

Internet Redundancy

with BGP



Internet Redundancy with BGP

Date: 01/27/10

ENGINEERED NETWORKS

www.engineered-net.com

Rev: 1.0

ssurdock@engineered-net.com

Page: 1

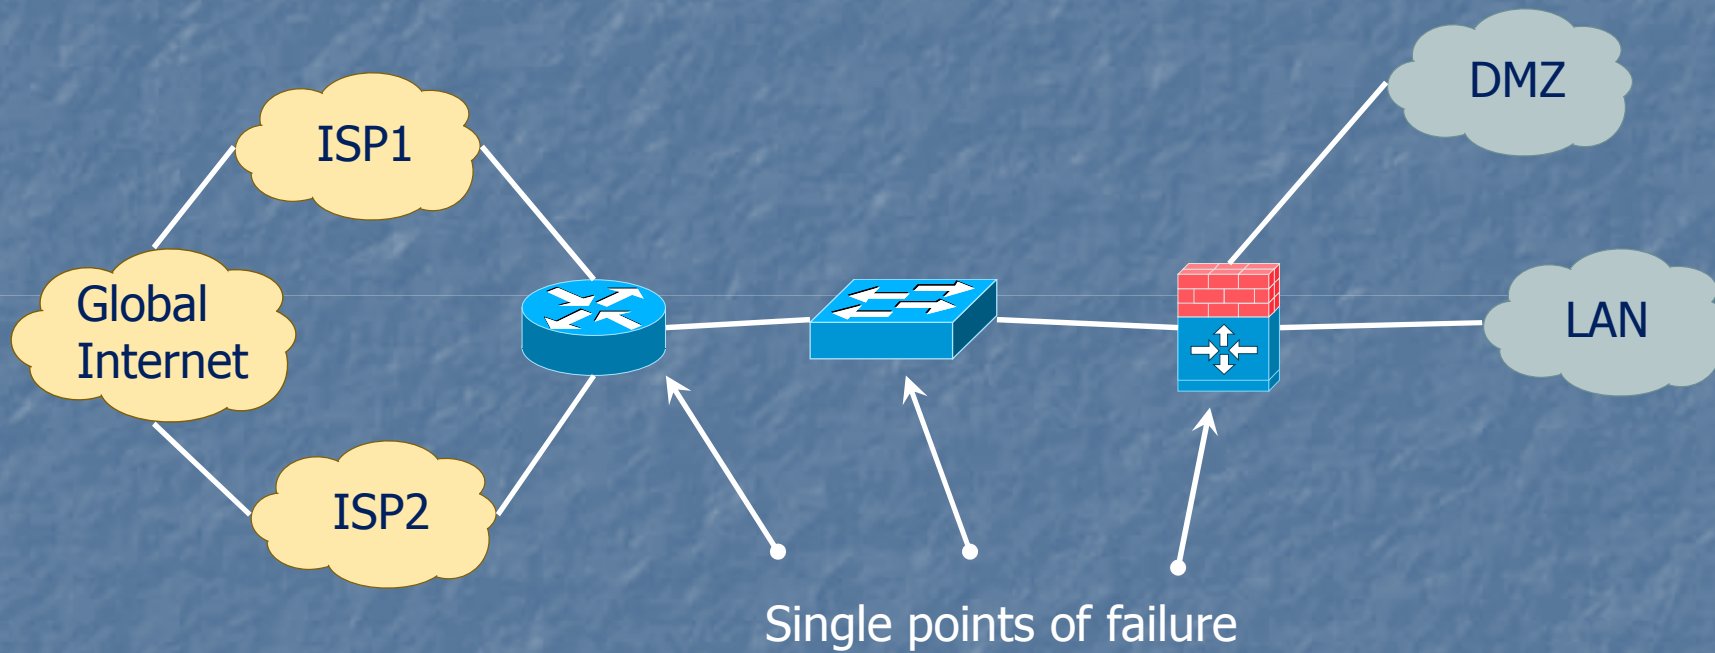
Design Goals

- No single point of failure.
 - Availability = $MTBF / (MTBF + MTTR)$
- Reasonable failover time.
- Efficient bandwidth usage.
- Vendor independent.
- Monitored, measured and adjustable.

