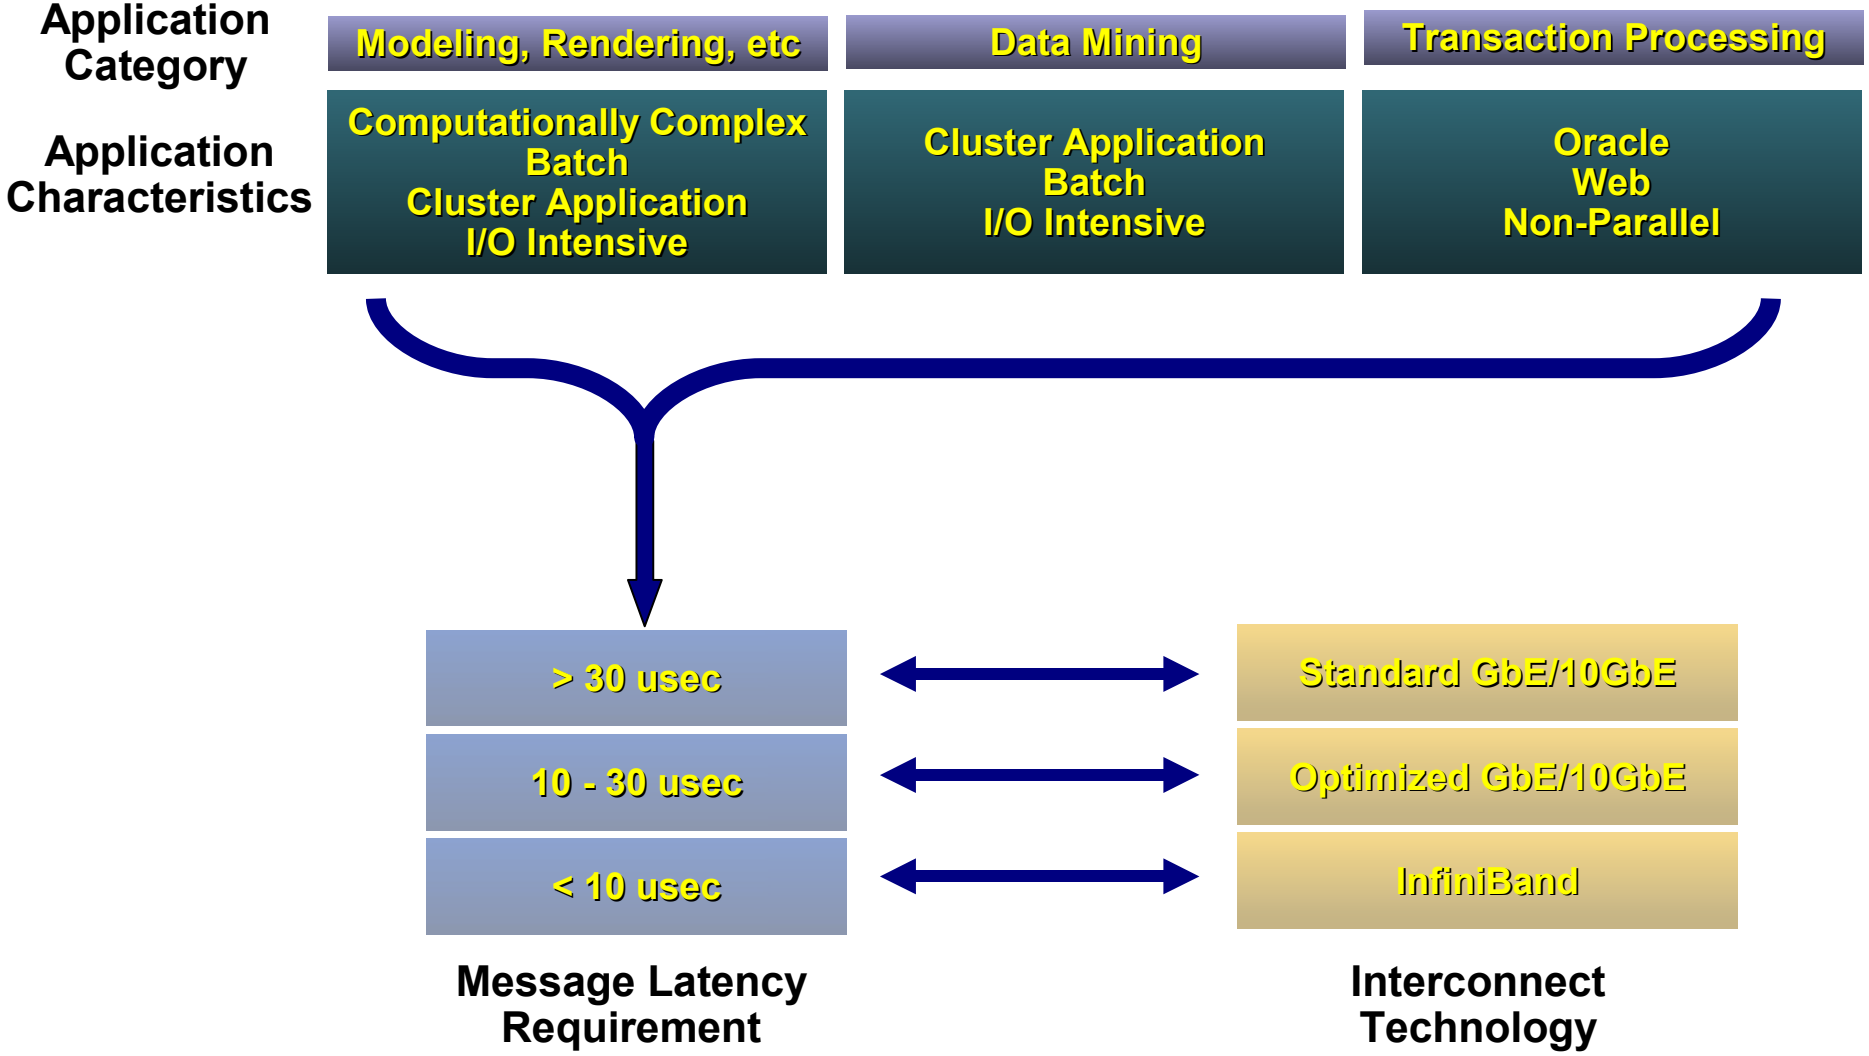

Price / Performance Comparative

InfiniBand Offers the Best Price / Performance for HPC

	InfiniBand PCI-Express	Myrinet D	Myrinet E	10GbE	GbE	GbE/RNIC
Data Bandwidth (Large Messages)	950MB/s	245MB/s	495MB/s	900MB/s	100MB/s	100MB/s
MPI Latency (Small Messages)	5us	6.5us	5.7us	50us	50us	18us
HCA Cost (Street Price)	$550	$535	$880	$2K-$5K	Free	$500
Switch Port	$250	$400	$400	$2K-$6K	$100-$300	$100-$300
Cable Cost (3m Street Price)	$100	$175	$175	$50	$25	$25

Note: MPI "User Space" to "User Space" latency – switch latency is less

** Myrinet pricing data from Myricom Web Site (Dec 2004) utilizing Myrinet's latest switches*
*** InfiniBand pricing data based on Topspin avg. sales price (Dec 2004)*
**** Myrinet, GigE, and IB performance data from independent June 2004 OSU study*
***** 10GigE and GigE Cost and Performance data from Cisco Internal document*

InfiniBand Performance
Measured Results

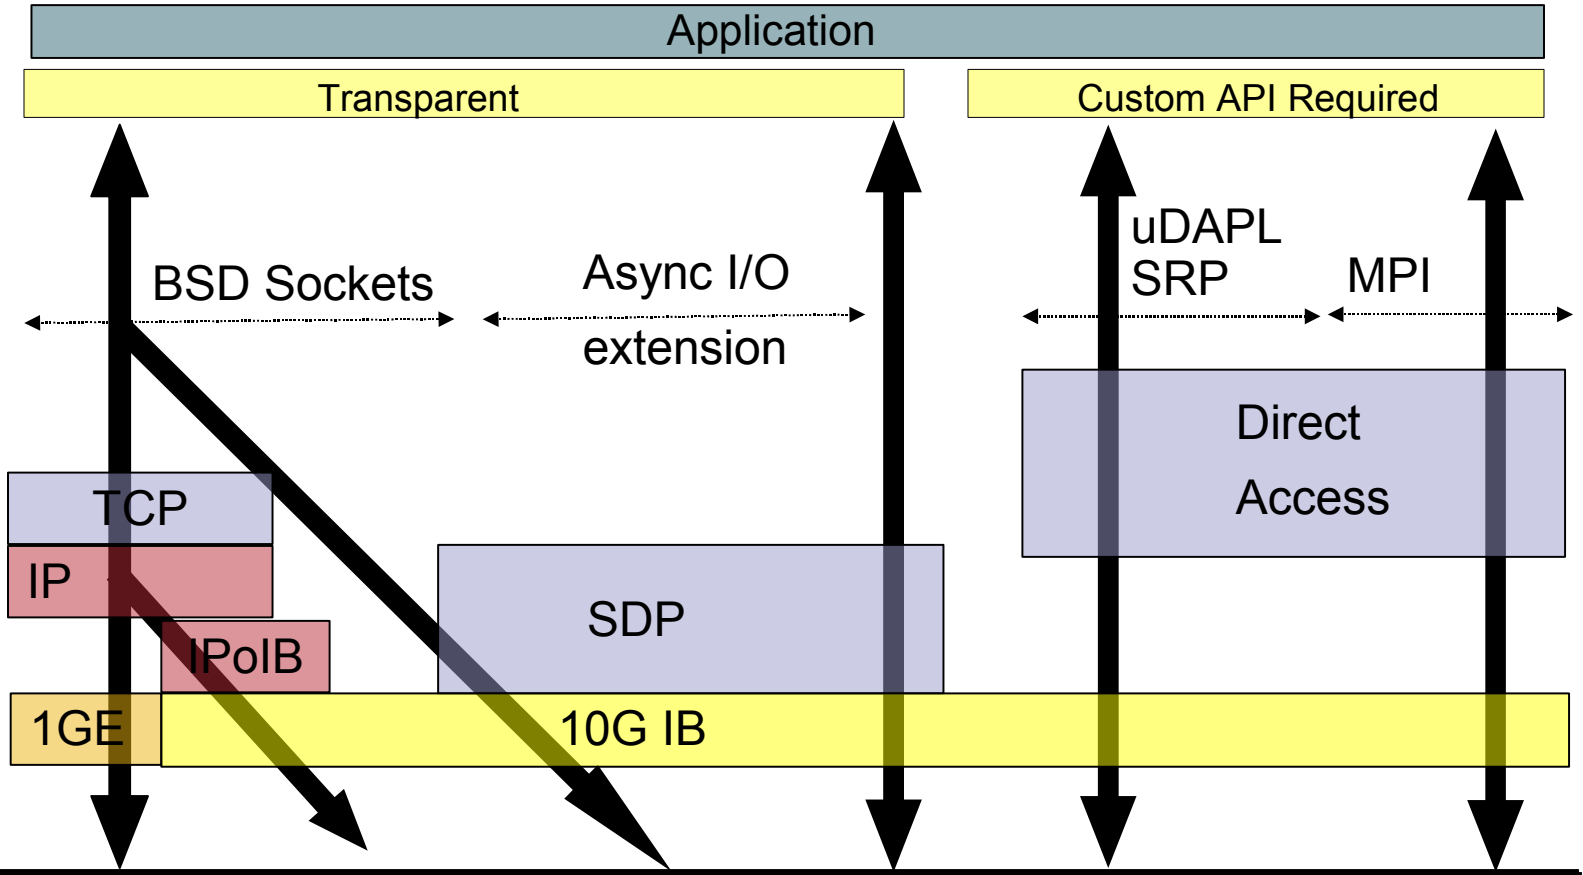

Throughput	1 Gb/s	4.1Gb/s	4.5 Gb/s	7.9Gb/s	8Gb/s	8Gb/s
Latency	40-60 usec	30 usec	18 usec	18 usec	8 usec	4.5 usec

