Application Networking Services

Rene Raeber, CSE

Central Advanced Technology Consulting

Agenda

- Introduction
- Application Protocol Differentiation
- WAN Optimization
- Datacenter SLB & L4-7 SEC
- Summary

Introduction

Today's Business Processes Are Complex

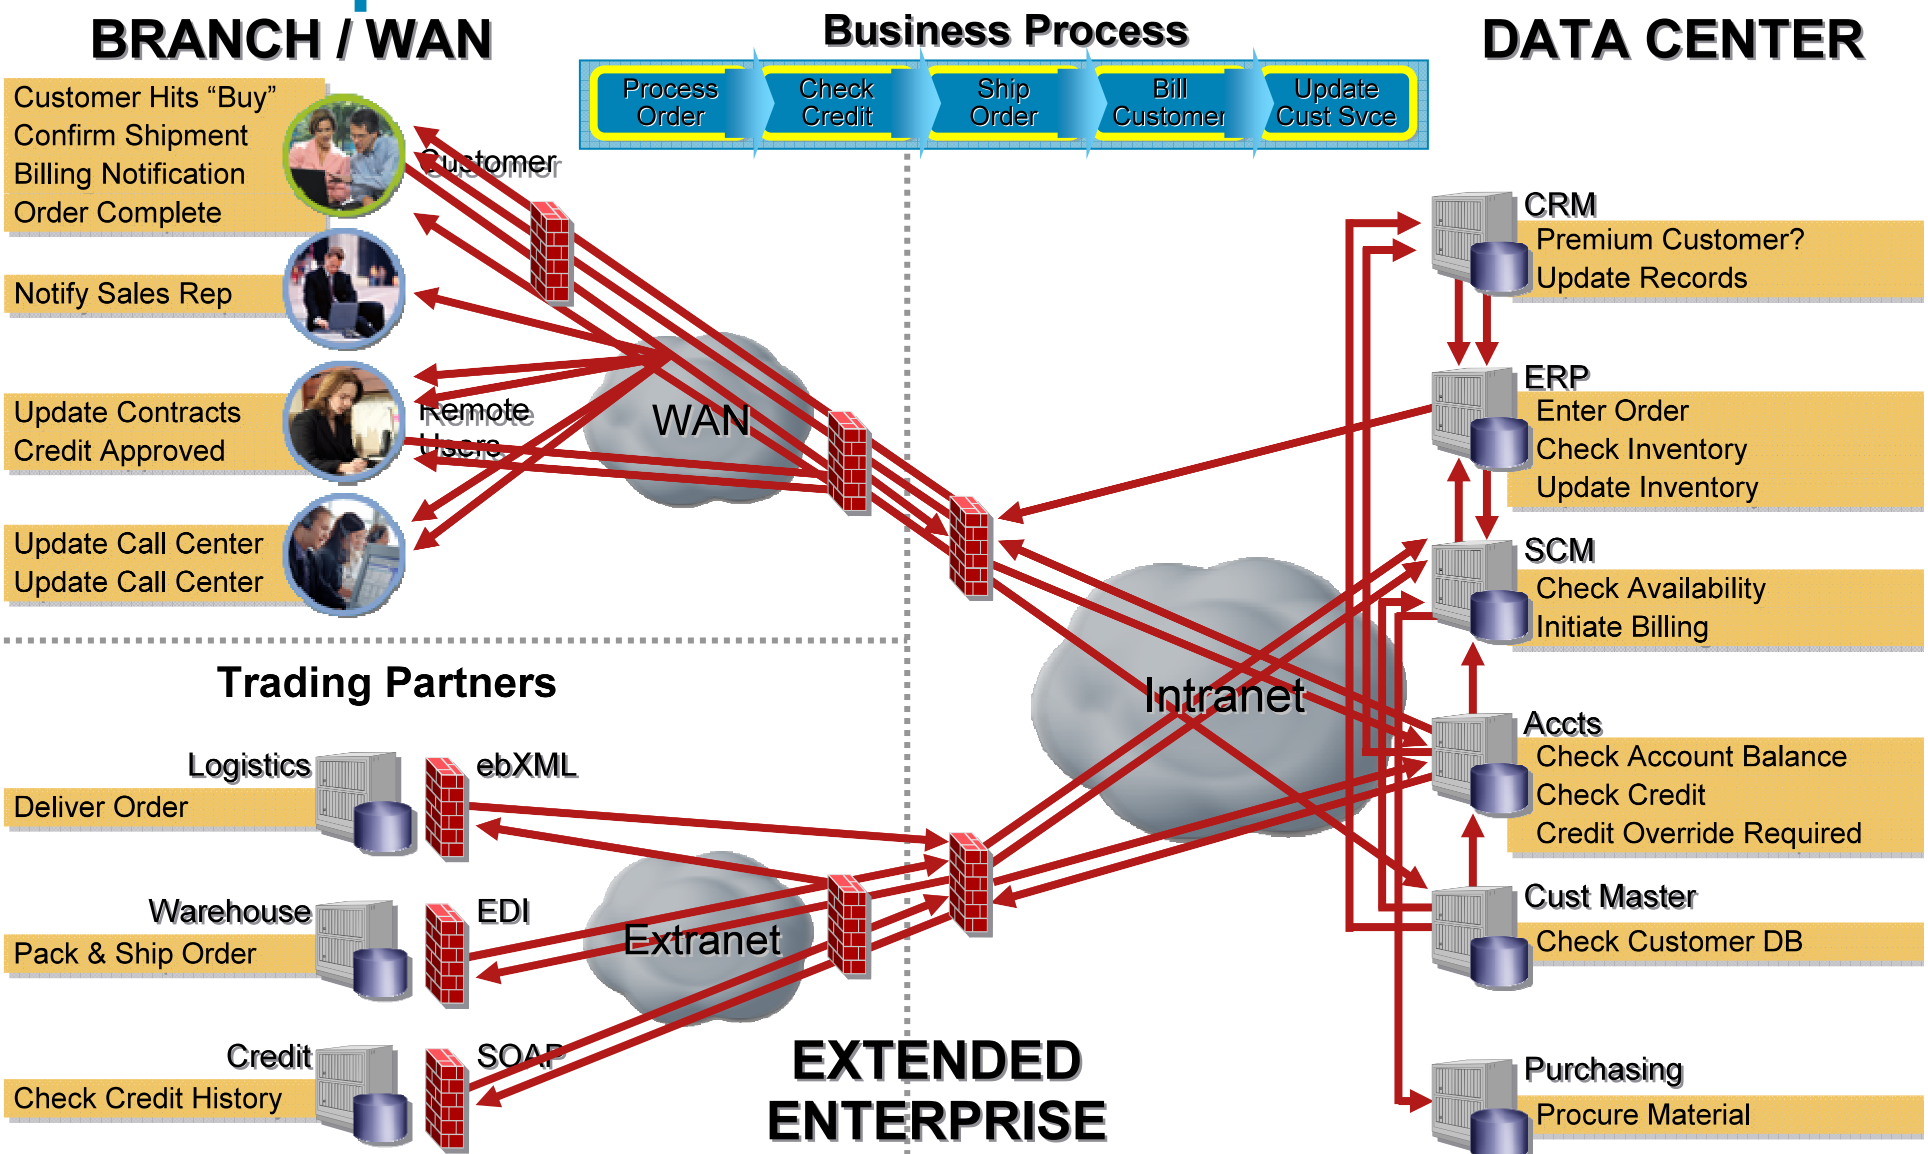

IT Megatrends: The Walls Are Coming Down

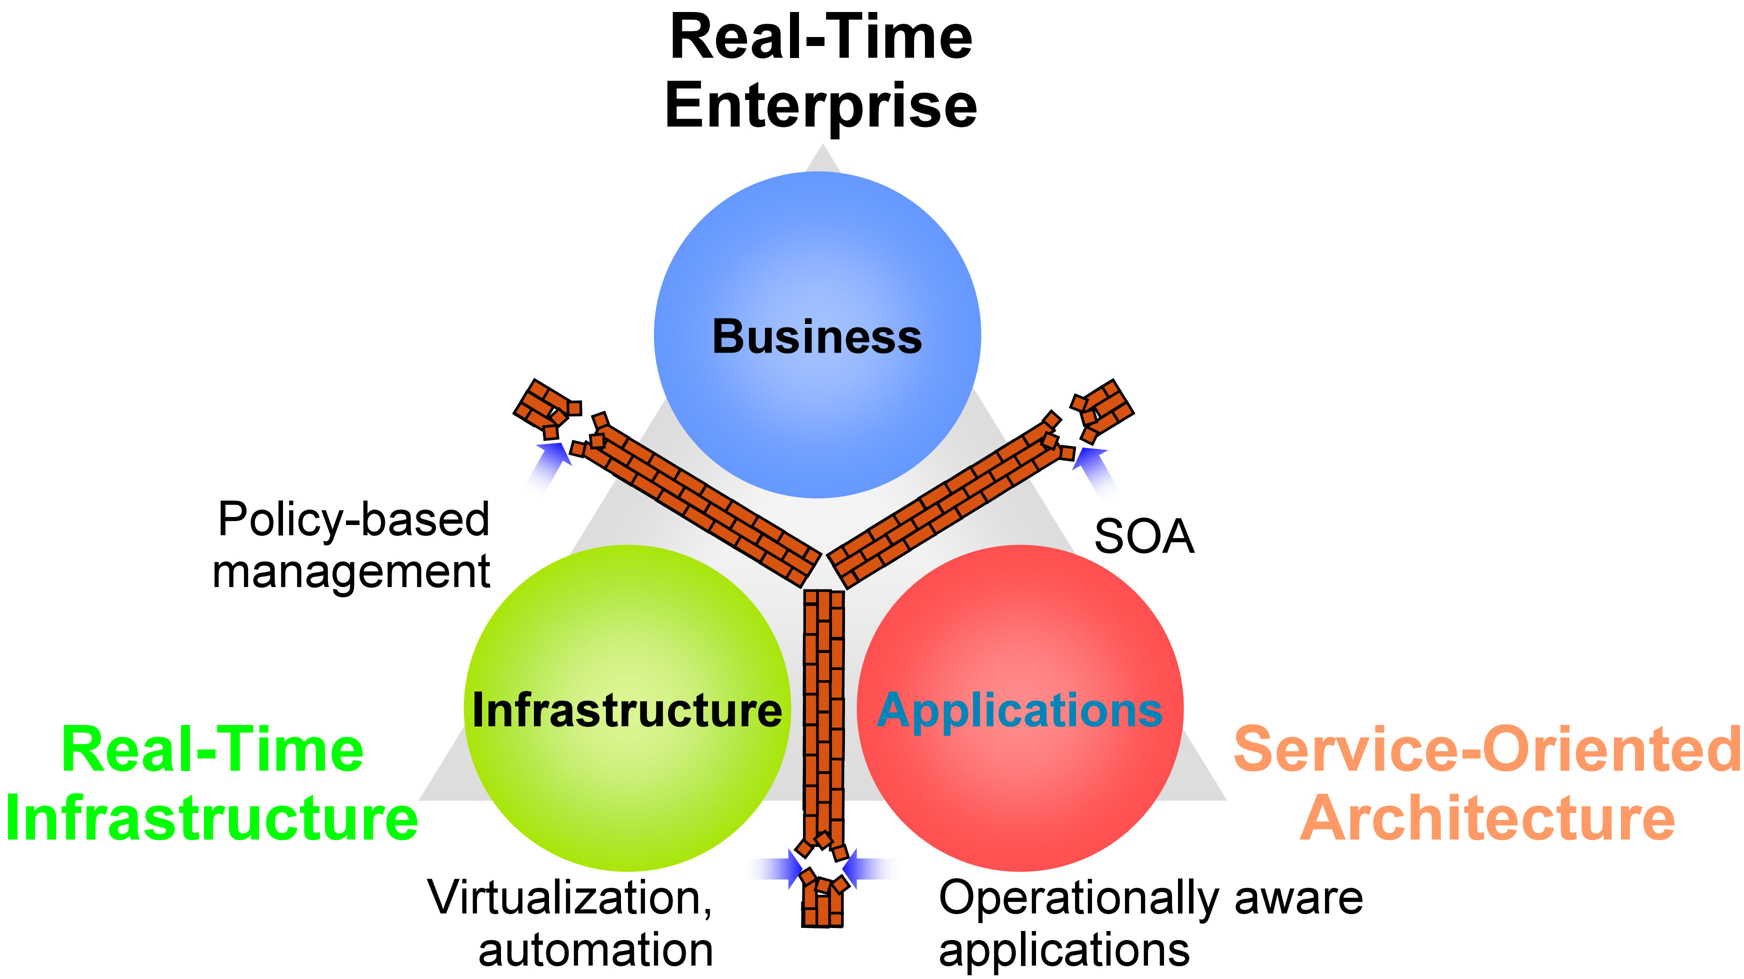

Compute Evolution

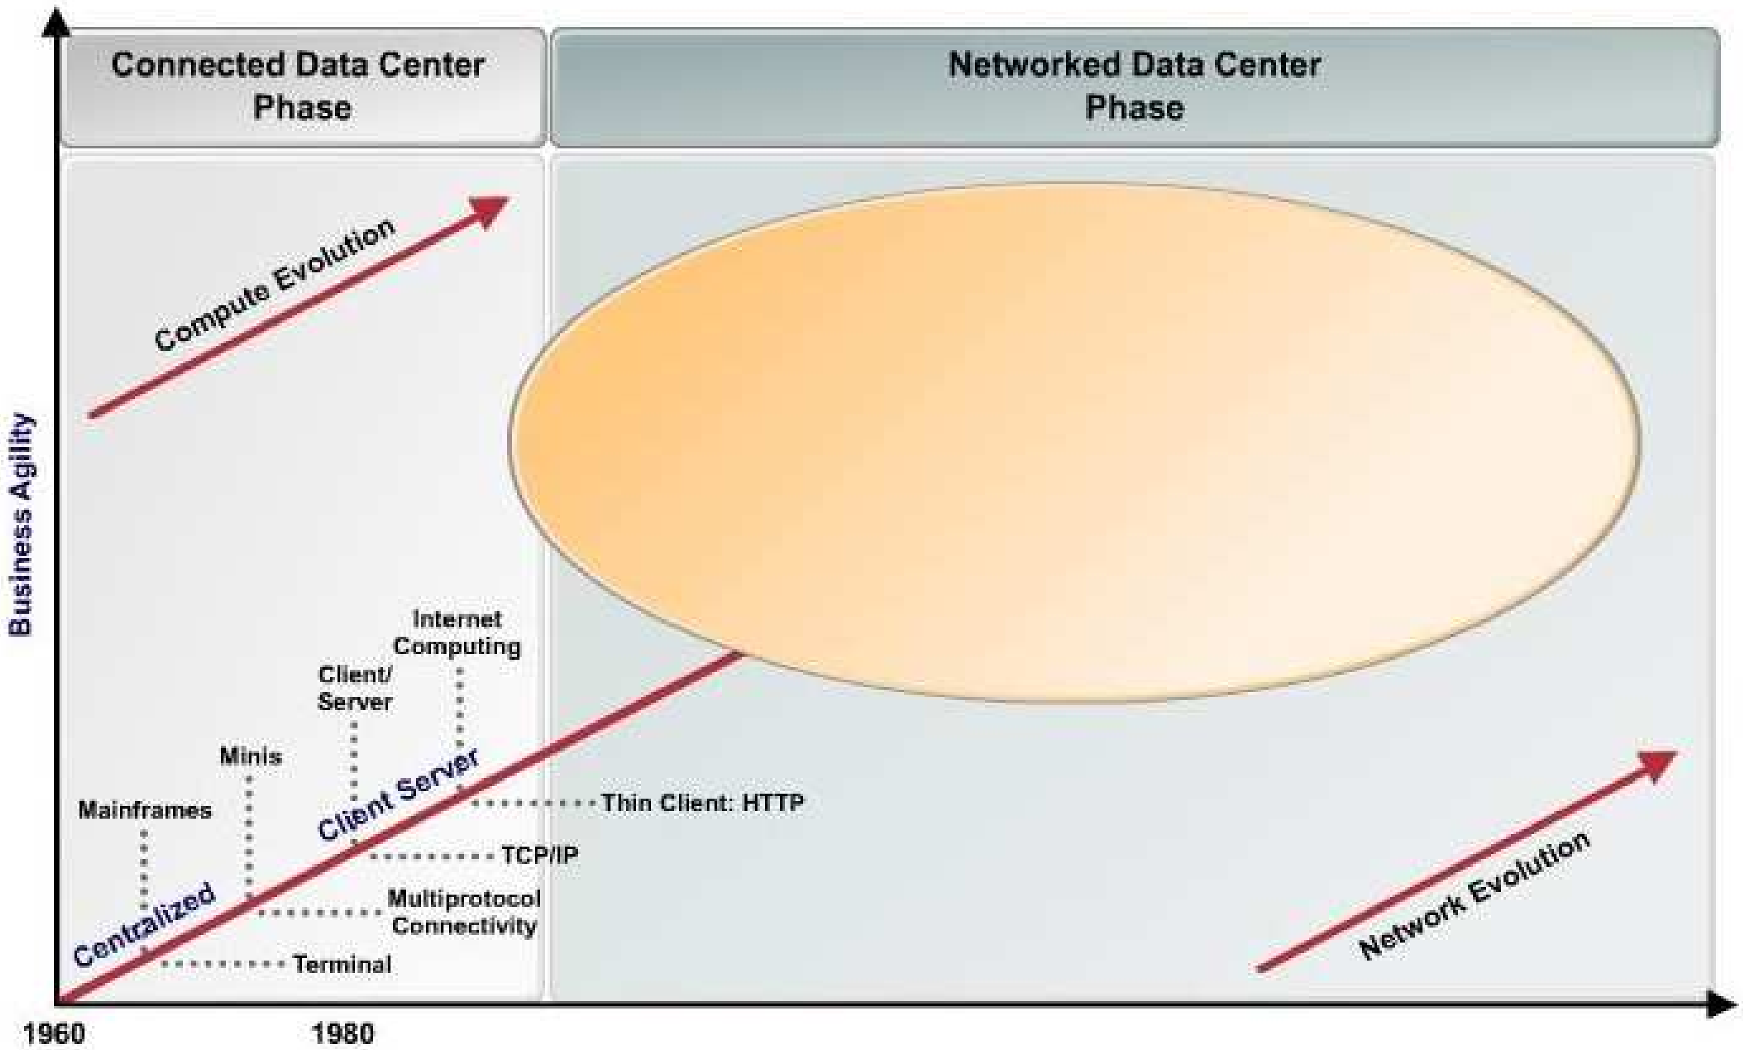

The Network Phase
The Intelligent Information Network

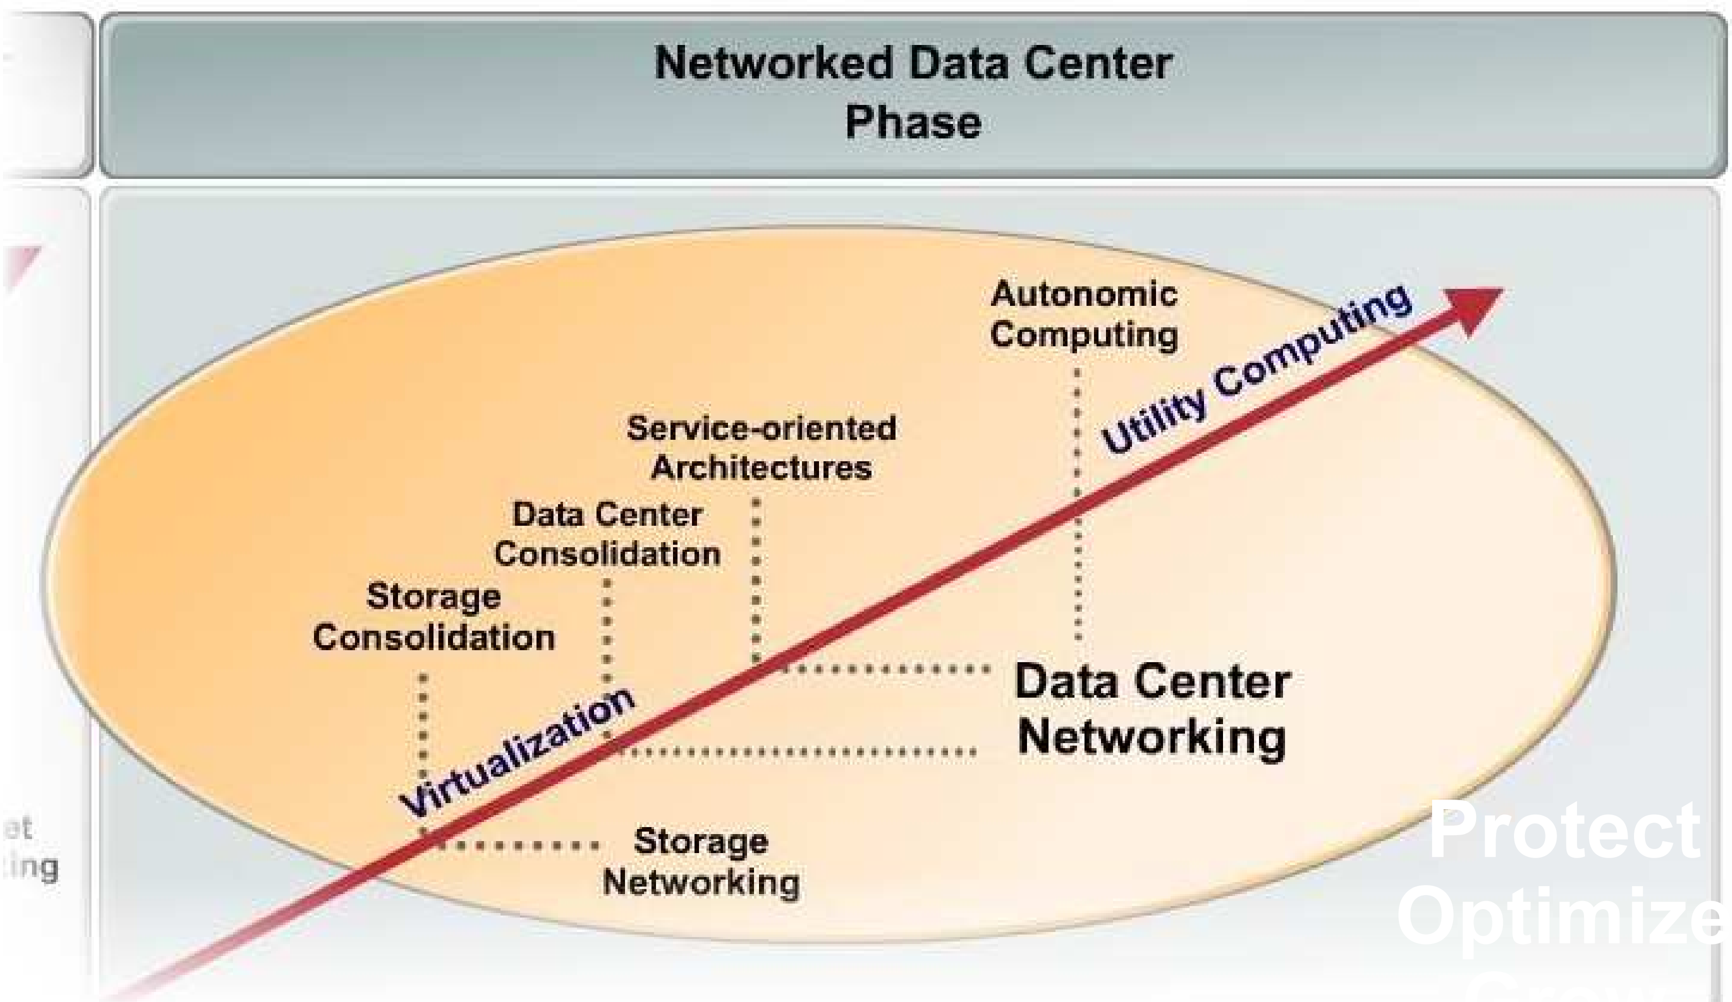

Data Center Architectural Evolution

Consolidated Multilayer Data Center → **Modular and Virtual Data Center** → **Automated Data Center**

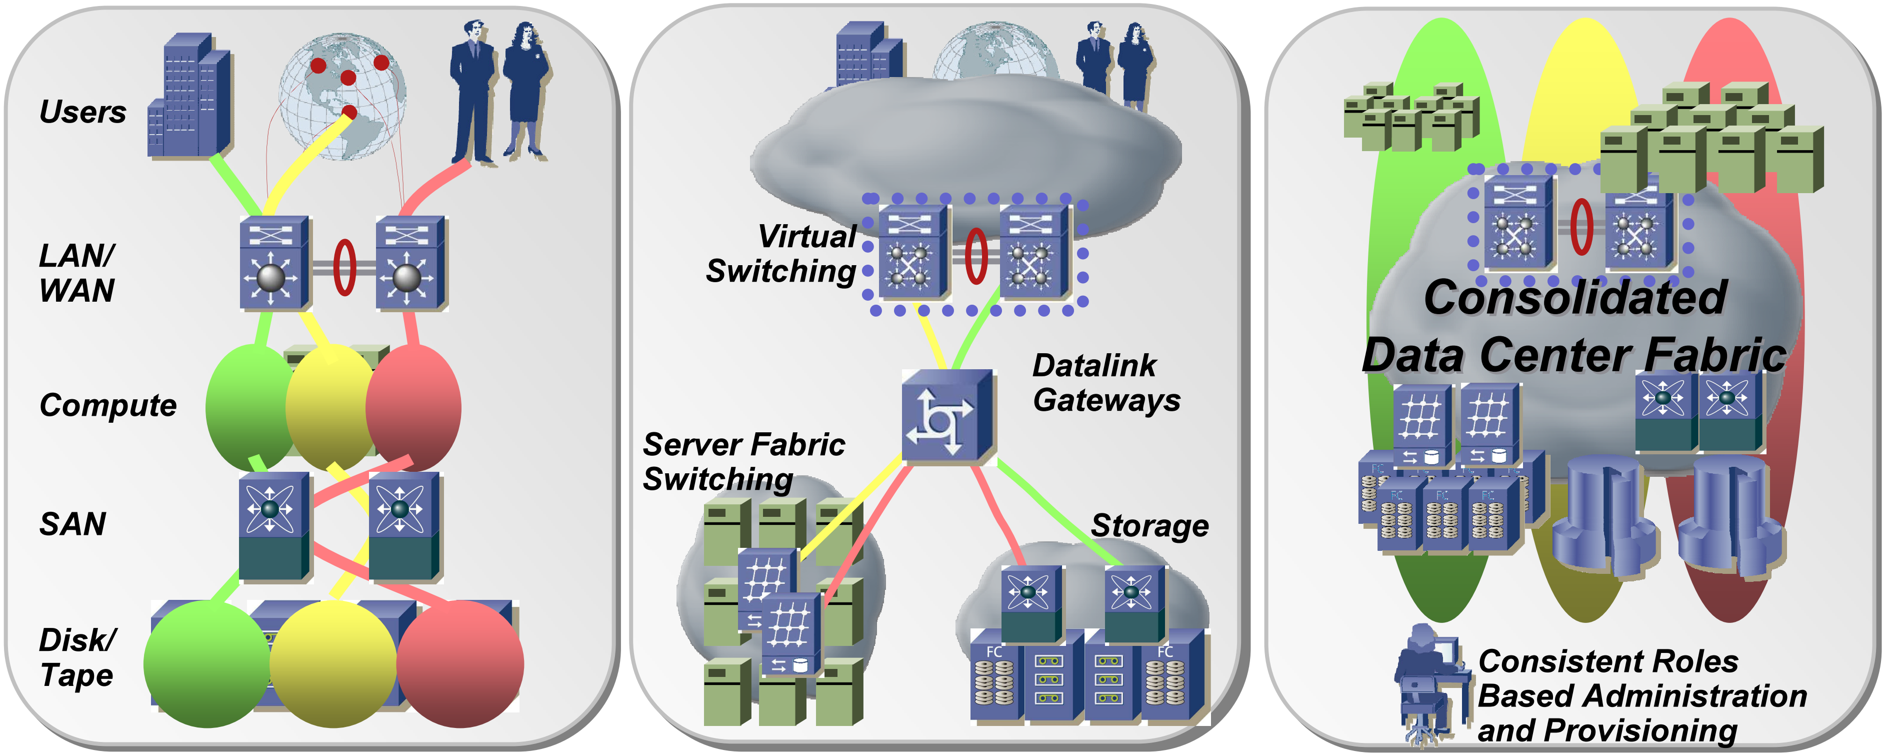

Today: Service Modules and VLANs/Routing

Next: Synchronous Virtual DC, Server Fabric Switching

Future: Transport, Service, Security, Server Provisioning

Data Center Architecture

Functional Layers and Services

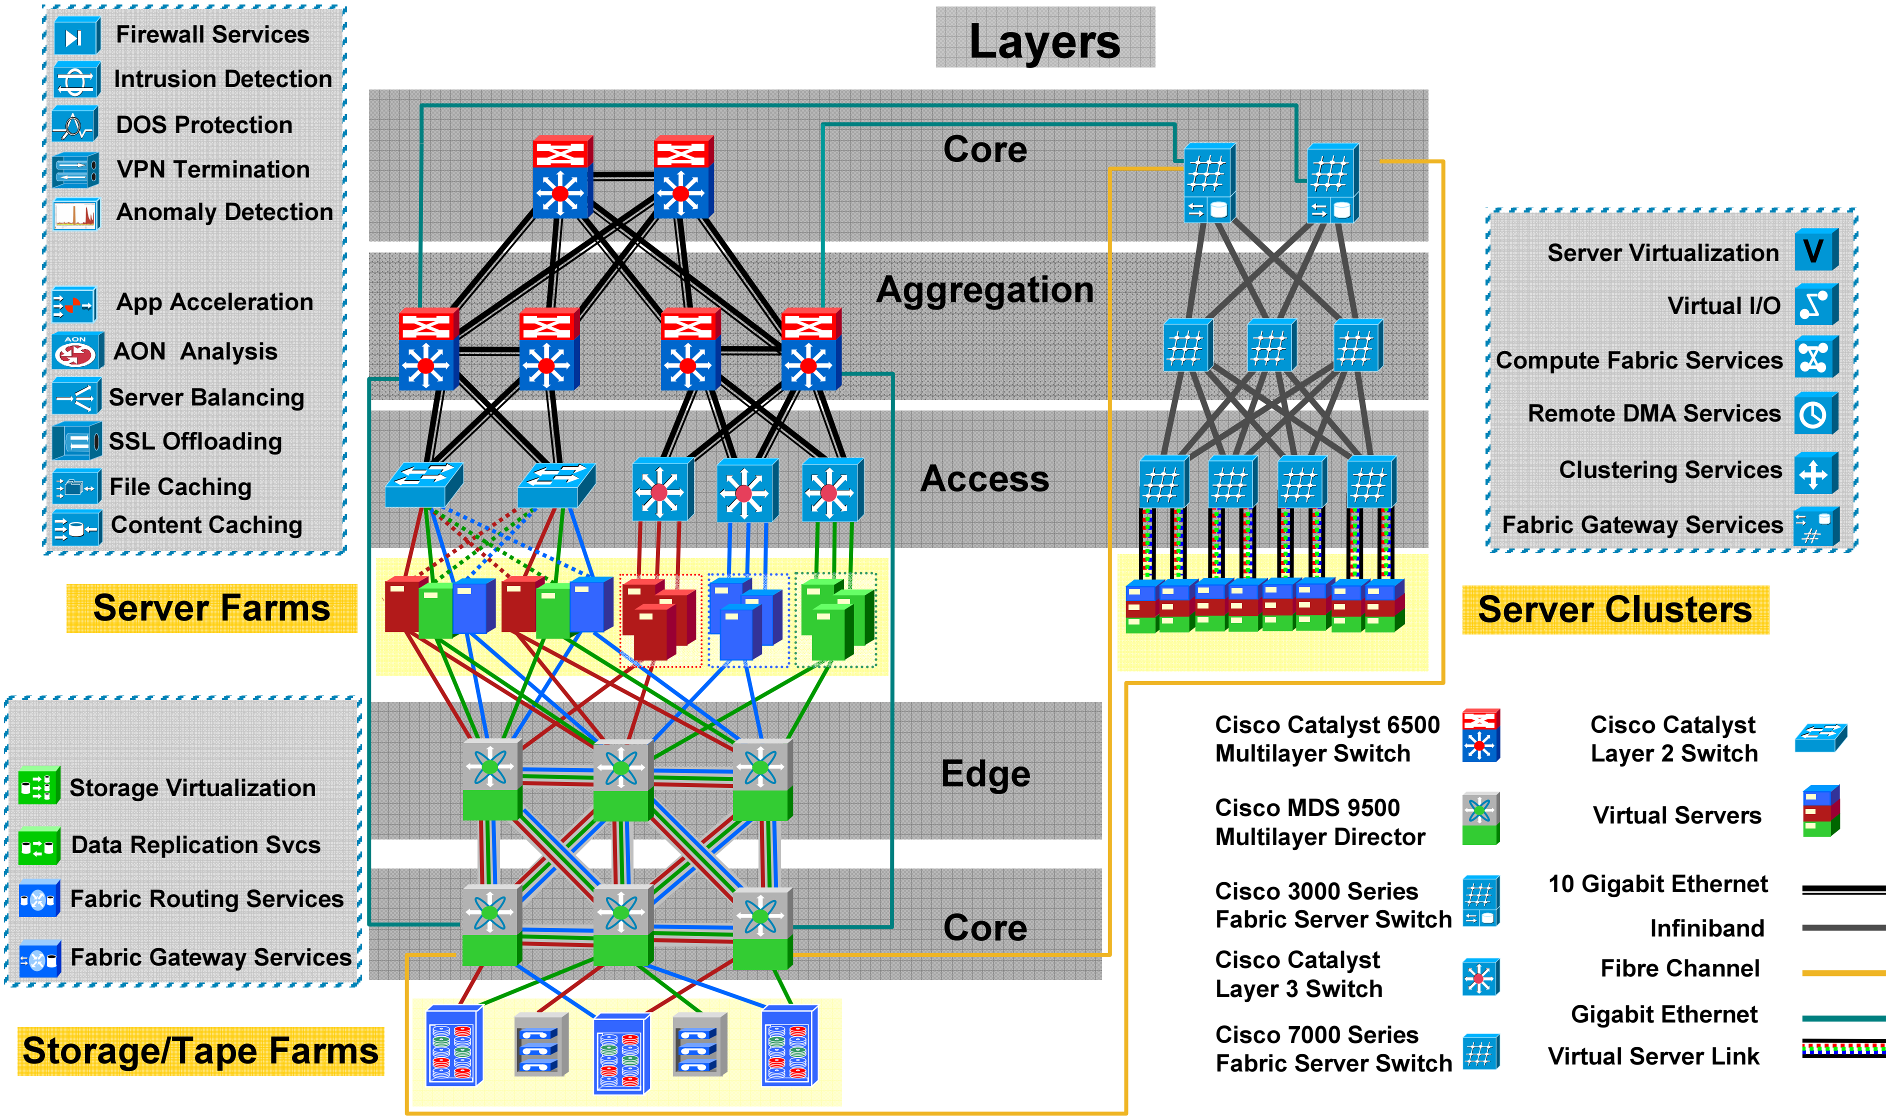

Application Architectures and Protocols

Application Optimization Infrastructure

Network Classification
- Quality of service
- Network-based app recognition
- Queuing, policing, shaping
- Visibility, monitoring, control

