



Data Center Network Architecture

Bandit Kiatsingnakorn
Systems Engineer
bandit@cisco.com

Data Center Architecture Evolution




Data Center Architectural Evolution

Phase I – Mainframe ‘Fortress’

Cisco.com

Mainframe Centricity



Centralized Computing Model

Closed System and Application Development Environment

Very Expensive Hardware and Ongoing Services

Largely a Single-Vendor System/Solution

Cisco Value was in Multi-protocol WAN and LAN connectivity – removing parallel networks and making Mainframe Terminals ‘routed’

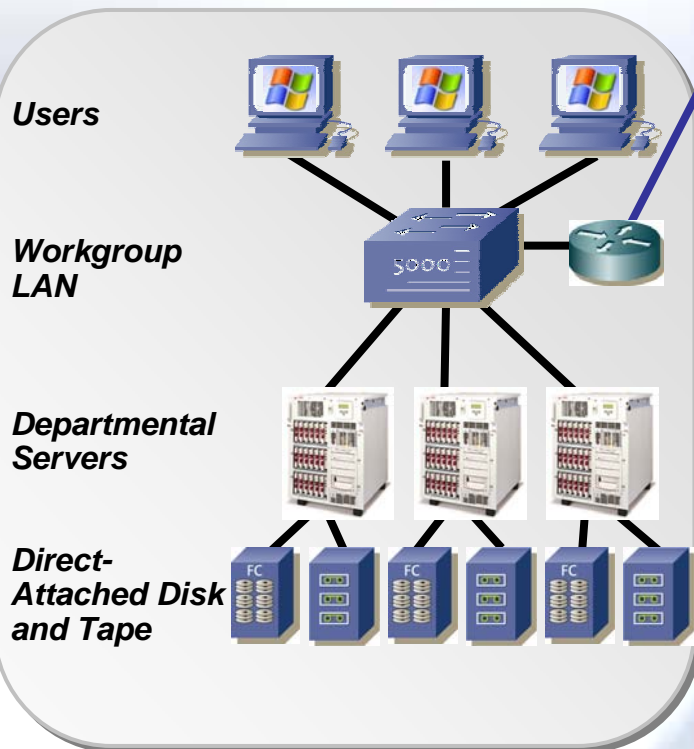
1970’s through mid-1990’s

Data Center Architectural Evolution

Phase II – Departmental Servers and Computing

Cisco.com

Departmental Computing



1995 - 2000

MS-Windows based Desktop computing came of age

Departments deployed their own 'workgroup' or 'departmental' servers based on Novell Netware or MS-Windows NT

Large compute servers cost \$70k - \$100k each-size was dictated by storage requirement

IT maintained the 'legacy' mainframes and datacenters

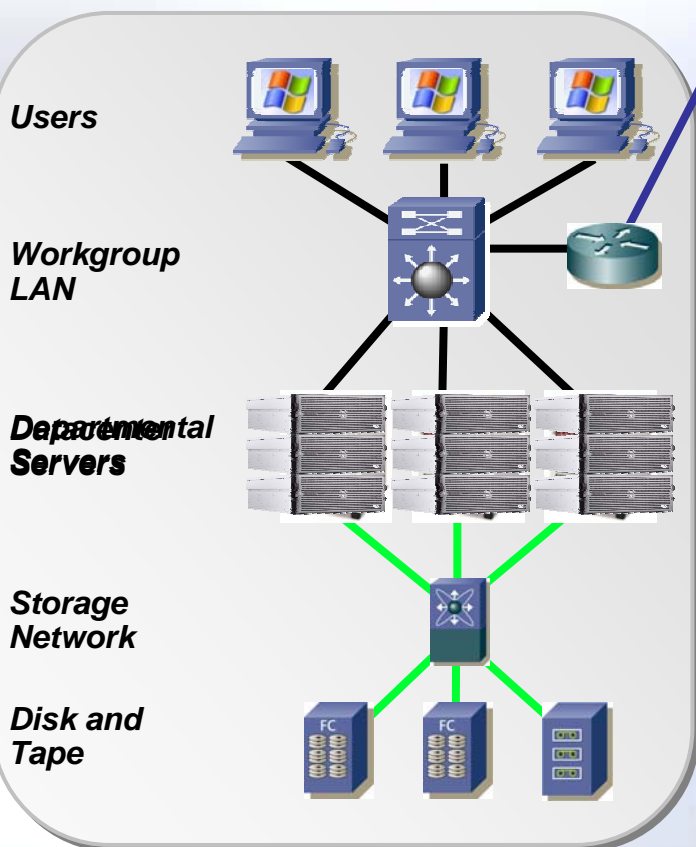
Cisco Value was providing IP/IPX/Appletalk LAN/WAN and workgroup LAN switching – usually Catalyst 5000 series.

Data Center Architectural Evolution


Phase IIIa - SANs and Storage Consolidation

Cisco.com

Data Center Computing



2002 and Ongoing



Storage Networks allowed Disk and Tape consolidation and sharing across multiple servers

Storage Utilization was driven from ~20% to ~40-50%.

Reduced storage capital outlay by >3x from a Direct-Attached Only Model

**Allowed rapid server provisioning via net-boot
10RU Servers (\$100k) were replaced with multiple 4RU servers (\$15k)**



Increased compute capacity, network dependency/utilization, and loosely-coupled application message transfer

Data Center Architectural Evolution

Phase IIIb – Branch Compute Consolidation

Cisco.com

Data Center Computing



The average organization has between 50 – 150 branch office locations with each branch having between 3 and 5 servers

There are over 150 – 750 branch office servers, on average, that can be consolidated to centralized datacenters.

Cisco Wide-Area File Services, integrated with the ISR series routers allows consolidation of branch office servers into centralized data center facilities







Central Archival, Backup, Compliance, with remote Print/File/Login serving through the network

2002 and Ongoing

Data Center Network Architecture

SERVER FARM NETWORK

EMBEDDED NETWORK SERVICES

- DDOS Guard 
- Intrusion Prevention 
- SSL Off-load 
- Firewall Services 
- Server Load Balancing 
- Application Message Services 