InfiniBand Protocol Summary

Protocol / Application	Summary	Application Example
IPoIB (IP over InfiniBand)	Enables IP-based applications to run over InfiniBand transport.	Standard IP-based applications. When used in conjunction with Ethernet Gateway, allows connectivity between IB network and LAN.
SDP (Sockets Direct Protocol)	Accelerates sockets-based applications using RDMA.	Communication between database nodes and application nodes, as well as between database instances.
SRP (SCSI RDMA Protocol)	Allows InfiniBand-attached servers to utilize block storage devices.	When used in conjunction with the Fibre Channel gateway, allows connectivity between IB network and SAN.
uDAPL (Direct Access Programming Library)	Enables maximum advantage of RDMA flexible programming API.	Used for IPC communication between cluster nodes for Oracle 10G RAC.
MPI (Message Passing Interface)	Low latency protocol used widely in HPC environments.	HPC applications.

The InfiniBand Driver Architecture

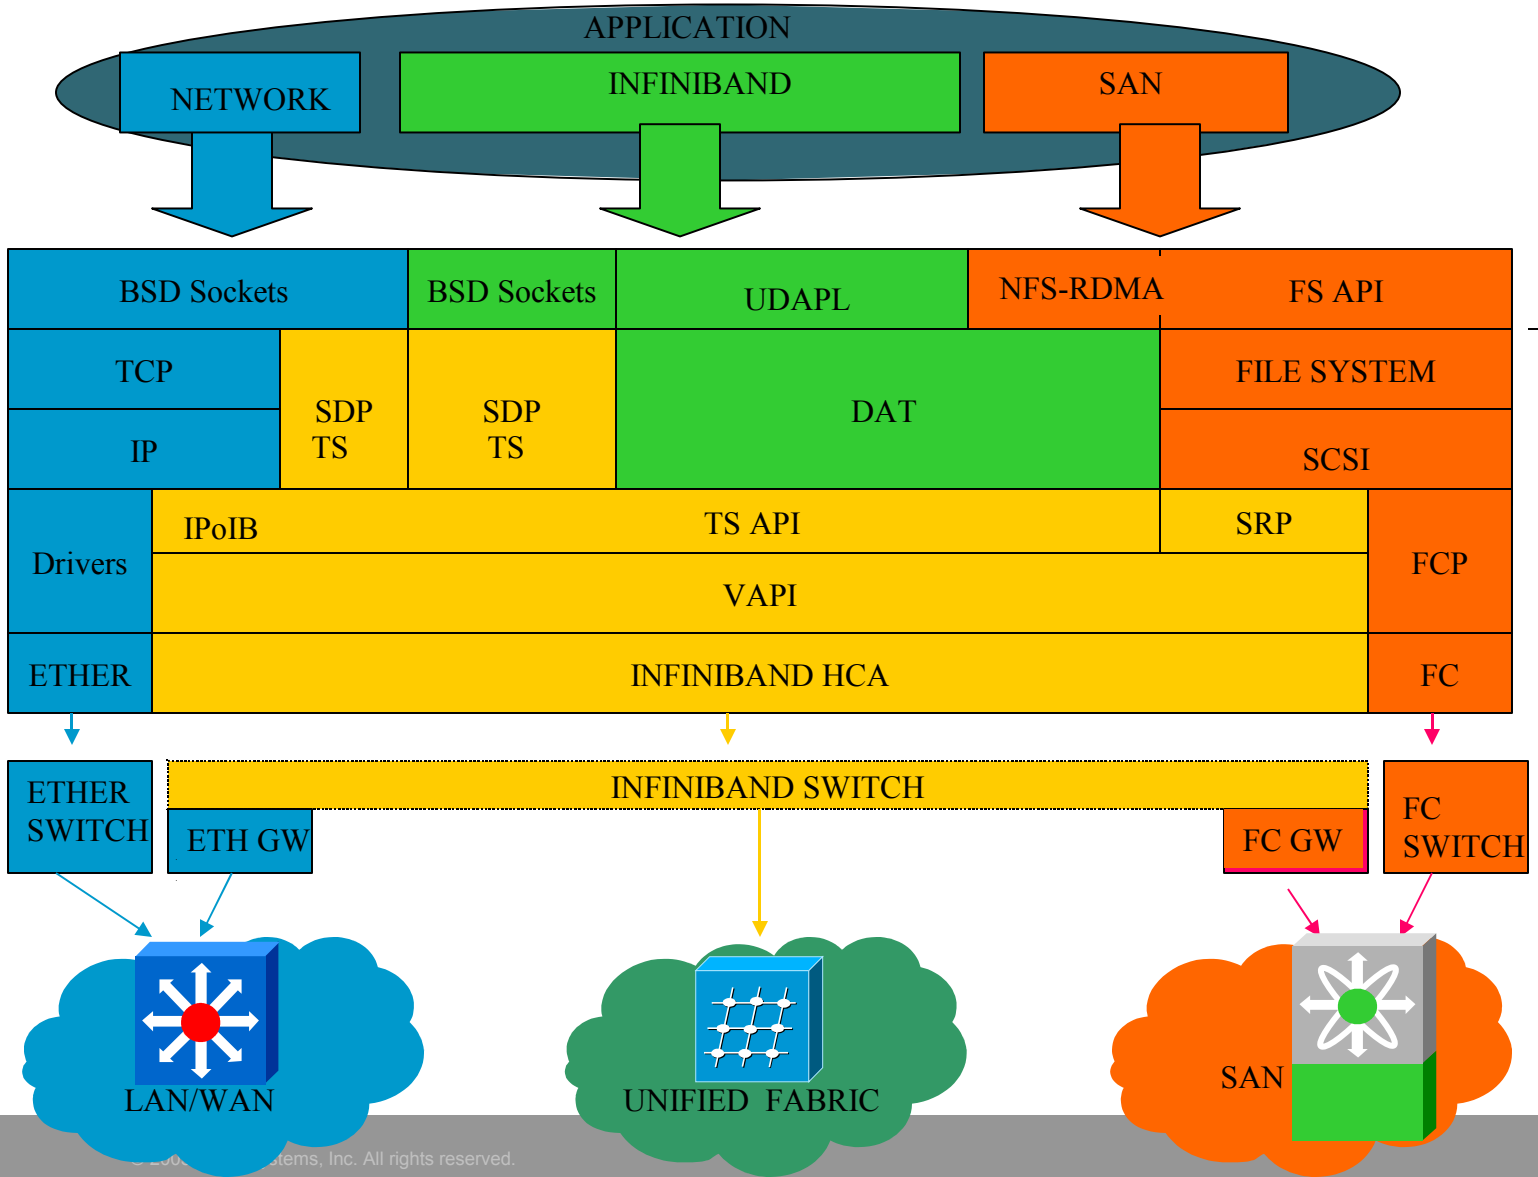

IP over InfiniBand

- Transmission of IP over Infiniband
  - Use IB as a link layer for IP
  - Define data link and link layer address
  - Encapsulation for ARP, IPv4 and IPv6
  - Address resolution
  - Transport IP multicast over IB
- Provides highest level of application compatibility.
- Applications do not need to be re-written or re-compiled
- Standard IP utilities and applications work as usual:
  - Ifconfig, ping, telnet, File sharing (NFS, CIFS); Login access (ssh, telent, etc); Cluster heartbeat
  - DHCP over IB
  - IP over InfiniBand MIB

