An Architectural Approach to Campus Networks

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Enterprise Architectures & The Network

TOGAF

Zachman Framework

IAF - CapGemini

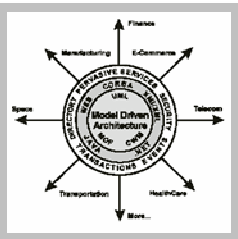

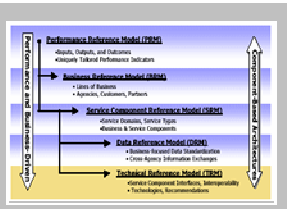

Federal Enterprise Architecture

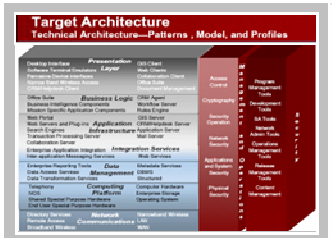

Homeland Security

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

More Than Just Connectivity

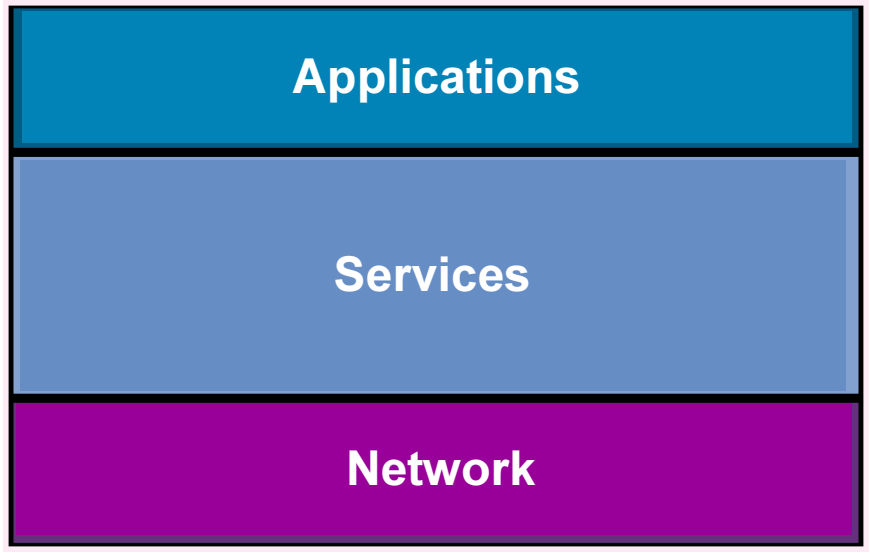

An Architectural Framework

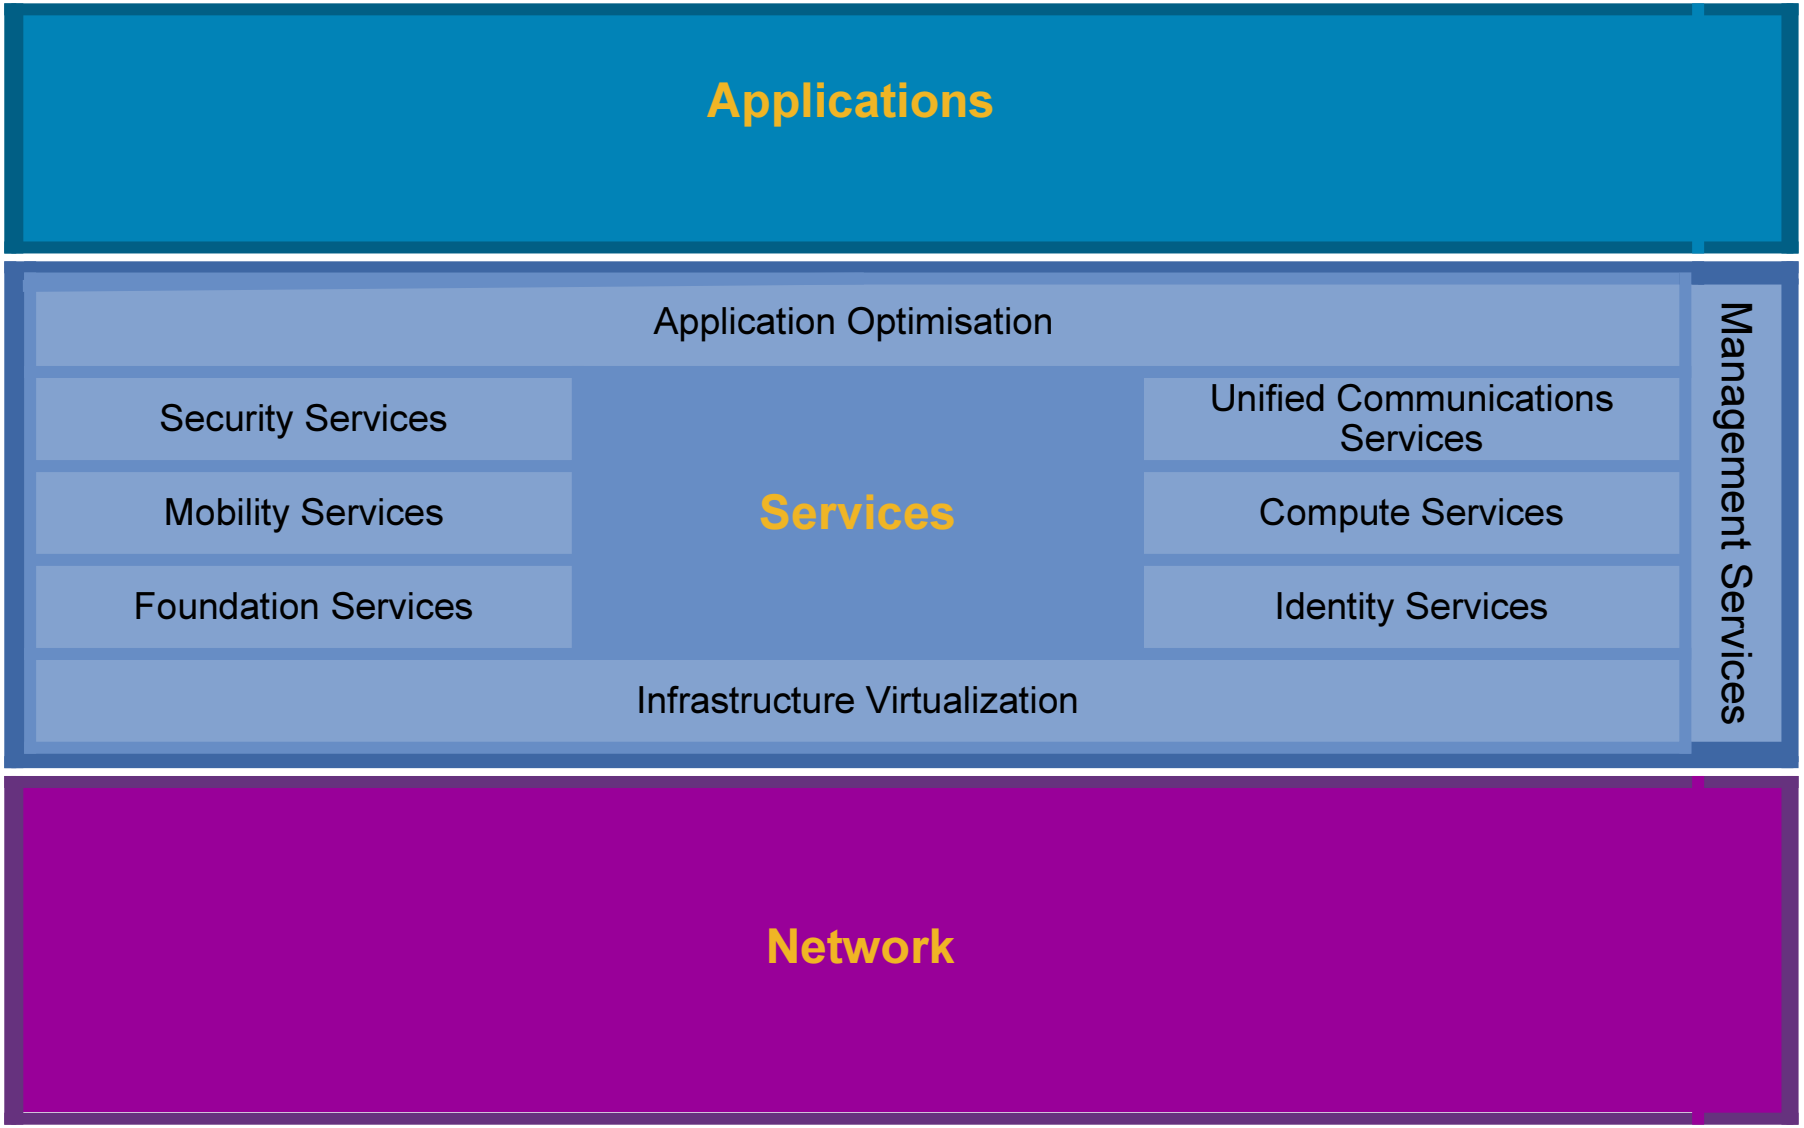

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

The Network as The Platform for …

The Network as The Platform for …

Application

Application Optimisation

Call Control,

Services:

Interaction & Integration

Application
Inpsection

Network

Ethernet
MPLS

Architecture Example
Distributed Application Intelligence

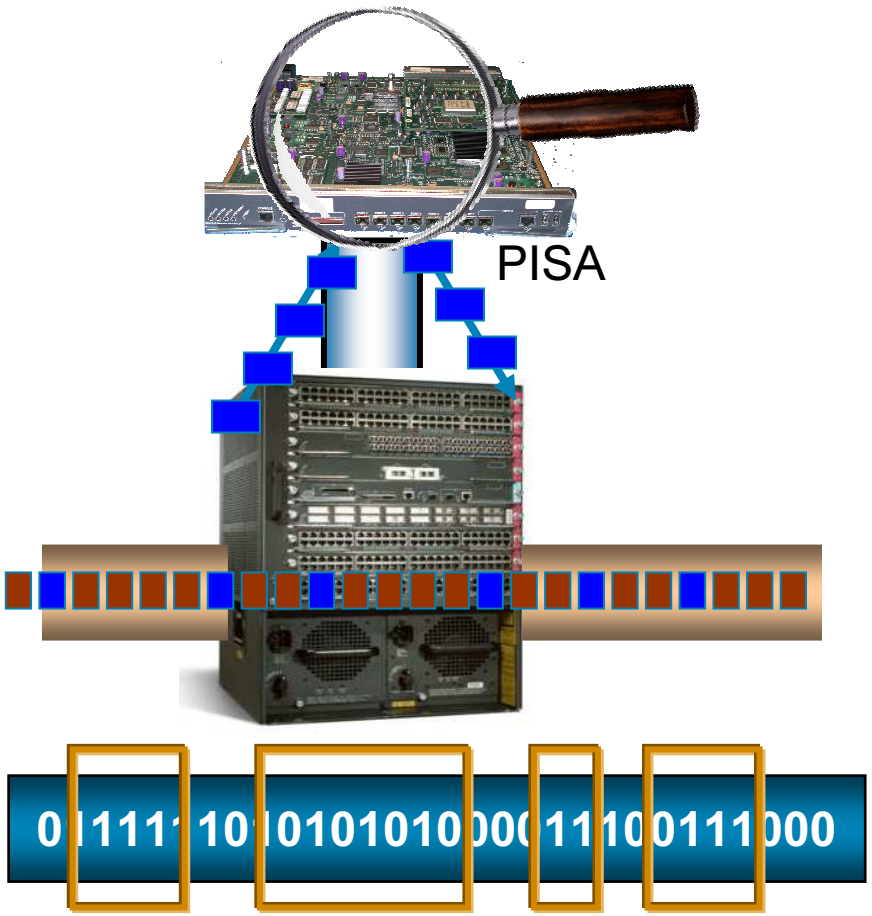

- Deep Packet Inspection for Application Recognition & Security
- Recognizes 100+ Applications
- Protection from misbehaving Applications
- Highly scalable distributed control of application flows
- Perimeter defense through flexible packet matching
- Programmatic Interface
- Policy enforcement and distributed policy enforcement

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Architecture Example
Distributed Application Intelligence

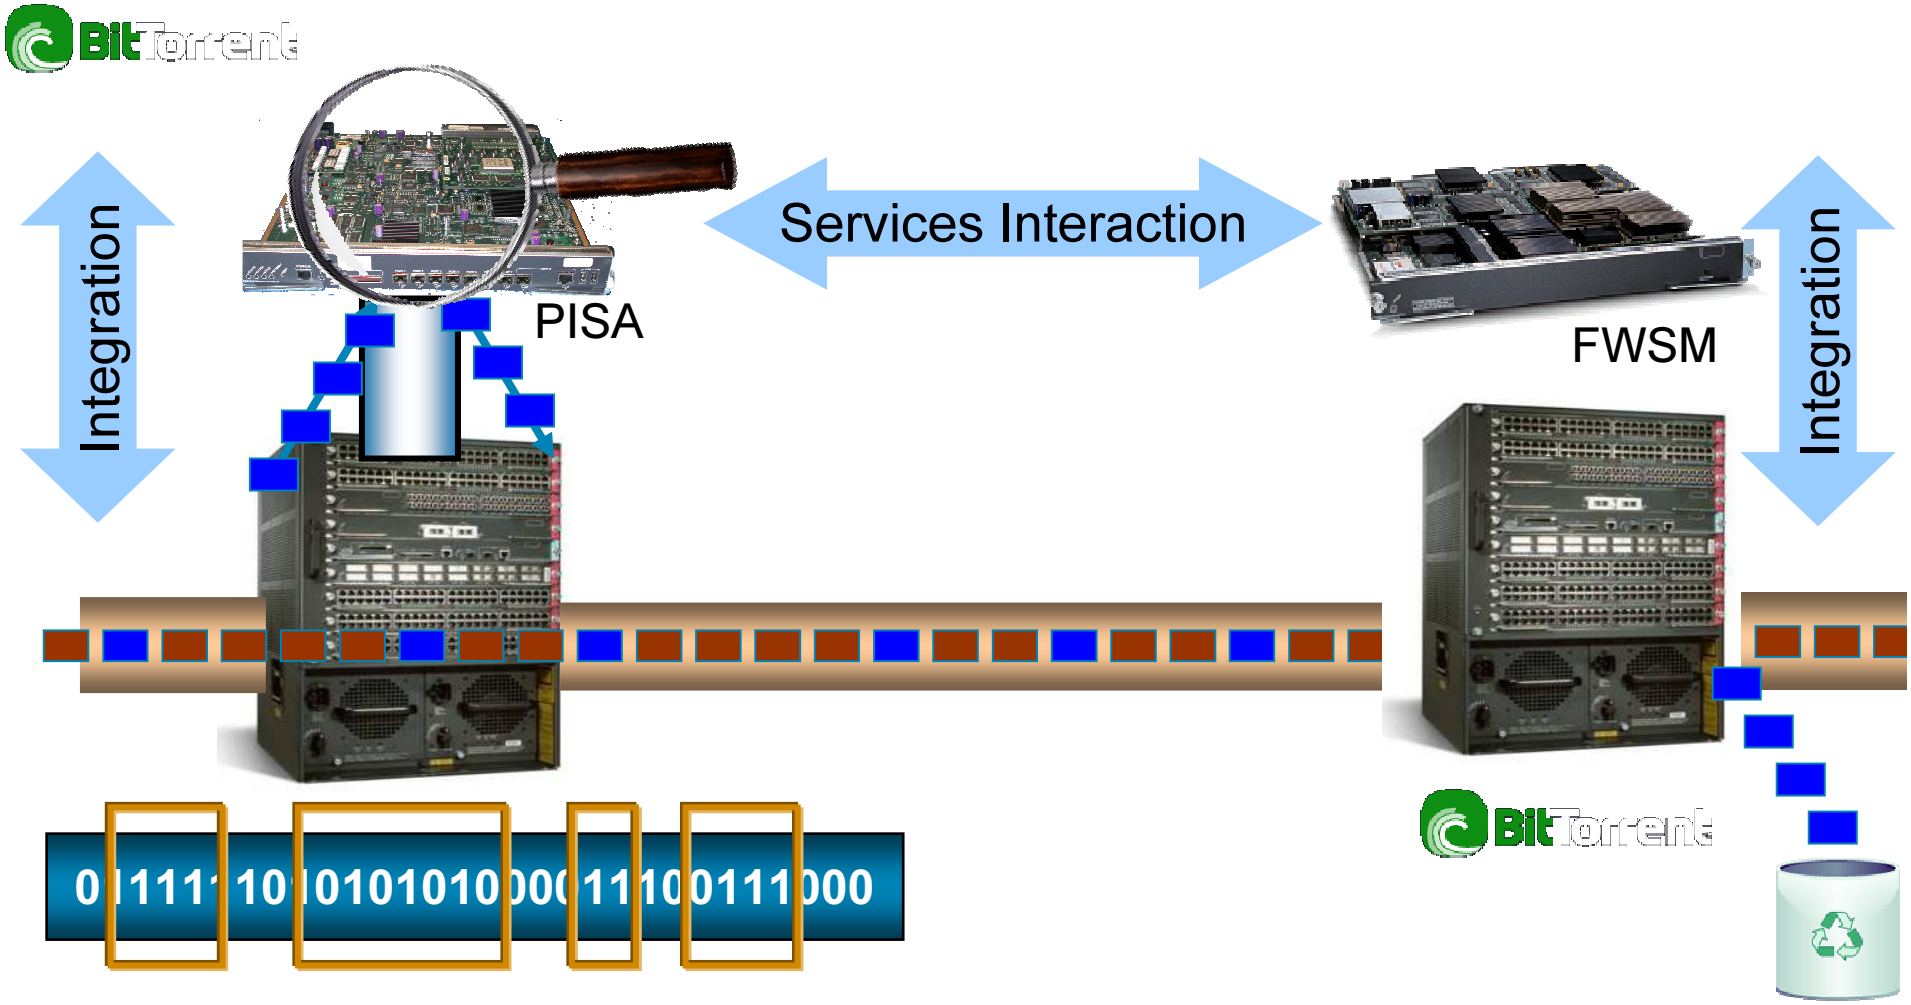

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Architecture Example
Application Acceleration

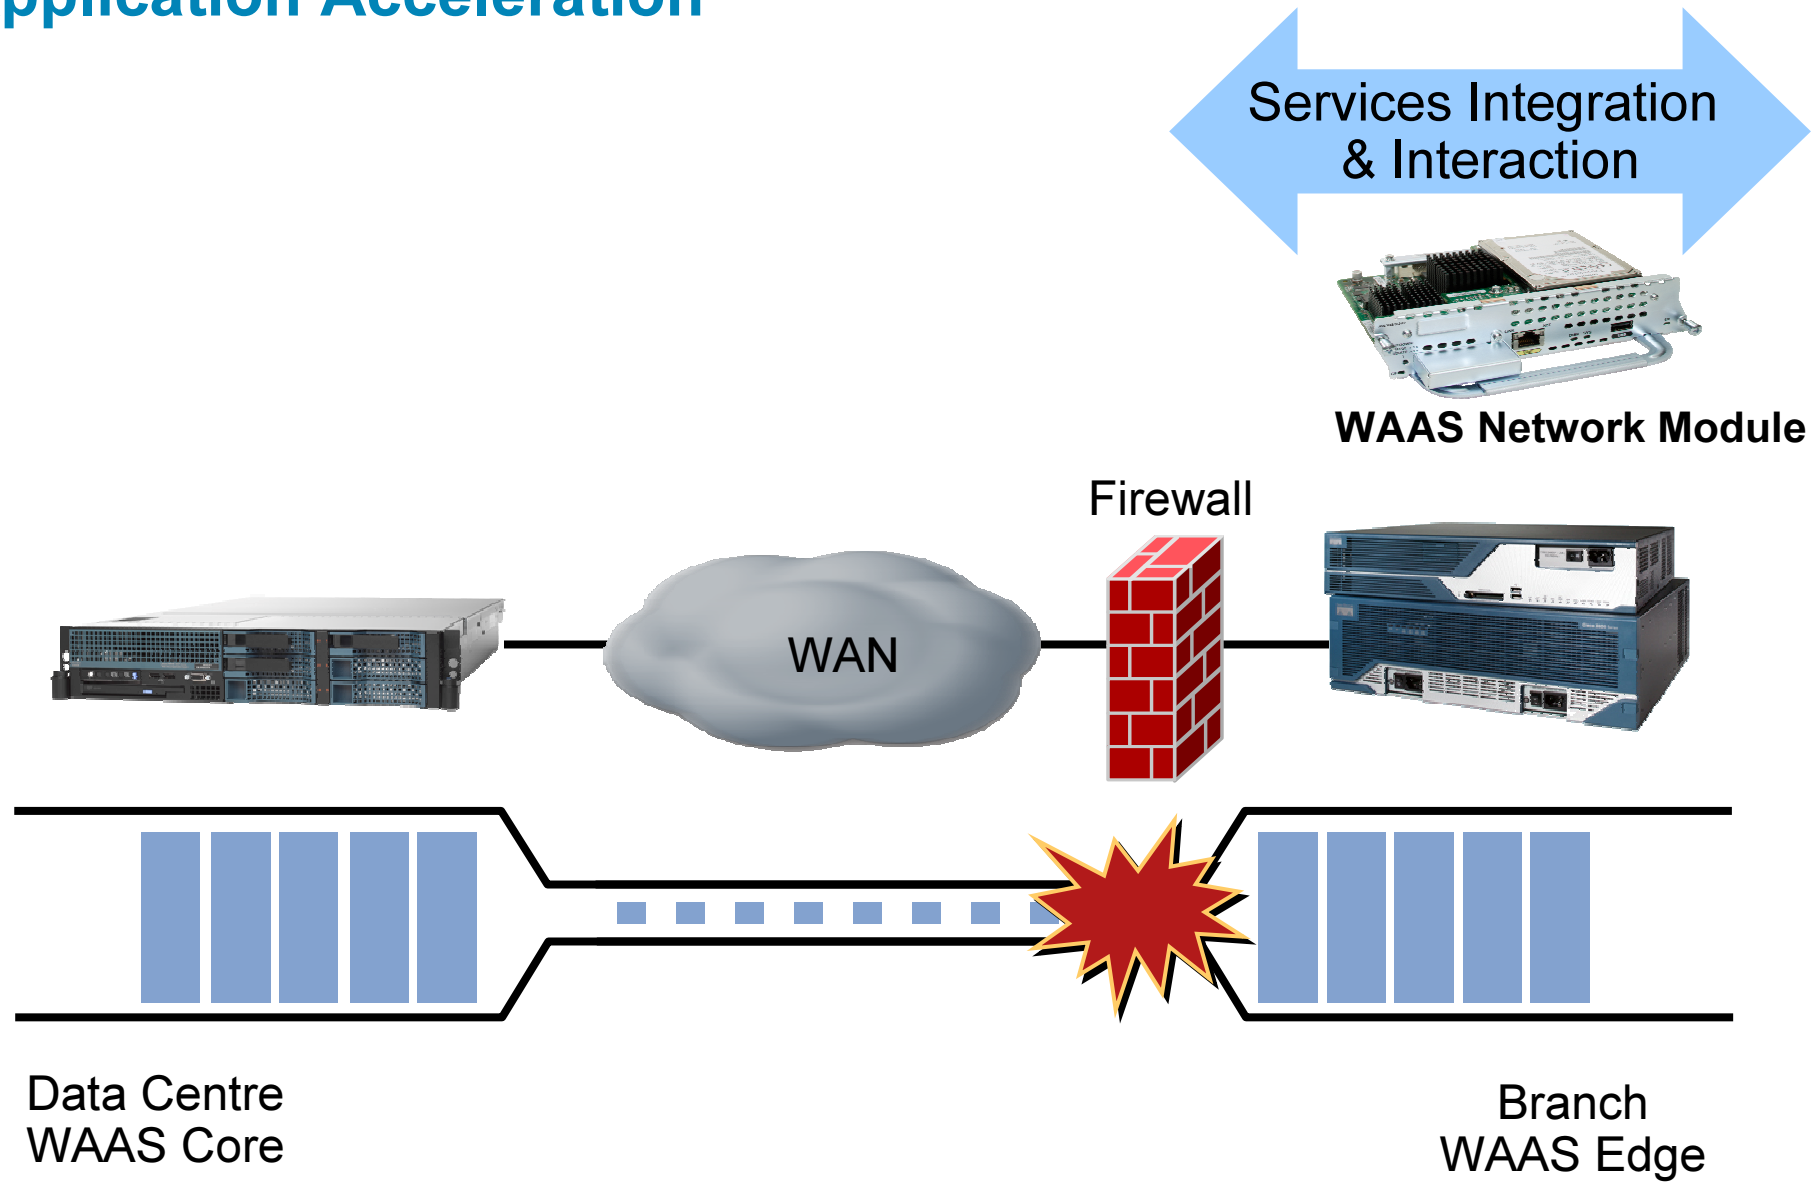

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Places In The Network

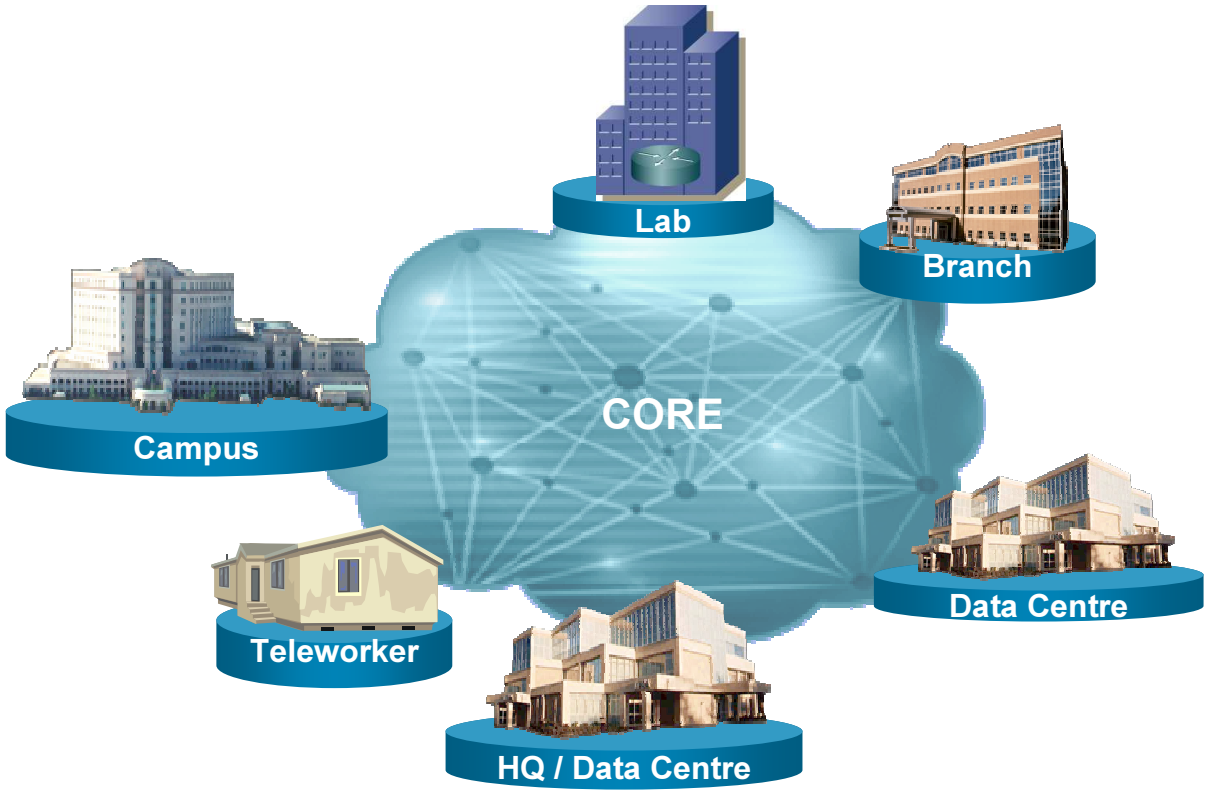

Cisco Validated Designs (CVD)

