



Multihoming Definition

- More than one link external to the local network
 - two or more links to the same ISP
 - two or more links to different ISPs
- Usually **two** external facing routers
 - one router gives link and provider redundancy only

Multihoming

- The scenarios described here apply equally well to end sites being customers of ISPs and ISPs being customers of other ISPs
- Implementation detail may be different

end site ® ISP	ISP controls config
ISP1 ® ISP2	ISPs share config

AS Numbers

- An Autonomous System Number is required by BGP
- Obtained from upstream ISP or Regional Registry
- Necessary when you have links to more than one ISP or exchange point

Configuring Policy

- Assumptions:
 - prefix-lists are used throughout
 - easier/better/faster than access-lists
- Three BASIC Principles
 - prefix-lists** to filter **prefixes**
 - filter-lists** to filter **ASNs**
 - route-maps** to apply **policy**

Originating Prefixes

- Basic Assumptions
 - MUST** announce assigned address block to Internet
 - MAY** also announce subprefixes - reachability is not guaranteed
 - RIR minimum allocation is /20 - several ISPs filter RIR blocks on this boundary - "Net Police"

Part of the “Net Police” prefix list

```
!! RIPE
ip prefix-list FILTER permit 62.0.0.0/8 ge 12 le 20
ip prefix-list FILTER permit 193.0.0.0/8 ge 12 le 20
ip prefix-list FILTER permit 194.0.0.0/7 ge 12 le 20
ip prefix-list FILTER permit 212.0.0.0/7 ge 12 le 20
!! APNIC
ip prefix-list FILTER permit 61.0.0.0/8 ge 12 le 20
ip prefix-list FILTER permit 202.0.0.0/7 ge 12 le 20
ip prefix-list FILTER permit 210.0.0.0/7 ge 12 le 20
!! ARIN
ip prefix-list FILTER permit 63.0.0.0/8 le 20
ip prefix-list FILTER permit 64.0.0.0/8 le 20
ip prefix-list FILTER permit 199.0.0.0/8 le 20
ip prefix-list FILTER permit 200.0.0.0/8 le 20
ip prefix-list FILTER permit 204.0.0.0/6 le 20
ip prefix-list FILTER permit 208.0.0.0/7 le 20
ip prefix-list FILTER permit 216.0.0.0/8 le 20
```

ISP/ISP Workshops © 2008, Cisco Systems, Inc. www.cisco.com

“Net Police” prefix list issues

- meant to “punish” ISPs who won’t and don’t aggregate
- impacts legitimate multihoming
- impacts regions where domestic backbone is unavailable or costs \$\$\$ compared with international bandwidth
- hard to maintain - requires updating when RIRs start allocating from new address blocks
- **don’t do it unless consequences understood**

ISP/ISP Workshops © 2008, Cisco Systems, Inc. www.cisco.com

Multihoming Options


Presentation, ID © 2008, Cisco Systems, Inc. www.cisco.com

Multihoming Scenarios

- Stub network
- Multi-homed stub network
- Multi-homed network
- Configuration Options

ISP/ISP Workshops © 2008, Cisco Systems, Inc. www.cisco.com


Stub Network



- No need for BGP
- Point static default to upstream ISP
- Upstream ISP advertises stub network
- Policy confined within upstream ISP’s policy

ISP/ISP Workshops © 2008, Cisco Systems, Inc. www.cisco.com


Multi-homed Stub Network



- Use BGP (not IGP or static) to loadshare
- Use private AS (ASN > 64511)
- Upstream ISP advertises stub network
- Policy confined within upstream ISP’s policy

ISP/ISP Workshops © 2008, Cisco Systems, Inc. www.cisco.com

Multi-Homed Network



- Many situations possible
 - multiple sessions to same ISP
 - secondary for backup only
 - load-share between primary and secondary
 - selectively use different ISPs

ISP/ISP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 13

Multiple Sessions to ISPs


- Planning and some work required to achieve load sharing
 - Point default towards one ISP
 - Learn selected prefixes from second ISP
 - Modify the number of prefixes learnt to achieve acceptable load sharing
- No magic solution

ISP/ISP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 14

Private-AS - Application

Applications

ISP with single-homed customers
corporate network with several regions and connections to the Internet only in the core




ISP/ISP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 15

Two links to the same ISP

One link primary, the other link backup only

Two links to the same ISP



- AS109 removes private AS and any customer subprefixes from Internet announcement