How IP over InfiniBand works

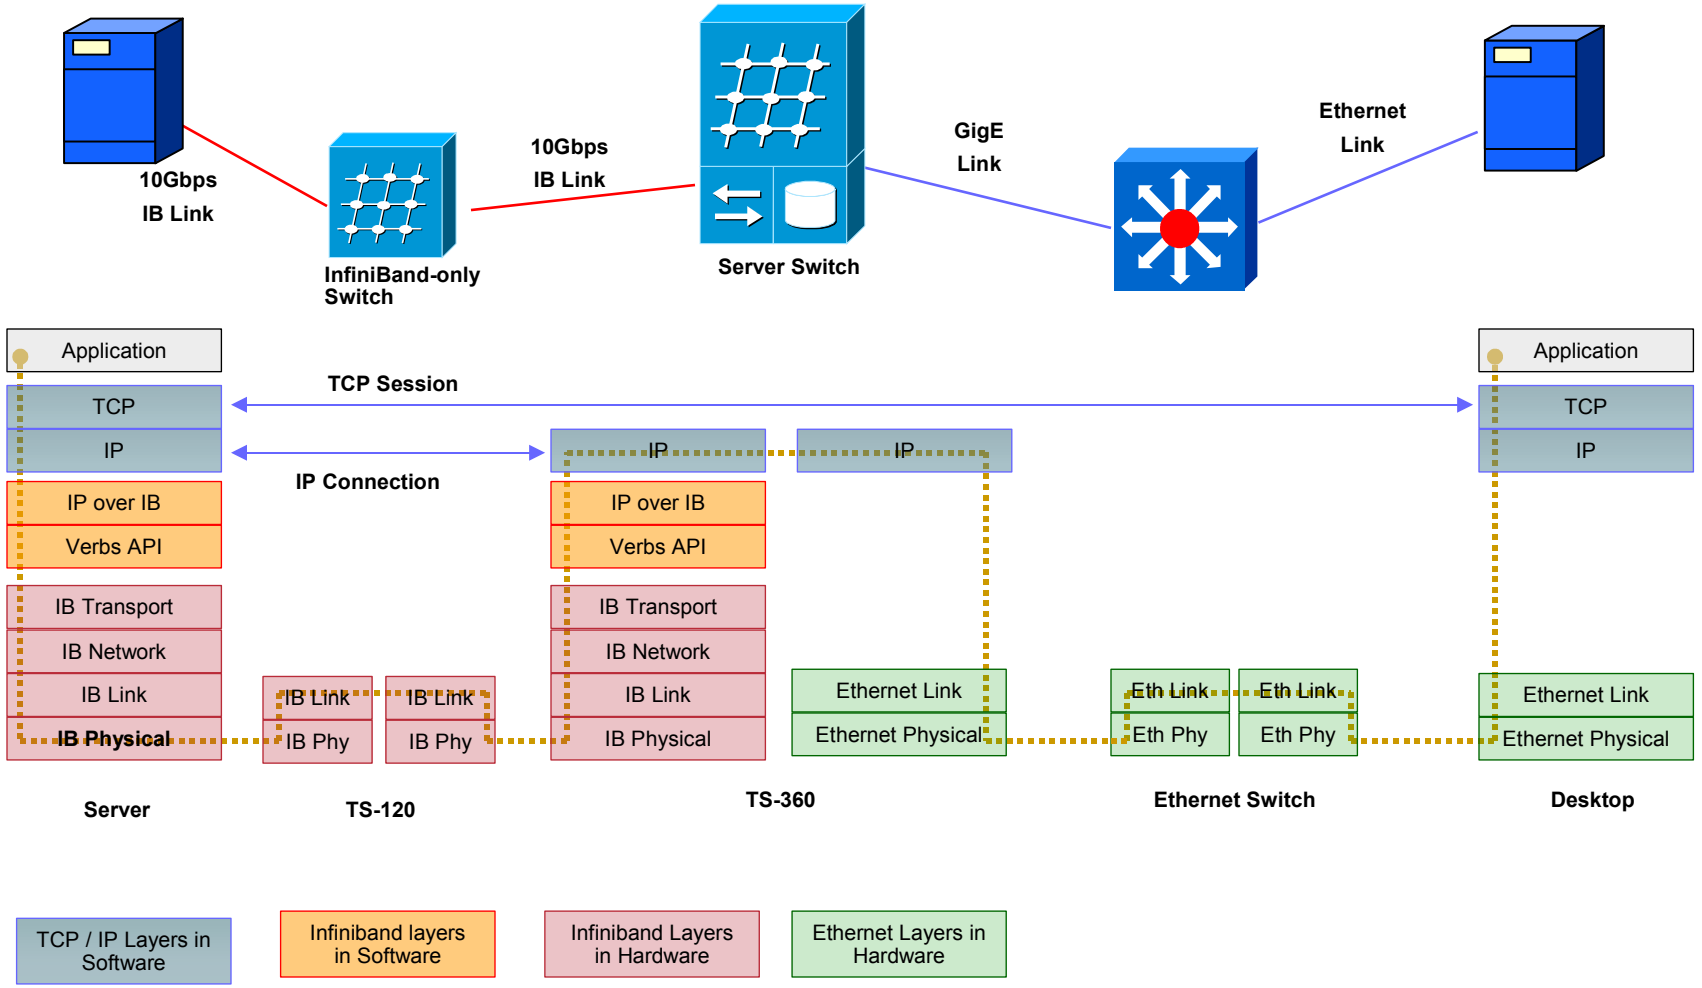

* Notes: Uses standard Berkeley TCP/IP libraries

Sockets Direct Protocol

- Sockets Direct Protocol
- Runs socket based TCP/IP traffic with TCP and copy offload
- Highly configurable:
  - By process
  - By port
  - By destination
  - By environment variable
- No application recompile or rework necessary
- Zero copy capability using Asynchronous I/O (AIO)

Tangible Benefits of Server Fabric Switching

- Purchase 20X more compute power for same dollars
  (pay as you grow, moore's law, expense vs capitalization)
- 50% cost savings from resource consolidation delivers instantaneous ROI
  (Single Server Fabric- eliminates adapter, cables, ports)
- Dramatically Reduce TCO - Manage enterprise-wide Server GRID centrally
  (wire once, *control* servers over the network)
- Provision New Servers in seconds, not days (or weeks)
- Help Eliminate Server Downtime
  (Failover provision, add/remove I/O or storage bandwidth on the fly)
- Control Ballooning Investments in Real Estate, Power & Cooling

  (capitalize on dense server packaging and Blade architectures)
- Political Power and "Self-Rule" for the Server Team
  (Eliminate dependence on other teams to get apps provisioned quickly)

Cisco DNA Impact : Improved Server Utilization

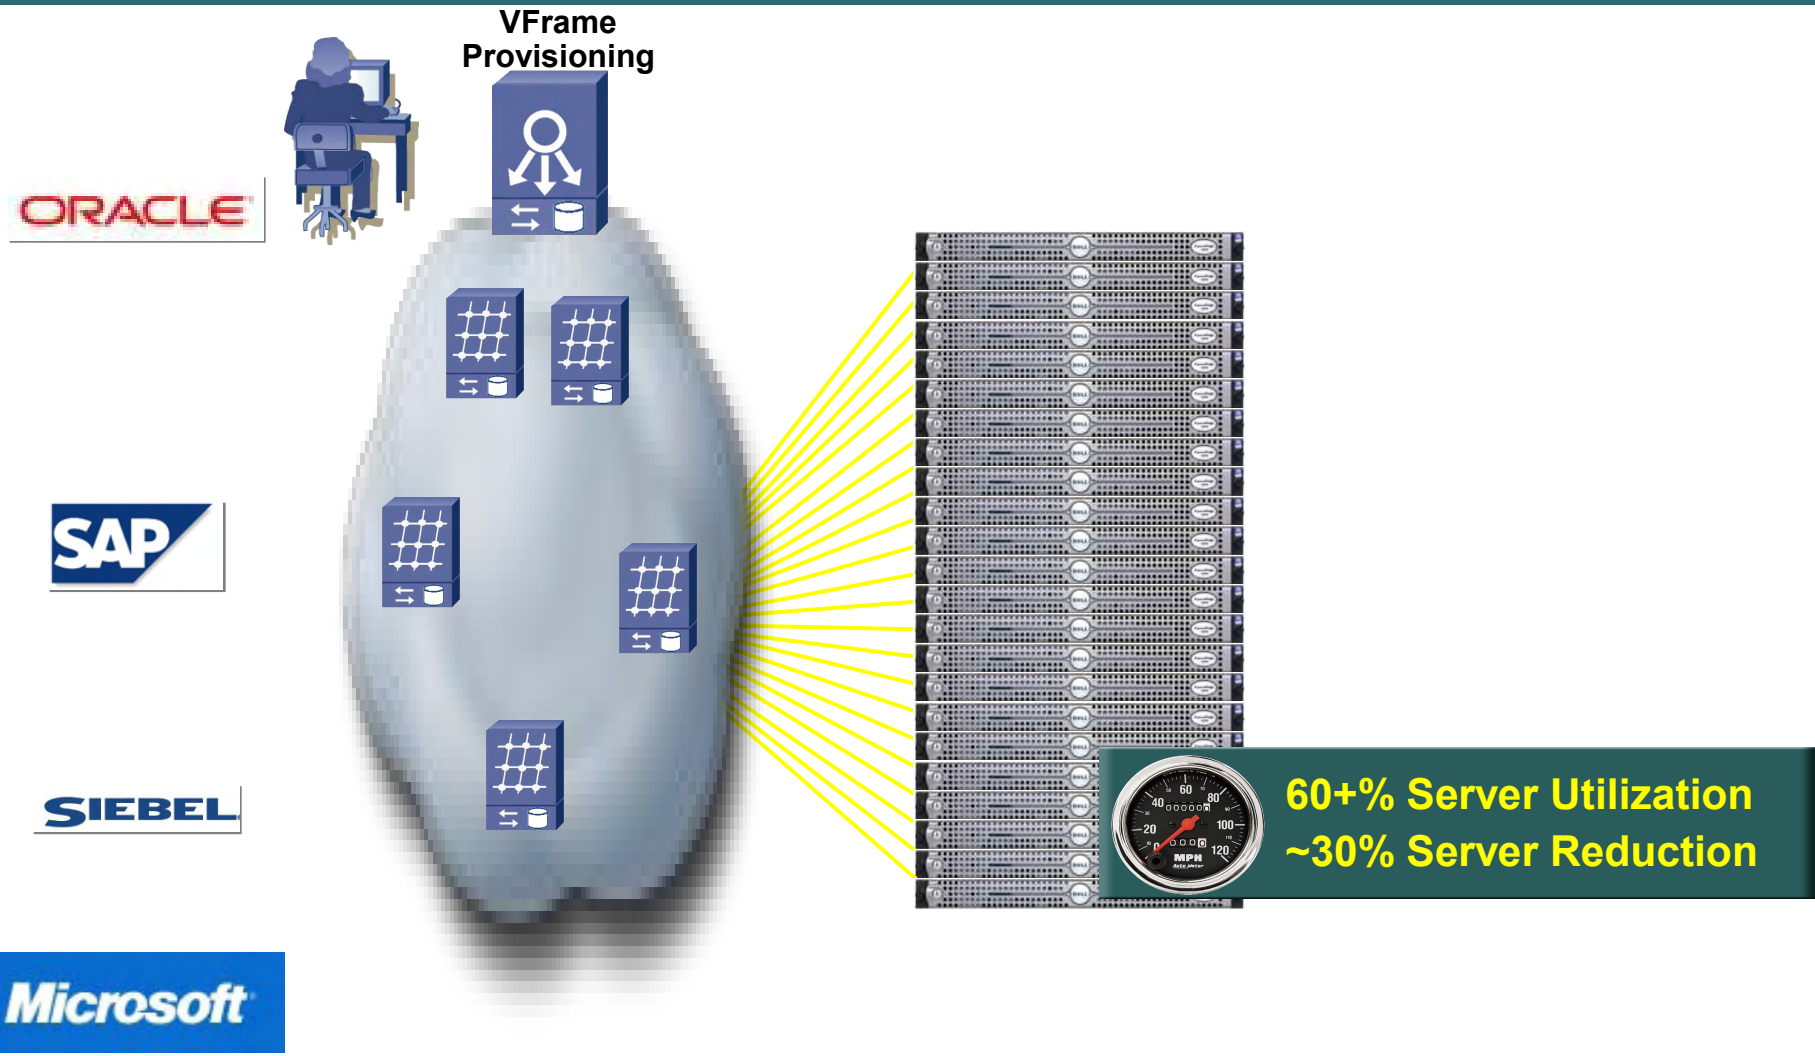

Dramatic Cisco DNA Impact: Application Acceleration
Improvement on Response Times