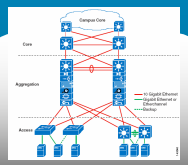
Data Centre Architecture

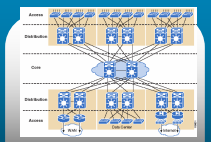
Campus Architecture

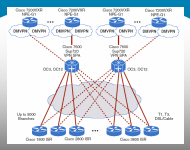
WAN/Branch Architecture

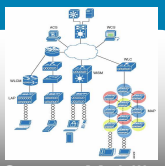
Secure Mobility Architecture

Unified Communications

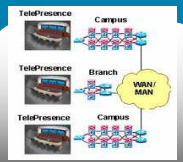

Telepresence Architecture

Architectural Approach
Campus Communications Fabric

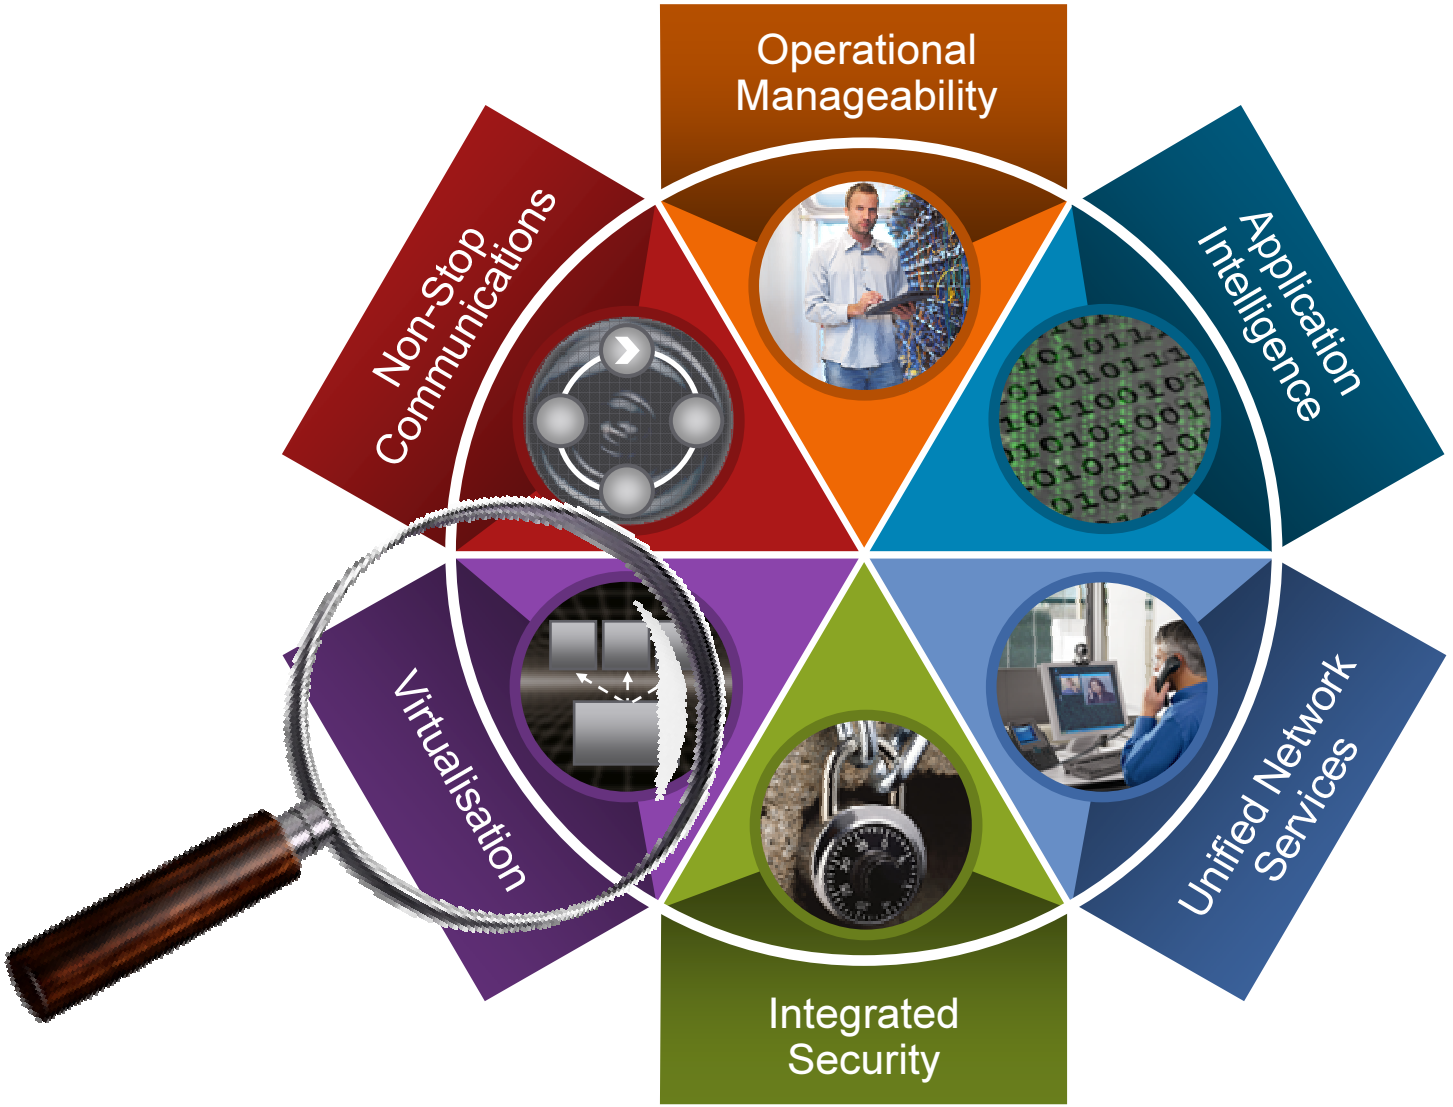

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation History

- Physically separate workgroups of hub-based users
- As the hub islands grew, so did the requirements to connect disparate users with some form of security to police traffic between them – from this the router was born…

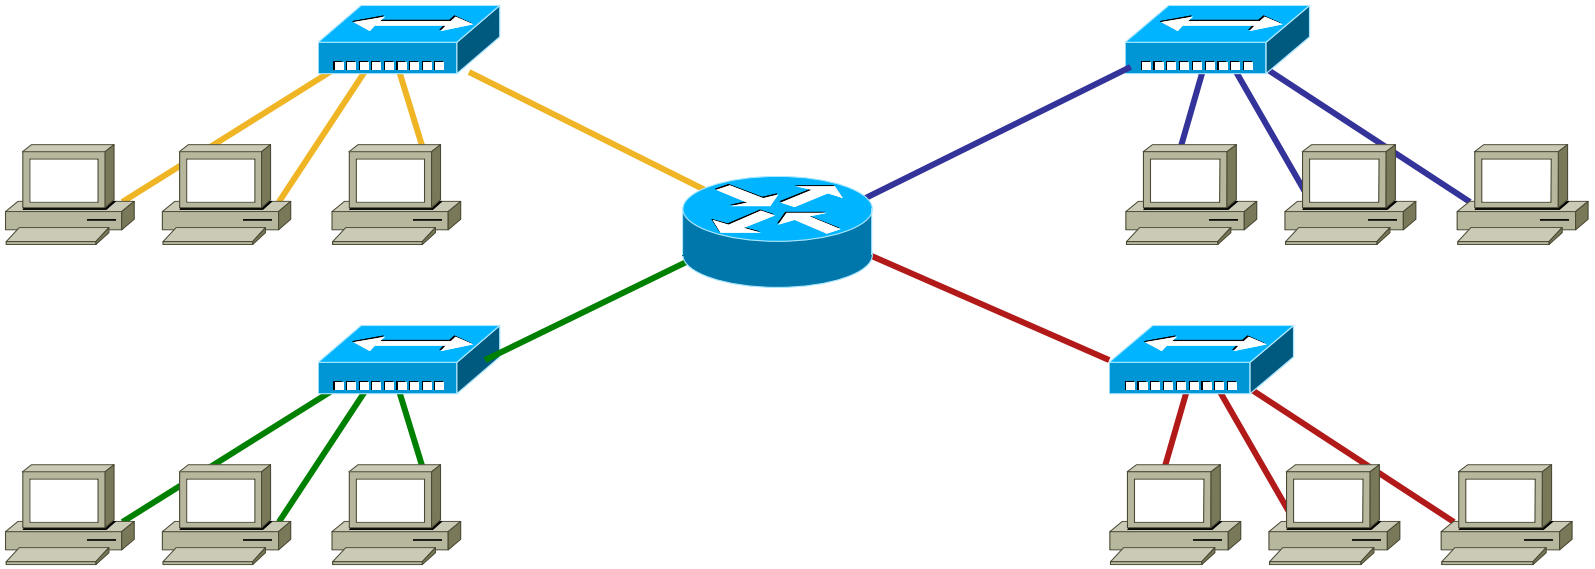

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation History

- Switches with trunking provided the next evolutionary step.
- Trunks implemented to allow transport of VLANs across multiple devices – this provided Virtualisation at Layer 2…
- Router physically ‘Virtualised’ in switch

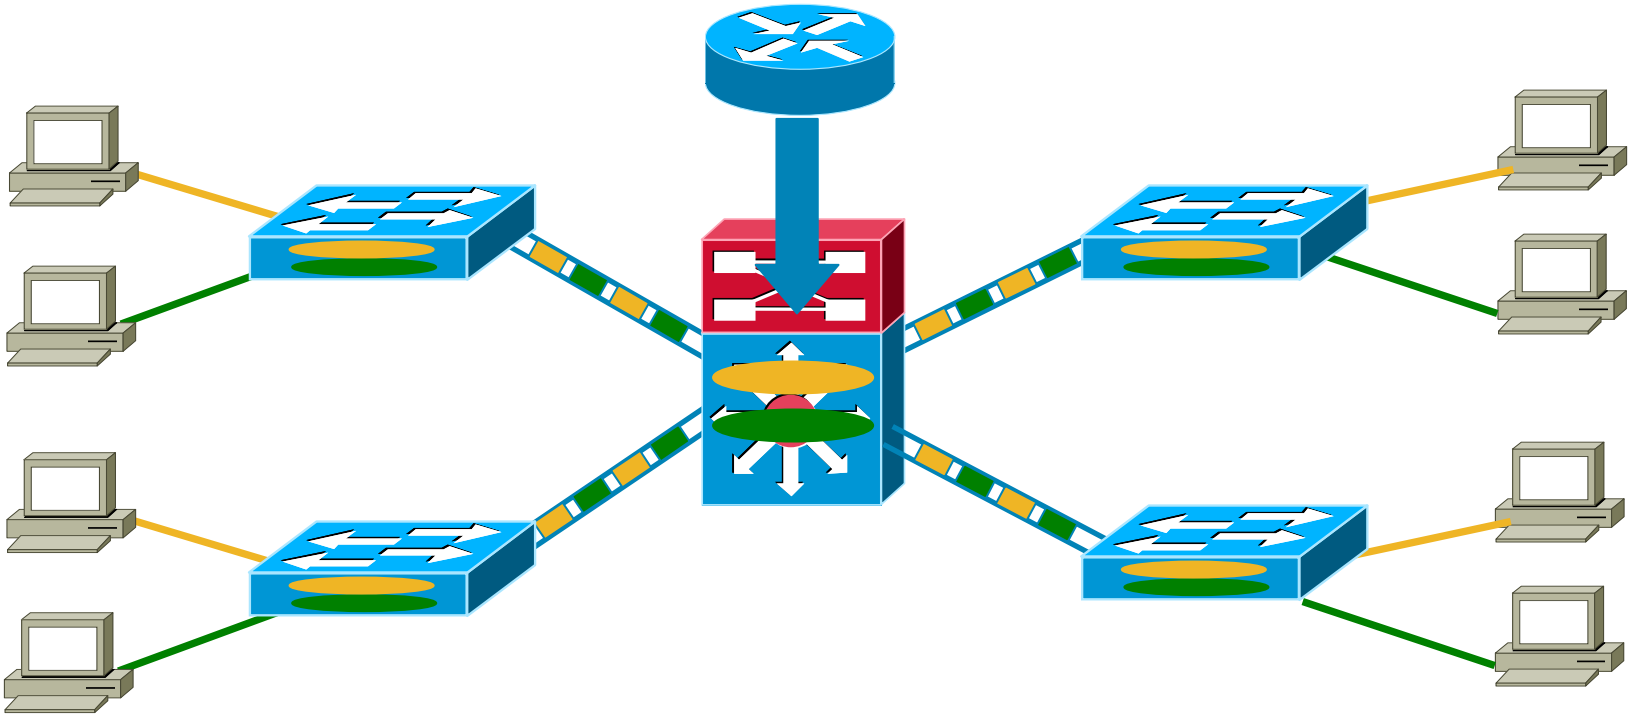

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation History

- **Layer 3 Virtualisation appeared in the form of Virtual Routing and Forwarding (VRF) which allowed a Layer 3 device to virtualize routing tables giving each VRF each it's logical view of the network …**

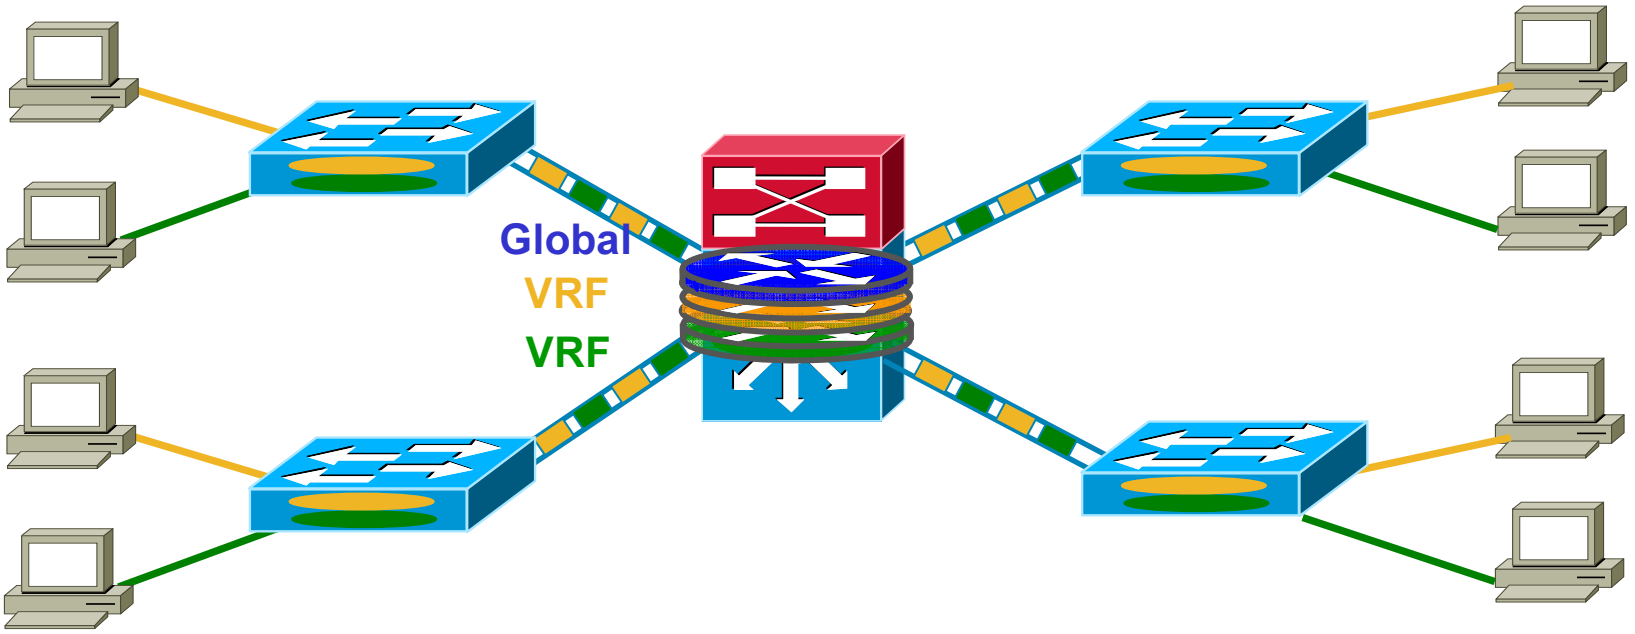

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation History

- **High Availability requirements meant dual homing and associated complexities.**

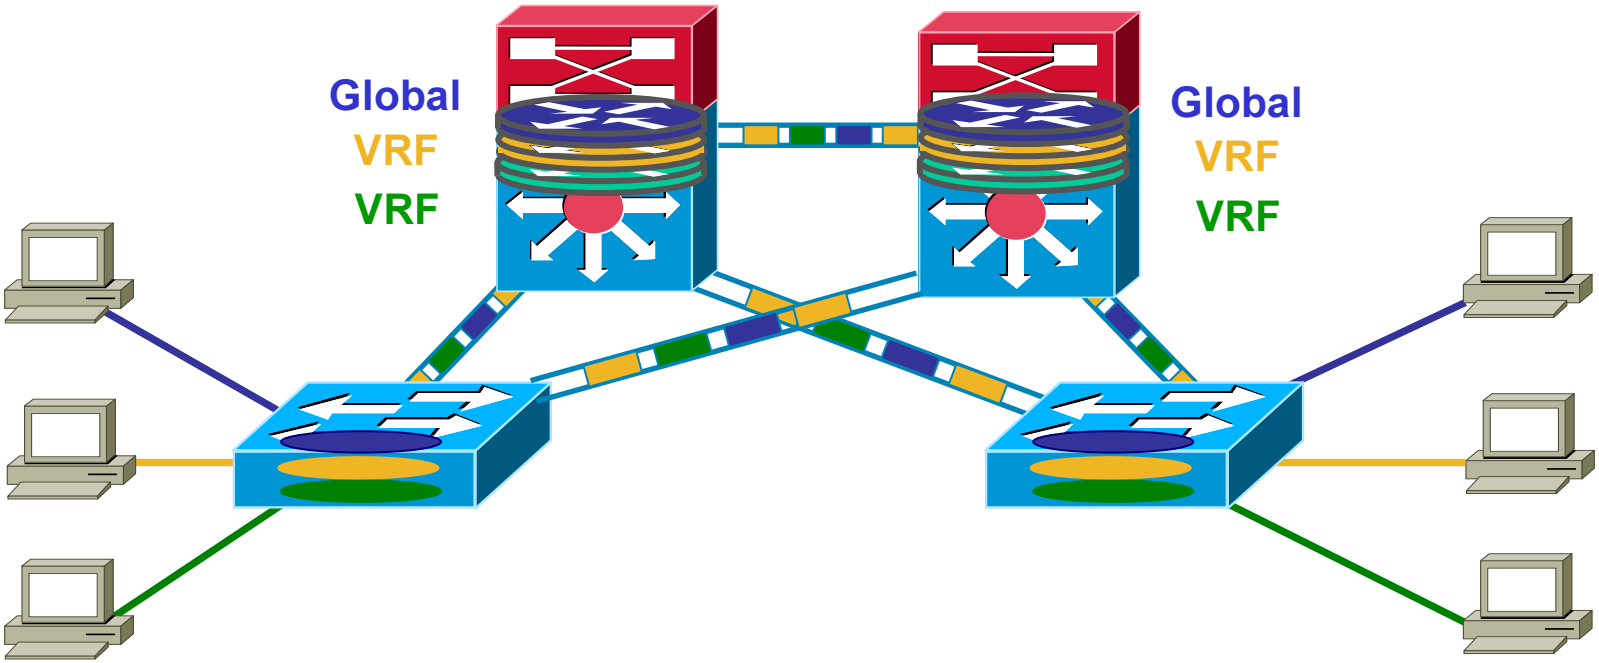

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation History

- **Virtual Switching is the latest addition to the Virtualisation story allowing 2 devices to operate as a single logical network device**

