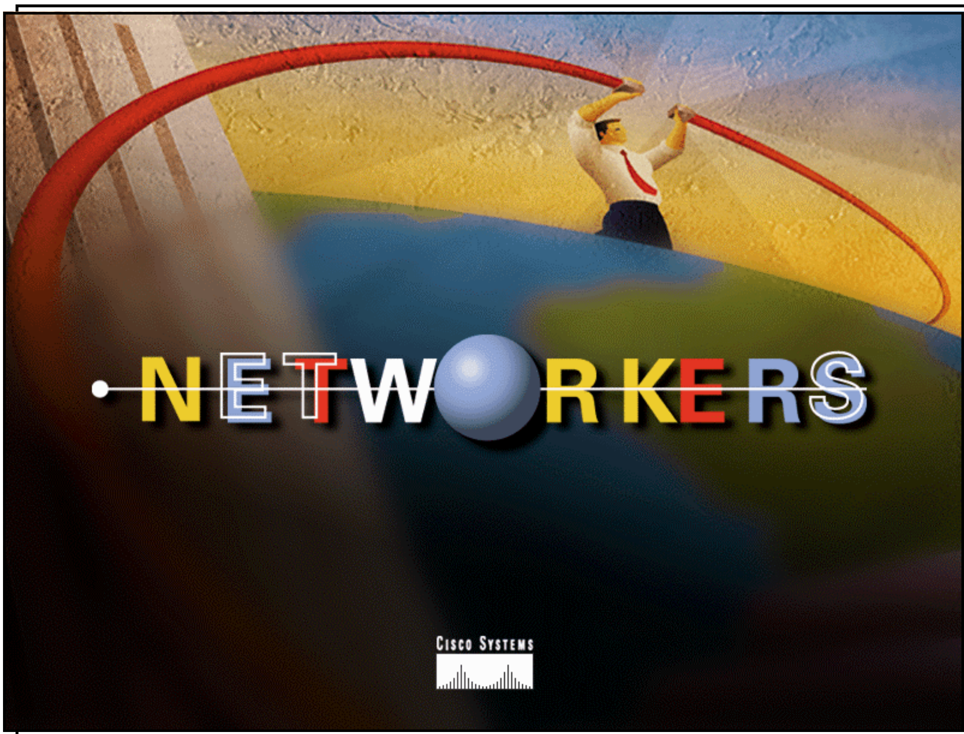
NETWORKERS
Cisco Systems


Deploying BGP:
Scalable, Stable, Simple
Session 309
Cisco Systems

Overview

- Using BGP Attributes
- Deploying IBGP
- Deploying EBGP
- Focus on Stability, Scalability, and Example Configurations

www.cisco.com

Recap of BGP

Why Use It?

www.cisco.com

Bottom Line?

- Implementation of routing policies that are:
  - Scalable
  - Stable
  - Simple (we hope!)

www.cisco.com

More Details…

- You need to scale your IGP
- You're a multihomed ISP customer
- You need to transit full Internet routes
- You need to implement route policy
- Enable QoS policy propagation

www.cisco.com

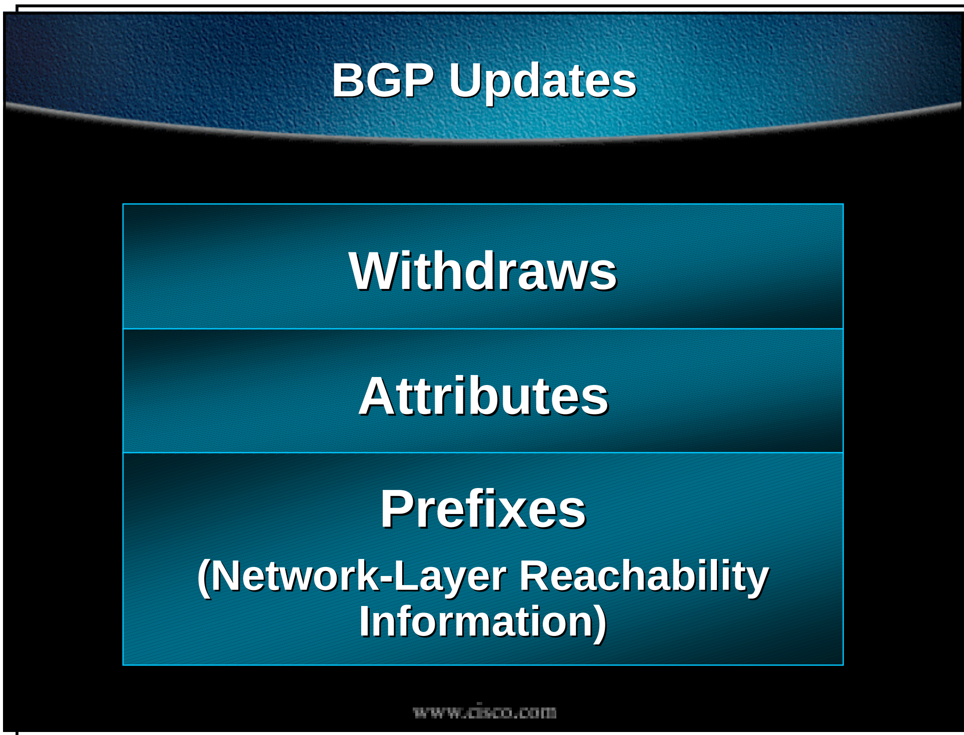
BGP Updates
Withdraws
Attributes
Prefixes
(Network-Layer Reachability Information)
www.cisco.com

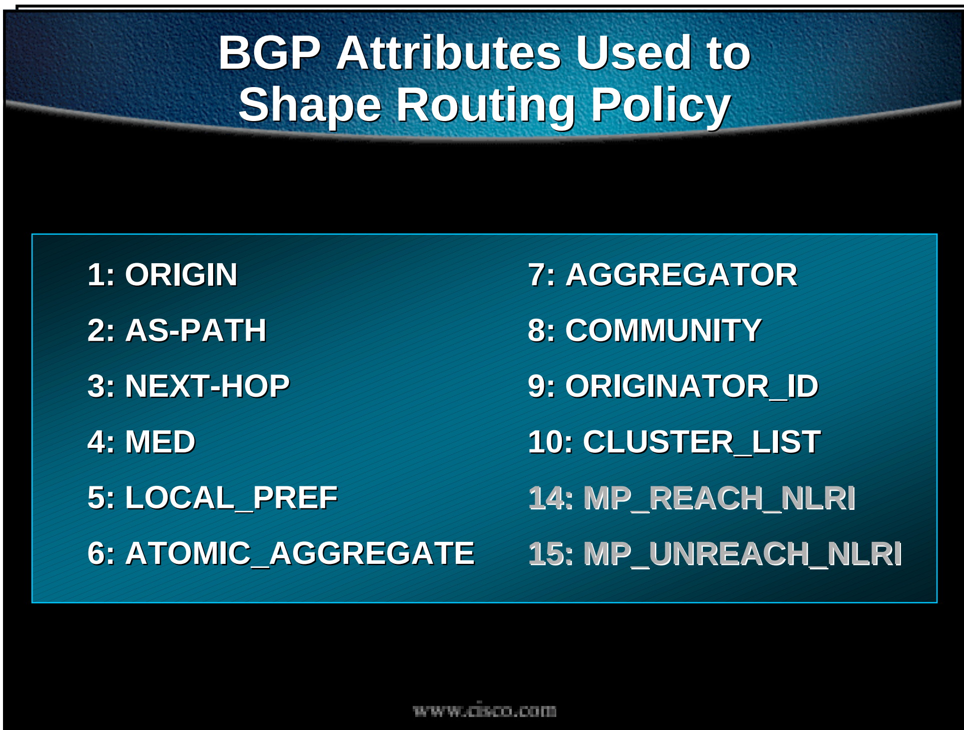
BGP Attributes Used to Shape Routing Policy
1: ORIGIN
2: AS-PATH
3: NEXT-HOP
4: MED
5: LOCAL_PREF
6: ATOMIC_AGGREGATE
7: AGGREGATOR
8: COMMUNITY
9: ORIGINATOR_ID
10: CLUSTER_LIST
14: MP_REACH_NLRI
15: MP_UNREACH_NLRI
www.cisco.com

External BGP

- Between BGP speakers in different AS
- Usually directly connected
- Usually sets next-hop to self

Router B

router bgp 110
neighbor 131.108.10.1 remote-as 109

Router A

router bgp 109
neighbor 131.108.10.2 remote-as 110

neighbor 131.108.10.2 route-map X {in|out}
:
route-map X permit 10
{set | match} *attribute*