Design Options

- Outsource
 - ✗ Still need to get to data from corporate office.
- NAT-based
 - ✗ Not good with “many” inbound services. Limited protection from ISP routing issues.
 - ✗ Handles many types of ISP failures.
- BGP

Design



Internet Redundancy with BGP

ENGINEERED NETWORKS


www.engineered-net.com
ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

Page: 4

Updated Design



No single points of failure

Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

Page: 5

Next Steps – Layer 8

- You need a /24 from one of your providers.
- You need a Autonomous System Number (ASN) from ARIN.
- Letter of Authorization for upstreams.
- Procure hardware.

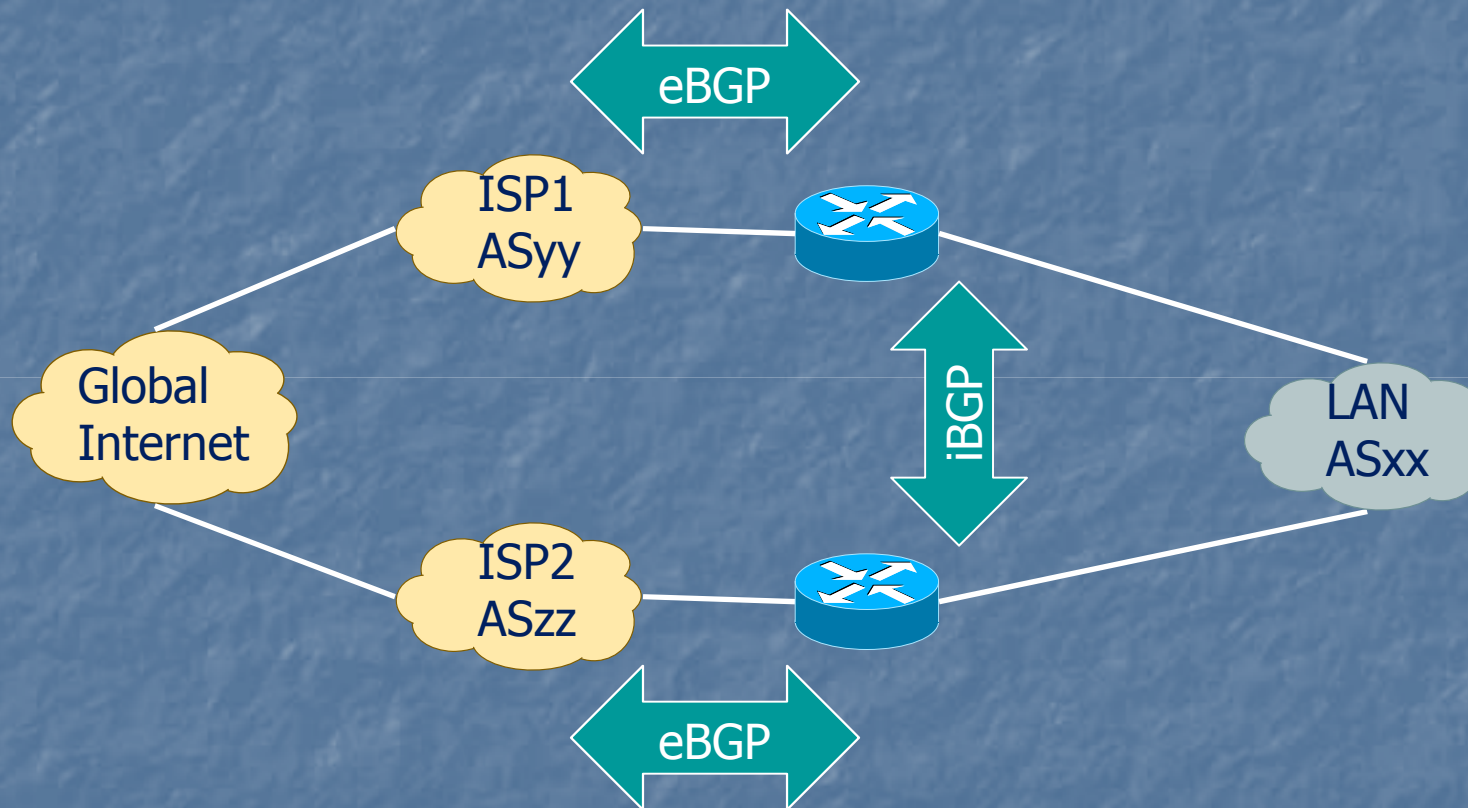
What Kind of Hardware?