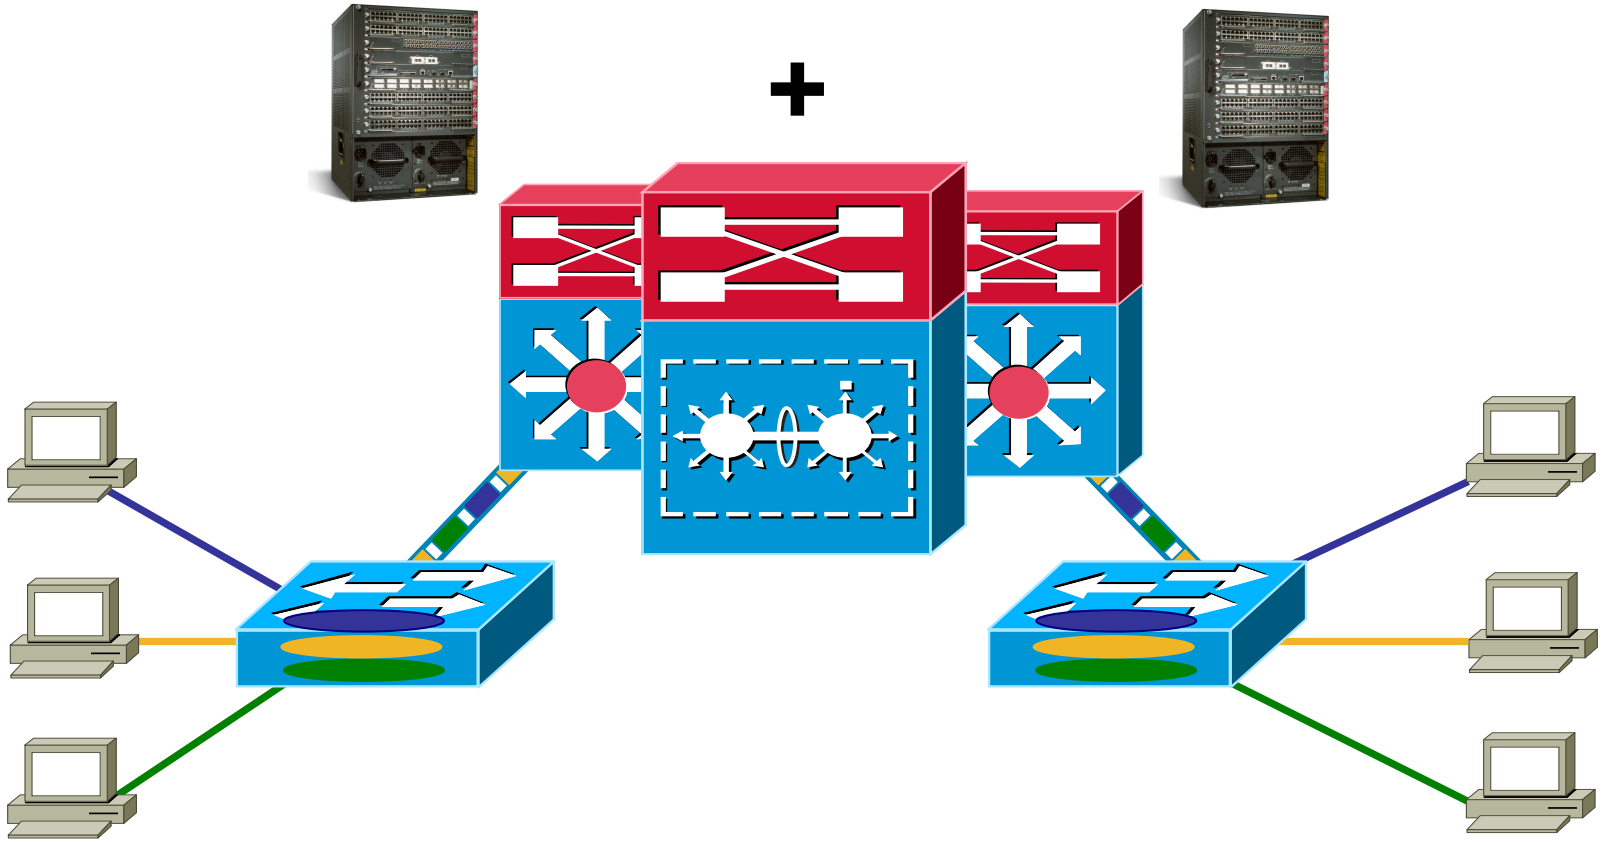

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

High Availability Campus Design

Access

Distribution

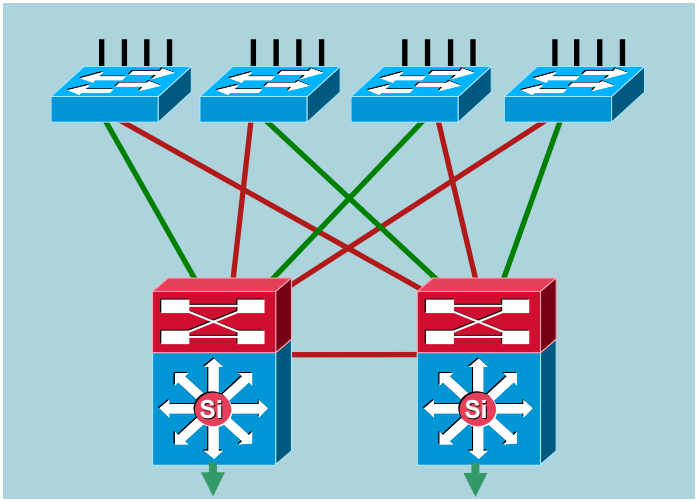

Complex suite of protocols – Spanning Tree, HSRP, VRRP

Non-deterministic, Stateless Recovery From Failure

Inefficient Resource Utilization

Increased Management Burden

High Availability Campus Design

Access

Distribution

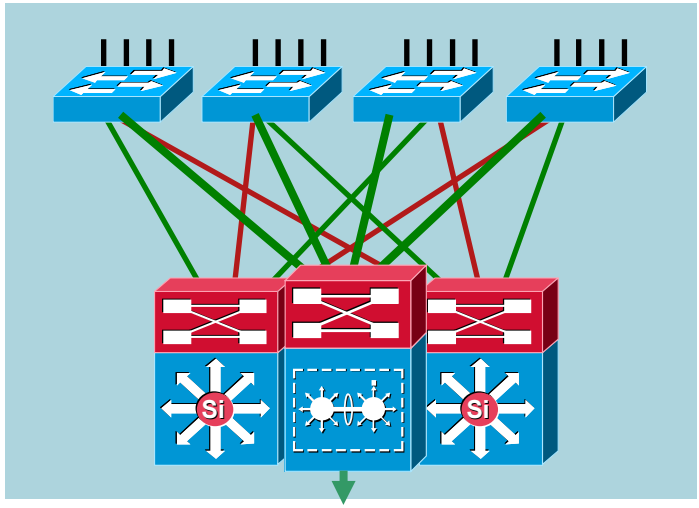

Eliminate complex suite of protocols

Stateful Deterministic Inter-chassis Recovery From Failure

Double Fabric Bandwidth

Reduced Management Burden

High Availability Campus Design

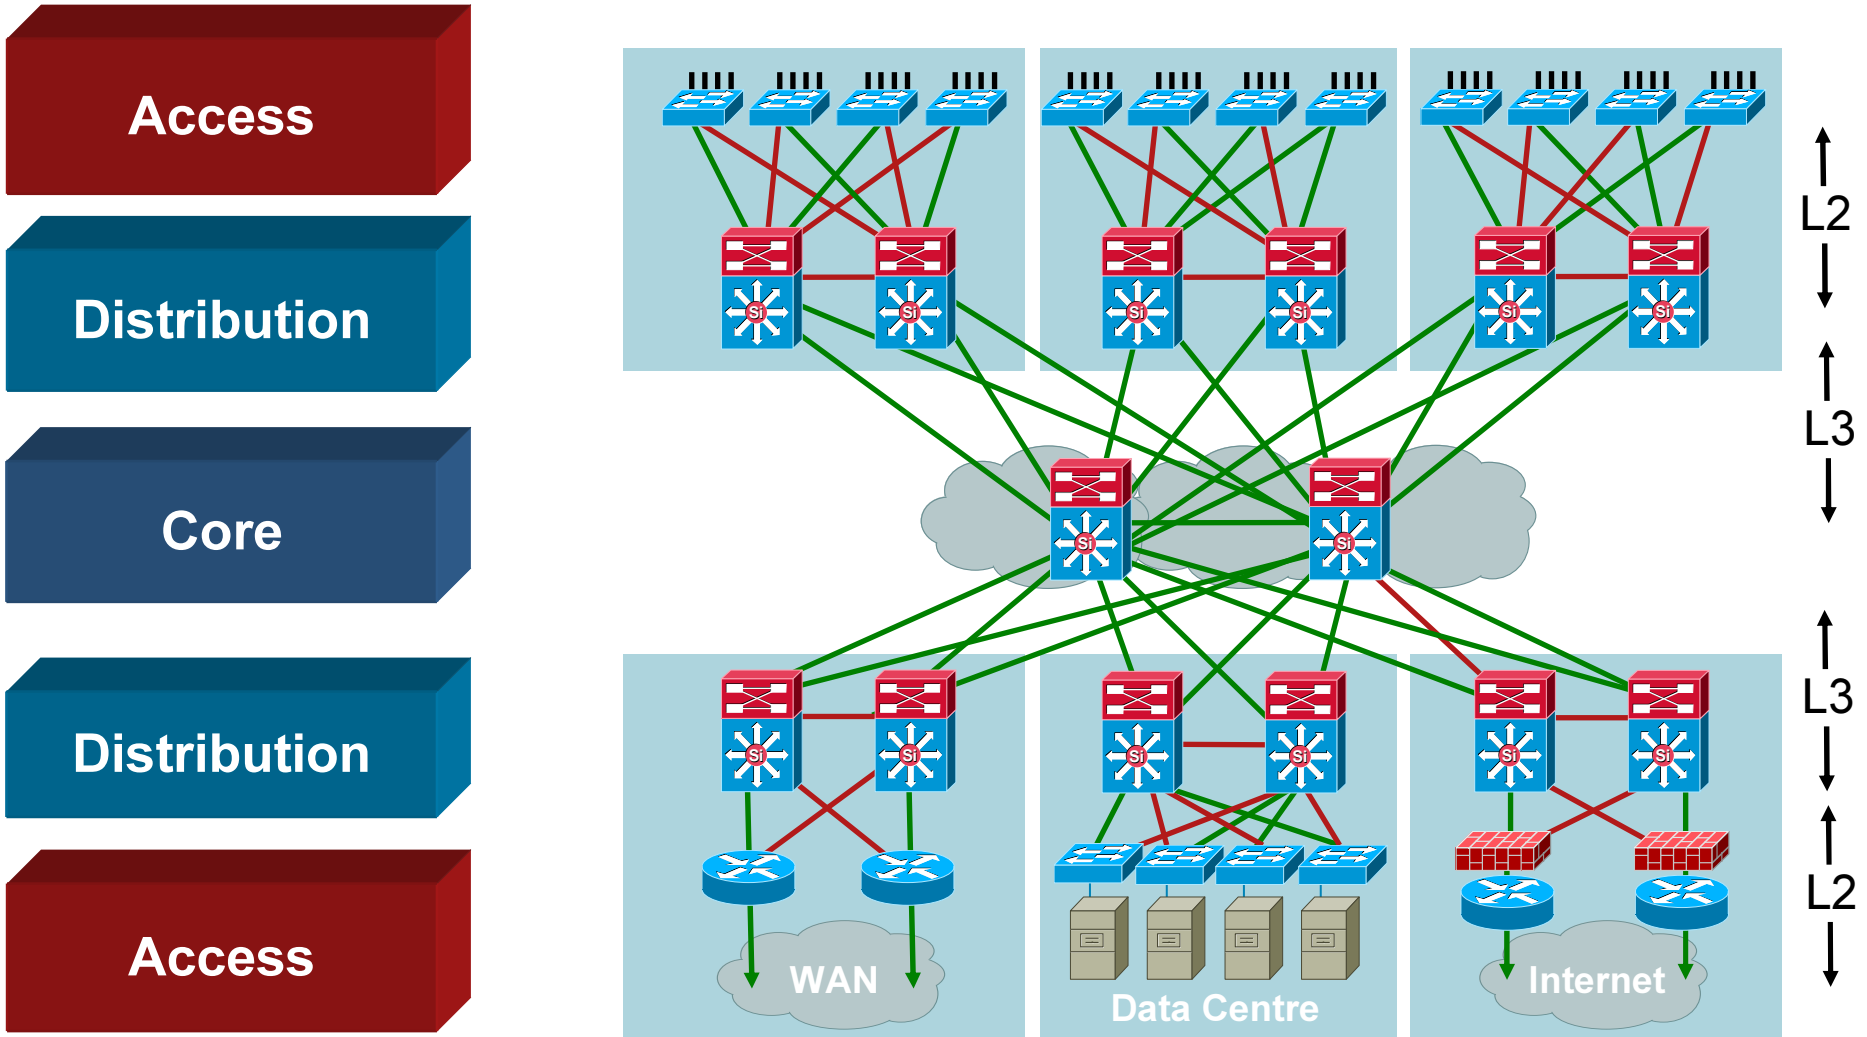

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

High Availability Campus Design

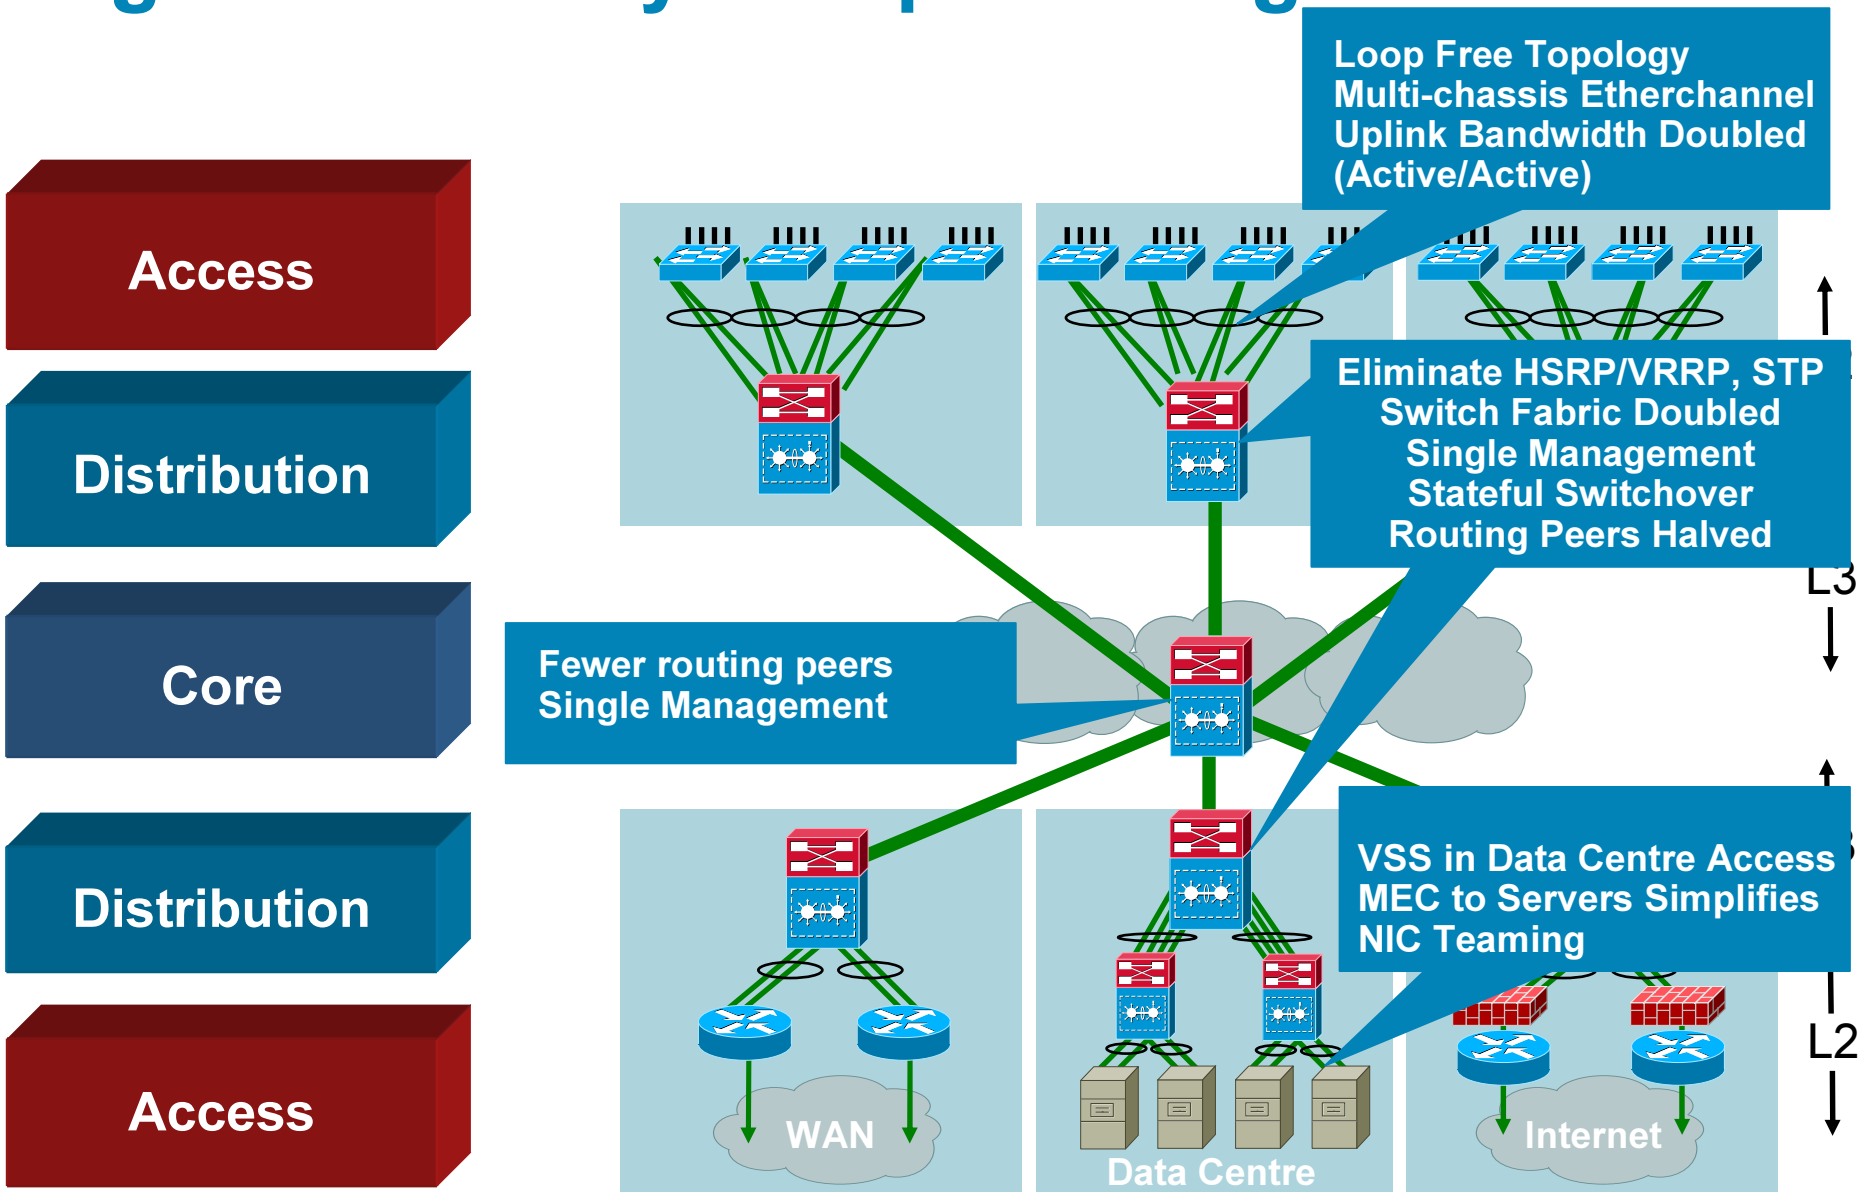

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System Concepts

Virtual Switch Domain

Defines 2 switches participating together as a single virtual switching system

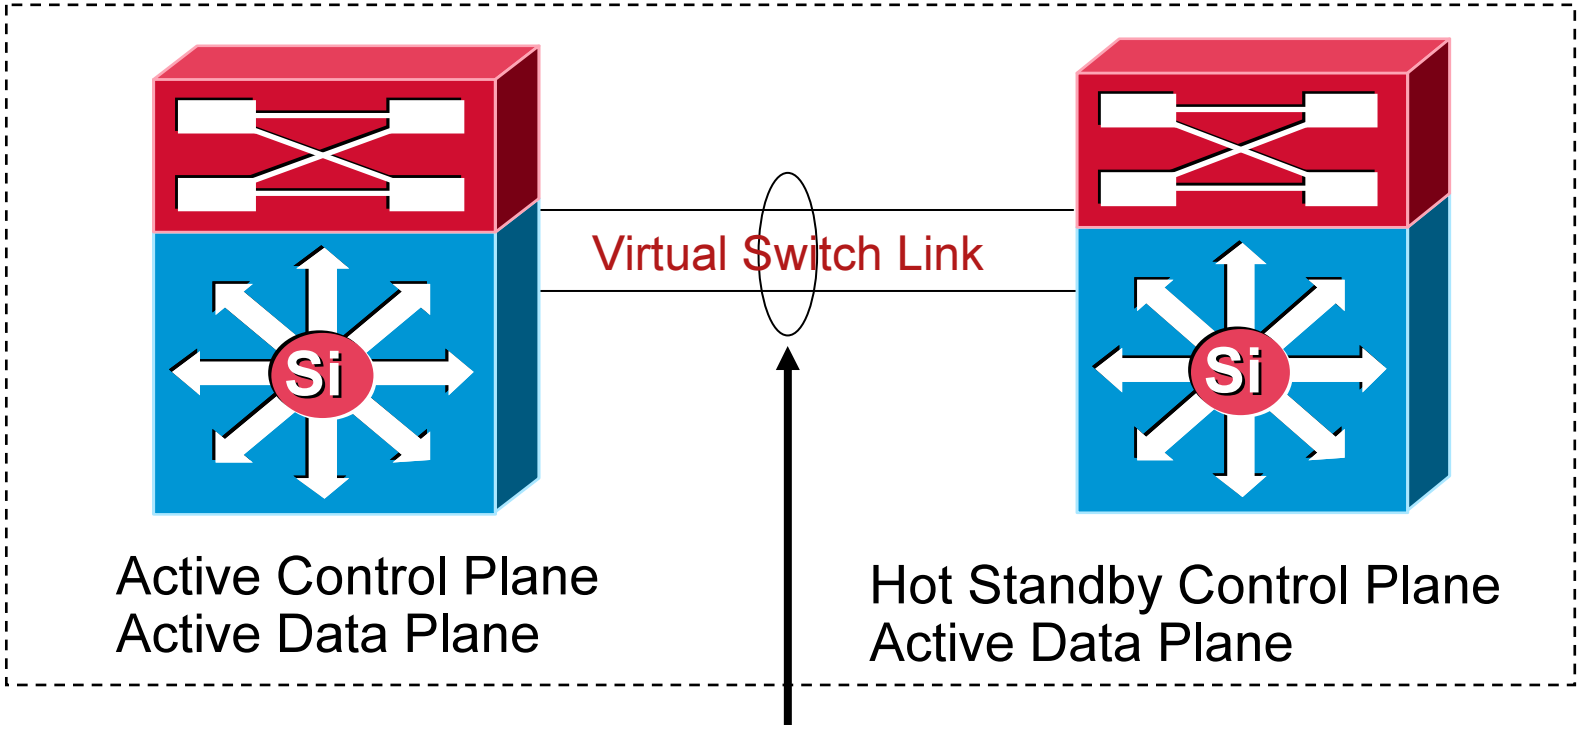

Special 10GbE link bundle joining the 2 switches allowing them to operate as a single device

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System Concepts
Virtual Switch Domain

A Virtual Switch Domain ID represents the logical grouping the 2 physical chassis within a VSS. It is possible to have multiple self-contained VS Domains throughout the network…

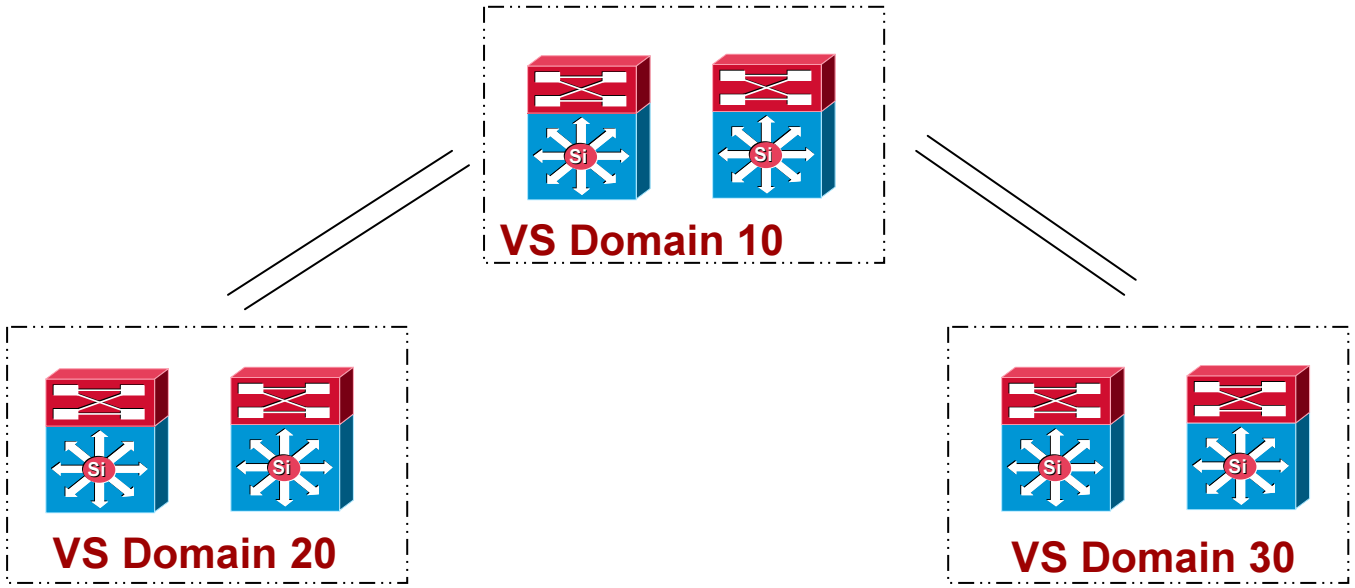

The configurable values for the domain ID are 1-255. It is always recommended to use a unique VS Domain ID for each VS Domain throughout the network…

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System Concepts
Control Plane

Both data planes are active and forwarding, but only one control plance – hence there is only a single management point

Virtual Switch Active
Active Control Plane
Active Data Plane

Virtual Switch Standby
Standby Control Plane
Active Data Plane

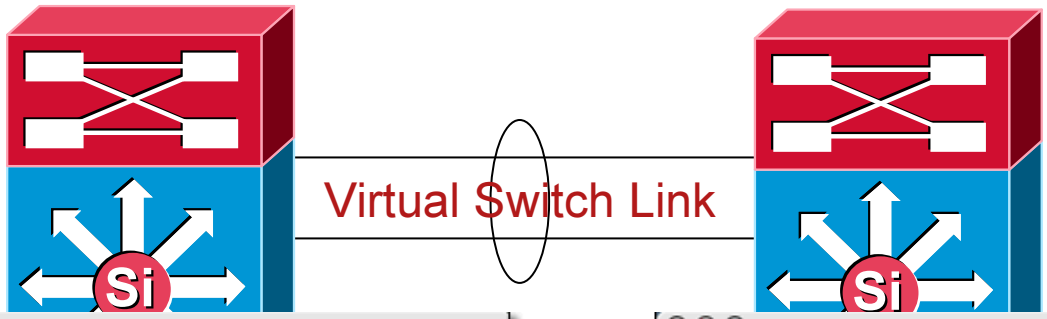


```
vs-vsl#sh mod
Mod Ports Card Type                                Model              Serial No.
--- ----- ---------------------------------------- ------------------ -----------
  1    8  CEF720 8 port 10GE with DFC               WS-X6708-10GE      SAD110100PF
  5    5  Supervisor Engine 720 10GE (Active)      WS-S720-10G        SAD104707BB
  7   48  CEF720 48 port 10/100/1000mb Ethernet    WS-X6748-GE-TX     SAL0943435M
  8   24  CEF720 24 port 1000mb SFP                WS-X6724-SFP       SAL09158ZZT

Mod MAC addresses                       Hw    Fw           Sw           Status
--- ----------------------------------- ------ ------------ ------------ -------
  1  001a.6d65.e668 to 001a.6d65.e66f   1.3   12.2(18r)S1  12.2(SIERRA_ Ok
  5  000f.8f3a.8b7c to 000f.8f3a.8b83   0.423 8.5(1)       12.2(SIERRA_ Ok
  7  0015.f90a.d3a0 to 0015.f90a.d3cf   2.3   12.2(14r)S5  12.2(SIERRA_ Ok
  8  0013.1a23.463c to 0013.1a23.4653   2.2   12.2(14r)S5  12.2(SIERRA_ Ok

Mod  Sub-Module                  Model              Serial       Hw     Status
---- --------------------------- ------------------ ------------ ------ -------
  1  Distributed Forwarding Card WS-F6700-DFC3C     SAD110208FY  1.0    Ok
  5  Policy Feature Card 3       WS-F6K-PFC3CXL     SAD104704VB  0.203  Ok
  5  MSFC3 Daughterboard         WS-F6K-MSFC3       SAD104607U3  0.100  Ok
  7  Centralized Forwarding Card WS-F6700-CFC       SAL09486L84  2.0    Ok
  8  Centralized Forwarding Card WS-F6700-CFC       SAL09222493  2.0    Ok

Mod  Online Diag Status

vs-vsl#
```

```
Trying
Connected to
Escape character is '^]'.

vs-vsl-sdby>en
Standby console disabled

vs-vsl-sdby>
```

NOTE: The Standby Console is disabled and cannot be used for normal operational purposes

Virtual Switch System Concepts
Data Plane

Data Planes in both switches are active - each has a full copy of the forwarding tables and Security/QOS policies in hardware such that each can make a fully informed local forwarding decision…

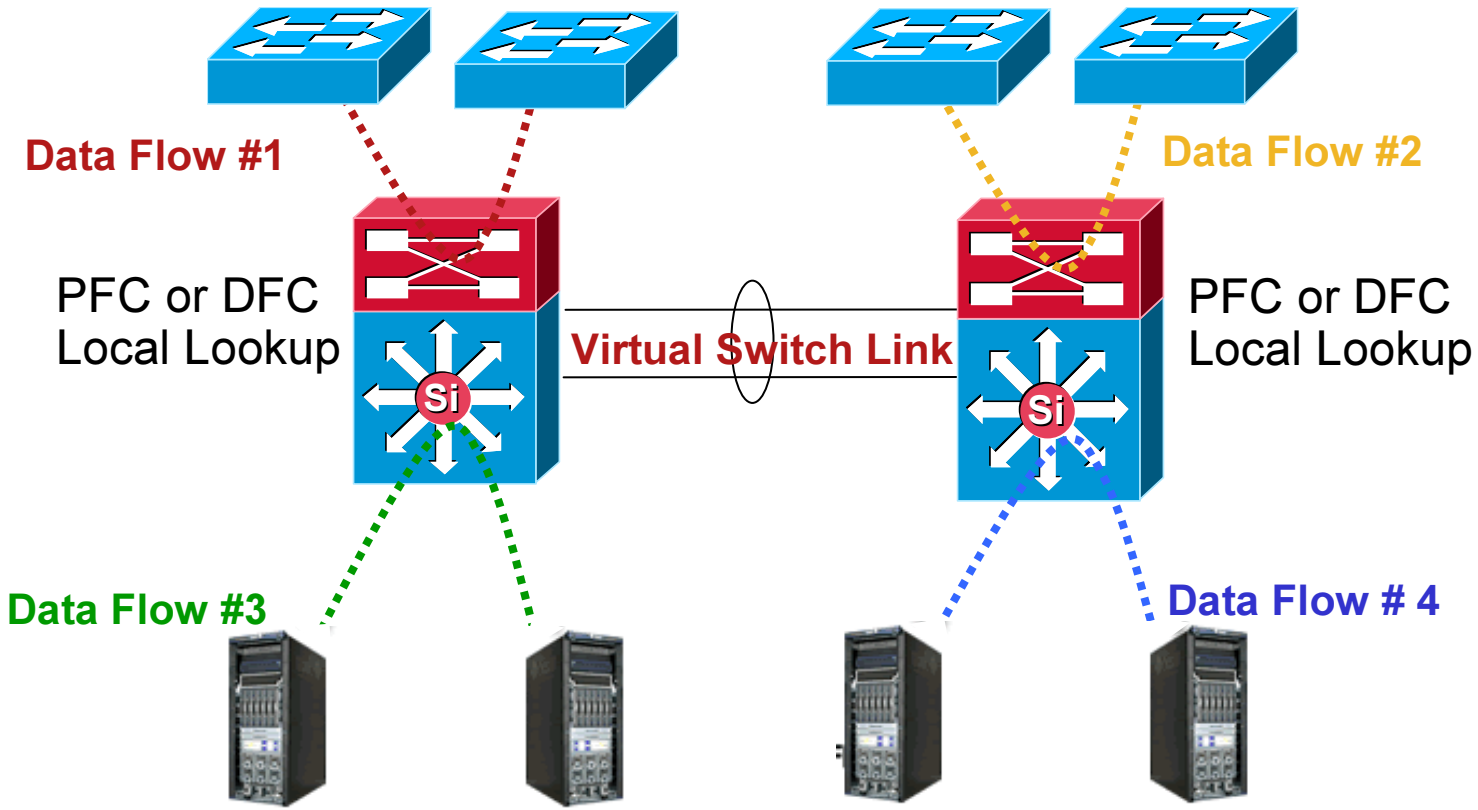

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch Architecture
Virtual Switch Link