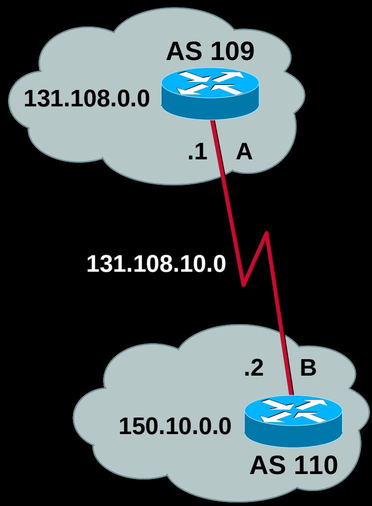

www.cisco.com

Internal BGP

- Neighbor in same AS
- Next-hop unchanged
- May be several hops away
- Don't forward other IBGP routes

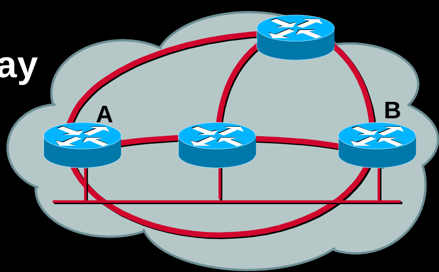

Router B:

router bgp 109
neighbor 131.108.30.2 remote-as 109

Router A:

router bgp 109
neighbor 131.108.20.1 remote-as 109

www.cisco.com

BGP Attributes: Next-Hop

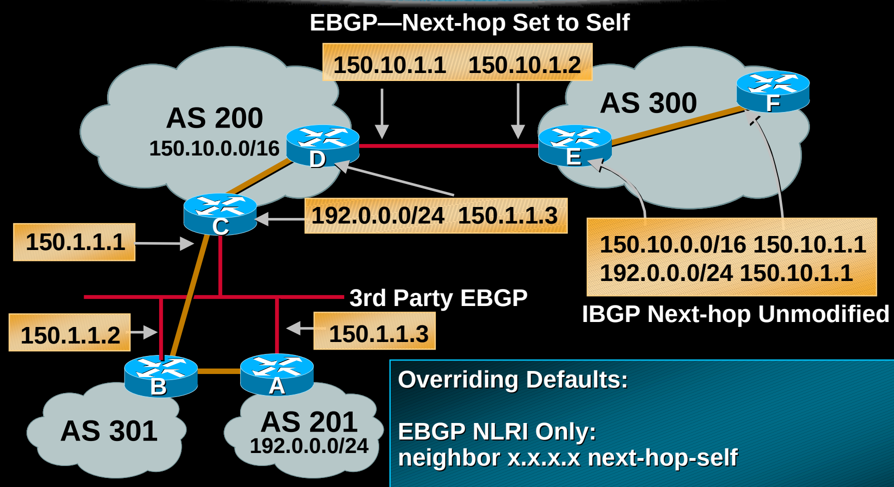

Overriding Defaults:

EBGP NLRI Only:
neighbor x.x.x.x next-hop-self

route-map:
set ip next-hop { A.B.C.D | peeraddress}

www.cisco.com

Problem: Loop Detection, Policy Solution: As-Path

- As sequence

  List of AS's that a route has traversed

- As set

  Summarizes contributing sequence

  Sequence ordering is lost

- route-map prepend: set as-path

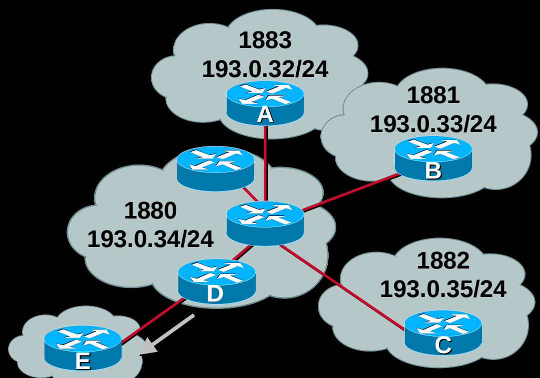

A: 193.0.33/24 1880 1881
B: 193.0.34/24 1880
C: 193.0.32/24 1880 1883
E: 193.0.32/22 1880 {1881, 1882,1883}

www.cisco.com

Problem: Indicate Best Path into AS Solution: MED

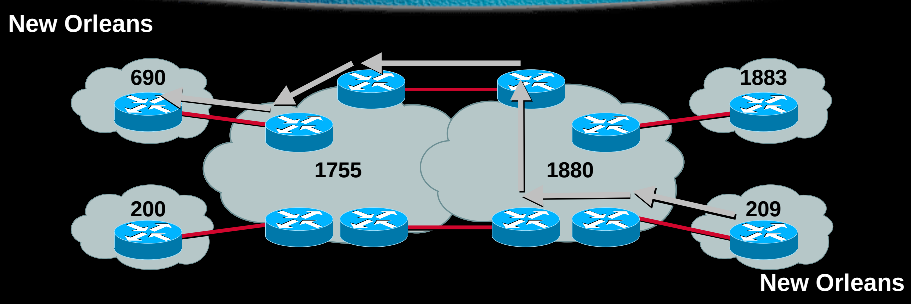

- Conveys the relative preference of entry points
- Comparable if paths are from same AS
  - Unless "bgp always-compare-med" configured
- Non-transitive attribute
- route-map: set metric
  set metric-type internal

www.cisco.com

Problem: Override As-path/MED? Solution: Local Preference

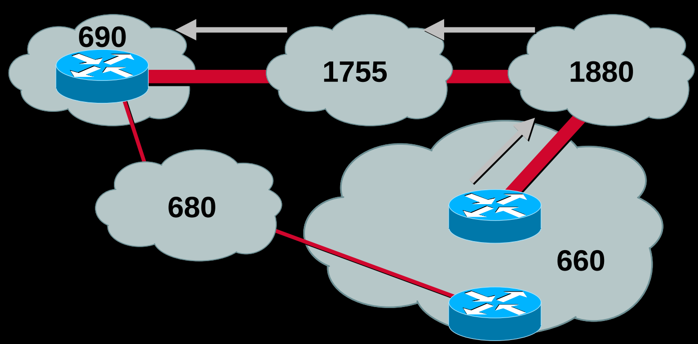

- Attribute local to AS—mandatory for IBGP updates
- route-map: set local-preference

www.cisco.com

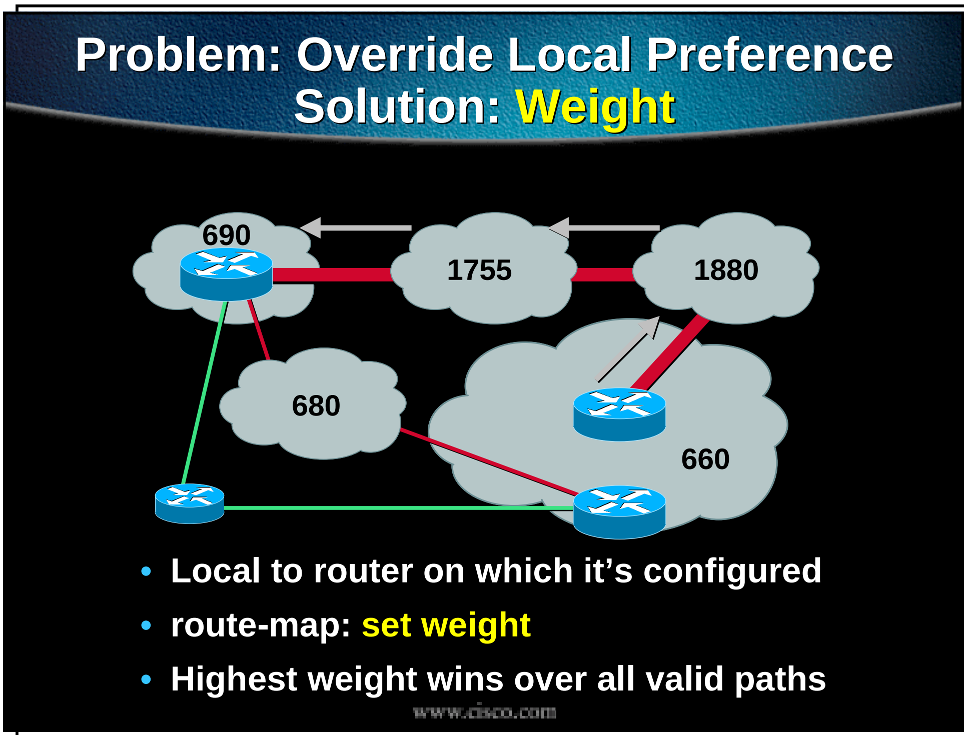
Problem: Override Local Preference
Solution: Weight
690
1755
1880
680
660
• Local to router on which it's configured
• route-map: set weight
• Highest weight wins over all valid paths
www.cisco.com

