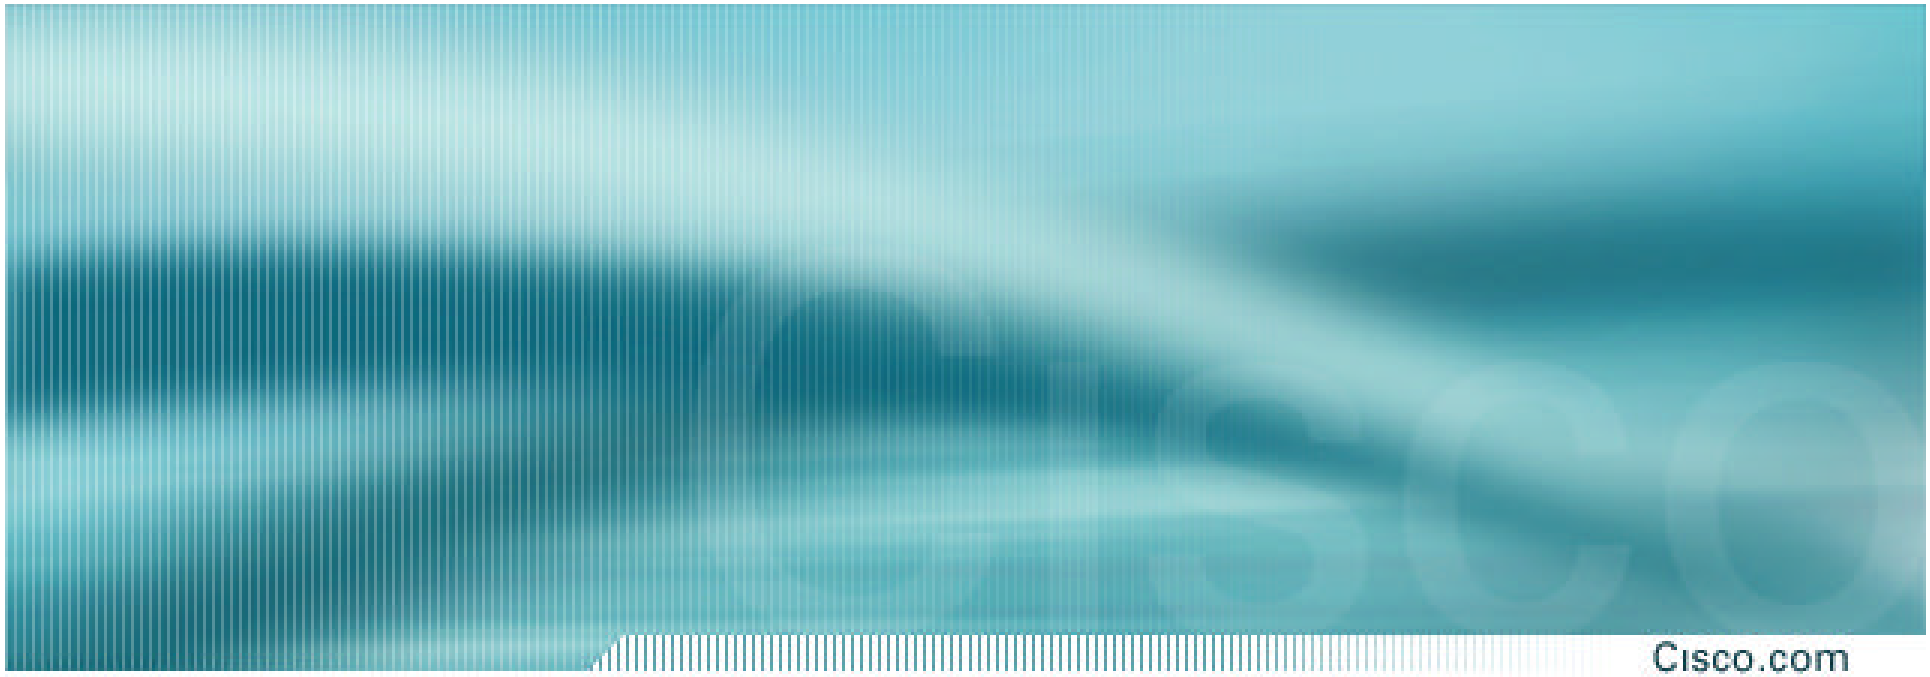

BGP enhancements for IPv6

Philip Smith <pfs@cisco.com>

1st South Asian IPv6 Summit/SANOG 3

15-22 January 2004

Bangalore, India

Preliminaries

- **Thanks to Cisco IPv6 team for the content**
- **Presentation slides available on**

