

### Transparent Network Caching

- Connection initiated from web-browser

© 1999, Cisco Systems, Inc. 7

### Transparent Network Caching

- Connection initiated from web-browser
- Router intercepts flow and redirects it to new location (the original packet is encapsulated unchanged within a GRE frame)

© 1999, Cisco Systems, Inc. 8

### Transparent Network Caching

- Connection initiated from web-browser
- Router intercepts flow and redirects it to new location (the original packet is encapsulated unchanged within a GRE frame)
- Device that flow is redirected to can choose what to do with flow --
  - send somewhere else
  - masquerade as real server

© 1999, Cisco Systems, Inc. 9

### Transparent Network Caching

- Connection initiated from web-browser
- Router intercepts flow and redirects it to new location (the original packet is encapsulated unchanged within a GRE frame)
- Device that flow is redirected to can choose what to do with flow --
  - send somewhere else
  - masquerade as real server
- Cache Engine will serve flow (in case of hit), will initiate second flow if a miss

© 1999, Cisco Systems, Inc. 10

### Transparent Network Caching -- subsequent requests

- Connection initiated from web-browser

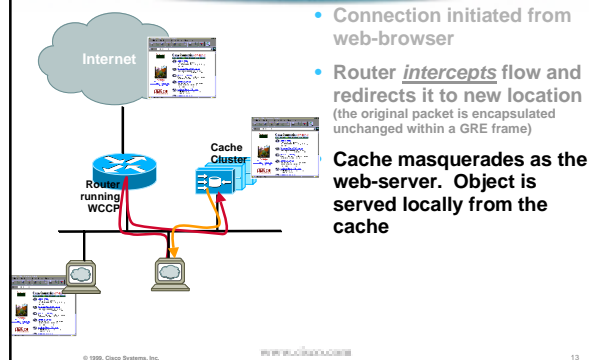
© 1999, Cisco Systems, Inc. 11

### Transparent Network Caching -- subsequent requests

- Connection initiated from web-browser
- Router intercepts flow and redirects it to new location (the original packet is encapsulated unchanged within a GRE frame)

© 1999, Cisco Systems, Inc. 12

## Transparent Network Caching -- subsequent requests



## Network Caching Benefits


- Minimizes redundant traffic on WAN links (help make room for voice)
- Significant bandwidth savings
- Deliver faster network response time
- Significant network cost savings:

$$\text{Monthly savings} = \text{Bandwidth (Mbps)} \times \text{Monthly BW cost/Mbps} \times \text{Percent HTTP traffic} \times \text{Byte hit rate}$$

$$\text{Payback period} = \frac{\text{purchase price}}{\text{monthly savings}}$$


## Enterprise Challenges

- Deployment of new web-based apps
- Significant increases in amount of data traffic (VoIP, ...)
- Increasing operational complexity
- Bandwidth costs
- Slow Web access




## Enterprise Challenges

- Deployment of new web-based apps
- Minimised increase in amount of data traffic
- Network no more complex than before
- Delay expensive bandwidth purchases
- Web is as fast as it can be




## Service Provider Challenges

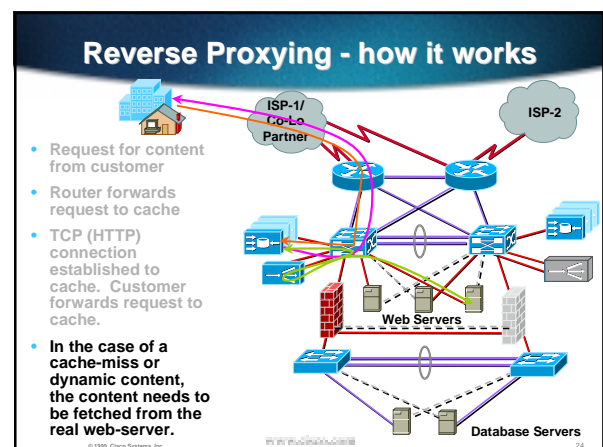
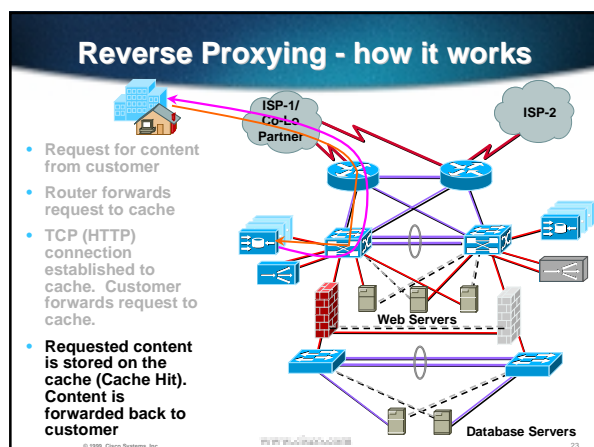
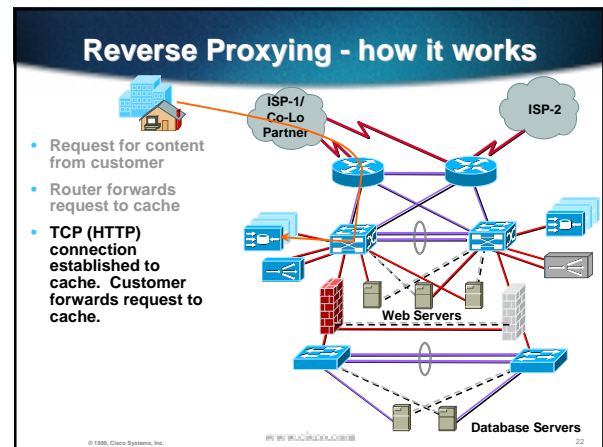
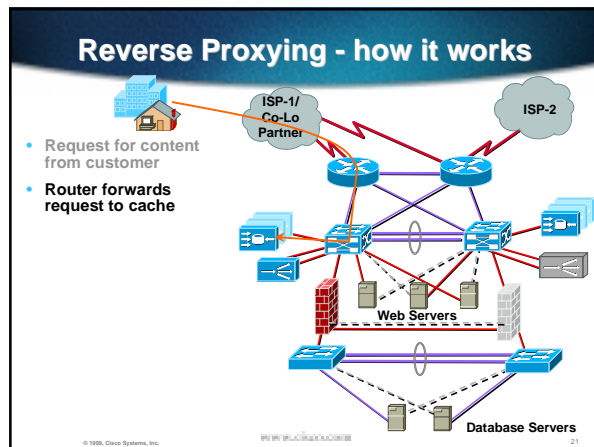
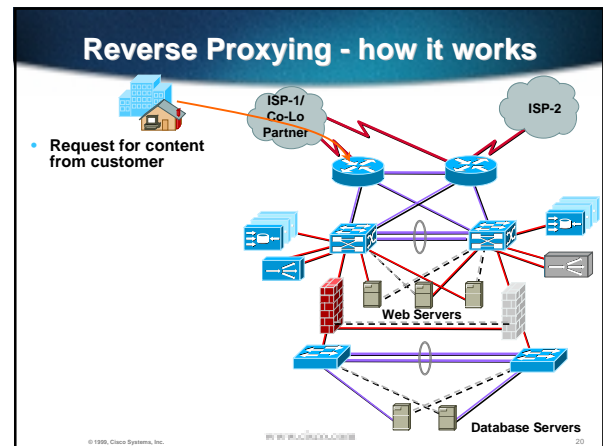
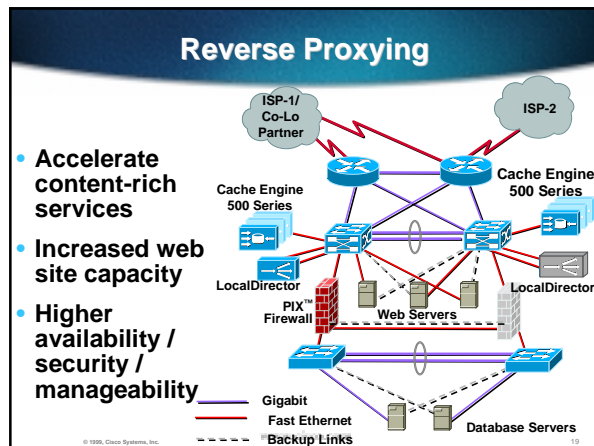
- Exponential Growth
- Ever-increasing demand for bandwidth
- Ever-increasing network complexity