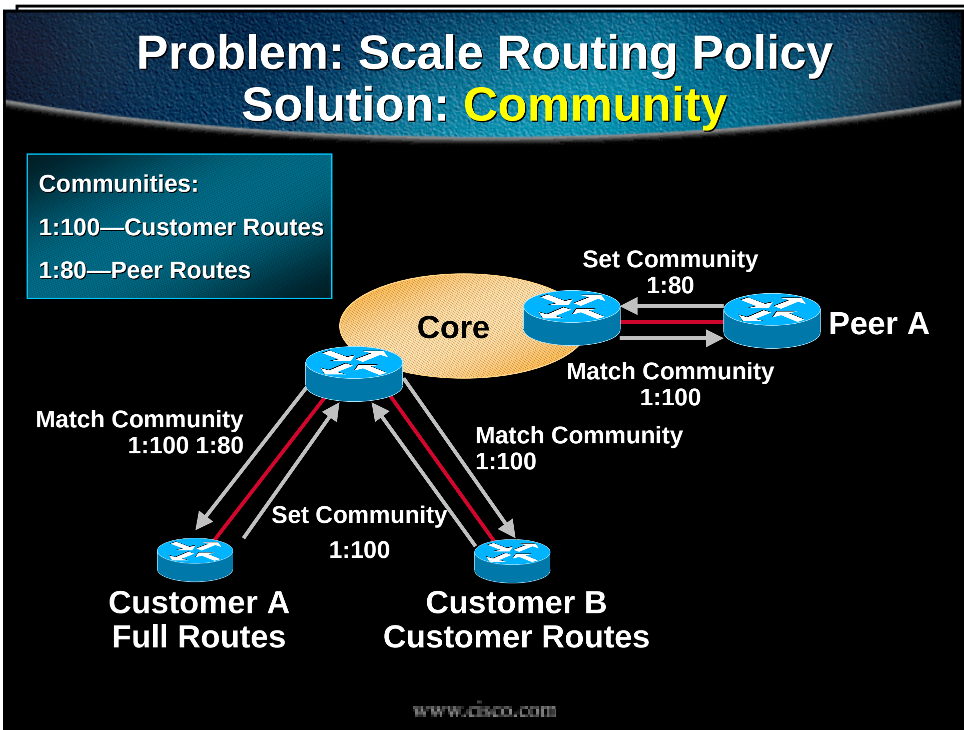
Problem: Scale Routing Policy
Solution: Community
Communities:
1:100—Customer Routes
1:80—Peer Routes
Set Community
1:80
Core
Peer A
Match Community
1:100
Match Community
1:100 1:80
Match Community
1:100
Set Community
1:100
Customer A
Full Routes
Customer B
Customer Routes
www.cisco.com

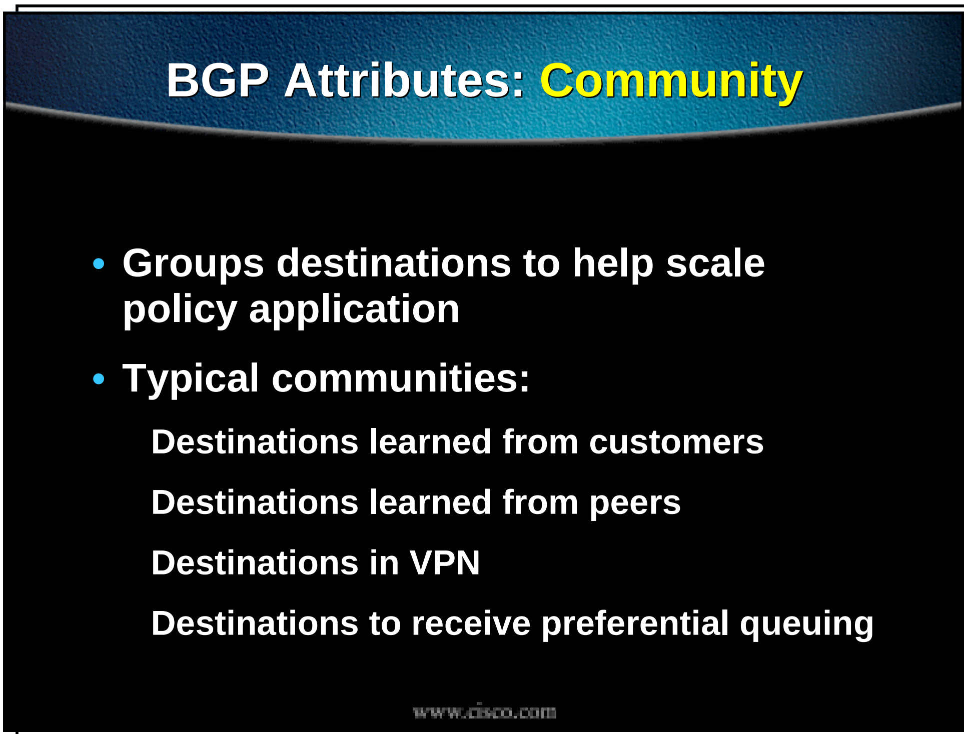
BGP Attributes: Community
• Groups destinations to help scale policy application
• Typical communities:
Destinations learned from customers
Destinations learned from peers
Destinations in VPN
Destinations to receive preferential queuing
www.cisco.com

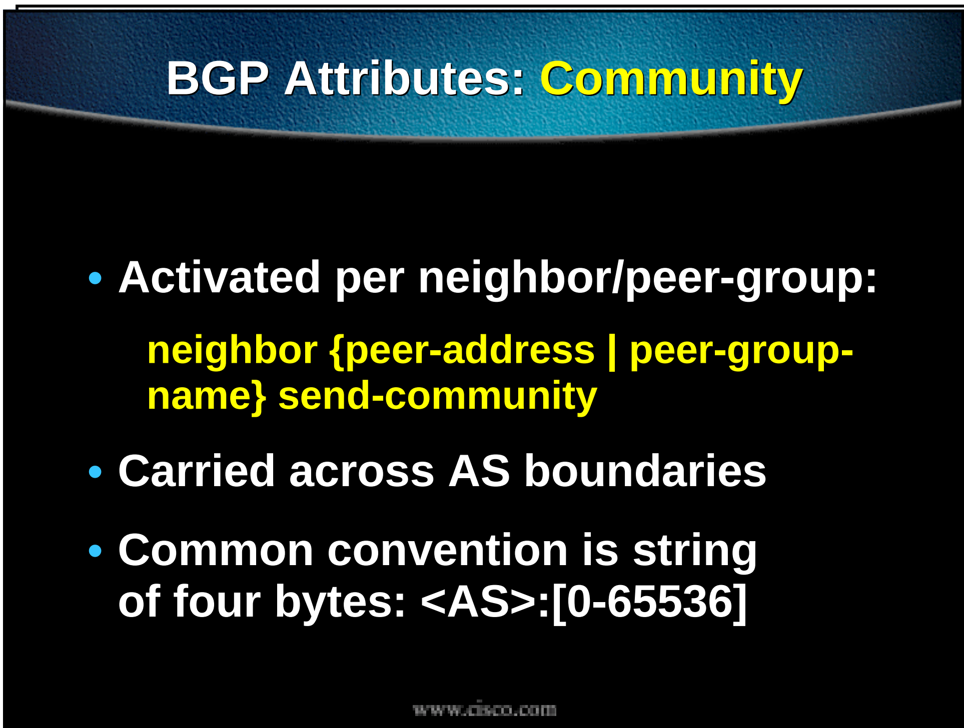
BGP Attributes: Community
• Activated per neighbor/peer-group:
neighbor {peer-address | peer-group-name} send-community
• Carried across AS boundaries
• Common convention is string of four bytes: <AS>:[0-65536]
www.cisco.com

BGP Attributes: Community (Cont.)