SERVER FABRIC NETWORK

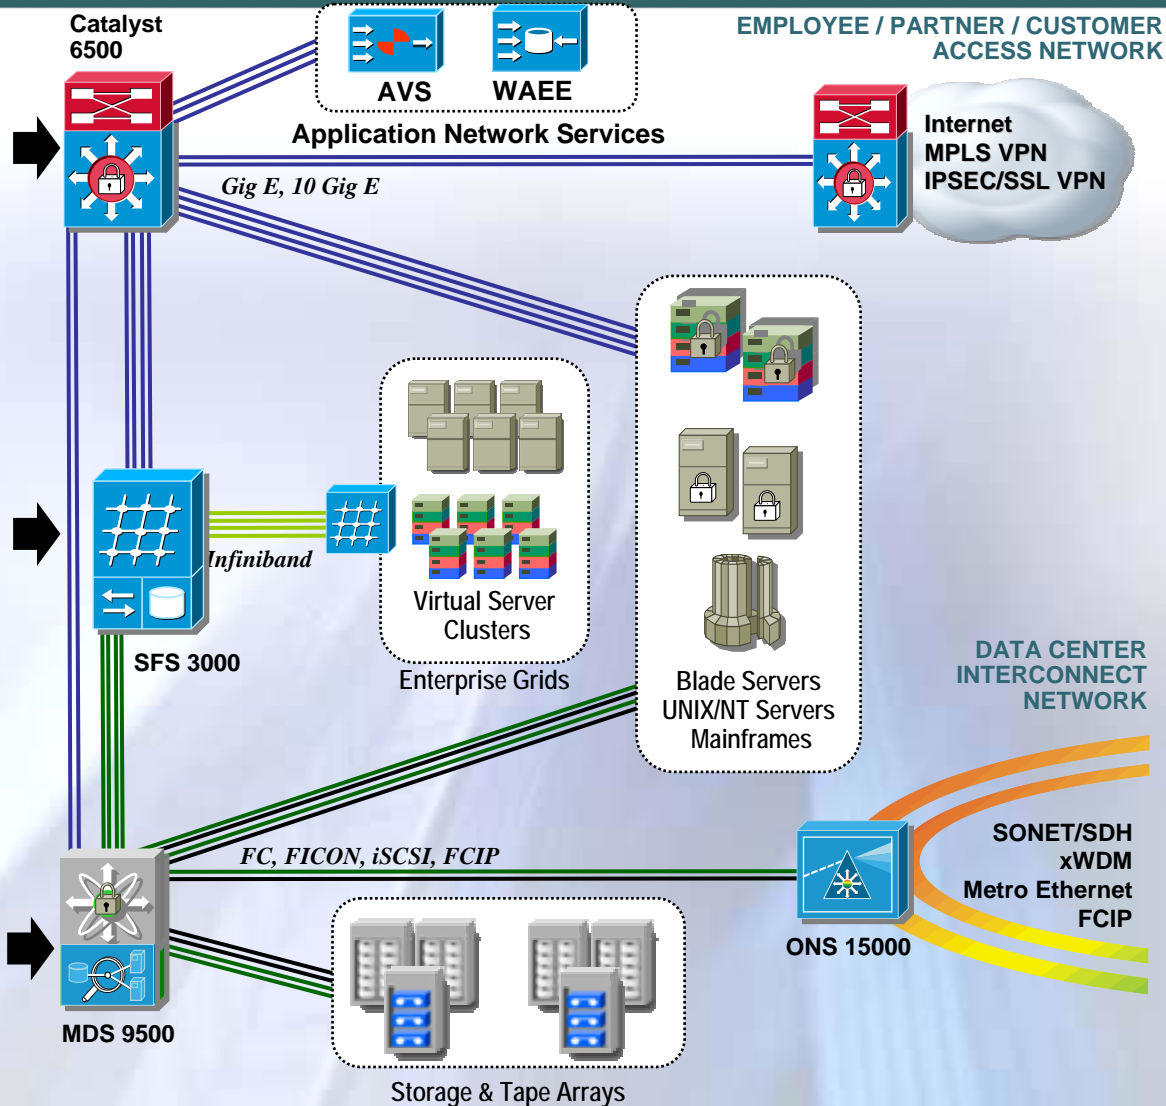
EMBEDDED VIRTUALIZATION SERVICES

- Grid/Utility Computing 
- Low Latency RDMA 
- Server Virtualization 
- Virtual I/O 

STORAGE AREA NETWORK

EMBEDDED STORAGE SERVICES

- Virtual Fabrics 
- Fabric Assisted Applications 
- Storage Virtualization 
- Data Replication Services 



SAN Consolidation




Scalability – Reduced CapEx and Simplified Management


Cisco.com

- **75% efficient port utilization**
 - 64 total ports required – 16 ISLs
- **Large number of elements to manage**
- **Large datacenter footprint**
- **Complex cable management**

64 Port SAN Built w/
16-Port SAN Switches



64 Port SAN Built w/
224-Port Expandable
Directors



- **100% efficient port utilization**
 - 64 ports required – no ISLs
- **Small number of elements to manage**
- **Smaller datacenter footprint**
- **Simplified cable management**


Case Study – Major Insurance Company


SAN Consolidation

Cisco.com

Customer Reference

- One of the largest insurance and financial services companies in the world
- Migrated storage infrastructure which includes several hundred Terabytes from several SAN islands to a consolidated MDS 9000-based SAN designed for availability, recoverability, and growth



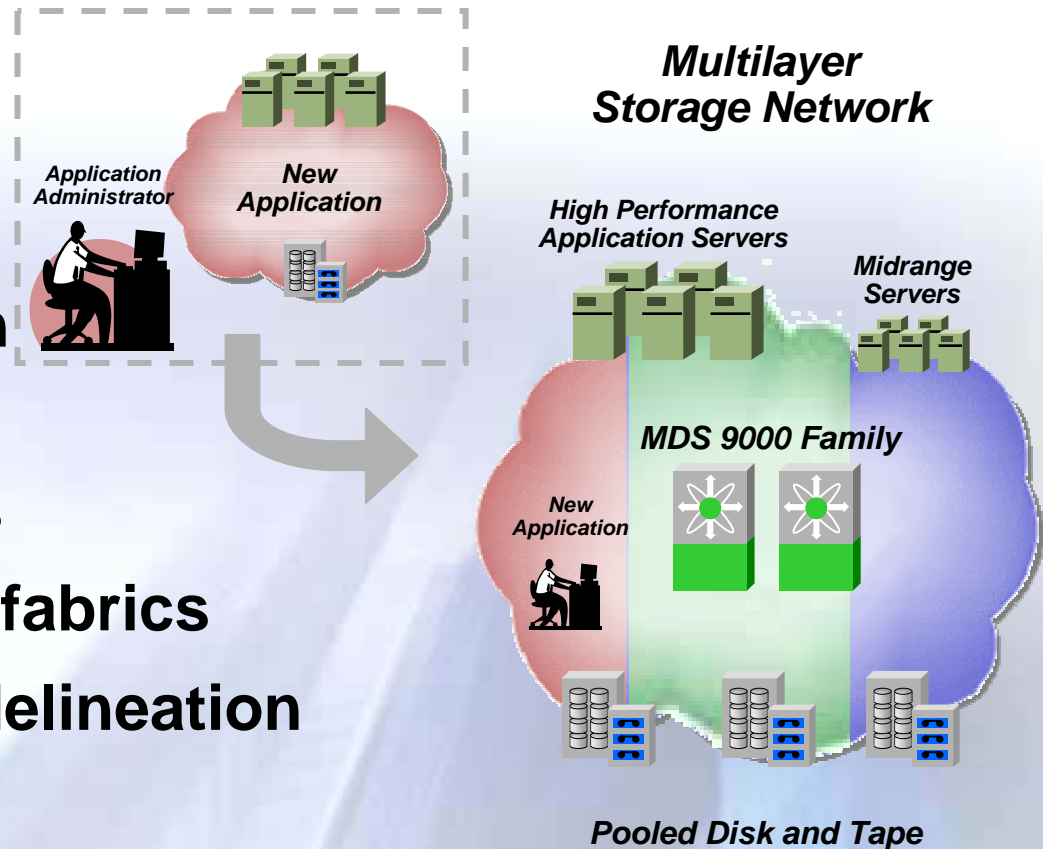
- 
- Converted 24 (competitive) fabrics to 4 fabrics over two production data centers
 - Consolidated 102 legacy switches to 20 MDS directors
 - Completed project in 90 days