Application Scalability
- Server load-balancing
- Site selection
- SSL termination and offload
- Video delivery

Application Networking
- Message transformation
- Protocol transformation
- Message-based security
- Application visibility

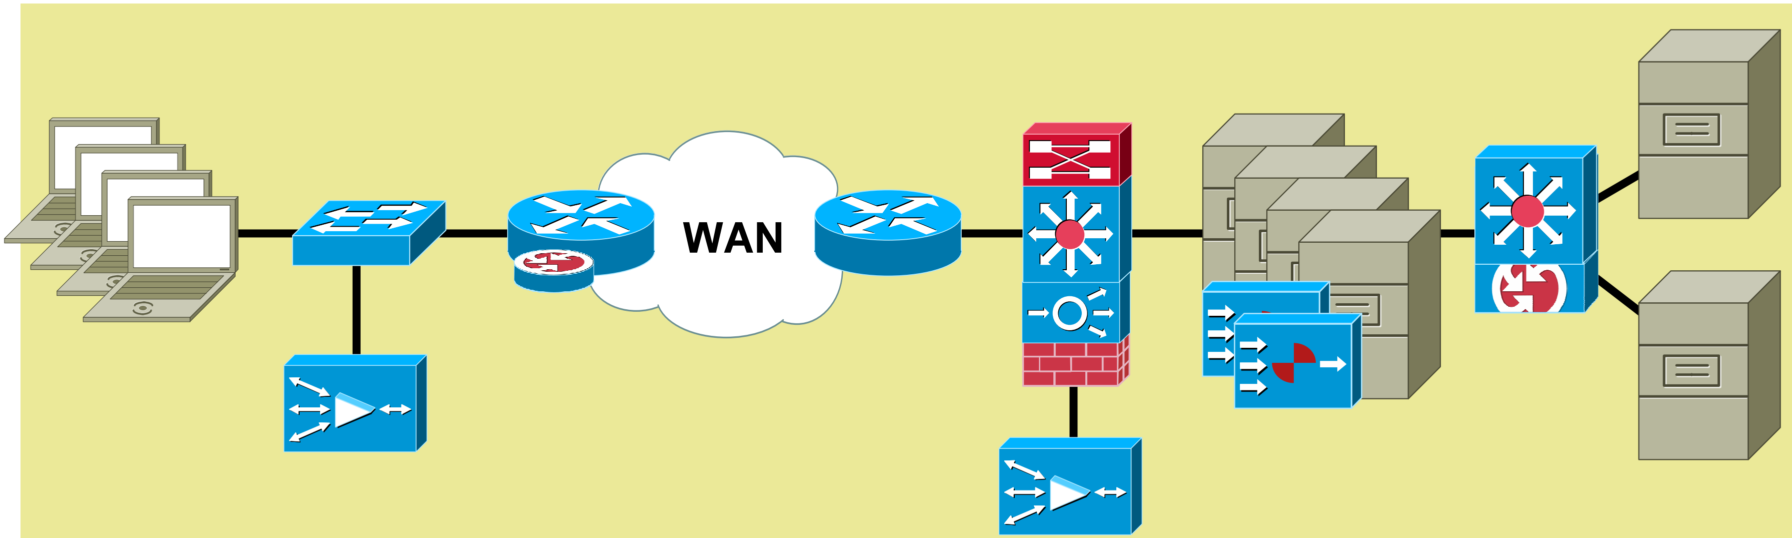

Application Acceleration
- Latency mitigation
- Application data cache
- Meta data cache
- Local services

WAN Acceleration
- Data redundancy elimination
- Window scaling
- LZ compression
- Adaptive congestion avoidance

Application Optimization
- Delta encoding
- FlashForward optimization
- Application security
- Server offload

Different Types of Traffic Have Different Needs

- Real-time applications especially sensitive
  - Interactive voice
  - Videoconferencing
- Causes of degraded performance
  - Congestion
    - Convergence
    - Peak traffic load
  - Link speed & capacity differences

➢ Set application service level objectives

Application Examples	Sensitivity		
	Delay	Jitter	Packet Loss
Interactive Voice and Video	Y	Y	Y
Streaming Video	N	Y	Y
Transactional/ Interactive	Y	N	N
Bulk Data Email File Transfer	N	N	N

Traffic Differentiation Overview

Traffic Differentiation Is a Network-Integrated Function That Ensures Network Administrators Can:

- Identify how the network is being used and by who
- Identify (classify) flows and applications on the network
- Apply network prioritization and specialized handling
- Verify performance, report, and adapt accordingly

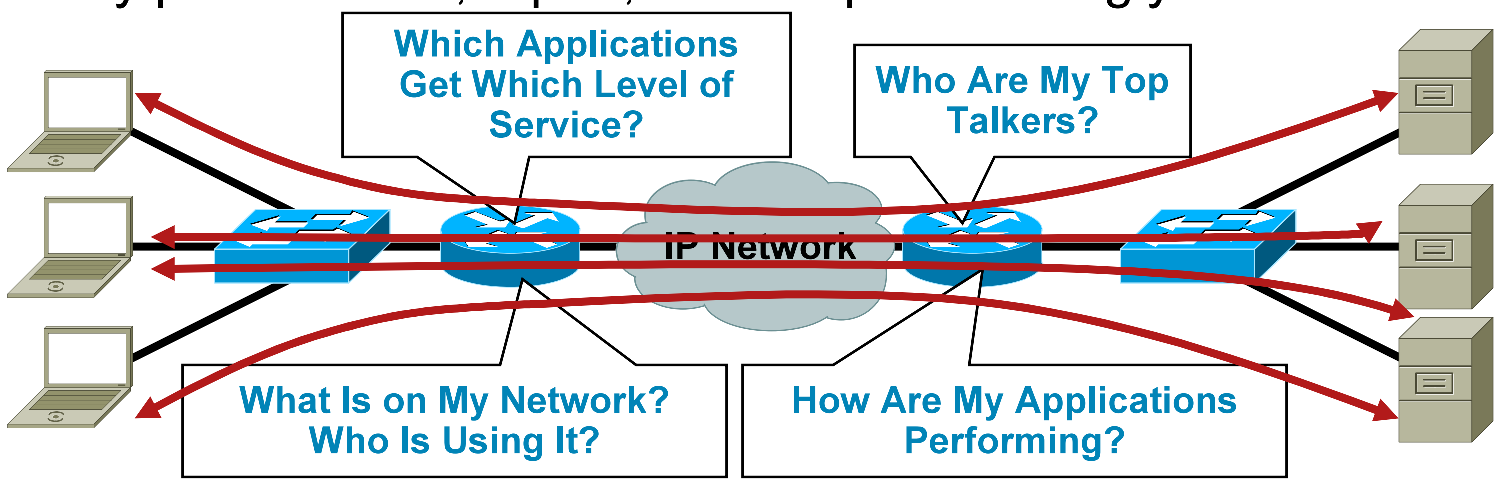

Traffic Differentiation Overview

The Primary Functional Components of Network-Integrated Traffic Differentiation Include:

- **NetFlow:** understand how the network is being used and who is using the network
- **Quality of service:** provide differentiated services based on business priority
- **IP Service Level Agreements (IP SLAs):** ensure network is performing as expected

NetFlow Analyzes IP Flows

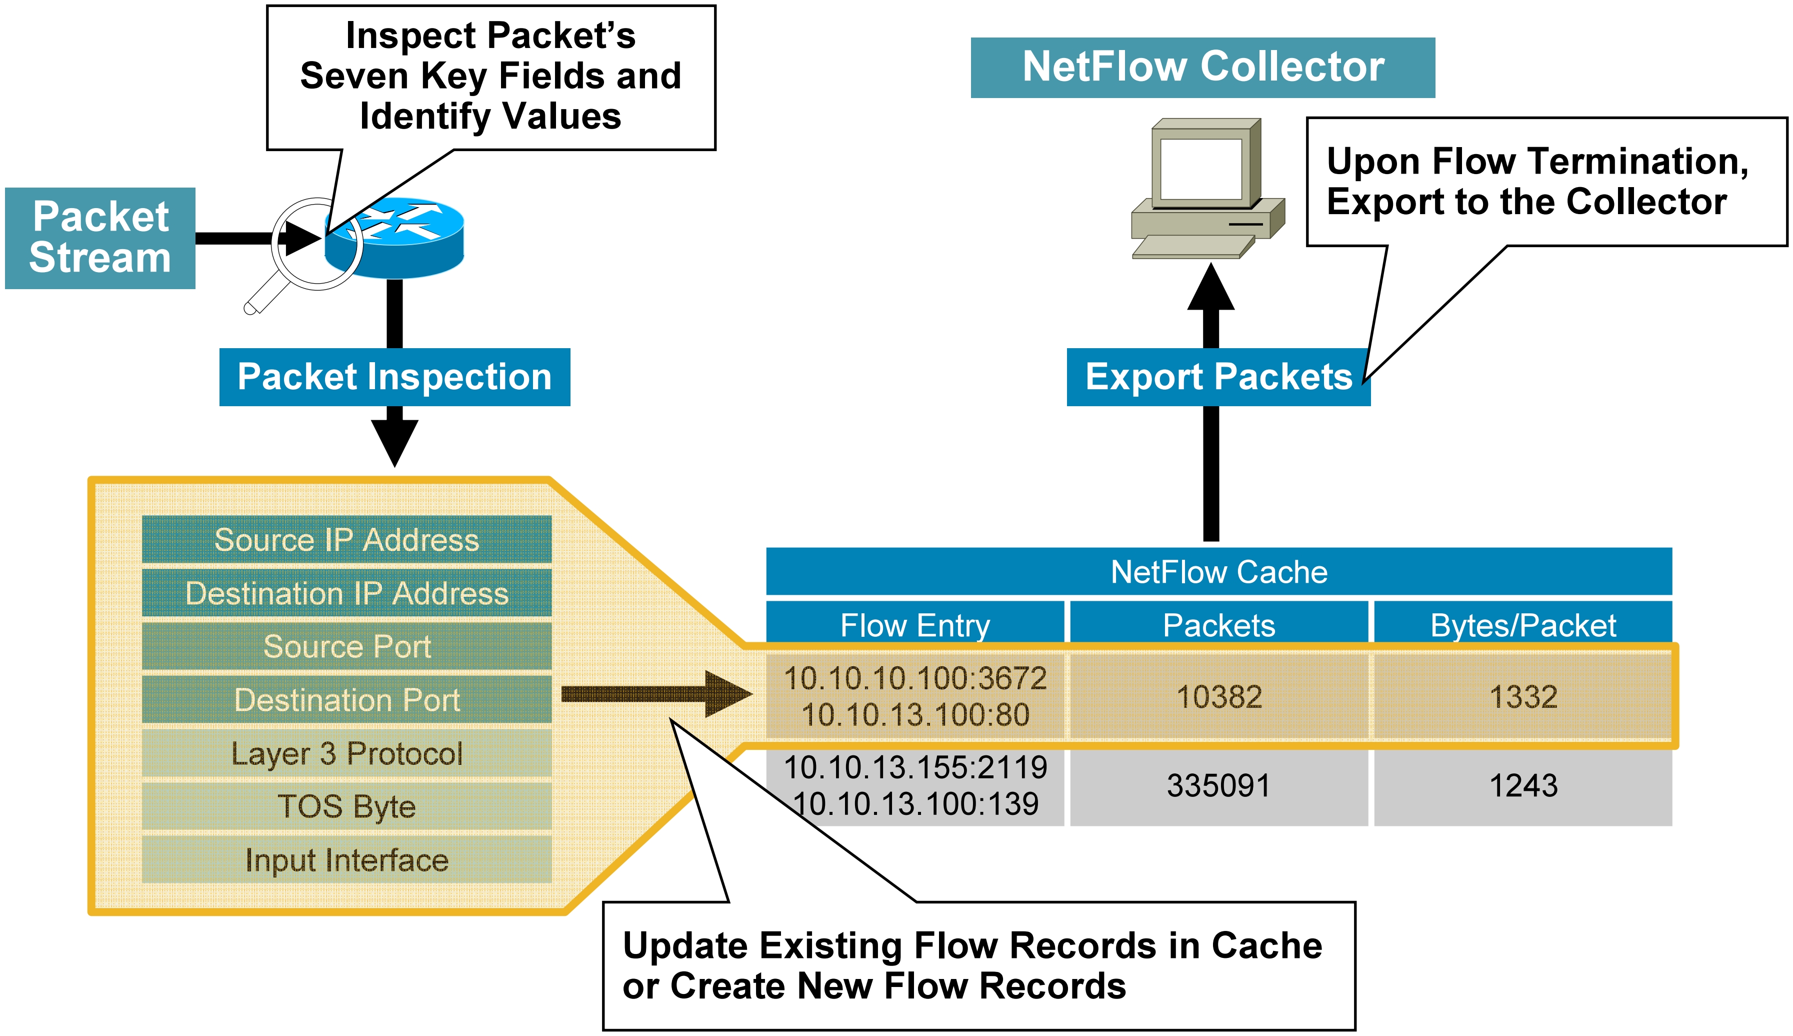

What Information Does NetFlow Provide?

1. Create and Update Flows in NetFlow Cache

Key Fields in Blue; Non-Key Fields Green

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	Pkts	Src Port	Src Msk	Src AS	Dst Port	Dst Msk	Dst AS	NextHop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	00A2	/24	5	00A2	/24	15	10.0.23.2	1528	1745	4
Fa1/0	173.100.3.2	Fa0/0	10.0.227.12	6	40	0	2491	15	/26	196	15	/24	15	10.0.23.2	740	41.5	1
Fa1/0	173.100.20.2	Fa0/0	10.0.227.12	11	80	10	10000	00A1	/24	180	00A1	/24	15	10.0.23.2	1428	1145.5	3
Fa1/0	173.100.6.2	Fa0/0	10.0.227.12	6	40	0	2210	19	/30	180	19	/24	15	10.0.23.2	1040	24.5	14

2. Expiration

- Inactive Timer Expired (15 Sec Is Default)
- Active Timer Expired (30 Min (1800 Sec) Is Default)
- Netflow Cache Is Full (Oldest Flows Are Expired)

SrcIf	SrcIPadd	DstIf	DstIPadd	Protocol	TOS	Flgs	Pkts	Src Port	Src Msk	Src AS	Dst Port	Dst Msk	Dst AS	NextHop	Bytes/ Pkt	Active	Idle
Fa1/0	173.100.21.2	Fa0/0	10.0.227.12	11	80	10	11000	00A2	/24	5	00A2	/24	15	10.0.23.2	1528	1800	4

4. Export Version

Non-Aggregated Flows—Export Version 5 or 9

5. Transport Protocol

30 Flows per 1500 Byte Export Packet

Export Packet

Header	Payload (Flows)

"Best Effort" Quality of Service

- Without QoS policies, traffic is served with "best effort"
  - No distinction between high and low priority
    - Business critical vs. background
  - No allowances for different application needs
    - Real-time voice/video vs. bulk data transfer
- No problem, until congestion occurs

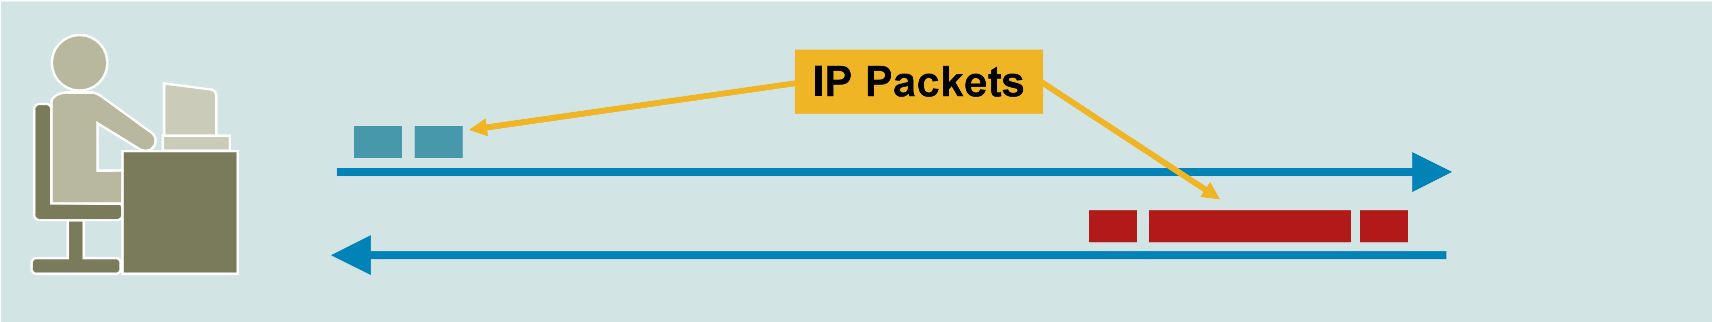

No Congestion—No Problem

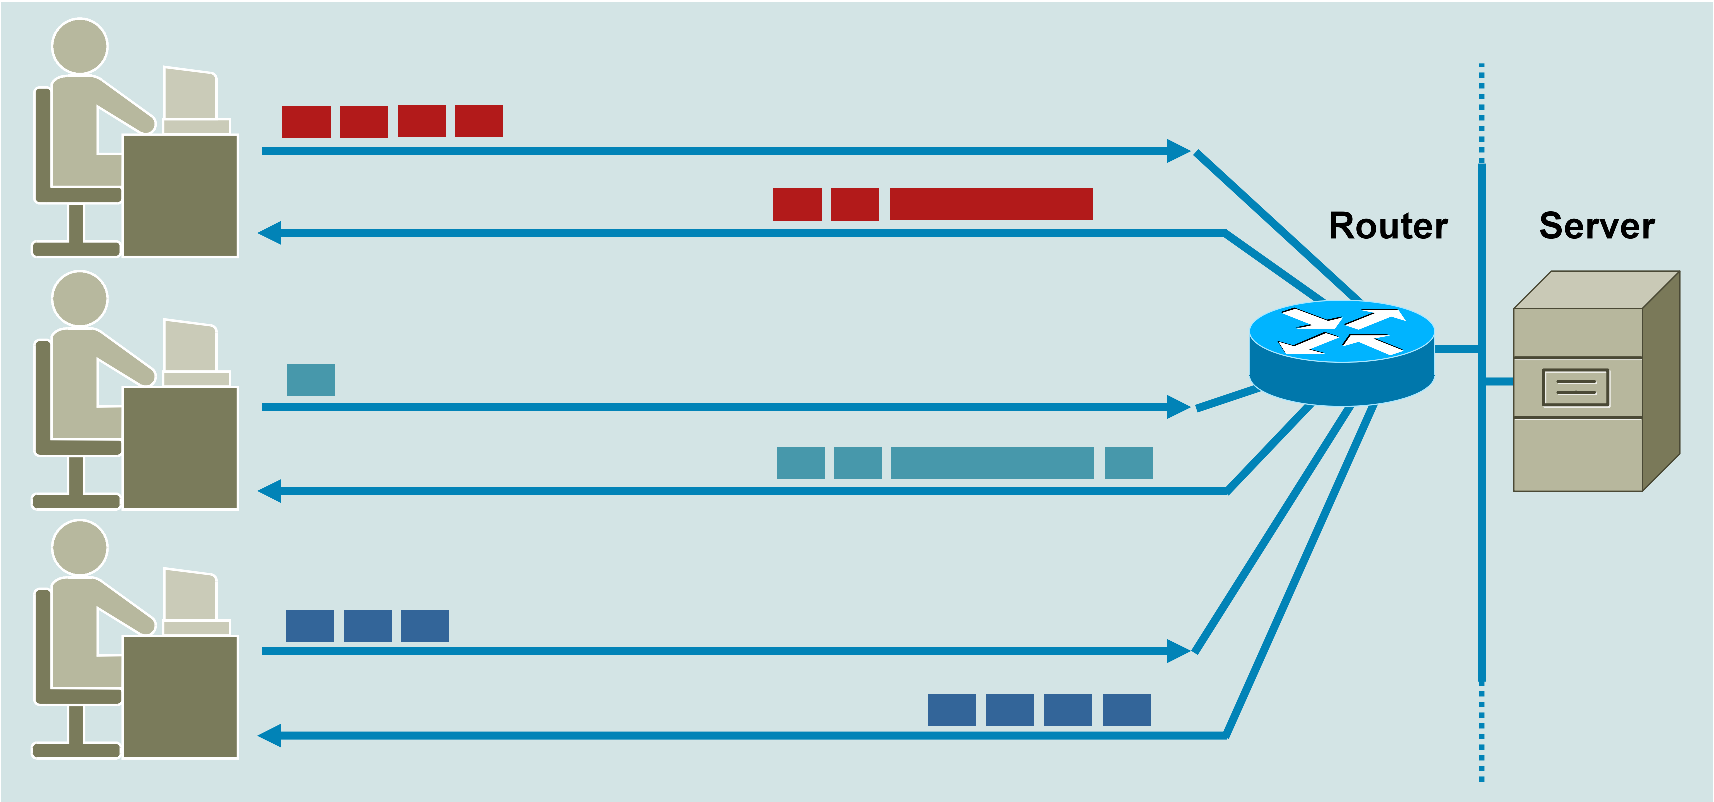

- On serial links, longer packets take longer to transmit
- Smaller packets can be delayed behind longer ones

Congestion Without QoS Policies

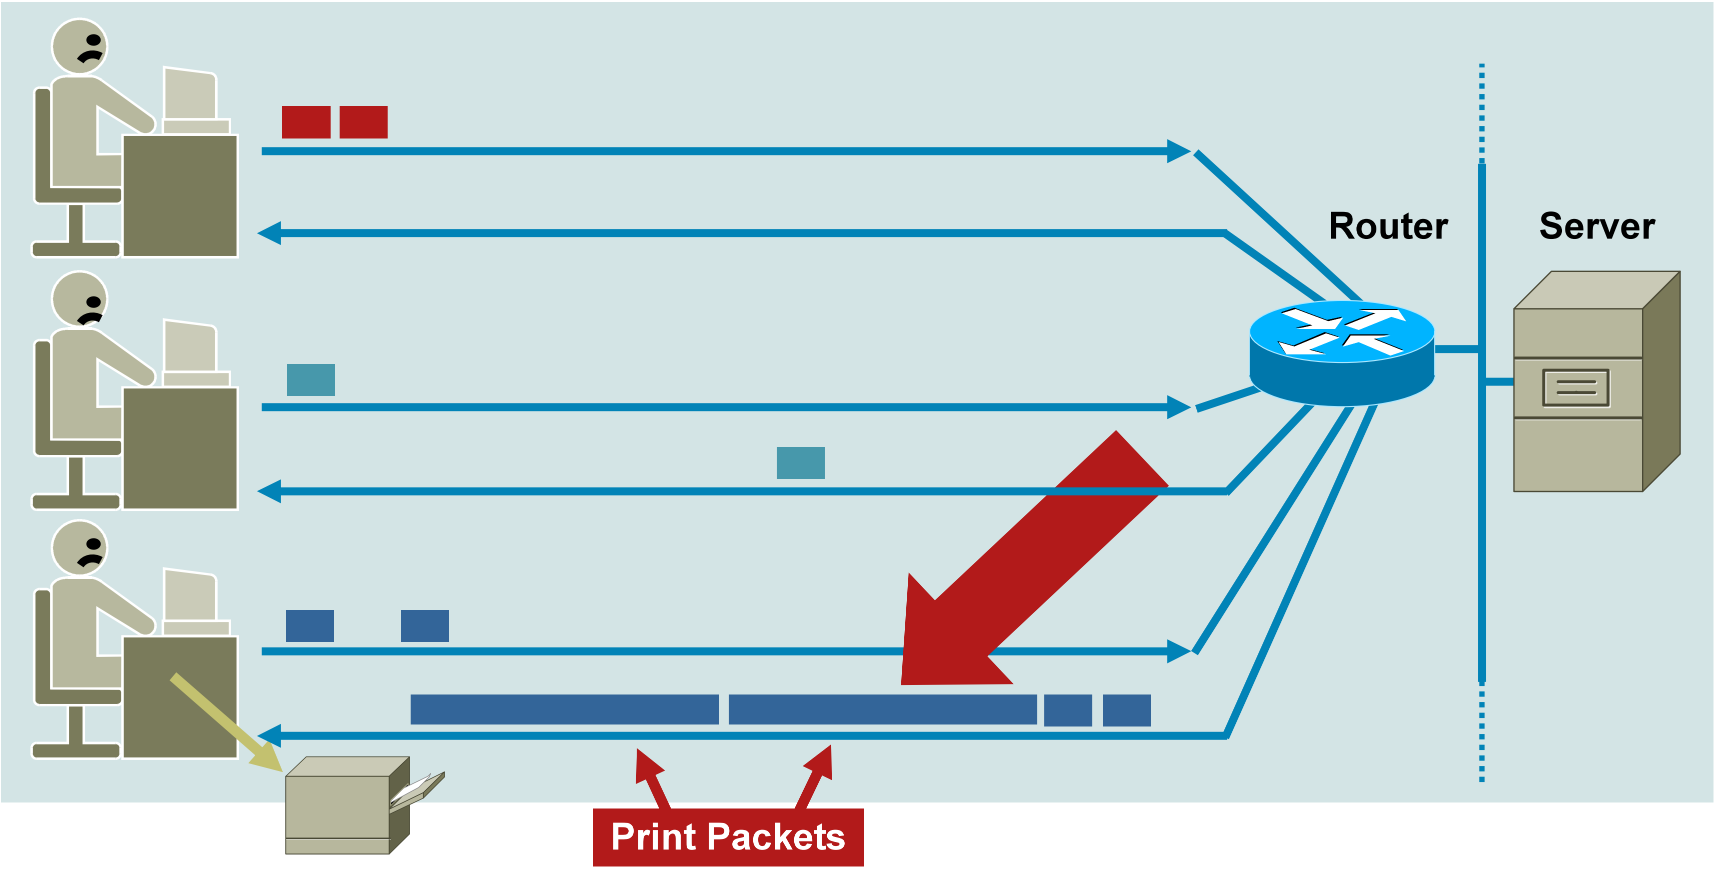

- Print packets dominate available bandwidth
- Response time slows for all users

Congestion with QoS Policies Configured

•Non-printing users not affected

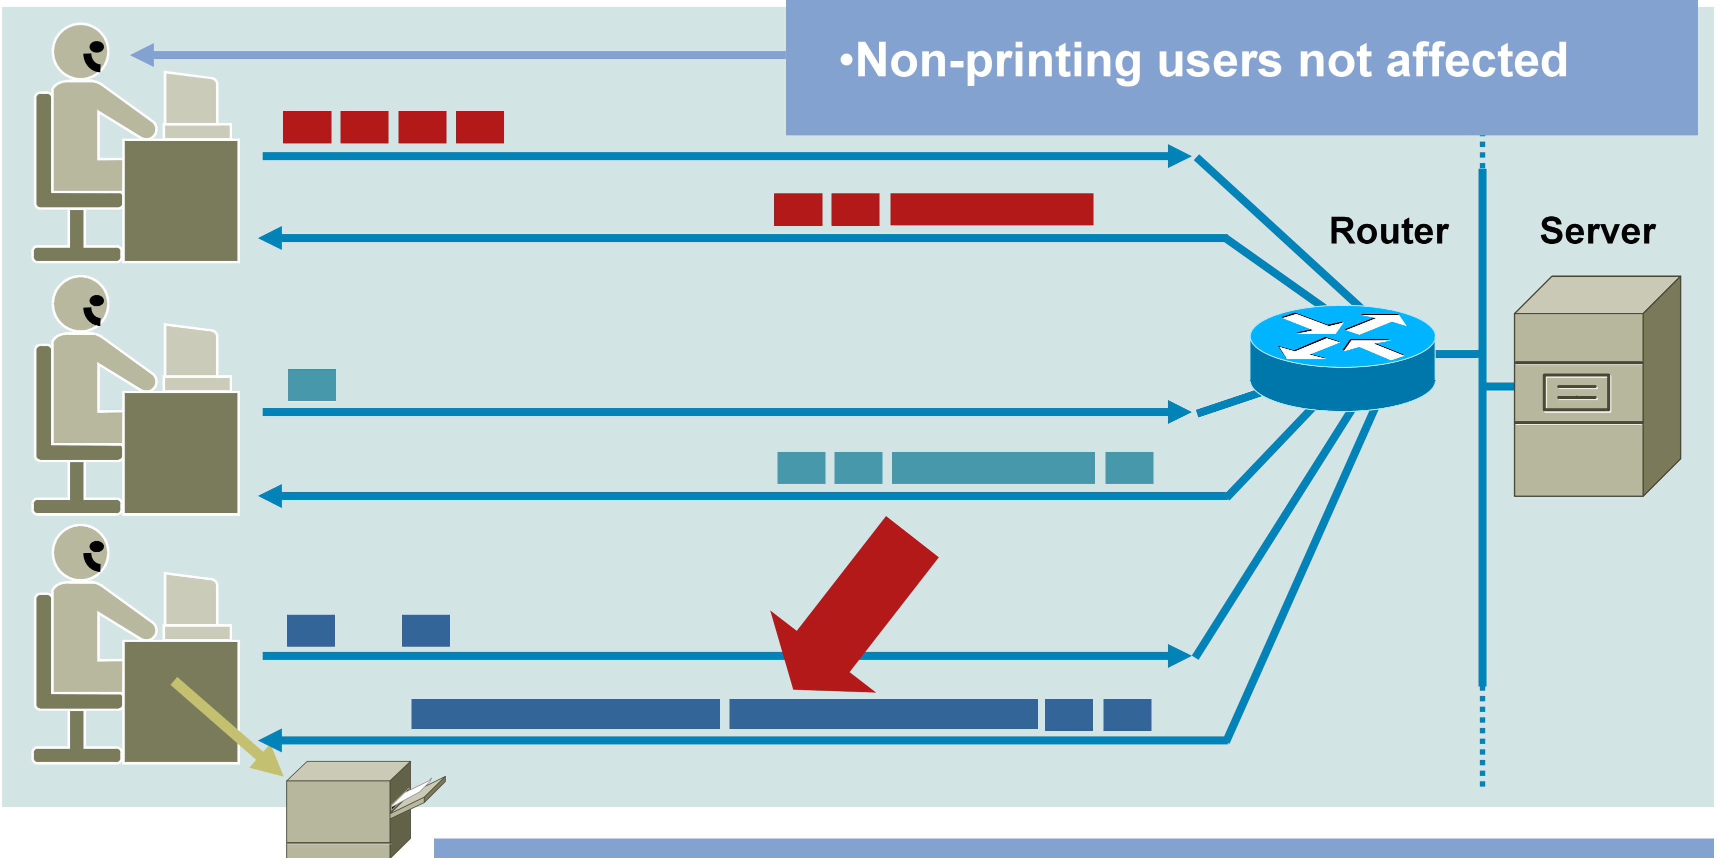

- While printing, lower priority given to that session

Cisco IOS QoS Behavioral Model

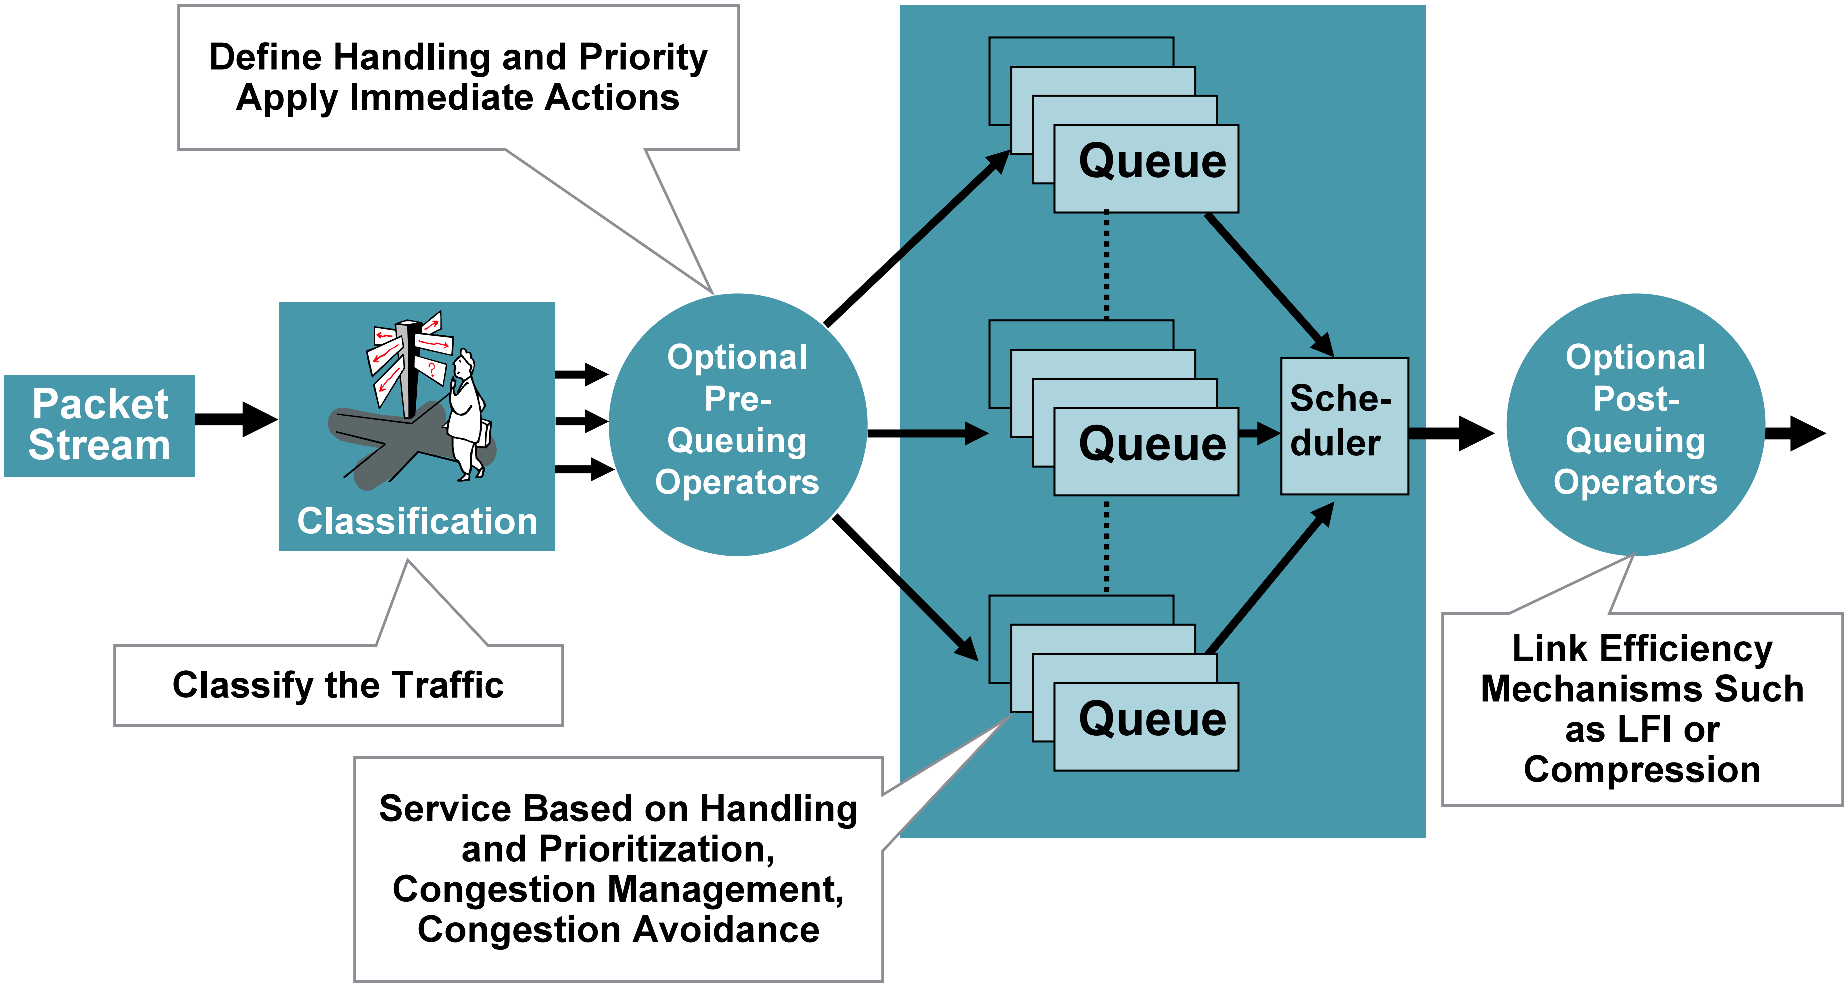

Operators for Traffic Classification and QoS Policy Actions

<table>
<tr><th>Match Conditions
Keyword:
class-map</th><th colspan="3">Policy Actions
Keyword: policy-map</th></tr>
<tr><td>Classification</td><td>Pre-Queuing</td><td>Queuing and Scheduling</td><td>Post-Queuing</td></tr>
<tr><td>Classify Traffic</td><td>Immediate Actions</td><td>Congestion Management and Avoidance</td><td>Link Efficiency Mechanisms</td></tr>
<tr><td>Match One or More Attributes (Partial List):
• ACL list
• COS
• DSCP
• Input-interface
• MAC address
• Packet length
• Precedence
• Protocol
• VLAN</td><td>• Mark (set QoS values)
• Police
• Drop
• Count
• Estimate bandwidth</td><td>• Queue-limit
• Random-detect
• Bandwidth
• Fair-queue
• Priority
• Shape</td><td>• Compress header
• Fragment
(Link fragmentation and interleaving, layer 2)</td></tr>
</table>