## Service Provider Challenges

- Controlled Growth (cut by at least 30%)
- Network no more complex than before (via the use of WCCP)





## The Problem ...


- The last 5 years were about e-enabling relatively simple applications
- The next 5 years will be about moving more demanding applications on the network
  - e-learning, e-commerce, e-communication
- The applications will require:
  - High bandwidth
  - High reliability
  - Large scale

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

25

## Traditional Web Growth



© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

26

## Solution: Content Delivery Network (CDN)


- Network infrastructure that cost-efficiently scales e-business applications
- Principles:
  - Distributes content to the edge of the network
  - Redirects content requests to the optimal content engine for accelerated content delivery
  - Centrally controls content and CDN components
- Eliminates strain of rich content on infrastructure
- Increases reliability while scaling to address larger audiences

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

27

## 5 Key Service Elements of Content Delivery Networks




© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

28

## 5 Key Service Elements Content Distribution & Mgmt




- Global centralized provisioning for content distribution
- Automatically distribute content to network edge
- Management service for CDN monitoring
- Configuration management for delivery nodes at network edges

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

29

## 5 Key Service Elements Content Routing



- Reliably route user requests for content to the best site across a CDN
- Metrics include: presence of content, geographic proximity, network conditions, POP load, Content Engine load measure of performance and usage
- Adaptive routing around failures/congestion

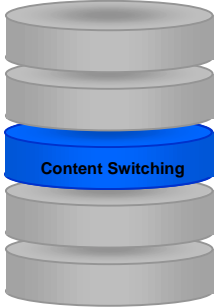
© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

30



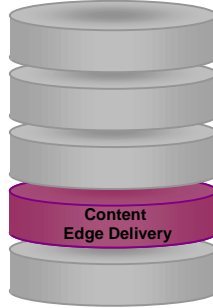
### 5 Key Service Elements Content Switching



- Intelligently switch traffic across origin servers for load balancing
- Hot-spot detection for intelligent global load distribution
- Content verification services to ensure content validity
- Determine availability of content and load on server
- Server acceleration features to enhance user response time

© 1999, Cisco Systems, Inc. 31

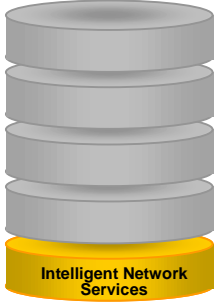
### 5 Key Service Elements Content Edge Delivery



- Import content and maintain copies at edge of network
- Seamlessly deliver content of any type from network edge to end user
- Provide content delivery service, streaming, and transparent caching all in one
- "Appliance" technology for ease of use and reduced administration

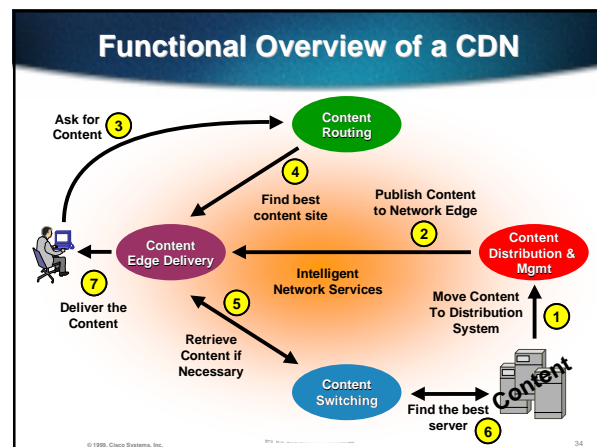
© 1999, Cisco Systems, Inc. 32

### 5 Key Elements Intelligent Network Services

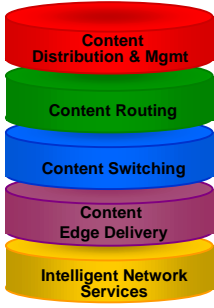


- Cisco IOS Software: Intelligence within the Cisco IP infrastructure
- Required to build reliable and scalable CDNs
- A solid layer 2-7 intelligent and application-aware infrastructure
- Intelligent network services include QoS, VPNs, Security, and Multicast

© 1999, Cisco Systems, Inc. 33



### What's New?



- 3 Content Distribution Managers (CDM)  
+ CDN Software, Enterprise 2.0  
+ CDN Software, Service Provider 1.0
- 3 Content Routers (CR)  
+ CR 4400 + CR 7200  
+ CR 4450
- 3 Enhanced Content Switching Solutions  
+ CSS-11000 with WebNS 4.0  
+ Catalyst 4840G SLB Switch  
+ LocalDirector with 4.1 software
- 5 New Content Engines (CE)  
+ CE507, 507AV, 560, 590  
+ CE7320  
With CDN, Web Caching, and Streaming


© 1999, Cisco Systems, Inc. 35

### Content Distribution Managers


**CDM**

**CDM 4600 Series:**

- Management for CDN
- CDN policy & configuration database
- Content, customer registration & domain assignment
- CDN Network monitoring
- CDN Accounting interface
- Enterprise content registry
- Peering/Billing (Future)



**CDM 4650 / 4670**



**CDM 4630**

© 1999, Cisco Systems, Inc. 36

## Content Routers

CR

- Distributed Director Server Algorithm
  - Routing based on routing metrics
  - BGP Route maps (policy)
  - Dynamic Feedback Protocol support
- SightPath SODA Server Algorithm
  - Scales to >10K nodes
  - DNS hierarchical mechanism
  - Adapts to changing network conditions
- Boomerang Server Algorithm
  - DNS 'Triangulation' Algorithm
  - Transparent Server Selection
  - Server health check w/ HTTP GET



DD 7200



CR 4450



CR 4400

© 1999, Cisco Systems, Inc.

www.cisco.com

37

## Content Engines

CE

### Content Engine 7300 Series:


- Service Provider & Large Enterprise Deployment
- Caches Enabled Content
- Populated via Pull or Push
- Serves Requested Content



CE 7300 Series

### Content Engine 500 Series:

- Enterprise Deployment
- Transparent caching & content delivery software modules
- Static files, Real G2, MPEG 1 & 2 support



CE 500 Series

© 1999, Cisco Systems, Inc.

www.cisco.com

38

## Datacenter Acceleration with LocalDirector

Future Enhancements in Local Director 4.1 and beyond:  
URL Load Balancing  
Content Routing Agents

SSL Acceleration  
TCP Aggregation  
XML Services (peering)

LocalDirector 430/ 416



Dynamic Web Page Acceleration

Predictive Content Replication

### Proven Reliability

- Stateful Failover
- Layer 7 Persistence
- Content Routing Protocol Support
- Sophisticated Content Verification

© 1999, Cisco Systems, Inc.

www.cisco.com

39

## Content Switches - Layer 2-4

CS

### Catalyst 6500 Family:

- Up to 256 Gbps switching - modular
- Full internet routing protocol support
- HW-acceleration with Cisco IOS features
- WAN Interfaces (T1/E1, POS, ATM)
- Advanced Security Services (ACL / IDS)



Catalyst 6500

### Catalyst 4840G Switch:

- 22 Gbps switching / 9 Mpps - fixed
- Wire Speed NAT
- HW-acceleration with Cisco IOS Features



Catalyst 4840G

© 1999, Cisco Systems, Inc.

www.cisco.com

40

## Content Switches - Layer 7

CS

- Scalable product family
  - CSS 11050 - 8 FE, 1 GBE
  - CSS 11150 - 12-16 FE, 2 GBE
  - CSS 11800 - 64 FE, 32 GBE
  - Up to 11 billion hits/day
- Smart
  - Full content based switching using content, application, network and user level info
  - Hot-spot or 'Flash Crowd' support
  - Intelligent Web Cache Bypass
- Fast
  - Up to 130,000 HTTP transactions per second



Content Switch  
11000 Family

© 1999, Cisco Systems, Inc.