Fabric Consolidation - Cisco Solution Foundation

Cisco.com

- Provide cost effective solution retaining benefits of SAN islands
- Flexible provisioning with port-based granularity
- Multi-protocol extensions
- Central administration of fabrics
- Secure hardware-based delineation between virtual fabrics
- Fabrics are added and scaled dynamically
- Enabled through Cisco's VSAN technology

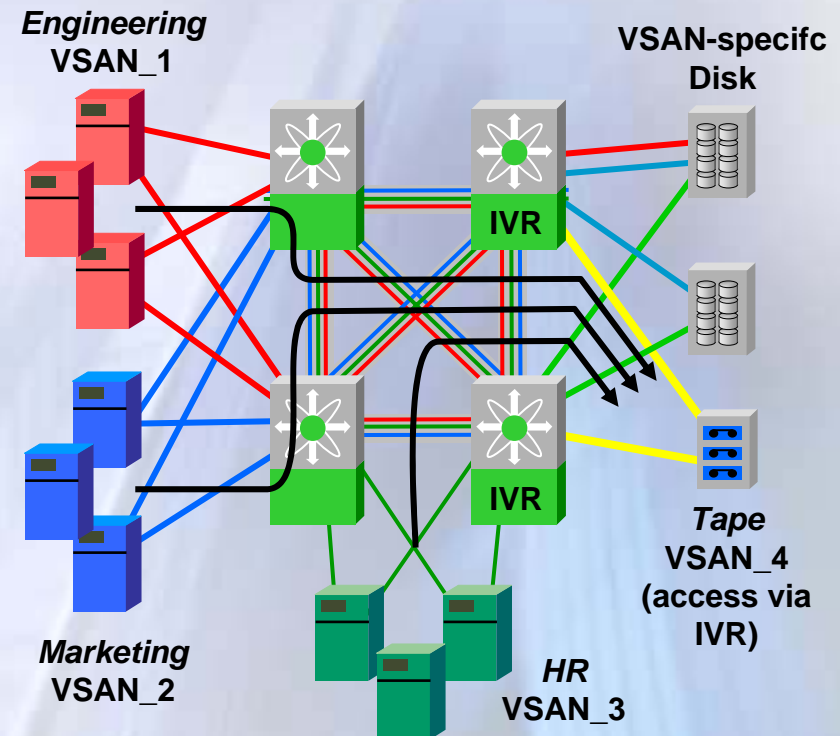
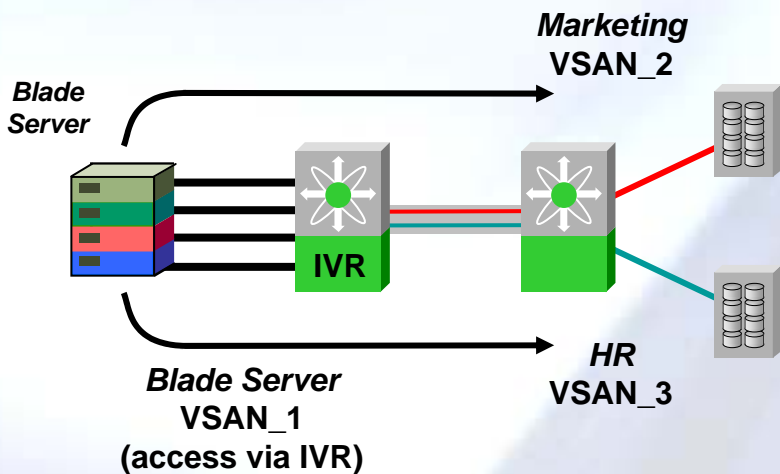


Inter-VSAN Routing (IVR)

Sharing Resources Across VSANs

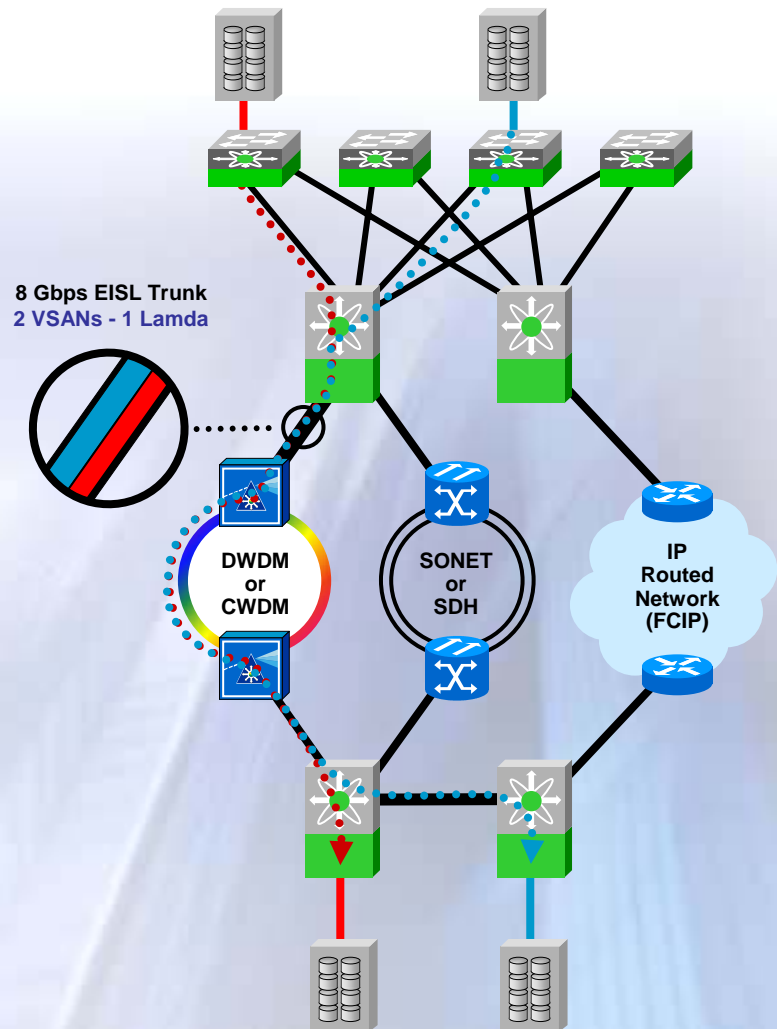
Cisco.com

- Allows sharing of centralized resources such as tape and disks across VSANs – without merging separate fabrics (VSANs)
 - Provides high fabric resiliency and VSAN-based manageability
 - Distributed, scalable, and highly resilient architecture
- Transparent to third-party switches
- Enables blade-per-VSAN architecture for blade servers