- Each destination can be a member of multiple communities
- Using a route-map: **set community**

<1-4294967295>	community number
aa:nn	community number in aa:nn format
additive	Add to the existing community
none	No community attribute
local-AS	Do not send to EBGP peers (well-known community)
no-advertise	Do not advertise to any peer (well-known community)
no-export	Do not export outside AS/confed (well-known community)

www.cisco.com

Least Useful Attribute Award: Origin

- IGP—**network** statement under **router bgp**
- EGP—redistributed from EGP
- Incomplete—redistribute IGP under **router bgp**
- TIP: always use route-map override: **set origin igp**

www.cisco.com

Route Map “Set” Capabilities

• as-path	Prepend string for a BGP AS-path attribute
• comm-list	set BGP community list (for deletion)
• community	BGP community attribute
• dampening	Set BGP route flap dampening parameters
• local-preference	BGP local preference path attribute
• metric	Metric value for destination routing protocol
• origin	BGP origin code
• weight	BGP weight for routing table
• ip next-hop	{ A.B.C.D | peer-address }

www.cisco.com

BGP Attributes

```
75k1#sh ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/24, version 139267814
Paths: (1 available, best #1)
  Not advertised to any peer
  ! AS-PATH                                        AS      ID
  65000 64000 {100 200},  (aggregated by 64000 16.0.0.2)
  !  NEXT-HOP   IGP METRIC        PEER-IP   PEER-ID
    10.0.10.4        (metric 10)    from 10.0.0.1 (10.0.0.2)
      Origin IGP, metric 100, localpref 230, valid, aggregated
      internal (or external or local),
      atomic-aggregate, best
      Community: 64000:3 100:0 200:10
      Originator: 10.0.0.1, Cluster list: 16.0.0.4, 16.0.0.14
```

www.cisco.com

Basic Decision Algorithm

Consider only (synchronized) routes with no AS loops and valid next-hop, then prefer:

- Highest WEIGHT
  - Highest LOCAL PREFERENCE
    - LOCALLY ORIGINATED (eg network/aggregate)
      - Shortest AS-PATH
        - Lowest ORIGIN (IGP < EGP < incomplete)
          - Lowest MED
            - EBGP
              - IBGP
                - Lowest IGP METRIC to next-hop

www.cisco.com

Synchronization

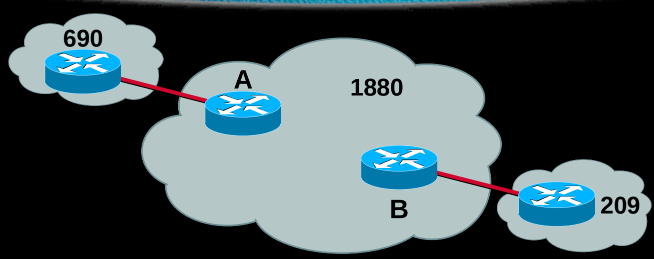

- Router A won't advertise the prefixes from AS209 until the IGP of AS 1880 converges
- Ensure IBGP next-hops are reachable via IGP, then:

```
router bgp 1880
no synchronization
```

www.cisco.com

General Issues

- Synchronization: not required if you have a full IBGP mesh
- => EBGP distance should be less than IGP distance
- auto-summary: use aggregation commands instead

```
router bgp 100
no synchronization
no auto-summary
distance 200 200 200
```

www.cisco.com

So Far…

- Apply policy on a per AS basis
- Group routes into even large communities
- Choose exit and entry points for large policy groups using med/local preference

Will Your Policies Scale?

www.cisco.com

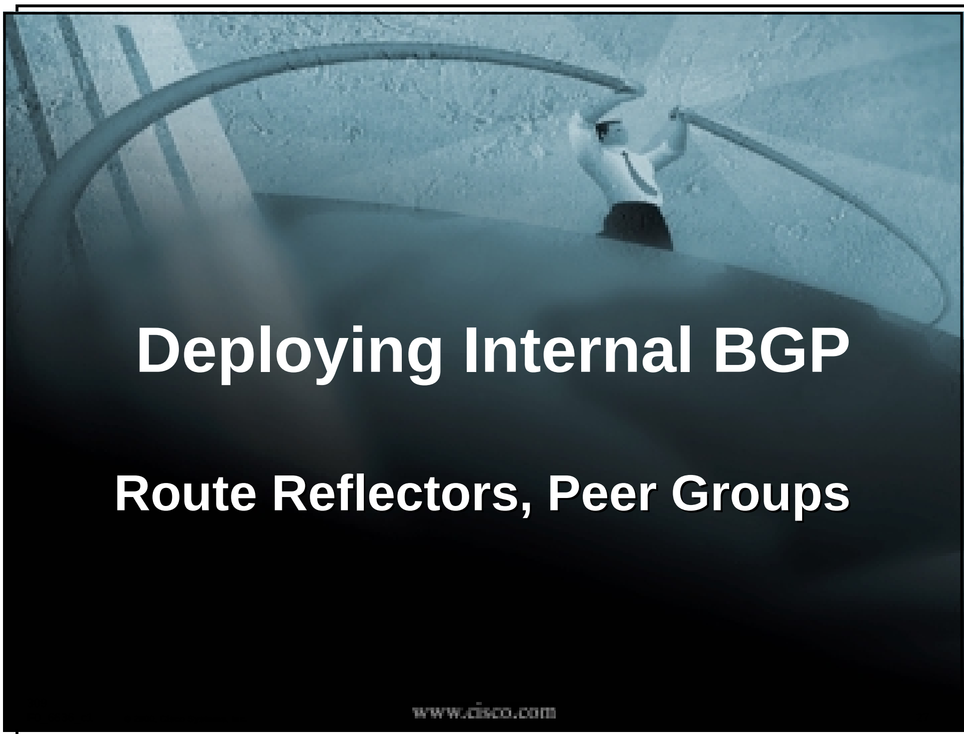
Deploying Internal BGP
Route Reflectors, Peer Groups
www.cisco.com

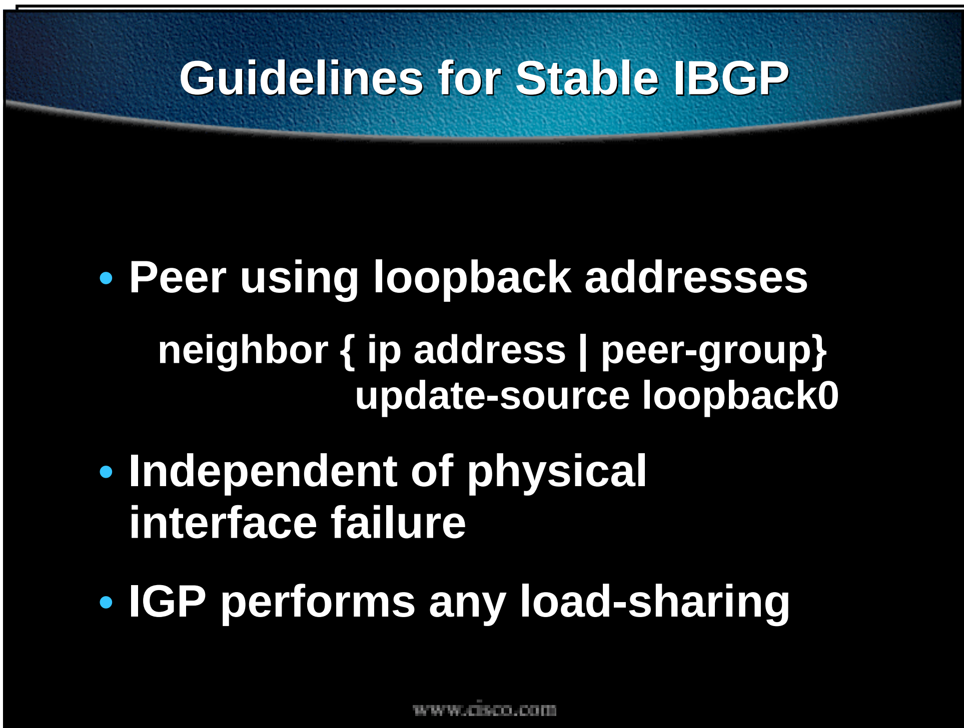
Guidelines for Stable IBGP
• Peer using loopback addresses
neighbor { ip address | peer-group}
update-source loopback0
• Independent of physical interface failure
• IGP performs any load-sharing
www.cisco.com

Guidelines for Scaling IBGP

- Use peer groups and RRs
- Carry only next-hops in IGP
- Carry full routes in BGP only if necessary
- Do not redistribute BGP into IGP

Using Peer Groups


```
router bgp  100
neighbor internal peer-group
neighbor internal description ibgp peers
neighbor internal remote-as 100
neighbor internal update-source Loopback0
neighbor internal next-hop-self
neighbor internal send-community
neighbor internal version 4
neighbor internal password 7 03085A09
neighbor 10.0.0.1 peer-group internal
neighbor 10.0.0.2 peer-group internal
```

What Is a Peer Group?

- All peer-group members have a common outbound policy
- Updates generated once per peer group
- Simplifies configuration
- Members can have different inbound policy

Why Route Reflectors?

Avoid n(n-1)/2 IBGP mesh

n=1000 => Nearly Half a Million IBGP Sessions!