IP SLAs Ensure Application Performance

IP Service Level Agreements (SLAs) Ensure IP Service Levels, Proactively Verifies Network Operation, and Accurately Measures Network Performance

- Actively generates traffic to measure and monitor the network and performance
- Measures network characteristics such as latency, packet loss, and jitter
- Generate notifications or trigger other IP SLAs to remedy

Example: Multiprotocol Measurement and Management with Cisco IOS IP SLAs

Uses

Availability	Network Performance Monitoring	VoIP Monitoring	Service Level Agreement (SLA) Monitoring	Network Assessment	Multiprotocol Label Switching (MPLS) Monitoring	Trouble Shooting

Measurement Metrics

Latency	Packet Loss	Network Jitter	Dist. of Stats	Connectivity

Operations

Jitter	FTP	DNS	DHCP	DLSW	ICMP	UDP	TCP	HTTP	LDP	H.323	SIP	RTP	Radius	Video

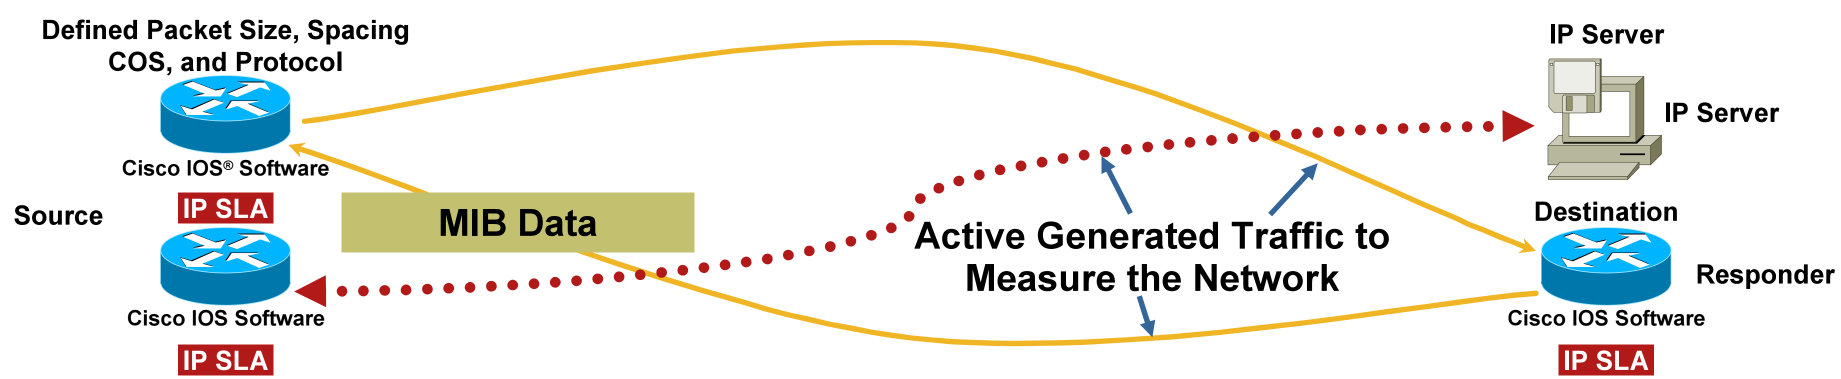

Summary of Traffic Differentiation

Traffic Differentiation Technologies Allow Administrators to Configure the Network to Align with Business and Application Requirement

- First, understand who and what are using the network by using NetFlow
- Second, apply policy based on business priority using Quality of Service (QoS) and Network-Based Application Recognition (NBAR)
- Third, ensure network is performing and reacting according to application requirements

WAN Optimization

Agenda

- Introduction to Cisco WAAS
- Mitigating Application Latency
- Managing Bandwidth Utilization
- Improving Transport Performance
- WAAS Integration and Deployment

Cisco WAAS Enables Consolidation

- Cisco Wide Area Application Services (WAAS)
  - Transparent integration
  - Robust optimizations
  - Auto discovery
- Infrastructure consolidation
  - Remote costly servers
  - Centralize data protection
  - Save WAN resources
- Application acceleration
  - Application adapters
  - Advanced compression
  - Throughput optimizations
  - Policy-based configuration

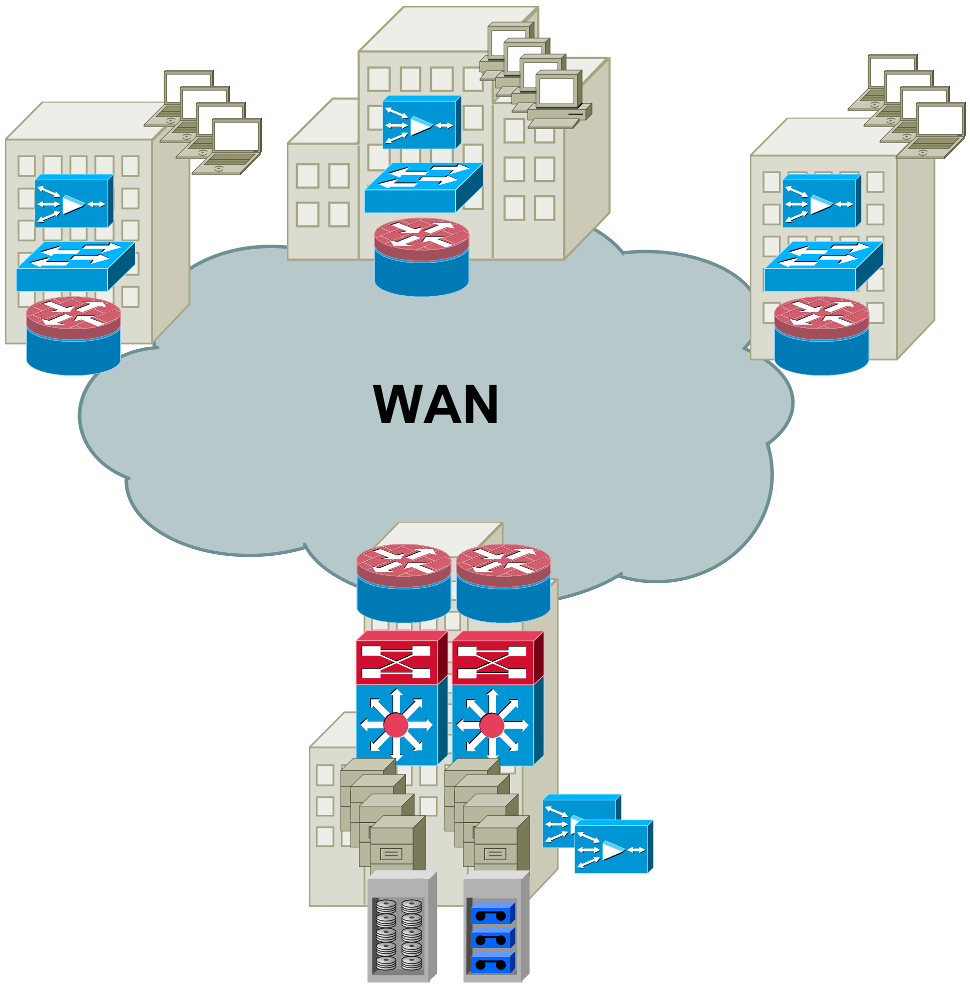

WAAS Addresses WAN Performance Impact

Problem	Solution	Cisco IOS/WAAS Technology
Latency Mitigation	• Reduced roundtrips from chatty application protocols • Faster connection setup	• Intelligent Protocol Proxies • Transport Flow Optimizations (TFO)
Bandwidth Management	• Offload the WAN by preventing requests from going to the WAN • Improve application response time on congested links by reducing the amount of data sent across the WAN	• Caching • Data Redundancy Elimination (DRE) • Persistent Session-Based Compression • Content Distribution & Pre-positioning
Link Throughput Improvement	• Improve network throughput by reducing TCP-related errors	• Transport Flow Optimizations (TFO)
Traffic Prioritization	• Prioritize selected jitter-sensitive traffic (e.g. VoIP, Video) over the packet network	• Cisco IOS • QoS, NBAR, NetFlow
Local Services	• Replacement for services that branch office servers provide	• Centrally managed remote services interface • Local print services

WAAS Accelerates Broad Range of Applications

Application	Protocol	Typical Improvement
File Sharing	• Windows (CIFS) • UNIX (NFS)	• 2X–100X
E-mail	• Exchange (MAPI) • SMTP/POP3, IMAP • Notes	• 2X–50X
Internet and Intranet	• HTTP, HTTPS, WebDAV	• 2X–50X
Data Transfer	• FTP	• 2X–50X
Software Distribution	• SMS • Altiris	• 2X–100X
Database Applications	• SQL • Oracle • Notes	• 2X–10X
Data Protection	• Backup applications • Replication applications	• 2X–10X
Other	• Any TCP-based application	• 2X–10X

*** Performance Improvement Varies Based on User Workload, Compressibility of Data, and WAN Characteristics and Utilization. Actual Numbers Are Case-Specific and ResultsMay Vary.**

Cisco WAAS Optimization Architecture

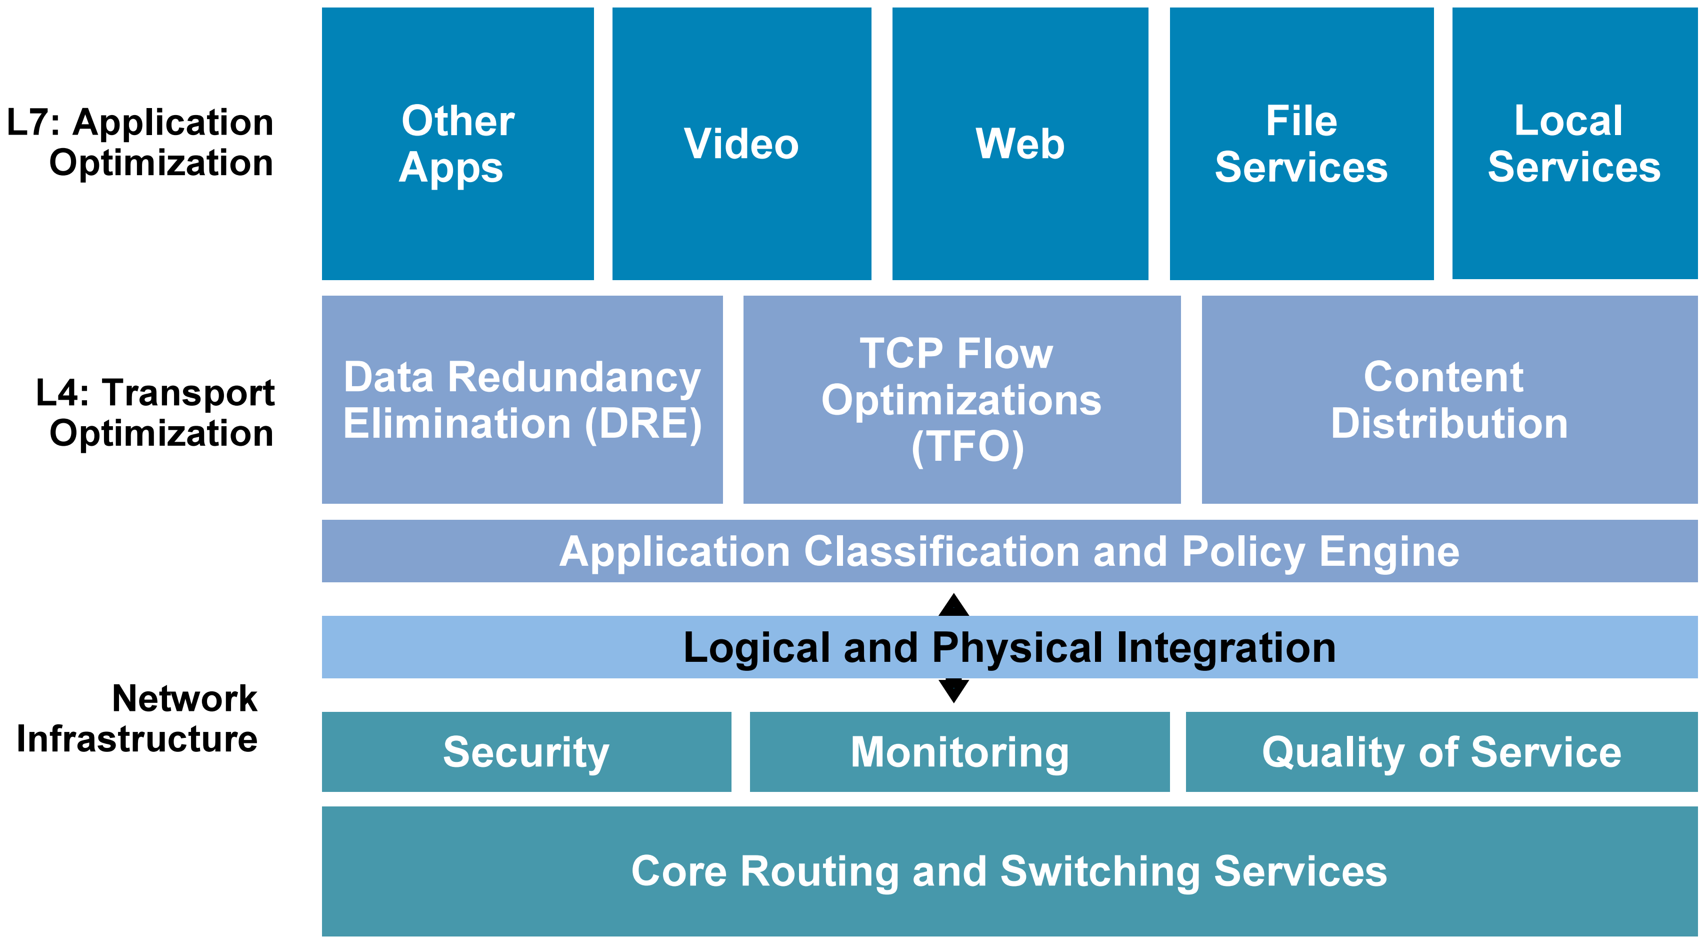

Cisco WAAS Deployment Architecture

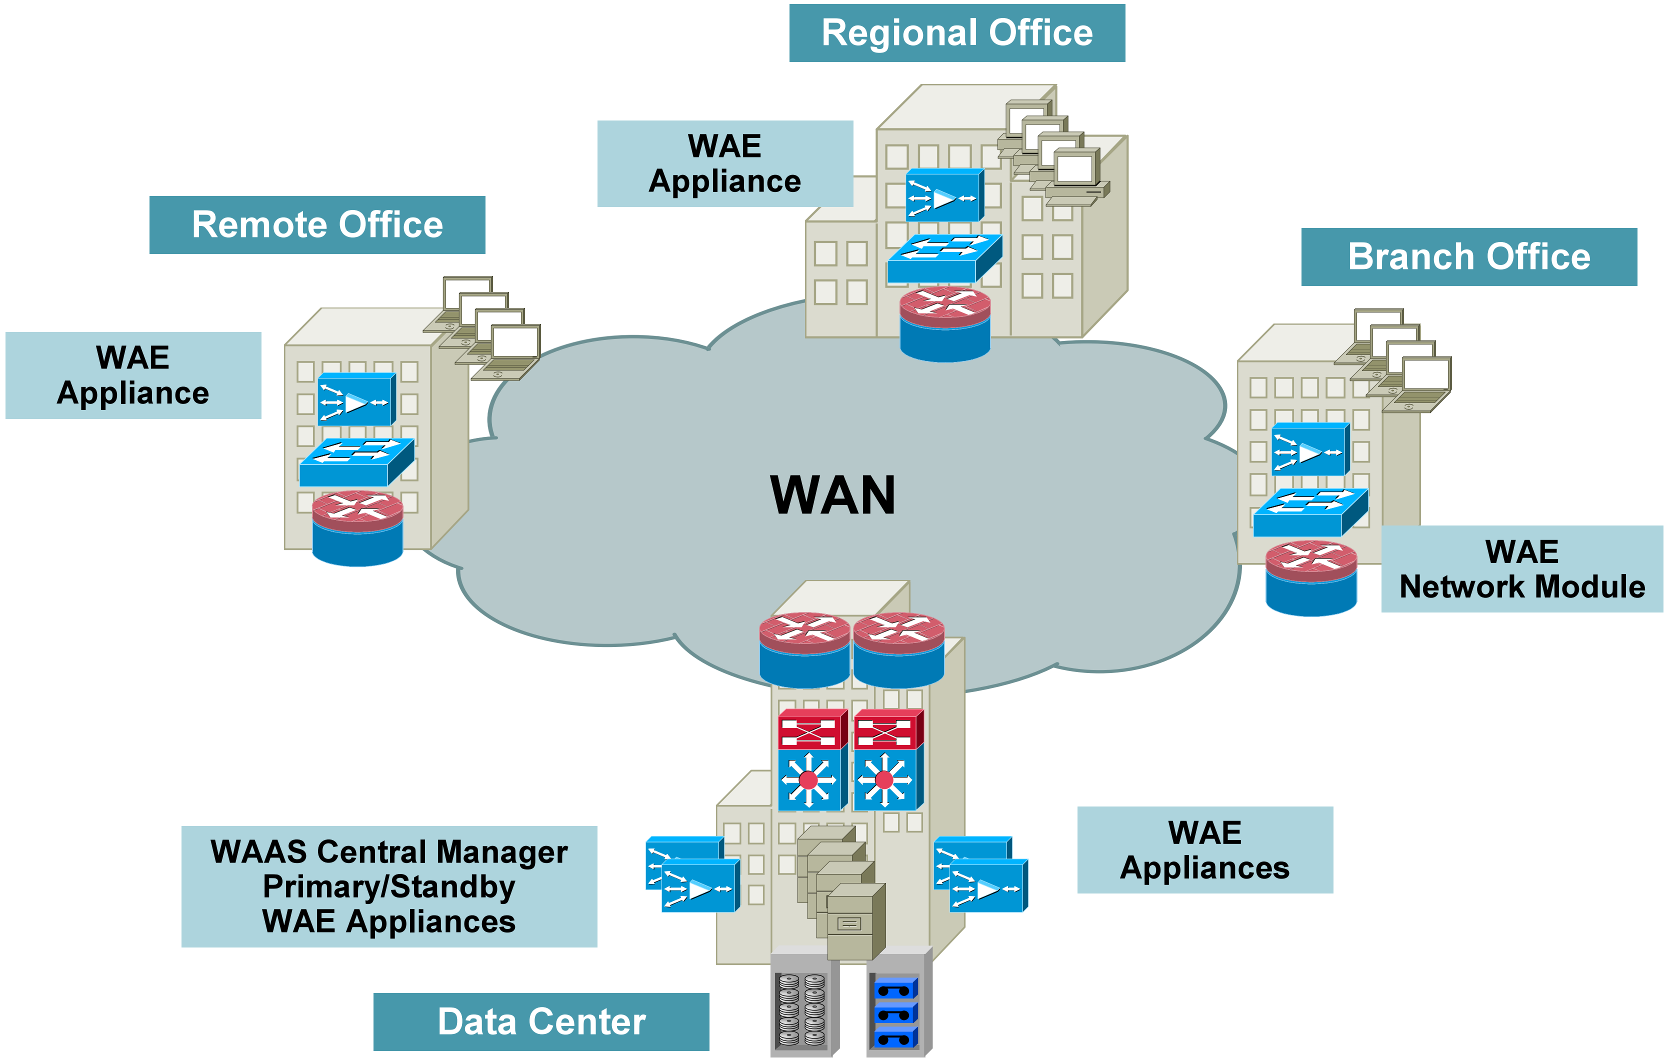

Network Interception

Network Attached Optimizations Rely on Devices Physically Attached to the Network at Strategic Locations

- Generally deployed at network entry/exit points
- Rely on network interception to supply flows to optimize

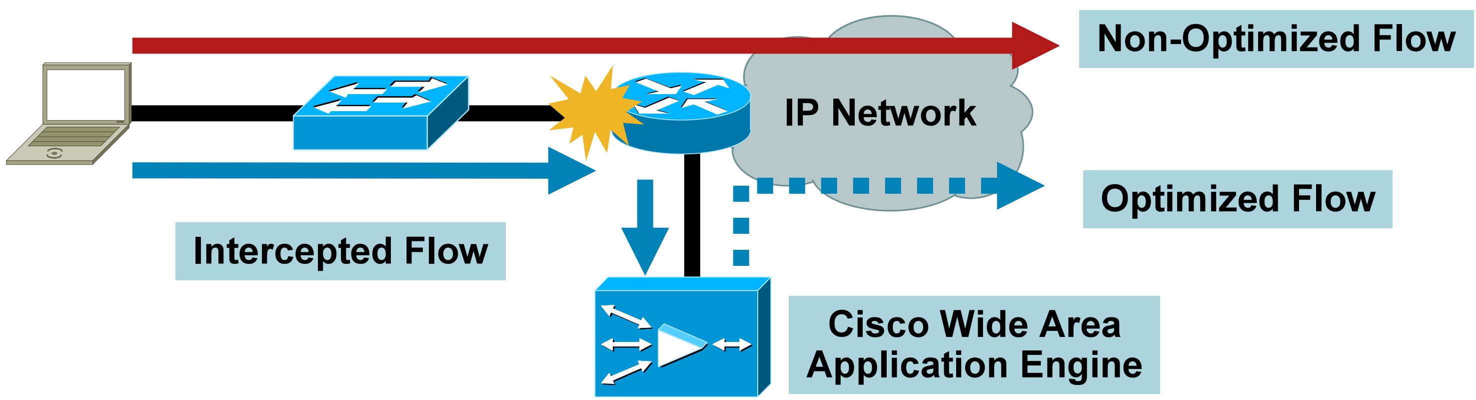

Flexible Acceleration Policies

Application Acceleration Must Provide Users with Flexible Configuration of Optimizations—Not All Flows Are Created Equal

- Low layer implementation to ensure high performance
- Default policies provided but able to be modified

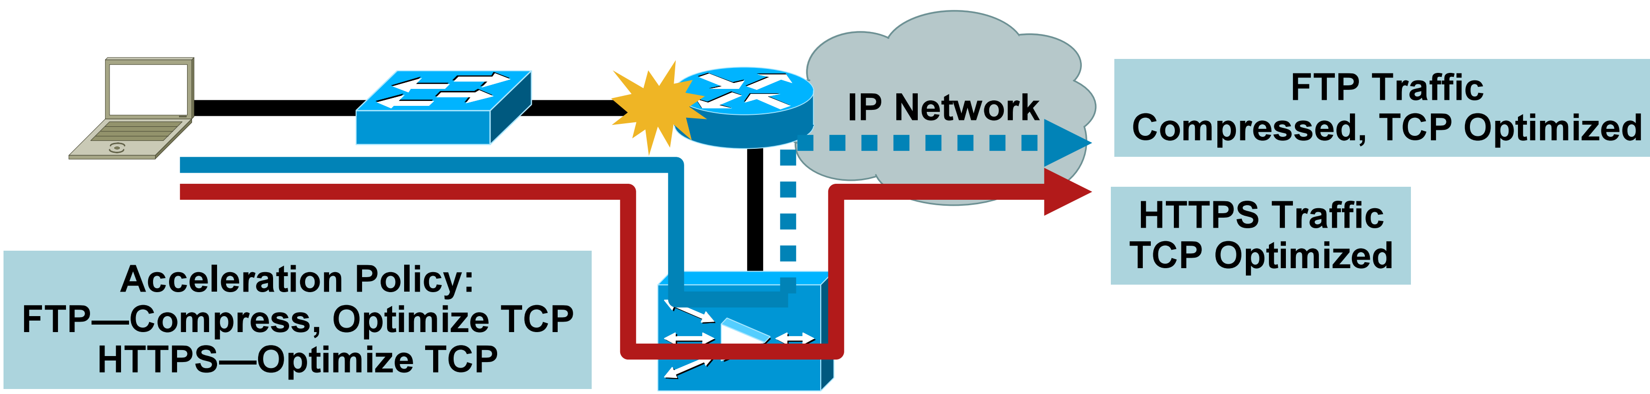

Application Acceleration Transparency

- Packet network transparency (L3/L4 headers) allows application acceleration components to maintain compliance with existing network features
  - Quality of Service (QoS), NBAR
  - NetFlow, monitoring, reporting
  - Security functions (ACLs, firewall policies)
- If source/destination L3/L4 information is not preserved, these features may need to be reconfigured to support application acceleration

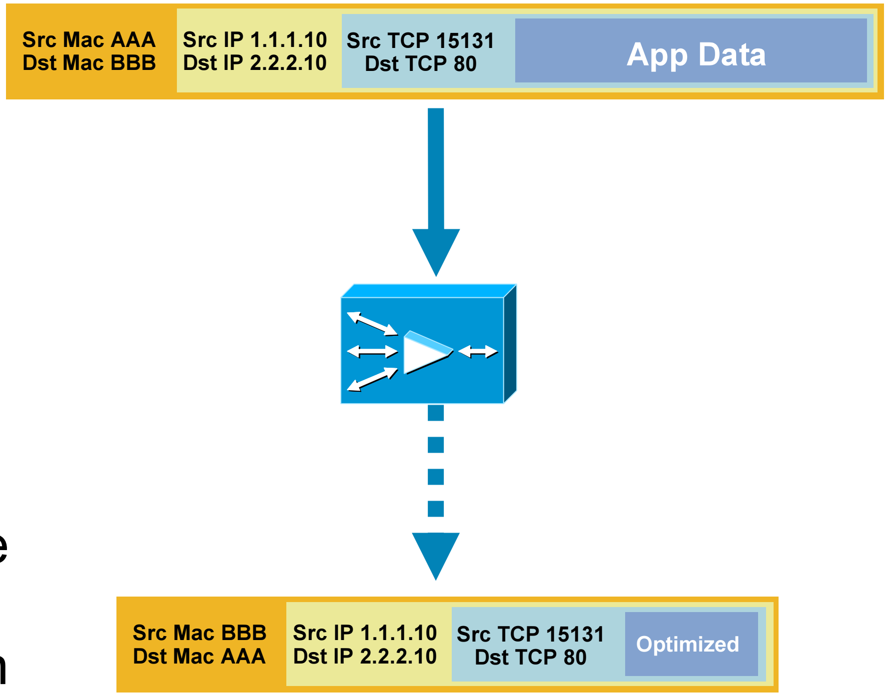

Application Latency Example—CIFS

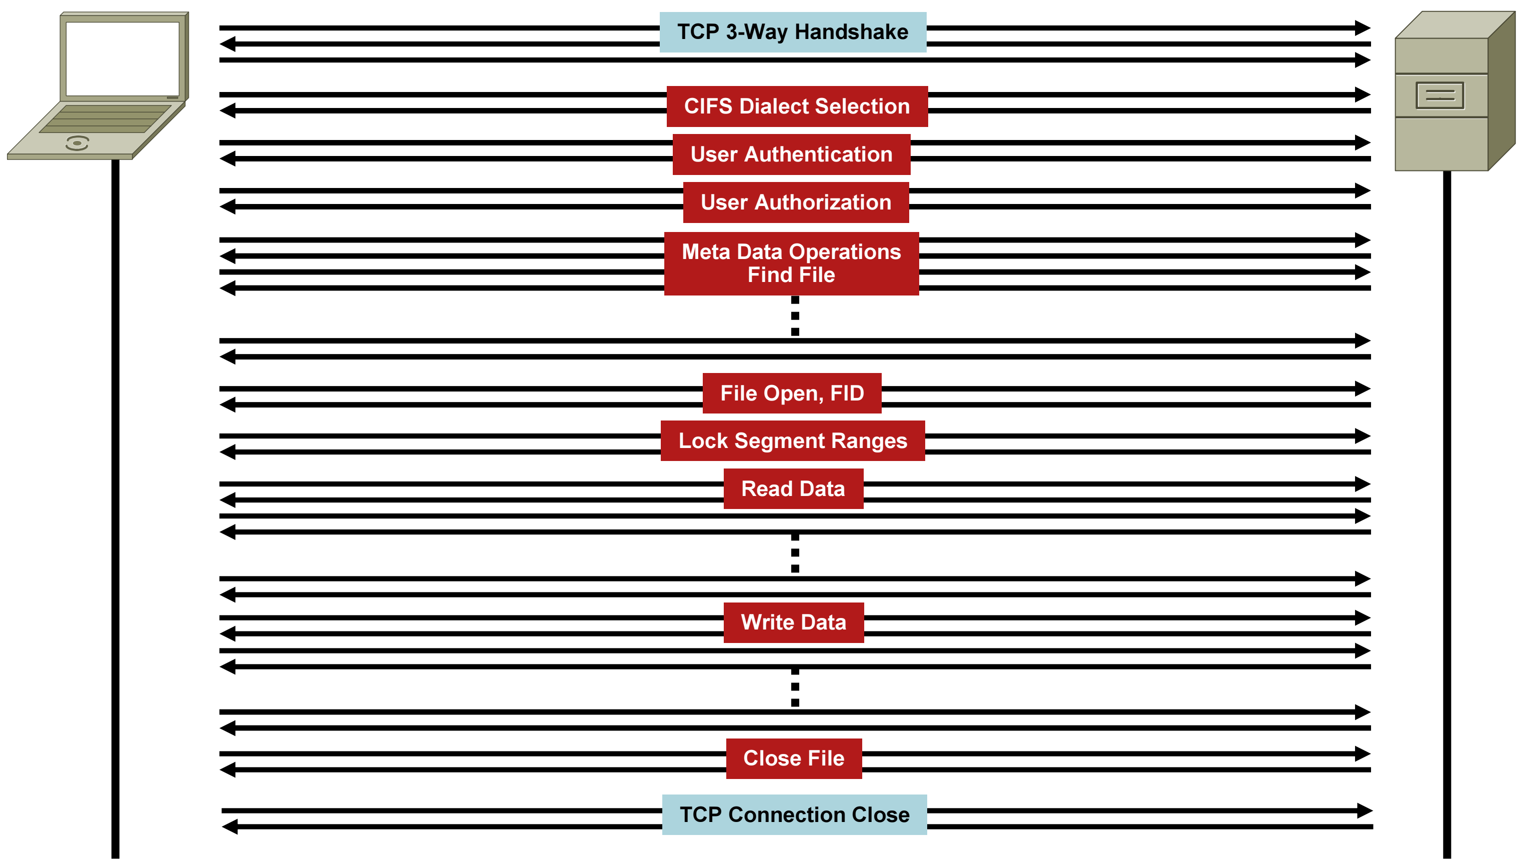

Application Latency Example—CIFS

- In this simple example of a 1MB Word document open, over 1,000 messages are exchanged
- With a 40mS RTT WAN, this equates to over 52 seconds of “wait” time before the document is usable

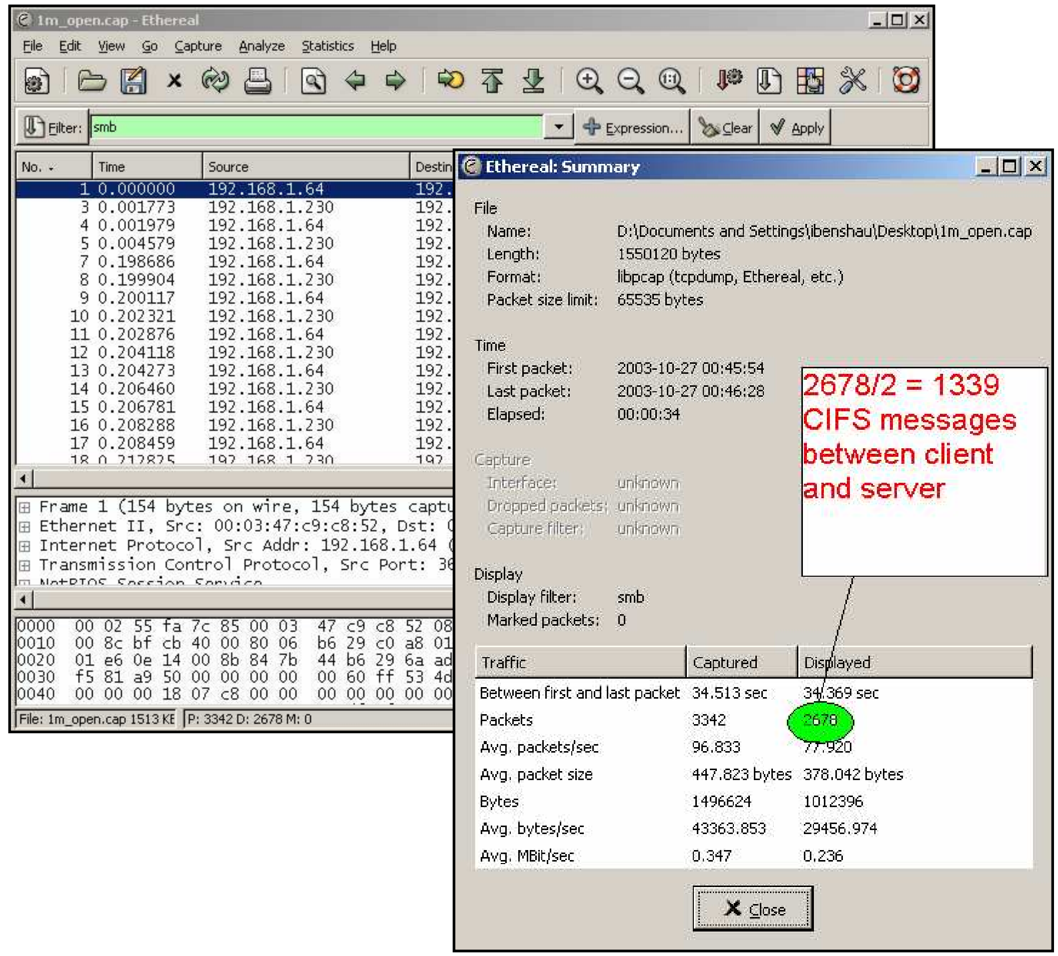

Compress or TCP Optimize?