The Virtual Switch Link extends the out of band channel allowing the active control plane to manage the hardware in the second chassis…

A Virtual Switch Link Bundle can consist of up to 8 x 10GbE Links

All traffic traversing VSL is encapsulated in 32-byte header containing ingress & egress switch port indexes, CoS, VLAN ID & other important L2, L3 information

Control plane uses VSL for CPU to CPU communications while data plane uses VSL to extend internal switch fabric to external chassis

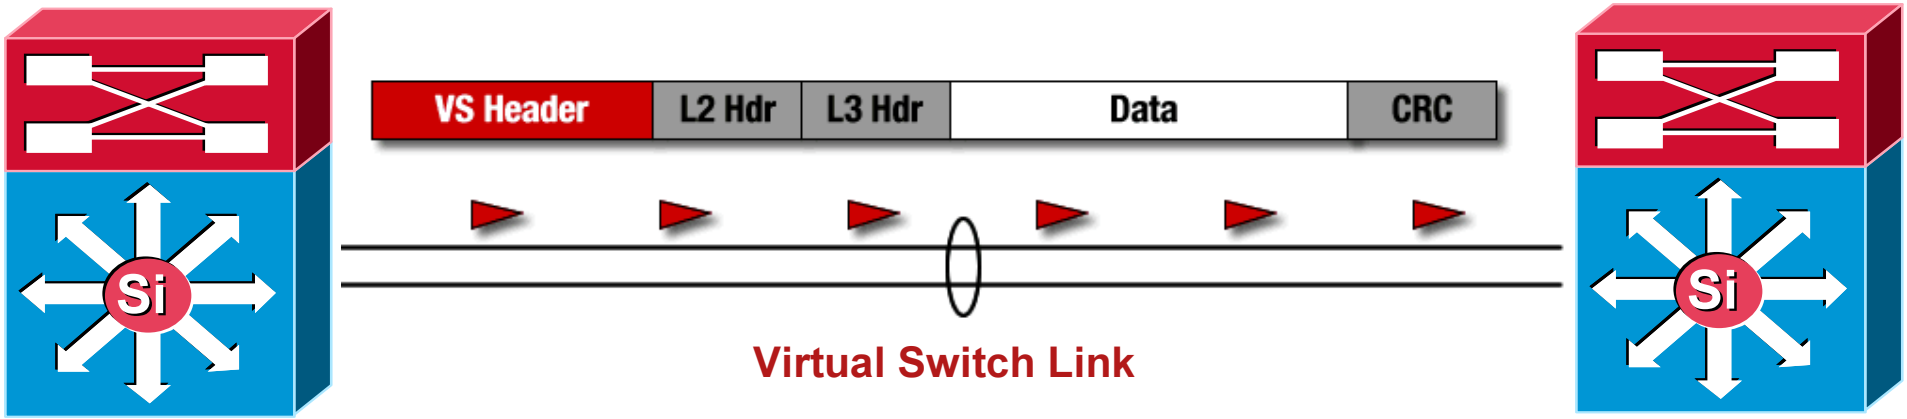

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch Architecture
Multichassis EtherChannel (MEC)

In a Virtual Switch environment, 2 physical switches form a single logical entity. Etherchannels can now also be extended across the 2 physical chassis

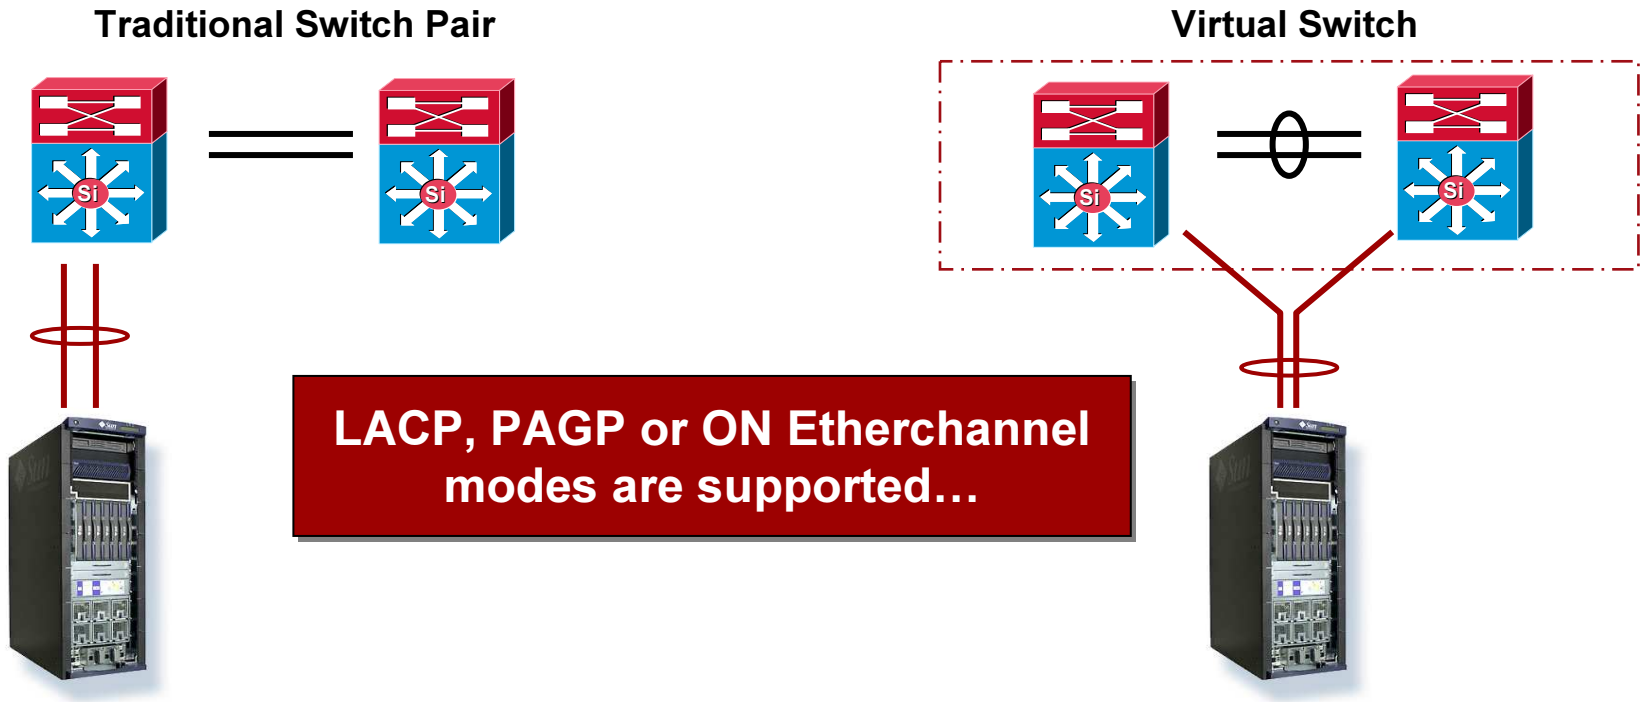

Regular Etherchannel on single chassis

Multichassis EtherChannel across 2 VSL-enabled Chassis

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch Architecture
Multi Chassis Etherchannel (MEC)

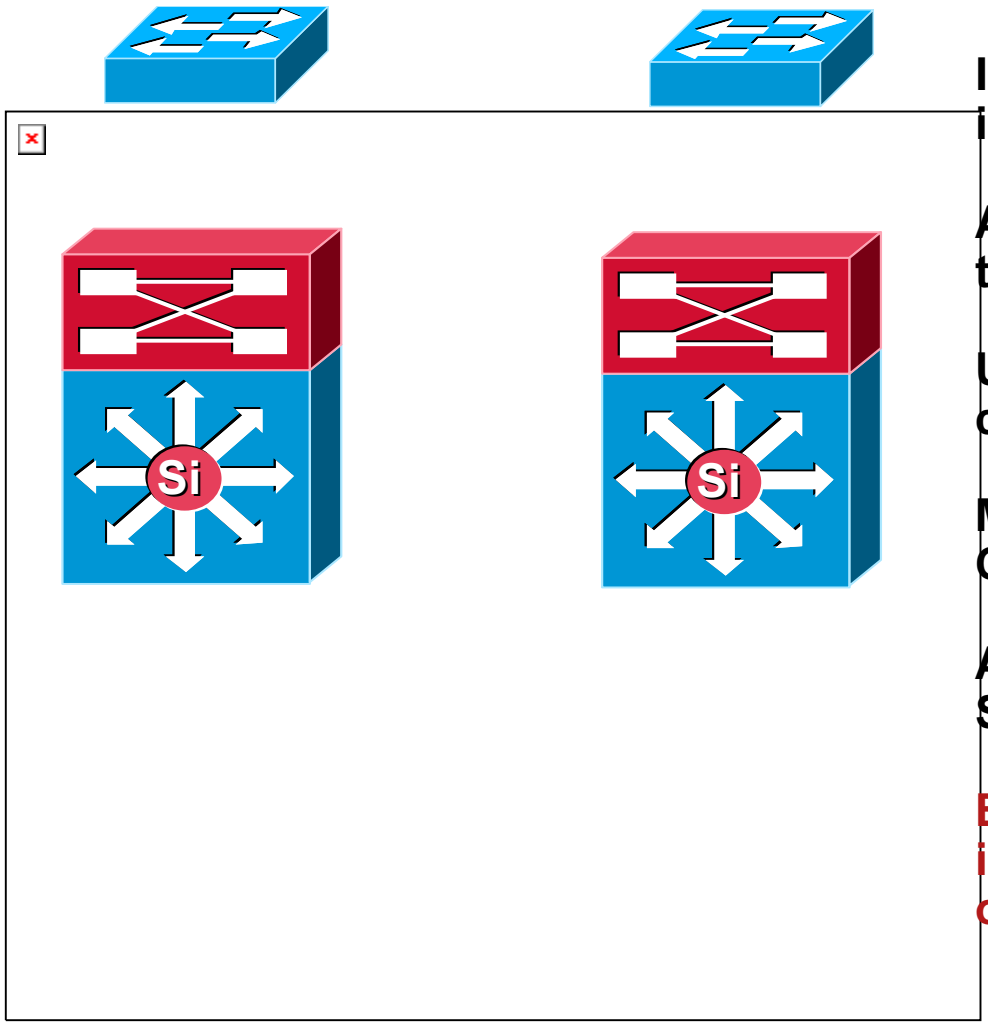

Multi Chassis Etherchannel (MEC)

Introduces a new deployment option for improving link resiliency

Allows an Etherchannel link to be terminated across 2 physical chassis

Up to 8 links support in a single multi-chassis Etherchannel

MEC is supported with both 802.3ad and Cisco PAGP

Attached host sees other end (Virtual Switch) as a single device

Etherchannel hash has been modified so in any bundle we always prefer local link over a link on other chassis

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch Architecture
Etherchannel Hash for MEC

Deciding on which link of a Multi-chassis Etherchannel to use in a Virtual Switch is skewed in favour of 'switch-local' links in the bundle to avoid overloading the Virtual Switch Link (VSL) with unnecessary traffic loads…

Blue Traffic destined for the Switch will result in Link A1 in the MEC link bundle being chosen as the destination path…

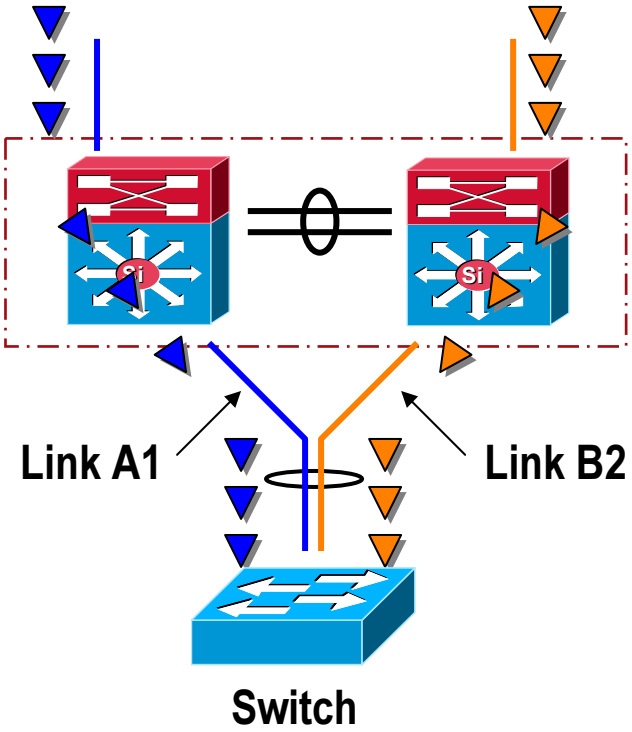

Orange Traffic destined for the Switch will result in Link B2 in the MEC link bundle being chosen as the destination path…

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch Architecture
Inter Chassis NSF/SSO

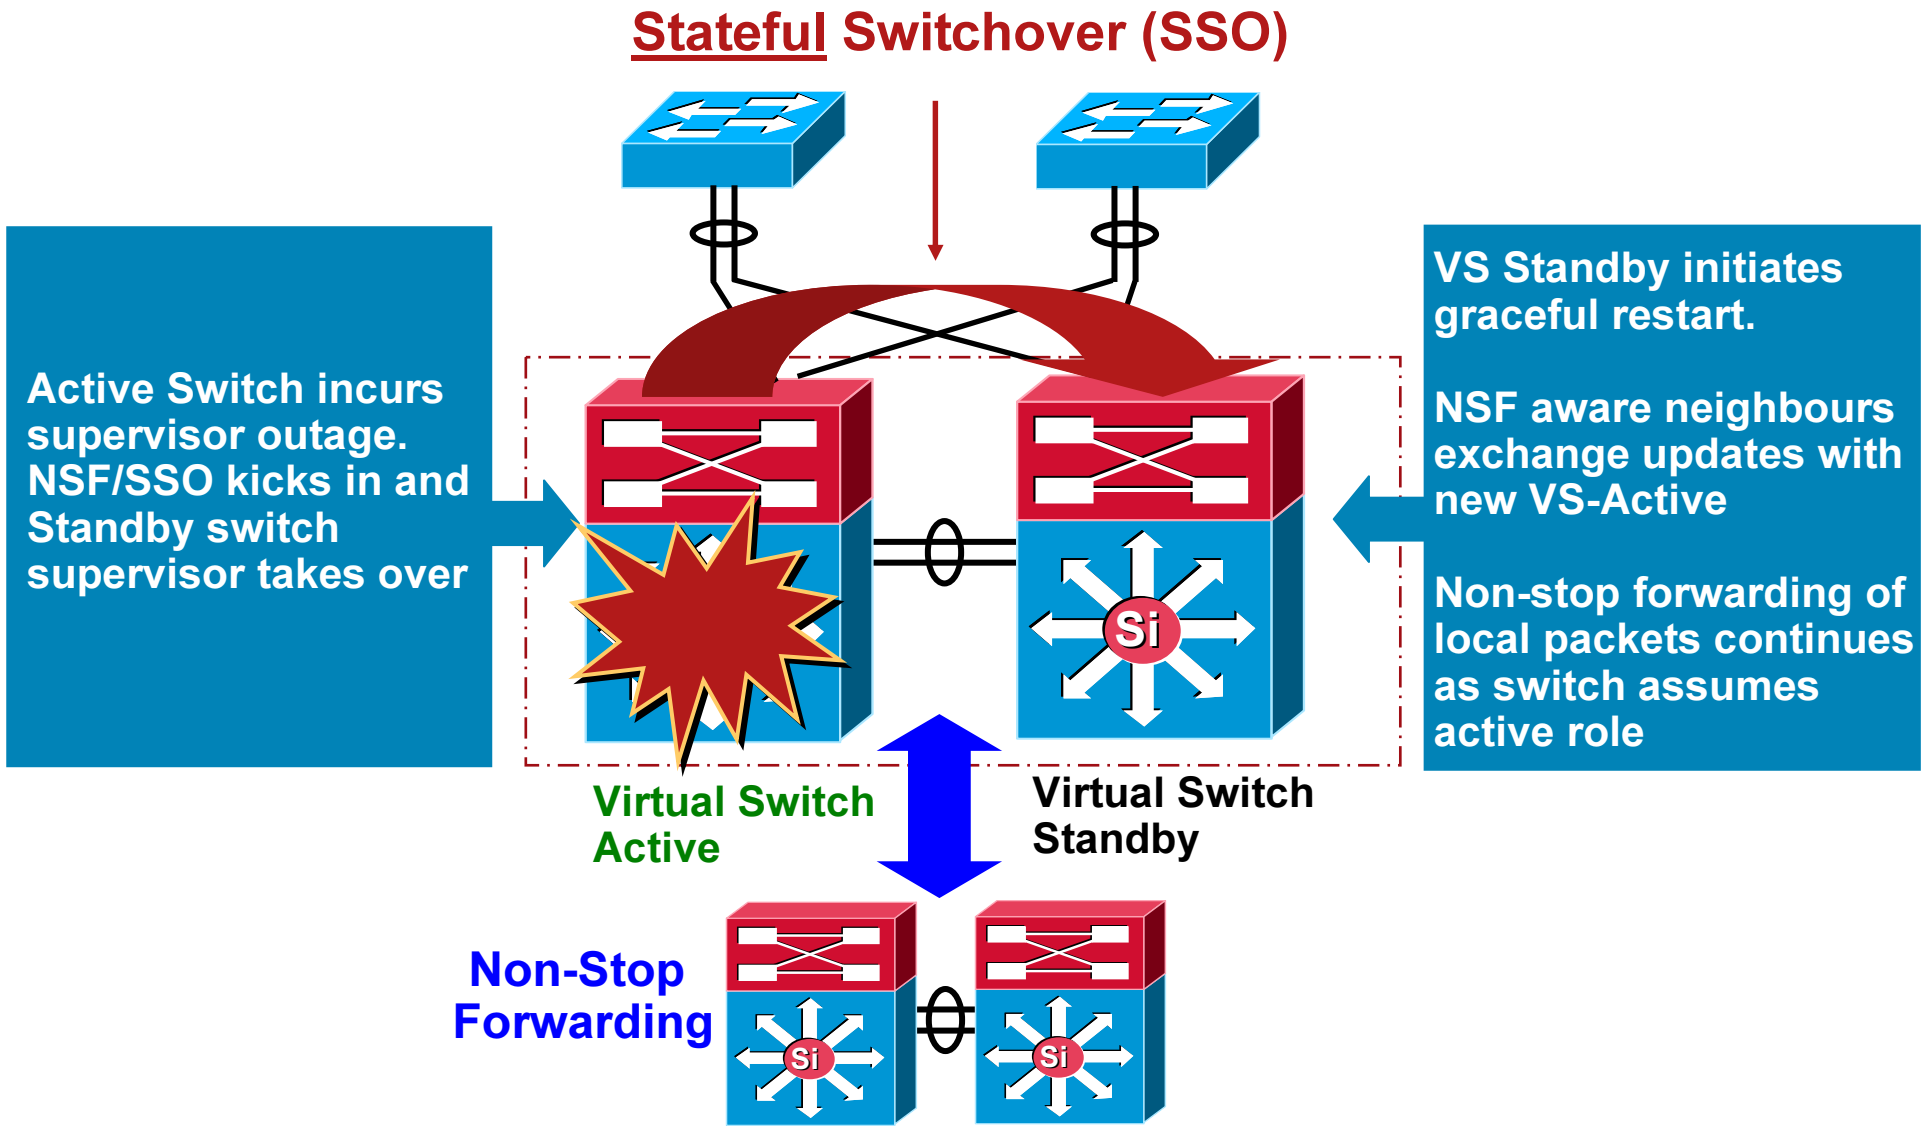

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System
Inter Chassis NSF/SSO

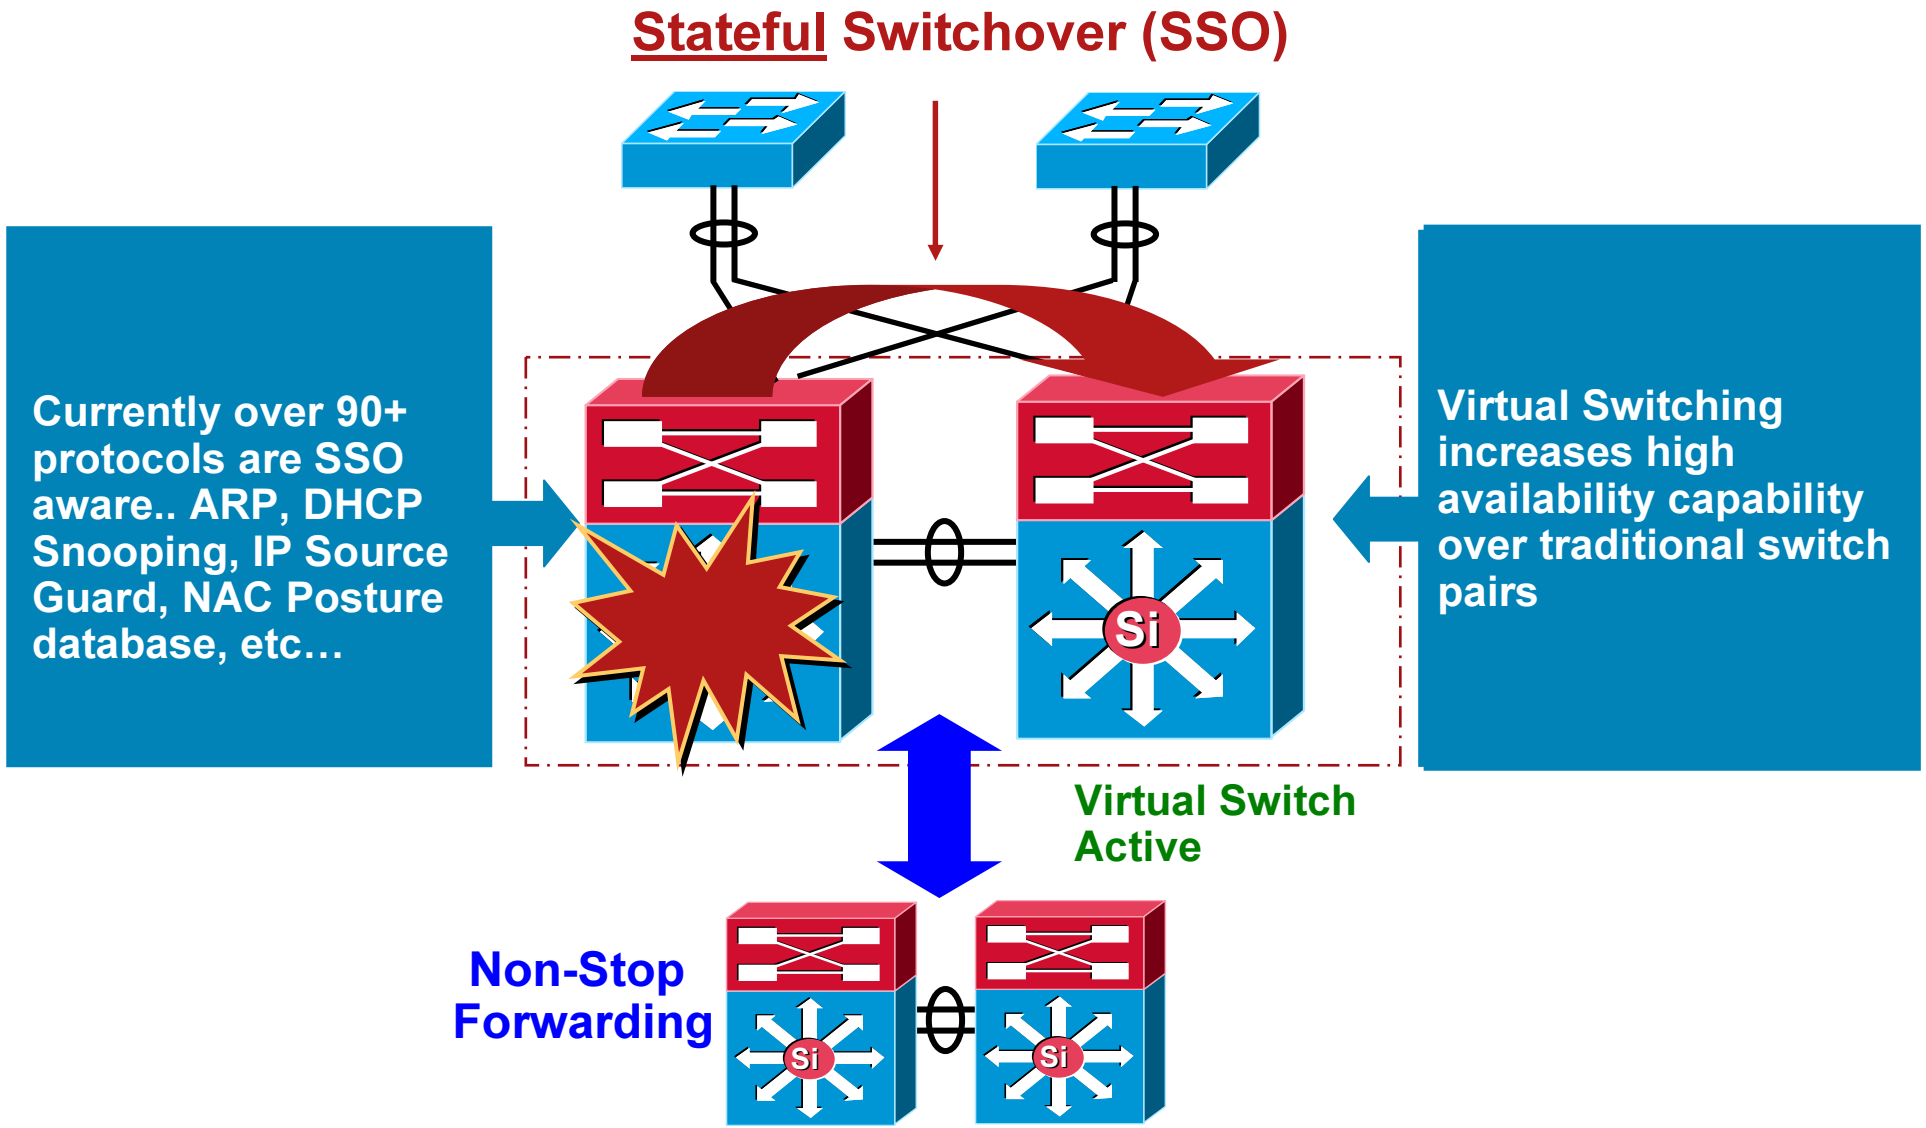

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System
Dual Active

Two mechanisms to recover if Virtual Switch Link Fails: PAGP +, IP-BFD
Either mechanism can be used to detect active peer.