www.cisco.com

41

## The Cisco Difference

Content  
Delivery  
Networks

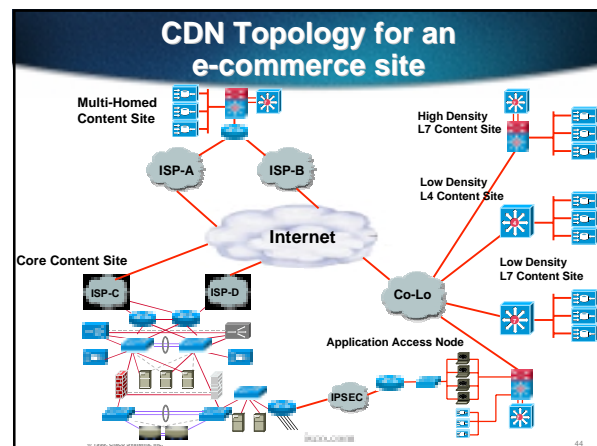
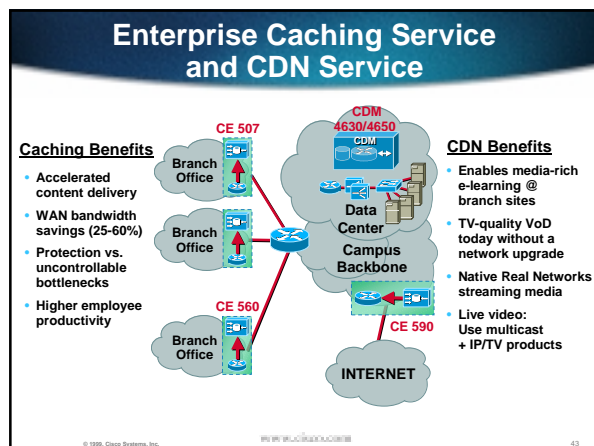


- Cisco is the first and only vendor to deliver a complete system for Content Delivery Networks and Next Generation Content Based Services
- Enables Enterprises and Dot.coms to grow brand equity and increase customer loyalty by improving end-user experience
- Enables Service Providers to build value into their networks and to deploy highly profitable content delivery services

© 1999, Cisco Systems, Inc.

www.cisco.com

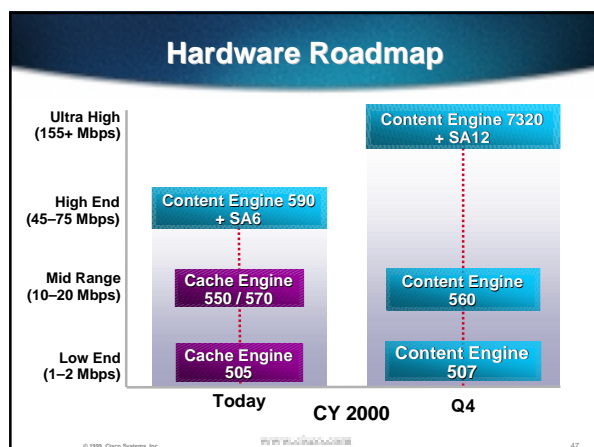

42



## Enterprise Intranet/Extranet Solutions

Problem	Service Solution	Cisco Products
Internet Access not Fast Enough, WAN Bandwidth Scalability	Caching	Content Engine WCCP Router/Switch
Unable to Deliver High-Quality Video on Demand to Clients Today	CDN	Content Engine Content Distrib. Mgr
Objectionable or Non-Business Internet Content	Employee Internet Management (EIM)	Content Engine Websense Software

© 1999, Cisco Systems, Inc. 45



## Cisco Content Engine Product Line

Customer	Solution	HTTP Performance	Rack Units	Storage (GB)	Base Price
Large Service Providers	CE 7320 (a)	155+ Mbps	4-7	180-396	\$90,000
Service Providers, Enterprises	CE 590 (a)	45+ Mbps	1-4	36-252	\$30,000
Regional Sites	CE 560 (a)	20 Mbps	1-4	36-144	\$15,000
Small Branch Sites	CE 507 CE 507AV	T1/E1 T1/E1	1 1	18-36 18	\$5,500 \$6,995


All models have (2) 10BaseT/100BaseTX interfaces. CE 7320 also has (2) GigE (fiber) ports. AC and DC versions will be offered for the CE 590 and CE 7320.

(a) An external Cisco Storage Array is required for optimal performance.  
 Cisco Storage Array 6: 6x18=108 GB: \$13,000 (3 RU)  
 Cisco Storage Array 12: 12x18=216 GB: \$21,500 (3 RU)

© 1999, Cisco Systems, Inc. 48



## Cisco Content Engines



### Content Services Platform



© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

49

## Content Engine Solutions

Solution	Benefit	Orderability
Enterprise Caching	Speed, Bandwidth Savings	Today
Enterprise CDN	VoD today	Q4 CY '00
Enterprise EIM	Content Control	Today
E-Commerce Caching	Scalability	Today
E-Commerce CDN	Scalability/Redundancy	Q4 CY '00
SP Caching	Speed, Bandwidth Savings	Today
SP CDN	New Hosting Service	December

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

50

## Software Releases

- Content Engines: CDN or Cache software**
  - CDN software: CDN service
  - Cache 2.x, 3.x software: Caching/EIM services
- Cache Engines: Cache 2.x software**
  - Cache 2.x sw: Caching/EIM services (final feature release in Q4 CY '00)

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

51

## CDN and Caching

Capability	Caching	CDN
Content Source	All Content Requested	Only Assigned Domains and Files
Content Population	Pull or Cache-Driven Prepopulation	Centralized Prepopulation, Based on Policies
Request Routing	Local—Intercept Traffic and Deliver From Cache	Global—Redirect Request to Optimal Node
Load Balancing	Local Clustering	Global, Across Entire CDN
Content Management	Decentralized—Automated on CEs	Centralized Content Management
Content Retention	Store Frequently Requested Files	Policy-Driven. Can Ensure 100% Hit Rate

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

52

## Software Releases

	Software Release			
	Cache 2.X	Cache 3.X	CDN Enterprise	CDN Service Provider
Content Engine 7320	—	3.0+	—	1.0+
Content Engine 590	2.2+	3.0+	—	1.0+
Content Engine 560	2.3.4+	3.1+	Future	—
Content Engine 507	2.3.4+	3.1+	2.0+	—
Content Engine 507AV	—	—	2.0+	—
Cache Engine 505/550/570	Yes	—	—	—

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

53

## Content Engine Summary

- Leading overall performance and density
  - Density = performance / RU
- Cisco is the only vendor that offers the complete CDN solution
- Most advanced **transparent** caching
- Powerful Employee Internet Management services
- Scale from super PoPs to branch sites

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

54

# Leading Overall Performance

## The Second IRCache Web Cache Bakeoff Results

Raw Data from IRCache Report


Rankings						Raw Data from IRCache Report		
Place	Product	Overall Performance Ranking	Throughput Ranking	Latency Ranking	Byte Hit Rate Ranking	Throughput (URLs/sec)	Latency (sec)	Byte Hit Rate
1	Cisco CE 7200	3.3	2	6	2	2304	1.48	57%
2	Compaq C2500	3.7	1	5	5	2400	1.45	55%
3	Mimic-5000	5.0	4	1	10	1453	1.35	47%
4	Cisco CE 590	5.3	7	7	2	951	1.50	57%
5	Squid-2.4	6.3	16	2	1	160	1.41	58%
6	Lucent-1505	6.3	9	9	2	771	1.53	57%
7	Network Appliance	9.0	17	3	7	151	1.44	53%
8	IBM-5600	9.3	6	10	12	1906	1.65	43%
9	Lucent-100	9.7	9	13	7	675	1.71	53%
10	Swell-1000	10.0	20	3	7	77	1.44	53%

(a) Overall performance ranking is the average of the throughput, latency, and byte hit rate rankings.  
(b) Cacheflow pulled out of this competition because their Polymix-2 benchmark numbers were lower than the 2300 TPS that they achieved with the older DataComm-1 benchmark. Results obtained on different benchmarks are not comparable.  
(c) Inktomi failed to register for the bakeoff due to scalability issues.

Source: <http://polygraph.ircache.net/Results/bakeoff-2/>  
Polymix-2 Benchmark: <http://polygraph.ircache.net/sources/>

© 1998, Cisco Systems, Inc.

55



- ### Browser Function
- Go to [www.cisco.com](http://www.cisco.com)
  - Pull down index (HTML) file
  - Fetch individual objects specified in the index file
  - HTTP 1.0: one object/TCP connection
  - HTTP 1.1: multiple objects/connection
- © 1999, Cisco Systems, Inc.

- ### Dynamic Content
- Banner ads, stock quotes
    - Marked as noncacheable to overcome every browser's cache
  - Active server pages (.asp)
    - Dynamically builds a Web page of objects
    - Index file is noncacheable
    - But most objects will be static
- © 1999, Cisco Systems, Inc.