ISP/ISP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 17

Two links to the same ISP (one as backup only)

- Announce /19 aggregate on each link
 - primary link makes standard announcement
 - backup link increases metric on outbound, and reduces local-pref on inbound
- When one link fails, the announcement of the /19 aggregate via the other link ensures continued connectivity

ISP/ISP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 18

Two links to the same ISP (one as backup only)

- Router A Configuration

```
router bgp 65534
network 221.10.0.0 mask 255.255.224.0
neighbor 222.222.10.2 remote-as 109
neighbor 222.222.10.2 description RouterC
neighbor 222.222.10.2 prefix-list aggregate out
neighbor 222.222.10.2 prefix-list default in
!
ip prefix-list aggregate permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
```

Two links to the same ISP (one as backup only)

- Router B Configuration

```
router bgp 65534
network 221.10.0.0 mask 255.255.224.0
neighbor 222.222.10.6 remote-as 109
neighbor 222.222.10.6 description RouterD
neighbor 222.222.10.6 prefix-list aggregate out
neighbor 222.222.10.6 route-map routerD-out out
neighbor 222.222.10.6 prefix-list default in
neighbor 222.222.10.6 route-map routerD-in in
!
..next slide
```

Two links to the same ISP (one as backup only)

```
ip prefix-list aggregate permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map routerD-out permit 10
match ip address prefix-list aggregate
set metric 10
route-map routerD-out permit 20
!
route-map routerD-in permit 10
set local-preference 90
!
```

Two links to the same ISP (one as backup only)

- Router C Configuration (main link)

```
router bgp 109
neighbor 222.222.10.1 remote-as 65534
neighbor 222.222.10.1 default-originate
neighbor 222.222.10.1 prefix-list Customer in
neighbor 222.222.10.1 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
```

Two links to the same ISP (one as backup only)

- Router D Configuration (backup link)

```
router bgp 109
neighbor 222.222.10.5 remote-as 65534
neighbor 222.222.10.5 default-originate
neighbor 222.222.10.5 prefix-list Customer in
neighbor 222.222.10.5 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
```

Two links to the same ISP (one as backup only)

- Router E Configuration

```
router bgp 109
neighbor 222.222.10.17 remote-as 110
neighbor 222.222.10.17 remove-private-AS
neighbor 222.222.10.17 prefix-list Customer out
!
ip prefix-list Customer permit 221.10.0.0/19
```


- Router E removes the private AS and customer's subprefixes from external announcements

- Private AS still visible inside AS109

Two links to the same ISP

With Redundancy and Loadsharing

Two links to the same ISP (with redundancy)



- AS109 removes private AS and any customer subprefixes from Internet announcement

Loadsharing to the same ISP

- Announce /19 aggregate on each link
- Split /19 and announce as two /20s, one on each link
 - basic inbound loadsharing
 - assumes equal circuit capacity and even spread of traffic across address block
- Vary the split until “perfect” loadsharing achieved

Two links to the same ISP

Router A Configuration

```
router bgp 65534
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.0.0 mask 255.255.240.0
 neighbor 222.222.10.2 remote-as 109
 neighbor 222.222.10.2 prefix-list routerC out
 neighbor 222.222.10.2 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list routerC permit 221.10.0.0/20
ip prefix-list routerC permit 221.10.0.0/19
!
ip route 221.10.0.0 255.255.240.0 null0
ip route 221.10.0.0 255.255.224.0 null0
```

Two links to the same ISP

Router B Configuration

```
router bgp 65534
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.16.0 mask 255.255.240.0
 neighbor 222.222.10.6 remote-as 109
 neighbor 222.222.10.6 prefix-list routerD out
 neighbor 222.222.10.6 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list routerD permit 221.10.16.0/20
ip prefix-list routerD permit 221.10.0.0/19
!
ip route 221.10.0.0 255.255.224.0 null0
ip route 221.10.16.0 255.255.240.0 null0
```

Loadsharing to the same ISP

Default route for outbound traffic?