- Chatty application protocols require exchange of many messages to ensure proper operation of the applications that are using the protocol
- Many of these messages are zero-byte length
  - Hard to compress zero-byte messages ☺
  - Messages still must be exchanged
- The transport (TCP) is rarely the limiting factor
  - Improving TCP does not mitigate message exchange

Managing Bandwidth Utilization

The Need for Compression

- Advanced compression technologies allow customers to virtually increase WAN capacity
- Allows customers to leverage existing WAN capacity and may mitigate the need for a costly bandwidth upgrade

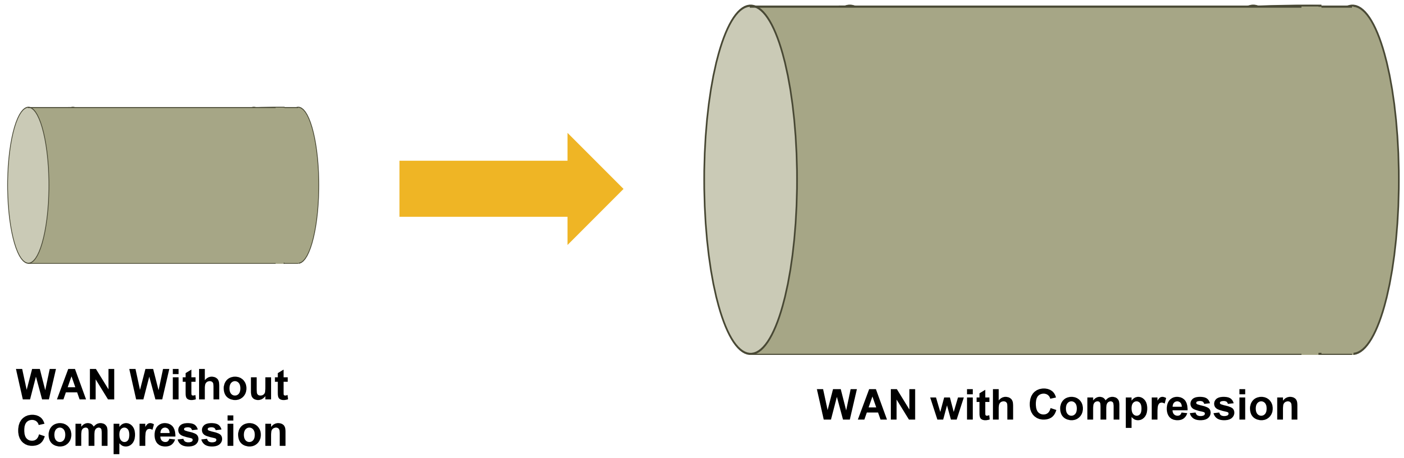

The Need for Compression

- Some data sets are not good candidates for compression unless adaptation is first performed

    Previously-compressed data—no additional compression provided by computational compression, good candidate for data suppression

    Previously-encrypted data—minimal additional compression provided by computation compression,
    good candidate for data suppression if not using session-based encryption (i.e., non-repeatable data)

- Such adaptation could include local termination of encryption, apply compression, then re-encrypt

Advanced Compression Overview

Two Forms of Compression (Together) Enable Significant Savings of WAN Bandwidth

- Data suppression (DRE): store chunks of TCP traffic patterns in loosely-synchronized contexts to suppress transmission of redundant chunks
- Standards-based compression: i.e., Lempel-Ziv, deflate

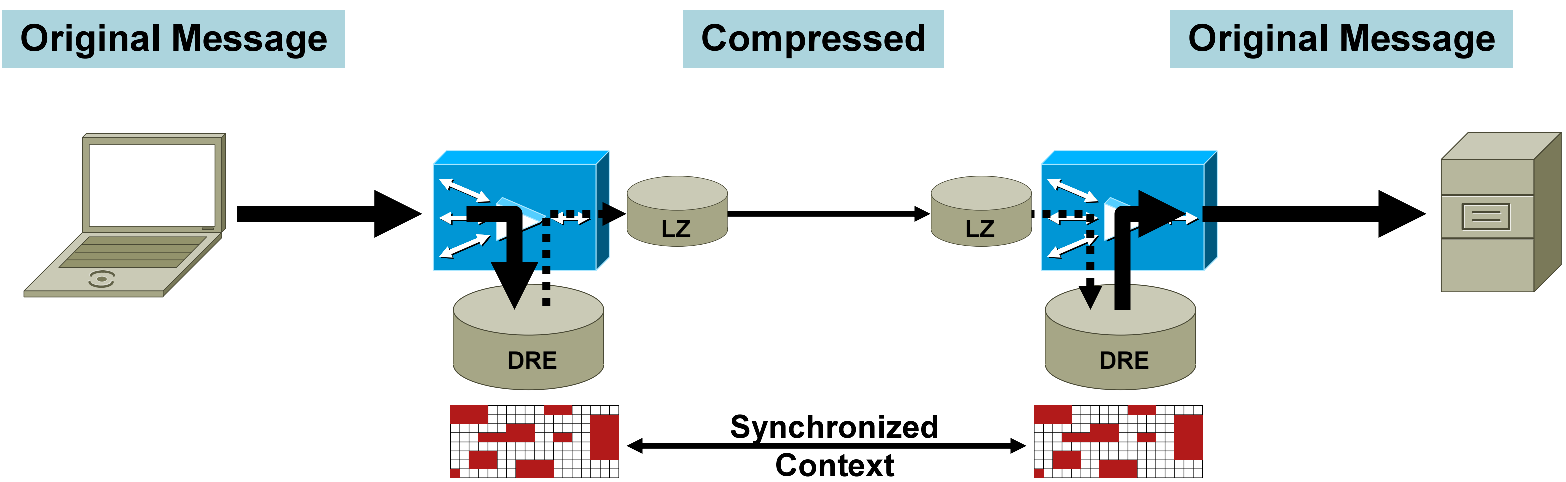

Advanced Compression Block Diagram

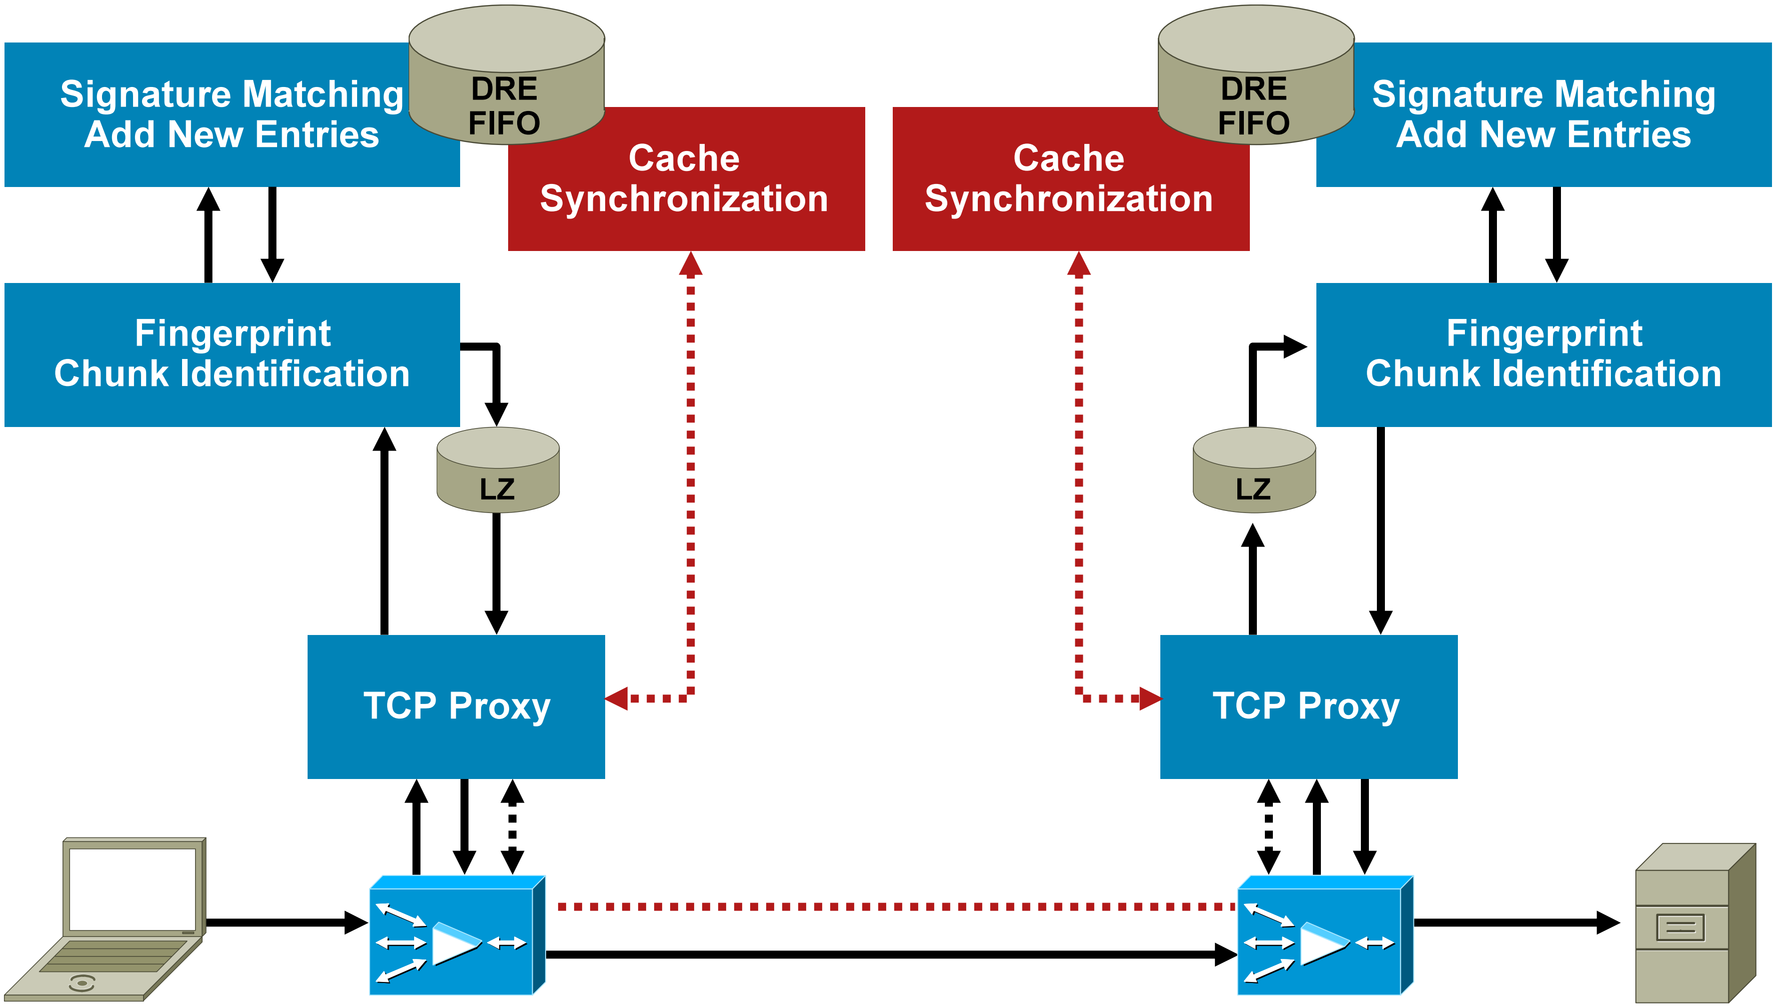

DRE Encoding—Pattern Matching

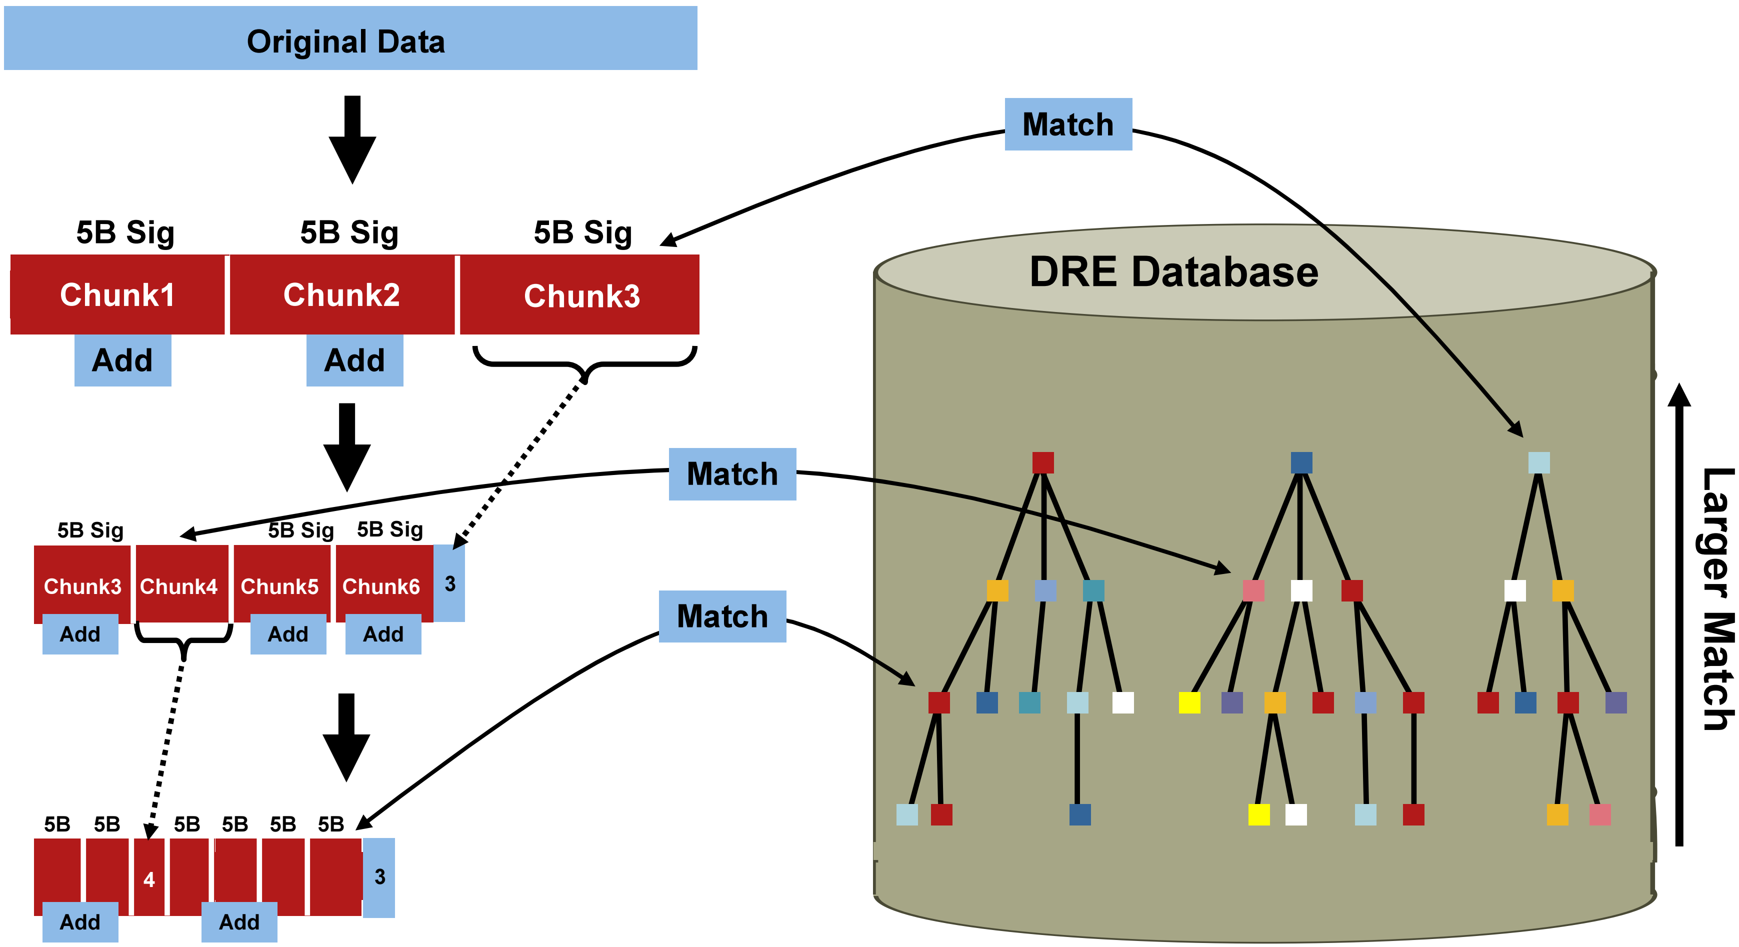

DRE Encoding—Resultant Message

DRE Sender, Cont.

- A fully encoded message will contain:
  - Signatures only for previously-seen patterns
  - Signatures, data for non-redundant patterns (update adjacent WAE)
  - 16-byte MD5 hash of original message to verify integrity after rebuild
- Message is passed to LZ compression (based on policy) and to TCP proxy to return to the network

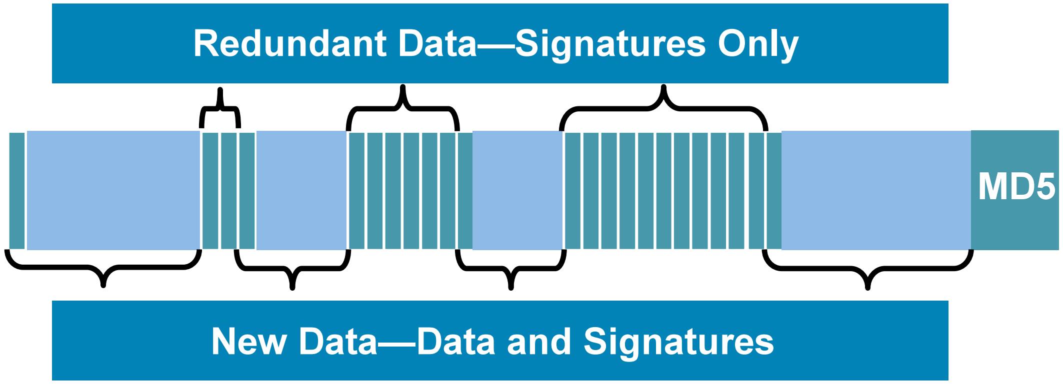

DRE Synchronization: Example 1

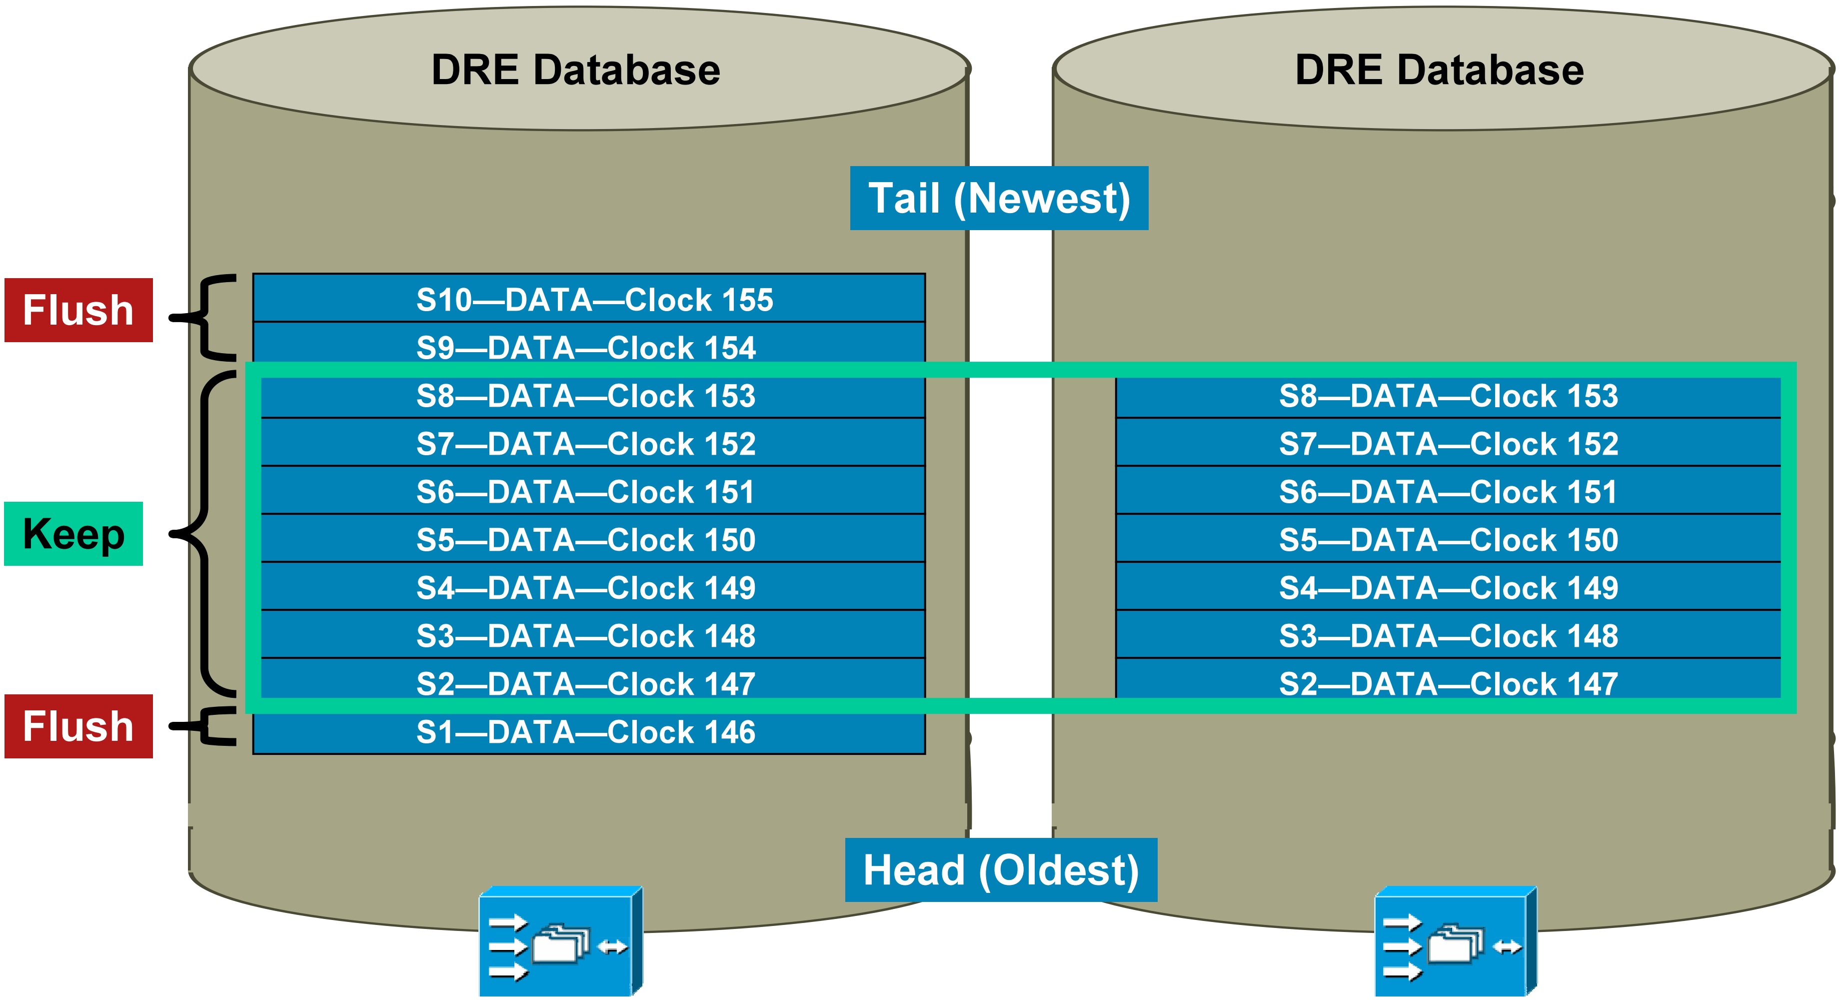

DRE Synchronization: Example 2

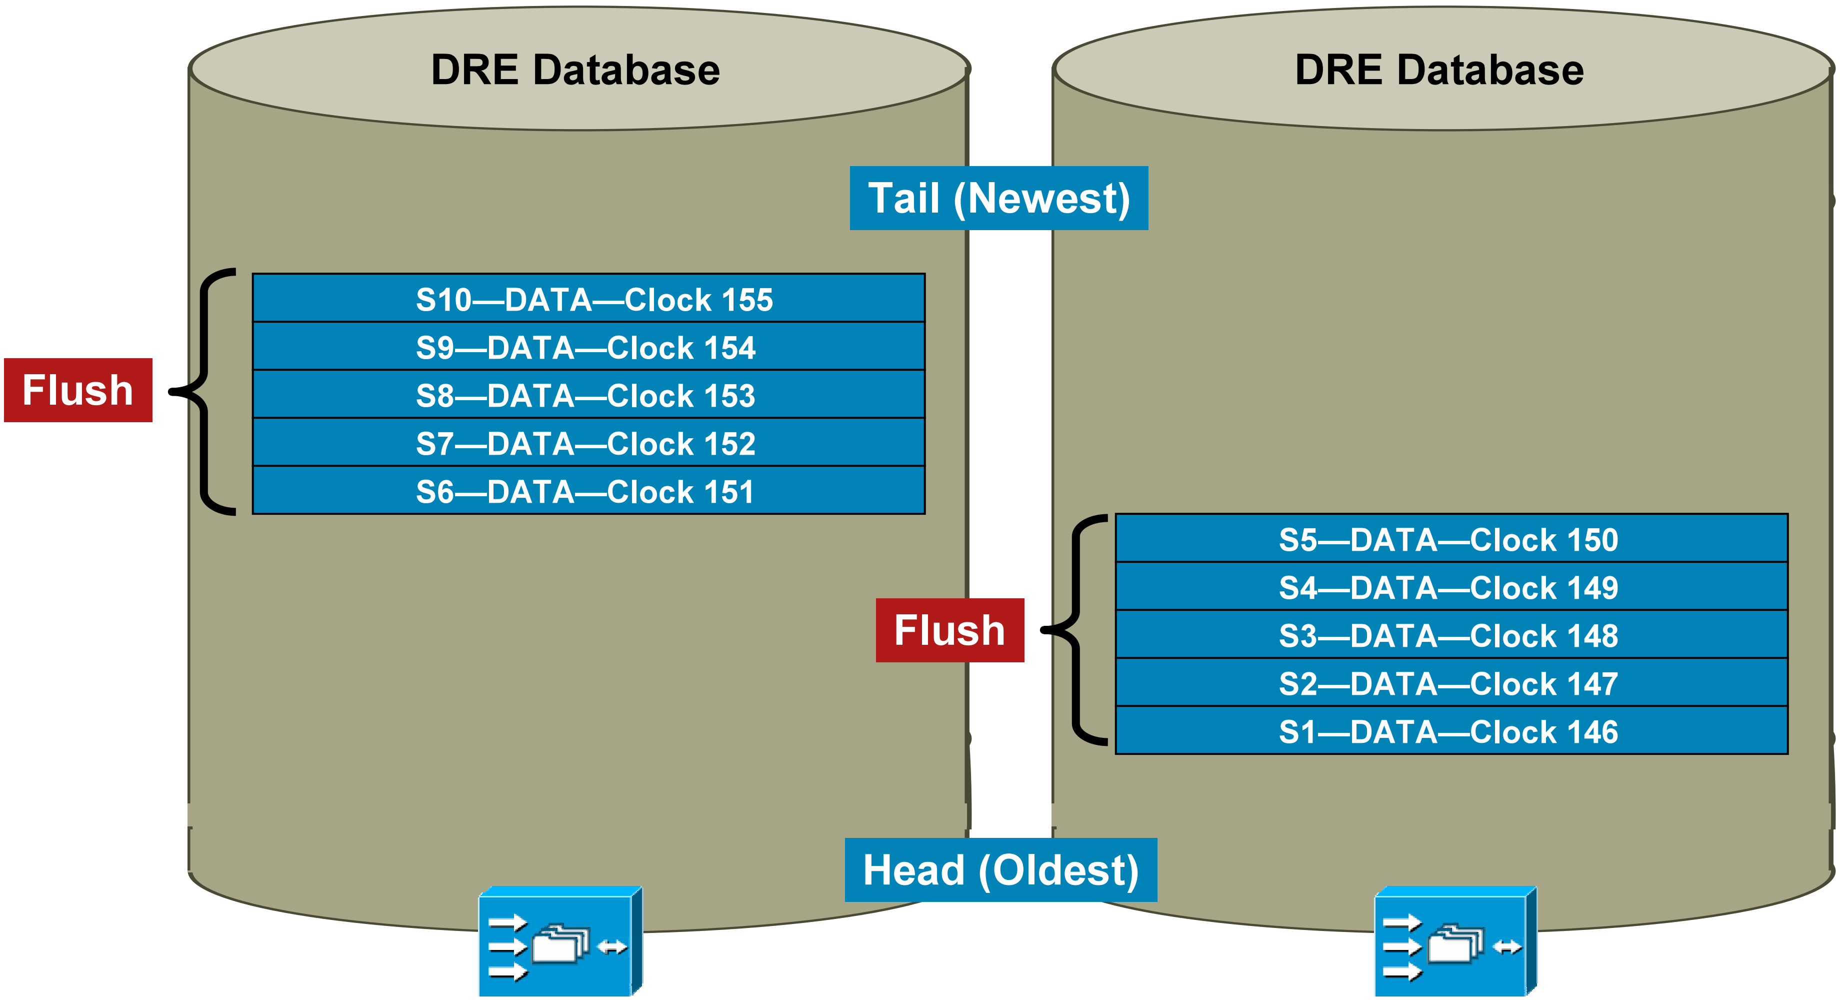

Impact of Advanced Compression

- Advanced compression can significantly minimize the amount of data that traverses the WAN
- Flows are safely rebuilt in their entirety at the distant end, allowing large amounts of application data to traverse the network

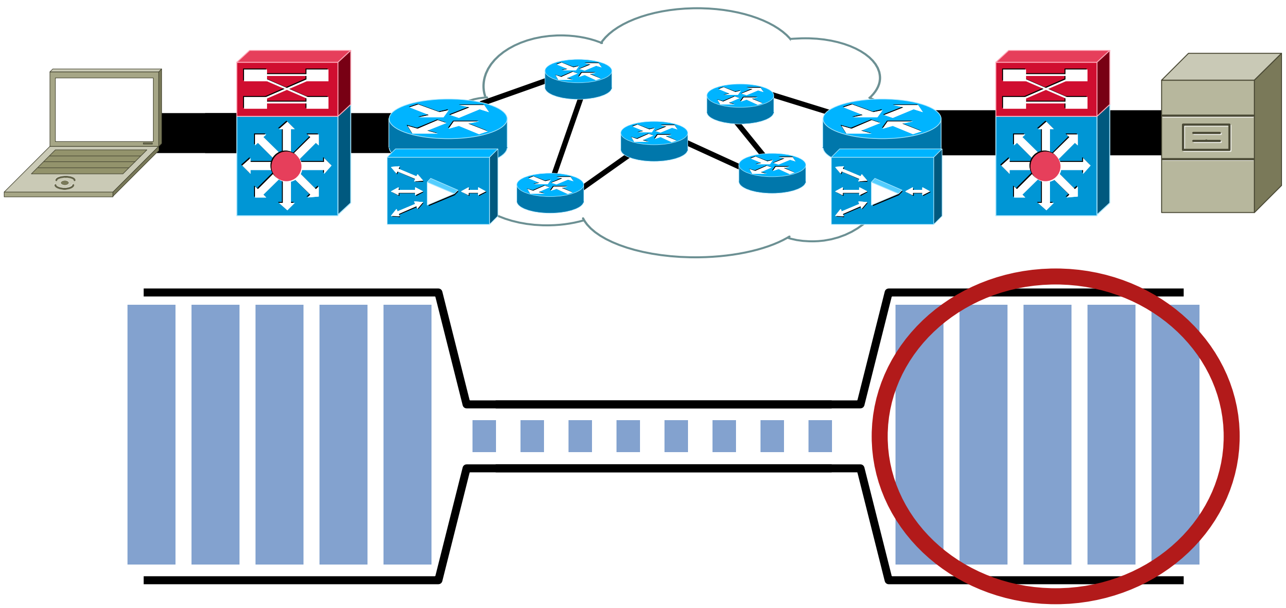

Improving Transport Performance

Challenge

- Common TCP implementations on client and server operating systems can be bottlenecks to application performance
  - Inability to fill-the-pipe, i.e., utilize available bandwidth
  - Inefficient recovery from packet loss, retransmission
  - Bandwidth starvation for short-lived connections
- Cisco WAAS Transport Flow Optimization (TFO) utilizes industry-standard TCP optimizations to remove these application performance barriers

TCP Maximum Window Size (MWS)

- MWS (maximum window size) determines the maximum amount of data that can be in transit and unacknowledged at any given time
- BDP (bandwidth delay product) defines the amount of data that can be contained within a network at any given time

  If MWS > BDP, then application may not be throughput bound (i.e., application can “fill the pipe”)

  If BDP > MWS, then application will not be able to fully utilize the network capacity (i.e., application can not “fill the pipe”)
- Does not account for application-layer (L7) latency such as found with protocol-specific messaging