- What kind of Packets Per Second (PPS)?
 - Less than 200 kpps (100 Mbps) *could* use PC hardware.
- How many routes will the router hold?
 - One Default Route?
 - ISP Customer Routes?
 - All 300,000+ Internet Routes?

What is BGP

- Border Gateway Protocol (version 4)
- Distance/Path Vector routing protocol
- TCP Session between “peers”
- Internal AS (iBGP) and External AS (eBGP) sessions
- Exchange route information with attributes

BGP Sessions



Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

Page: 9

(Basic) Router Config

ROUTER-A

```
router bgp xx
  no synchronization
  network my.class.C.0 mask 255.255.255.0
  !
  neighbor isp.1.ip.addr remote-as yy
  neighbor rtr.b.ip.addr remote-as xx
ip route my.class.C.0 255.255.255.0 FastEthernet0
```


ROUTER-B

```
router bgp xx
  no synchronization
  network my.class.C.0 mask 255.255.255.0
  !
  neighbor isp.2.ip.addr remote-as zz
  neighbor rtr.b.ip.addr remote-as xx
ip route my.class.C.0 255.255.255.0 FastEthernet0
```

Routing

INBOUND traffic

What I tell (advertise to) my ISP via BGP affects inbound traffic.



What my ISP tells (advertises to) me via BGP affects outbound traffic.