Use default-information originate for the IGP and rely on IGP metrics for nearest exit

e.g. on router A:

```
router ospf 65534
 default-information originate metric 2 metric-type 1
```

Two links to the same ISP

- Router C Configuration

```
router bgp 109
neighbor 222.222.10.1 remote-as 65534
neighbor 222.222.10.1 default-originate
neighbor 222.222.10.1 prefix-list Customer in
neighbor 222.222.10.1 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19 le 20
ip prefix-list default permit 0.0.0.0/0
```

- Router C only allows in /19 and /20 prefixes from customer block

Two links to the same ISP

- Router D Configuration

```
router bgp 109
neighbor 222.222.10.5 remote-as 65534
neighbor 222.222.10.5 default-originate
neighbor 222.222.10.5 prefix-list Customer in
neighbor 222.222.10.5 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19 le 20
ip prefix-list default permit 0.0.0.0/0
```

- Router D only allows in /19 and /20 prefixes from customer block

Two links to the same ISP

- Router E is AS109 border router
 - removes subprefixes in the private AS from external announcements
 - removes the private AS from external announcement of the customer /19

Two links to the same ISP (with redundancy)

- Router E Configuration

```
router bgp 109
neighbor 222.222.10.17 remote-as 110
neighbor 222.222.10.17 remove-private-AS
neighbor 222.222.10.17 prefix-list Customer out
!
ip prefix-list Customer permit 221.10.0.0/19
```

- Private AS still visible inside AS109


Loadsharing to the same ISP

- Loadsharing configuration is only on customer router
- Upstream ISP has to
 - remove customer subprefixes from external announcements
 - remove private AS from external announcements
- Could also use BGP communities

Two links to the same ISP

Multiple Dualhomed Customers (RFC2270)

Multiple Dualhomed Customers (RFC2270)



- AS109 removes private AS and any customer subprefixes from Internet announcement

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 37

Multiple Dualhomed Customers

- Customer announcements as per previous example
- Use the **same** private AS for each customer
 - documented in RFC2270
 - address space is not overlapping
 - each customer hears default only
- Router **A_n** and **B_n** configuration same as Router A and B previously

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 38

Two links to the same ISP

- Router A1 Configuration

```
router bgp 65534
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.0.0 mask 255.255.240.0
 neighbor 222.222.10.2 remote-as 109
 neighbor 222.222.10.2 prefix-list routerC out
 neighbor 222.222.10.2 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list routerC permit 221.10.0.0/20
ip prefix-list routerC permit 221.10.0.0/19
!
ip route 221.10.0.0 255.255.240.0 null0
ip route 221.10.0.0 255.255.240.0 null0
```

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 39

Two links to the same ISP

- Router B1 Configuration

```
router bgp 65534
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.16.0 mask 255.255.240.0
 neighbor 222.222.10.6 remote-as 109
 neighbor 222.222.10.6 prefix-list routerD out
 neighbor 222.222.10.6 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list routerD permit 221.10.16.0/20
ip prefix-list routerD permit 221.10.0.0/19
!
ip route 221.10.0.0 255.255.224.0 null0
ip route 221.10.16.0 255.255.240.0 null0
```

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 40

Multiple Dualhomed Customers

- Router C Configuration

```
router bgp 109
 neighbor bgp-customers peer-group
 neighbor bgp-customers remote-as 65534
 neighbor bgp-customers default-originate
 neighbor bgp-customers prefix-list default out
 neighbor 222.222.10.1 peer-group bgp-customers
 neighbor 222.222.10.1 description Customer One
 neighbor 222.222.10.1 prefix-list Customer1 in
 neighbor 222.222.10.9 peer-group bgp-customers
 neighbor 222.222.10.9 description Customer Two
 neighbor 222.222.10.9 prefix-list Customer2 in
```

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 41

Multiple Dualhomed Customers

```
neighbor 222.222.10.17 peer-group bgp-customers
neighbor 222.222.10.17 description Customer Three
neighbor 222.222.10.17 prefix-list Customer3 in
!
ip prefix-list Customer1 permit 221.10.0.0/19 le 20
ip prefix-list Customer2 permit 221.16.64.0/19 le 20
ip prefix-list Customer3 permit 221.14.192.0/19 le 20
ip prefix-list default permit 0.0.0.0/0
```