Link Utilization and MWS, BDP

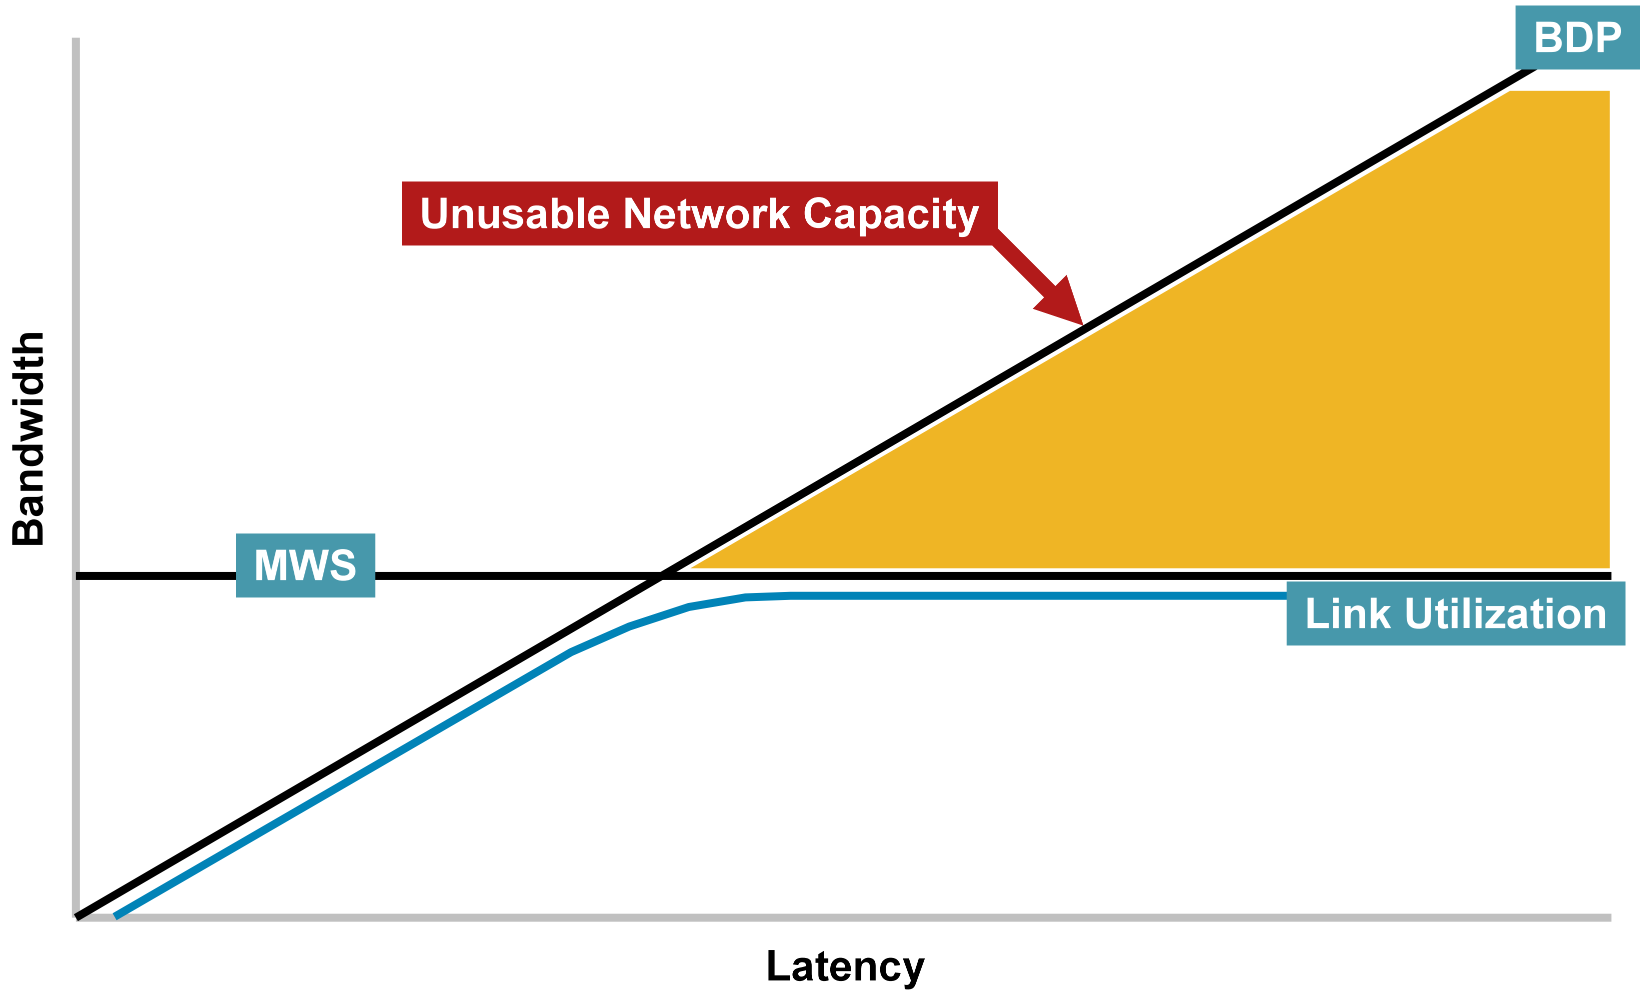

Standard TCP Congestion Avoidance

- Standard TCP implementations employ an exponential slow start to increase throughput to the slow start threshold
- From the slow start threshold, the congestion window is increased linearly by one packet per round-trip until packet loss is encountered
- Upon encountering packet loss, the congestion window is cut in half to return to a throughput level safe given the congested environment
- The net result is “saw-tooth” throughput, and return to maximum throughput can take hours for long-lived connections and LFNs

WAAS TCP Flow Optimization Techniques

- **Windows Scaling**
  - RFC 1323—TCP Performance Extensions—defines the use of a TCP option to scale the TCP window beyond the standard 16-bit limitation (64KB)
  - Cisco WAAS provides window scaling up to 2MB per optimized TCP connection
- **Large Initial Windows**
  - 80% of connections are Short-lived connections
  - Short lived connections transmit smaller numbers of packets and are torn down before ever leaving the slow-start phase of TCP
  - Cisco WAAS Large Initial Windows, based on RFC3390, increases initial window size to expedite entry into congestion avoidance mode for high throughput
- **Selective Acknowledgement**
  - With standard implementations, receipt acknowledgement is done once the entire window has been received and therefore with standard TCP, a Loss of a packet causes retransmission of the entire TCP window, causing performance degradation as the window becomes larger
  - WAAS selective acknowledgement, Improves acknowledgement of transmitted data and improves delivery of missing segments and in turn minimizes unnecessary retransmission
- **Binary Increase congestion (BIC)**
  - Uses a binary search to adaptively increase the congestion window, resulting in a stable and timely return to higher levels of throughput

Link Utilization After Window Scaling

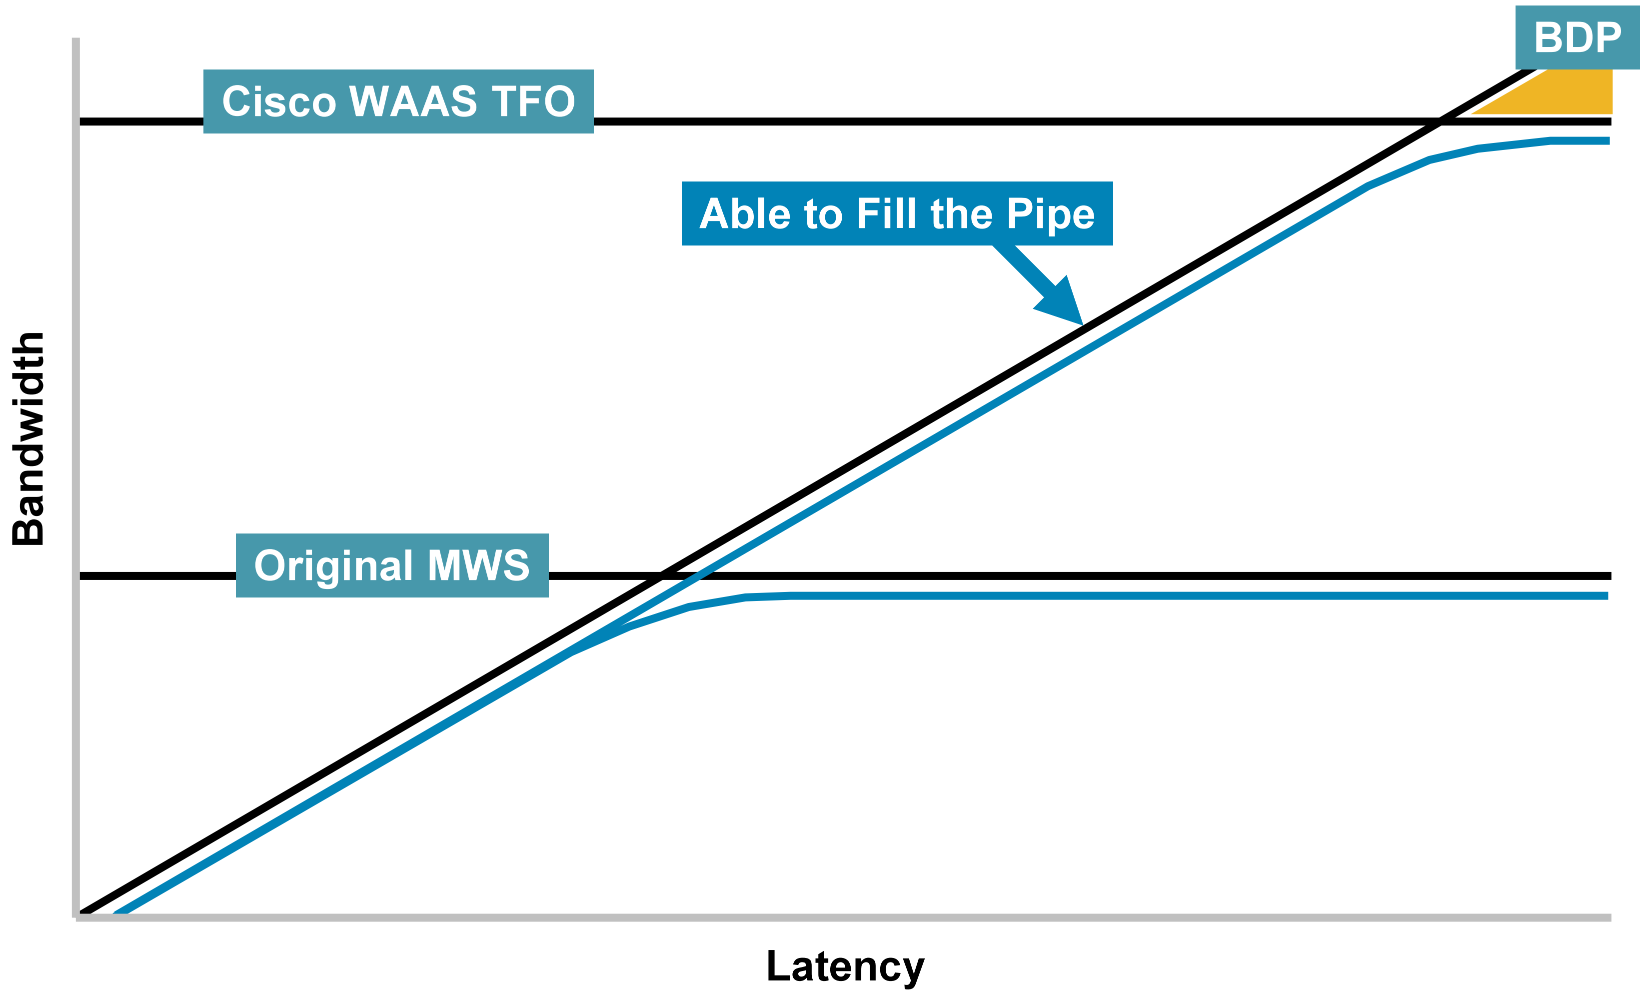

Selective Acknowledgement

- Standard TCP

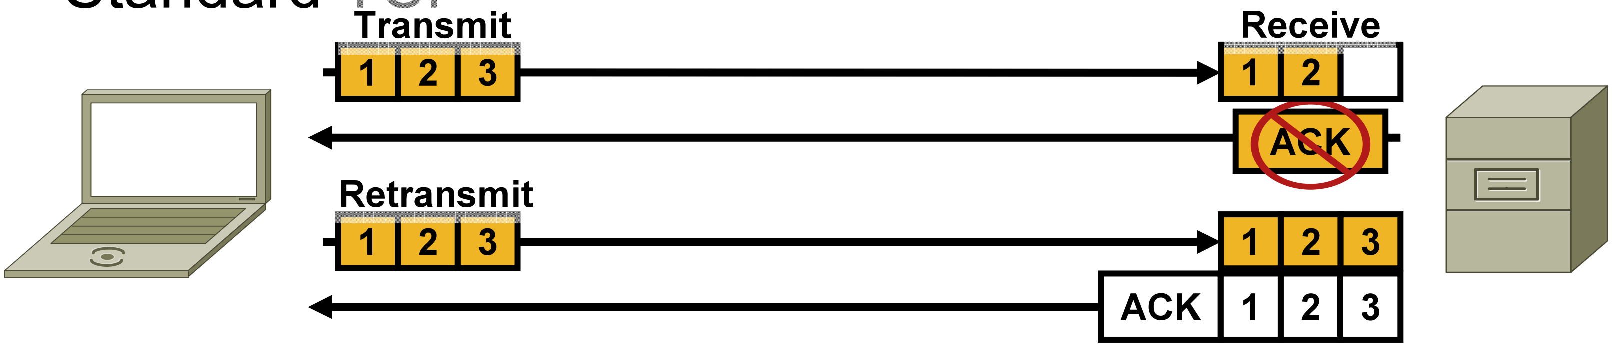

- WAAS Selective Acknowledgement

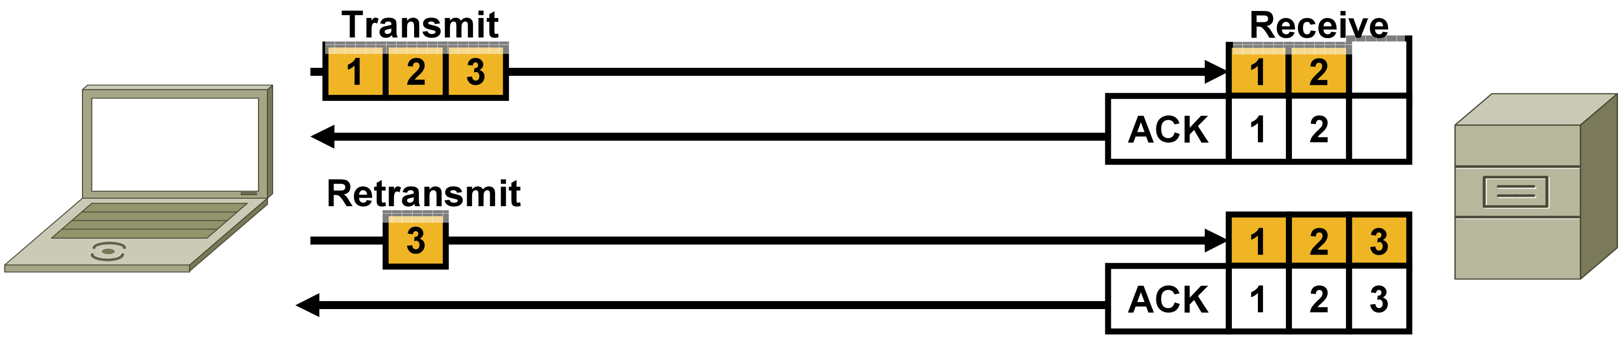

WAAS Throughput and Congestion Avoidance

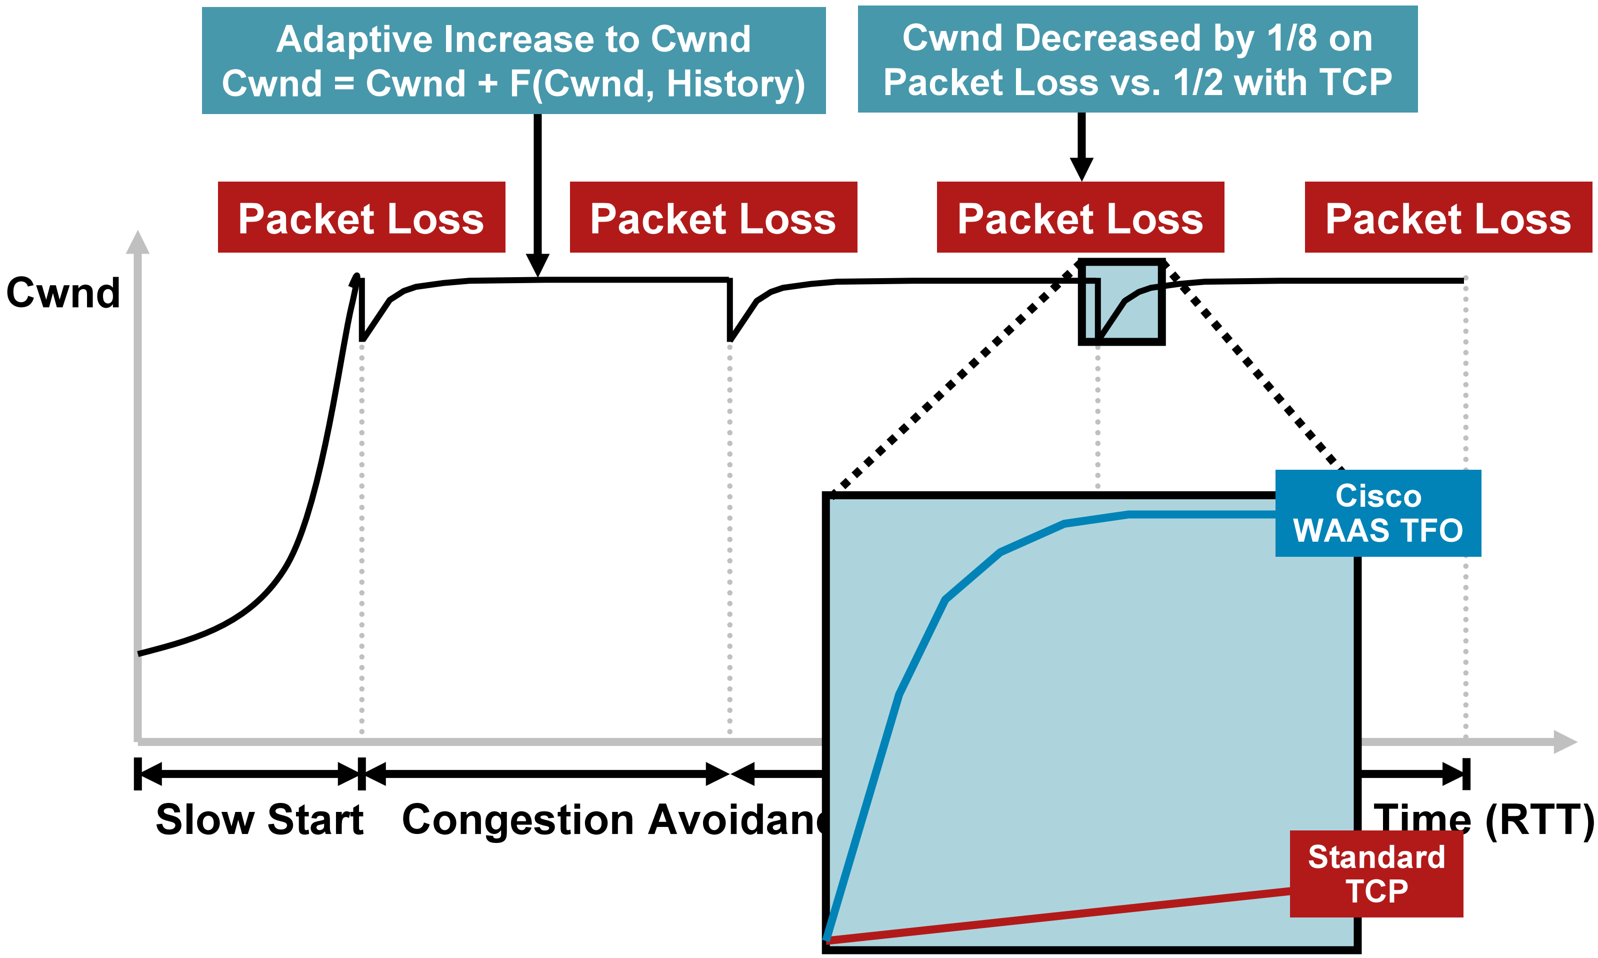

Comparing TCP and TFO

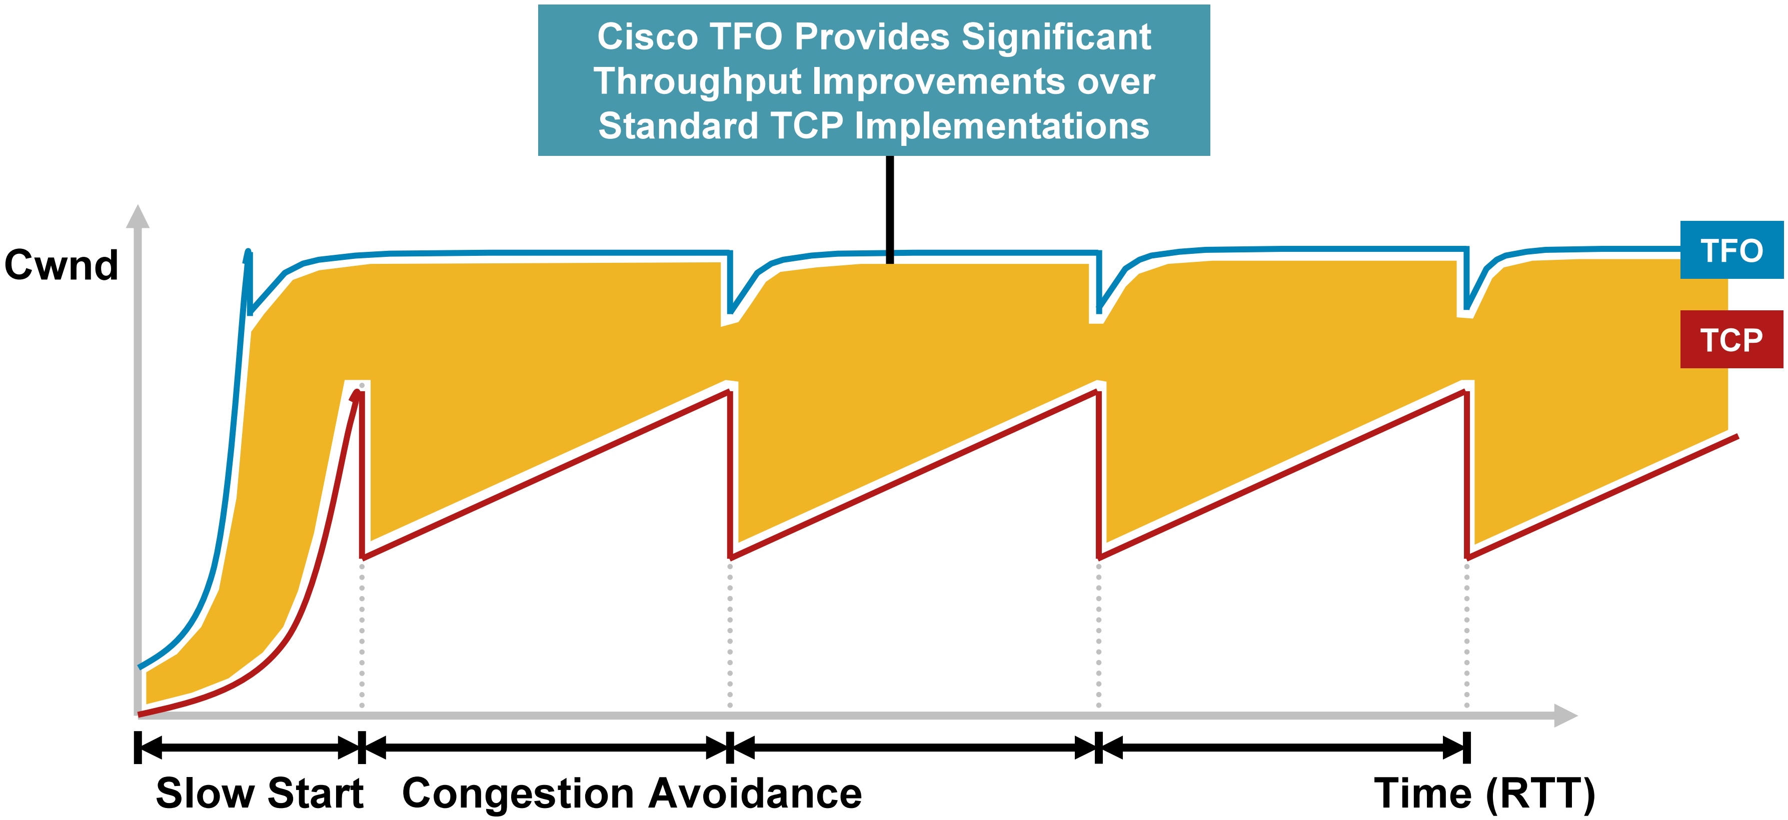

Impact of Transport Flow Optimizations

- TFO overcomes TCP performance bottlenecks
- Shields nodes connections from WAN conditions
  - Clients experience fast acknowledgement
  - Minimize perceived packet loss
  - Eliminate need to use inefficient congestion handling

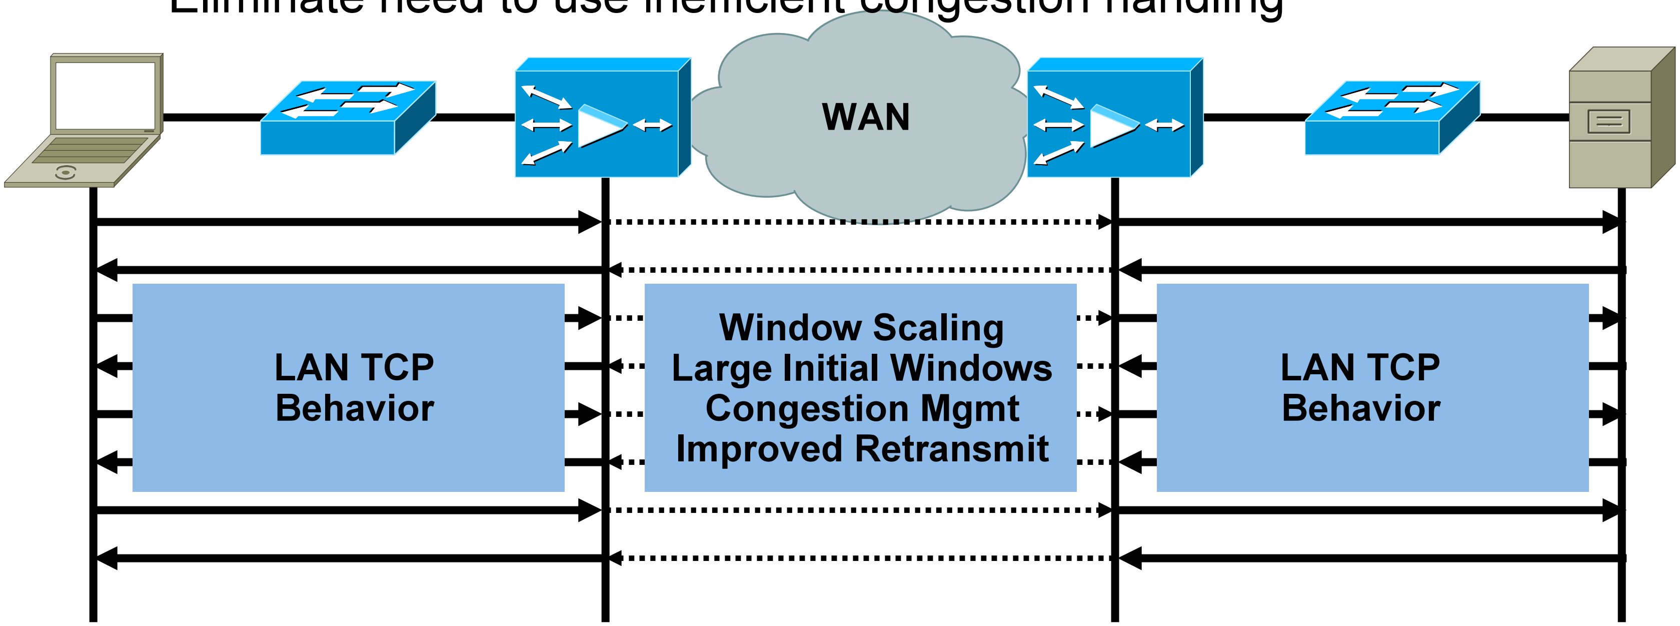

WAAS Integration and Deployment

WAAS Interception methods

- PBR
- WCCP
- Inline
- CSM/ACE

Cisco WAE PBR Deployment

- Policy-Based Routing (PBR)
  - Out-of-path with redirection of flows to be optimized (all flows or selective via access-list)
  - WAE treated as a next-hop router
- High availability
  - Failover capability allows a secondary WAE to be used should the primary WAE fail
  - IP SLAs ensure availability by tracking WAE liveliness
- Seamless integration
  - Transparency and automatic discovery
  - Supported on all WAE platforms

Note:

- Lacks load balancing capabilities
- Failover efficiency depends on tracking protocols such as CDP, IP SLA

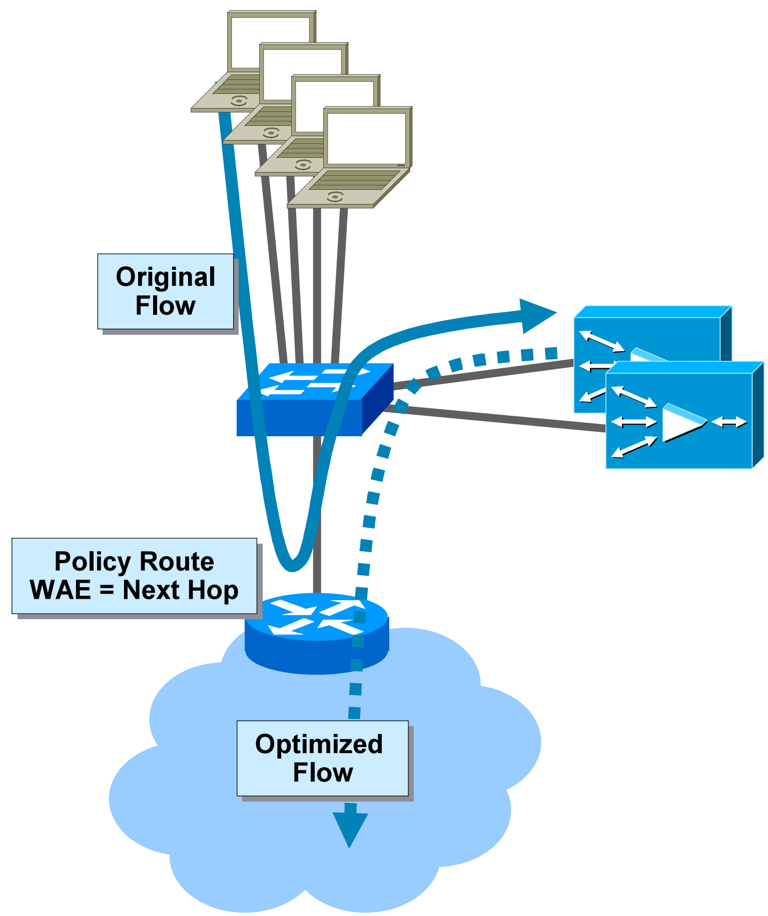

Cisco WAE WCCPv2 Deployment

- WCCPv2 interception
  - Out-of-path with redirection of flows to be optimized (all flows or selective via redirect-list)
  - Automatic load-balancing, load redistribution, fail-over, and fail-through operation
- Scalability and high availability
  - Up to 32 WAEs within a service group and up to 32 routers
  - Linear performance and scalability increase as devices are added
- Seamless integration
  - Transparency and automatic discovery
  - Supported on all WAE platforms