  ftp://ftp-eng.cisco.com/pfs/seminars/SANOG3-BGP-IPv6.pdf

Adding IPv6 to BGP…

- **RFC2858 (obsoletes RFC2283)**

  Defines *Multi-protocol Extensions for BGP4*

  Enables BGP to carry routing information of protocols other than IPv4

    e.g. MPLS, IPv6, Multicast etc

  Exchange of multiprotocol NLRI must be negotiated at session startup

- **RFC2545**

  Use of BGP Multiprotocol Extensions for IPv6 Inter-Domain Routing

Adding IPv6 to BGP…

- **New optional and non-transitive BGP attributes:**

  MP_REACH_NLRI (Attribute code: 14)

  "Carry the set of reachable destinations together with the next-hop information to be used for forwarding to these destinations" (RFC2858)

  MP_UNREACH_NLRI (Attribute code: 15)

  Carry the set of unreachable destinations

- **Attribute contains one or more Triples:**

  AFI Address Family Information

  Next-Hop Information (must be of the same address family)

  NLRI Network Layer Reachability Information

Adding IPv6 to BGP…

- **Address Family Information (AFI) for IPv6**
  - **AFI = 2 (RFC 1700)**
  - **Sub-AFI = 1 Unicast**
  - **Sub-AFI = 2 (Mulitcast for RPF check)**
  - **Sub-AFI = 3 for both Unicast and Mulitcast**
  - **Sub-AFI = 4 Label**
  - **Sub-AFI= 128 VPN**

BGP Considerations

Cisco.com

- **Rules for constructing the NEXTHOP attribute:**

  When two peers share a common subnet the NEXTHOP information is formed by a global address and a link local address

  Redirects in IPv6 are restricted to the usage of link local addresses

Routing Information

- **Independent operation**
  - **One RIB per protocol**
    - **e.g. IPv6 has its own BGP table**
  - **Distinct policies per protocol**
- **Peering sessions can be shared when the topology is congruent**

BGP next-hop attribute

- **Next-hop contains a global IPv6 address (or potentially a link local address)**
- **Link local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised)**

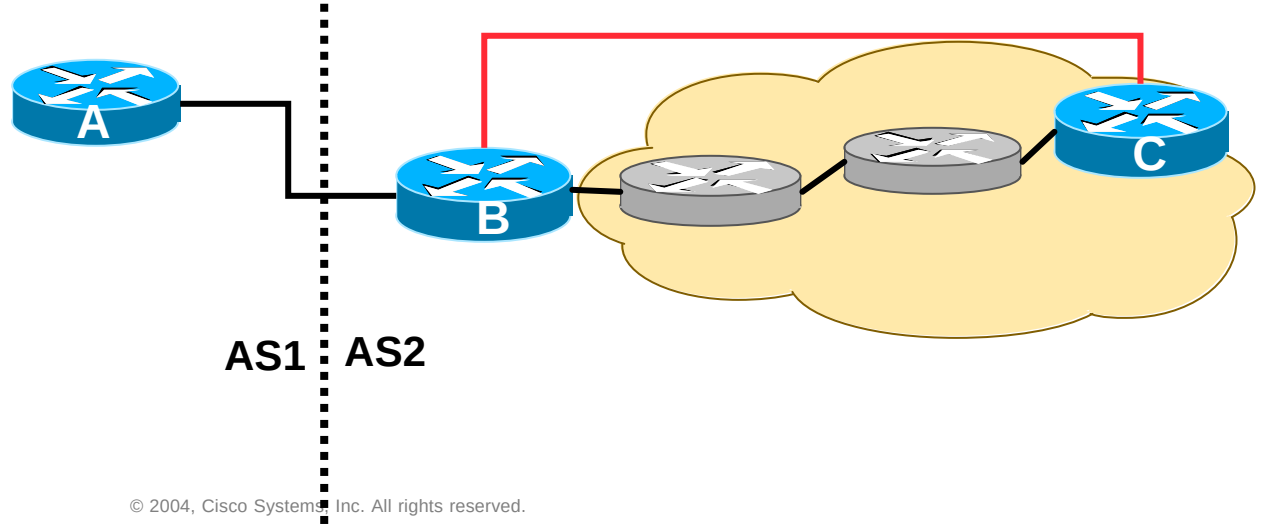

More BGP considerations

Cisco.com

- **TCP Interaction**

  BGP runs on top of TCP

  This connection could be setup either over IPv4 or IPv6

- **Router ID**

  When no IPv4 is configured, an explicit bgp router-id needs to be configured

    BGP identifier is a 32 bit integer currently generated from the router identifier – which is generated from an IPv4 address on the router