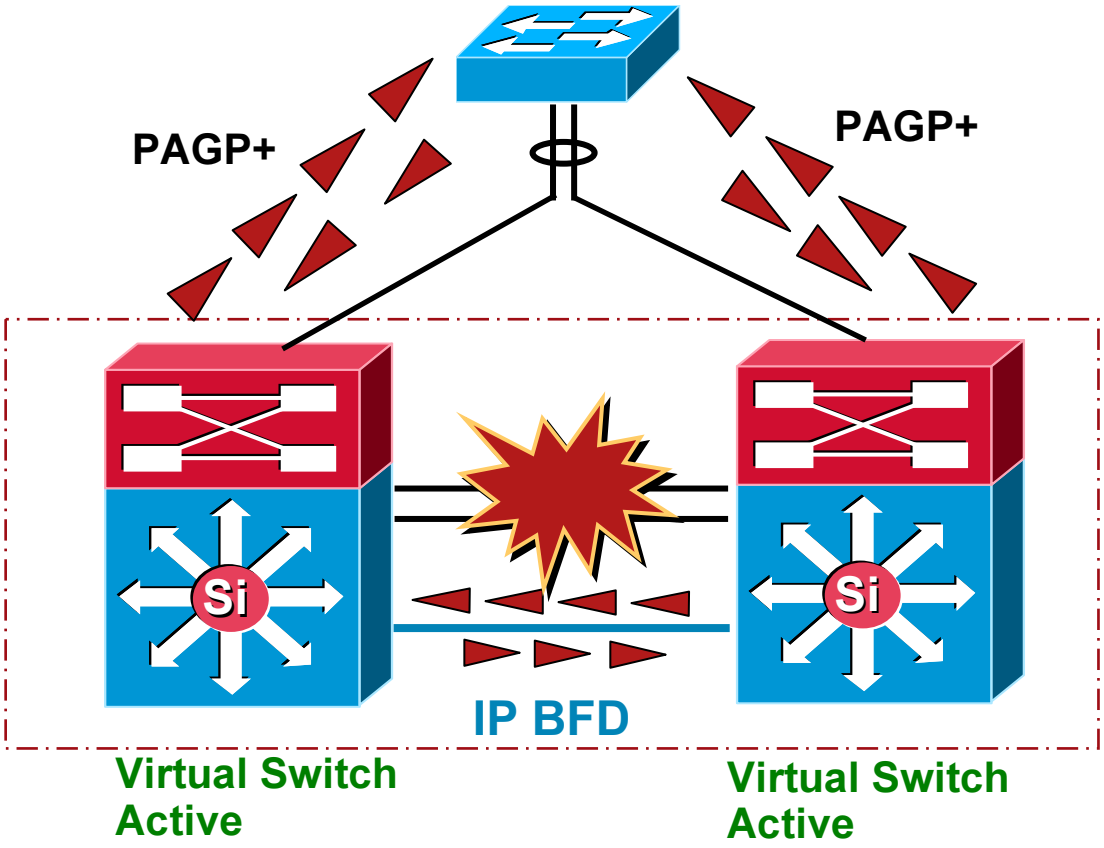

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Virtual Switch System
Deployment Considerations

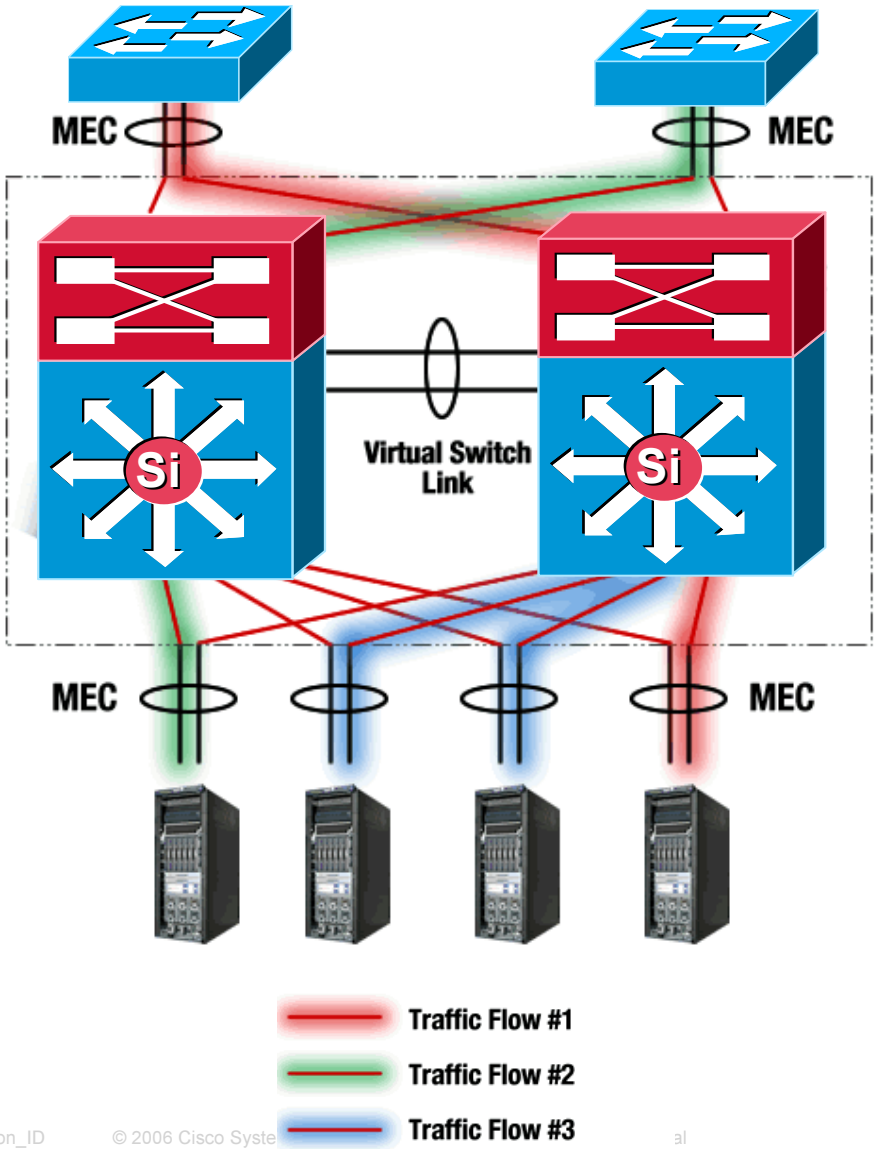

Deployment Considerations

Connected nodes should ALWAYS be dual homed

Etherchannel hash modified so local link always preferred over remote link in same bundle

ECMP has been modified to choose local link over remote link in same bundle

Up to 128 Etherchannel link bundles in virtual switch domain

Minimal traffic crosses VSL

Size number of links in VSL to meet local redundancy and traffic requirements

Hardware/Software Dependencies

Architectural Approach

Campus Communications Fabric

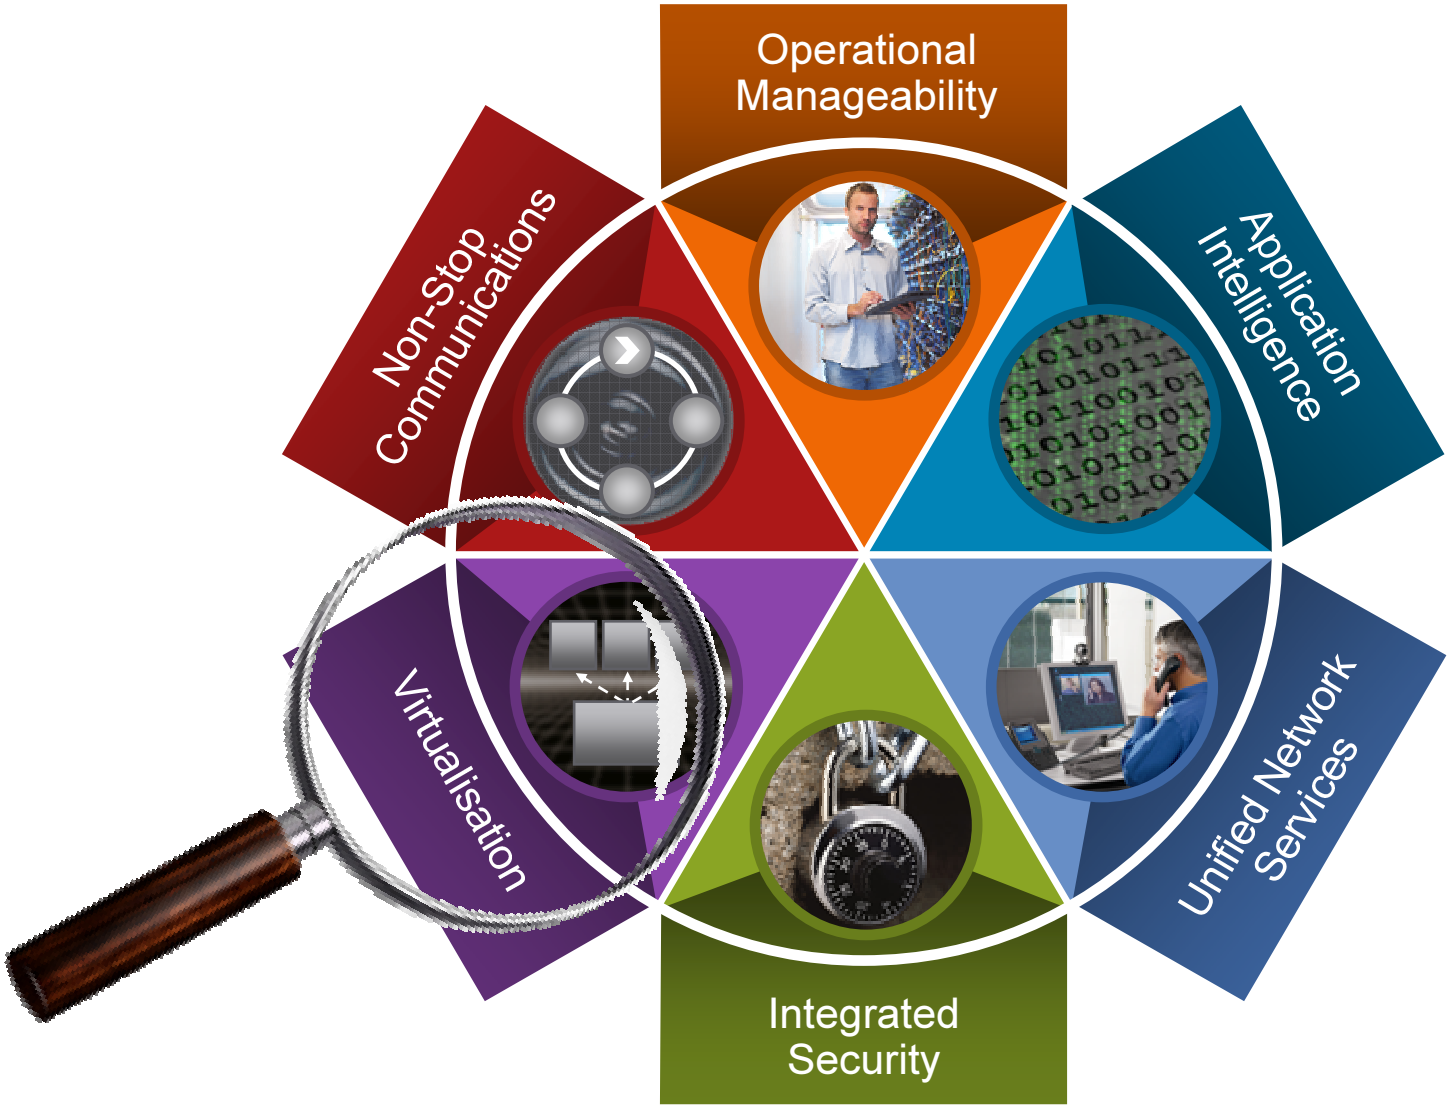

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation
Many to One

One network consolidates many physical networks

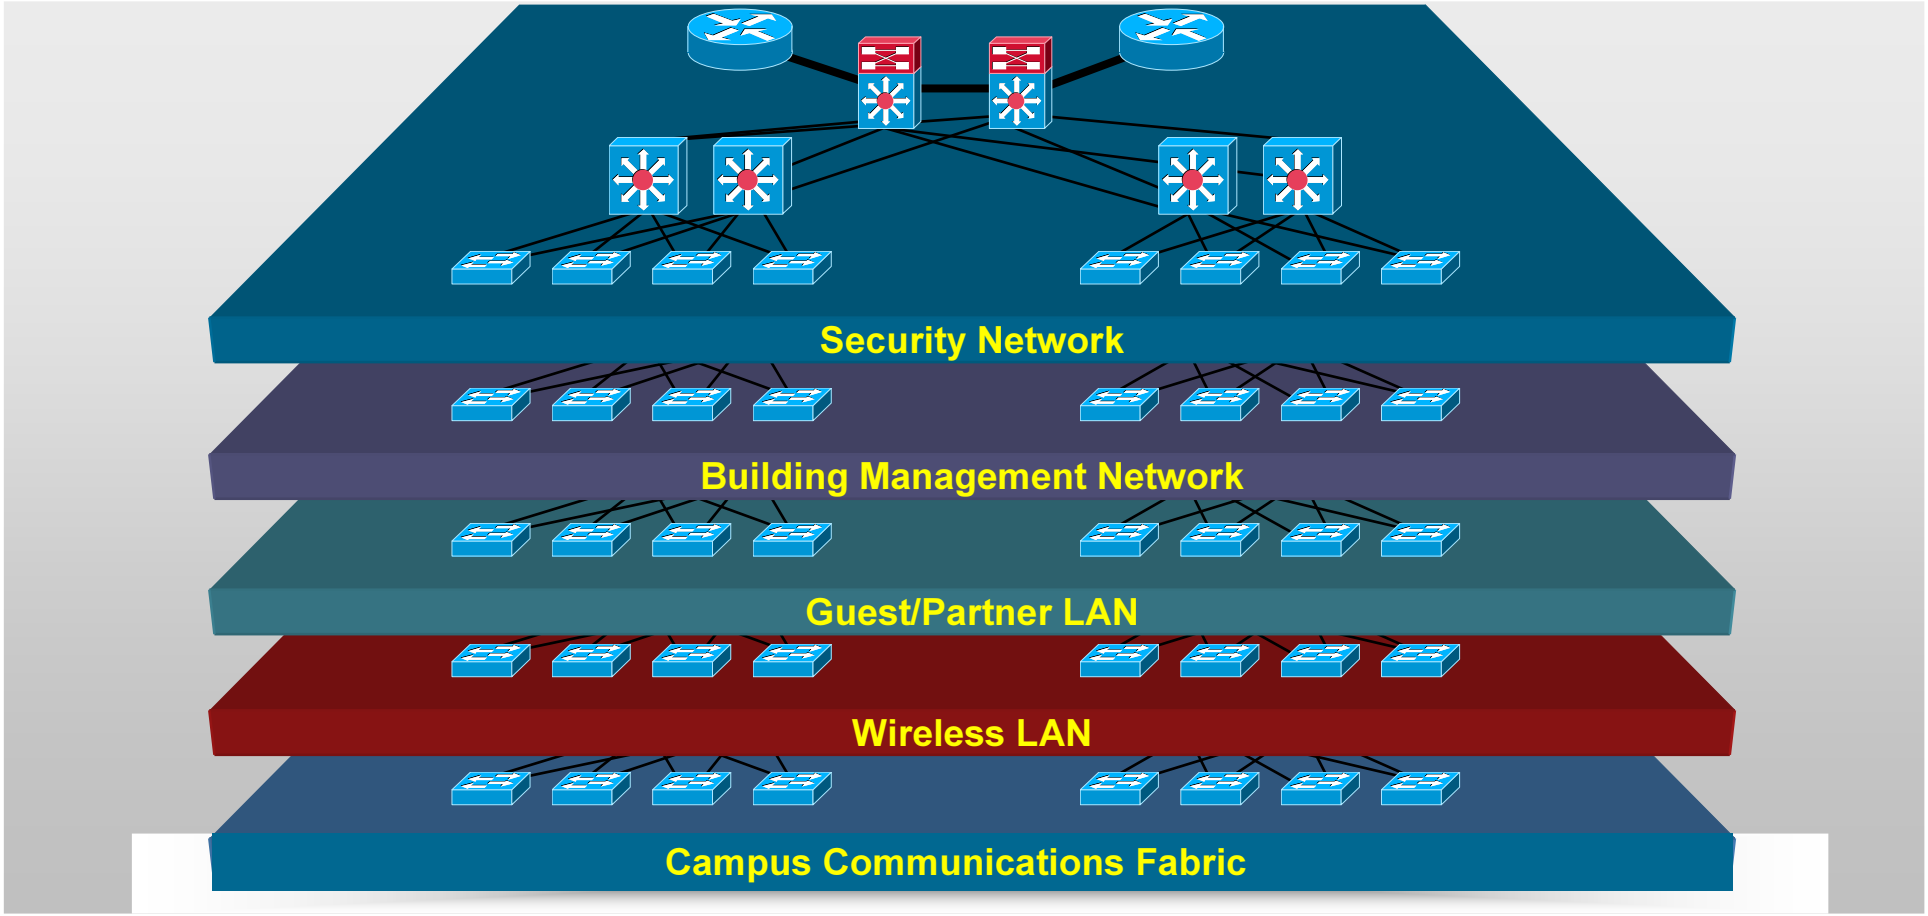

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Virtualisation
One to Many

One network supports many virtual networks

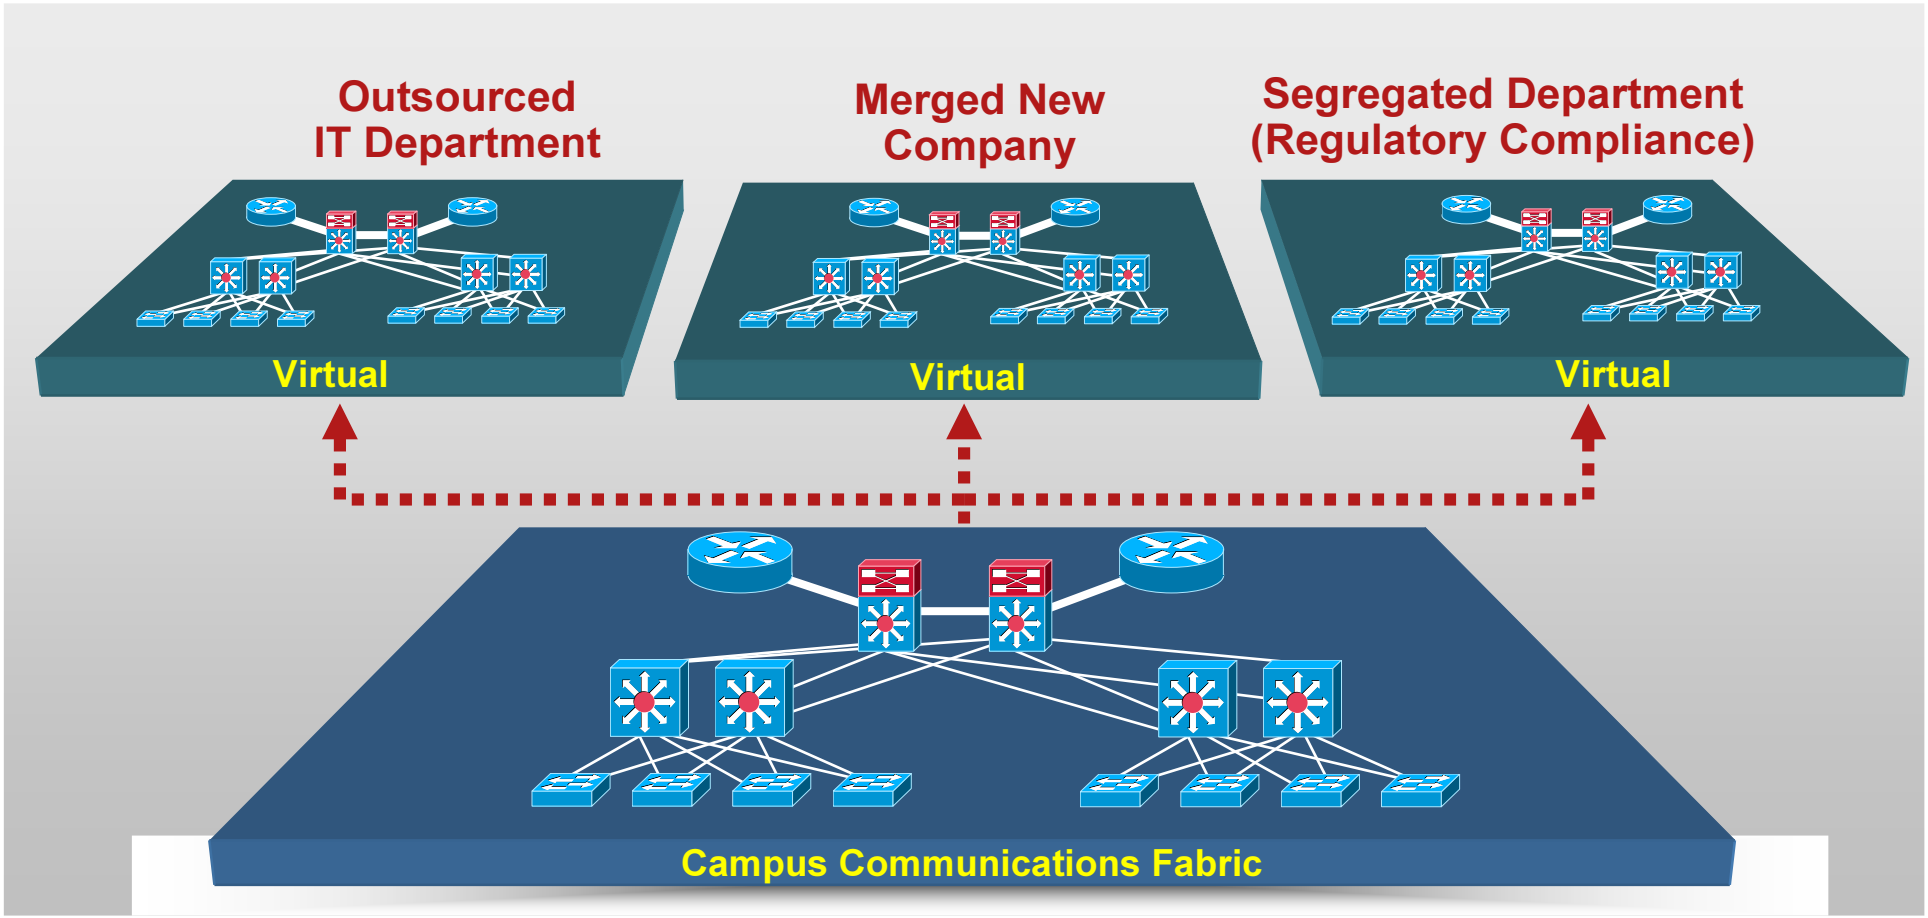

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

History Repeated

Network Virtualization History

- The Layer 3 Switch with t
  evolutionary step – to acl
  connectivity throughout t
  implemented to allow the
  devices – this provided v

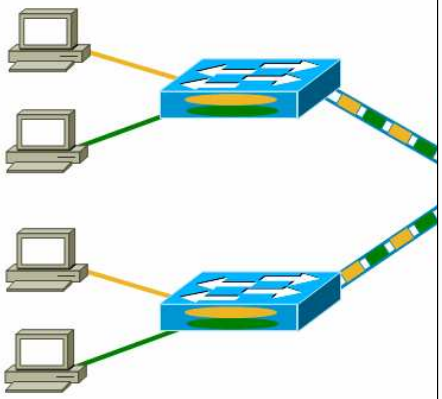

Network Virtualisation History

- High Availability requirements meant dual homing and associated complexities.

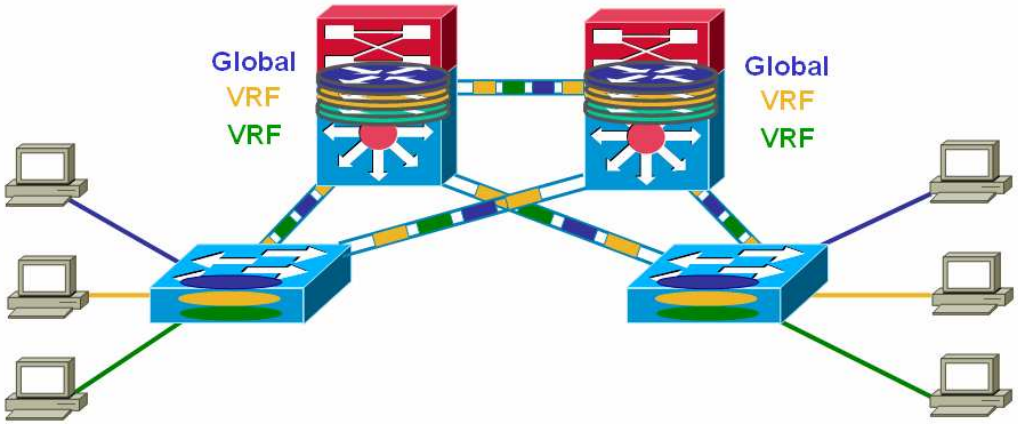

Data Path Virtualisation Techniques

Single Hop Data Path Virtualisation

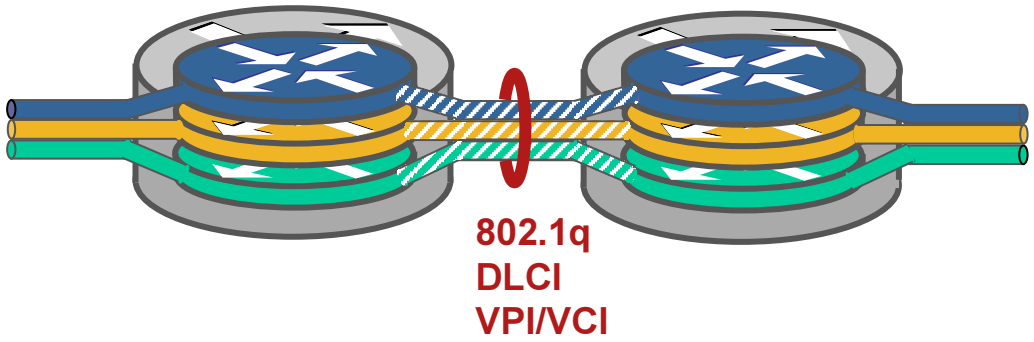

- Tags
  - 802.1q
  - Others (DSCP etc.)
  - More about Tags Later !
- Virtual circuits
  - ATM
  - Frame Relay
  - AToM L2 Circuits

Multi-Hop Data Path Virtualisation

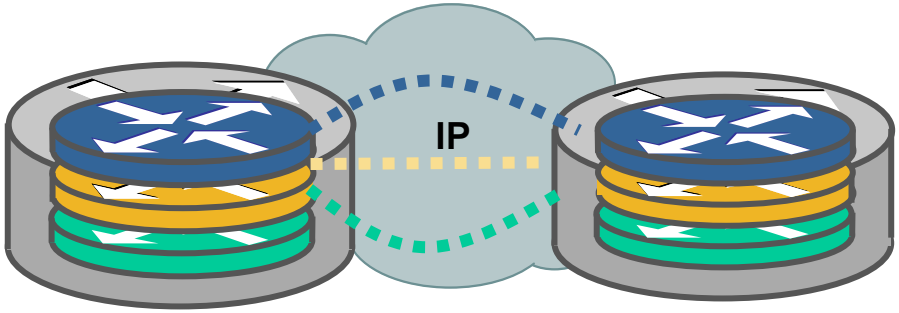

- Tunnels
  - GRE/mGRE
  - L2TPv3
  - Label Switched Paths—LSP (MPLS)
- Future Tag Transports
  - More about Tag Transports Later !