## Content Freshness

- HTTP 1.0: cache/no cache header
- HTTP 1.1: expiration tag header
- When object is cached,  
Time To Live (TTL) value =

**(Today—Modified) \* Freshness Factor**

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

61

## Content Freshness

- Revalidate content freshness with If-Modified-Since (IMS) feature

Issued upon a request for expired content or when client browser issues an IMS

```
(to web-server)
GET /index.html HTTP/1.1
Server: www.cisco.com
Connection: keep-alive
If-Modified-Since: Tue 12 Sep 2000 10:07:04 GMT
Accept: */*
```

```
(web-server response)
304 Not Modified OR 200 OK
(end-of-request)      (response headers)
                      (data)
                      (end-of-request)
```

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

62

## CE Configuration

Presentation\_ID

© 1999, Cisco Systems, Inc.

www.cisco.com

63

## Cache Engine Configuration

Cache just like any other Cisco device as far as setup

- Initial cache configuration in less than 10 commands.
- User can have Cache operational in less than 15 minutes.

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

64

## Initial Cache installation

- Power on cache. Cache boots up ..  
Initial configuration dialog on console:

```
---- System Configuration Dialog ----
You may use Ctrl-D to abort Configuration Dialog at any prompt.
Would you like to enter the initial configuration? [yes]: yes
host name: sj-cache1
domain name: mysite.com
admin password: reallysecret
cache IP address: 192.168.10.40
IP network mask: 255.255.255.0
gateway IP address: 192.168.10.1
DNS server: 171.69.2.132
---- System Configuration Complete ----
```

- Cache services may now be configured:

```
Username: admin
Password: reallysecret
sj-cache1# conf t
Enter configuration commands, one per line. End with CNTL/Z
sj-cache1(config)# wccp router-list 1 192.168.10.1
sj-cache1(config)# wccp web-cache router-list-num 1 password secretcache
sj-cache1(config)# ^Z
```

- Cache operational in only 9 commands!

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

65

## WCCP transparent proxying



```
ROUTER
ip wccp version {1,2}
ip wccp web-cache password <pass>
!
interface <internet-lf/ace>
ip wccp web-cache redirect out
!
```


```
CACHE(s)
wccp version {1,2}
wccp router-list 1 <router ip addr>
wccp web-cache router-list-num 1 pass <pass>
```

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

66

## WCCP transparent proxying on ports other than 80



```
ROUTER
ip wccp version 2
ip wccp 98 password <pass>
!
interface <internet-l/face>
ip wccp 98 redirect out
!
```

```
CACHE(s)
wccp version 2
wccp router-list 1 <router ip addr>
wccp custom-web-cache router-list-num 1
port 8080 pass <pass>
```

© 1999, Cisco Systems, Inc.

W W W C C P

67

## Cache Engine Configuration

- Configurable via IOS CLI (telnet, console port)
- Configurable via web-based GUI (http://cache:8001/)
- Monitoring via SNMP, SNMP traps, syslog, web-based GUI, CLI

© 1999, Cisco Systems, Inc.

W W W C C P

68

## [new] Rules Template Feature

- Client request pattern match (pick one)
  - URL regular expression
  - Destination domain name
  - Source/destination IP address
  - Destination port number
  - MIME type
- Action (pick one)
  - Do not cache
  - Only cache objects that match the pattern, nothing else
  - On a cache miss, go to a specified server or proxy
  - Do not use upstream proxy, do internal DNS resolution
  - Revalidate content via IMS
  - Do not authenticate
  - Block it

© 1999, Cisco Systems, Inc.

W W W C C P

69

## CE Deployment Scenarios with WCCP

© 1999, Cisco Systems, Inc.

www.cisco.com

70

## Design Objectives for the ISP

- Transparent *Redirection* of a IP flow based on source, destination, and/or port number.
- Transparent *Integration* - no rebuilding the POP to add this service.
- Failed open - if the service fails, it should not effect the core IP service nor any other services.


© 1999, Cisco Systems, Inc.

W W W C C P

71

## Design Objectives for the ISP

- Not to effect the primary packet flow of the POP - if not redirected - then is CEF/dCEF Switched!




© 1999, Cisco Systems, Inc.

W W W C C P

72


## Design Objectives for the ISP

- Work with the multi-level L2/L3 redundancy of the ISP POP. Equal paths in the IGP + CEF leads packet asymmetry.



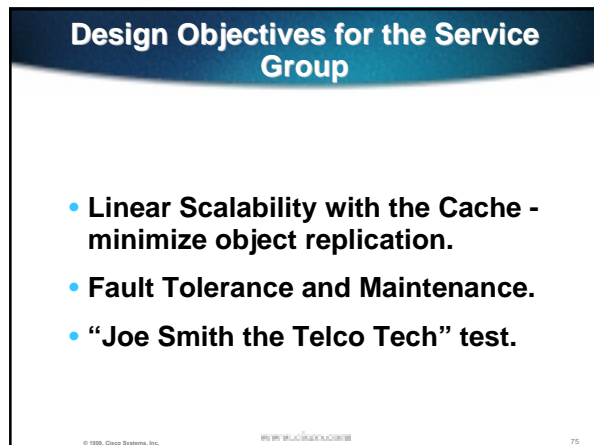
## Design Objectives for the ISP

- Provide the ISP with Flexibility on the point of redirection.




## Design Objectives for the Service Group

- Linear Scalability with the Cache - minimize object replication.
- Fault Tolerance and Maintenance.
- "Joe Smith the Telco Tech" test.




## Basic Transparency



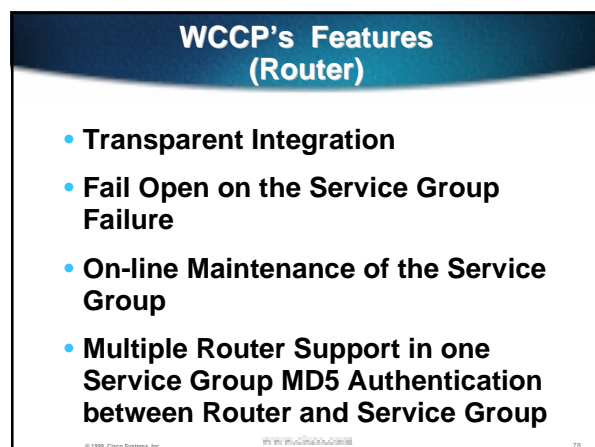
## WCCP's Features

- WCCP's Features are in two parts:  
Router Based - Benefiting the operation of the ISP Edge (POP) or Enterprise Gateway.  
Service Group - Benefiting the applications WCCP is servicing



## WCCP's Features (Router)

- Transparent Integration
- Fail Open on the Service Group Failure
- On-line Maintenance of the Service Group
- Multiple Router Support in one Service Group MD5 Authentication between Router and Service Group





## WCCP's Features (Router)

- CEF and dCEF Switched
- Multiple Service Groups
- Options on where the redirections happens

© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

79

## WCCP's Features (Service Group)

- Fault Tolerance of the Service Group
- On-line Maintenance of the Service Group
- Linear Scalability of the of the Service Group
- WCCP Slow Start

© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

80

## WCCP's Features (Service Group)

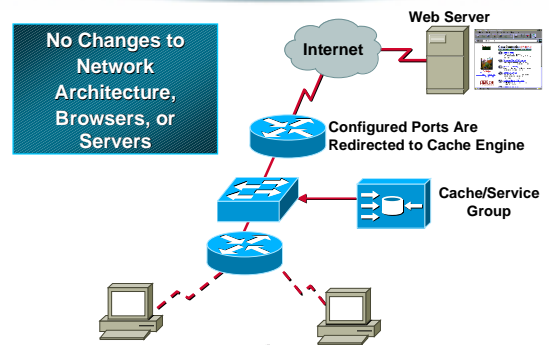
- Fault Prevention - Packet Return Feature (Overload and Bypass)
- Load Distribution (Hot Spots)
- Fail Open on the Service Group Failure
- Authentication By-pass

© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

81

## Transparent Integration




© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

82

## Fail Open




© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

83

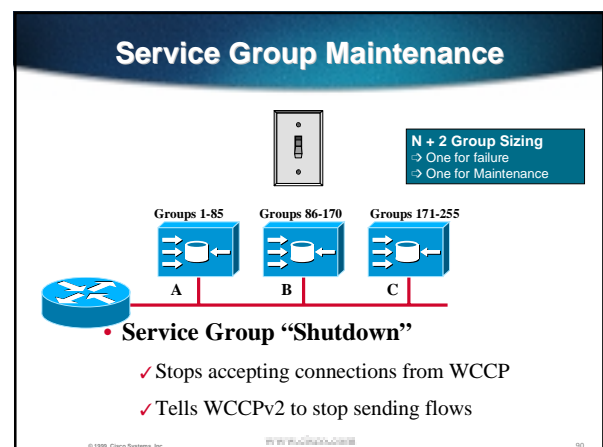
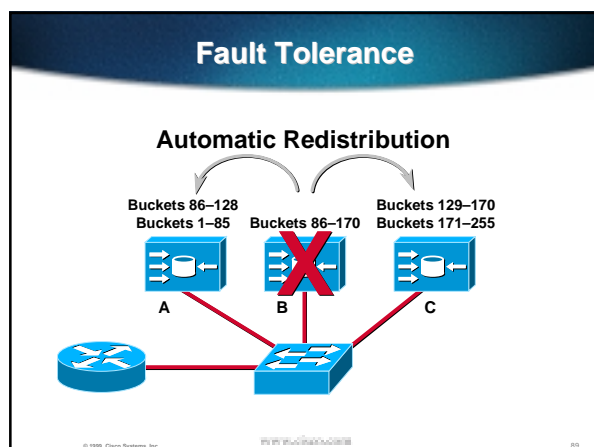
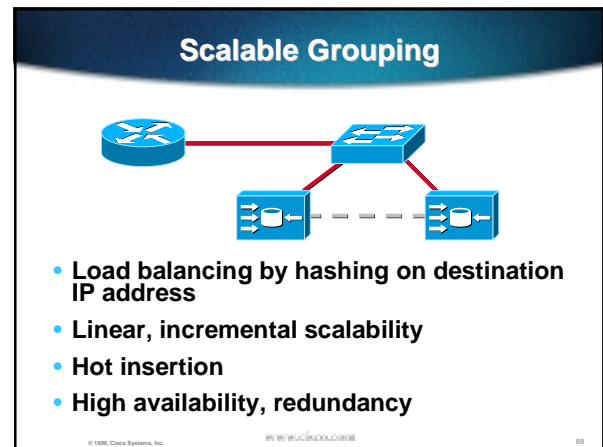
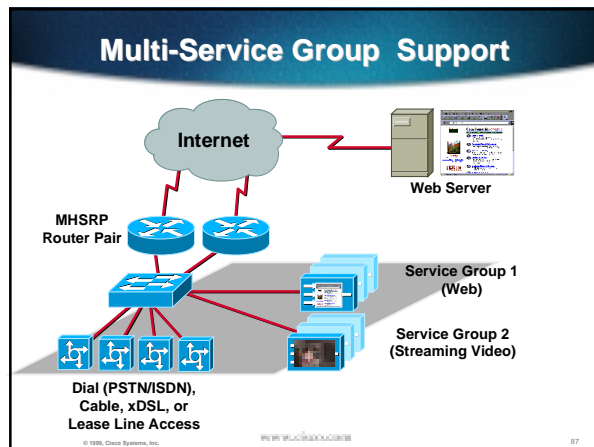
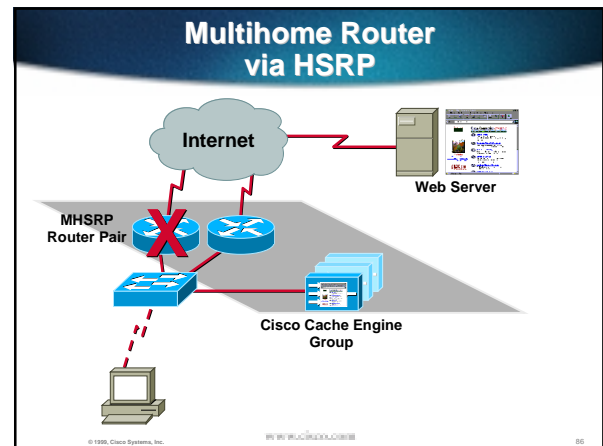
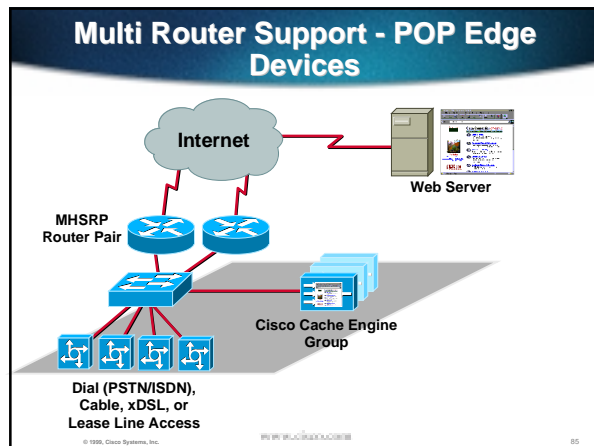
## Multi Router - POP Border Routers

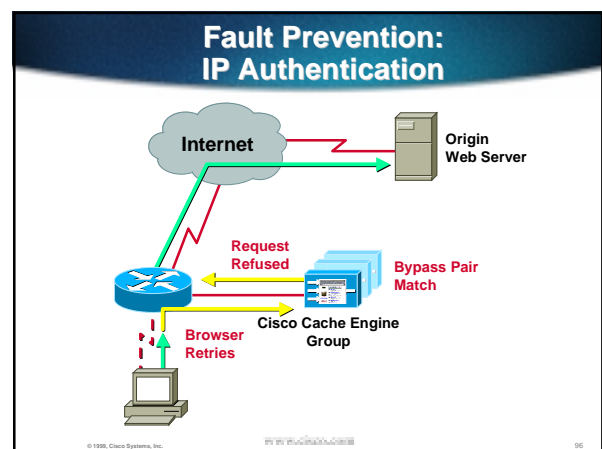
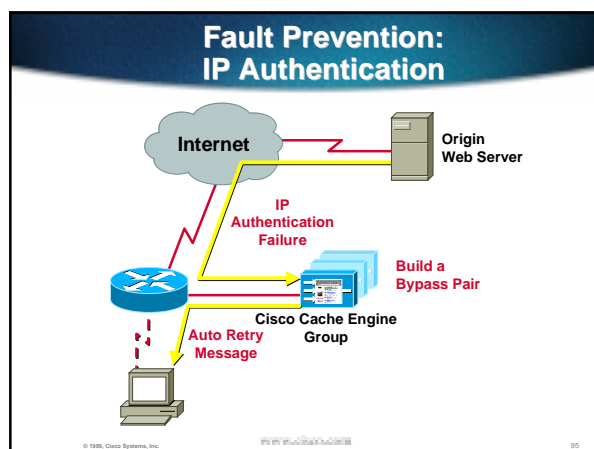
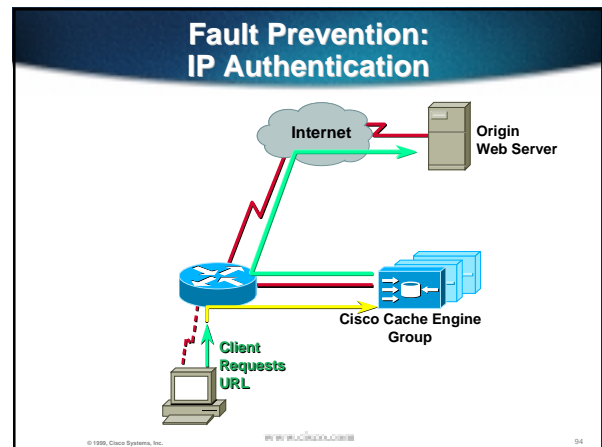
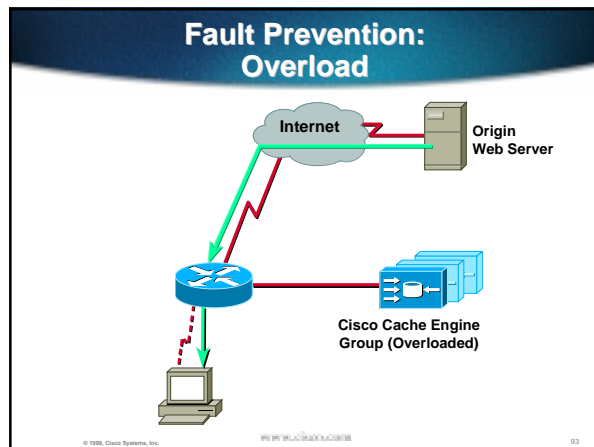
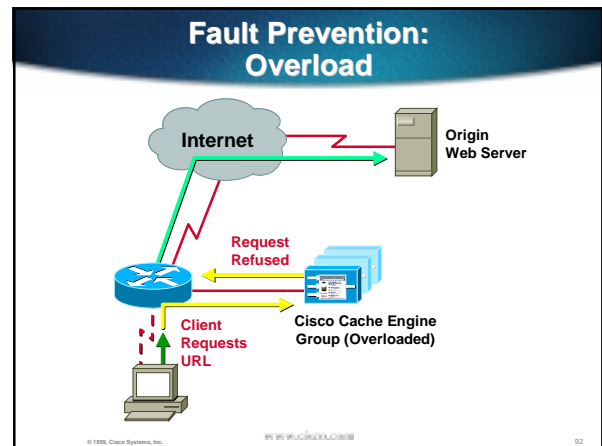
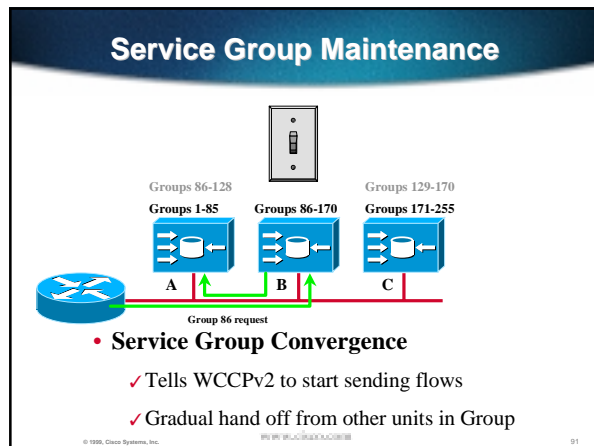