VSANs Allow Sharing of DR Facilities

- Virtual SANs can be carried between data centers over various transport facilities
- Cost savings through consolidation of DR facilities while still maintaining isolation
- Various wide and metro area facilities can be used securely: FCIP (PoS, ATM, Metro Ethernet), SONET, DWDM, or CDWM, etc.
- Cisco MDS 9000 family provides traffic statistics per VSAN for chargeback possibility
- Full fabric discovery per-VSAN through Cisco Fabric Manager



MDS 9000 Business Continuance & Disaster Recovery



SAN Extension for Business Continuance

Distance Summary with MDS SAN Switch

Cisco.com



Data Center Campus Metro Regional National Global

Optical

FC over Dark Fiber



250km
255 BB_Credits at 2Gbps

500km
255 BB_Credits at 1Gbps

FC over CWDM




Using BB_Credit spoofing on ONS15454 SL Linecard

3500 BB_Credits with SAN-OS 2.0 & MPS-14/2 → 3500km @2G


FC over DWDM

ONS15530, ONS15540, ONS15454




FC over SONET/SDH

ONS15454



IP

MDS9000 FCIP

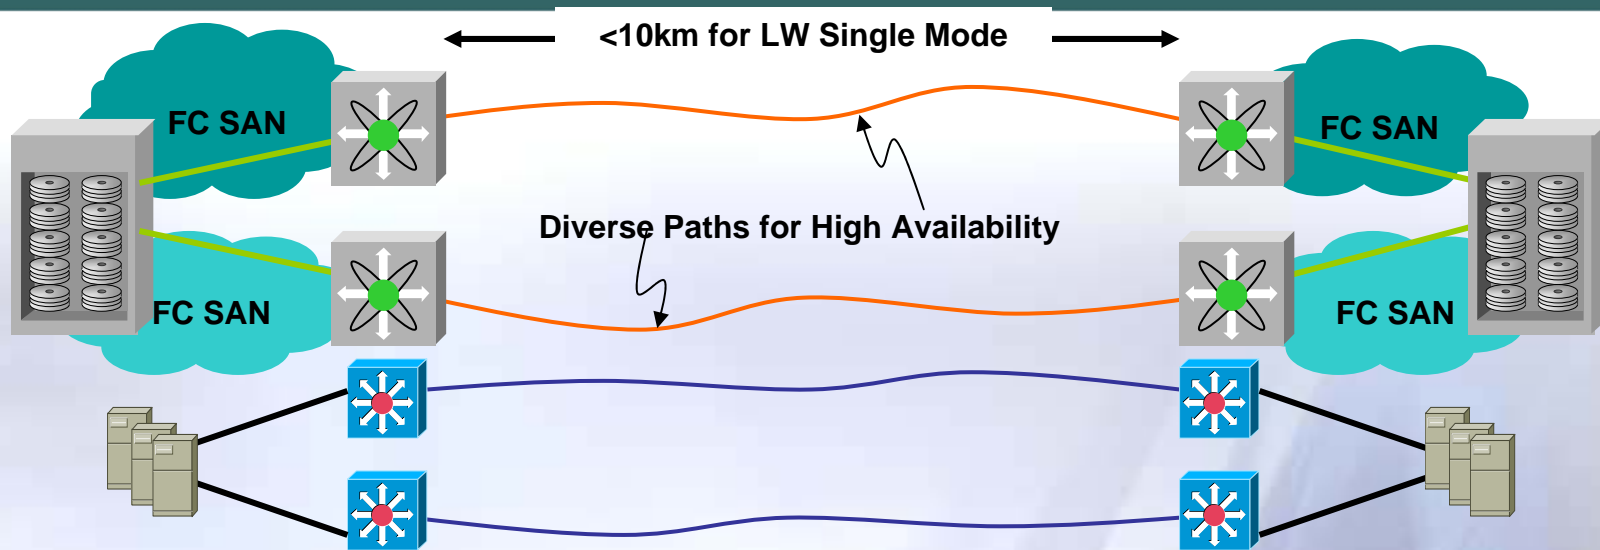


Conversion costs 75us per end

Dark Fiber

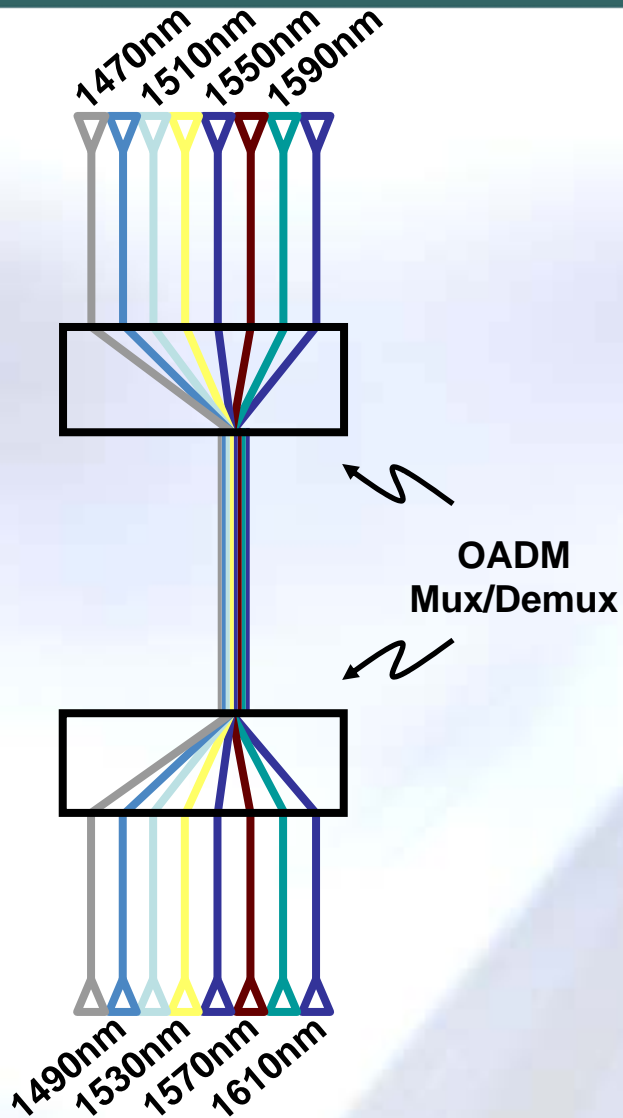
SAN Extension & Ethernet Extension

Cisco.com



- **Single 1 or 2 Gbps link per fiber pair**
 - SW (850nm) 300m over 62.5/125µm Multimode
 - SW (850nm) 500m over 50/125µm Multimode
 - LW (1310nm) 10km over 9/125µm Single Mode
- **“Client Protection”—ULP (SAN or Application) responsible for failover protection**

Coarse Wavelength Division Multiplexing (CWDM)