  This is needed as a BGP identifier, this is used as a tie breaker, and is send within the OPEN message

BGP Configurations
Non Link Local Peering

Router A

```
router bgp 1
 no bgp default ipv4 unicast
 bgp router-id 1.1.1.1
 neighbor 3ffe:b00:ffff:2::2 remote-as 2
!
 address-family ipv6
 neighbor 3ffe:b00:ffff:2::2 activate
 network 2003:3:2::/64
 network 2003:3:3::/64
!
```

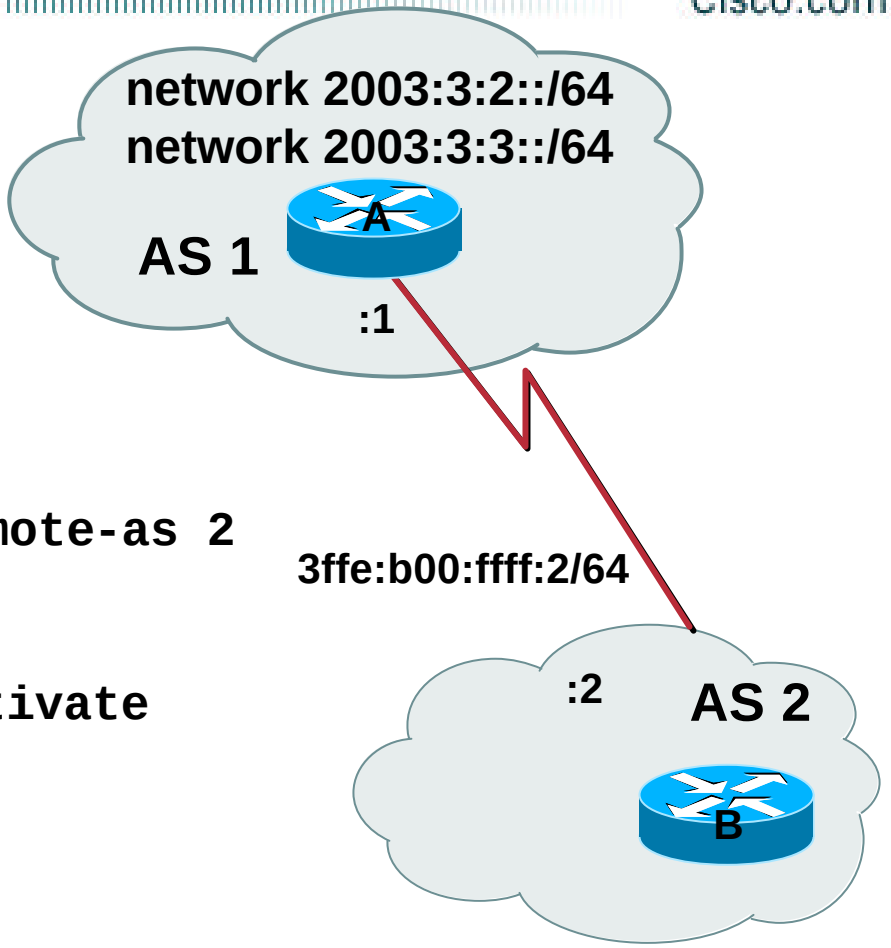

BGP Configurations
Link Local Peering

Router A

```
interface e2
 ipv6 address 2001:412:ffco:1::1/64
!
router bgp 1
 no bgp default ipv4 unicast
 bgp router-id 1.1.1.1
 neighbor fe80::260:3eff:c043:1143 remote-as 2
 neighbor fe80::260:3eff:c043:1143 update source e0
address-family ipv6
 neighbor fe80::260:3eff:c043:1143 activate
 neighbor fe80::260:3eff:c043:1143 route-map next-hop out
!
route-map next-hop
 set ipv6 next-hop 2001:412:ffco:1::1
!
```