Application	Software	Before	After	Transaction Time Reduction
Call center	PeopleSoft	63 sec	23 sec	63% (270%)
“JIT” manufacturing	SAP	76 sec	22 sec	71% (350%)
Store management	IBM WebSphere	46 sec	16 sec	66% (290%)
Claims management	IBM WebSphere	42 sec	19 sec	55% (220%)
Collaboration	Lotus iNotes	90 sec	28 sec	68% (320%)
Employee portal	Plumtree	204 sec	59 sec	71% (350%)
Portal consolidation	SunOne, Vignette	43 sec	6 sec	85% (670%)
Employee portal	PeopleSoft	103 sec	32 sec	69% (320%)
CRM	Siebel	389 sec	133 sec	66% (290%)

Notes: Bandwidth reduction averages 80-90%
All timings are customer-verified using either LoadRunner, FineGround AppScope, or in-house customer tools

Topspin Building Blocks

Switches

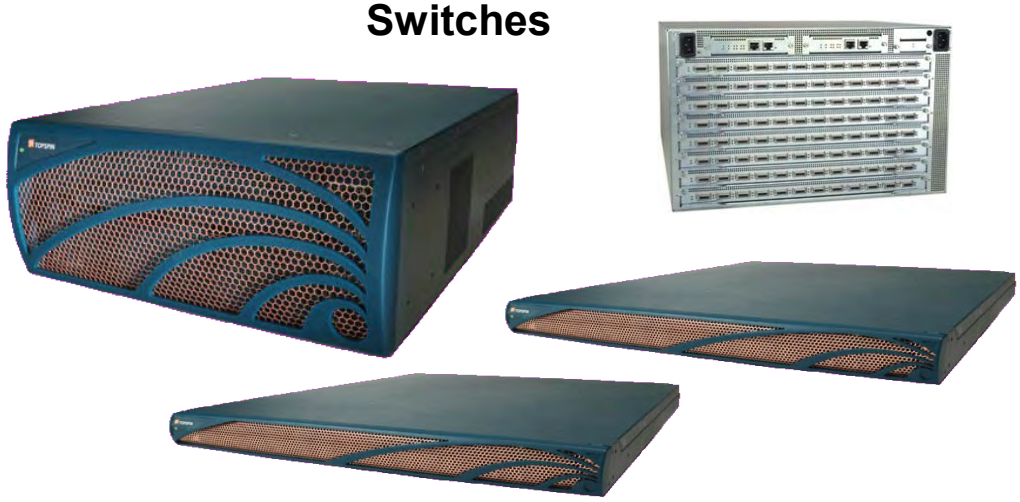

Gateway Modules
- InfiniBand to Ethernet
- InfiniBand to Fibre Channel

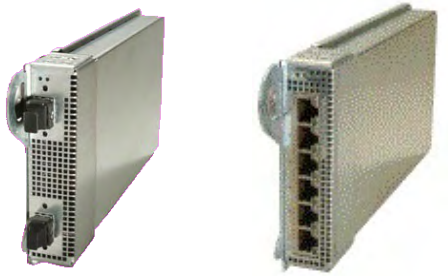

Host Channel Adapter (HCA)
With upper layer protocols

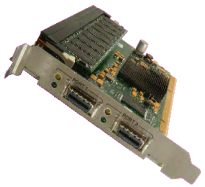

- SRP
- SDP
- uDAPL
- MPI
- IPoIB

Integrated System and Fabric management

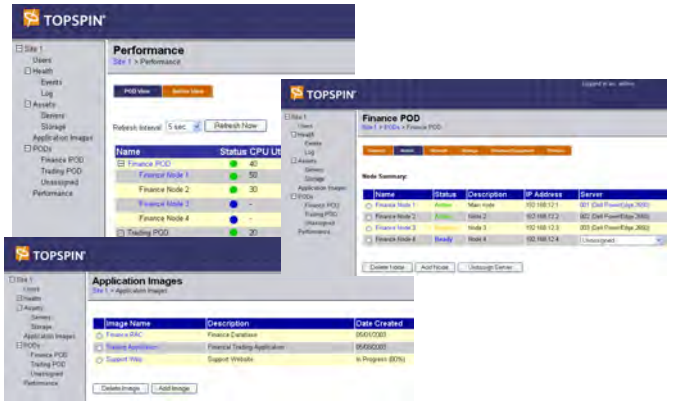

Linux and Windows driver support

The Cisco SFS Product Line

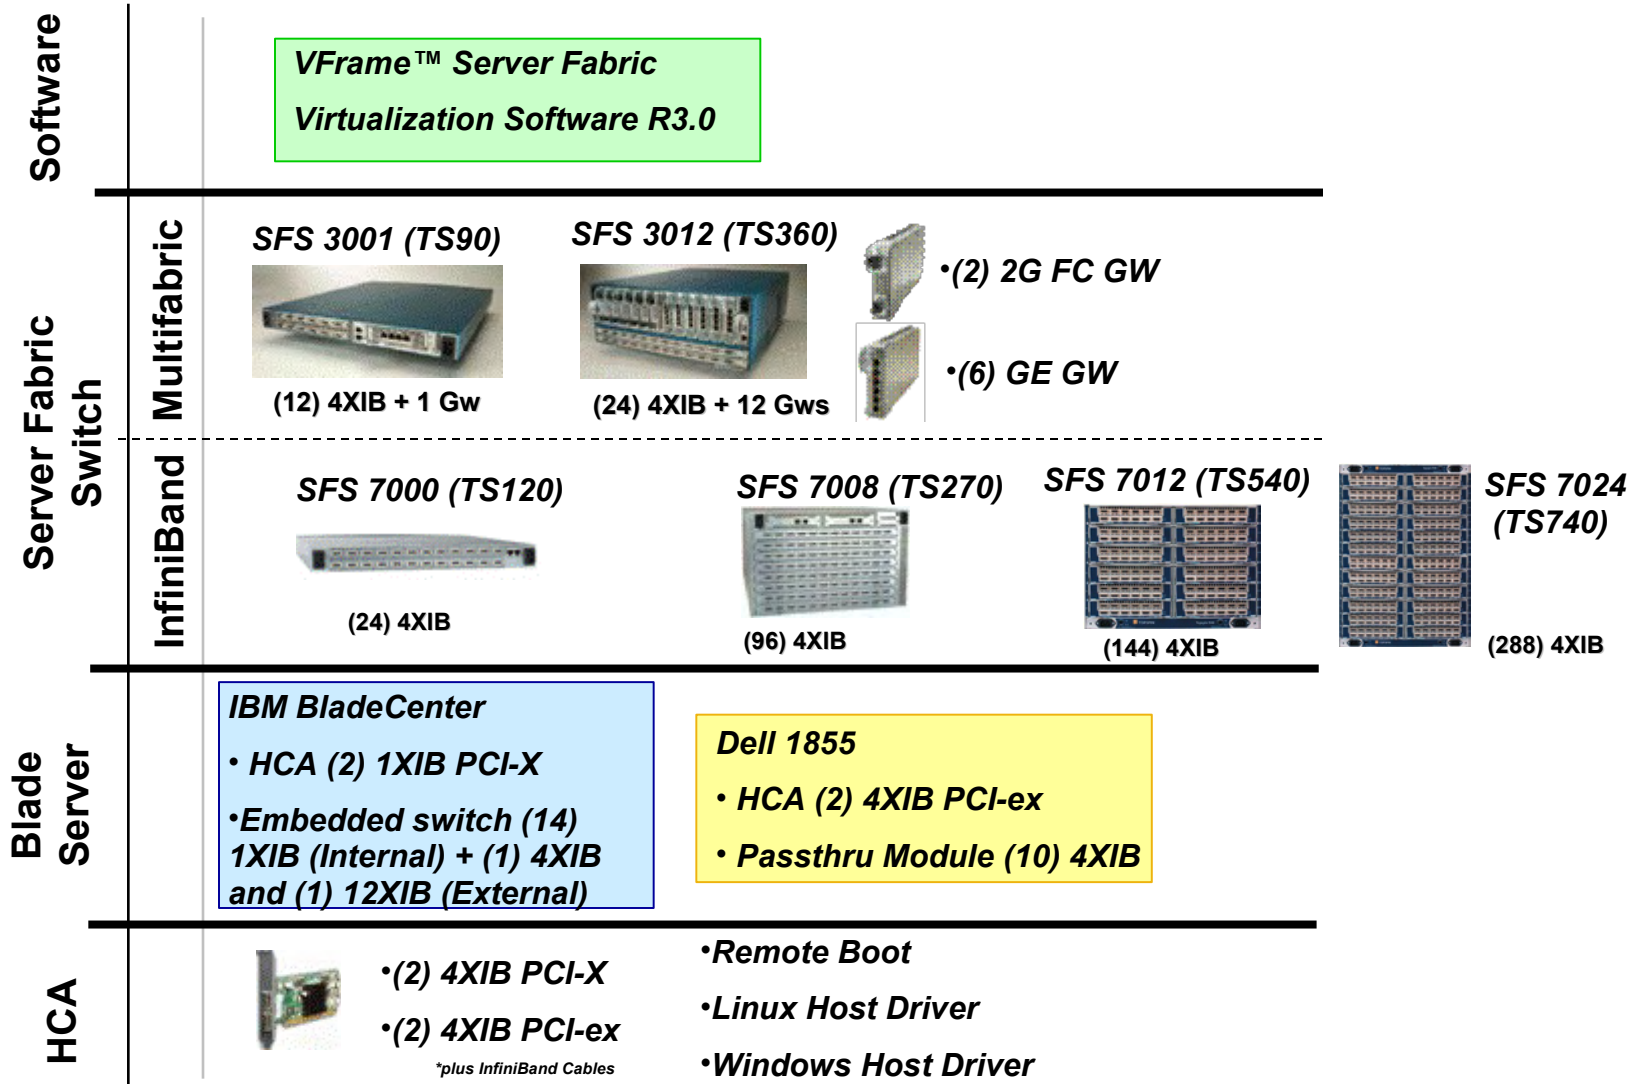

Cisco InfiniBand Blade Switch Modules

For IBM eServer BladeCenter

BladeCenter chassis

IB switch module

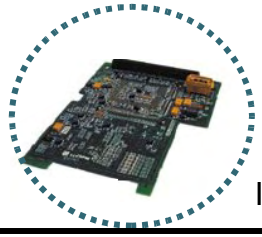

IB daughter card

- Plug one card into each server blade
- Plug one or two switch modules into chassis
- Each server blade gets one or two 1x IB (2.5Gbps) connections
- Target markets: HPC, Multifabric I/O (MFIO), On-Demand data centers
- See IBM Redbook for more details:

http://www.redbooks.ibm.com/redpieces/abstracts/REDP3949.html?Open

For Dell PowerEdge 1855 Blades

IB pass through module

1855 chassis

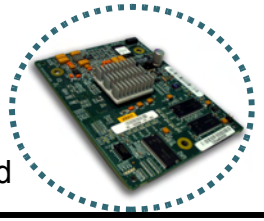

IB daughter card

- Plug one daughter card into each server blade
- Plug one or two pass through modules into chassis
- Each server blade gets one or two 4x (10Gbps) IB connections
- Target markets: HPC, Multifabric I/O (MFIO), Scalable Enterprise centers