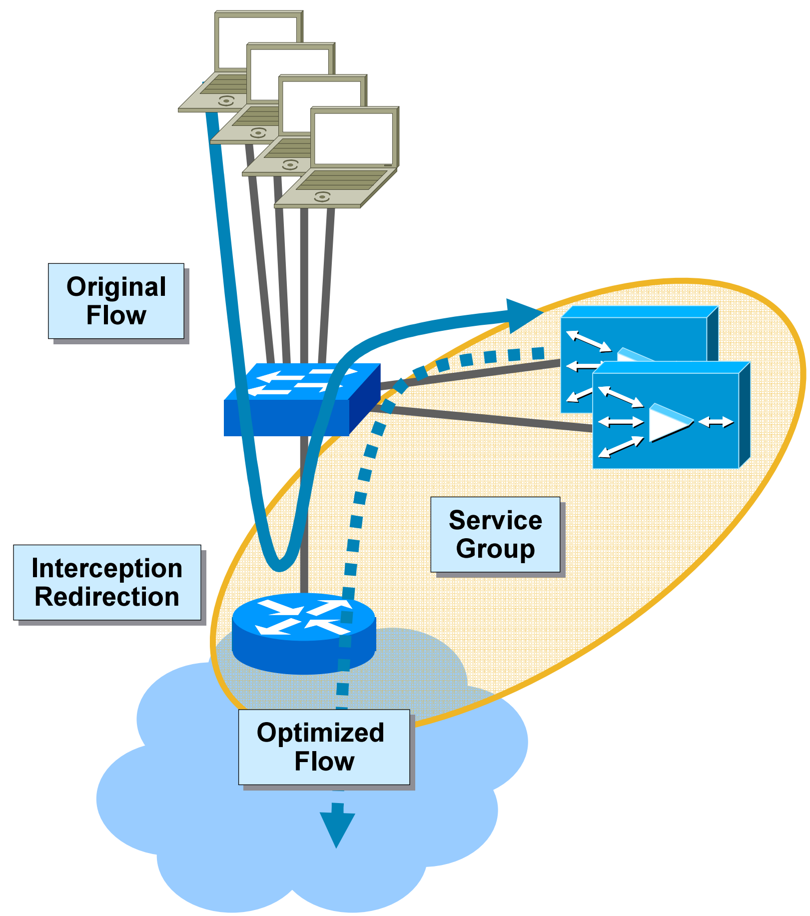

WCCPv2 Overview

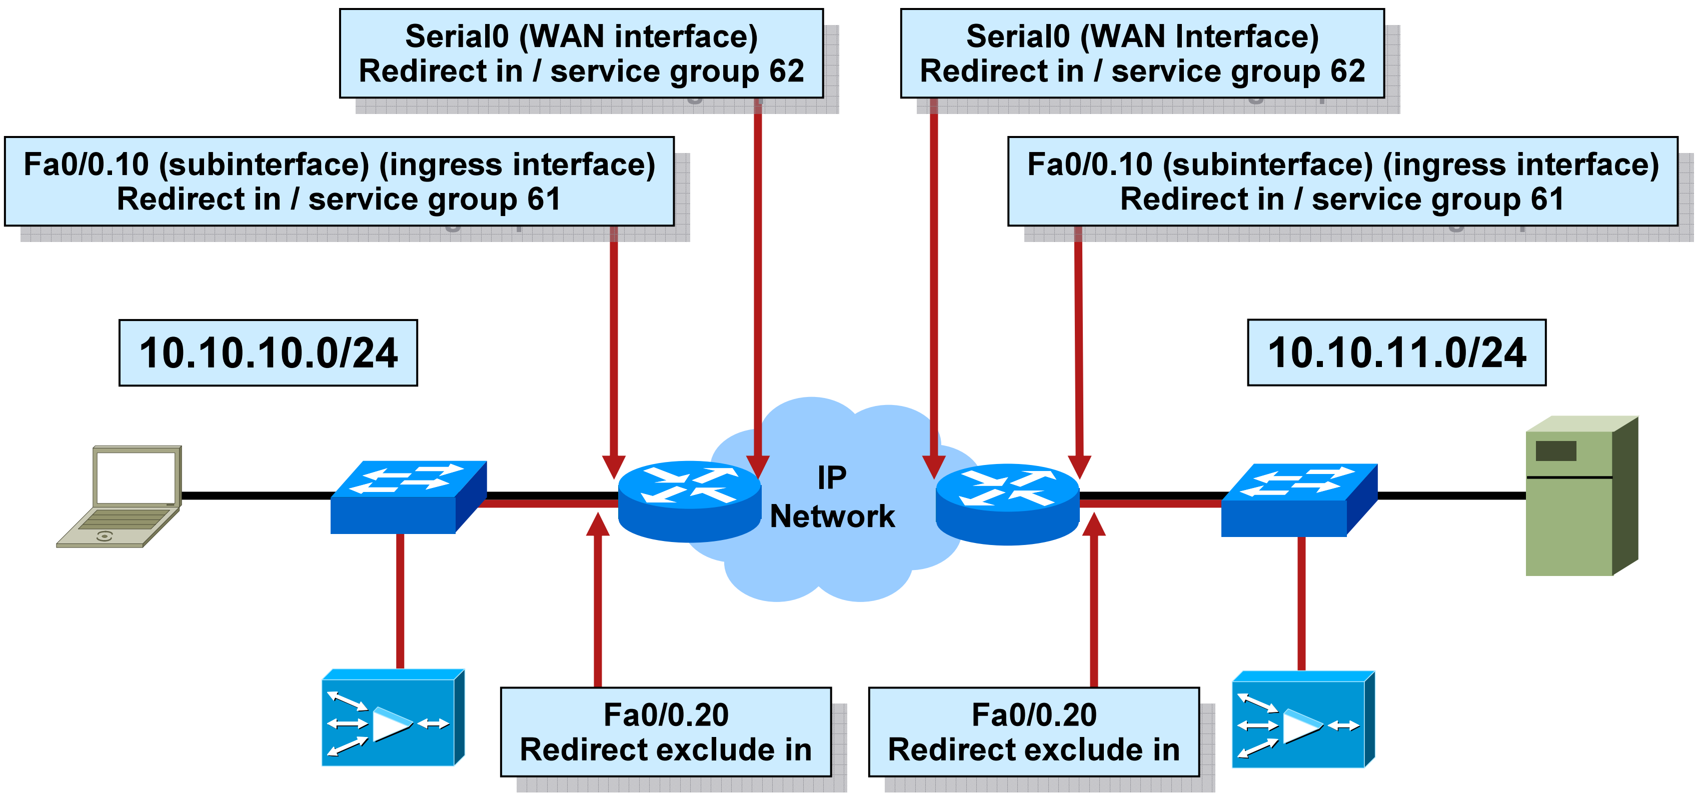

Inline Interception Deployment Modes

In-path, single WAE, redundant WAN links

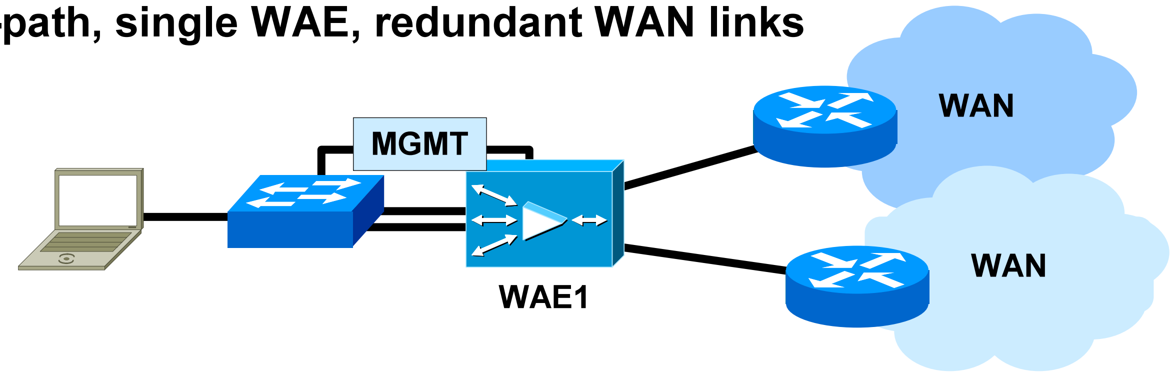

In-path cluster, redundant WAN links

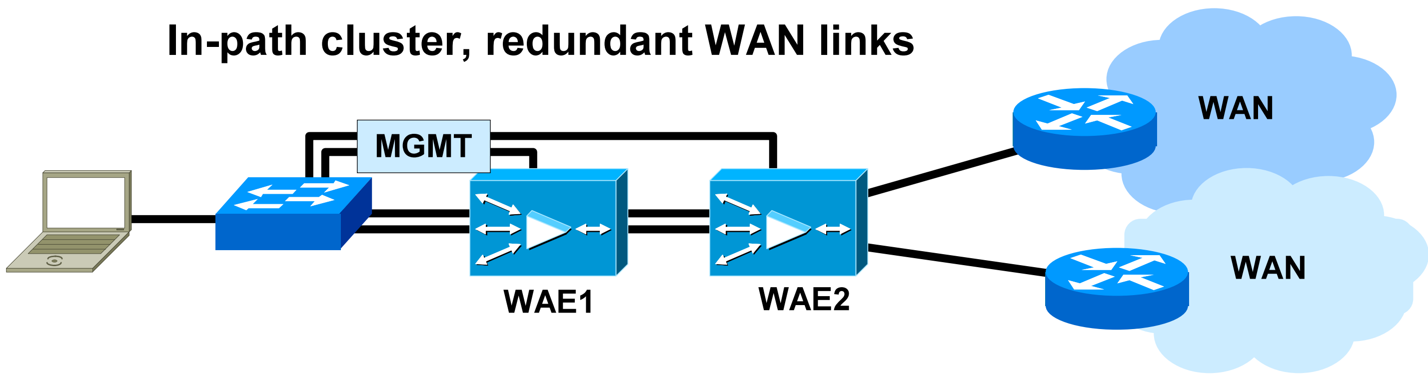

WAAS Redirection with ACE

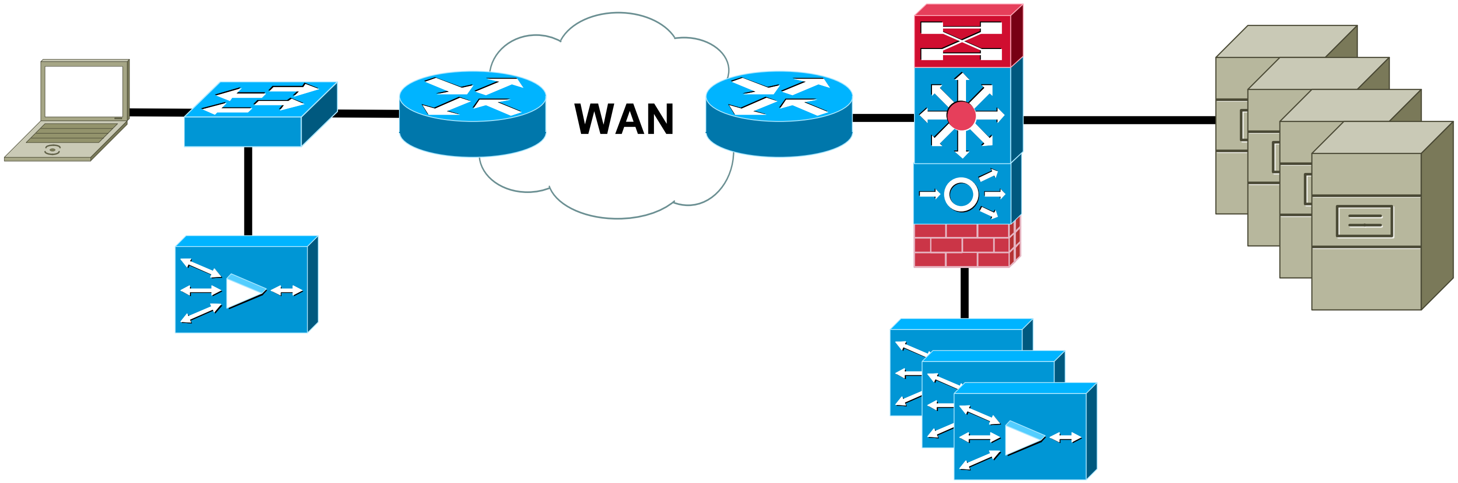

Which Interception Method to Use?

	WCCPv2	Inline	CSM/ACE	PBR
Number of Active WAEs	32	2 (serial cluster, tested limit)	16000 (not practical but possible)	1
Maximum Number of WAEs	32	2 (serial cluster, tested limit)	16000 (not practical but possible)	8 (IOS dependent)
Maximum Number of TCP Connections **(with WAE-7326)**	240K	15K	4M	7.5K
Maximum Throughput	Up to 32Gbps (platform dependent)	Up to 2Gbps (two inline pairs)	Up to 16Gbps (platform dependent)	Up to 1Gbps
Recommended Use	Generally Recommended	Only if WCCPv2 can not be used (SP managed or low-end router)	Very large scale data center deployments	Last resort

Cisco WAE Family Performance and Scalability

Platform	Mem (GB)	Max Drives	Drive Capacity/ Max Capacity	Max Optimized TCP Conns	Max Edge CIFS Sessions	WAN Link Capacity (Mbps)	Max Optimized Throughput (Mbps)	CM Scalability (Devices Managed)	Core Fan-out (Number of Peers)
Current Generation Platforms									
WAE-512-1GB	1	2	250/250	750	750	8	100	500	5
WAE-512-2GB	2	2	250/250	1500	1500	20	150	1000	10
WAE-612-2GB	2	2	300/300	2000	2000	45	250	2000	20
WAE-612-4GB	4	2	300/300	6000	2500	155	350	2500	30
WAE-7326	4	6	300/900	7500	2500	310	450	n/a	50

WAAS Product Portfolio

Cisco WAE Hardware Positioning

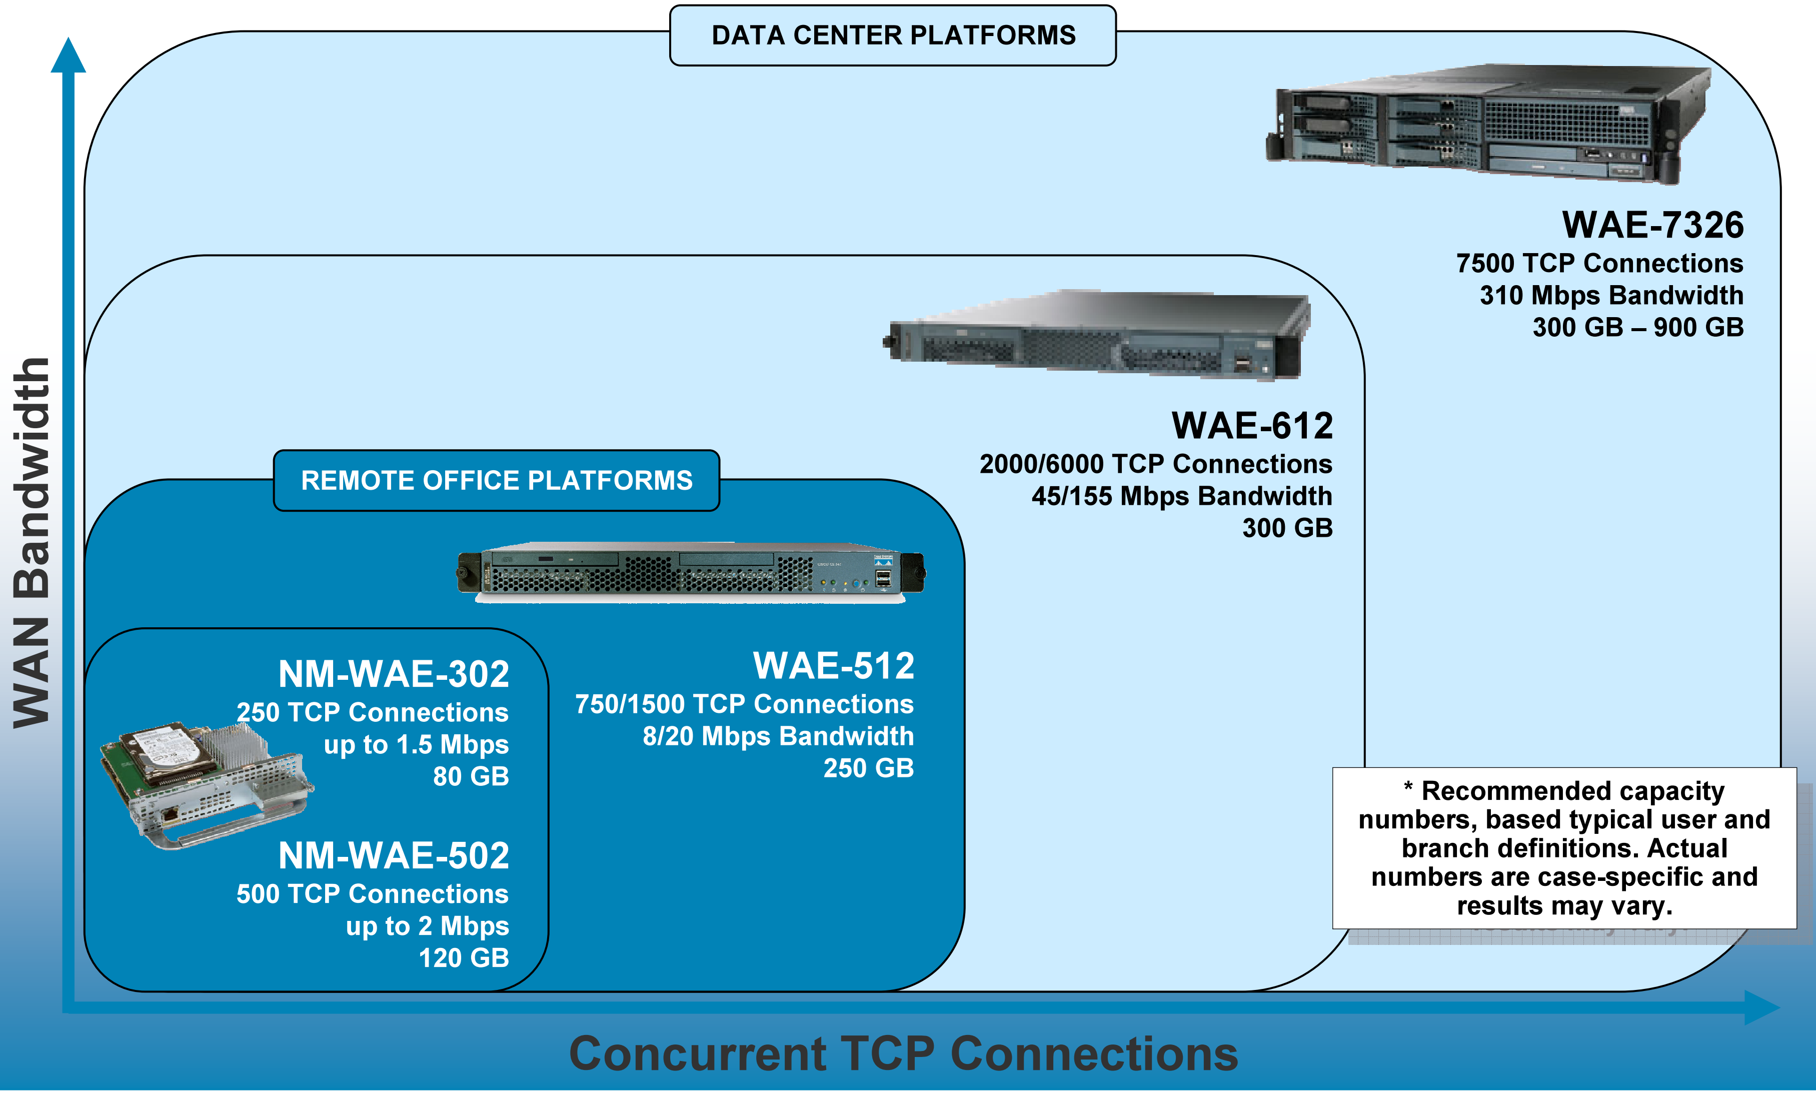

Cisco WAE Physical Inline Deployment

- Physical inline interception:
  - Physical in-path deployment between switch and router or firewall
  - Mechanical fail-to-wire upon hardware, software, or power failure
  - Requires no router configuration
- Scalability and high availability:
  - Two two-port groups
  - Serial clustering with load-sharing and fail-over
  - Redundant network paths and asymmetric routing
- Seamless integration:
  - Transparency and automatic discovery
  - 802.1q support, configurable VLANs
  - Supported on all WAE appliances

Cisco WAE 4-port inline card

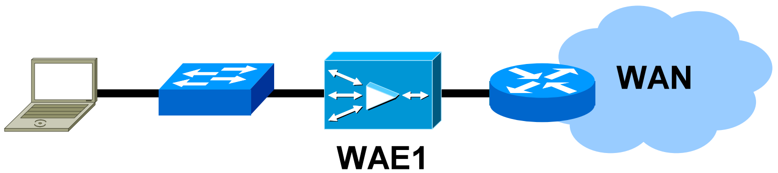

Summary

- "Chatty" applications can severely affect latency
- Standard transport protocols such as TCP are prone to inefficiencies and can be highly optimized
- WAAS makes use of multiple technologies to provide optimization to applications and transport protocols
- There are multiple interception methods to redirect traffic to the WAE's
- Consideration is needed for WAAS placement in the Data Center

Datacenter Server Load Balancing

Application Optimization Infrastructure

Network Classification
- Quality of service
- Network-based app recognition
- Queuing, policing, shaping
- Visibility, monitoring, control

Application Scalability
- Server load-balancing
- Site selection
- SSL termination and offload
- Video delivery

Application Networking
- Message transformation
- Protocol transformation
- Message-based security
- Application visibility

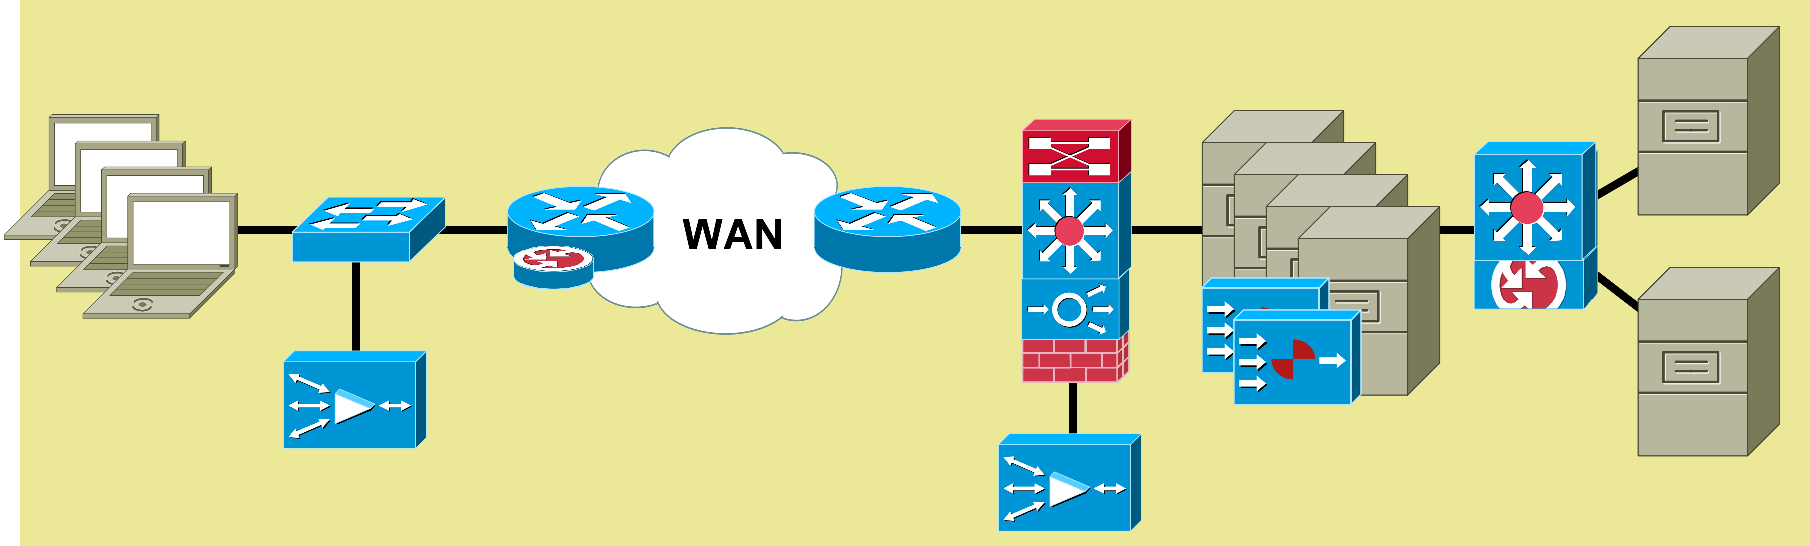

Application Acceleration
- Latency mitigation
- Application data cache
- Meta data cache
- Local services

WAN Acceleration
- Data redundancy elimination
- Window scaling
- LZ compression
- Adaptive congestion avoidance

Application Optimization
- Delta encoding
- FlashForward optimization
- Application security
- Server offload

Agenda

- Introduction
- Hardware
- Modular Policy CLI and Role-Based Access Control
- Virtual Partitioning
- Application Delivery and Security Features
- Redundancy
- Deployments
- Introduction to ACE Version 2.0

Application Load Balancing
Content Switching Requirements

- High-availability
- No single point of failure
- Disaster recovery
- High and scalable performance
- Intelligent content-based decisions
- Transaction assurance
- Security

The Evolution of L4 to 7 Services

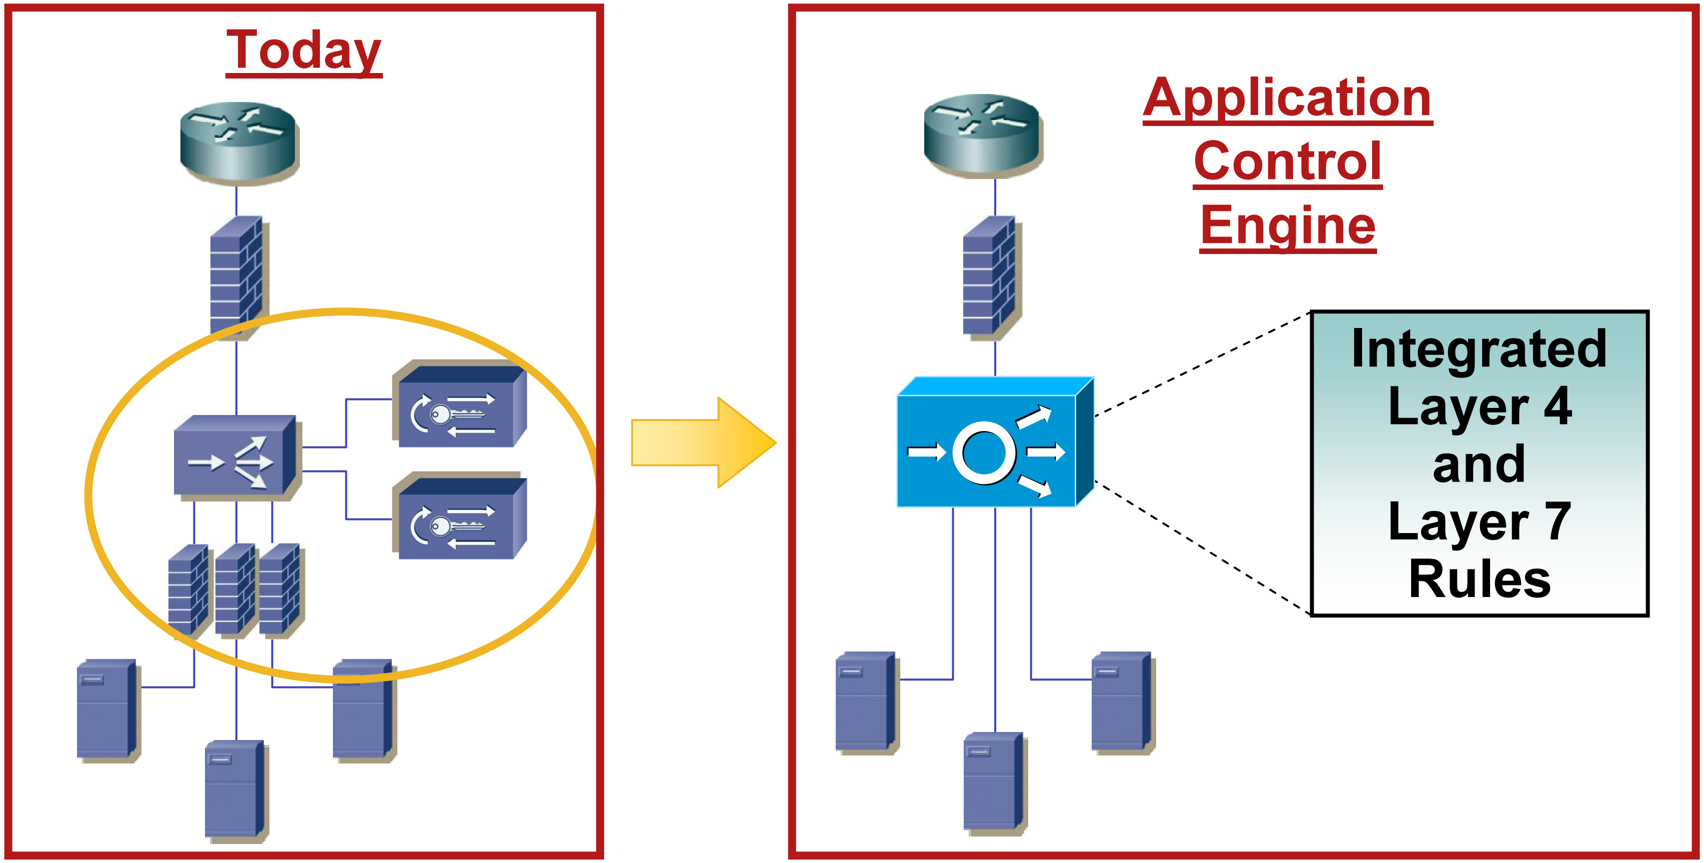

- Infrastructure simplification with L4–7 Services integration
- Converged policy creation, management, and troubleshooting
- Reduced latency (single TCP termination for all functions)

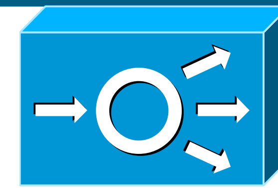

What Is ACE ?

Application Control Engine

- Brand new product line in the Cisco ANS portfolio
- Infrastructure simplicity in a single hardware platform, ACE integrates

    Content switching

    SSL offload

    Data center security features

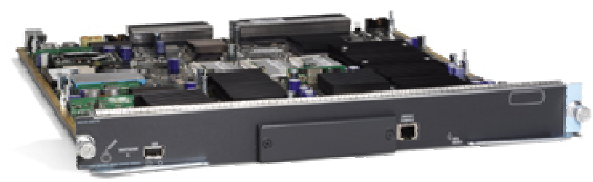

- The first ACE product is a Cisco Catalyst® 6500 service module, which comes in three flavours: 4Gbps, 8Gbps, and 16Gbps
- The hardware supports two field-replaceable daughtercards for future hardware-accelerated application delivery functionality like HTTP compression
- It delivers application infrastructure control, with features like virtual partitions and native role based access control (RBAC)

The Application Control Engine At-a-Glance

Application Infrastructure Control

- Virtual Partitioning
- Hierarchical Management Domains
- Role-Based Access Control

Application Performance

- High Throughput (16Gbps)
- Maximum Scalability (350K CPS)
- Multi-tiered reliability, availability, and scalability
- Server Load Balancing
- Content Switching (L7 decisions and advanced stickiness)

Application Security

- Protocol-layer inspection
- TCP/IP Normalization
- Hardware-accelerated Protocol Control
- High Performance NAT (1M xlates)
- Access Control List (ACL) (up to 256K ACEs)
- DDoS Protection

Infrastructure Simplification

- Layer 2–7 Network Integration
- Functional Consolidation
- Application Network Management solution
- TCP Offload
- SSL Termination
- XML API

Hardware

Application Control Engine

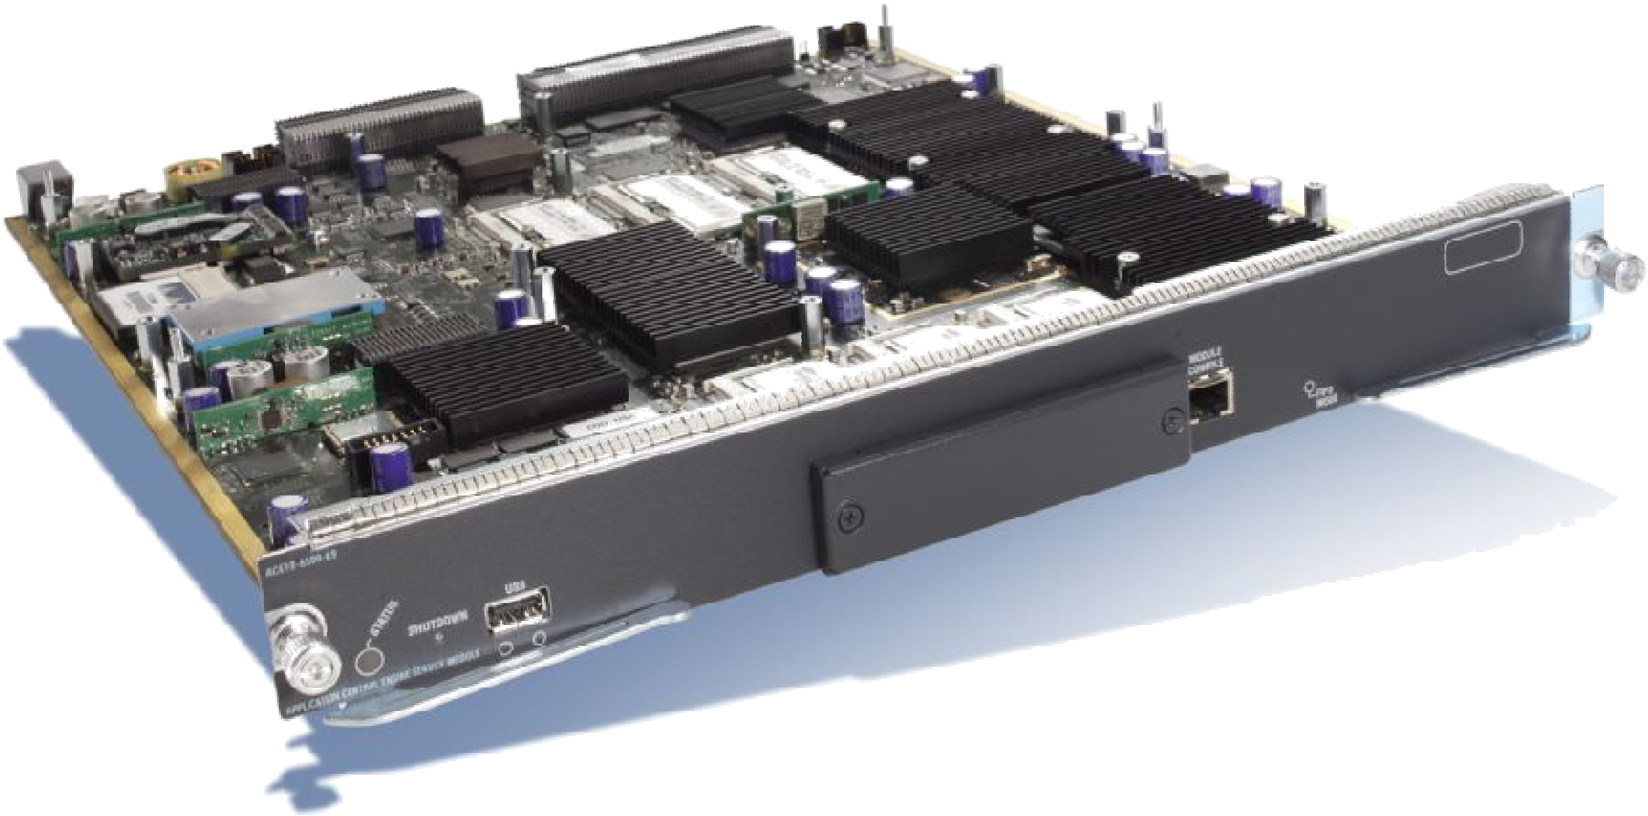

Parallel network-processor based hardware with separate control and data-path CPUs