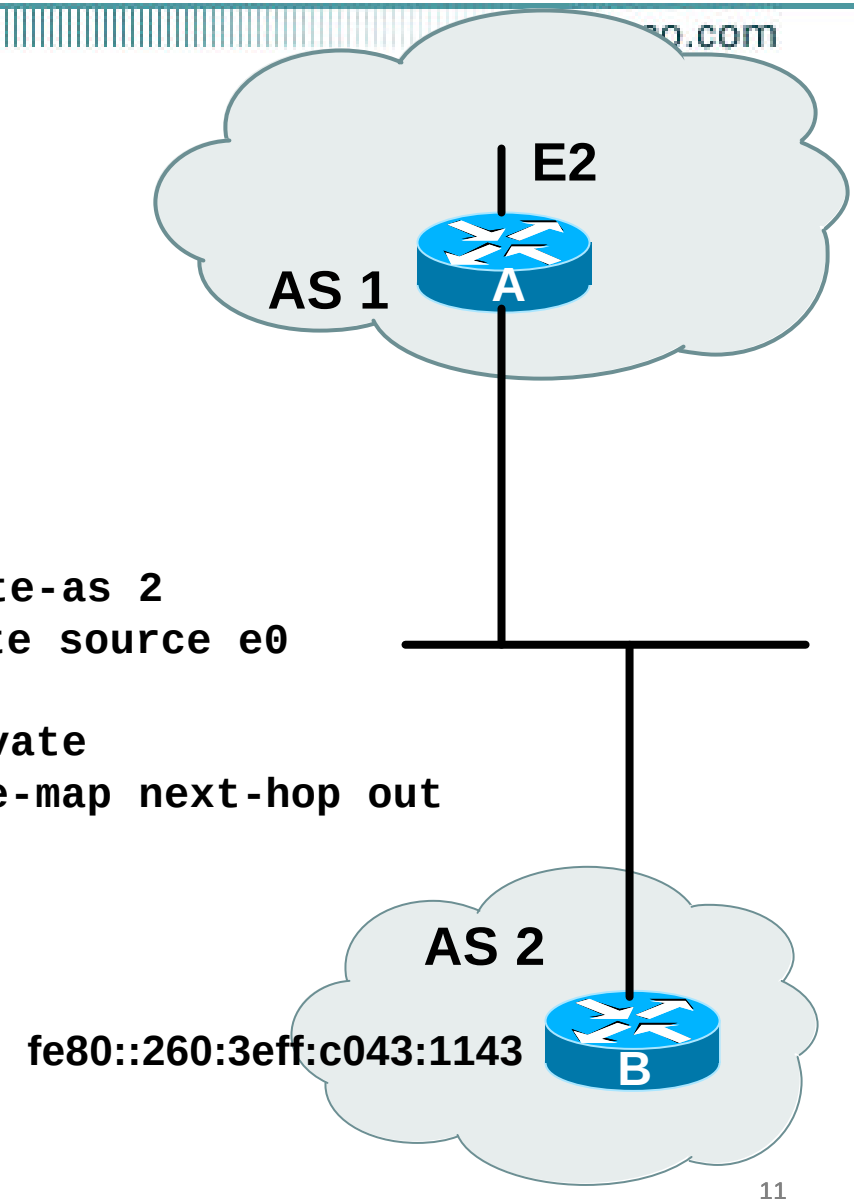

BGP Configuration
Filtering Prefixes

Cisco.com

- IOS Prefix-list is used for filtering prefixes in IPv4
  - And for IPv6 too!
- Example:

```
ipv6 prefix-list in-filter seq 5 permit 3ffe::/16 le 32
ipv6 prefix-list in-filter seq 6 permit 2001::/16 le 48
```

- Apply to the BGP neighbor:

```
router bgp 1
 no bgp default ipv4 unicast
 bgp router-id 1.1.1.1
 neighbor 3ffe:b00:ffff:2::2 remote-as 2
address-family ipv6
 neighbor 3ffe:b00:ffff:2::2 activate
 neighbor 3ffe:b00:ffff:2::2 prefix-list in-filter in
```