Case Study: Leading Research Facility
High Performance Computing Cluster

- **Application:**
  - **– High Performance Computing Cluster**
  - **– Compute time outsourced to Commercial Enterprises (major oil & gas)**
- **Environment:**
  - **– 520 Dell Servers**
  - **– 3:1 Blocking ratio**
  - **– 6x SFS 7008 (TS270)**
  - **– 29x SFS 7000 (TS120)**
- **Benefits**
  - **– Compelling Price & Performance**
  - **– Measured MPI latency 5.2µs**

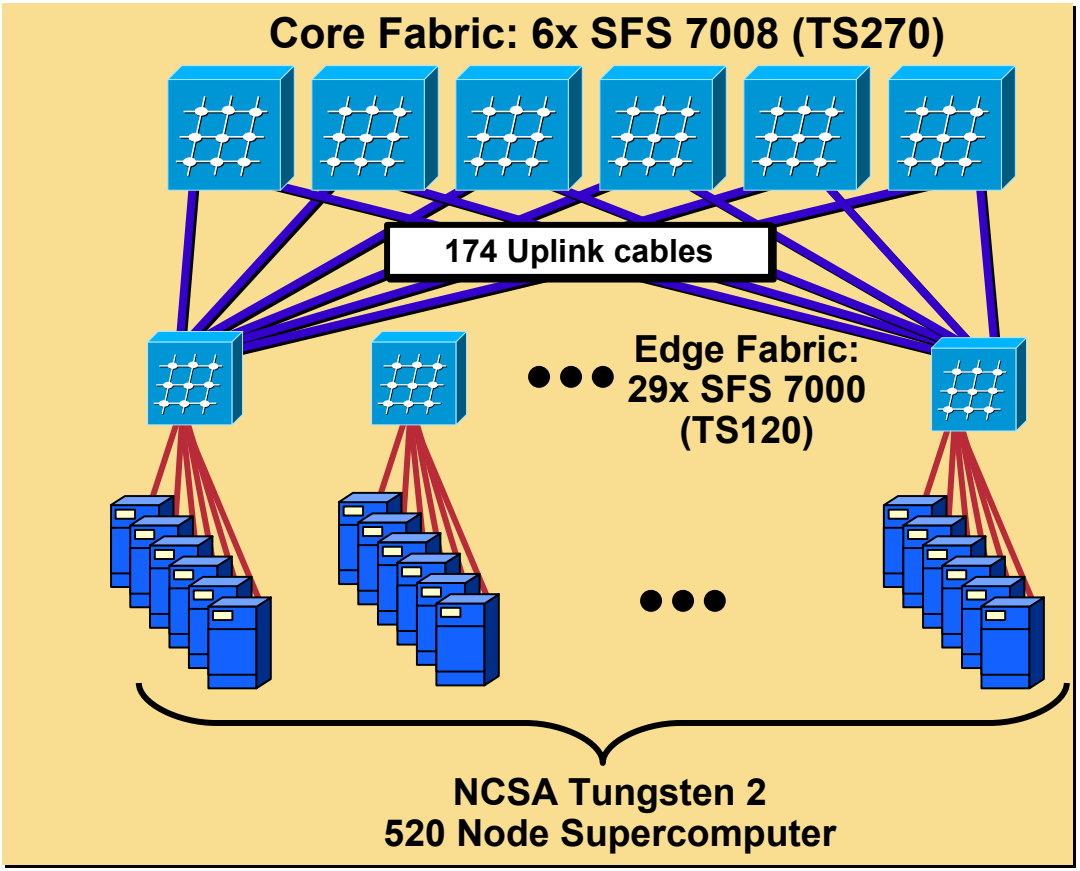

Case Study: Large Wall Street Bank
Enterprise Grid Computing

- **Application:**
  - Replace proprietary platforms with standards-based components
  - Build scalable “on-demand” compute grid for financial applications
- **Environment:**
  - 500+ Intel Servers per slice
  - Topspin Server Switch with Ethernet and Fibre Channel Gateways
  - Hitachi RAID Storage
  - SAN Switches
  - Ethernet Switches
- **Benefits:**
  - 20X Price/Performance Improvement over four years
  - 30-50% Application Performance Improvement
  - Standards-based solution for on-demand computing
  - Environment that scales using 500-node building blocks

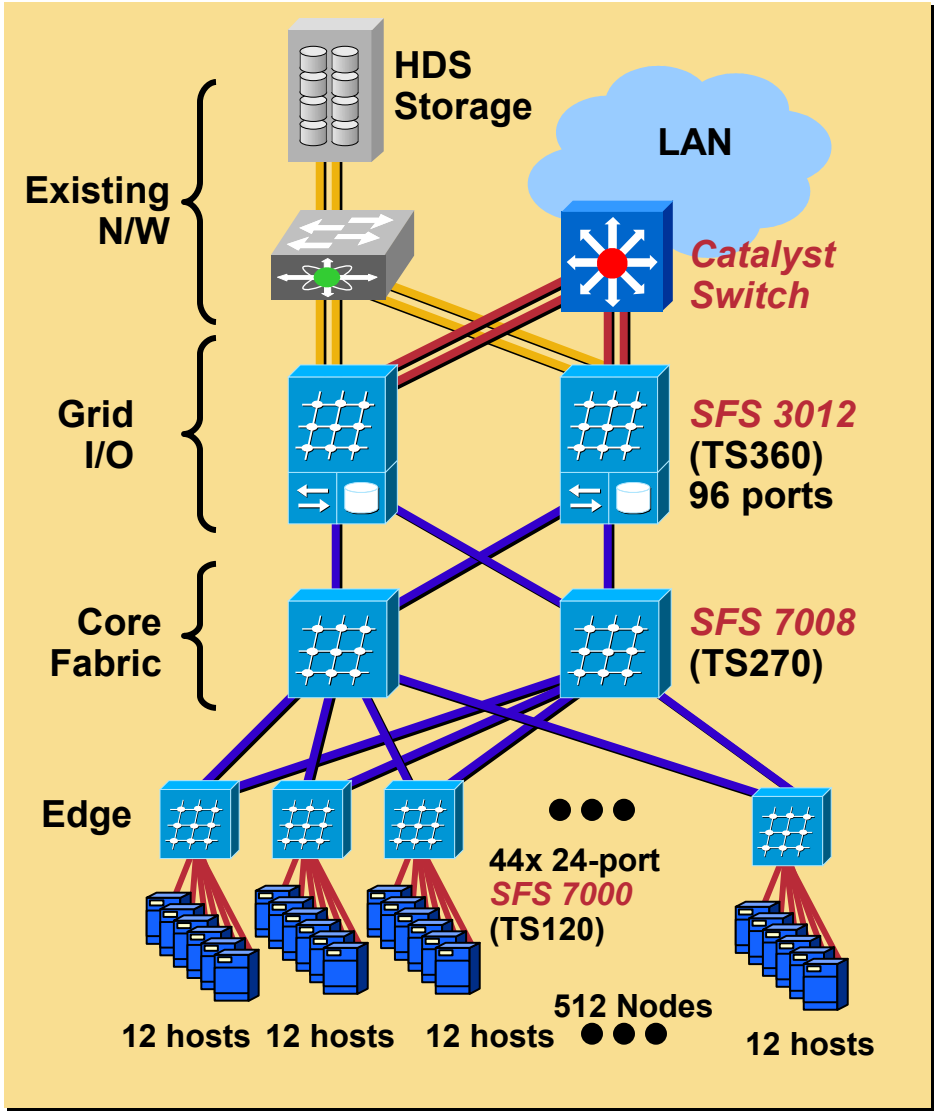

Case Study: Major System Vendor
Utility Computing Service

- Application:
  - Build scalable “on-demand” compute service for enterprise customers (license $/CPU)
  - Key initiatives around Financial Services and Energy verticals
- Environment
  - 1024x Sun V20z Nodes
  - 34 TS270 Server Fabric Switches
  - Non-Blocking CLOS Network
  - 8 TS360’s with Gateways
  - Sun Storage
  - Enterprise-Class Reliability
- Benefits:
  - Ability to outsource computing services to many customers with common infrastructure