ACE—Hardware Architecture

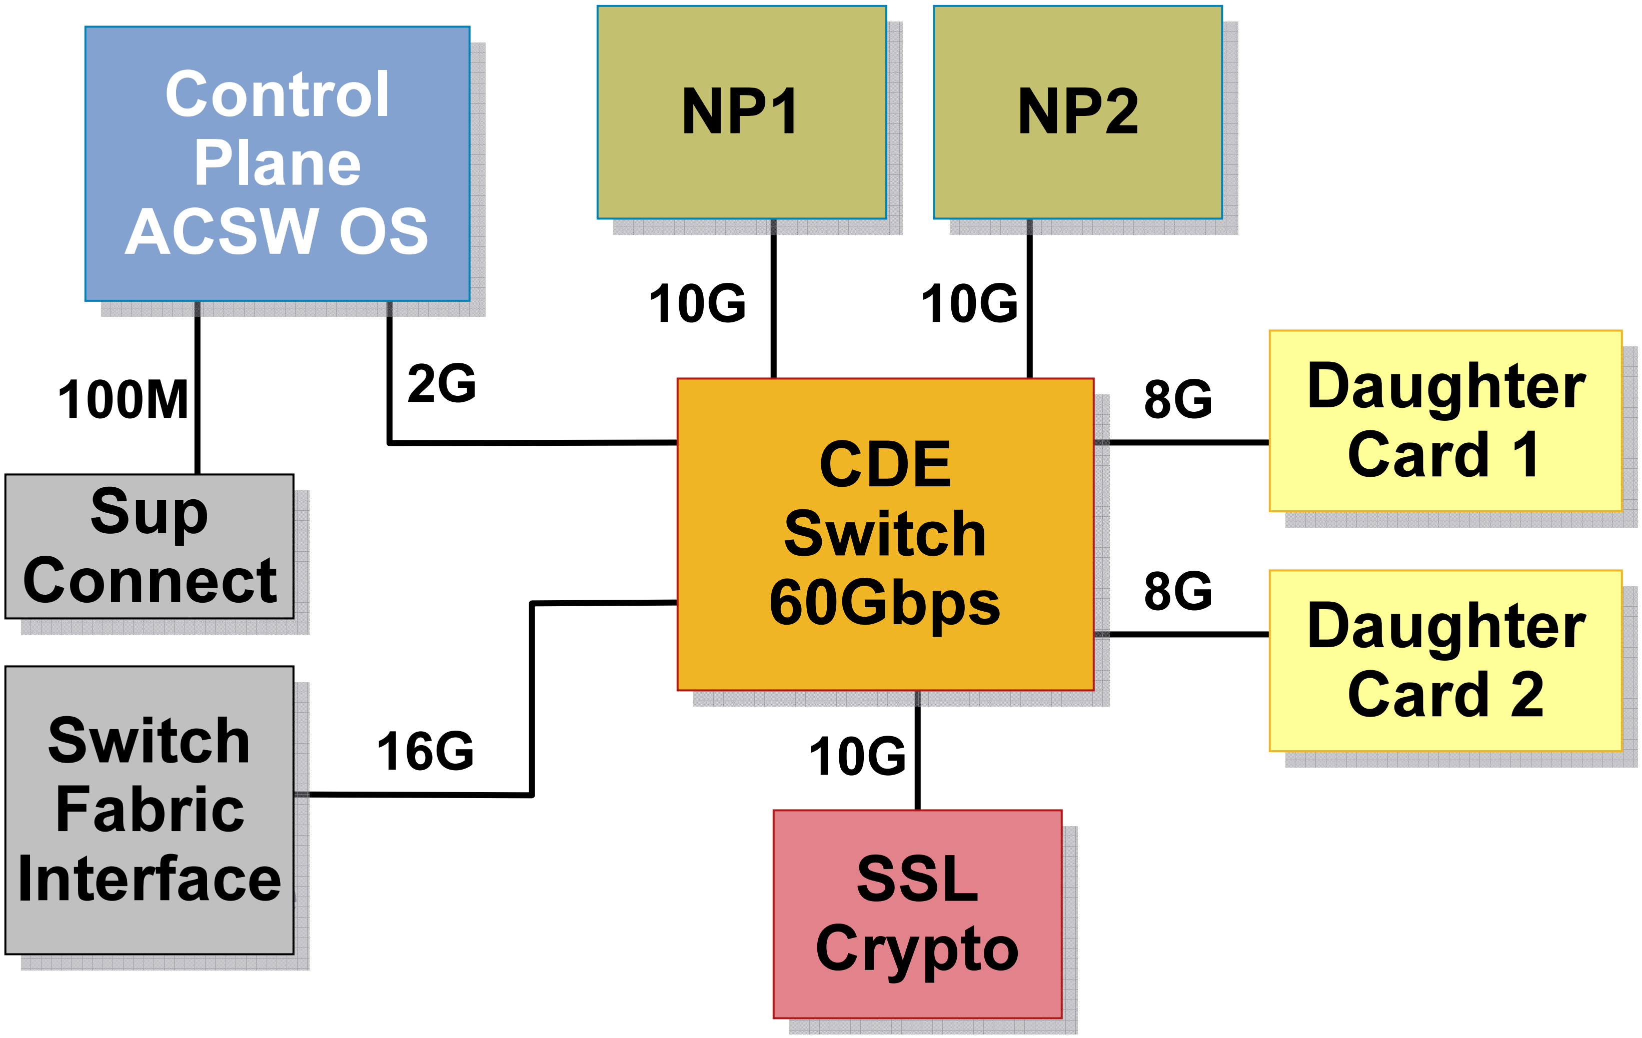

ACE—Detailed Hardware Architecture

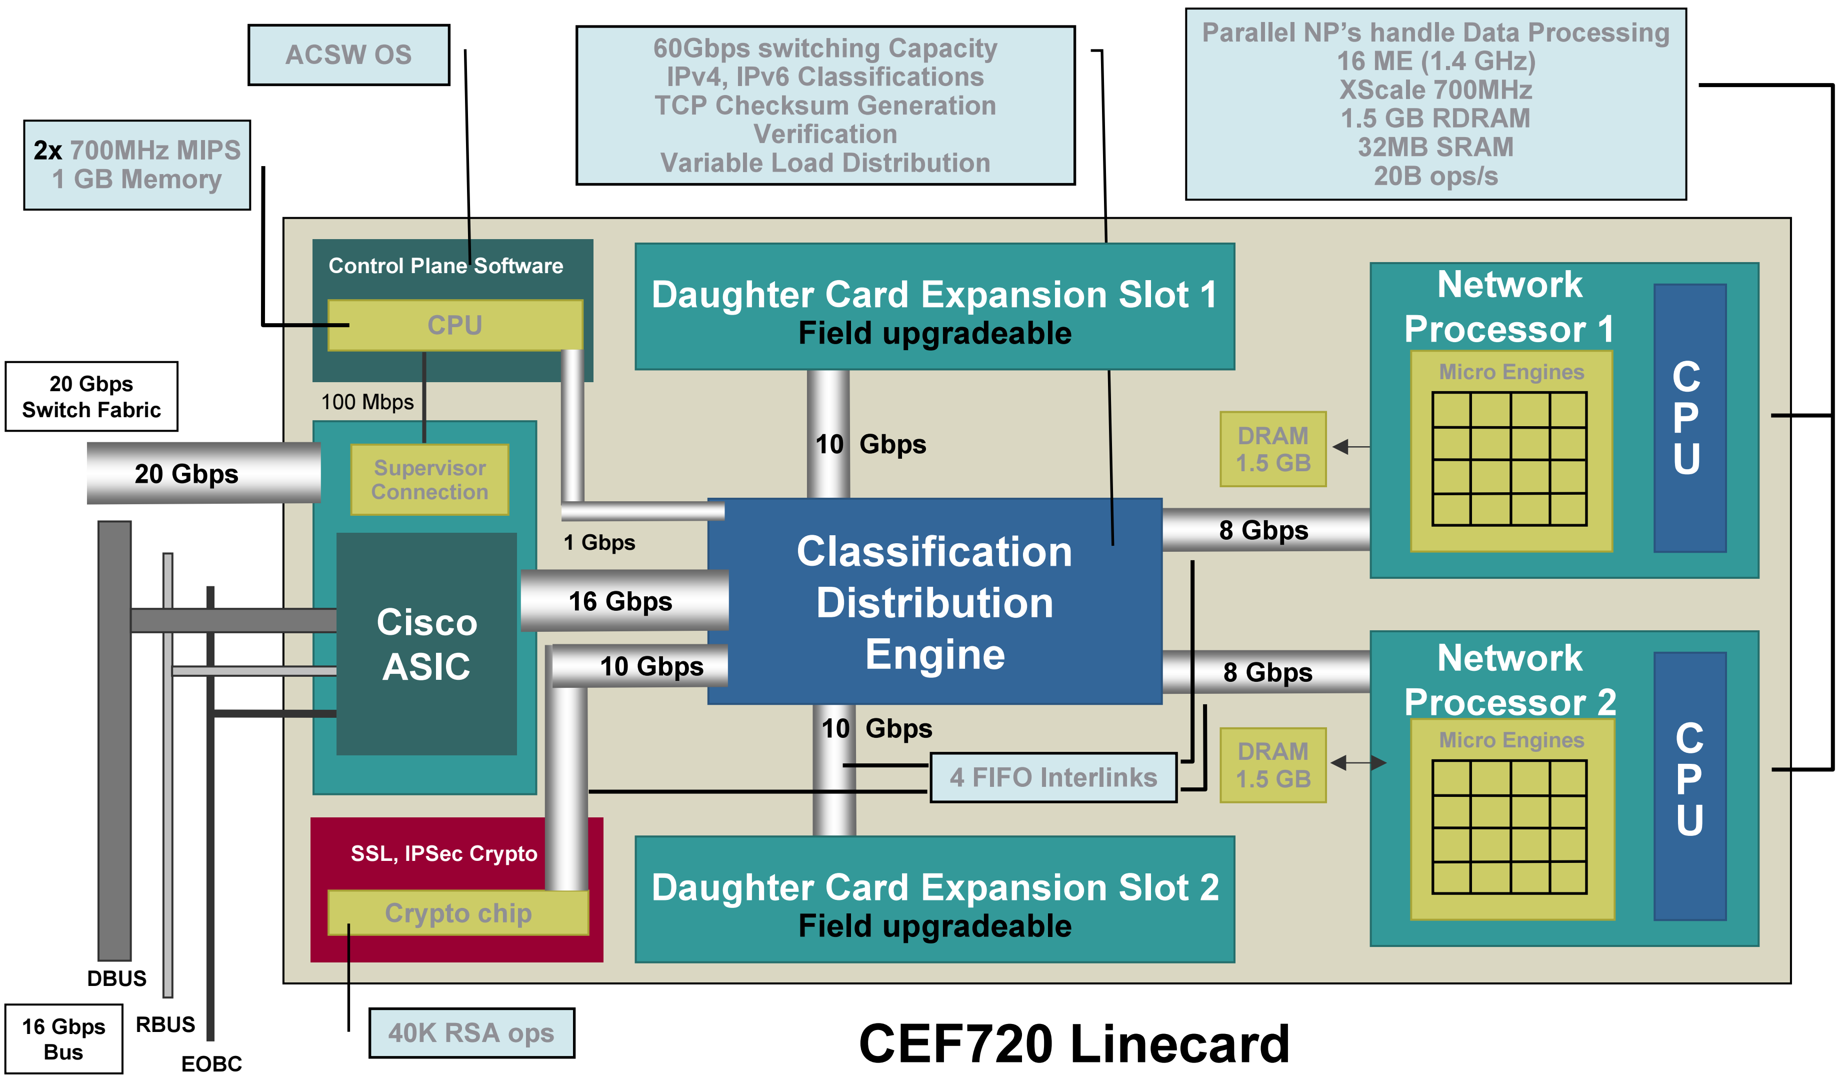

Data-Path / Control-Path Separation

- Control-Path
  - Device control
  - Configuration manager (CLI, XML API, SSH, …)
  - Server health monitoring (native probes, TCL scripts)
  - SYSLOGs, SNMP, …
  - ARP, DHCP relay
  - High-Availability

- Data-Path
  - Connection management
  - TCP termination
  - Access lists
  - SSL offload
  - Regular expression matching
  - Load Balancing & forwarding

Control path and data path run on separate processors

Modular Policy CLI

Modular Policy CLI in ACE

- ACE CLI is based on C3PL (Cisco Common Class-based Policy Language)
- Provides a common CLI framework across security implementations in-order to define consistent CLI across platforms
- The CLI aims at seamless integration in terms of configuring SLB, SSL and Security features
- No need to session in or enter a sub-mode of configuration for the different features
- Traffic classification is the core functionality for all delivery and security features

Policy CLI Overview

1. Define match criteria
2. Associate actions to match criteria
3. Activate the classification-action rules on either an interface or “globally”

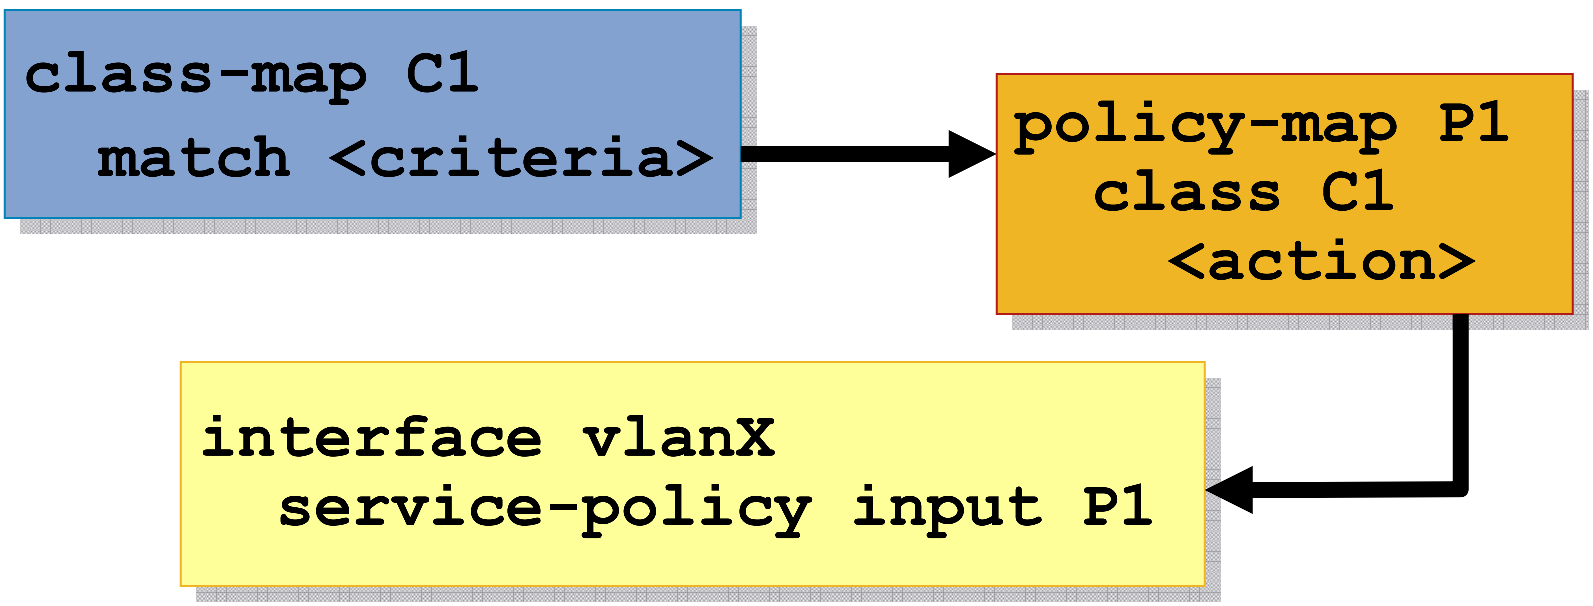

Policy Lookup Order

- There can be many features applied on a given interface, so feature lookup ordering is important
- The feature lookup order followed by datapath in ACE is as follows:
  1. Access-control (permit or deny a packet)
  2. Management traffic
  3. TCP normalization/connection parameters
  4. Server load balancing
  5. Fix-ups/application inspection
  6. Source NAT
  7. Destination NAT
- The policy lookup order is implicit, irrespective of the order in which the user configures policies on the interface

Virtual Partitioning

Design Considerations
Architecture Integration Considerations

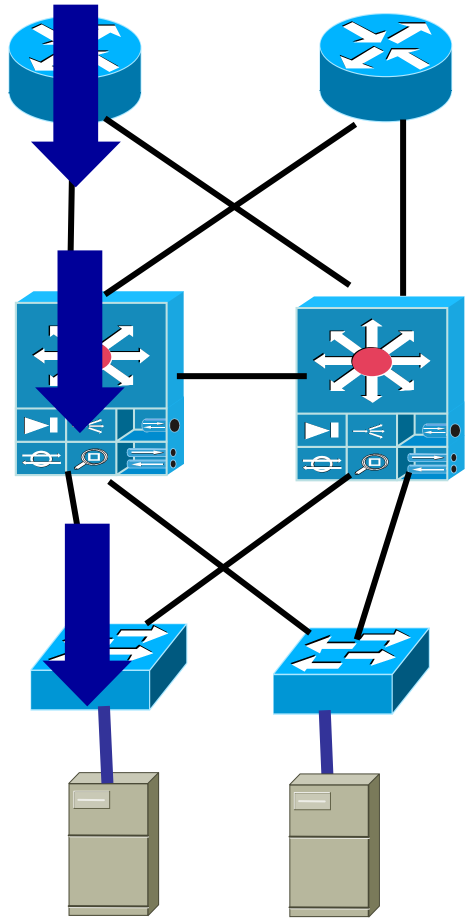

Design Considerations
Architecture Integration Considerations

- L2 sub-optimal path
- L3 sub-optimal path
- Higher Interswitch-link bandwidth
- Unuse of valuable resources

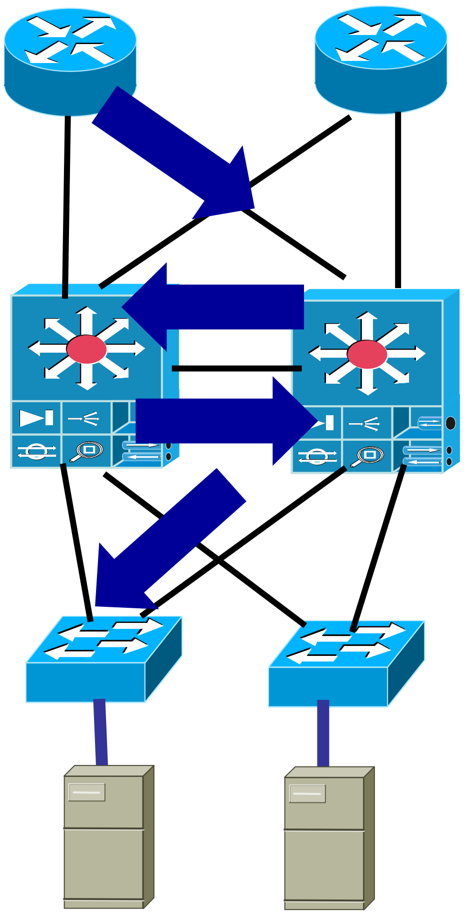

Virtual Partitioning

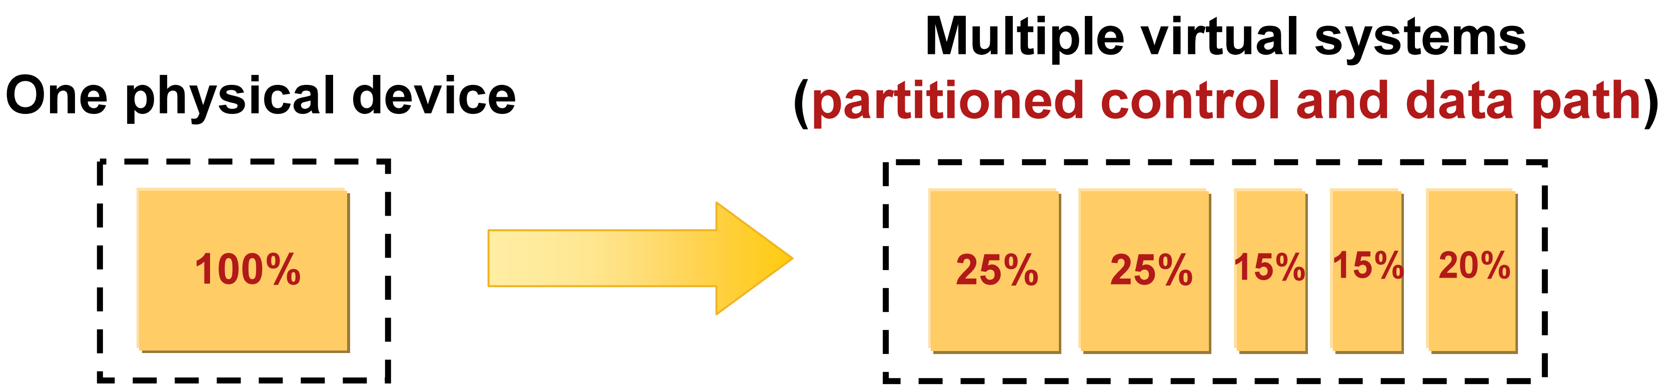

<u>Traditional device</u>

Single configuration file

Single routing table

Limited RBAC

Limited resource allocation

<u>Cisco Application Services Virtualization</u>

Distinct configuration files

Separate routing tables

RBAC with Contexts, Roles, Domains

Management and data resource control

Independent application rule sets

Global administration and monitoring

Virtual Partitioning Resource Control

Per Context Control

- Resource levels for each context
- Support for over-subscription

Rates

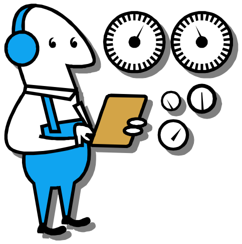

Memory

Rates:

Bandwidth
Data connections/sec
Management connections/sec
Ssl-bandwidth
Syslogs/sec

Memory:

Access Lists
Regular Expressions
Data connections
Management connections
SSL connections
Xlates
Sticky entries

ACE Virtual Partitioning Deployments

Customer Deployments

1. Isolate departments or customers

    Provide direct configuration access

    Reduce exposure to critical config components

    Provide consistent access across GUI, API, CLI

    Dedicated resources

2. Isolate applications

    Guarantee resources to critical applications

    Isolate from impact of other app roll outs

    Central config file for managing policy change

    Reduced complexity of security/application rules

    Possiblity to have a parallel test environment with no impact to production

Design Configuration
ACE Virtualization

- Provides means to partition one physical unit into independently managed logical engines
  - Provisions resource per logical device
  - Almost every feature subsystem is virtualized including Linux kernel
- Logical devices are called virtual contexts
  - Each with independent resource allocation and policies
- Default context called 'Admin' context is available initially
  - Customers who do not wish to use virtualization can perform all operations from within 'Admin' context
- 250 contexts + Admin context supported for phase 1

Design Configuration
ACE Service Virtualization

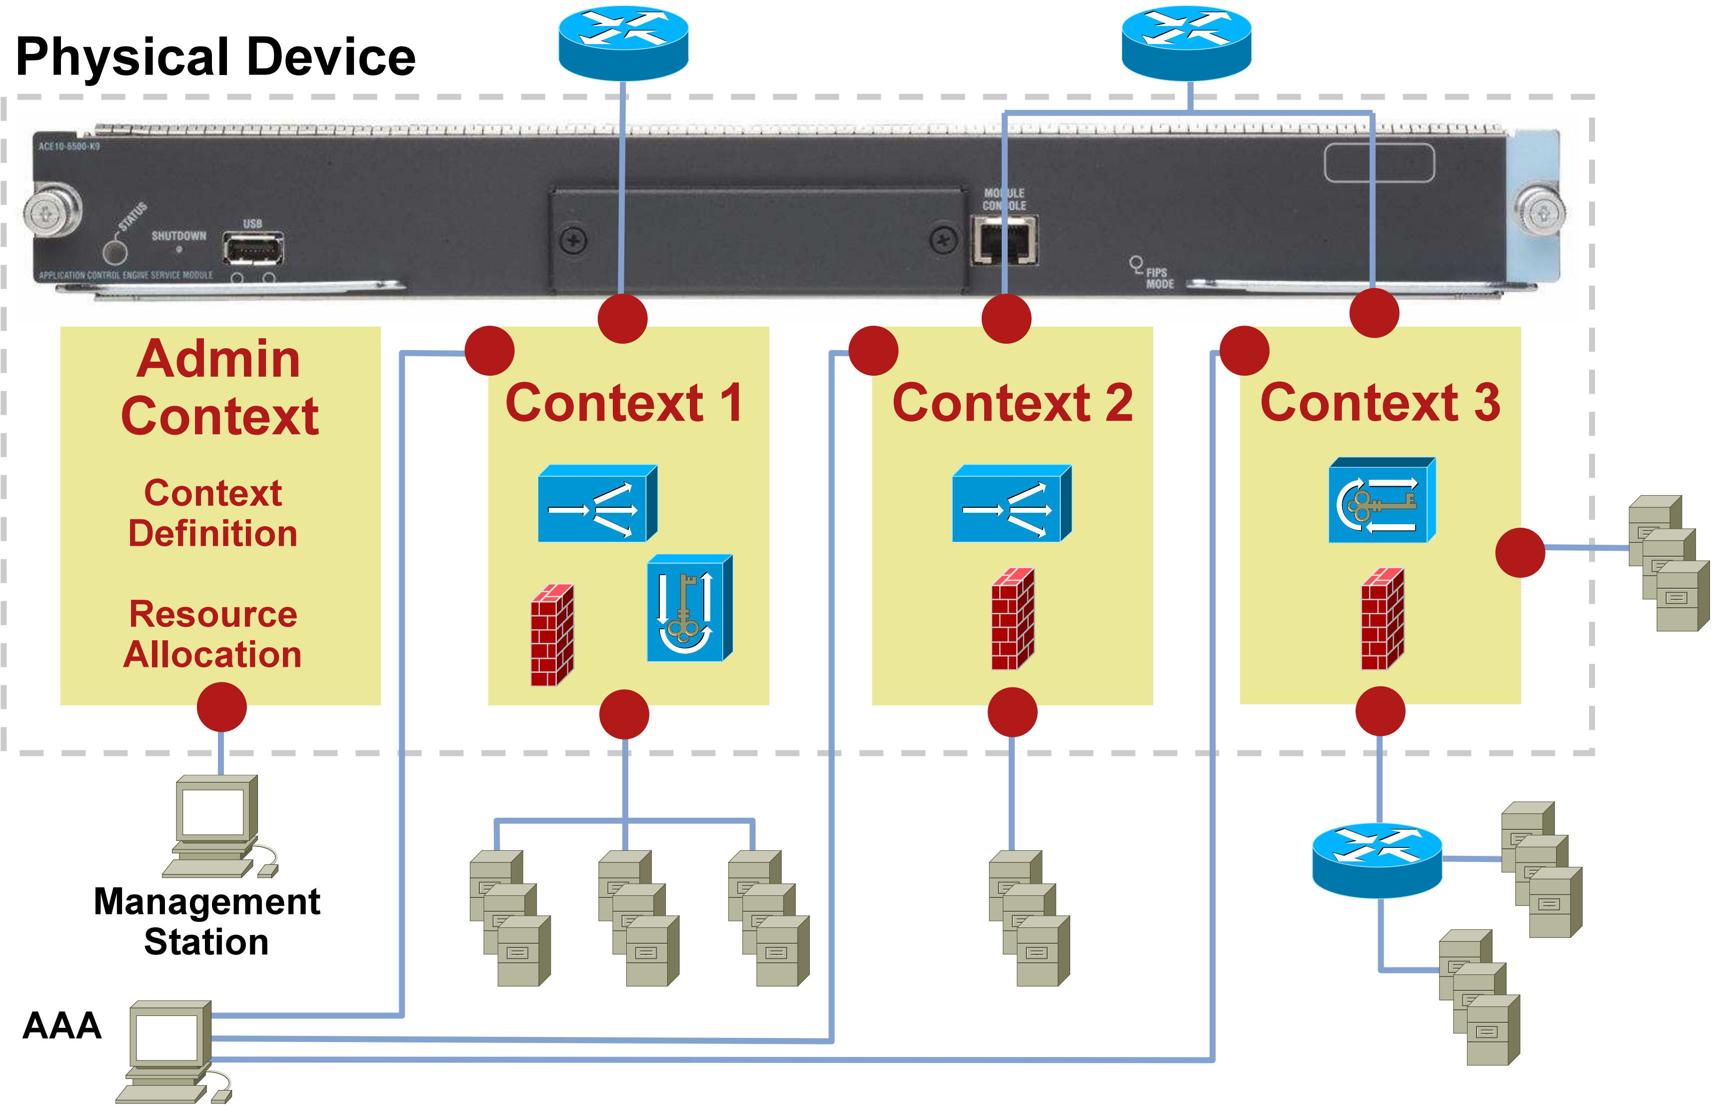

ACE Virtual Patitioning and App Security in Action

Multi-tier Applications

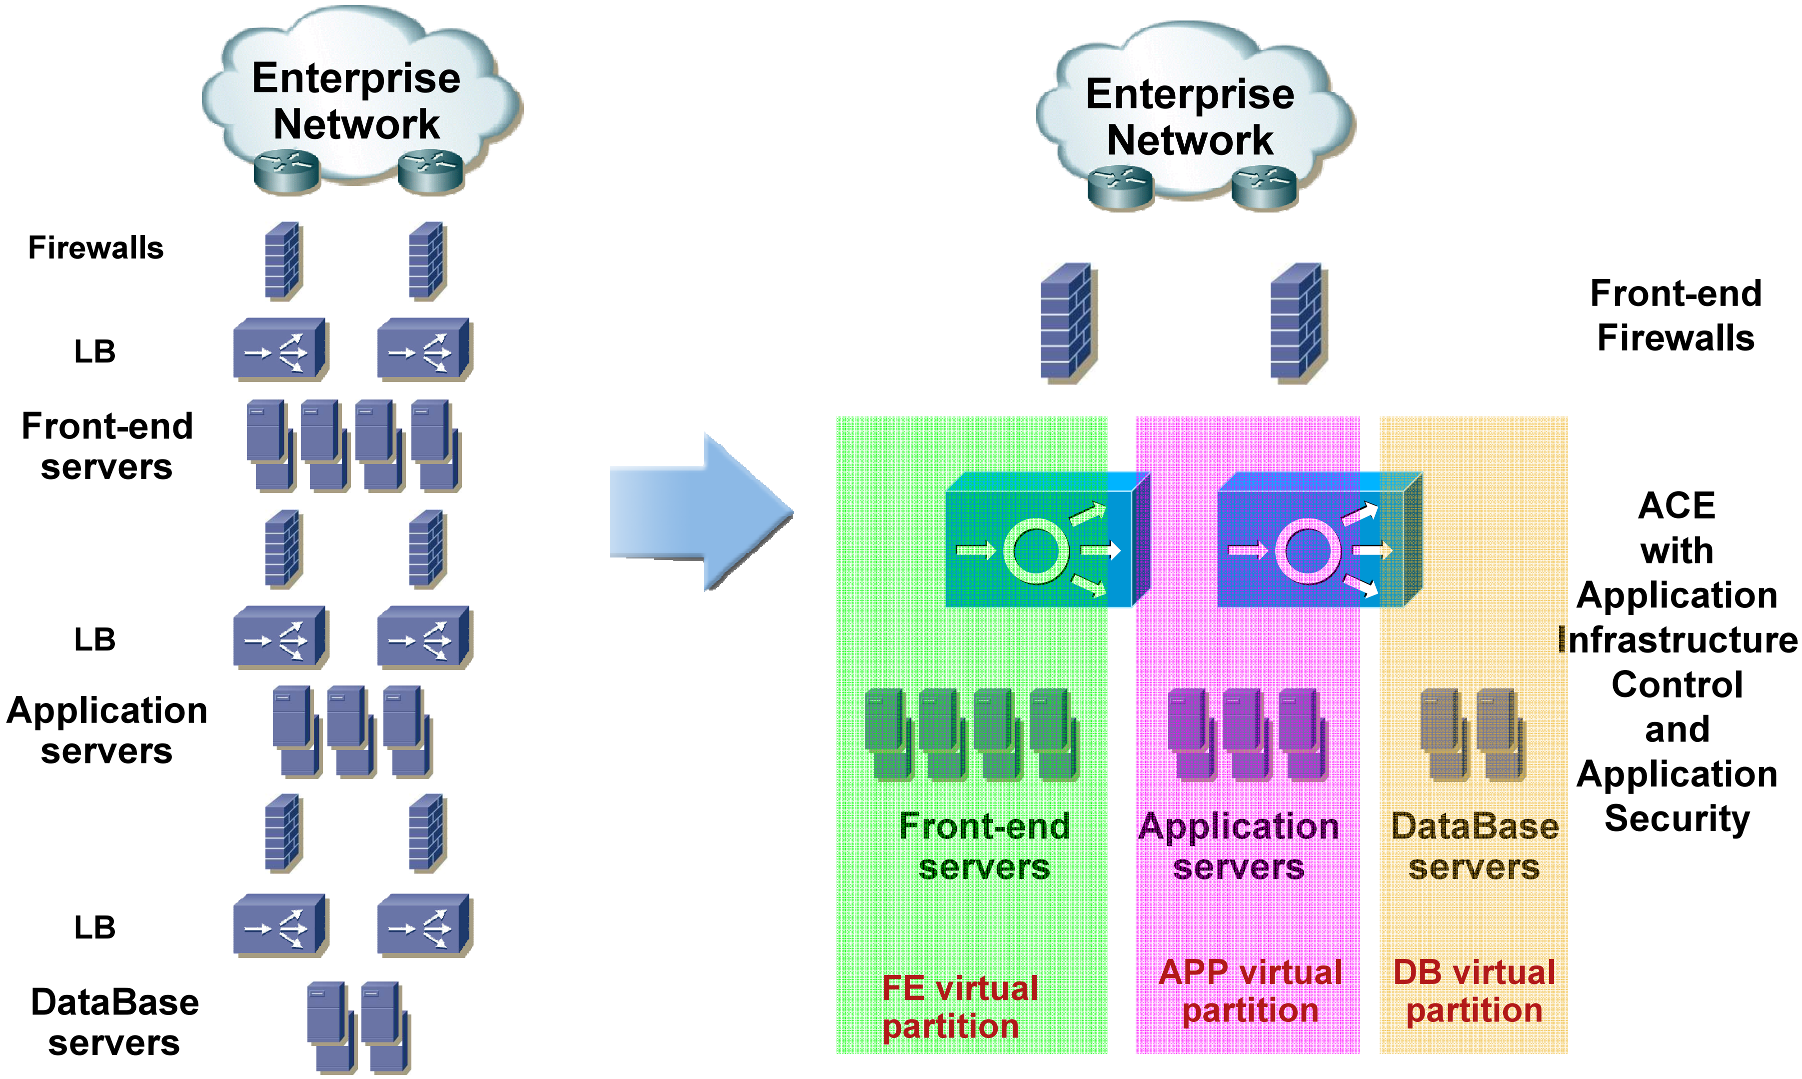

Application Delivery & Security Features

TCP Reuse (Offload)

- Offload TCP (HTTP) setup processing from server
- TCP connections to the server are kept open (HTTP 1.1 Persistence)
- Client requests multiplexed to existing server connections
- TCP Reuse can be enabled on per virtual server basis
- Creates a connection pool on the reals [ip:port] associated to the virtual server
  - Per rserver per serverfarm
  - Client connections matched to server connections based on TCP options
    - Sack
    - timestamp
    - window_scale
    - MSS

TCP Reuse (Offload)

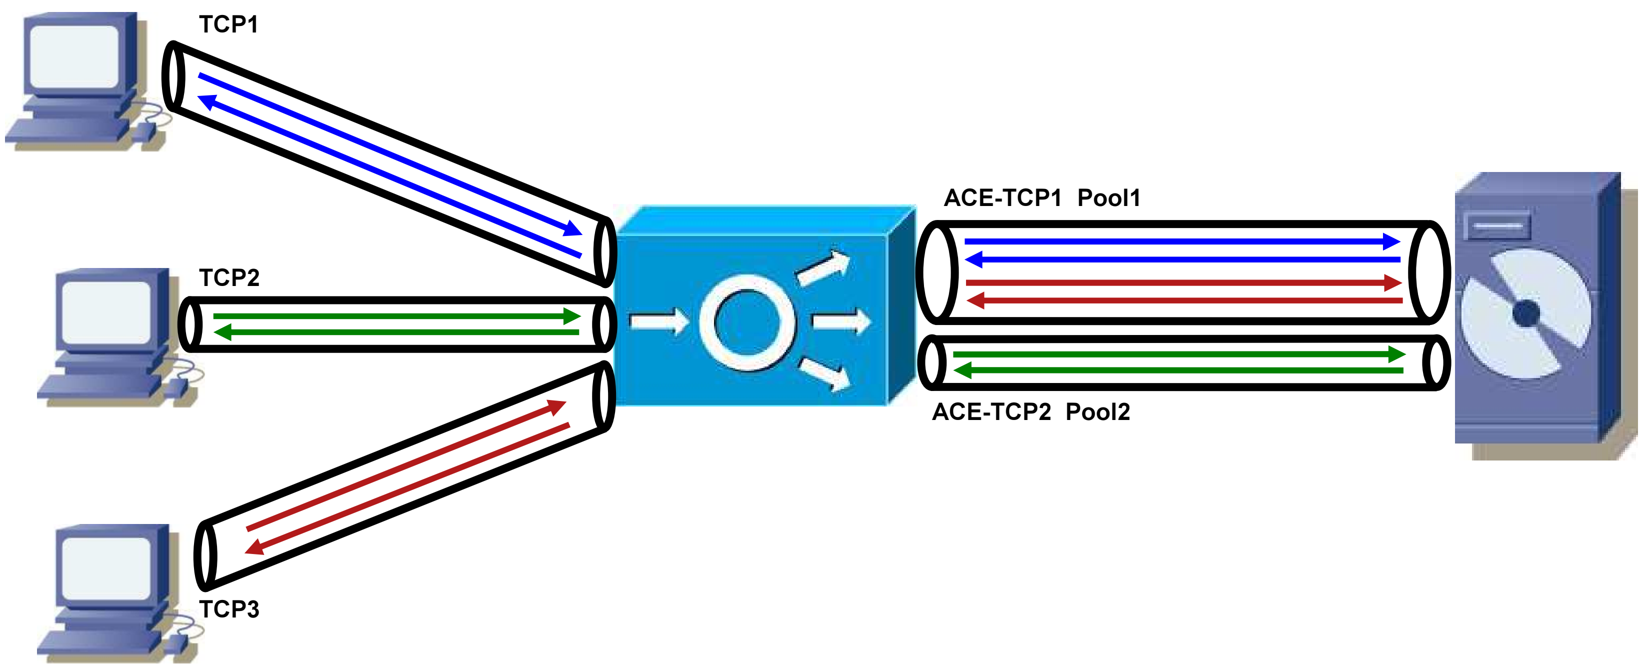

SSL Offloading Combined Solution

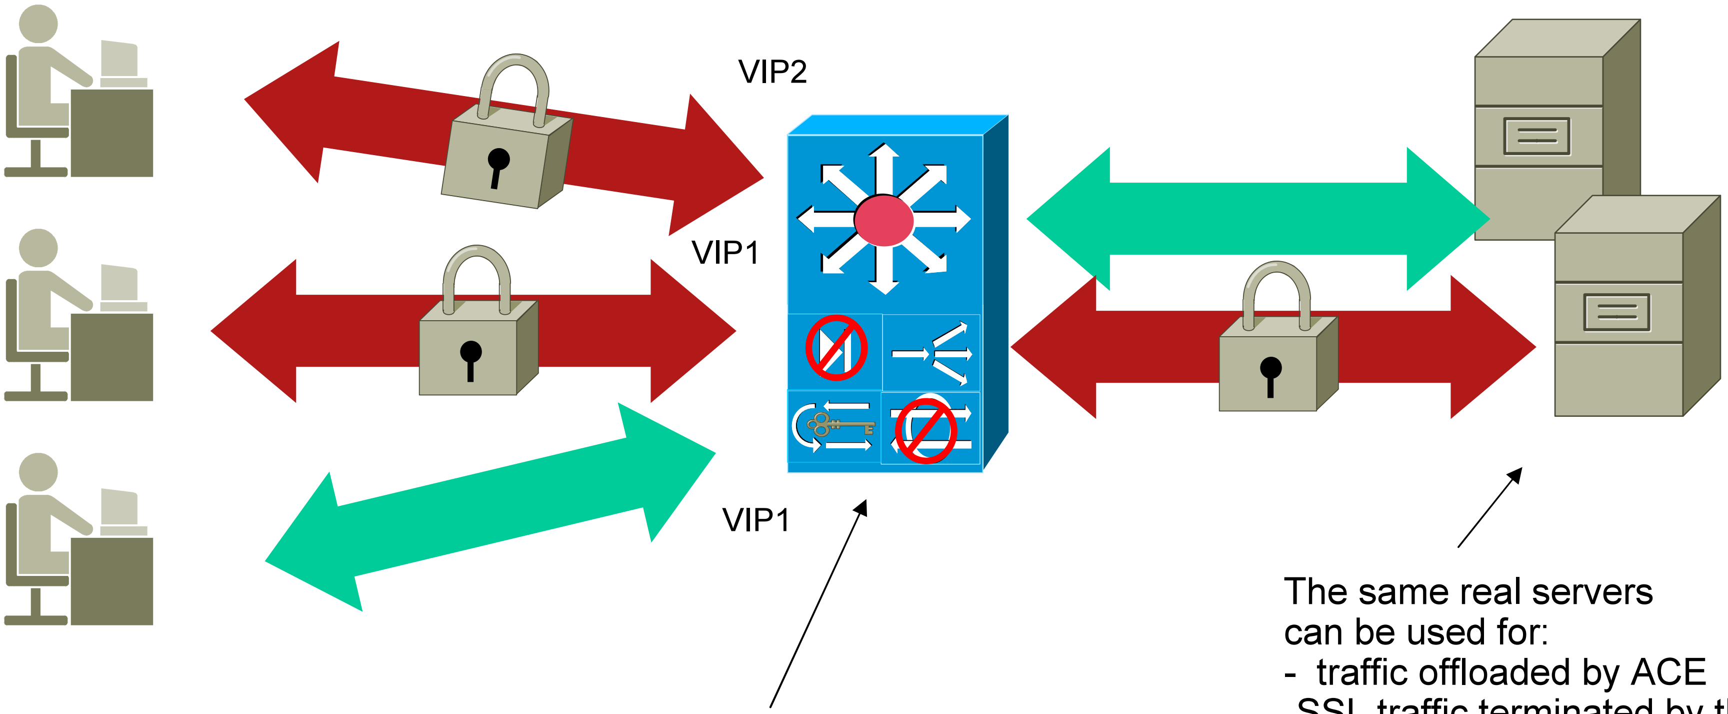

HA Solution:

Stateful failover cannot be achieved for SSL termination but server persistence is guarranted.