© 1999, Cisco Systems, Inc.

W W W C C P . C O . C O

84





## Reverse Proxying via WCCP

**Routers**

```
ip wccp 99 password pass1
interface fastethernetX/Y
ip wccp 99 redirect out
```

**Cache**

```
..
wccp router-list-num 1 10.64.70.1 10.64.70.2
! both WCCP-enabled routers listed
wccp reverse-proxy router-list-num 1 \
password pass1
! enable reverse-proxy service
..
rule enable
rule use-proxy 10.64.70.18 80 \
domain site.com
! send cache misses to
! LocalDirector
! VIP 10.64.70.18
http proxy outgoing 10.64.70.19 80
! send all other (non-listed)
! domains to a cache-all VIP
! on LD that tells user that
! the domain does not exist
```

97

## Reverse Proxying via WebNS

**WebNS**

```
owner www.site.com
content reverse-proxy-rule
vip address 10.64.70.19
protocol tcp
port 80
url /*/* eq1 cachable
add service cache1
add service cache2
add service cache3
```

**Cache**

```
no bypass load enable
! we always want to serve requests
! offered to us
http 14-switch enable
! we wish to operate in transparent mode
rule enable
rule use-proxy 10.64.70.18 80 \
domain site.com
! send cache misses to
! WebNS VIP 10.64.70.18
http proxy outgoing 10.64.70.17 80
! send all other (non-listed)
! domains to a cache-all VIP
! that tells user that
! the domain does not exist
```

98

## Sizing correct cache

**3 important parameters for sizing a cache:**

- Transactions/sec (aka req/sec, URLs/sec)
- Concurrent # TCP conns.
- Cache Disk capacity

All are sized appropriately in CE560/590/7320 for typical "Service Provider" traffic

If traffic patterns are not "typical", then consult Deployment Guide.  
(see <http://bock-bock/-ltd/cache/>)

Ultra High (100+ Mbps)

High End (40-75 Mbps)

Mid Range (10-25 Mbps)

Low End (1-2 Mbps)

CY 2000

99

## Accelerated WCCPv2 for Catalyst 6x00

**WCCPv2: GRE Encap, CEF/DCEF Switched  
Accelerated WCCPv2: L2 Rewrites, No GRE**

	Supervisor 1				Supervisor 2	
	MSFC 1 GRE	MSFC 1 L2	MSFC 2 GRE	MSFC 2 L2	MSFC 2 GRE	MSFC 2 L2
Conns/sec	50K	50K	150K	150K	150K	1M+
Throughput	170 Kpps	15 Mpps	510 Kpps	15 Mpps	510 Kpps	30 Mpps

**Caches need to be L2-connected to switch**

100

## Accelerated WCCPv2 for Catalyst 6x00

- **Today: Supervisor 1 support**  
CE 590+SA6: Cache 2.2 software  
Catalyst 6000: Catalyst OS 5.5, MSFC: Cisco IOS 12.1(2)E
- **Q1 CY '01: Supervisor 2 support**  
CE 7320+SA12 or CE 590+SA6: Cache 3.2 software  
Catalyst 6000: Catalyst OS 6.1, MSFC: Cisco IOS 12.1(4+)E

101

## WCCP: where to enable it?

- CEF switched in 12.0(4)T
- In 12.0(11)S  
Input interface feature  
7500+VIP DCEF switched  
Flow acceleration  
BGP Policy Propagation  
=> **Minimal CPU increase**
- Run WCCP2 on access routers and attach Content Engines to a switch
- Load is distributed across multiple routers, no increase in number of packets

102

## WCCP in 12.0(11)S

- Ability to define traffic which can be intercepted via route-map.

Very powerful -- provides for selective inclusion in cache eligibility

- 'Premium' hosting

Service Providers can offer transparent backbone caching. Peers/customers can choose to participate by setting bgp community/MED

- Cache-only-dial-pool

Provider only wants to cache dial or DSL pool, yet address space is segregated.

- Selective intercept based on administrative pref

Only cache traffic which is due to go out an expensive path (eg. International)

© 1999, Cisco Systems, Inc.

103

## Policy Propagation with WCCP

### Using MTRES vs ACLs

Presentation\_ID © 1999, Cisco Systems, Inc.

www.cisco.com

104

## BGP Policy Propagation for WCCP

- Problem: Caching is an *operational* savings. What ISPs and Co-Lo Providers are looking for is a new revenue stream - CDNs
- Problem: How to maintain redirection ACLs and Route-Maps that will point redirected packets to the correct CDN service? (think 1000s of devices w/ ACLs)

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

105

## BGP Policy Propagation for WCCP

- Answer - use the FIB!

The FIB has the capability to add extra fields to describe a prefix.

Currently (12.0(11)S) there are four extra FIB fields - precedence, qos\_group, traffic\_index, and wccp\_tag


Features would use a MTRE look-up in the FIB to get information on what to redirect.

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

106

## BGP Policy Propagation for WCCP



© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

107

## Example - Step 1

Step 1- Router R2 (or another Router) mark the prefix with a community

```
!
router bgp 210
 neighbor 210.210.14.1 remote-as 210
 neighbor 210.210.14.1 route-map comm-relay-prec out
 neighbor 210.210.14.1 send-community
!
ip bgp-community new-format
!
access-list 1 permit 210.210.1.0 0.0.0.255
!
route-map comm-relay-prec permit 10
 match ip address 1
 set community 210:5
!
route-map comm-relay-prec permit 20
 set community 210:0
!
!
```

© 1999, Cisco Systems, Inc.

WWW.CISCO.COM

108



## Example - Step 2

Step 2 - Use the BGP Update to match the community and set the value in the FIB


```
!
router bgp 210
  table-map precedence-map
  neighbor 200.200.14.4 remote-as 210
  neighbor 200.200.14.4 update-source Loopback0
!
ip bgp-community new-format
!
ip community-list 1 permit 210:5
!
route-map precedence-map permit 10
  match community 1
  set ip wccp 50
!
route-map precedence-map permit 20
!
!
```

© 1999, Cisco Systems, Inc.

W W W C C P

109

## Example - Status



© 1999, Cisco Systems, Inc.

W W W C C P

110

## Example - Step 3

Step 3 - WCCP used the a FIB lookup to get the WCCP\_TAG. It then redirected based on the WCCP\_TAG value.

```
!
ip wccp version 2
ip wccp web-cache password <pass> policy source 50
!
interface <xyz>
ip wccp web-cache redirect in
!
```

© 1999, Cisco Systems, Inc.