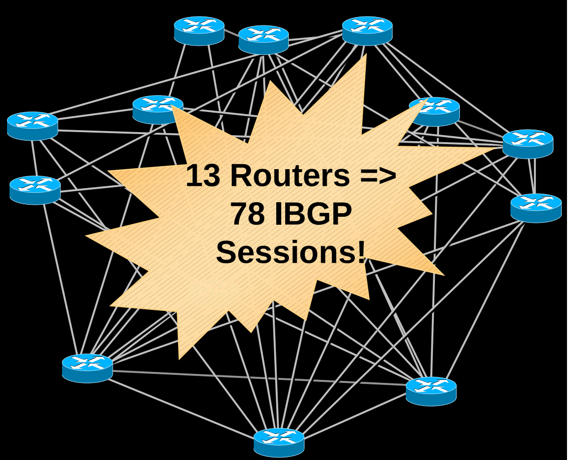

Using Route Reflectors

Golden Rule of RR Loop Avoidance: RR Topology Should Follow Physical Topology

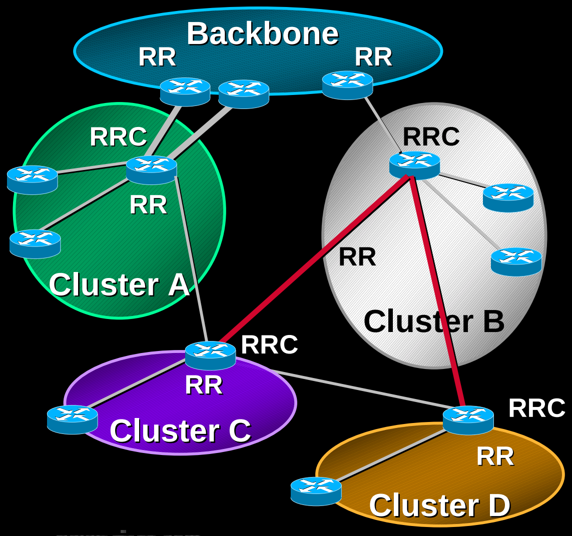

www.cisco.com

What Is a Route Reflector?

- Reflector receives path from clients and non clients
- If best path is from a client, reflect to clients and non-clients
- If best path is from a non-client, reflect to clients

www.cisco.com

Configuration of RR Peer Groups

```
router bgp  100
neighbor rr-client peer-group
neighbor rr-client description RR clients
neighbor rr-client remote-as 100
neighbor rr-client update-source Loopback0
neighbor rr-client route-reflector-client
neighbor rr-client next-hop-self
neighbor rr-client send-community
neighbor rr-client version 4
neighbor rr-client password 7 03085A09
neighbor 10.0.1.1 peer-group rr-client
neighbor 10.0.2.2 peer-group rr-client
```

This line on RRs only RRCs use still use internal peer group

www.cisco.com

Deploying Route Reflectors

- Divide backbone into multiple clusters
- Each cluster contains at least one RR (multiple for redundancy), and multiple clients
- RRs are fully meshed via IBGP
- Still use single IGP—next-hop unmodified by RR

www.cisco.com

Hierarchical Route Reflector

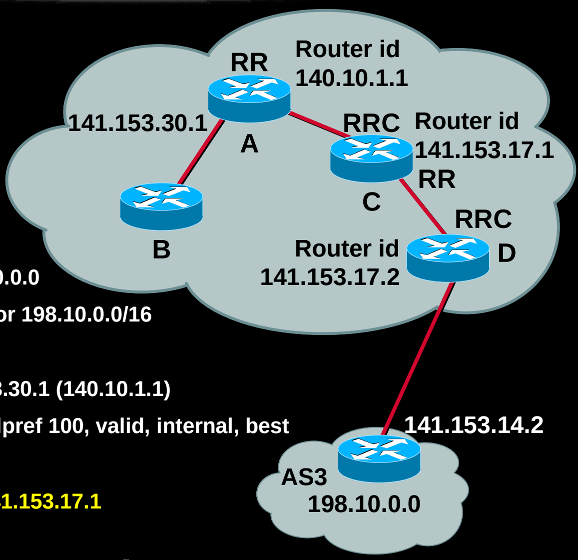

- **Example:**

  RouterB>sh ip bgp 198.10.0.0

  BGP routing table entry for 198.10.0.0/16

  3

  141.153.14.2 from 141.153.30.1 (140.10.1.1)

  Origin IGP, metric 0, localpref 100, valid, internal, best

  Originator: 141.153.17.2

  Cluster list: 144.10.1.1, 141.153.17.1

www.cisco.com

BGP Attributes: ORIGINATOR_ID

- **ORIGINATOR_ID**

  Router ID of IBGP speaker that reflects RR client routes to non-clients

  Overridden by: bgp cluster-id x.x.x.x

- **Useful for troubleshooting and loop detection**

www.cisco.com

BGP Attributes: CLUSTER_LIST

- CLUSTER_LIST

  String of ORIGINATOR_IDs through which the route has passed

- Useful for troubleshooting and loop detection

www.cisco.com

So Far…

- Is IBGP peering **S**table?

  Use loopbacks for peering

- Will it **S**cale?

  Use peer groups

  Use route reflectors

- **S**imple, hierarchical config?

www.cisco.com

Deploying External BGP

Customer Issues
ISP Issues—Optional

www.cisco.com

Customer Issues

- Steps

  Configure BGP (use session passwords!)

  Generate a stable aggregate

  Set inbound policy

  Set output policy

  Configure loadsharing/multihoming

www.cisco.com

Connecting to an ISP

- AS 100 is a customer of AS 200 (usually directly connected)

```
Router B:
router bgp 100
aggregate-address 10.60.0.0 255.255.0.0 as-set summary-only
neighbor 10.200.0.1 remote-as 200
neighbor 10.200.0.1 description ISP connection
neighbor 10.200.0.1 remove-private-AS
neighbor 10.200.0.1 version 4
neighbor 10.200.0.1 prefix-list ispout out
neighbor 10.200.0.1 route-map ispout out
neighbor 10.200.0.1 route-map ispin in
neighbor 10.200.0.1 password 7 020A0559
neighbor 10.200.0.1 maximum-prefix 65000
```

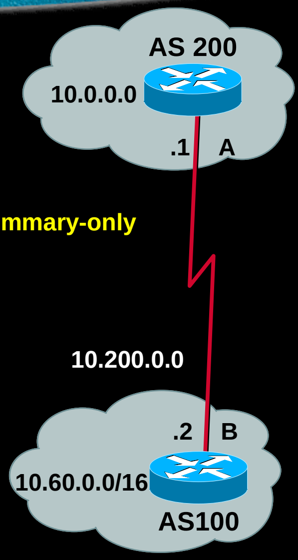

www.cisco.com

What Is Aggregation?

- Summarization based on specifics **from the BGP routing table**

  10.60.1.0 255.255.255.0

  10.60.2.1 255.255.255.240

  => 10.60.0.0 255.255.0.0

www.cisco.com

How to Aggregate

- Aggregate-address 10.60.0.0 255.255.0.0 {as-set} {summary-only} {route-map}
- Use as-set to include path and community info from specifics
- Summary-only suppresses specifics
- Route-map sets other attributes

www.cisco.com

Why Aggregate?

- Reduce number of Internet prefixes
- Increase stability—aggregate stays even specifics come and go
- Stable aggregate generation:

```
router bgp 100
aggregate-address 10.60.0.0 255.255.0.0 as-set summary-only
network 10.1.0.0 255.255.0.0
ip route 10.1.0.0 255.255.0.0 null0 254
```

www.cisco.com

BGP Attributes Atomic Aggregate

- Indicates loss of AS-PATH information
- Must not be removed once set
- Set by: aggregate-address x.x.x.x
- Not set if as-set keyword is used, however, AS-SET and COMMUNITY then carries information about specifics

www.cisco.com

BGP Attributes: Aggregator

- AS number and IP address of router generating aggregate
- Useful for troubleshooting

www.cisco.com

Aggregate Attributes

NEXT_HOP = local (0.0.0.0)

WEIGHT = 32768

LOCAL_PREF = none (assume 100)

AS_PATH = AS_SET or nothing

ORIGIN = IGP

MED = none

www.cisco.com

Why Inbound Policy?

- Apply a recognizable community to use in outbound filters or other policy
- Possibly adjust local-preference to override default of 100
- Multihoming loadsharing—more later

```
route-map ISPin permit 10
set local-preference 200
set community 100:2   ; routes from ISP
```

www.cisco.com

Why Outbound Policy?