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

End to End Virtualised Groups
Overlay Tunnels

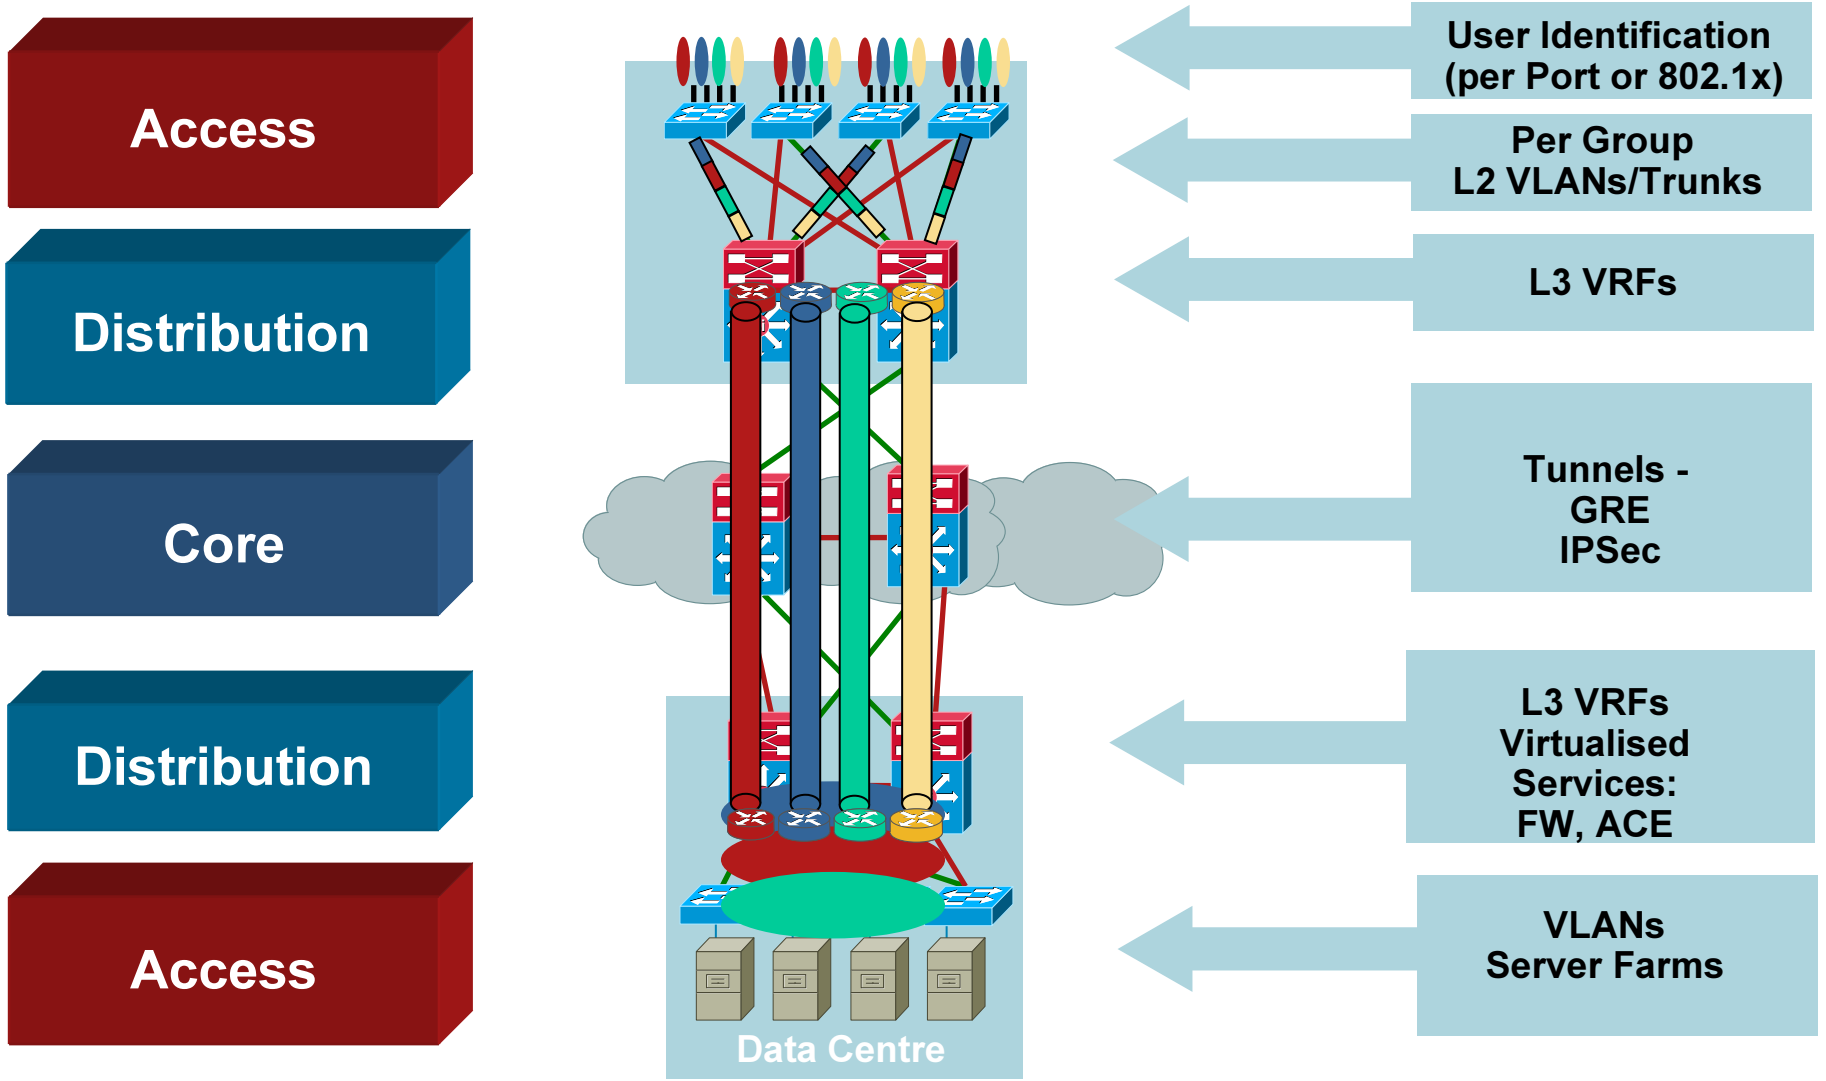

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

End to End Virtualised Groups
Hop By Hop

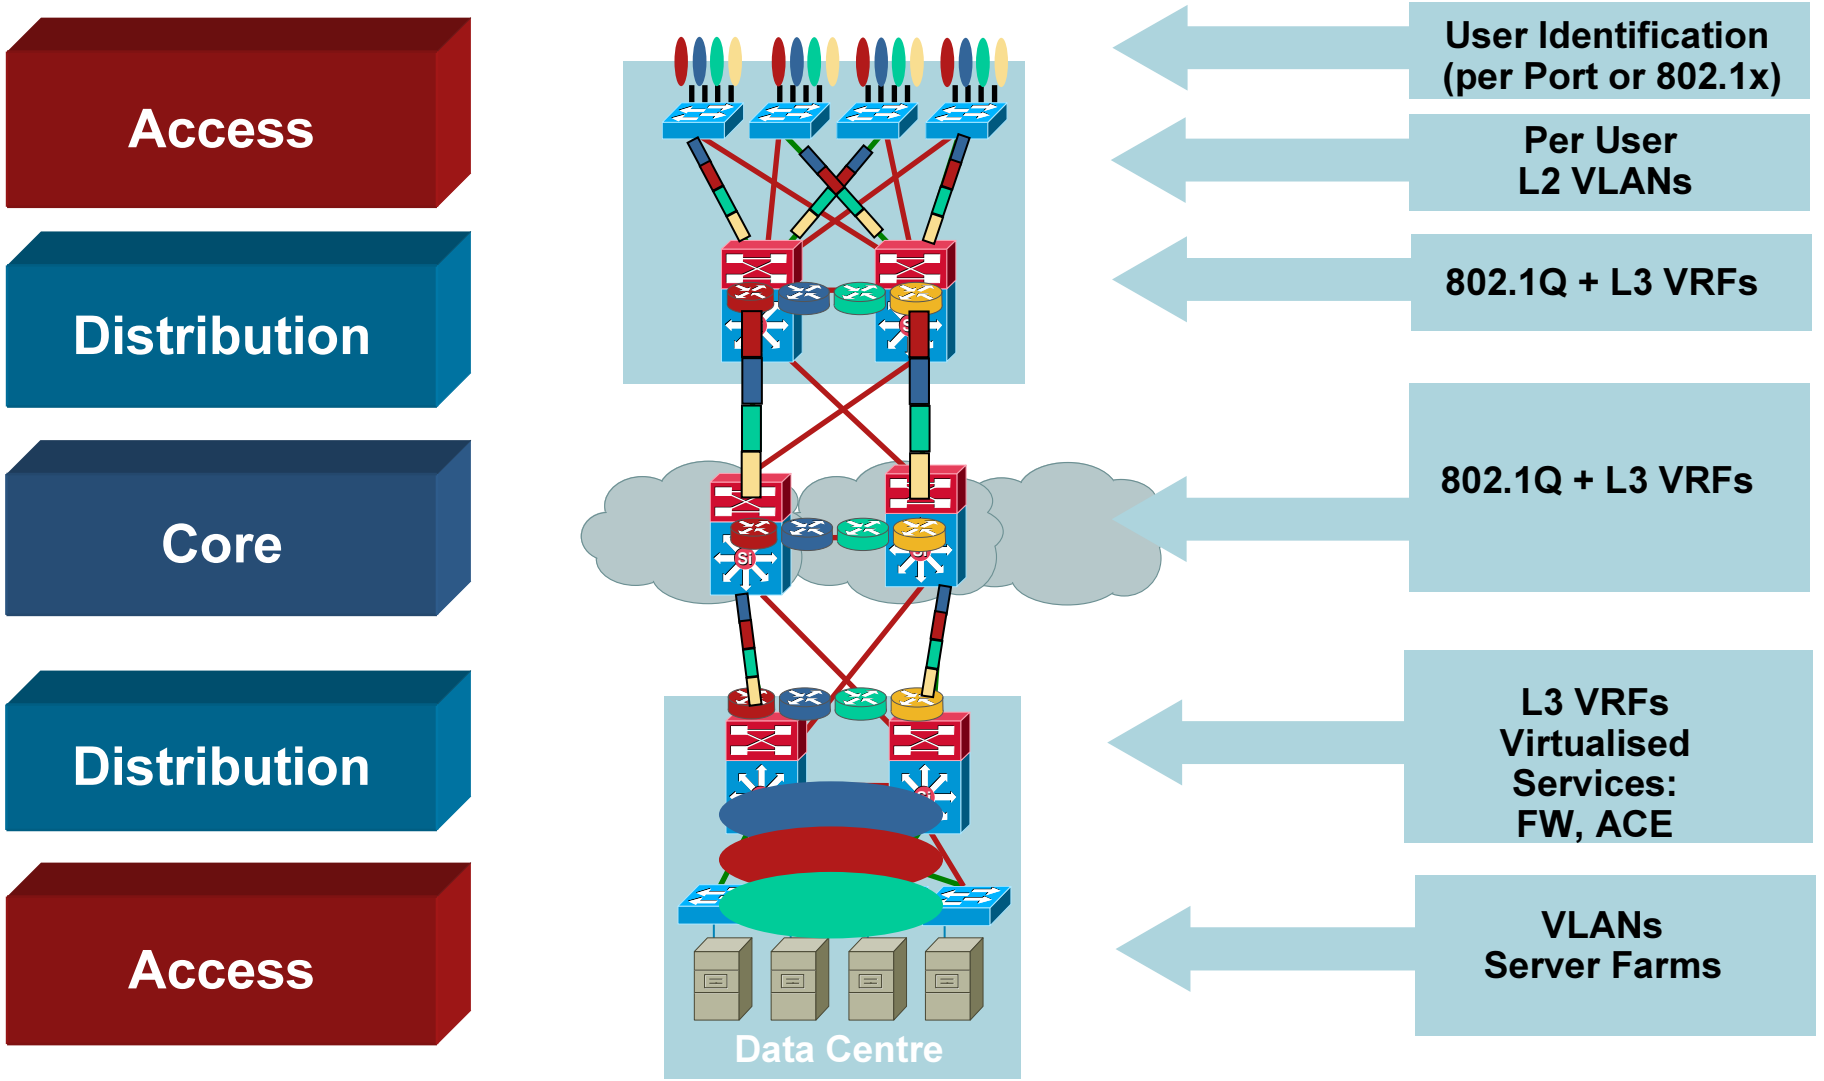

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

End to End Virtualised Groups
Distributed Access Control Lists

Access

Distribution

Core

Distribution

Access

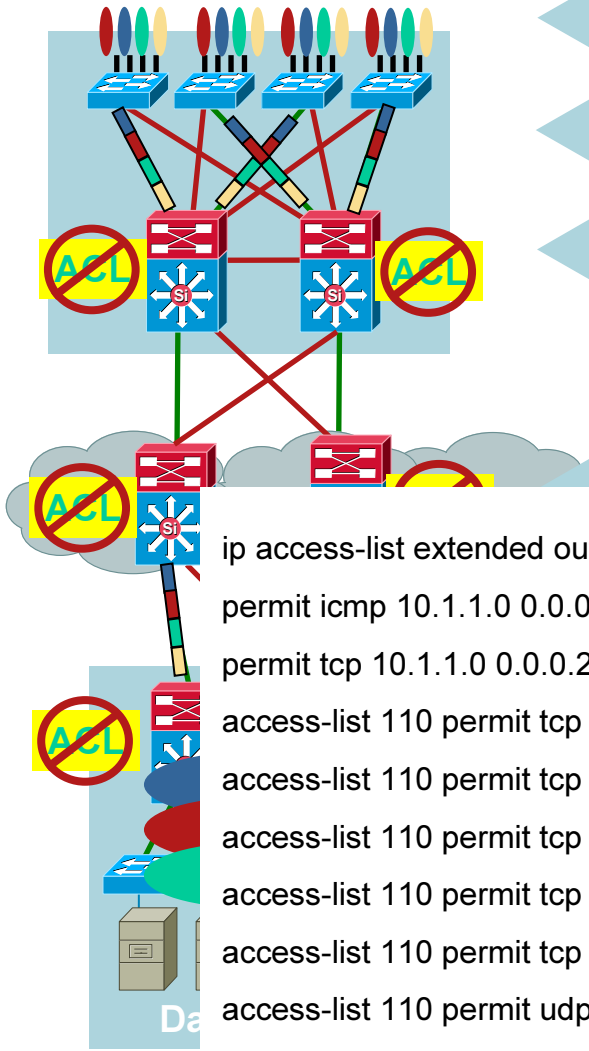

User Identification (per Port or 802.1x)

Per User L2 VLANs

Distributed ACL

Distributed ACLs

```
ip access-list extended outboundfilters
permit icmp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255
permit tcp 10.1.1.0 0.0.0.255 172.16.1.0 0.0.0.255 reflect
access-list 110 permit tcp any 63.67.50.0 0.0.0.255 eq pop3
access-list 110 permit tcp any 63.67.50.0 0.0.0.255 eq 143
access-list 110 permit tcp any 63.67.50.0 0.0.0.255 eq ftp-
access-list 110 permit tcp any 63.67.50.0 0.0.0.255 eq ftp
access-list 110 permit tcp any 63.67.50.0 0.0.0.255 gt 1023
access-list 110 permit udp any 63.67.50.0 0.0.0.255 gt 1023
access-list 110 deny ip 63.67.50.0 0.0.0.255 any
```


© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidenti

End to End Virtualised Groups
Static Topology Bound

Access

Distribution

Core

Distribution

Access

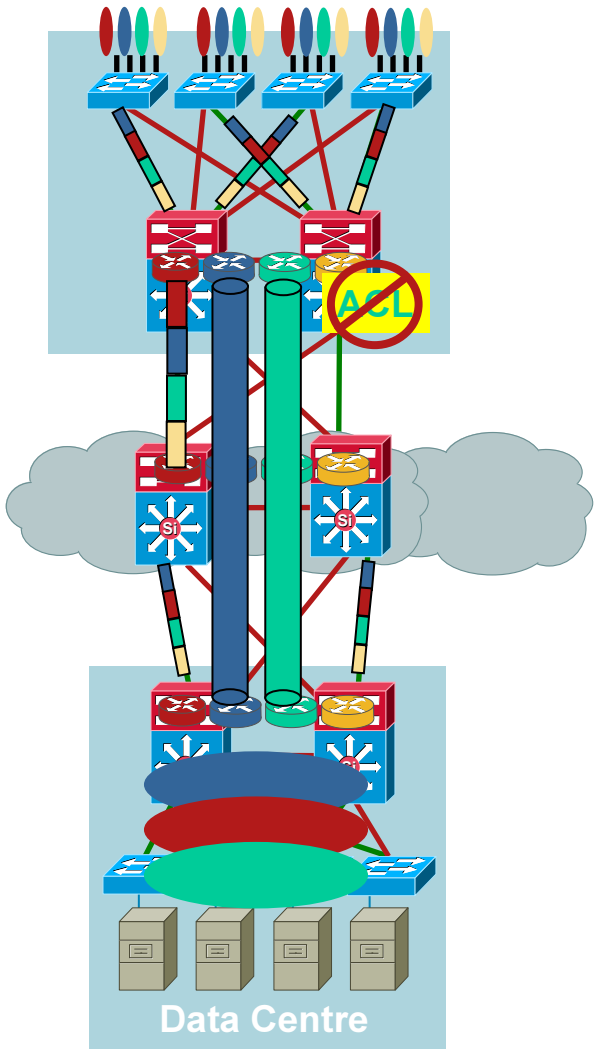

All VLANs must appear on all access switches for mobility
Prone to configuration error
Not scalable manageable

VLANs may have to be transported manually across L3 boundaries per VRF

Overlay tunnels obfuscate traffic and bypass policy

Any changes to policy requires changes to entire network

May suit specific requirements

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

End to End Virtualised Groups
Static Topology Bound

Access

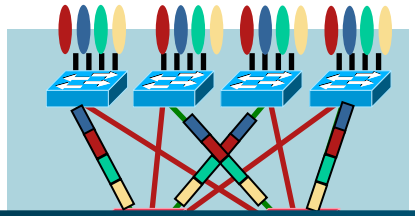

All VLANs must appear on all access switches for mobility
Prone to configuration error
Not scalable manageable

VLANs may have to be

Static Topology-Bound Separation & Policy Enforcement

Distribution

Access

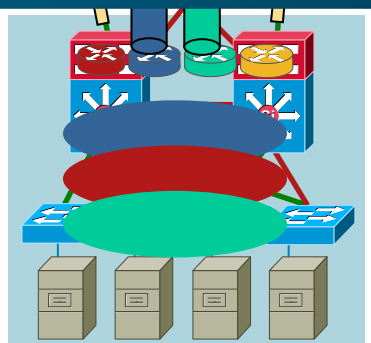

Any changes to policy requires changes to entire network

May suit specific requirements

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Architectural Approach to Policy Enforcement

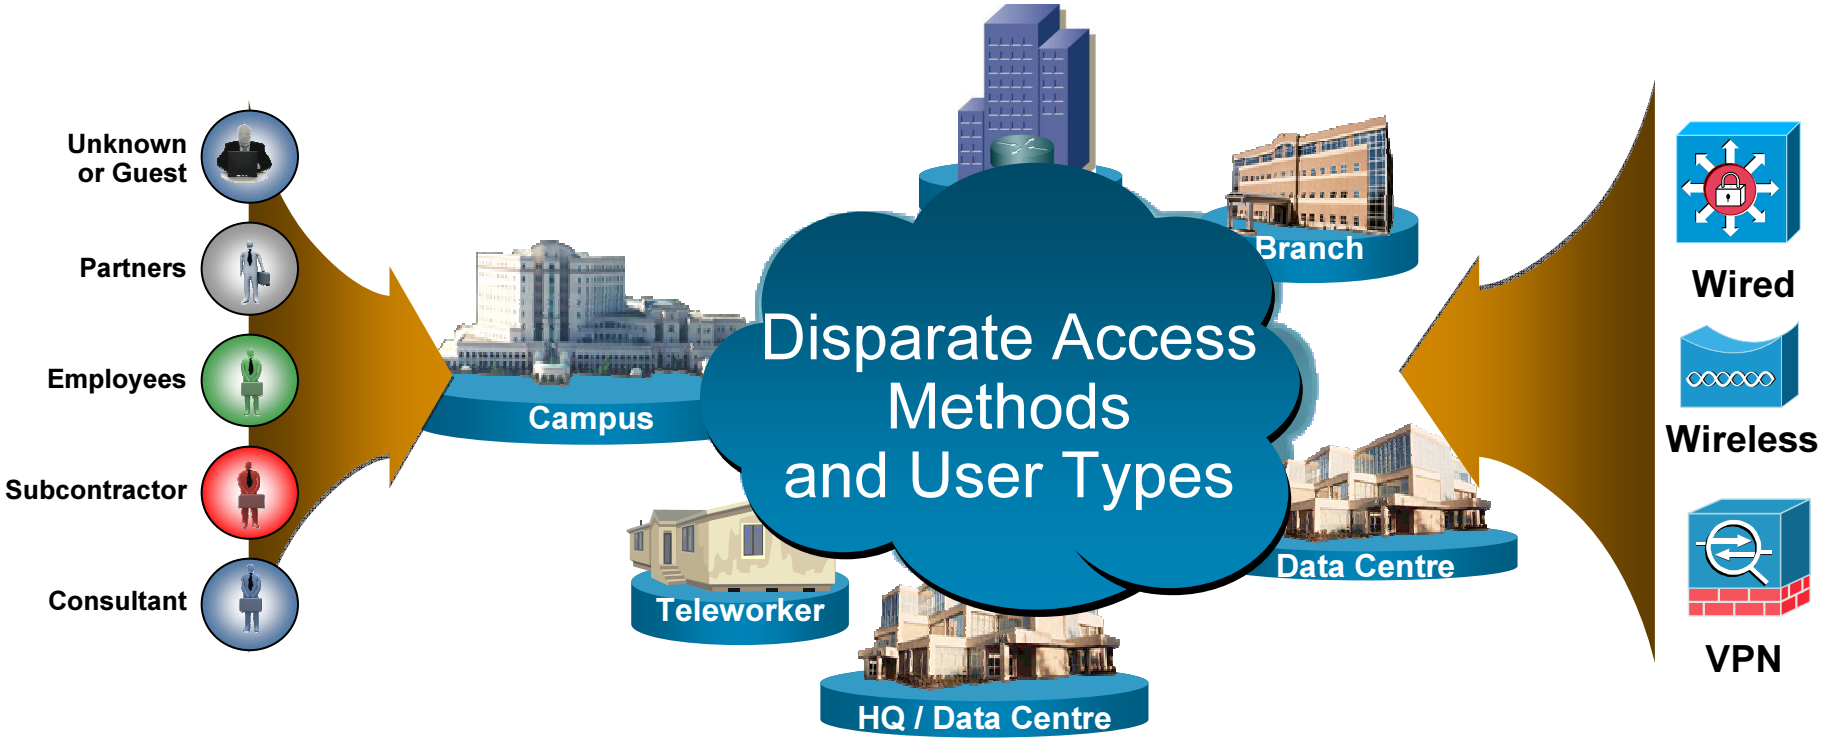

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Architectural Approach

Campus Communications Fabric

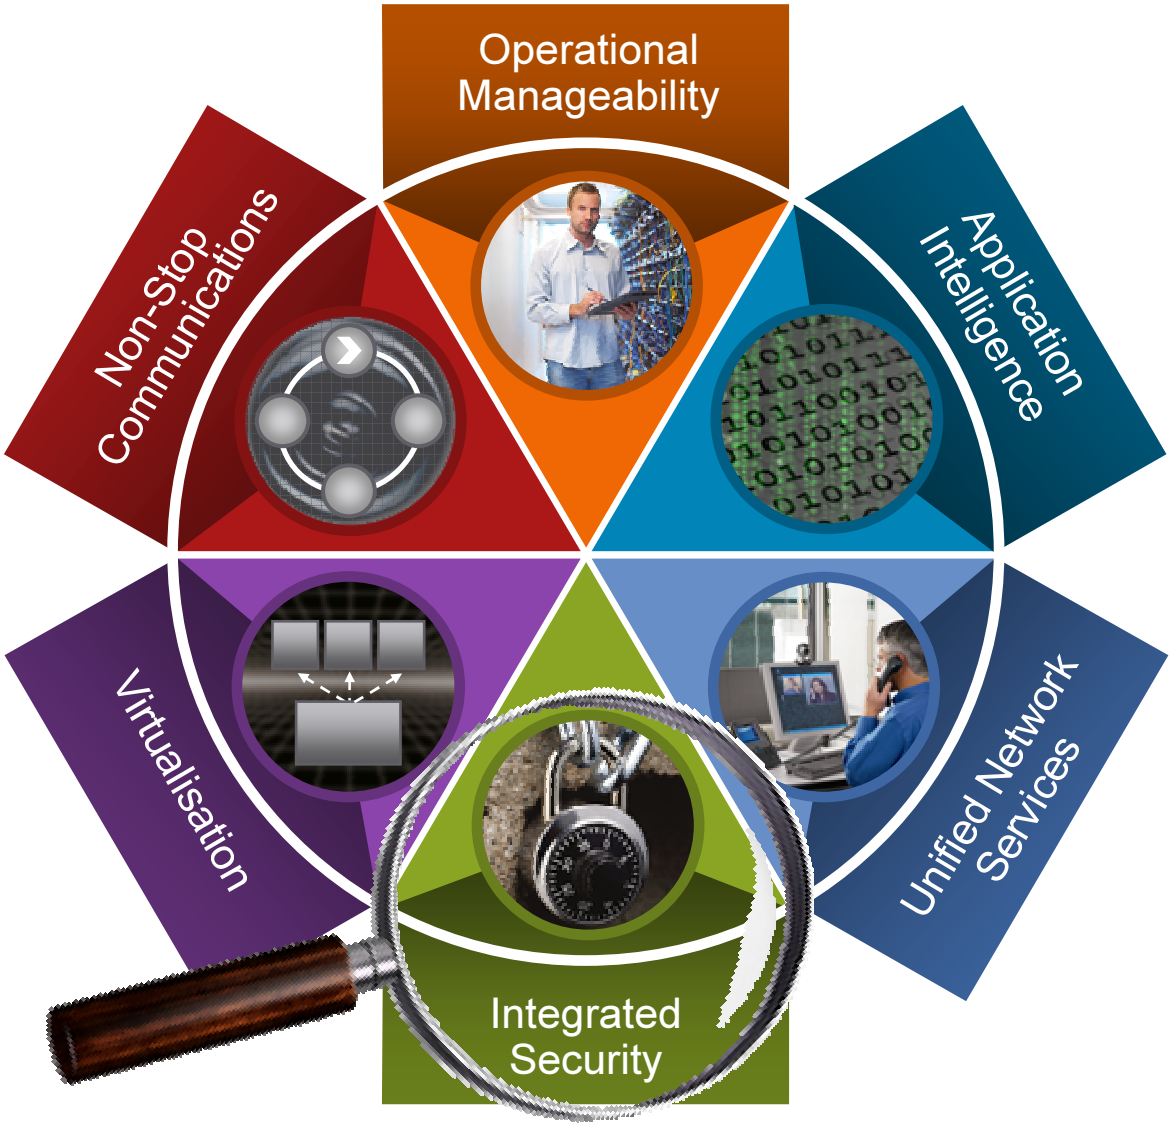

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Cisco TrustSec (CTS) Vision
An Architectural Approach