- Router C only allows in /19 and /20 prefixes from customer block

ISP/DP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 42

Multiple Dualhomed Customers

- Router D Configuration

```
router bgp 109
  neighbor bgp-customers peer-group
  neighbor bgp-customers remote-as 65534
  neighbor bgp-customers default-originate
  neighbor bgp-customers prefix-list default out
  neighbor 222.222.10.5 peer-group bgp-customers
  neighbor 222.222.10.5 description Customer One
  neighbor 222.222.10.5 prefix-list Customer1 in
  neighbor 222.222.10.13 peer-group bgp-customers
  neighbor 222.222.10.13 description Customer Two
  neighbor 222.222.10.13 prefix-list Customer2 in
```

Multiple Dualhomed Customers

```
neighbor 222.222.10.21 peer-group bgp-customers
neighbor 222.222.10.21 description Customer Three
neighbor 222.222.10.21 prefix-list Customer3 in
!
ip prefix-list Customer1 permit 221.10.0.0/19 le 20
ip prefix-list Customer2 permit 221.16.64.0/19 le 20
ip prefix-list Customer3 permit 221.14.192.0/19 le 20
ip prefix-list default permit 0.0.0.0/0
```

- Router D only allows in /19 and /20 prefixes from customer block

Multiple Dualhomed Customers

- Router E Configuration is as previously assumes customer address space is not part of upstream's address block

```
router bgp 109
  neighbor 222.222.10.17 remote-as 110
  neighbor 222.222.10.17 remove-private-AS
  neighbor 222.222.10.17 prefix-list Customers out
!
ip prefix-list Customers permit 221.10.0.0/19
ip prefix-list Customers permit 221.16.64.0/19
ip prefix-list Customers permit 221.14.192.0/19
```

- Private AS still visible inside AS109

Multiple Dualhomed Customers

- If customers' prefixes come from ISP's address block
do **NOT** announce them to the Internet
announce ISP aggregate only
- Router E configuration:

```
router bgp 109
  neighbor 222.222.10.17 remote-as 110
  neighbor 222.222.10.17 prefix-list my-aggregate out
!
ip prefix-list my-aggregate permit 221.8.0.0/13
```

Multihoming Summary

- Use private AS for multihoming to upstream
- Leak subprefixes to upstream only to aid loadsharing
- Upstream router E configuration is uniform across all scenarios



Two links to different ISPs

Two links to different ISPs


- Use Public ASes
or use private AS if agreed with the other ISP
- Address space comes from both upstreams (PA space) **or** Regional Internet Registry (PI space)
- Configuration concepts very similar

Two links to different ISPs With Redundancy

Two links to different ISPs (with redundancy)

- Example for PI space
ISP network, or large enterprise site
- Announce /19 aggregate on each link
- Split /19 and announce as two /20s, one on each link
basic inbound loadsharing
- When one link fails, the announcement of the /19 aggregate via the other ISP ensures continued connectivity

Two links to different ISPs (with redundancy)



Two links to different ISPs (with redundancy)

- Router A Configuration
- ```

router bgp 107
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.0.0 mask 255.255.240.0
 neighbor 222.222.10.1 remote-as 109
 neighbor 222.222.10.1 prefix-list firstblock out
 neighbor 222.222.10.1 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
!
ip prefix-list firstblock permit 221.10.0.0/20
ip prefix-list firstblock permit 221.10.0.0/19

```

## Two links to different ISPs (with redundancy)

- Router B Configuration
- ```

router bgp 107
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.16.0 mask 255.255.240.0
 neighbor 220.1.5.1 remote-as 108
 neighbor 220.1.5.1 prefix-list secondblock out
 neighbor 220.1.5.1 prefix-list default in
!
ip prefix-list default permit 0.0.0.0/0
!
ip prefix-list secondblock permit 221.10.16.0/20
ip prefix-list secondblock permit 221.10.0.0/19
    
```

Two links to different ISPs (with redundancy)

- Router C Configuration

```
router bgp 109
neighbor 221.10.1.1 remote-as 107
neighbor 221.10.1.1 default-originate
neighbor 221.10.1.1 prefix-list AS107cust in
neighbor 221.10.1.1 prefix-list default-out out
!
```

- Router C only announces default to AS 107
- Router C only accepts AS107's prefix block

Two links to different ISPs (with redundancy)

- Router D Configuration

```
router bgp 108
neighbor 220.1.5.1 remote-as 107
neighbor 220.1.5.1 default-originate
neighbor 220.1.5.1 prefix-list AS107cust in
neighbor 220.1.5.1 prefix-list default-out out
!
```

- Router D only announces default to AS 107
- Router D only accepts AS107's prefix block