- Main filter based on communities
- Adding a prefix filter helps protect against mistakes (can apply as-path filters too)
- Send community based on agreements with ISP (remember to add send-community line to config)
- Multihoming loadsharing policy

www.cisco.com

Outgoing Policy Config

```
ip prefix-list ISPout seq 5 permit 10.60.0.0 255.255.0.0
:
ip community-list 1 permit 100:1:
:
route-map ISPout permit 10
match community 1    ;  Internet transit community
set community 100:3 additive    ; something agreed with ISP
```

www.cisco.com

Load-Sharing—Single Path

Router A:
interface loopback 0
ip address 10.60.0.1 255.255.255.255
!
router bgp 100
neighbor 10.200.0.2 remote-as 200
neighbor 10.200.0.2 update-source loopback0
neighbor 10.200.0.2 ebgp-multi-hop 2

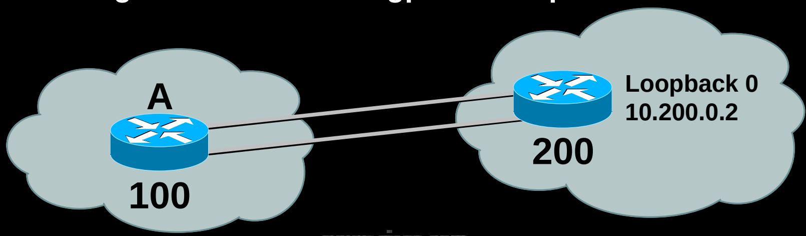

Load-Sharing—Multiple Paths from Same AS

Router A:
router bgp 100
neighbor 10.200.0.1 remote-as 200
neighbor 10.100.0.1 remote-as 200
maximum-paths 2 ; can configure up to 6

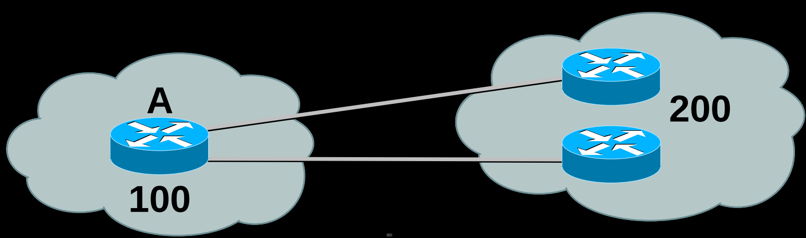

BGP over Simplex Links

- Many Internet link utilisations are asymmetric:
  - Data centers
  - International links

www.cisco.com

BGP over Simplex Links

- Two approaches:
  - EBGP—connecting to an upstream provider
  - IBGP—connecting to your own router at international exchange

www.cisco.com

EBGP over Simplex Links
(www.telstra.net/ops/satellite/satellite/satellite.html)

Router A:
interface loopback 0
ip address 100.0.0.2 255.255.255.255
router bgp 100
neighbor 200.0.0.2 remote-as 200
neighbor 200.0.0.2 update-source loopback0
neighbor 200.0.0.2 ebgp-multihop 255
neighbor 200.0.0.2 route-map set-nh in
neighbor 200.0.0.2 timers <keepalive> <holdtime>
ip route 200.0.0.2 255.255.255.255 hssi1/0

route-map set-nh permit 10
set ip next-hop peer-address

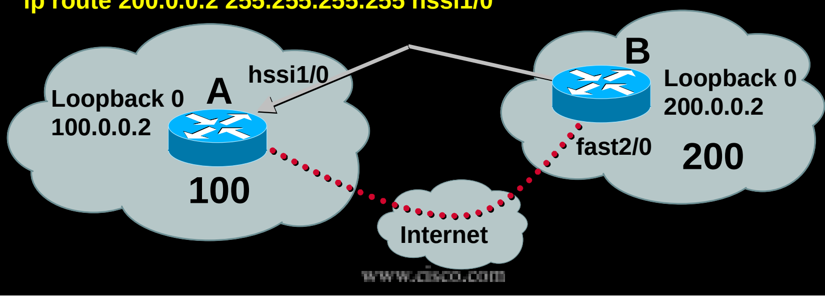

EBGP over Simplex Links
(www.telstra.net/ops/satellite/satellite/satellite.html)

Router B:
interface loopback 0
ip address 200.0.0.2 255.255.255.255
router bgp 200
neighbor 100.0.0.2 remote-as 200
neighbor 100.0.0.2 update-source loopback0
neighbor 100.0.0.2 ebgp-multihop 255
neighbor 100.0.0.2 route-map set-nh in
neighbor 100.0.0.2 timers <keepalive> <holdtime>
ip route 100.0.0.2 255.255.255.255 FastEthernet2/0

route-map set-nh permit 10
set ip next-hop peer-address

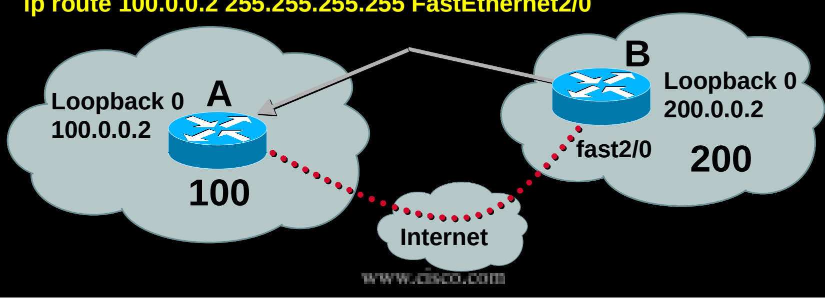

IBGP over Simplex Links

LOCAL ROUTER:
interface loopback 0
ip address 100.0.0.2 255.255.255.255
router bgp 100
neighbor 200.0.0.2 remote-as 200
neighbor 200.0.0.2 update-source loopback0
neighbor 200.0.0.2 route-map set-nh in
ip route 200.0.0.2 255.255.255.255 serial2/0
ip route 200.0.0.2 255.255.255.255 serial2/1

route-map set-nh permit 10
set ip next-hop peer-address

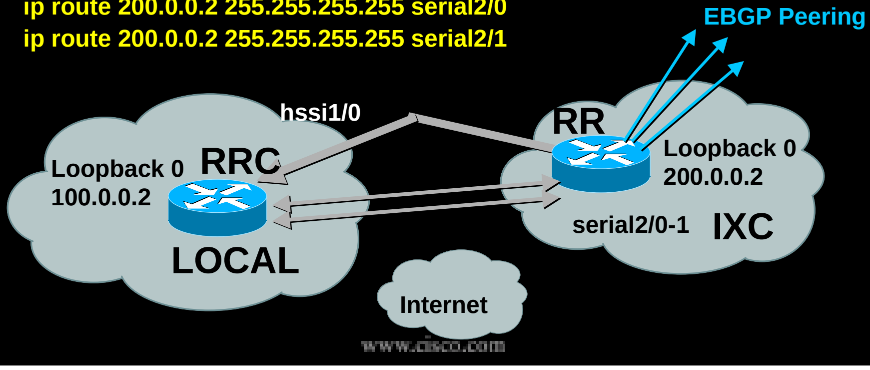

IBGP over Simplex Links

IXC ROUTER:
interface loopback 0
ip address 200.0.0.2 255.255.255.255
router bgp 100
neighbor 100.0.0.2 remote-as 200
neighbor 100.0.0.2 update-source loopback0
neighbor 100.0.0.2 route-reflector-client
neighbor 100.0.0.2 route-map set-nh in
ip route 100.0.0.2 255.255.255.255 hssi1/0

route-map set-nh permit 10
set ip next-hop peer-address

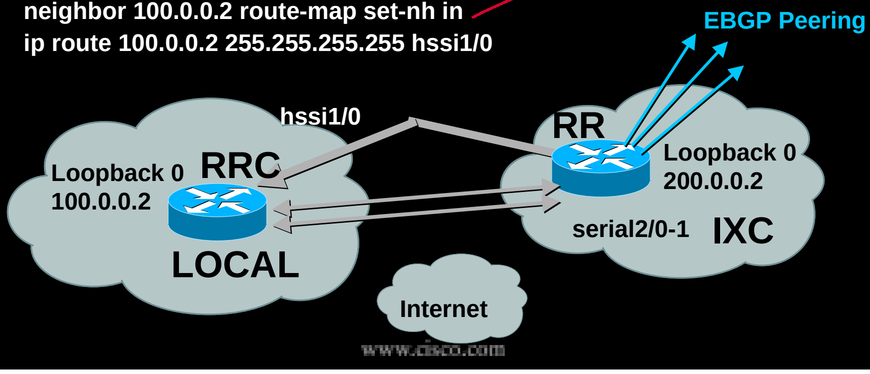

What Is Multihoming?