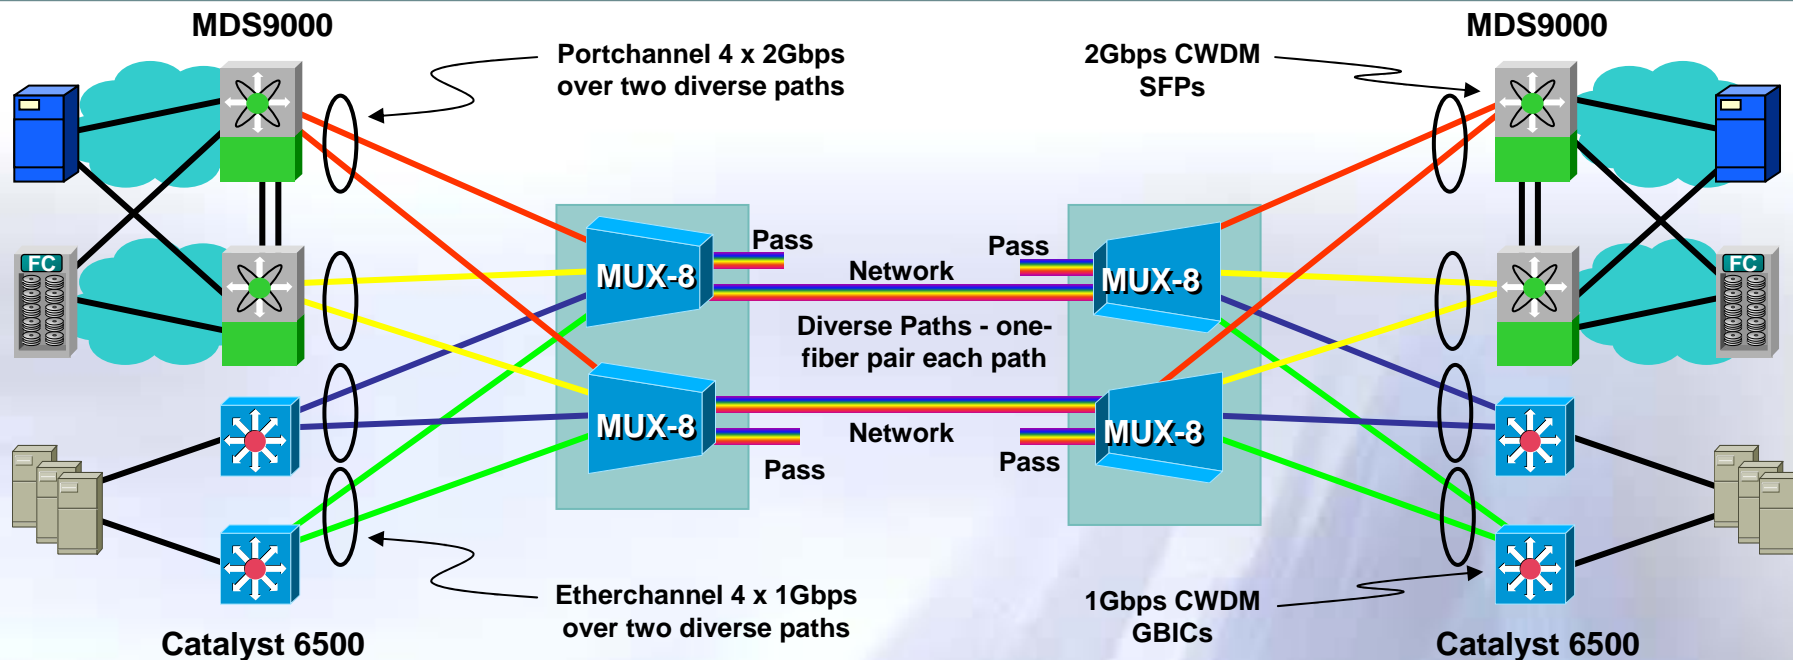


- 8-channel WDM at 20nm spacing (cf DWDM at <1nm spacing)
 - 1470, 1490, 1510, 1530, 1550, 1570, 1590, 1610nm
- “Colored” CWDM SFPs (or GBICs) used in FC switches (no transponder required)
- Optical multiplexing done in CWDM OADM (optical add/drop multiplexer)
 - Passive (unpowered) device—just mirrors and prisms
- Up to 30dB power budget (36dB typical) on SM fiber
 - ~100km point-to-point or ~40km ring

2-Site CWDM Network

SAN Extension & Ethernet Extension

Cisco.com




- **HA Resilience against fiber cut—“client” protection**
 - 4-member Portchannel/Etherchannel—2 x 2 diverse paths
 - Portchannel/Etherchannel appears as single logical link
 - Fiber cut will halve capacity from 8Gbps to 4Gbps (SAN) and 4Gbps to 2Gbps (Ethernet)

DWDM - Protocol Independent and Scalable



Legacy Services



Data Networking



Storage




- Integrate data networking, storage and information streaming
- Connectionless high bandwidth optical infrastructure
- Any packet, any wavelength, any platform

DWDM Protection Alternatives


SAN Extension & Ethernet Extension

Optical Splitter Protection



- Single transponder required
- Protects against Fiber Breaks
- Failover causes Loss of Light (and Fabric Change if only link)

Linecard or Y- cable Protection




- Dual transponders required
 - More expensive than Splitter-based protection
- Transmits over both circuits, but only one accepted