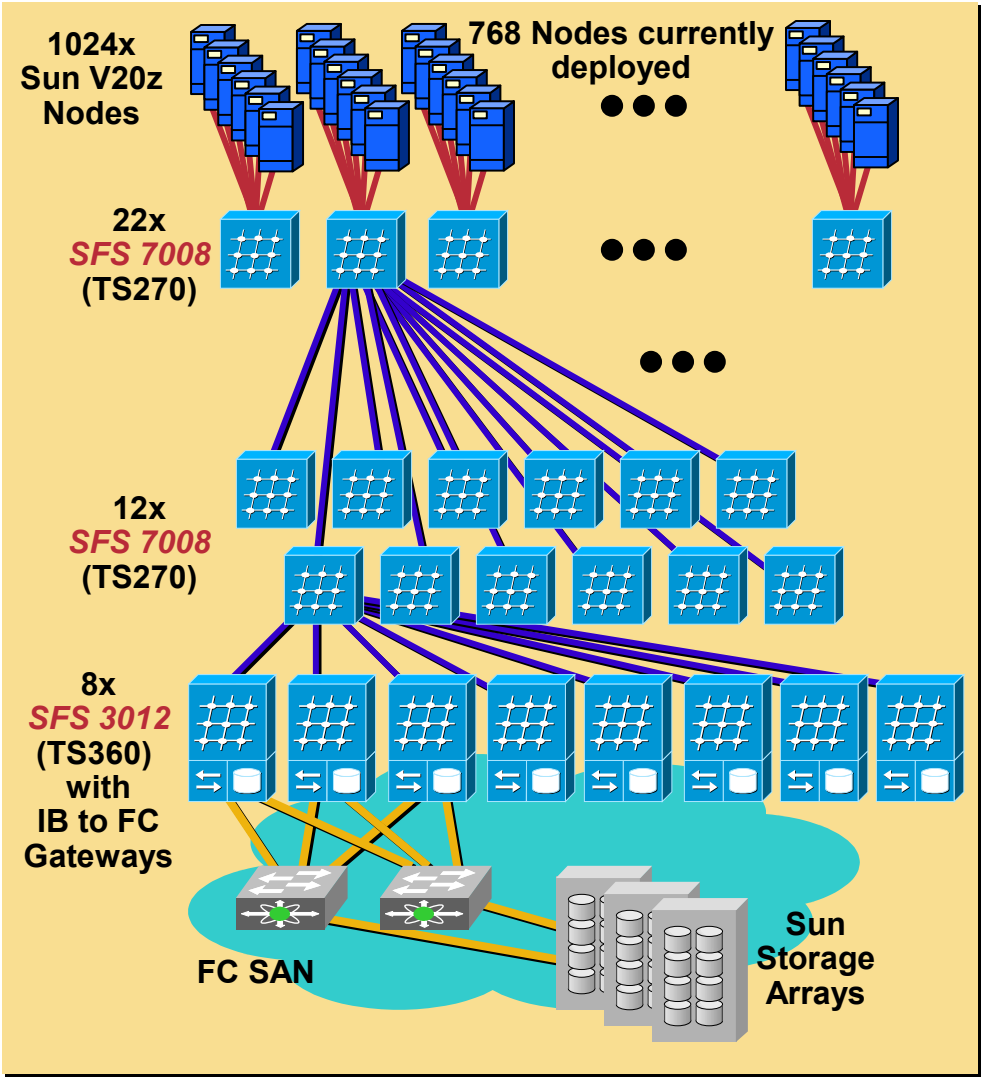

Case Study- Large Government Lab
Worlds 2nd Largest Super Computer

- Application:
  - High Performance SuperComputing Cluster
- Environment:
  - 4096 Dell Servers
  - 50% Blocking Ratio
  - 8 TS 740s
  - 256 TS120s
- Benefits:
  - Compelling Price/Performance
  - Largest Cluster Ever Built (by approx. 2X)
  - Expected to be 2nd Largest Supercomputer in the world

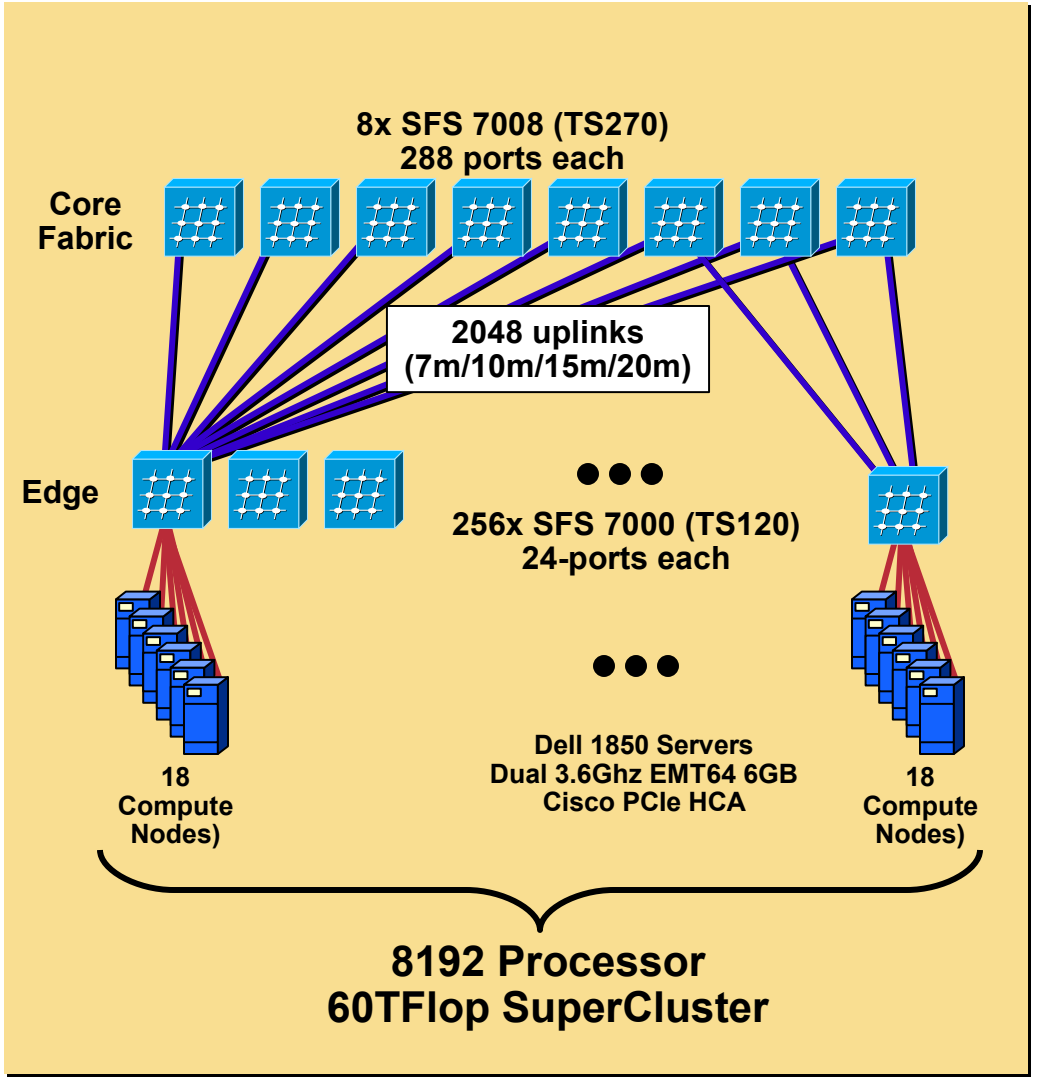

Leading UK Telecom Provider
Oracle 10g Deployment (BT)

- Broadband billing application
- 12 identical regional deployments
- Running Oracle 10g
- No single-point of failure
- Each server has a single HCA, with ports dual connected to two TS90s

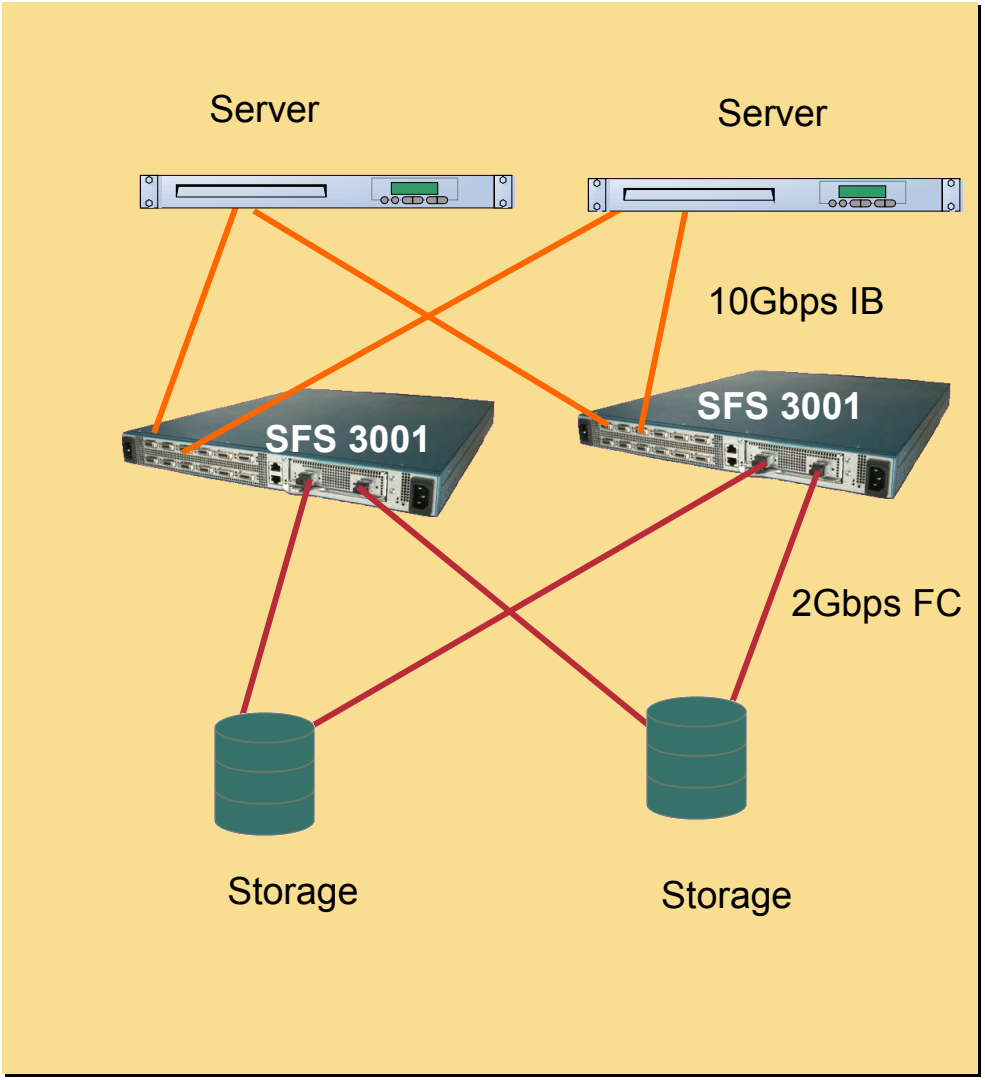

Sandia National Labs – 4096 Nodes Cluster

- Application:
  - High Performance SuperComputing Cluster
- Environment:
  - 4096 Dell Servers
  - 50% Blocking Ratio
  - 8 SFS 7048
  - 256 SFS 7000's
- Benefits:
  - Compelling Price/Performance
  - Largest IB Cluster ever built
  - Expected to be 3rd Largest Supercomputer in the world