- Connecting to two or more ISPs to increase:

  Reliability—one ISP fails, still OK

  Performance—better paths to common Internet destinations

www.cisco.com

Types of Multihoming

- Three common cases:

  Default from all providers

  Customer+default routes from all

  Full routes from all

www.cisco.com

Default from All Providers

- Low memory/CPU solution
- Provider sends BGP default => provider decided by IGP metrics to reach default
- You send all your routes to provider =>inbound path decided by Internet

  You can influence using AS-path prepend

www.cisco.com

Default from All Providers

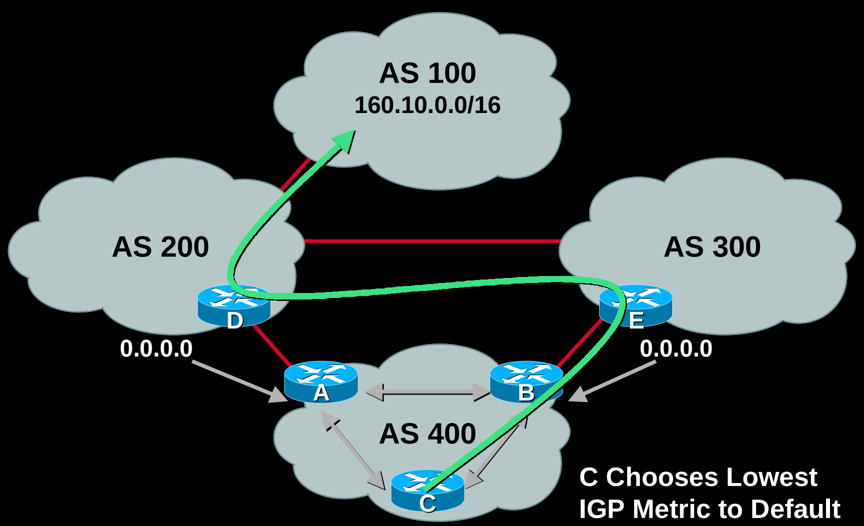

www.cisco.com

Customer+Default from All Providers

- Medium memory and CPU
- “Best” path—usually shortest AS-path
- Use local-preference to override based on prefix, as-path, or community
- IGP metric to default used for all other destinations

Customer Routers from All Providers

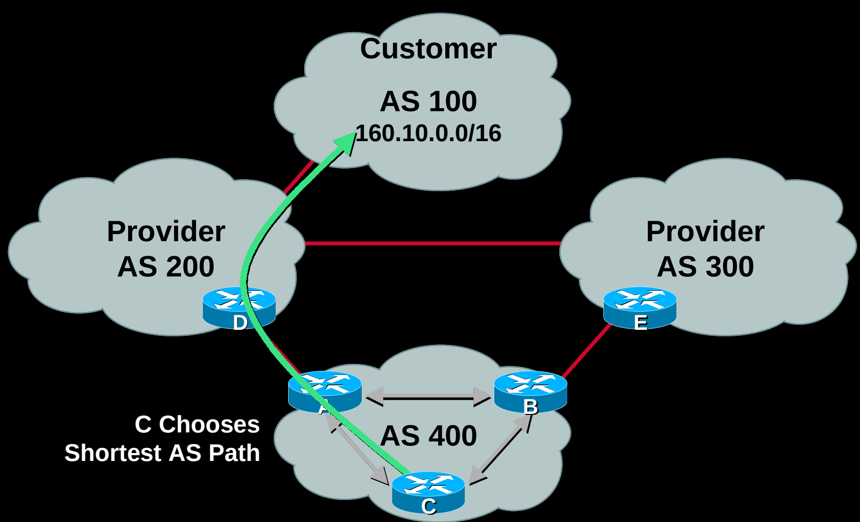

Customer Routes from All Providers

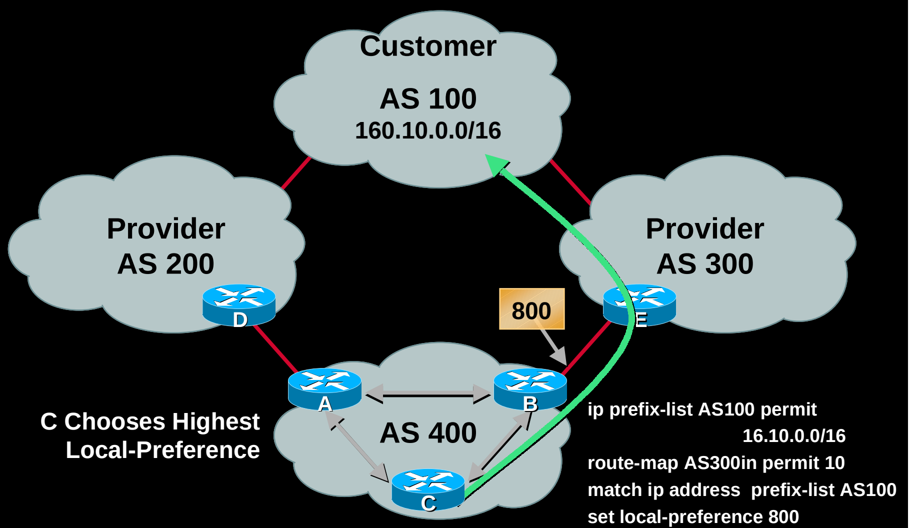

www.cisco.com

Full Routes from Both Providers

- Higher memory/CPU solution
- Reach all destinations by “best” path—usually shortest AS path
- Can still manually tune using local-pref and as-path/community/prefix matches

www.cisco.com

Full Routes from All Providers

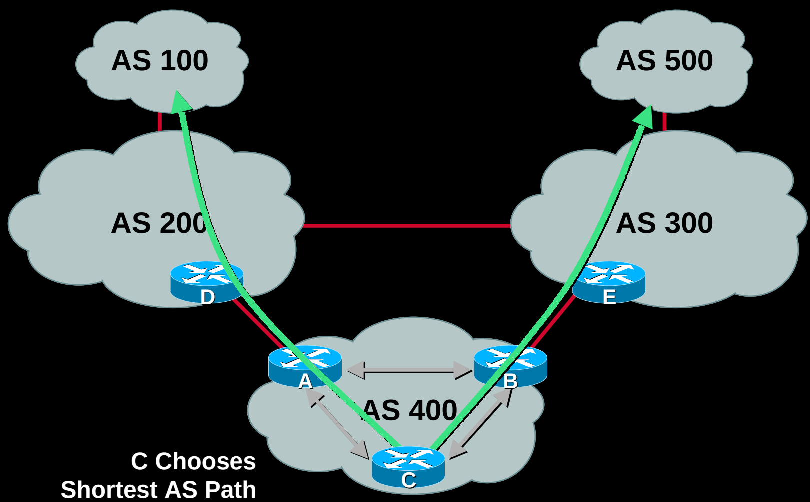

Controlling Inbound Traffic?

- Inbound is very difficult due to lack of transitive metric
- Can divide outgoing updates across providers, but what happens to redundancy?

Controlling Inbound Traffic? (Cont.)

- Bad Internet citizen:
  - Divide address space
  - Set as-path prepend
- Good Internet citizen
  - Divide address space
  - Use "advertise maps"

www.cisco.com

Using AS-PATH Prepend

10.1.0.0/16 300 400 400
10.1.0.0/16 200 400 (best)
10.2.0.0/16 300 400 (best)
10.2.0.0/16 200 400 400

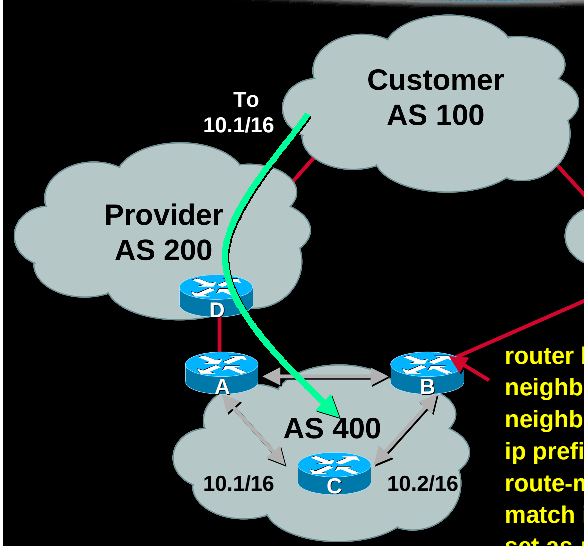


```
router bgp 400
neighbor 20.0.0.1 remote-as 300
neighbor 20.0.0.1 route-map AS300 out
ip prefix-list AS100 permit 10.1.0.0/16
route-map AS300out permit 10
match ip address prefix-list AS100
set as-path prepend 400
```

www.cisco.com

Using an Advertise-Map

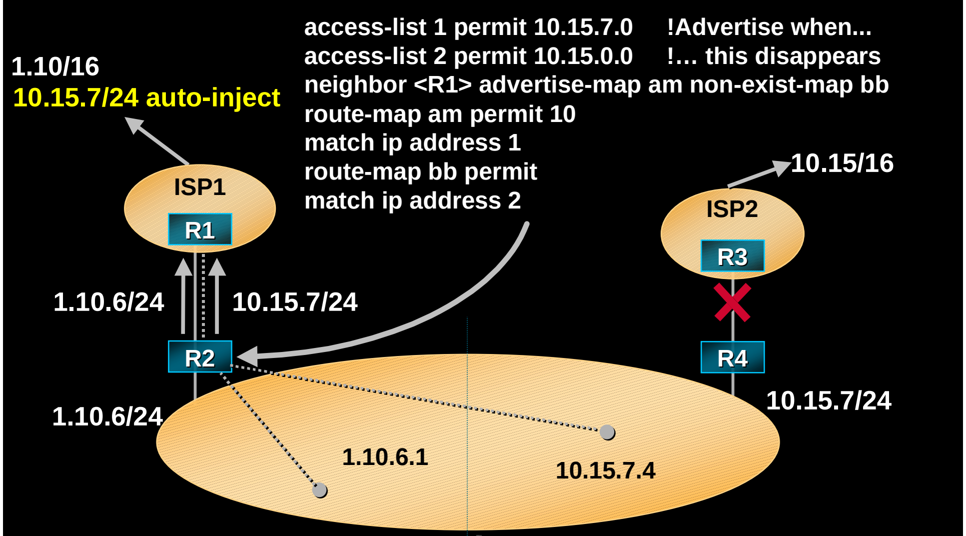

So Far…

- Stability through:
  - Aggregation
  - Multihoming
  - Inbound/outbound filtering
- Scalability of memory/CPU:
  - Default, customer routes, full routes
- Simplicity using “standard” solutions

Summary

- **Scalability:**
  - Use attributes, especially community
  - Use peer groups and route reflectors
- **Stability:**
  - Use loopback addresses for IBGP
  - Generate aggregates
  - Apply passwords
  - Always filter inbound and outbound

www.cisco.com

Summary (Cont.)