OUTBOUND traffic

Internet Redundancy with BGP

Date: 01/27/10

ENGINEERED NETWORKS

www.engineered-net.com

Rev: 1.0

ssurdock@engineered-net.com

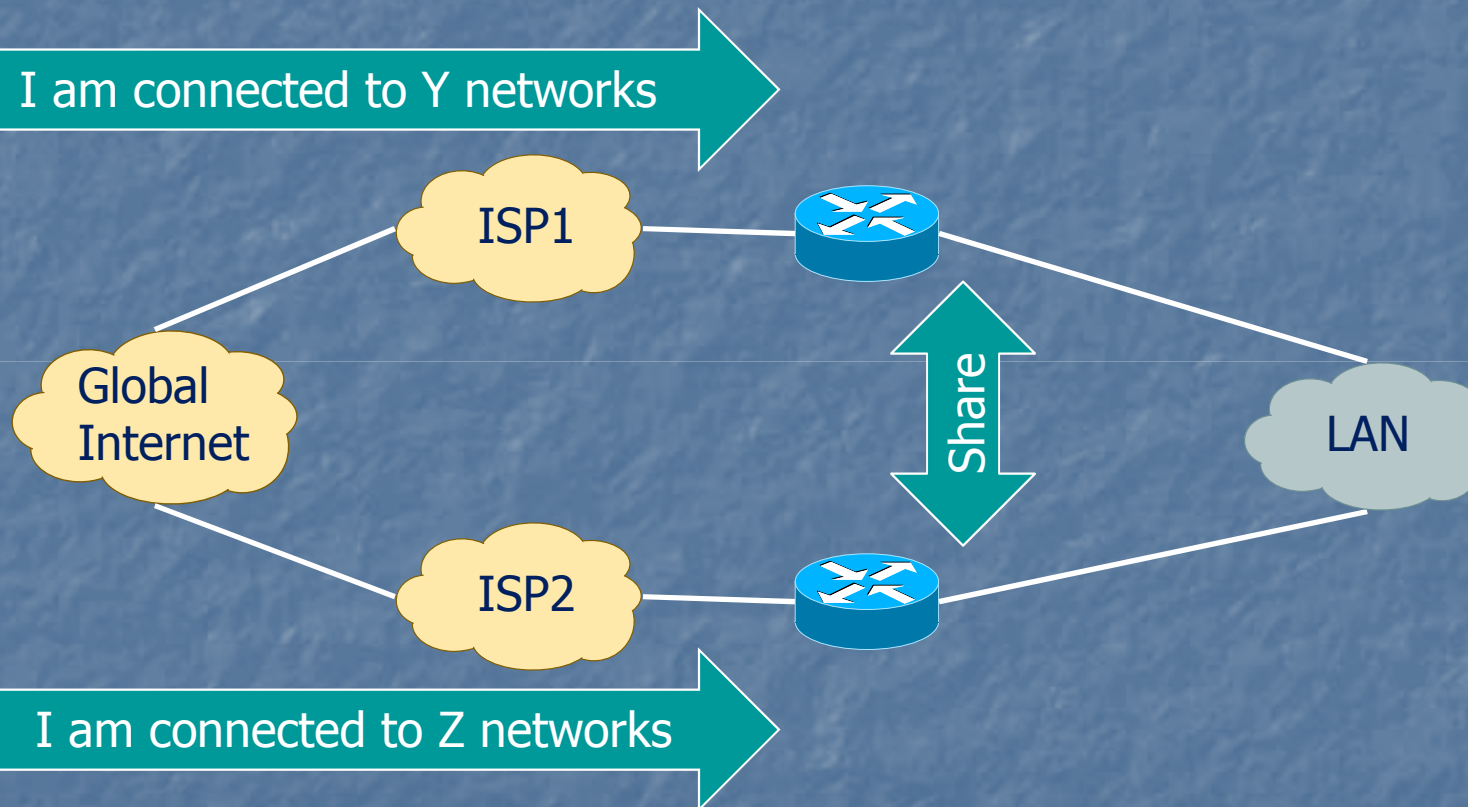
Page: 11

Routing - Outbound

- BGP is used to receive routes from your provider to make *outbound* routing decisions.
 - Default Only - Simplest case.
 - Customer + Default - Good compromise.
 - Full Tables - Best routing & traffic control
 - Route Maps used for Traffic Engineering