BGP Configuration
Manipulating Attributes

- **Prefer routes from AS 65002 (local preference)**

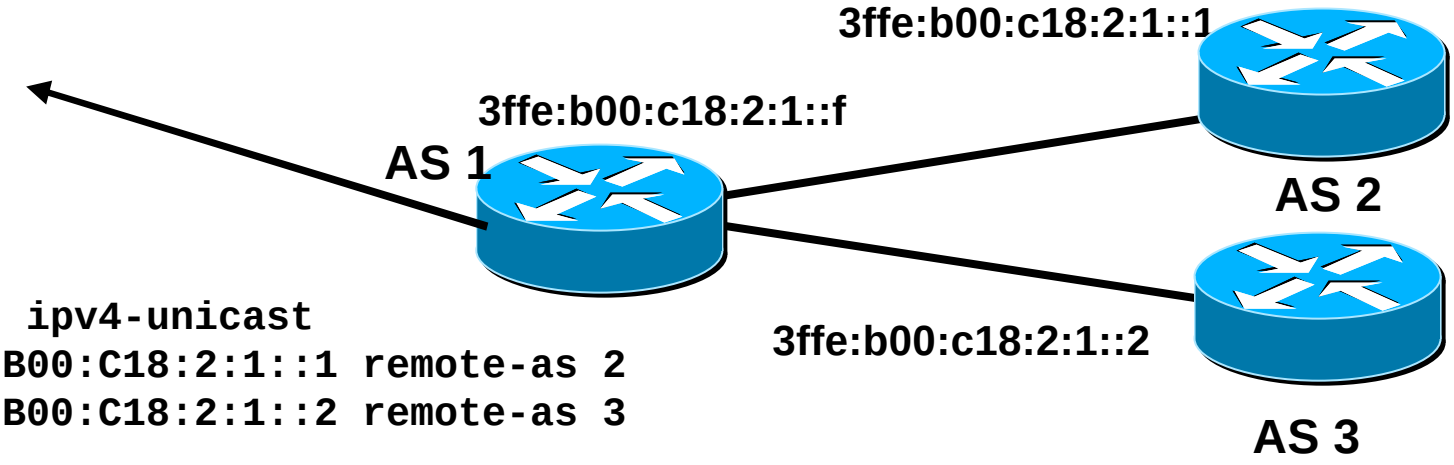


```
router bgp 1
 no bgp default ipv4-unicast
 neighbor 3FFE:B00:C18:2:1::1 remote-as 2
 neighbor 3FFE:B00:C18:2:1::2 remote-as 3
!
 address-family ipv6
 neighbor 3FFE:B00:C18:2:1::1 activate
 neighbor 3FFE:B00:C18:2:1::1 prefix-list in-filter in
 neighbor 3FFE:B00:C18:2:1::1 route-map fromAS2 in
 neighbor 3FFE:B00:C18:2:1::2 activate
 neighbor 3FFE:B00:C18:2:1::2 prefix-list in-filter in
 network 3FFE:B00::/24
 exit-address-family
!
route-map fromAS2 permit 10
 set local-preference 120
```

BGP Configuration
Carrying IPv4 inside IPv6 peering

- IPv4 prefixes can be carried inside an IPv6 peering

  Note that we need to “fix” the next-hop

- Example

```
router bgp 1
 neighbor 3ffe:b00:ffff:2::2 remote-as 2
!
address-family ipv4
 neighbor 3ffe:b00:ffff:2::2 activate
 neighbor 3ffe:b00:ffff:2::2 route-map ipv4 in
!
route-map ipv4 permit 10
 set ip next-hop 131.108.1.1
```

BGP Status Commands

Cisco.com

- **IPv6 BGP show commands take *ipv6* as argument**

  show bgp ipv6 *parameter*

```
Router1#show bgp ipv6 2017::/96
BGP routing table entry for 2017::/96, version 11
Paths: (1 available, best #1)
Local
    3FFE:B00:C18:2:1::1 from 3FFE:B00:C18:2:1::1 (10.10.20.2)
      Origin incomplete, localpref 100, valid, internal, best
```

BGP Status Commands

show bgp ipv6 summary

Displays summary information regarding the state of the BGP neighbours

```
RouterA# show bgp ipv6 summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 69046, main routing table version 69046
92 network entries and 92 paths using 17756 bytes of memory
826 BGP path attribute entries using 43108 bytes of memory
703 BGP AS-PATH entries using 19328 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
745 BGP filter-list cache entries using 8940 bytes of memory
BGP activity 22978/18661 prefixes, 27166/22626 paths, scan interval 15 secs

Neighbor        V    AS  MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
3FFE:B00:FFFF:2::2
                4    2   84194   14725    69044    0    0  3d08h           92
```

Neighbour Information **BGP Messages Activity**

Conclusion

- **BGP extended to support multiple protocols**
  - **IPv6 is but one more address family**
- **Operators experienced with IPv4 BGP should have no trouble adapting**
  - **Configuration concepts and CLI is familiar format**