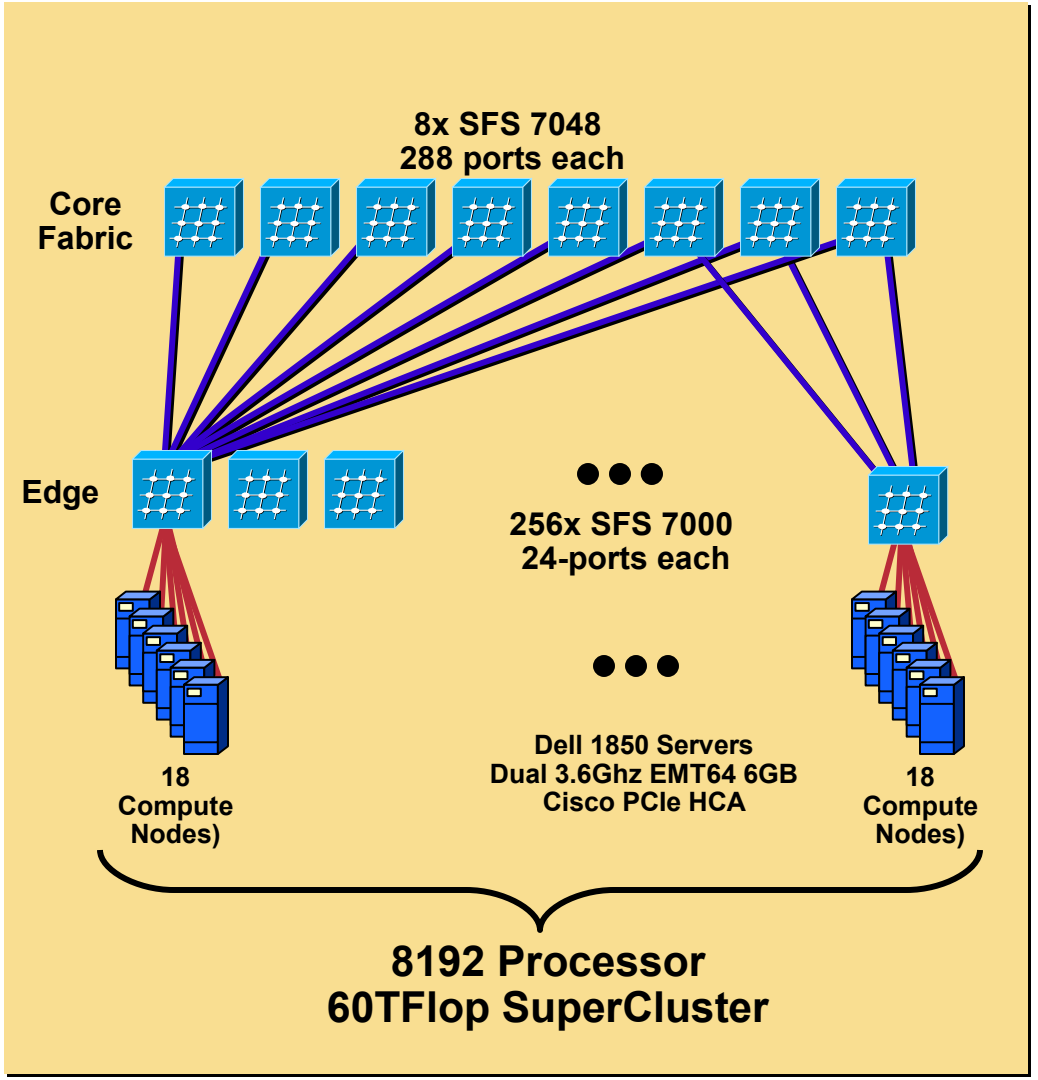

Oracle 10g: Broad Scope of IB Benefits

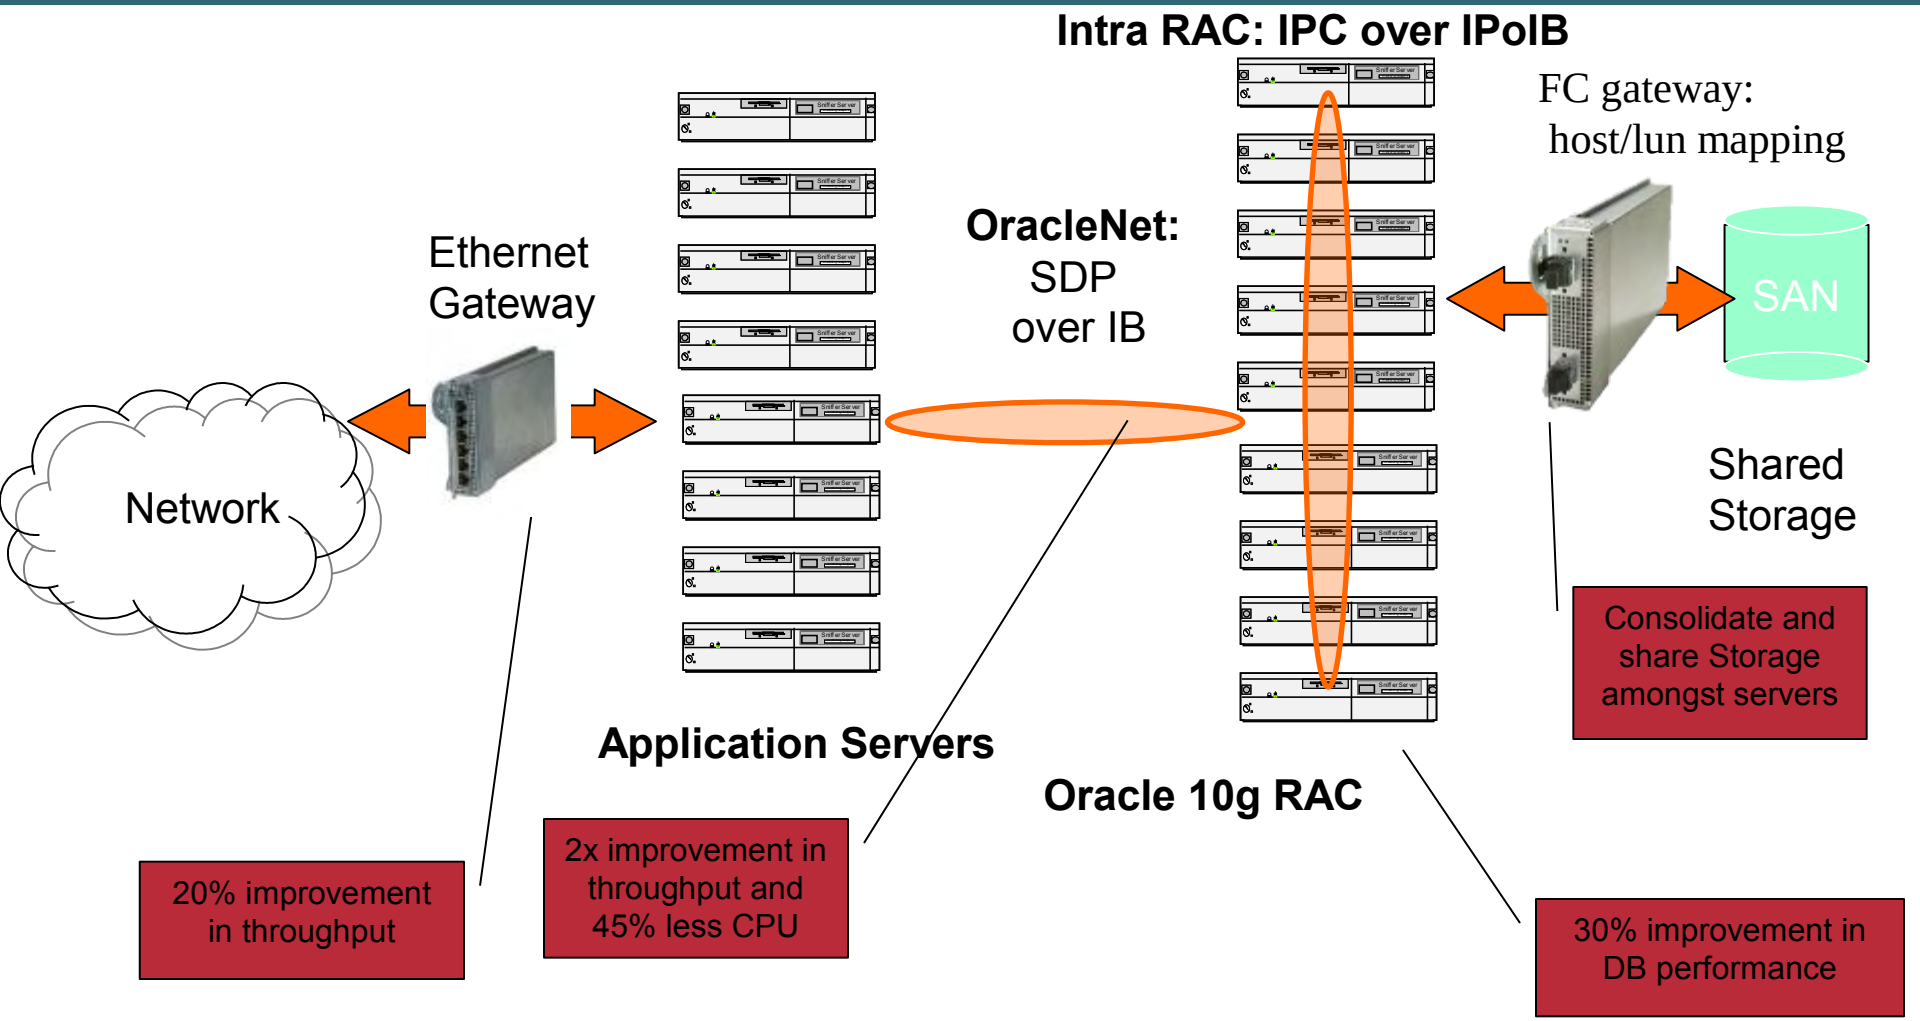

Bio-Informatics Cluster: 1,066 Node Supercomputer

1,066 Fully Non-Blocking Fault Tolerant IB Cluster

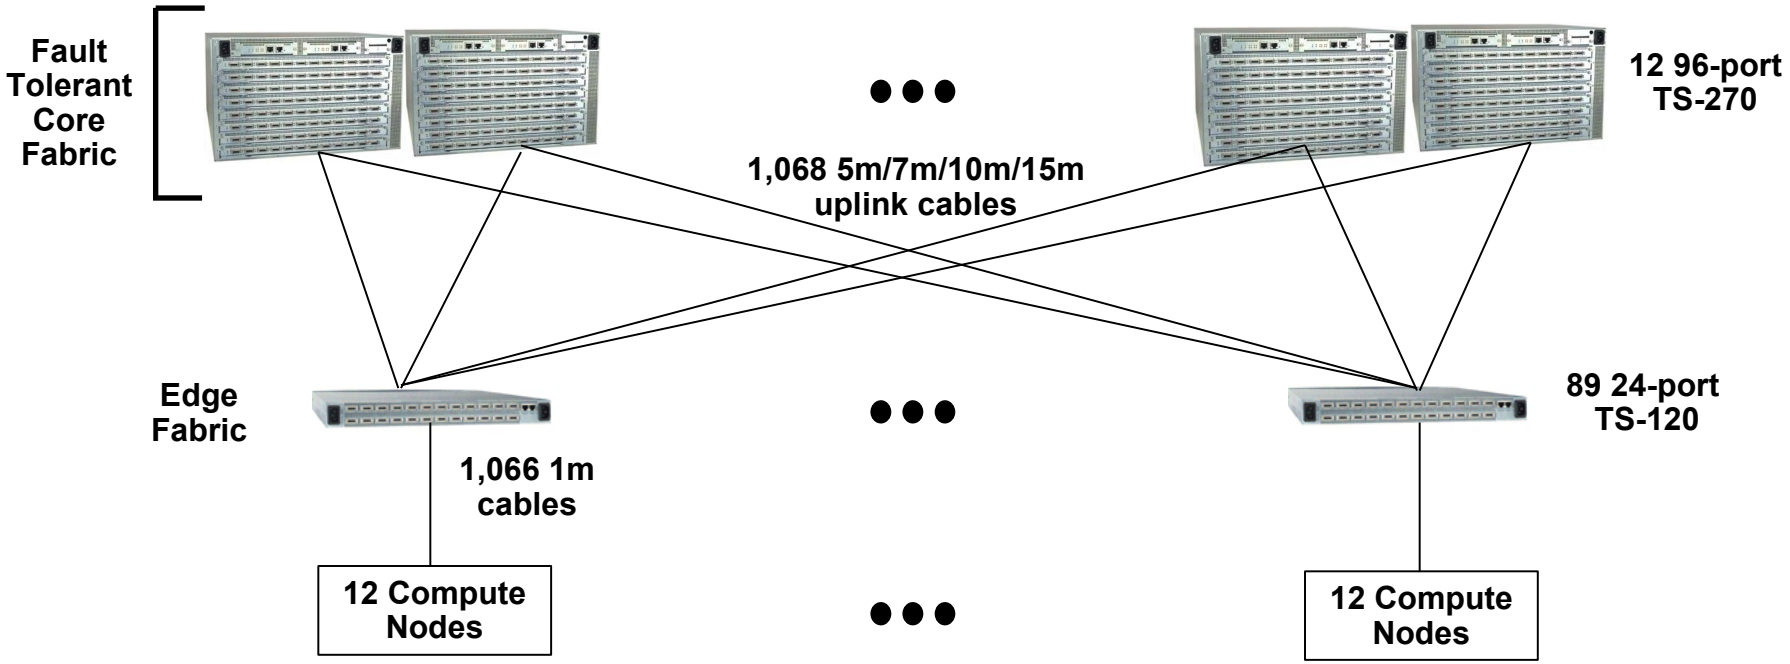

Key decision factors:

- Topspin benchmarked and tuned customer MPI application
- Best operational experience with large clusters – best references
- “Rapid Service” architecture proved 2-min vs. 2-day MTTR.

Cisco InfiniBand Landscape

Vendors working with Cisco Server Switches

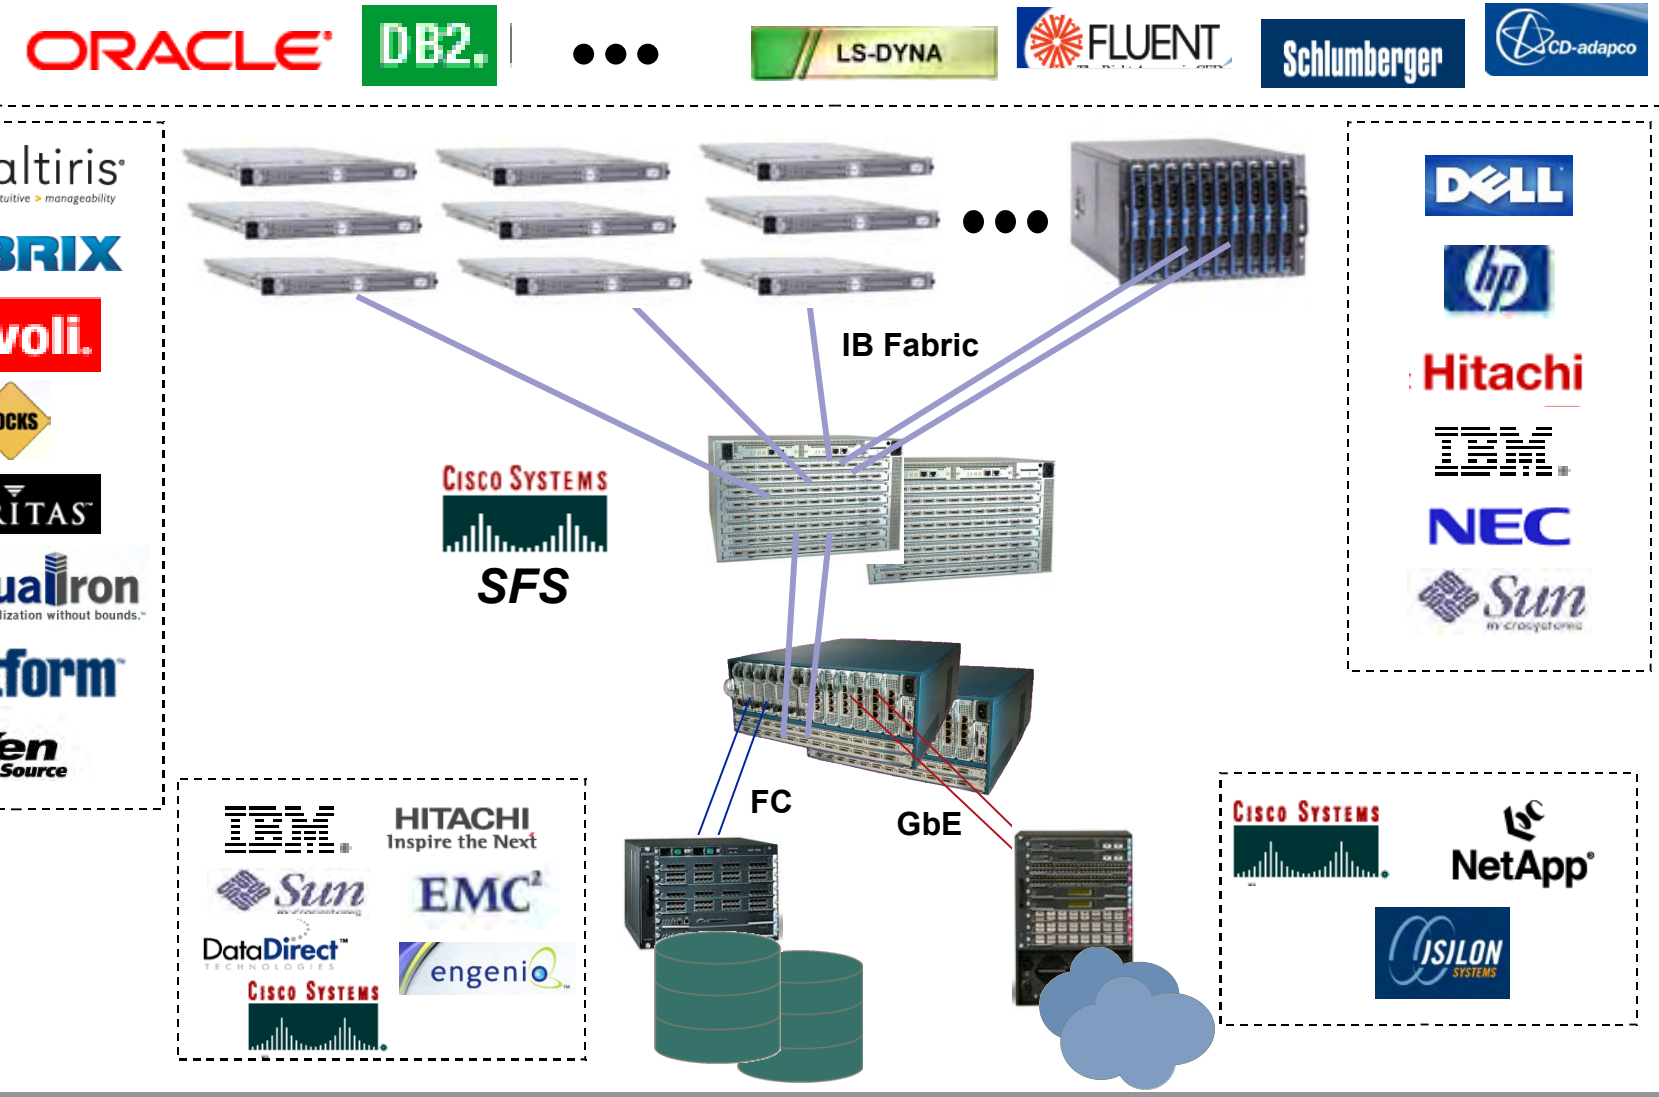

Topspin and Top Tier Server Vendors