Outbound – Customer Routes

BGP: I am connected to Y networks



Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

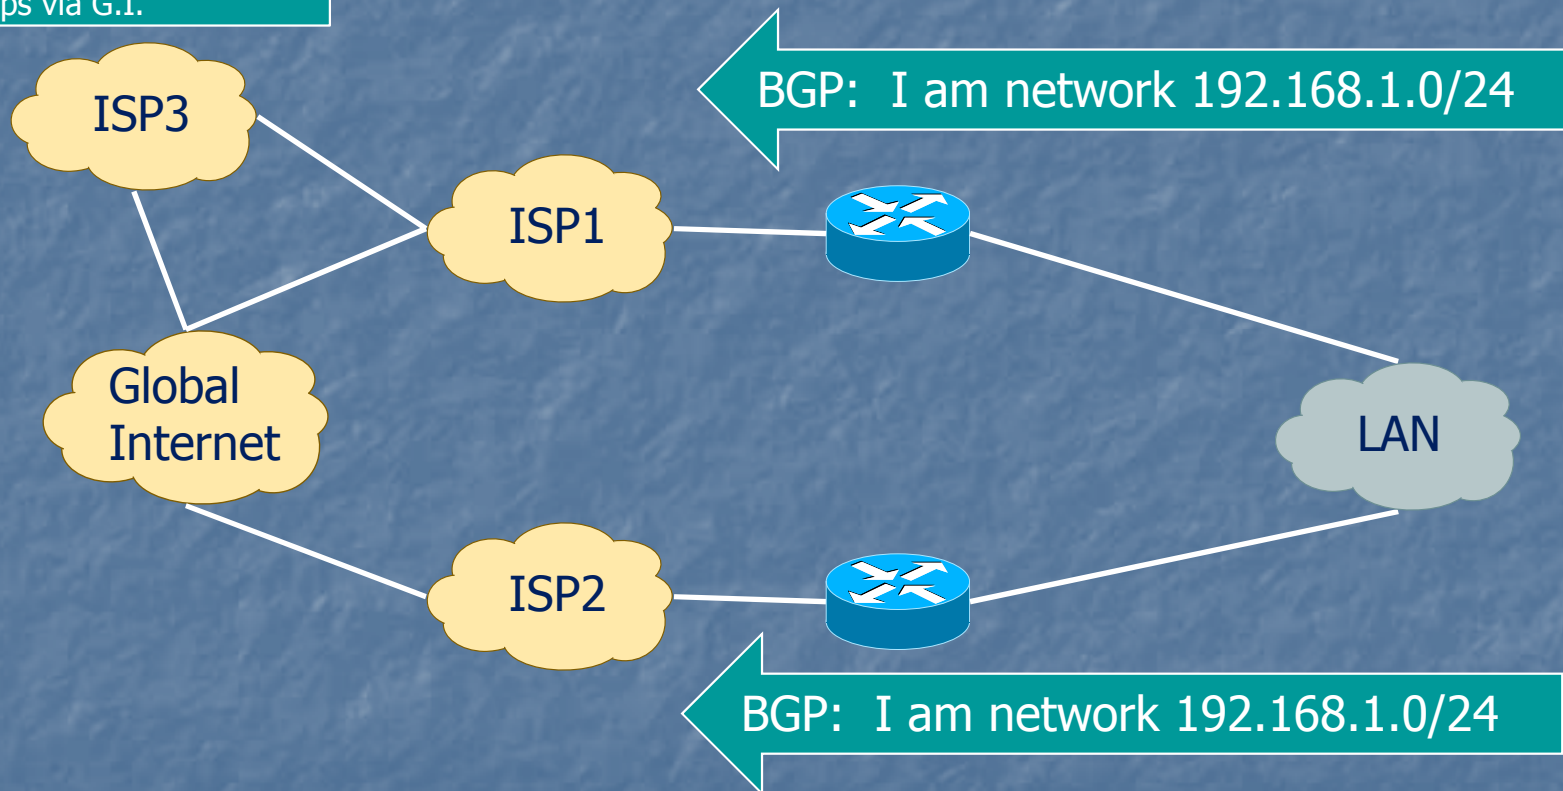
Page: 13

Routing - Inbound

- BGP used to announce your /24 to the Internet. This affects *inbound* routing decisions.
- Limited knobs available for tweaking the link preference.
 - AS Hops via prepending

Inbound – BGP Announcement

192.168.1.0/24 is ONE hop via ISP2 and TWO hops via G.I.



Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com


Date: 01/27/10

Rev: 1.0

Page: 15

Inbound – BGP Prepending

192.168.1.0/24 is
THREE hop via ISP2
and TWO hops via G.I.



Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com

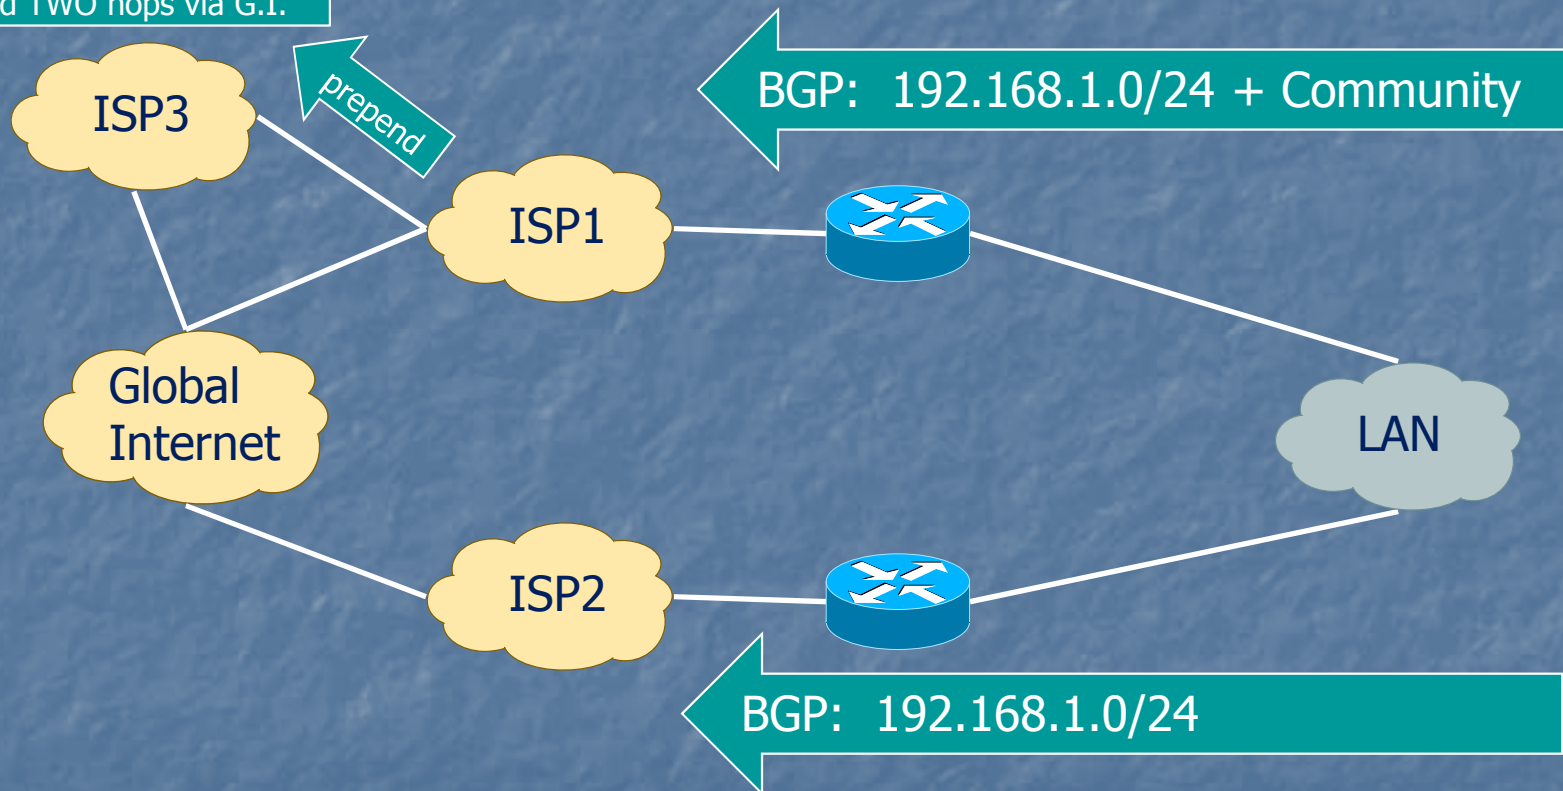
Date: 01/27/10

Rev: 1.0

Page: 16

Inbound – Selective Prepends

192.168.1.0/24 is
THREE hop via ISP2
and TWO hops via G.I.



Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com
ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

Page: 17

(More) Complex Config

```
ROUTER-A
interface Loopback0
  ip address rtr.a.ip.addr 255.255.255.255
!
router bgp xx
  no synchronization
  no auto-summary
  bgp log-neighbor-changes
  network my.class.C.0 mask 255.255.255.0
  neighbor isp.1.ip.addr remote-as yy
  neighbor isp.1.ip.addr description eBGP to ISP1
  neighbor isp.1.ip.addr soft-reconfiguration inbound
  neighbor isp.1.ip.addr version 4
  neighbor isp.1.ip.addr prefix-list bogons in
  neighbor isp.1.ip.addr prefix-list announce out
!
```

(More) Complex Config

ROUTER-A (cont.)

```
!  
neighbor rtr.b.ip.addr remote-as xx  
neighbor rtr.b.ip.addr description iBGP with ROUTER-B  
neighbor rtr.b.ip.addr password somefinecrypto  
neighbor rtr.b.ip.addr update-source Loopback0  
neighbor rtr.b.ip.addr soft-reconfiguration inbound  
neighbor rtr.b.ip.addr version 4  
neighbor rtr.b.ip.addr next-hop-self  
neighbor rtr.b.ip.addr prefix-list bogons in  
!  
ip route my.class.C.0 255.255.255.0 FastEthernet0  
!  
ip prefix-list announce description Our allowed routing announcements  
ip prefix-list announce seq 5 permit my.class.C.0/24  
ip prefix-list announce seq 10 deny 0.0.0.0/0 le 32
```

Internet Redundancy with BGP

ENGINEERED NETWORKS

www.engineered-net.com

ssurdock@engineered-net.com

Date: 01/27/10

Rev: 1.0

Page: 19

Validating and Testing

- Outbound
 - `'show ip neighbor isp.1.ip.addr advertised-routes'`
- Inbound
 - `'show ip neighbor isp.1.ip.addr received-routes'`
- External
 - Telnet to `'route-views.oregon-ix.net'`
 - `'show ip bgp my.class.C.0'`

Validating and Testing

```
route-views.oregon-ix.net> show ip bgp 205.217.83.0
BGP routing table entry for 205.217.64.0/19, version 4372857
Paths: (33 available, best #10, table Default-IP-Routing-Table)
  Not advertised to any peer
  3277 3216 3549 12129 12129 4208
    194.85.4.55 from 194.85.4.55 (194.85.4.16)
      Origin IGP, localpref 100, valid, external
      Community: 3216:3000 3216:3004 3277:3216 3549:4203 3549:30840
...
  701 4208
    157.130.10.233 from 157.130.10.233 (137.39.3.60)
      Origin IGP, localpref 100, valid, external, best
  3561 12129 4208
    206.24.210.102 from 206.24.210.102 (206.24.210.102)
      Origin IGP, localpref 100, valid, external
```

Next Steps

- Test periodically – Pull the plug.
- Monitoring – MRTG, Cacti, OpenNMS.
- Coarse tuning inbound with AS prepends.
- Fine tuning inbound with selective prepends.
- Fine tuning outbound with route-maps.