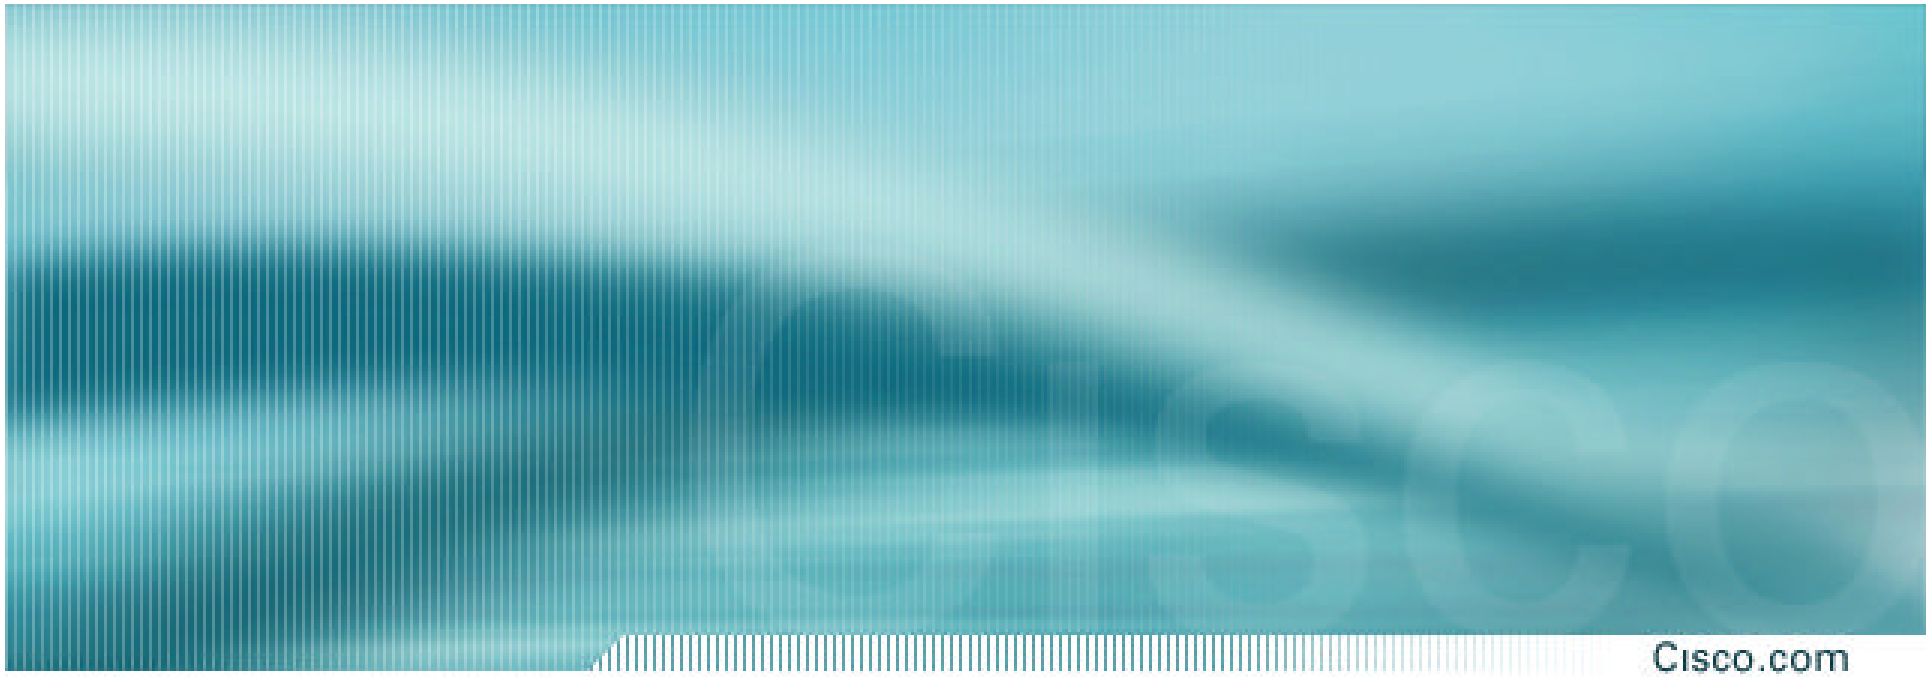

BGP enhancements for IPv6

Philip Smith <pfs@cisco.com>

1st South Asian IPv6 Summit/SANOG 3

15-22 January 2004

Bangalore, India