Comprehensive SAN Extension Solution

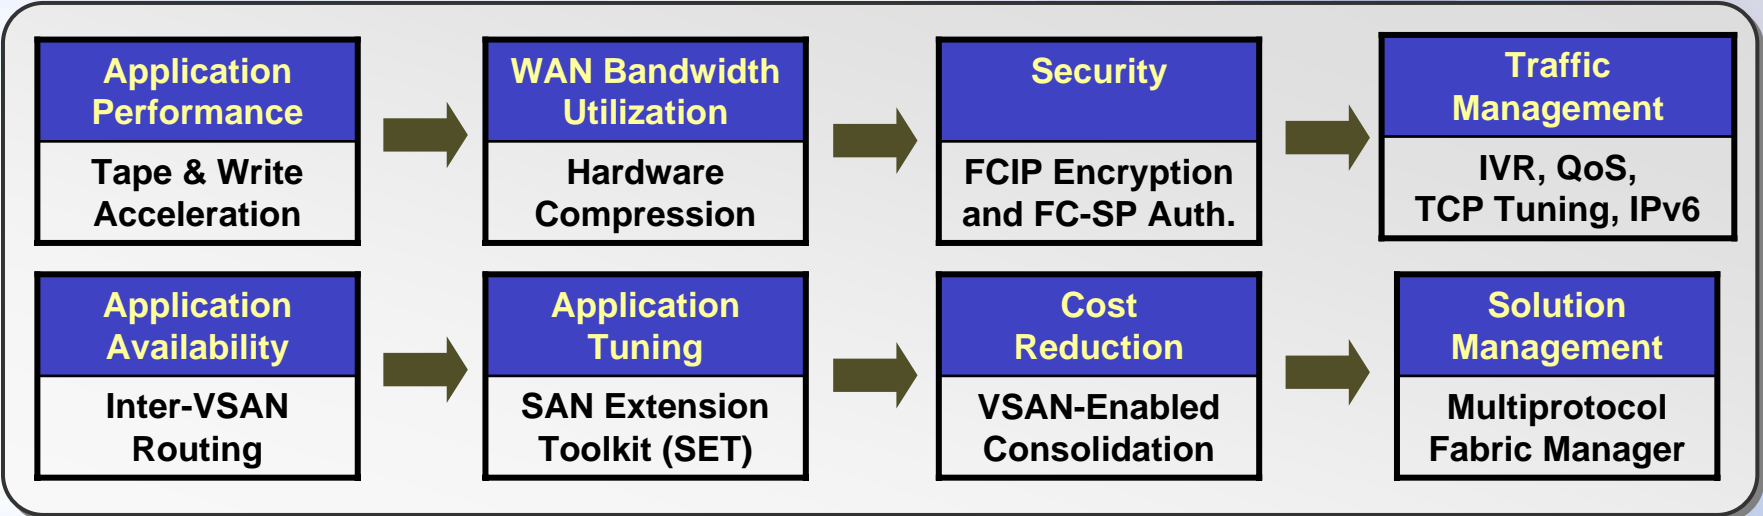

Primary Data Center

MDS 9500 with
MPS-14/2 Module



Backup Data Center

MDS 9216i



Enhances SAN Extension Capability Across All Transports

Cisco.com

- **FC over DWDM/CWDM**
 - ONS 15454, 15540, 15530 – DWDM
 - MDS 9000 – integrated CWDM
- **FC over SONET/SDH**
 - ONS 15454 with SL-Series modules
- **FCIP and FICON over IP**
 - MDS 9000 – integrated FCIP, FICON over IP
 - Cisco 7200, 7400 – FCIP with PAM



Service Virtualization Framework



Today's Enterprise Service Provisioning

A Scale-Out Example

Cisco.com

NetOps ensures Branch connectivity/ Routable Subnet




SecOps checks security policy, expands FW Port Range



SLB Admin Adds Server to Pool



NetOps connects Ethernet cabling, configures VLAN/Port Config



SysAdmin racks new server Loads O/S and Applications



StorageOps configures LUN, maps to Server



StorageOps provisions disk volume and resources




Assume you just want to add one server to a web-farm...

The challenge is one of 'coordination delays'. This type of simple scale-out of an existing serve often takes enterprises 90-days.

New service turn-ups, after the application has been developed, often take 180+ days.


VFrame is designed to eliminate these delays and automate the provisioning of services



Creating Virtual Compute Resources


Cisco.com

- Massive proliferation of server resources



- Server manager configures pools of compute, storage and I/O
- Based on application, process, or business objectives

Intelligent Interconnect Fabric



Resource Pool

Introducing the VFrame Data Center Solution

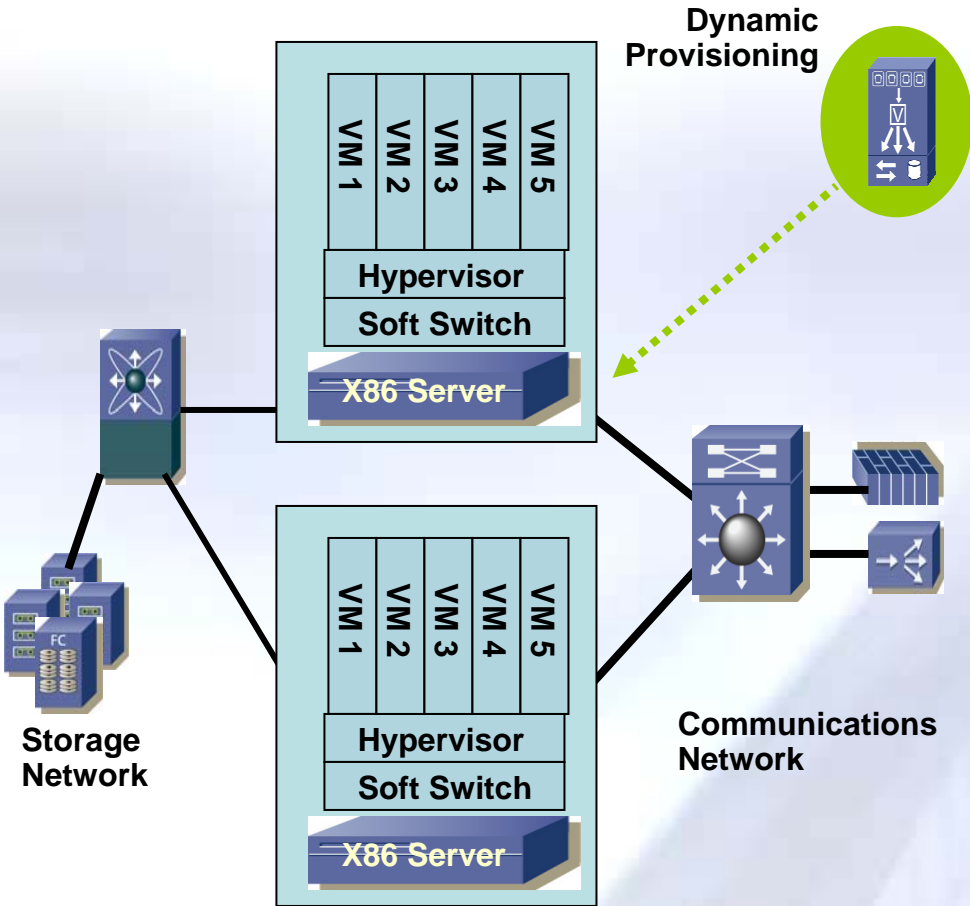
Cisco.com

- **VFrame is a framework service virtualization and provisioning in the data center**
- **Abstracts the physical topology from the policy-based “logical one”**
- **VFrame Components**
 - **VFrame Director – dedicated server which stores policies, logical topology and device information**
 - **VFrame agent – resides on server to feed back information to VFrame director (Windows or Linux support)**
- **Configurable Devices**
 - **Catalyst 6500**
 - **MDS 9500**
 - **Content Service Module**
 - **Firewall Service Module**




Virtualized Server Models


One Server to Many Virtualized Services



Many Servers to one Virtualized Service



Creating Virtual I/O over a Unified Fabric



Fibre Channel to InfiniBand gateway for storage access

- Create 10-Gbps virtual storage pipe to each server

Ethernet to InfiniBand gateway for LAN access


- Creates a virtual GigE IP connection to each server

VFrame Example—Failover

1. Physical server fails

VFrame running on Infiniband-attached Linux Server

JC1



Group 1

Group 2


Standby Pool



4. New server restarts in new group

2. VFrame detects fault

3. VFrame programs SFS to map standby physical server to LUN and virtual I/O



JC1

Changed: Linus Server


To: Linux Server

Joe Childs, 29-Oct-05

VFrame Example— Add Capacity on Demand

JC2

1. Application monitors issue triggers to VFrame



3. Standby servers restart in new group with new mappings

Group 1




Group 2



Standby Pool



2. VFrame programs SFS to add new servers from standby pool



Storage Network



Campus/
WAN/VPN

JC2 Spell out: "App" when there is room
Joe Childs, 29-Oct-05

Application Velocity System



Performance Impedes Web-based Business

Cisco.com



Applications Increasingly Represent the Business

- Pressure to move from staff to software for productivity
- Web promoted process automation across the Extended Enterprise
- Need to serve ALL users, everywhere