The same real servers can be used for:
- traffic offloaded by ACE
- SSL traffic terminated by the server
- Unencrypted traffic

Redundancy

Redundancy Model

- Redundancy groups (Fault Tolerance, FT groups) are configured based on virtual contexts
- Two instances of the same context (on two distinct ACE modules) form a redundancy group, one being active and the other standby
- The peer ACE can be in the same or different Cisco Catalyst 6k chassis
- Both ACE modules can be active at the same time, processing traffic for distinct contexts, and backing-up each other (stateful redundancy)

Example:
Two ACE modules
Four FT groups
Four Virtual Contexts (A,B,C,D)

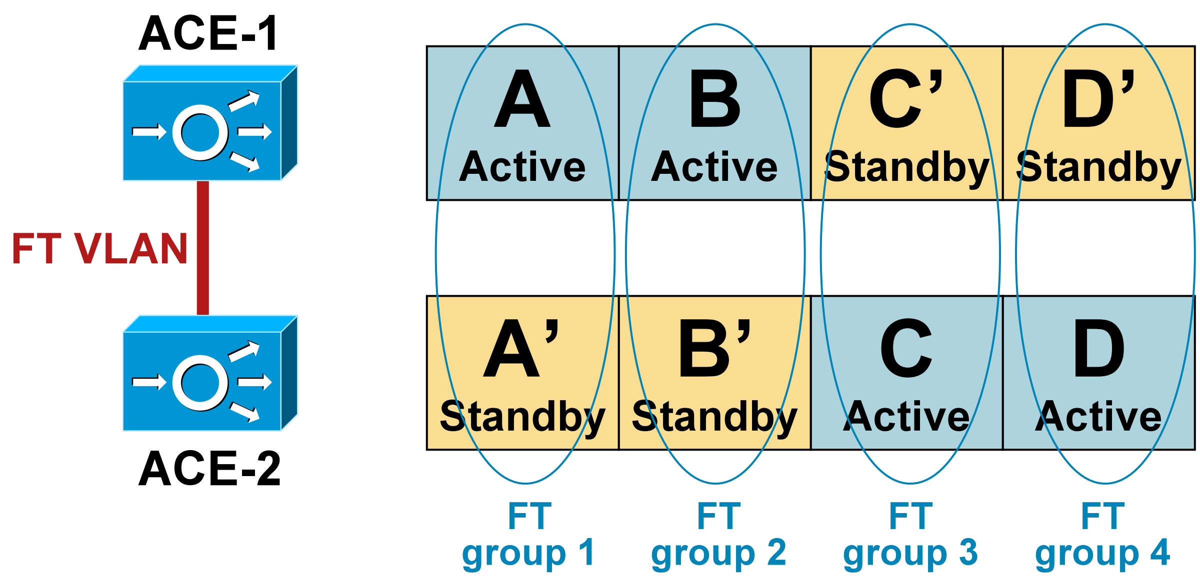

Configuration Sync

1. Bulk sync

   Entire config transfered from Active to Standby

   State during sync: ACTIVE/STANDBY_CONFIG

2. Incremental Sync

   Line-by-line config sync while Active is being configured

   State during sync: ACTIVE/STANDBY_HOT

- The user can choose to disable config sync through CLI; State would remain ACTIVE/STANDBY_HOT with no config sync
- If user re-enables config sync, HA would trigger a bulk sync to make sure both devices are again statefully in sync

Deployment

Typical Data Center Design with ACE

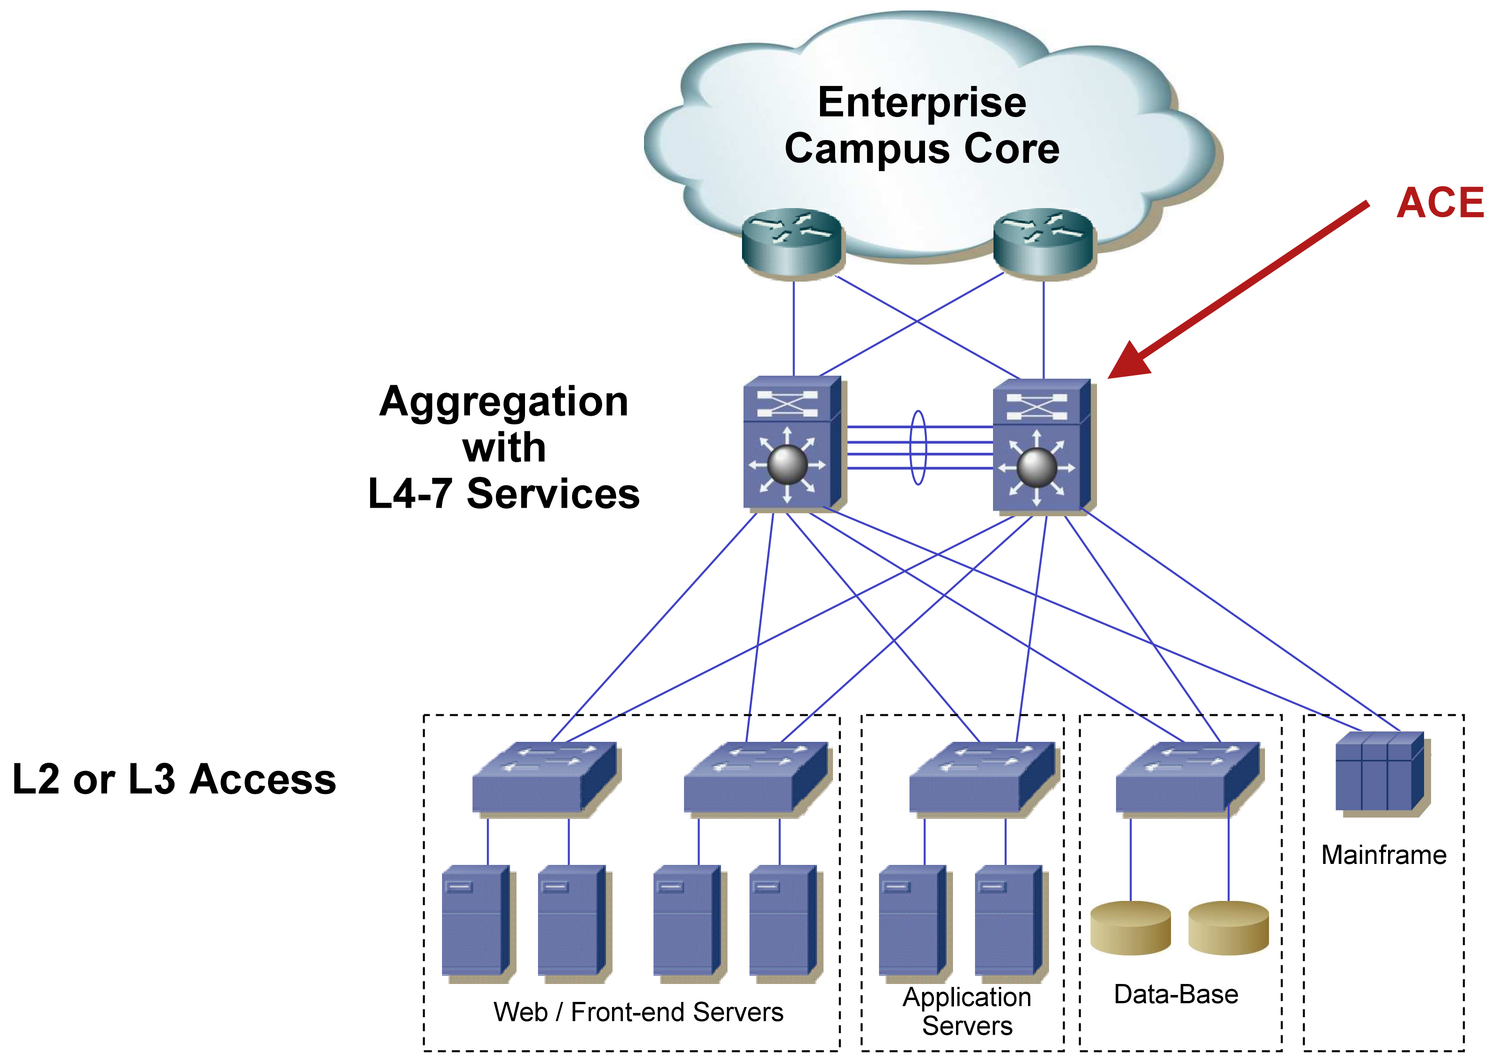

ACE Deployed in Router Mode

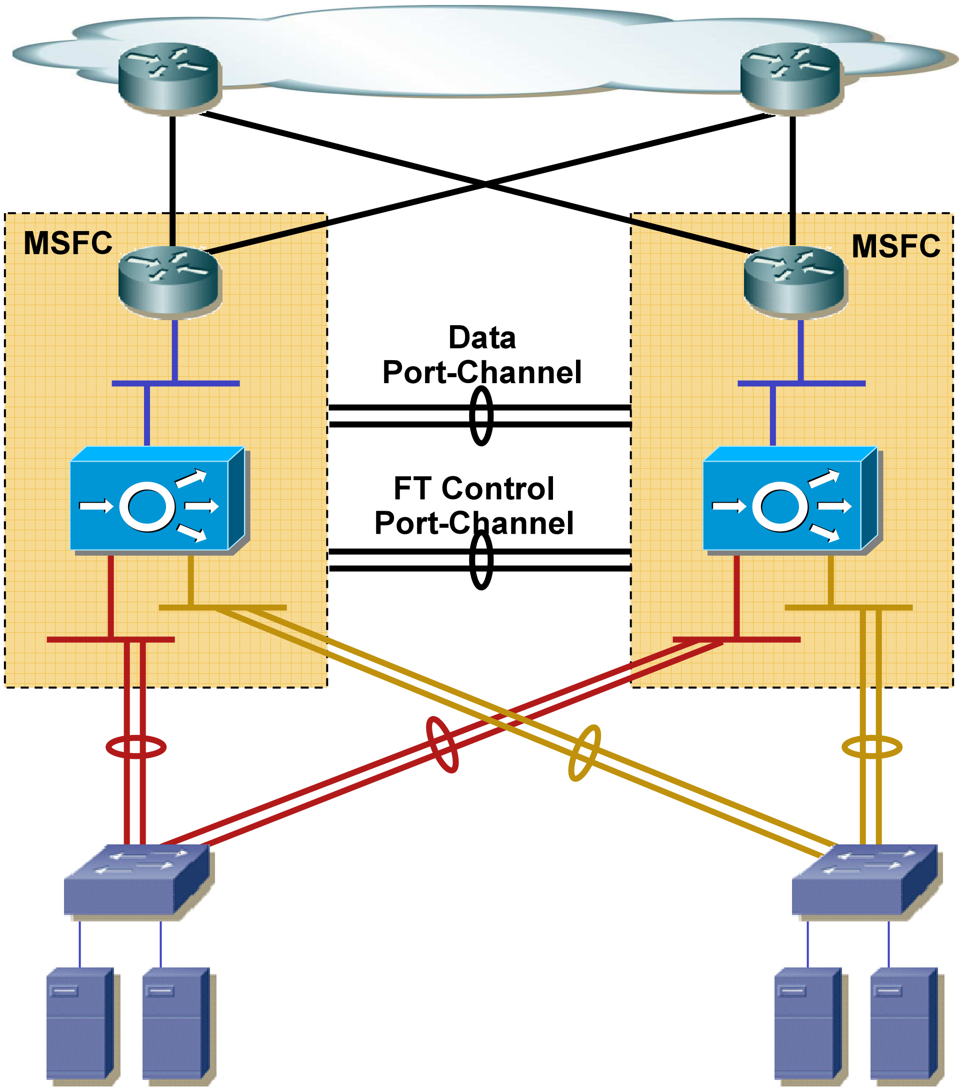

- Client VLAN and server VLANs on different IP subnets
- Servers' default gateway is ACE alias IP
- All data VLANs and FT VLAN carried over port-channels
- Each Cisco Catalyst has redundant physical links to each access switch
- Serverfarms can span multiple access switches
- Management access to servers requires access-list

ACE Deployed in Bridge Mode

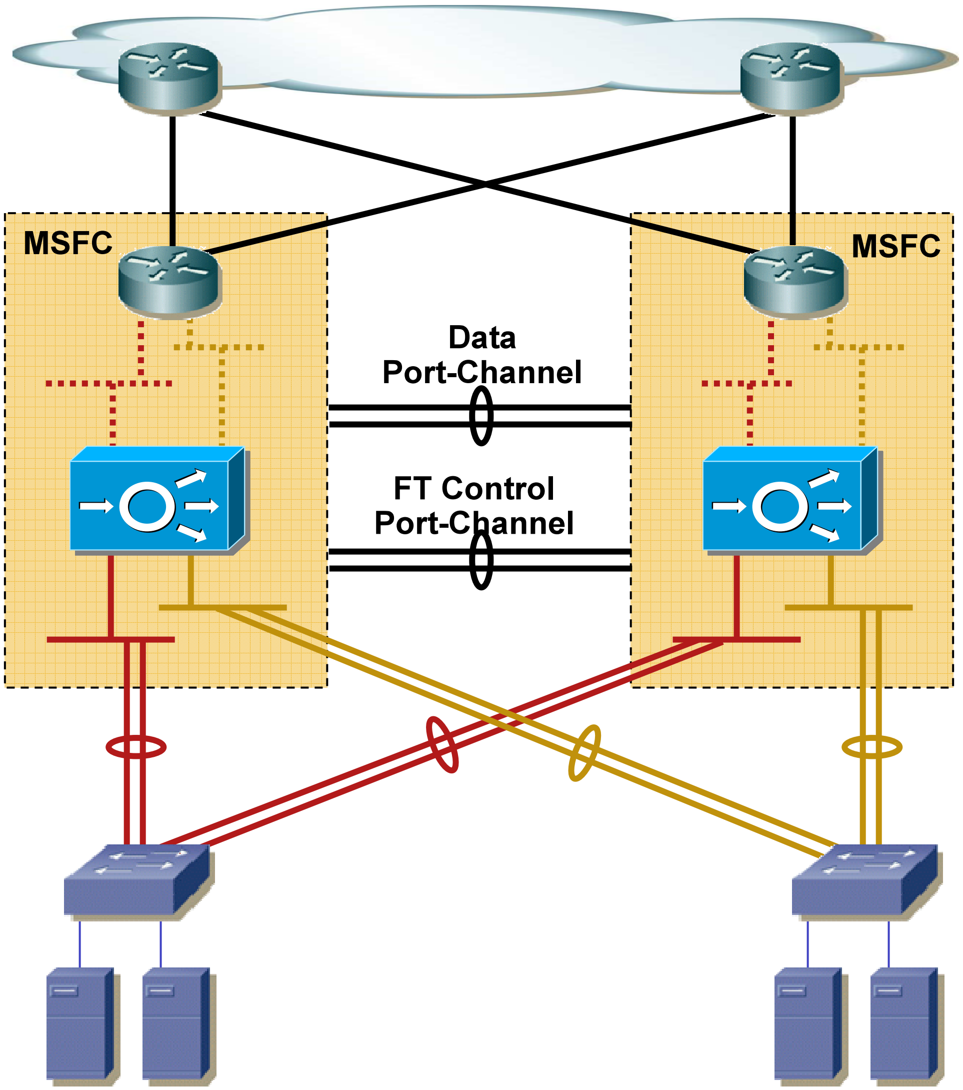

- Pairs of one client and one server VLAN on the same subnet (BVI used to "merge" the two VLANs)
- Limit of two VLANs in the same subnet
- Servers' default gateway is MSFC (or other router) HRSP virtual address
- All data VLANs and FT VLAN carried over port-channels
- Each Cisco Catalyst has redundant physical links to each access switch
- Serverfarms can span multiple access switches
- Management access to servers requires access-list

ACE Deployed in One-Arm Mode

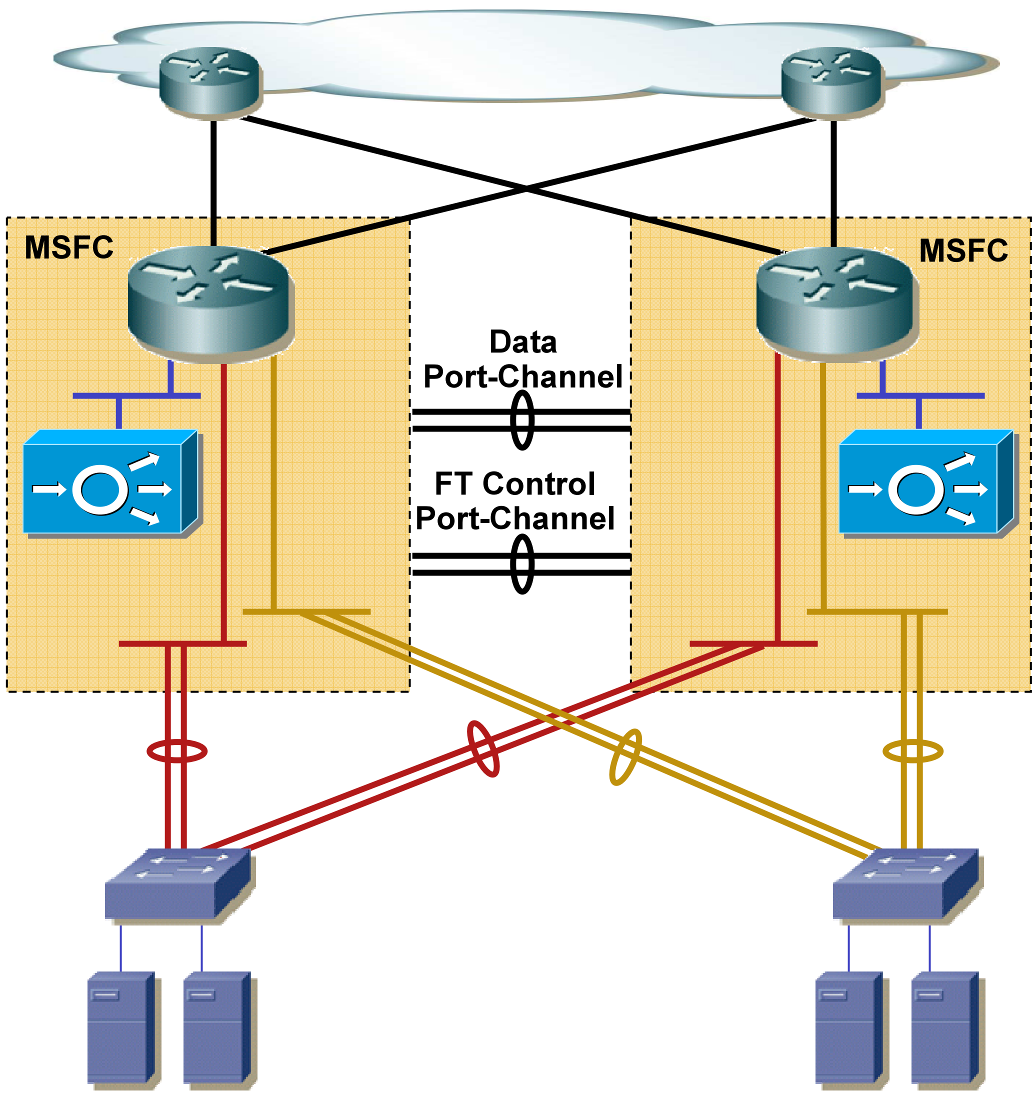

- Single VLAN on ACE
- Servers' default gateway is MSFC HSRP IP
- All data VLANs and FT VLAN carried over port-channels
- Each Cisco Catalyst has redundant physical links to each access switch
- Serverfarms can span multiple access switches
- Management access to servers bypass ACE

ACE Virtual Contexts "L3 Cascaded"

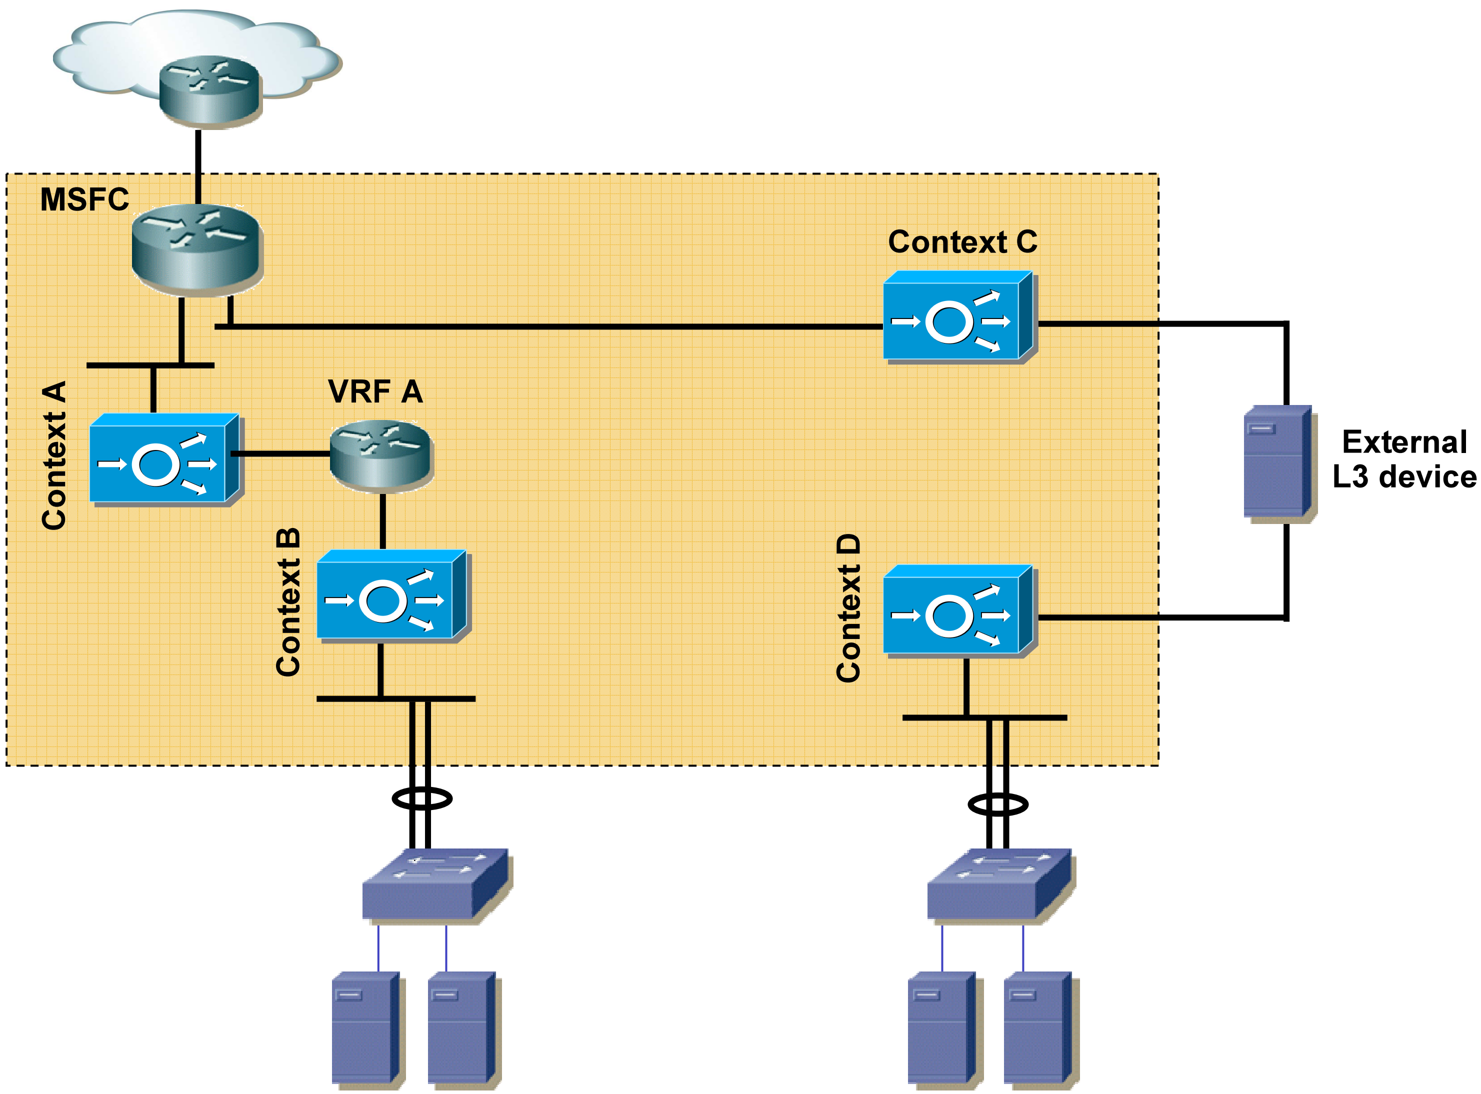

ACE Virtual Contexts Mapped to VRFs

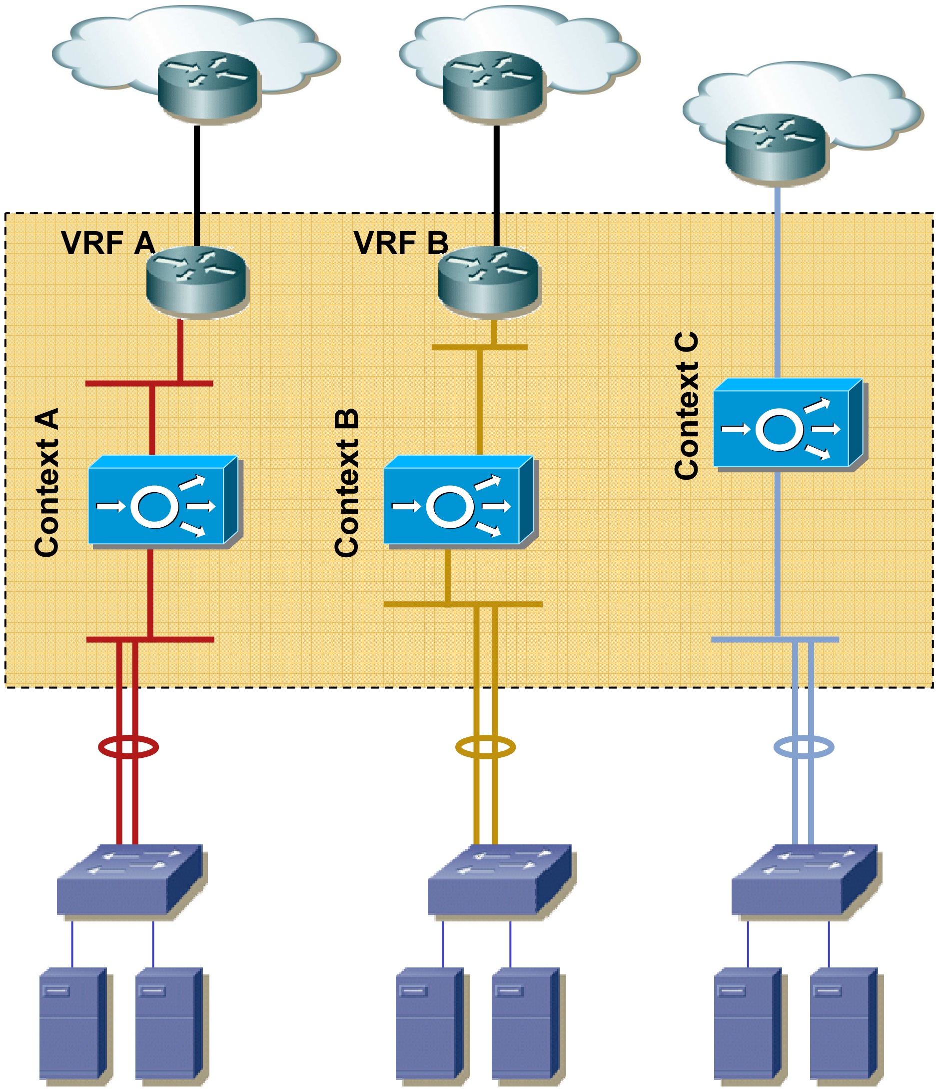

- Virtual Contexts can be mapped to VRFs on the MSFC
- Or directly to external routers
- VRF-aware Route Health Injection (add/remove routes to/from MSFC main routing table as well as VRF routing tables)

Application Control Engine (ACE)

ACE 2.0 Features Include:

Availability
Acceleration
Security

- Enhanced Application Algorithms
- Application-aware Load Balancing
- Generic Protocol Parsing
- Improved Application-level Inspection
- Extended SSL Functionality
- Denial-of-Service Protection
- Increased Health Monitoring Support
- Improved Usability

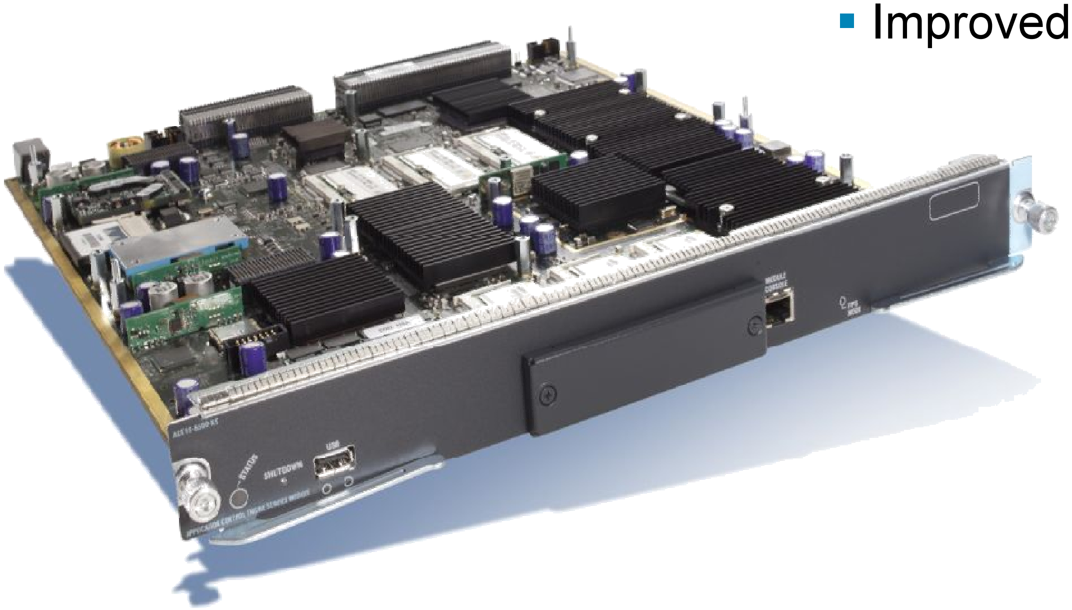

Summary
What's New in ACE ?

ACE

- **250 virtual partitions** (contexts) for L4–7
  Per-context active-standby
- **Role-Based Access Control**
  Predefined and user-configurable roles
- **Reflexive access lists, TCP and IP normalization**
  Protocol inspection
- **HTTP inspection**, RFC compliance, control on headers/payload
- **Modular Policy CLI** for integrated features configuration
- **Configuration rollback** for quick error recovery and testing
- Integrated **SSL termination**, with support for back-end SSL
- **TCP-reuse** for HTTP (aka TCP multiplexing or TCP offload)
- Command Line Interface objects **name completion**

Questions ?