Two links to different ISPs

One link primary, the other link
backup only

Two links to different ISPs (one as backup only)

- Announce /19 aggregate on each link
primary link makes standard
announcement
backup link lengthens the AS PATH by
using AS PATH prepend
- When one link fails, the announcement
of the /19 aggregate via the other link
ensures continued connectivity

Two links to different ISPs (one as backup only)



Two links to different ISPs (one as backup only)

- Router A Configuration

```
router bgp 107
network 221.10.0.0 mask 255.255.224.0
neighbor 222.222.10.1 remote-as 109
neighbor 222.222.10.1 prefix-list aggregate out
neighbor 222.222.10.1 prefix-list default in
!
ip prefix-list aggregate permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
```

Two links to different ISPs (one as backup only)

- Router B Configuration

```
router bgp 107
network 221.10.0.0 mask 255.255.224.0
neighbor 220.1.5.1 remote-as 108
neighbor 220.1.5.1 prefix-list aggregate out
neighbor 220.1.5.1 route-map routerD-out out
neighbor 220.1.5.1 prefix-list default in
neighbor 220.1.5.1 route-map routerD-in in
!
```

..next slide

ISP/DP Workshops

© 2008, Cisco Systems, Inc.

www.cisco.com

61

Two links to different ISPs (one as backup only)

- Router B Configuration

```
!
ip prefix-list aggregate permit 221.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map routerD-out permit 10
set as-path prepend 107 107
!
route-map routerD-in permit 10
set local-preference 80
```

ISP/DP Workshops

© 2008, Cisco Systems, Inc.

www.cisco.com

62

Two links to different ISPs

More Controlled Loadsharing

Loadsharing with different ISPs

- Announce /19 aggregate on each link

On first link, announce /19 as normal

On second link, announce /19 with longer AS PATH, and announce one /20 subprefix

controls loadsharing between upstreams and the Internet

- Vary the subprefix size and AS PATH length until "perfect" loadsharing achieved

- Still require redundancy!


ISP/DP Workshops

© 2008, Cisco Systems, Inc.

www.cisco.com

64

Loadsharing with different ISPs



ISP/DP Workshops

© 2008, Cisco Systems, Inc.

www.cisco.com

65

Loadsharing with different ISPs

- Router A Configuration

```
router bgp 107
network 221.10.0.0 mask 255.255.224.0
neighbor 222.222.10.1 remote-as 109
neighbor 222.222.10.1 prefix-list default in
neighbor 222.222.10.1 prefix-list aggregate out
!
```

```
ip prefix-list aggregate permit 221.10.0.0/19
```

ISP/DP Workshops

© 2008, Cisco Systems, Inc.

www.cisco.com

66

Loadsharing with different ISPs

- Router B Configuration

```
router bgp 107
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.16.0 mask 255.255.240.0
 neighbor 220.1.5.1 remote-as 108
 neighbor 220.1.5.1 prefix-list default in
 neighbor 220.1.5.1 prefix-list subblocks out
 neighbor 220.1.5.1 route-map routerD out
!
```

..next slide..

Loadsharing with different ISPs


```
route-map routerD permit 10
 match ip address prefix-list aggregate
 set as-path prepend 107 107
route-map routerD permit 20
!
```

```
ip prefix-list subblocks permit 221.10.0.0/19 le 20
ip prefix-list aggregate permit 221.10.0.0/19
```

Loadsharing Using Communities

4 links - Private AS

Private AS



- AS109 removes private AS and any customer subprefixes from Internet announcement

Private AS

- Announce /19 aggregate on each link
- Split /19 and announce as four /21s, one on each link
 - basic inbound loadsharing
 - assumes equal circuit capacity and even spread of traffic across address block
- Vary the split until "perfect" loadsharing achieved
 - use the no-export community for subprefixes

Private AS

- Router A Configuration

```
router bgp 65534
 network 221.10.0.0 mask 255.255.224.0
 network 221.10.0.0 mask 255.255.248.0
 neighbor 222.222.10.2 remote-as 109
 neighbor 222.222.10.2 send-community
 neighbor 222.222.10.2 prefix-list subblocks1 out
 neighbor 222.222.10.2 route-map routerC-out out
 neighbor 222.222.10.2 prefix-list default in
!
```