Chokepoint Evolves From Bandwidth to Latency

- Developers write for functionality
- Bandwidth issues are more limited
- Individual web page rendering can take 150 network calls

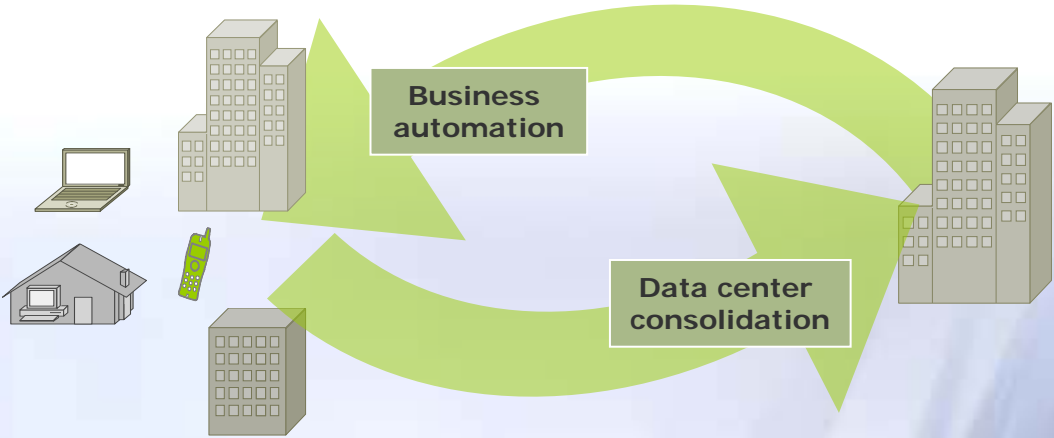


Applications Give Unprecedented Access to Critical Business Data

- Private Data, Confidential, Intellectual property
- Developers write for functionality
- Too expensive to fix all security bugs

The Global 2000 Application Delivery Problem

> **80%** of users across the extended enterprise



< **20%** of users at HQ



Conventional Approaches Fall Short

Network Team


- More bandwidth
- Point products
- Replicate data centers

Operations Team


- More servers
- More management tools
- Re-architect infrastructure

Applications Team

- More testing
- Rewrite applications
- Security patching



Cisco Application Delivery Solutions



Branch/WAN

- Wide Area Application Engine (WAE)
- Integrated Services Router (ISR)
- IOS with NetFlow, NBAR, QoS, IP-SLAs...

Data Center

- Application Velocity System (AVS)
- Content Switches (CSS/CSM)
- Catalyst 6500 switches

Industry Leading Acceleration and Optimization Features

Functional Areas	AVS Acceleration Features
Latency Reduction	<ul style="list-style-type: none">▪ FlashForwarding*▪ Browser TCP multiplexing*▪ PDF download optimization▪ Response redirection control*
Bandwidth Reduction	<ul style="list-style-type: none">▪ GZIP Compression▪ Delta encoding*▪ Dynamic browser caching*▪ Dynamic image optimization▪ Flexible processing rules
Server Offload	<ul style="list-style-type: none">▪ TCP Offload▪ SSL Offload▪ RAM Caching▪ Dynamic caching*▪ Load-based caching*▪ Lazy request evaluation*▪ Single sign-on optimizations▪ XML merging/transformation

AVS Delivers Real-World Value

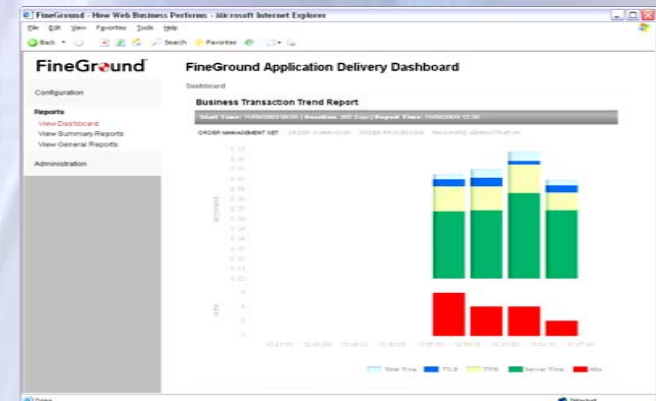
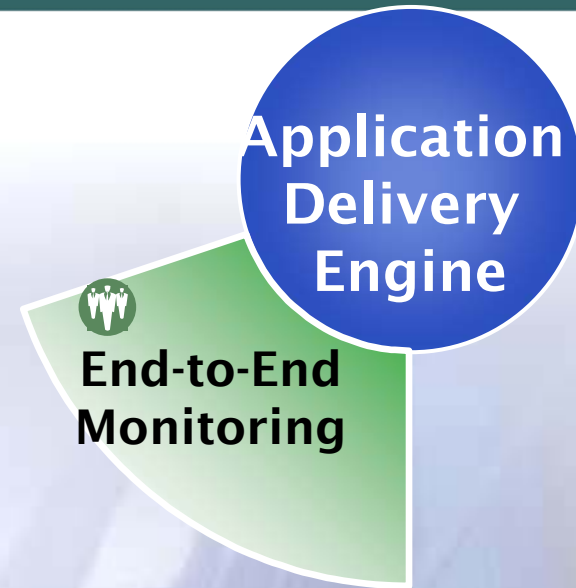
Application	Software	AVS Improvement	Cost of Likely Alternative	Business Impact
Call Center (High tech)	PeopleSoft	↑270%	\$4MM (2 New Overseas Data Centers)	<ul style="list-style-type: none"> Meet support goals with no additional staffing or costs
Purchasing (Manufacturing)	SAP	↑350%	\$5MM (Multiple Overseas Data Centers)	<ul style="list-style-type: none"> Increase procurement automation
Mortgage Origination (Financial)	Custom J2EE: WebSphere	↑300%	\$2MM (Reengineer Apps and Infrastructure)	<ul style="list-style-type: none"> 30% more transactions across same infrastructure
Claims Management (Insurance)	Custom J2EE: WebSphere	↑220%	\$3MM (No Reengineering)	<ul style="list-style-type: none"> Support “zero-footprint” branch
B2B Operations (Retail)	Plumtree	↑350%	\$500K annually (Upgrade 650 Sites)	<ul style="list-style-type: none"> Move all costly paper-based processes online
CRM (Financial)	Siebel	↑290%	\$2.4MM annually (Upgrade 200 U.S. Locations)	<ul style="list-style-type: none"> Immediate jump in CRM usage with Improved account retention

Application Monitoring

AVS AppScope Performance Monitor


Cisco.com

- **End-user response time monitoring**
 - Actual users and transactions
 - Business- and process-level aggregation
 - Full drill-down to page and location
- **“Drop-in” agentless deployment**
 - No changes to application or desktop
 - Data center installation
- **Delivery Dashboard and flexible reporting**
 - Wizard-based transaction builder
 - Support for Enterprise Consoles (BMC, Tivoli, OpenView...)
- **Benefits**
 - End-user visibility
 - First-line problem triage
 - Reduce mean-time-to-repair