Seamless Authentication for Various Access Types	Secure Campus Access Control	Converged Policy Framework	Converged Policy Definition Enforced via Network
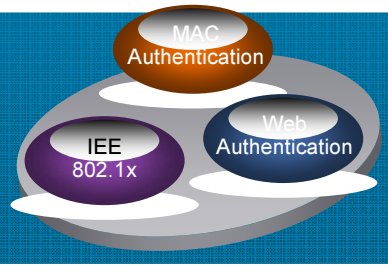			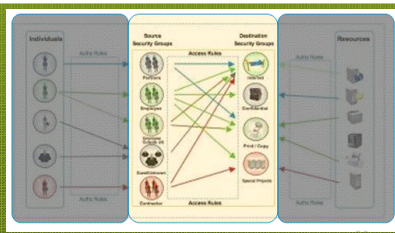
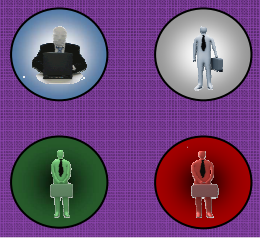			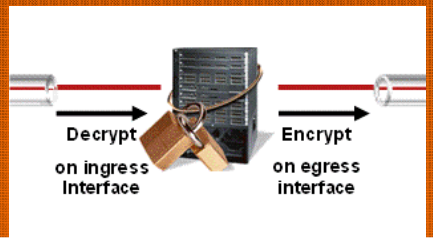
Transforming From Topology-aware to Role-aware Access Control	**Role Aware Network**	**Trust, Integrity & Confidentiality**	**Prevent Data Sniffing and Tampering with Line-rate Hop-by-hop Encryption**

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Policy Enforcement Challenges

Traditional Discretionary Access Control

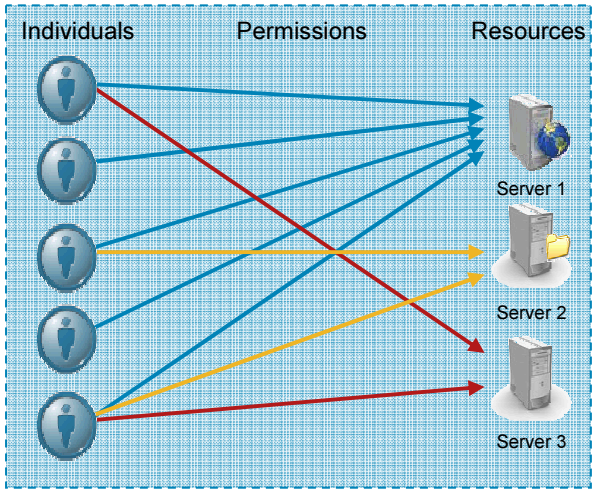

Access List for S1

```
access-list 101 permit tcp S1/32 D1/32 eq http
access-list 101 permit tcp S1/32 D1/32 eq https
access-list 101 permit tcp S1/32 D2/32 eq ftp
access-list 101 permit tcp S1/32 D2/32 eq http
access-list 101 permit tcp S1/32 D2/32 eq https
access-list 101 permit tcp S1/32 D2/32 eq ftp
access-list 101 permit udp S1/32 D1/32 gt 1023
access-list 101 permit udp S1/32 D2/32 gt 1023
```

Challenges

- Leads to ACE explosion
  - (# of sources) X (# of Destinations) X (# of permissions)= # ACEs
- IP-address based ACLs are challenging
  - Changes in addressing schemes
  - Use of DHCP
  - Proliferation of Wireless LAN devices
- Assumes relatively static placement of users/resources

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Cisco TrustSec User Authorization and Access Control

- Define Security Groups
- Users and Resources Sessions are Authorized via flexible ABAC model
- Access Control Policies are created without regards to Network Topology (No IP Addresses or subnets)
- Access Control Policies are mapped between source and destination Security Groups via a Matrix
- At runtime user's traffic carries the Security Group Tag (SGT) in every packet
- These SGTs are are filtered (i.e., SGACLs) processed at wirespeed on egress devices

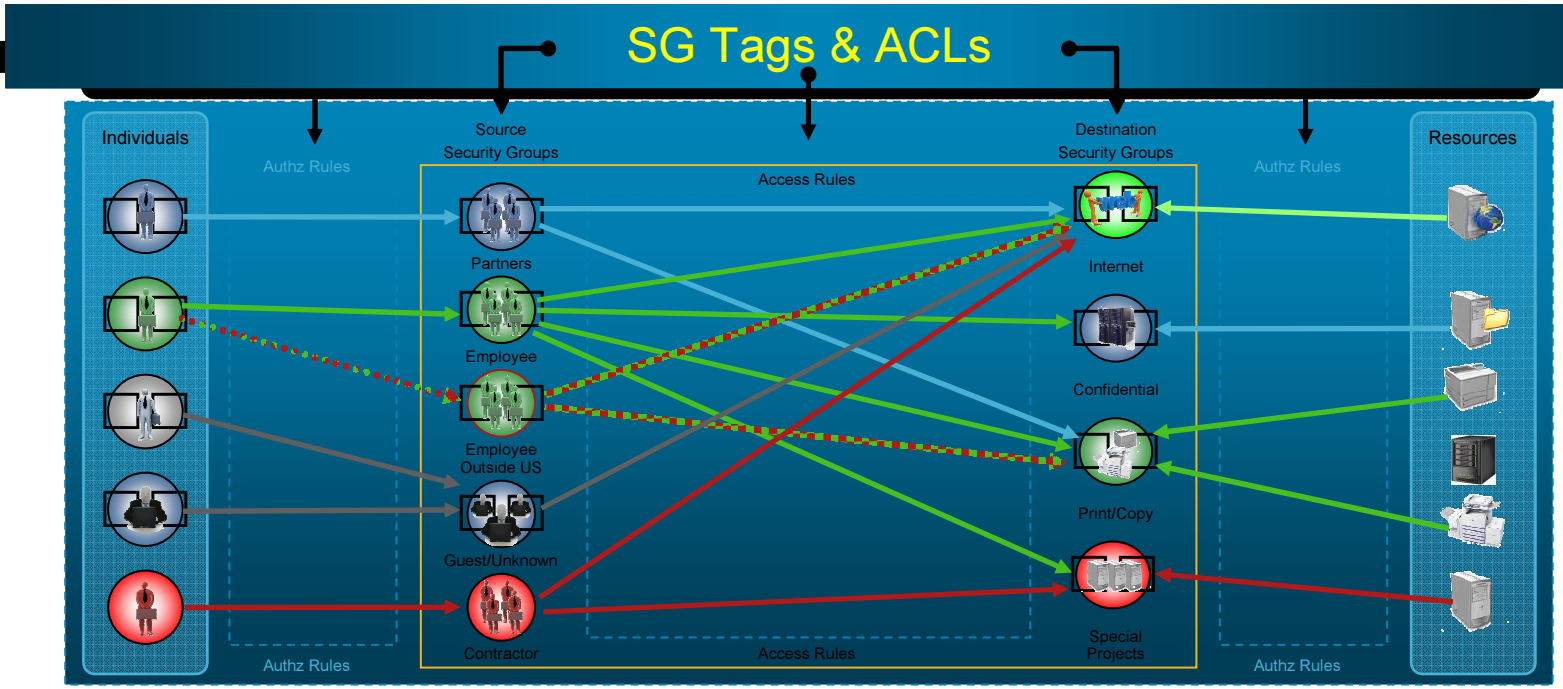

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Cisco TrustSec User Authorization Policy
Attribute Based Access Control

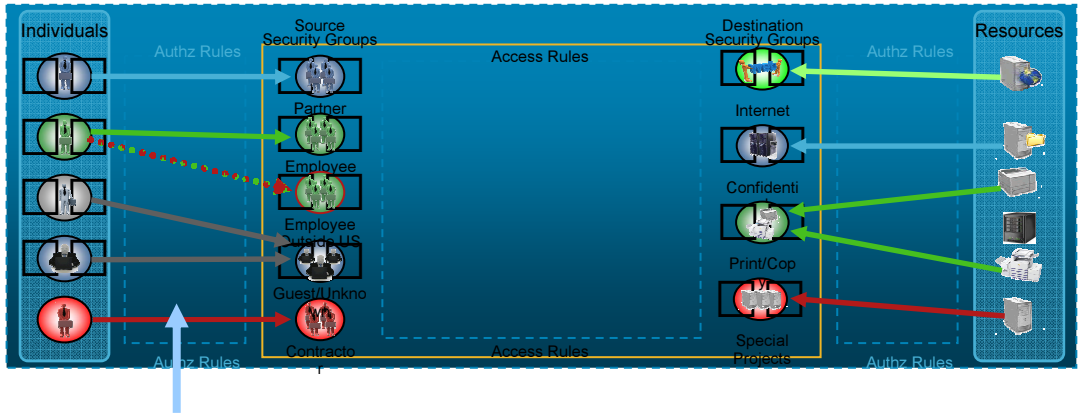

Users and Resources are grouped into logical *topology independent* security groups

- Ingress Session Authorization Example

Authorization Rule : if ((identity group = Teller) & (location = Sydney) & (access-type = LAN) apply SGT-4

Authorization Rule : if ((identity group = Teller) & (location = Sydney) & (access-type = WLAN) apply SGT-9

Authorization Rule : if ((identity group = Teller) & (location = Sydney) & (access-type = VPN) apply SGT-22

Authorization Rule : if ((identity group = Teller) & (location = Russia) & (access-type = LAN) apply SGT-50

Authorization Rule : if ((identity group = Teller & Manager) & (location = Russia) & (access-type = LAN) apply SGT-51

Authorization Rule: if ((identity group = LoanApp) & (location == DC-Sydney) then apply SGT-1

Access Control is then defined by ACL between Source & Destination Groups

RBACL1 : s(SGT-4) d(SGT-1) permission-list A

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Policy Enforcement Throughout the Network: Role Based Access Control Set-up

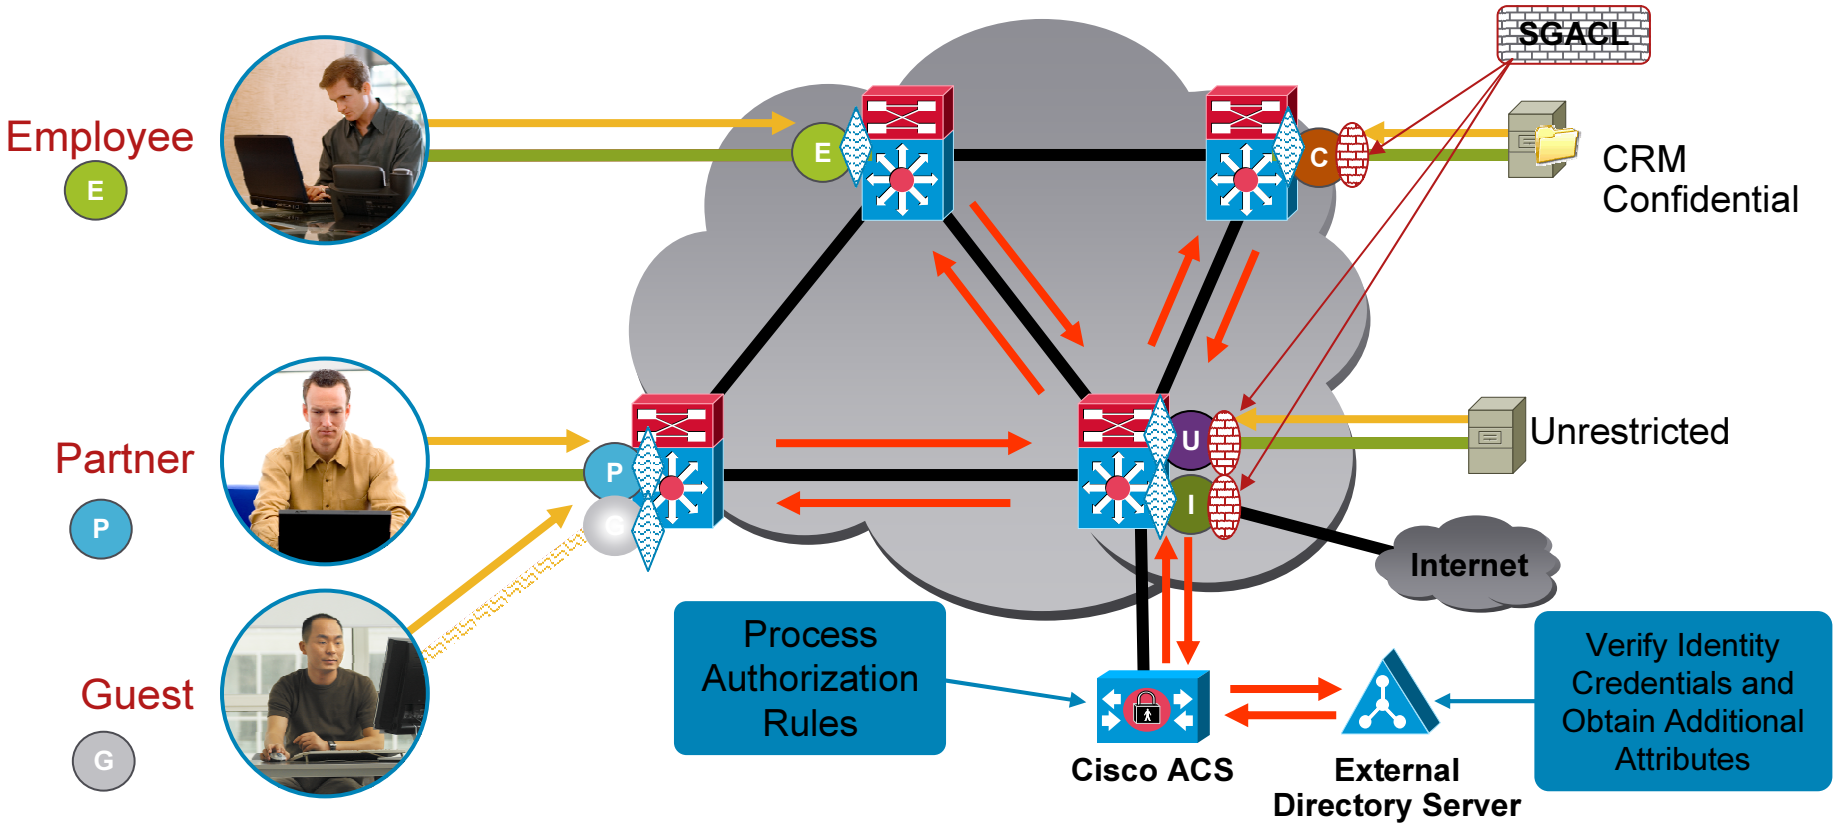

Legend

Link/Port Status		Security Group Classifications	
Unauthenticated	Ingress Tagging	E Employee Group	C Confidential Group
Failed Authentication		P Partner Group	U Unrestricted Group
Authenticated	Egress Filtering	Guest Group	I Internet Group
Shutdown			

1. **Authentication Request**
2. **Radius and AD Authc/Authz**
3. **Group Membership Dynamically Assigned**
4. **SGACL Dynamically Applied**
5. **Links Up**

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Policy Enforcement Throughout the Network: Role Based Access Control Deployment

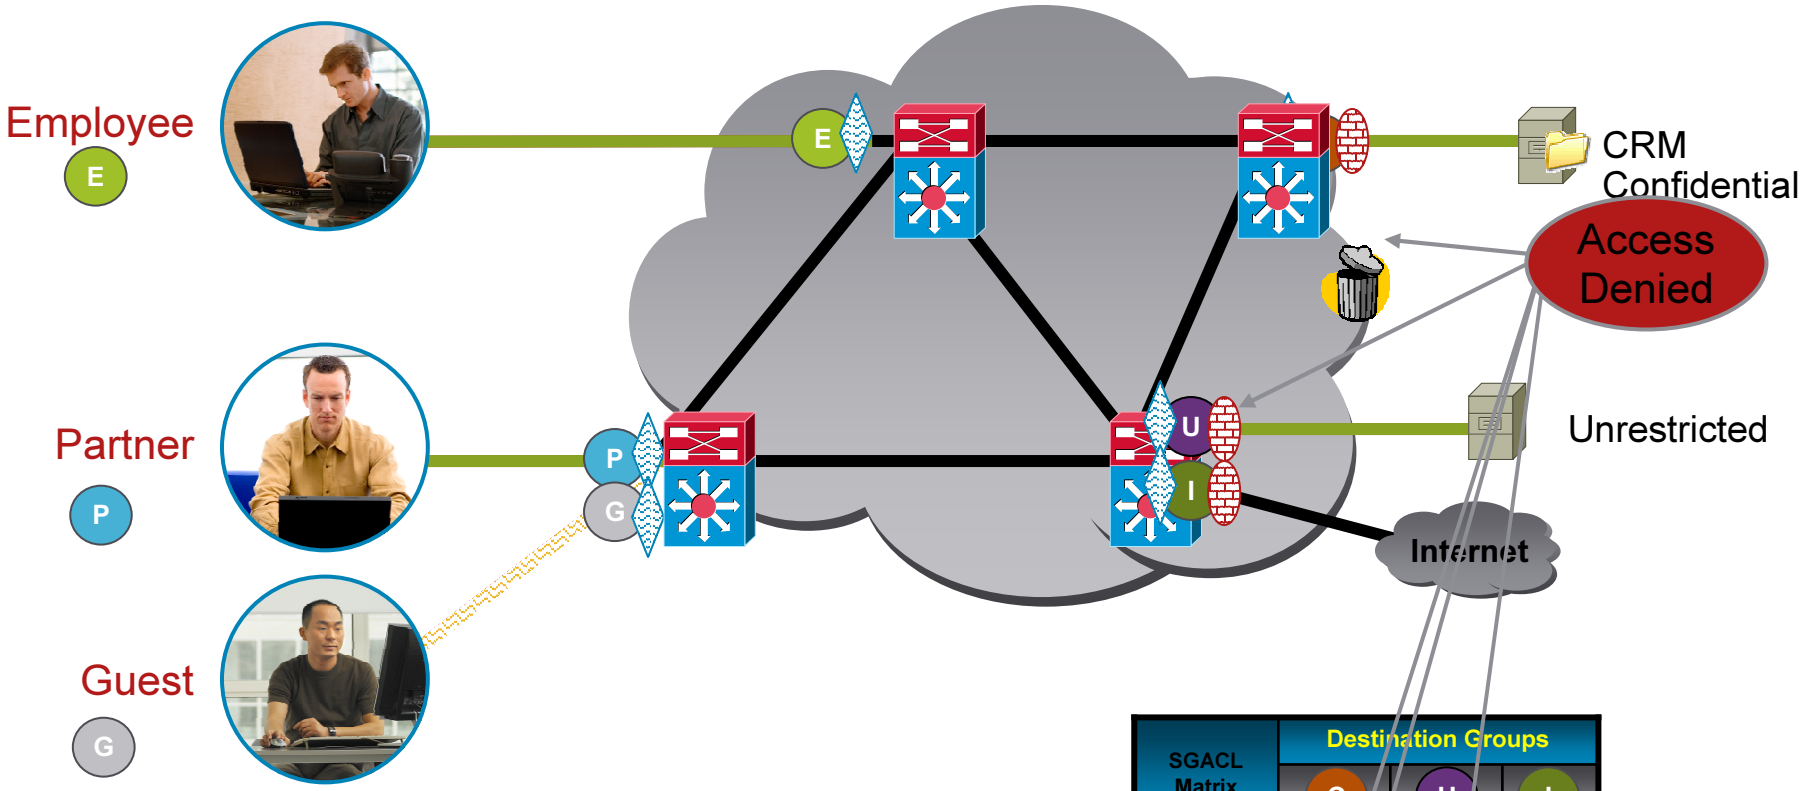

Legend

Link/Port Status		Security Group Classifications	
Unauthenticated	Ingress Tagging	E Employee Group	C Confidential Group
Failed Authentication		P Partner Group	U Unrestricted Group
Authenticated	Egress Filtering	Guest Group	I Internet Group
Shutdown			

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Network Device Admission Control (NDAC)

Authenticate network devices to trusted node via 802.1X prior to any network connectivity becoming available. Polices are then acquired as well as any applicable session keys or cipher suites…

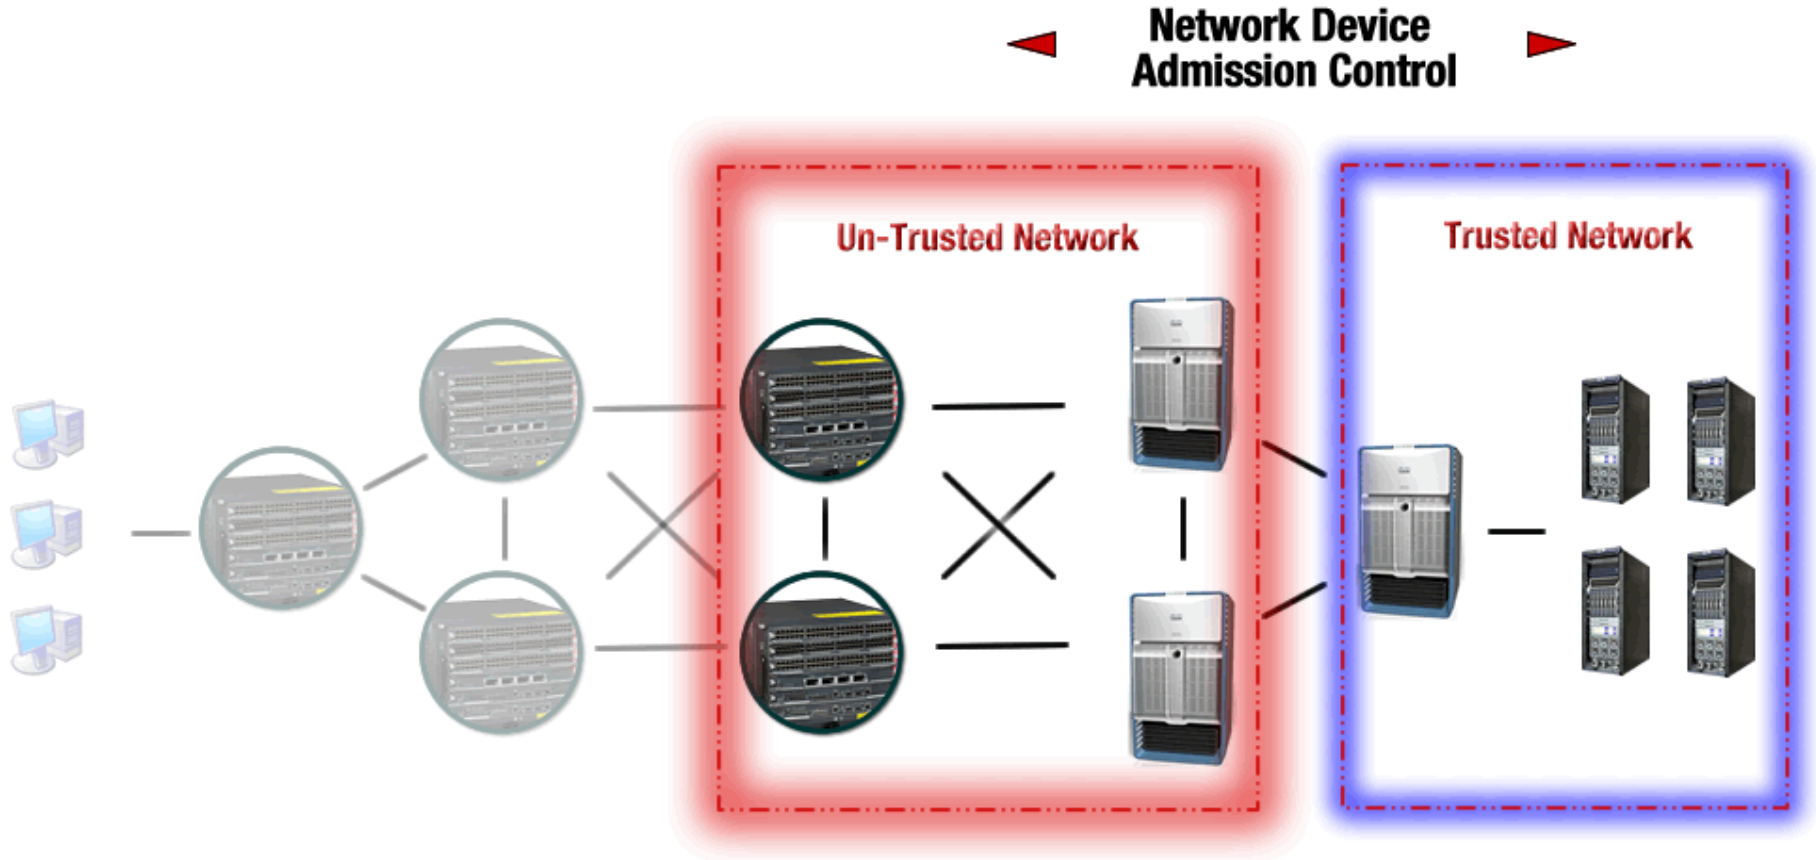

TrustSec
An Architectural Approach

CTS NDAC
Trusted Network Infrastructure

- Mutual Device Authentication
- Policy Acquisition by device pairs
- Environmental Variables

L2 TrustSec Provisioning
Security Association Protocol (SAP)
4-way handshake

- Negotiate Cipher Suite
- Exchange SCI, Offset for data traffic
- Derive Session Keys to encrypt the data traffic
- 'Ready to Receive' & periodic rekey

Data Path Processing

- Encrypt/decrypt on ingress/egress
- Secure Data Communications & SGT Transport
- ASIC-based 1000Mbps & 10GbE
- Built into the port ASICs on L2 switch ports