..next slide

Private AS

```
ip prefix-list subblocks1 permit 221.10.0.0/19
ip prefix-list subblocks1 permit 221.10.0.0/21
!
ip prefix-list firstblock permit 221.10.0.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map routerC-out permit 10
match ip address prefix-list firstblock
set community no-export
route-map routerC-out permit 20
```

Private AS

• Router B Configuration

```
router bgp 65534
network 221.10.0.0 mask 255.255.224.0
network 221.10.24.0 mask 255.255.248.0
neighbor 222.222.20.2 remote-as 109
neighbor 222.222.20.2 send-community
neighbor 222.222.20.2 prefix-list subblocks2 out
neighbor 222.222.20.2 route-map routerD-out out
neighbor 222.222.20.2 prefix-list default in
!
..next slide
```

Private AS

```
ip prefix-list subblocks2 permit 221.10.0.0/19
ip prefix-list subblocks2 permit 221.10.24.0/21
!
ip prefix-list secondblock permit 221.10.24.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map routerD-out permit 10
match ip address prefix-list secondblock
set community no-export
route-map routerD-out permit 20
```

Private AS

• Router E Configuration

```
router bgp 109
neighbor 222.222.10.17 remote-as 110
neighbor 222.222.10.17 remove-private-AS
!
```

- Router E removes the private AS from external announcements
- Router E automatically removes subprefixes with no-export community set
- Private AS still visible inside AS109


Private AS

- Router C and D configuration is as previously
- AS109 routers will not advertise prefixes marked with community no-export to other ASes
- AS109 routers still need to filter the private AS
- Only a single /19 prefix is announced to the Internet - no routing table bloat! :-)

Loadsharing Using Communities

4 links - Public AS

Public AS



- 4 links between AS107 and AS109

Public AS

- Announce /19 aggregate on each link
- Split /19 and announce as four /21s, one on each link

basic inbound loadsharing

assumes equal circuit capacity and even spread of traffic across address block

- Vary the split until “perfect” loadsharing achieved

use the no-export community for subprefixes

Public AS

- Router A Configuration

```
router bgp 107
network 221.10.0.0 mask 255.255.224.0
network 221.10.0.0 mask 255.255.248.0
neighbor 222.222.10.2 remote-as 109
neighbor 222.222.10.2 send-community
neighbor 222.222.10.2 prefix-list subblocks1 out
neighbor 222.222.10.2 route-map routerC-out out
neighbor 222.222.10.2 prefix-list default in
!
..next slide
```

Public AS

```
ip prefix-list subblocks1 permit 221.10.0.0/19
ip prefix-list subblocks1 permit 221.10.0.0/21
!
ip prefix-list firstblock permit 221.10.0.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map routerC-out permit 10
match ip address prefix-list firstblock
set community no-export
route-map routerC-out permit 20
```

Public AS

- Router B Configuration

```
router bgp 107
network 221.10.0.0 mask 255.255.224.0
network 221.10.24.0 mask 255.255.248.0
neighbor 222.222.20.2 remote-as 109
neighbor 222.222.20.2 send-community
neighbor 222.222.20.2 prefix-list subblocks2 out
neighbor 222.222.20.2 route-map routerD-out out
neighbor 222.222.20.2 prefix-list default in
!
..next slide
```

Public AS

```
ip prefix-list subblocks2 permit 221.10.0.0/19
ip prefix-list subblocks2 permit 221.10.24.0/21
!
ip prefix-list secondblock permit 221.10.24.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map routerD-out permit 10
match ip address prefix-list secondblock
set community no-export
route-map routerD-out permit 20
```

Public AS

- Router C Configuration

```
router bgp 109
neighbor 222.222.10.1 remote-as 107
neighbor 222.222.10.1 default-originate
neighbor 222.222.10.1 prefix-list Customer in
neighbor 222.222.10.1 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19 le 21
ip prefix-list default permit 0.0.0.0/0
```

Loadsharing to the same ISP

- Router D Configuration

```
router bgp 109
neighbor 222.222.10.5 remote-as 107
neighbor 222.222.10.5 default-originate
neighbor 222.222.10.5 prefix-list Customer in
neighbor 222.222.10.5 prefix-list default out
!
ip prefix-list Customer permit 221.10.0.0/19 le 21
ip prefix-list default permit 0.0.0.0/0
```