In A Bigger Picture: Cisco Application Delivery Solution in the Data Center

Cisco.com



**Application
Velocity
System (FG)**


Breakthrough Impact on User Response Times

Features

- Full application proxy
- Unique optimizations
- Integrated security/monitoring

Benefits

- Dramatic response time improvements
- Transparent deployment



**CSS/CSM/ACE
Content Switch**


The Market-leading L4-7 Switching Platform

Features

- Load balancing
- SSL encryption/decryption
- Session redundancy

Benefits

- Application availability
- Server offload
- Application continuity



**Catalyst
Switch**

The World's Leading LAN Switching Platform

Features

- Highest industry performance
- Modular architecture
- Multi-service approach

Benefits

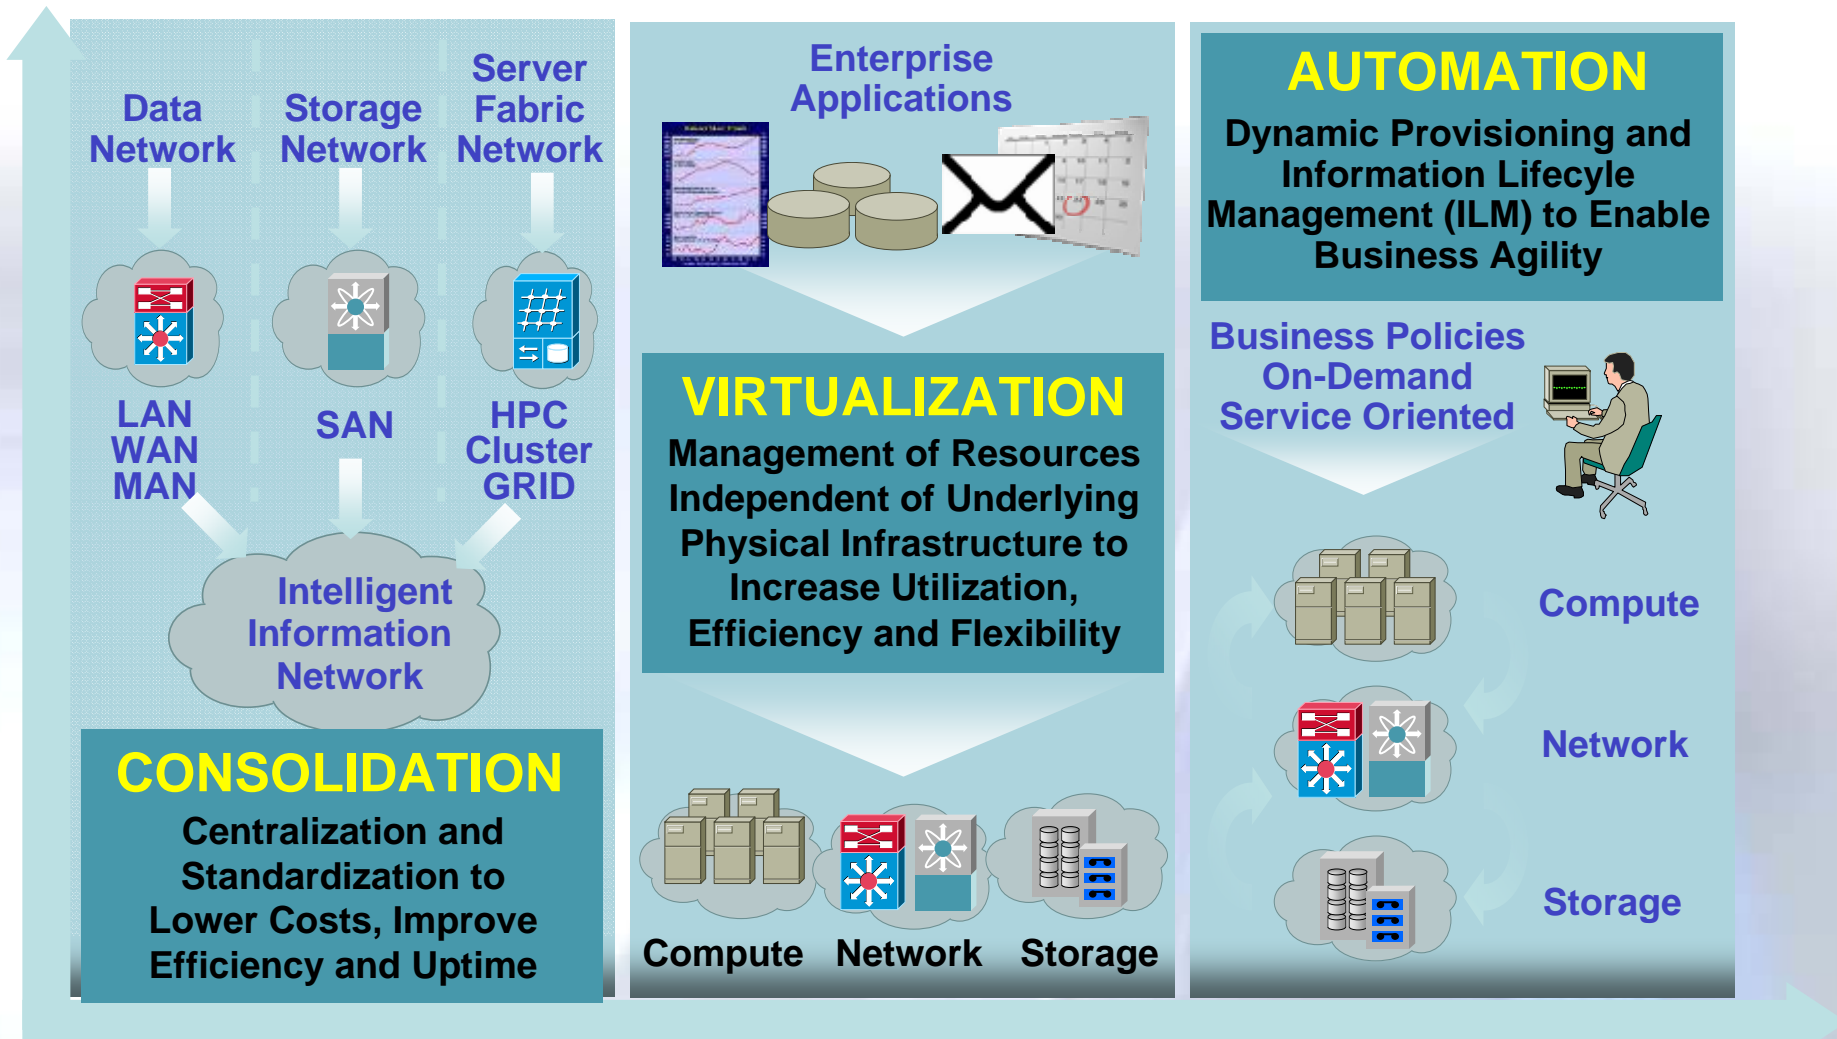
- Largest technology investments
- Single point of aggregation
- World class support

Summary



Cisco Vision in Data Centre Networking

Cisco.com



CISCO SYSTEMS