- **Simplicity—standard solutions:**
  - Three multihoming options
  - Group customers into communities
  - Apply standard policy at the edge
  - Avoid “special configs”
  - Script your config generation

www.cisco.com

For Further Reference:

- Case studies on www.cisco.com
  - http://www.cisco.com/warp/public/459/18.html
- Cisco Press
  - "Internet Routing Architectures"
  - "Advanced IP Network Design"
  - "Large-Scale IP Network Solutions"

www.cisco.com

For Further Reference:

- John Stewart, BGP4, Addison Wesley
- Geoff Huston, ISP Survival Guide,Wiley

www.cisco.com

Thank You!

- **Related presentations:**

  Advanced BGP and Troubleshooting

  Session 317

- **Questions?**

www.cisco.com

Cisco Systems

Empowering the Internet Generation[SM]

www.cisco.com

ISP Issues

- Scale BGP customer aggregation
- Offer a choice of route-feeds
- Peer with other providers
- Minimize BGP activity and protect against customer misconfigurations
- Provide a backup service
- Propagate QoS policy

www.cisco.com

Customer Aggregation Guidelines

- Define at least three peer groups:
  - cust-default—send default route only
  - cust-cust—send customer routes only
  - cust-full—send full Internet routes
- Identify routes via communities
  - 2:100=customers; 2:80=peers
- Apply passwords and an inbound prefix-list on a per neighbor basis

Customer Aggregation

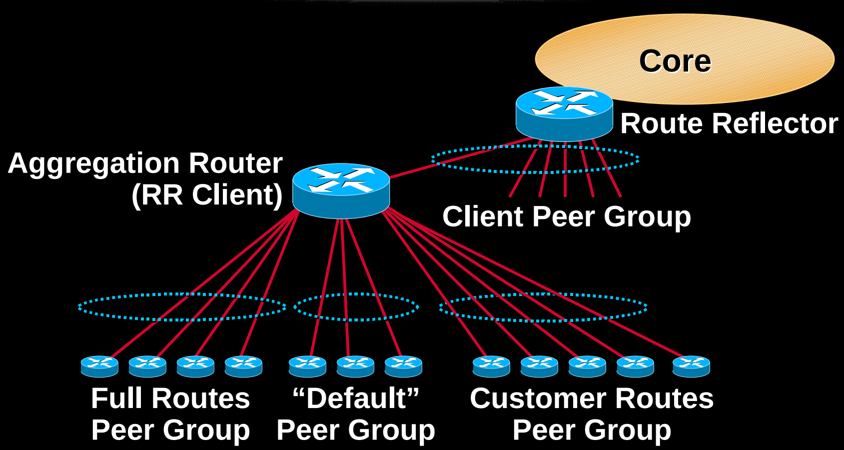

TIP: Apply Passwords and Inbound Prefix-List to **Each** Customer

cust-full peer-group

```
neighbor cust-full peer-group
neighbor cust-full description Send full Routes
neighbor cust-full remove-private-AS
neighbor cust-full version 4
neighbor cust-full route-map cust-in in
neighbor cust-full prefix-list cidr-block out
neighbor cust-full route-map full-routes out
.
ip prefix-list cidr-block seq 5 deny 10.0.0.0/8 ge 9
ip prefix-list cidr-block seq 10 permit 0.0.0.0/0 le 32
```

www.cisco.com

cust-full outgoing route-map

```
ip community-list 1 permit 2:100
ip community-list 80 permit 2:80
.
route-map full-routes permit 10
match community 1 80      ; customer & peers
set metric-type internal    ; MED = IGP metric
set ip next-hop peer-address ; our own
```

www.cisco.com

cust-in route-map

route-map cust-in permit 10

set metric 4294967294 ; ignore MED

set ip next-hop peer-address

set community 2:100 additive

cust-cust peer-group

neighbor cust-cust peer-group

neighbor cust-cust description customer routes

neighbor cust-cust remove-private-AS

neighbor cust-cust version 4

neighbor cust-cust route-map cust-in in

neighbor cust-cust prefix-list cidr-block out

neighbor cust-cust route-map cust-routes out

www.cisco.com

cust-routes route-map

route-map cust-routes permit 10

match community 1 ; customers only

set metric-type internal ; MED = igp metric

set ip next-hop peer-address ; our own

www.cisco.com

default route peer-group

neighbor cust-default peer-group
neighbor cust-default description Send default
neighbor cust-default default-originate
route-map default-route
neighbor cust-default remove-private-AS
neighbor cust-default version 4
neighbor cust-default route-map cust-in in
neighbor cust-default prefix-list deny-all out

ip prefix-list deny-all seq 5 deny 0.0.0.0/0 le 32

www.cisco.com

default-route route-map

route-map default-route permit 10

set metric-type internal ; MED = igp metric

set ip next-hop peer-address ; our own

www.cisco.com

Peer Groups for NAPs

- Similar to EBGP customer aggregation except inbound prefix filtering is rarely used (lack of registry)
- Use maximum-prefix and prefix sanity checking instead
- Still use per-neighbor passwords!

www.cisco.com

Peer Groups for NAPs (Cont.)

neighbor nap peer-group

neighbor nap description for peer ISPs

neighbor nap remove-private-AS

neighbor nap version 4

neighbor nap prefix-list sanity-check in

neighbor nap prefix-list cidr-block out

neighbor nap route-map nap-out out

neighbor nap maximum prefix 30000

www.cisco.com

Peer Groups for NAPs (Cont.)

route-map nap-out permit 10

match community 1 ; customers only

set metric-type internal ; MED = IGP metric

set ip next-hop peer-address ; our own

www.cisco.com

Peer Groups for NAPs: Sanity-Check Prefix-List

```
# FIRST - FILTER OUT YOUR IGP ADDRESS SPACE!!
ip prefix-list sanity-check seq 5 deny 0.0.0.0/32
# deny the default route
ip prefix-list sanity-check seq 10 deny 0.0.0.0/8 le 32
# deny anything beginning with 0
ip prefix-list sanity-check seq 15 deny 0.0.0.0/1 ge 20
# deny masks > 20 for all class A nets (1-127)
ip prefix-list sanity-check seq 20 deny 10.0.0.0/8 le 32
# deny 10/8 per RFC1918
ip prefix-list sanity-check seq 25 deny 127.0.0.0/8 le 32
# reserved by IANA - loopback address
ip prefix-list sanity-check seq 30 deny 128.0.0.0/2 ge 17
deny masks >= 17 for all class B nets (129-191)
ip prefix-list sanity-check seq 35 deny 128.0.0.0/16 le 32
# deny net 128.0 - reserved by IANA
ip prefix-list sanity-check seq 40 deny 172.16.0.0/12 le 32
# deny 172.16 as RFC1918
```


Peer Groups for NAPs: Sanity-Check Prefix-List

```
ip prefix-list sanity-check seq 45 deny 192.0.2.0/24 le 32
# class C 192.0.20.0 reserved by IANA
ip prefix-list sanity-check seq 50 deny 192.0.0.0/24 le 32
# class C 192.0.0.0 reserved by IANA
ip prefix-list sanity-check seq 55 deny 192.168.0.0/16 le 32
# deny 192.168/16 per RFC1918
ip prefix-list sanity-check seq 60 deny 191.255.0.0/16 le 32
# deny 191.255.0.0 - IANA reserved (I think)
ip prefix-list sanity-check seq 65 deny 192.0.0.0/3 ge 25
# deny masks > 25 for class C (192-222)
ip prefix-list sanity-check seq 70 deny 223.255.255.0/24 le 32
# deny anything in net 223 - IANA reserved
ip prefix-list sanity-check seq 75 deny 224.0.0.0/3 le 32
# deny class D/Experimental
```

www.cisco.com