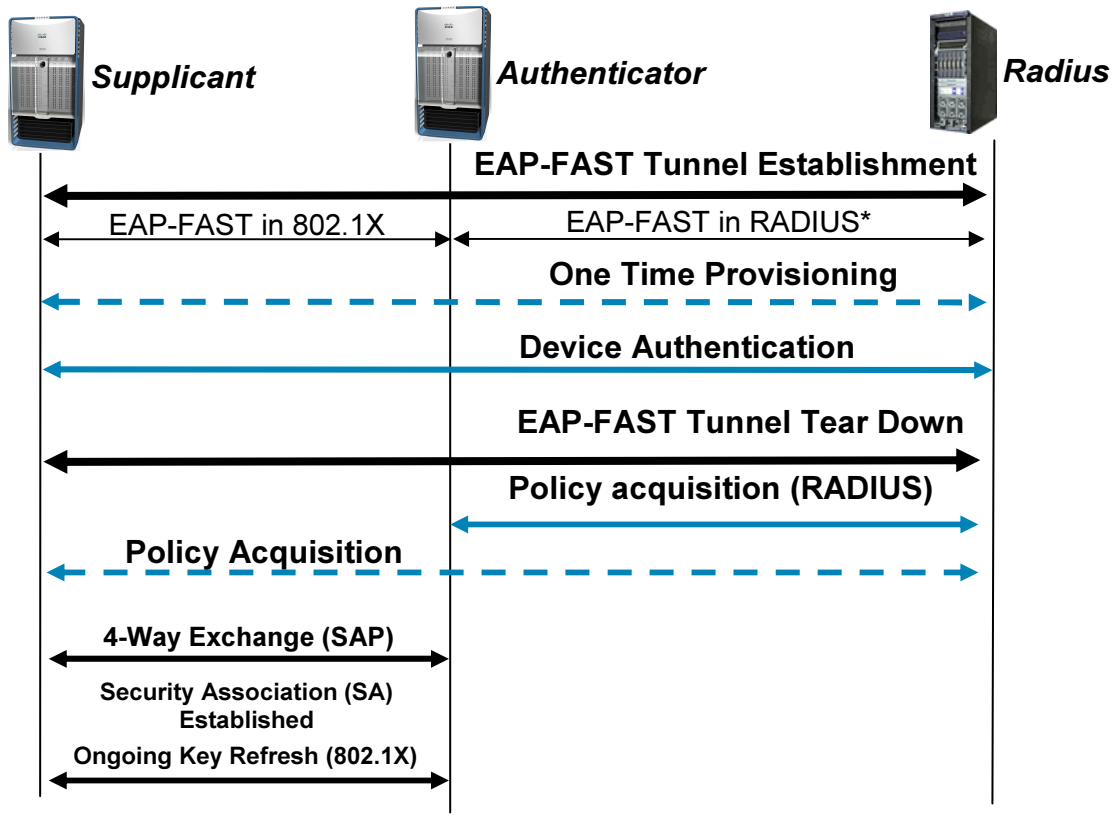

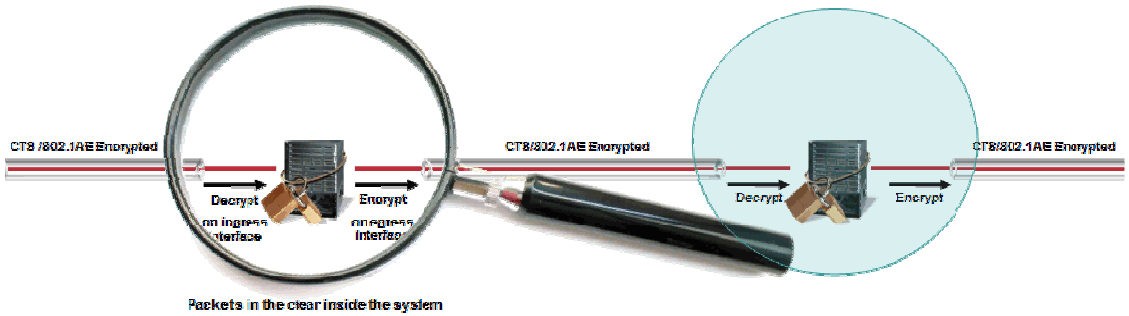

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Link Layer Encryption
An Architectural Approach

- Hop-by-Hop packet confidentiality and integrity via IEEE 802.1AE
- Packets are encrypted on egress, Packets are decrypted on ingress
- Packets are in the clear in the device allowing the network to continue to perform all the packet inspection features currently used
- Can be incrementally deployed depending on link vulnerability

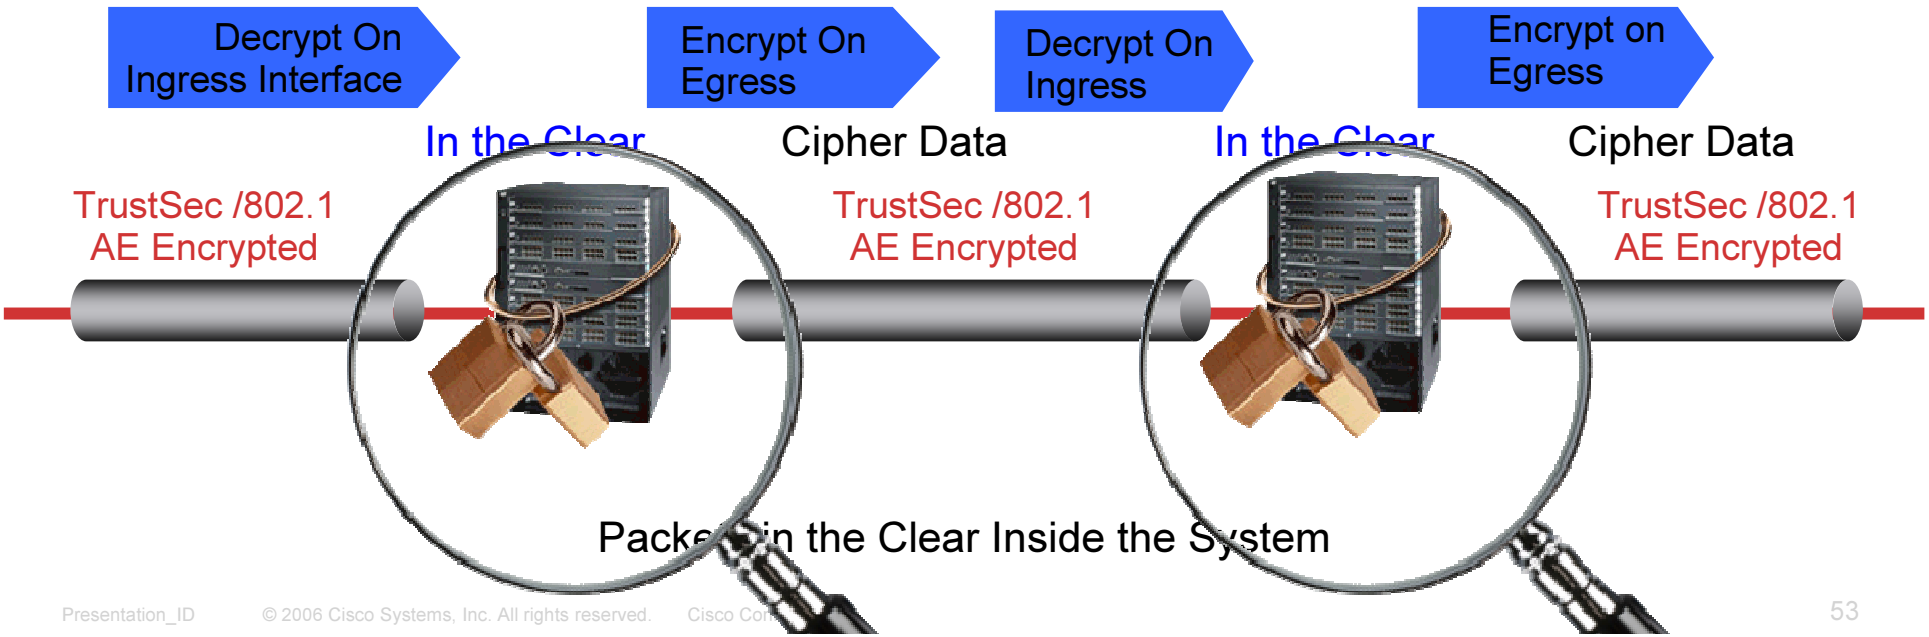

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Con

Security Group Tag Exchange Protocol (SXP)
Migration for Access-Layer

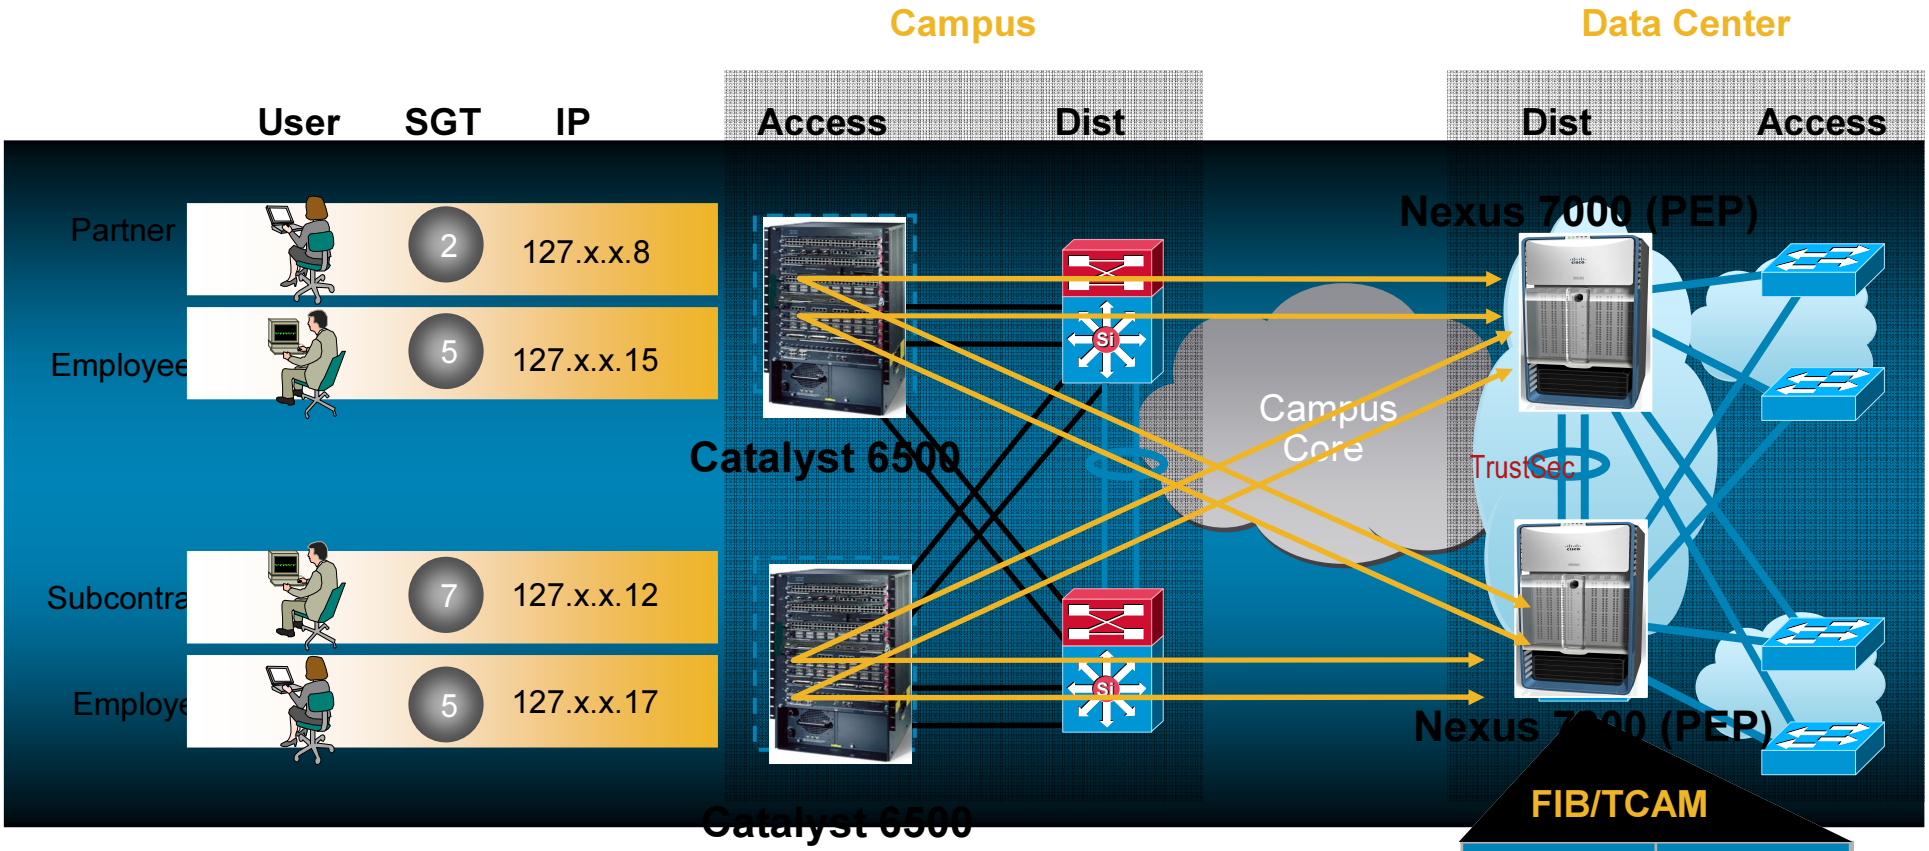

IP	SGT
127.x.x.8	2
127.x.x.15	5
127.x.x.12	7
127.x.x.17	5

Legend

- **SXP Flow**
- **TrustSec SW Only** (SGT/IP Relay only)
- **TrustSec HW Capable** (SGT Tagging, Fwd and Filtering)

Security Group Tag Exchange Protocol (SXP)
Migration for Access-Layer

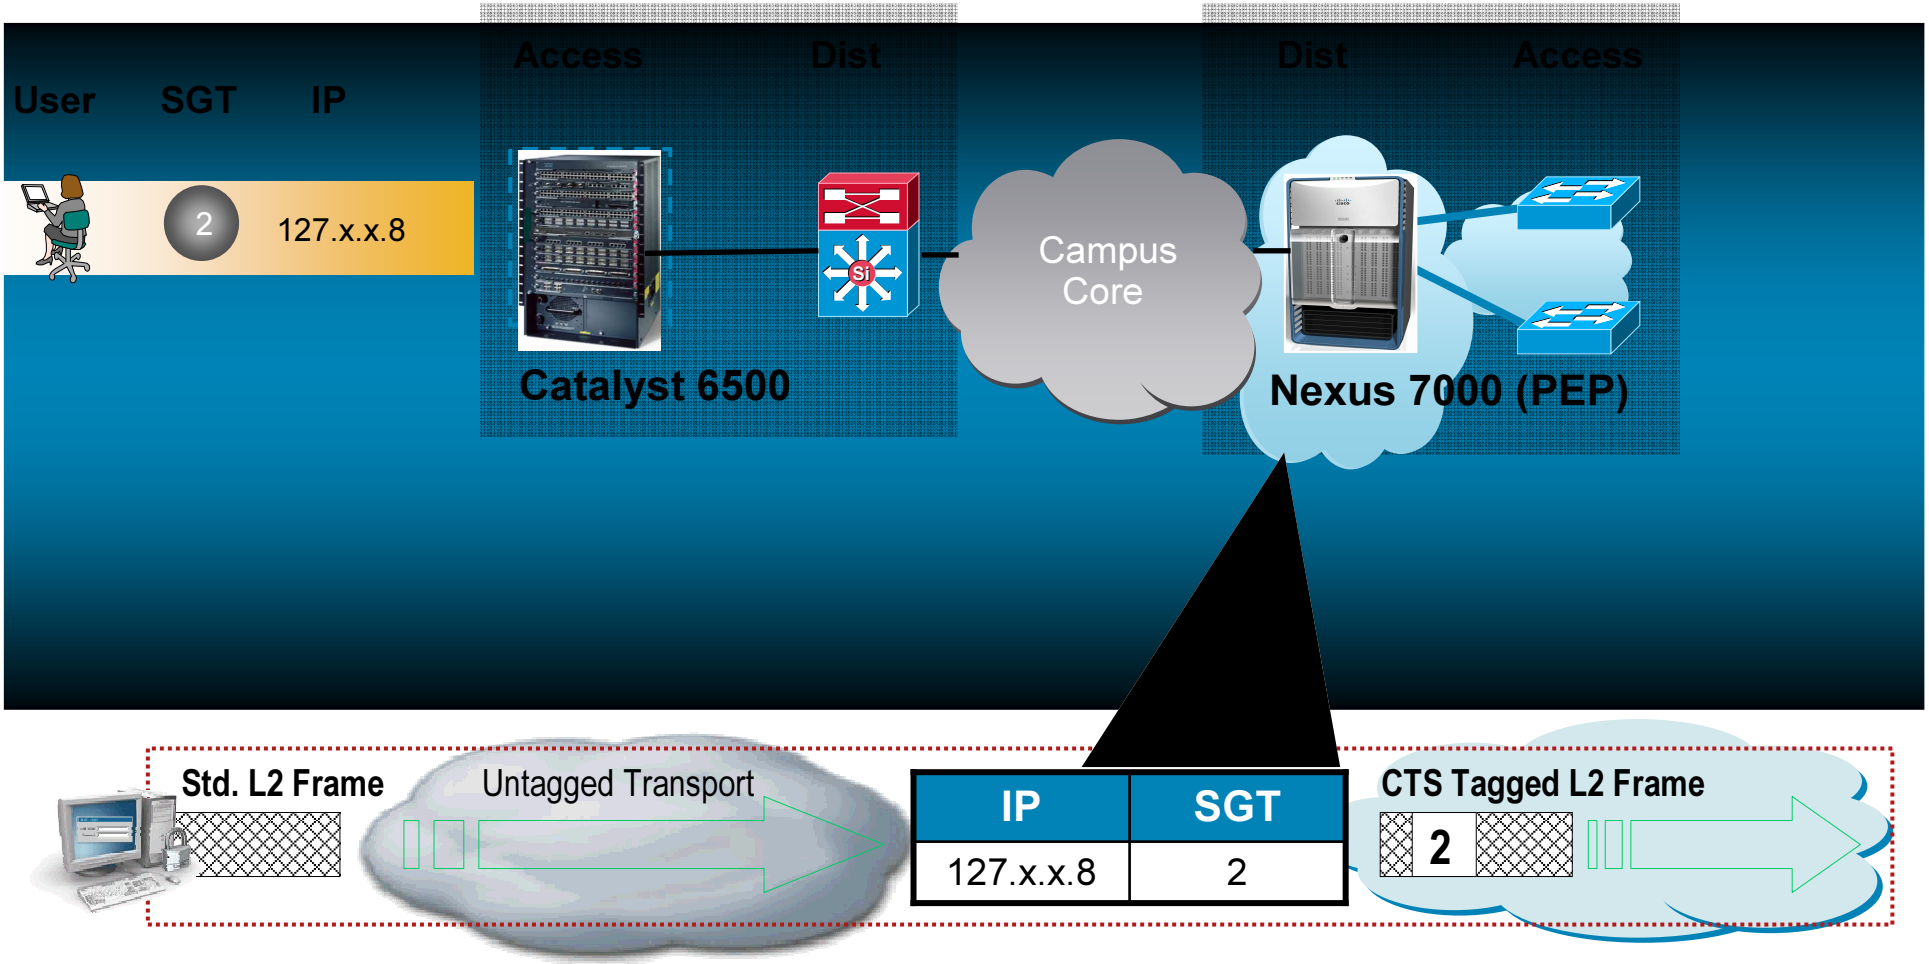

Legend

CTS SW Only (SGT/IP Relay only)

CTS HW Capable (SGT Tagging, Fwd & Filtering)

User SGT IP

Partner 2 **127.x.x.8**

© 2012 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Cisco TrustSec Summary

- Identity = Who are you?
- NAC = Are you healthy?
- CTS = What can you do on this trusted and secure network ?
- Next generation architectural framework for identity, policy and security in DC and Campus
- Innovation & Technology Leadership

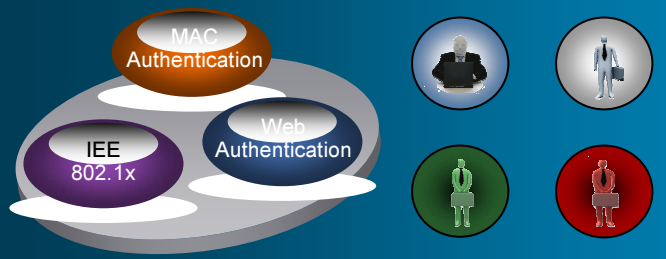

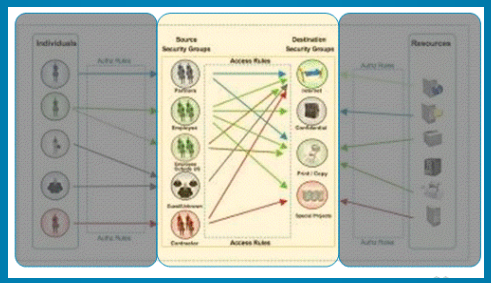

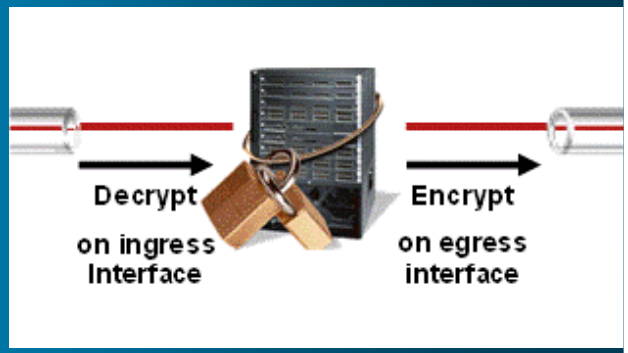

Summary

"Nothing is as dangerous in architecture as dealing with separated problems. If we split life into separated problems we split the possibilities to make good building art"

- Alvar Aalto 1898-1976

  Finnish Modernist Architect/Designer

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Summary

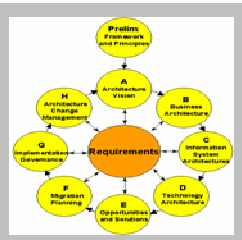

TOGAF

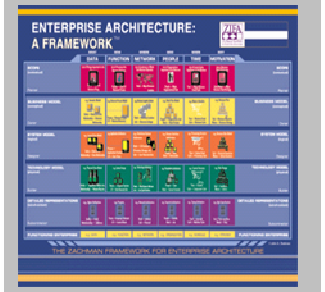

Zachman Framework

IAF - CapGemini

Applications

Services

Network

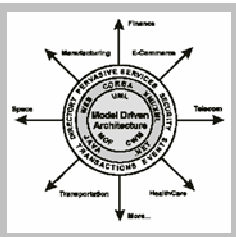

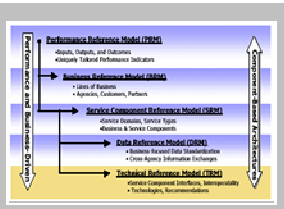

Federal Enterprise Architecture

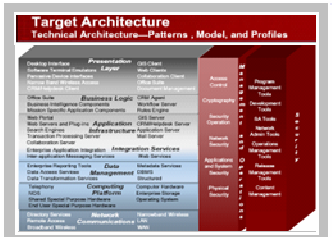

Homeland Security

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Summary

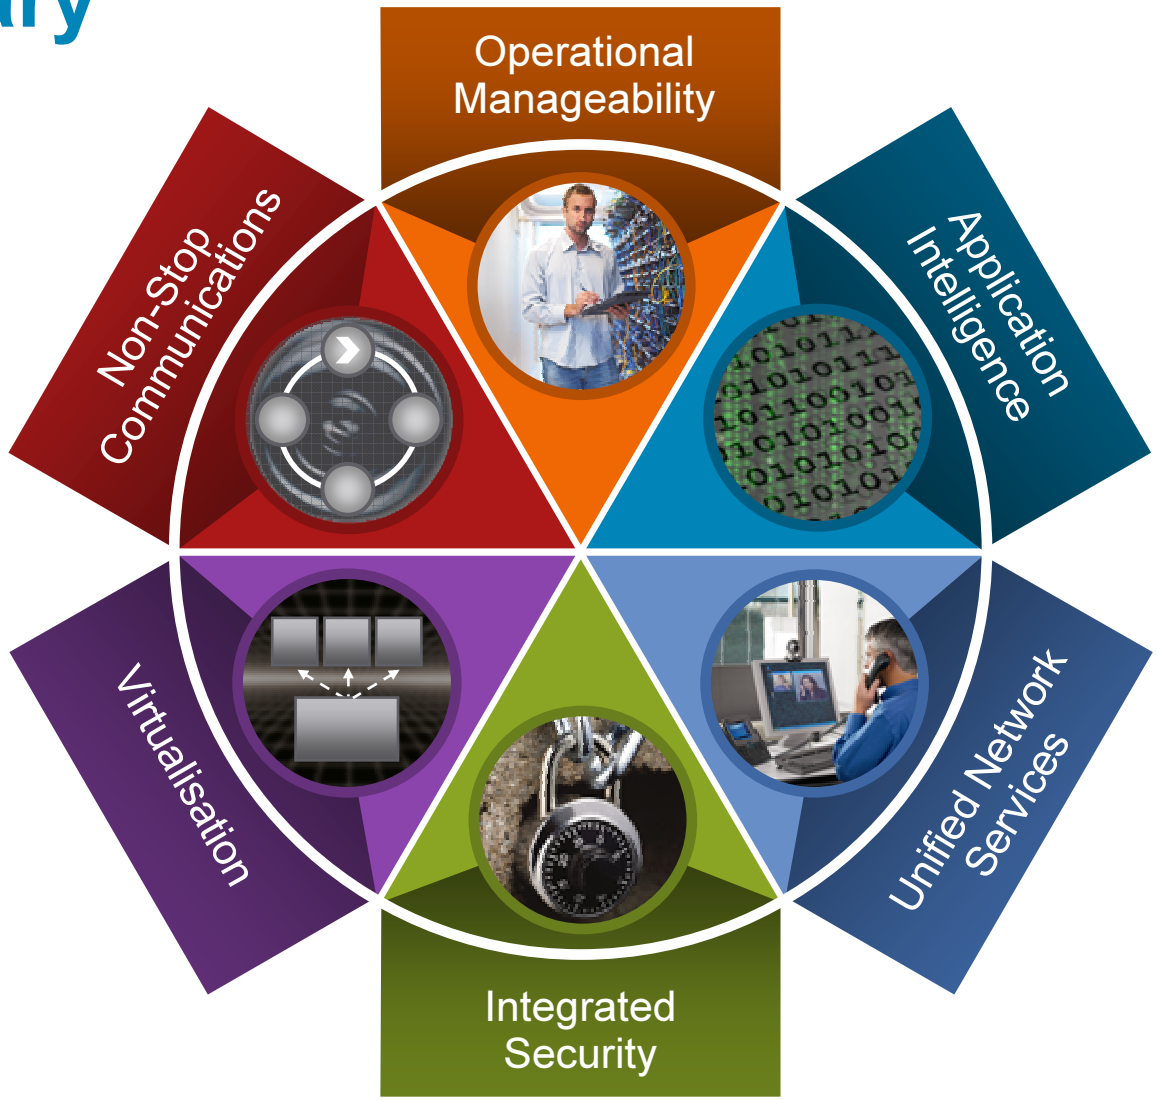

Campus Communication Fabric: http://www.cisco.com/go/ccf
Cisco TrustSec: http://www.cisco.com/go/switchsecurity

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

Thank You !

Please Complete Your Evaluation Form

© 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential

 © 2006 Cisco Systems, Inc. All rights reserved. Cisco Confidential