W W W C C P

111

## BGP Policy Propagation for WCCP

- Very powerful -- provides for selective inclusion in cache eligibility
  - 'Premium' hosting
    - Service Providers can offer transparent backbone caching. Peers/customers can choose to participate by setting bgp community/MED
  - Cache-only-dial-pool
    - Provider only wants to cache dial or DSL pool, yet address space is segregated.
  - Selective intercept based on administrative pref
    - Only cache traffic which is due to go out an expensive path (eg. International)
  - Redirects into CDN Services

© 1999, Cisco Systems, Inc.

W W W C C P

112

## Another Example

- The following example shows only "premium" traffic being cached.

"Premium" traffic is defined as traffic which has:

The policy defined below is:

- any traffic with community 4433:1050 set,
- any traffic with community 4433:1055 set,
- any traffic originating from directly-connected AS 65521,
- any traffic passing thru directly-connected AS 65522,
- any traffic passing thru AS 65523

is eligible for intercept.

Standard "web-cache" service is used -- which is a standard assignment of 'match tcp destination port 80', distribute traffic among participating caches as hashed by destination ip address.

© 1999, Cisco Systems, Inc.

W W W C C P

113

## Another Example

```
!
ip cef distributed      # ensure Distributed CEF is enabled
!
ip wccp version 2      # enable WCCPv2
ip wccp web-cache password <pass> policy source 50
!
interface <xyz>
  ip wccp web-cache redirect in  # incoming i/face # redirect on input traffic
!
router bgp XXXX
  table-map neighbor-xyz-in  # BGP Updates the FIB's WCCP_TAG field
!
ip bgp-community new-format
ip community-list 3 permit 4433:1050  # AS4433 community 1050 is premium
ip community-list 3 permit 4433:1055  # AS4433 community 1055 is premium
!
ip as-path access-list 121 permit ^65521$  # only traffic from AS65521 is premium
ip as-path access-list 121 permit ^65522$  # any traffic thru AS65522 premium
!
route-map neighbor-xyz-in permit 10  # incoming route filter on
  match as-path 121
  set ip wccp 50
!
route-map neighbor-xyz-in permit 15
  match community 3
  set ip wccp 50
!
```

© 1999, Cisco Systems, Inc.

W W W C C P

114

## The Caveat

- BGP Policy Propagation for WCCP was only committed to 12.0(11)S.  
Hence it is currently in 12.0(11)S and it's children - 12.0SC and 12.0SL
- Work is underway to have this committed to 12.1T and find ways for it to work on the EARL and GSR architecture (*issue is the MTRE for the source address*).

© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

115



© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

116

## Content Routing

3203

1346\_06\_2000\_02\_001

© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

117

## Content Routing


- The goal of content routing is to direct a user request to the closest, available content delivery site that has the content the user is requesting.

© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

118

## Content Routing




- DNS Based Technology
- Dynamically returns the "best" IP address for DNS queries
- Load Distribution Decisions include Network Topology and/or Link Latency

© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

119

## DNS Lookup Process



Authoritative DNS is delegated for subdomains by:

Inserting an NS record for [www.cnn.com](http://www.cnn.com) on the authoritative server for cnn.com

Inserting an A record for cdm.cnn.com into cnn.com

Adding an SOA between ns1.cnn.com and cdm.cnn.com

Root DNS for ".com"

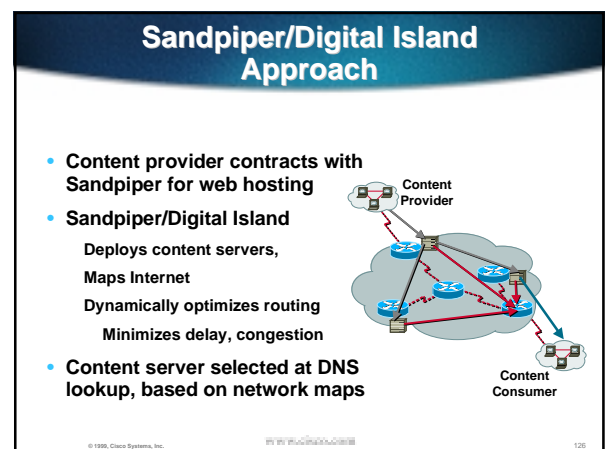
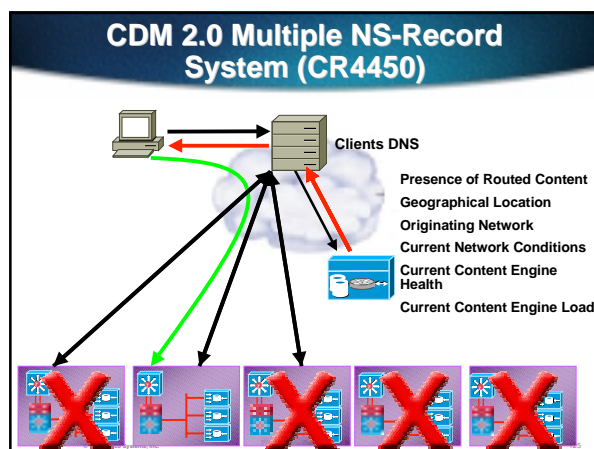

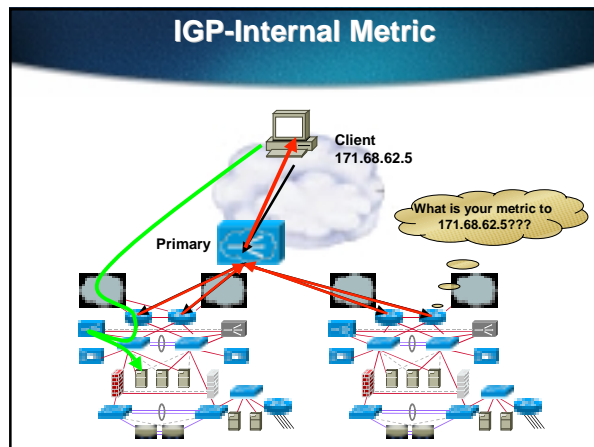
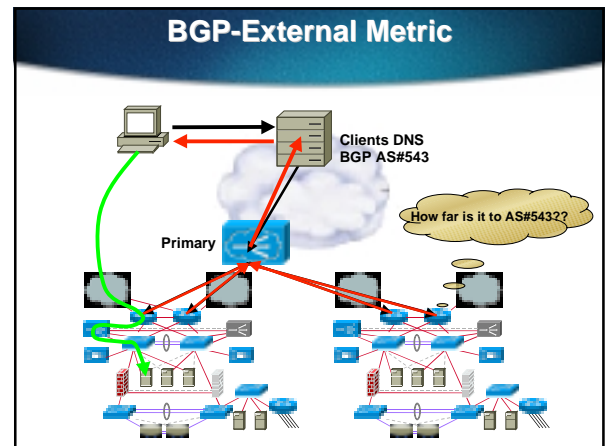
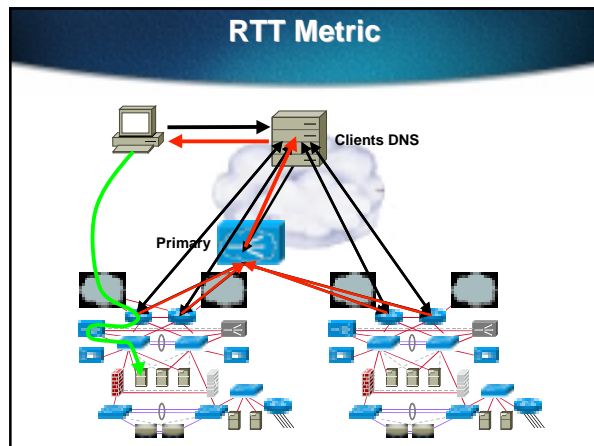
Authoritative DNS for "cnn.com" --ns1.cnn.com