Loadsharing to the same ISP

- Router E Configuration

```
router bgp 109
neighbor 222.222.10.17 remote-as 110
neighbor 222.222.10.17 filter-list 1 out
!
ip as-path access-list 1 permit ^107$
ip as-path access-list 1 permit ^$
```

- Router E only has to announce AS107 in the same way it announces other ASes

Public AS

- AS109 routers will not advertise prefixes marked with community no-export to other ASes
- AS109 ISP has no configuration work to do
AS107 ISP can control his own loadsharing
- Only a single /19 prefix is announced to the Internet - no routing table bloat! :-)

Enterprise Multihoming

Address Space from different ISPs


Enterprise Multihoming

- Common situation is enterprise multihoming
address space used by enterprise comes from both upstream ISPs
multihoming and loadsharing more difficult
want to avoid leaking subprefixes of upstream provider address space

Enterprise Multihoming

- New conditional advertisement feature in BGP
 - loadsharing under normal conditions
 - subprefixes only announced in failure scenarios
 - requires upstreams to announce **only one** prefix to enterprise border network

Steady State



Steady State

- ISP1 has 220.10.0.0/16 address block
- ISP2 has 222.5.0.0/16 address block
- Enterprise customer multihomes
 - upstreams don't announce subprefixes
 - can use private AS (ASN>64511)
 - R2 and R4 originate default in their IGP
 - outbound traffic uses nearest exit (IGP metrics)

Steady State

- Router2 configuration:


```
router bgp 65534
 network 220.10.4.0 mask 255.255.254.0
 network 222.5.64.0 mask 255.255.254.0
 neighbor <R1> remote-as 150
 neighbor <R1> prefix-list isp1-in in
 neighbor <R1> prefix-list isp1-out out
 neighbor <R1> advertise-map isp2-sb non-exist-map isp2-bb
 neighbor <R4> remote-as 65534
 neighbor <R4> update-source loopback 0
!
ip route 220.10.4.0 255.255.254.0 null0 250
..next slide
```

Steady State

```
ip route 222.5.64.0 255.255.254.0 null0 250
!
ip prefix-list isp1-out permit 220.10.4.0/23
ip prefix-list isp2-out permit 222.5.64.0/23
!
ip prefix-list isp1-in permit 220.10.0.0/16
ip prefix-list isp2-in permit 222.5.0.0/16
!
route-map isp2-sb permit 10
 match ip address prefix-list isp2-out
!
route-map isp2-bb permit 10
 match ip address prefix-list isp2-in
```

Steady State

- Router2 peers iBGP with Router4
 - hears ISP2's /16 prefix
- Router2 peers eBGP with Router1
 - hears ISP1's /16 prefix only
 - announces 220.10.4.0/23 only

Link Failure

The diagram illustrates a network topology with two Internet Service Providers (ISPs) and an Enterprise.

- ISP₁** (Left):
 - Cloud: 220.10.0.0/16 and 222.5.64.0/23 with "remove-private-AS"
 - Router **R₁**: 220.10.4/23 and 222.5.64/23 announced to ISP₁
 - Router **R₂**: 220.10.4.0/23
- ISP₂** (Right):
 - Cloud: 222.5.0.0/16
 - Router **R₃**: 222.5.0.0/16
 - Router **R₄**: 222.5.64/23
- Enterprise** (Center):
 - Router **R₂**: 220.10.16
 - Router **R₄**: 222.5.65.10
 - Internal IP: 220.10.5.250
 - Internal IP: 222.5.64/23 (marked with a red X)

A red 'X' is placed over the link between R₃ and R₄, indicating a link failure. Dashed arrows show the flow of traffic and announcements between the routers and their respective ISPs.

© 2009, Cisco Systems, Inc. www.cisco.com

- ISP/IXP Workshops © 2000, Cisco Systems, Inc. www.cisco.com 98

Enterprise Multihoming

- **Conditional advertisement useful when address space comes from both upstreams**
 - no subprefixes leaked to Internet unless in failure situation
- **Alternative backup mechanism would be to leak /23 prefixes with longer AS path**
 - routing table bloat, reachability issues

© 2006, Cisco Systems, Inc. www.cisco.com

Summary

- **Private vs Public ASes**
- **Multihoming to Same ISP**
- **Multihoming to Different ISPs**
- **Community based multihoming**
- **Enterprise multihoming**

ISP/IXP Workshops © 2008, Cisco Systems, Inc. www.cisco.com 100