“IBM and Topspin Communications Forge Key Agreement”

“Sun and Topspin Partner to Deliver Grid Computing Solutions”

“Dell Adds Topspin Switches to High-Performance Computing Clusters”

“Topspin Selected As NEC's Strategic InfiniBand Technology Provider”

“HP To Leverage Topspin Technology”

Who Owns the Datacenter?

I own

Applications

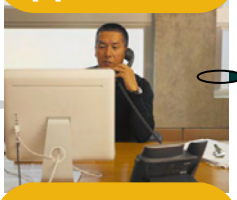

THE NETWORK

Network

I own

Server

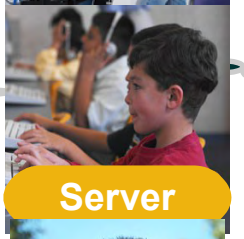

I own

Storage

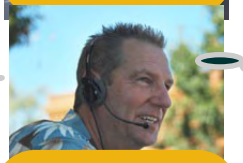

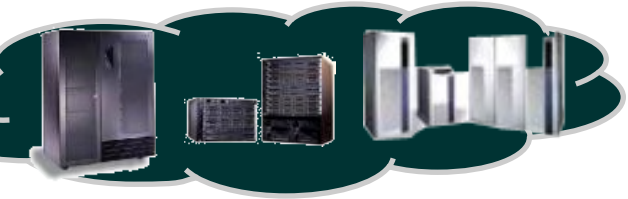

I own

Mainframe

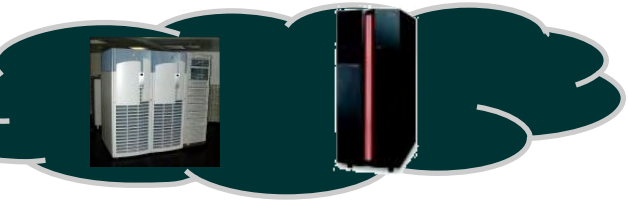

Thank You!