Authoritative DNS for [www.cnn.com](http://www.cnn.com) --cdm.cnn.com

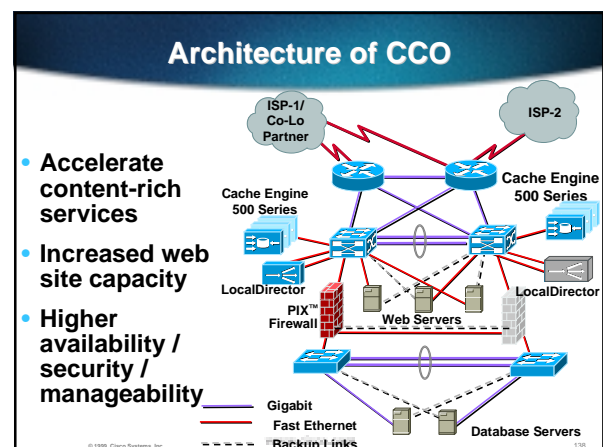
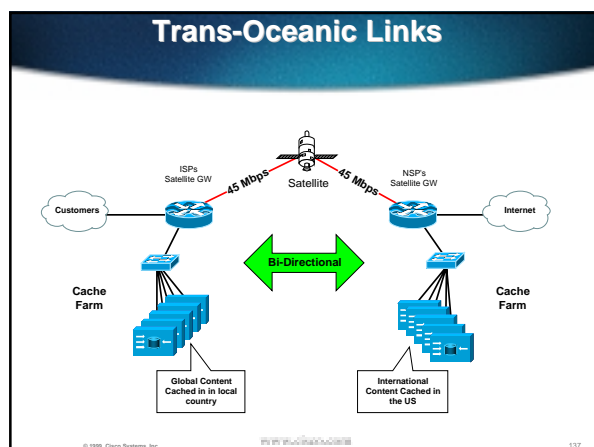
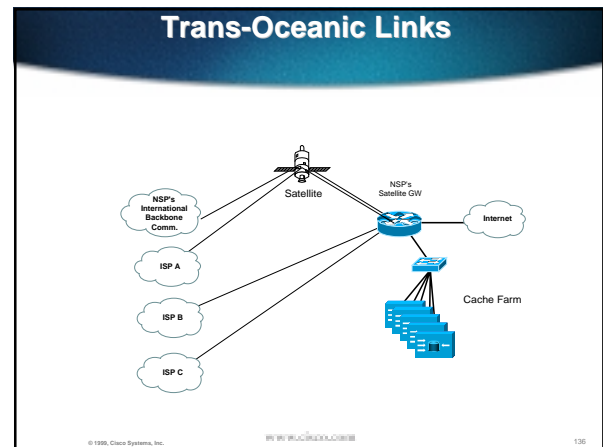
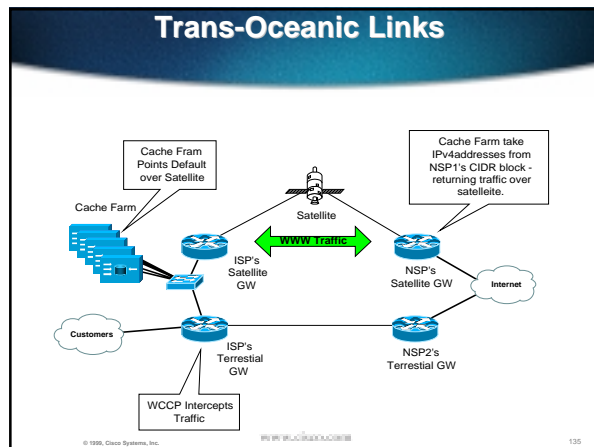
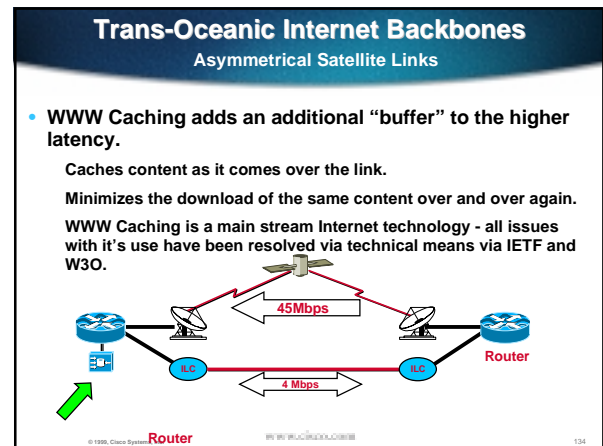

© 1999, Cisco Systems, Inc.

W W W C I S C O . C O M

120









```

CCO cache statistics

cco-cache-1#sh stat http s
Statistics - Savings
-----
Requests                               Bytes
Total:      11364633                    108602153913
Hits:       7998850                      44203825628
Miss:       3289486                      63671295119
Savings:    70.4 %                       40.7 %
cco-cache-1#sh stat http ims
Statistics - If-Modified-Since
-----
Total                                Fresh      %
Client to Cache:      11365078
IMS Received:         2864792
Served from cache:    2427581          2389596  98.4
Cache Miss:          412656          154745   86.0
Revalidated:         24555           0 0.0
Cache to Server:     4169649
IMS Issued:          880193
Due to Client IMS:    24555          24547 100.0
Due to Expiration:    855638          855502 100.0

If-None-Match Requests recv : 1711989
ETag Responses recv        : 3211662
IMM Etag match             : 1537227
IMM Etag mismatch          : 24096
  
```

# CCO effectiveness

---

**“The reliability of CCO improved considerably. Precisely when we turned on the cache engines, the server loads decreased to less than 20% of previous levels.**

**The webmasters/sysadmins downstairs can now concentrate on improving other aspects of CCO than worrying about [previous P1] scaling issues.”**

---

© 1999, Cisco Systems, Inc. [www.cisco.com](http://www.cisco.com)

140

# Cisco branch offices

- Most AsiaPac Cisco offices have 1 x CE550 in them.
- All content can be cached - including content requiring authentication  
("http cache-authenticated")
- Cache passes on authentication/authorisation credentials in a check ..

© 1999, Cisco Systems, Inc. WWW.BUSINESS.CISCO.COM

# Caching of Authenticated Content

- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE

The diagram illustrates the caching of authenticated content. It shows an Intranet cloud connected to a Router running WCCP. The Router is connected to a Cache Cluster. Two desktop computers are connected to the Router. A web browser window is shown on the Cache Cluster, displaying a list of files.

© 1999, Cisco Systems, Inc. [www.cisco.com](http://www.cisco.com) 142

# Caching of Authenticated Content

- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE
- Cache Hit, but force authentication re-check by sending an IMS request to web-server

The diagram illustrates the process of caching authenticated content. A web browser (labeled 'Intranet') sends an 'IMS request' to a 'Router running WCCP'. The router intercepts the flow and redirects it to a 'Cache Cluster'. The cache cluster then sends a response back to the browser. A second browser is shown at the bottom, also connected to the router. A screenshot of a web browser window is shown next to the router, displaying a list of items with checkboxes and a 'Log Out' button.

© 1999, Cisco Systems, Inc. 143


# Caching of Authenticated Content

- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE
- Cache Hit, but force authentication re-check by sending an IMS request to web-server

**Web-server challenges credentials**

© 1999, Cisco Systems, Inc. 148

## Caching of Authenticated Content




- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE
- Cache Hit, but force authentication re-check by sending an IMS request to web-server
- Web-server challenges credentials
- Cache passes challenge back to browser

© 1999, Cisco Systems, Inc.

W W WCCP/Cache/AS5300

145

## Caching of Authenticated Content




- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE
- Cache Hit, but force authentication re-check by sending an IMS request to web-server
- Web-server challenges credentials
- Cache passes challenge back to browser
- Browser/user respond with credentials

© 1999, Cisco Systems, Inc.

W W WCCP/Cache/AS5300

146

## Caching of Authenticated Content




- Connection initiated from web-browser for content stored in cache by another user, Router intercepts flow and redirects it to CE
- Cache Hit, but force authentication re-check by sending an IMS request to web-server
- Web-server challenges credentials
- Cache passes challenge back to browser
- Browser/user respond with credentials
- Cache verifies credentials with web-server (with IMS request). If correct, object is served (from cache)

© 1999, Cisco Systems, Inc.

W W WCCP/Cache/AS5300

147

## iinet example



- 8 x CE550 in 2 WCCP service-groups to cover 38 access-servers
- WCCPv2 on the *access edge* (AS5300's)
- Caches export transaction logs to a host at regular intervals -- volume-based billing
- Traffic to caches is routed via *satellite*. Caches have the *TCP-over-satellite* TCP stack enhancements enabled for maximum "goodput"

© 1999, Cisco Systems, Inc.

W W WCCP/Cache/AS5300

148