Bounding longest match considered

Problem Space

- Reduce growth in routing table
- Reduce BGP convergence time
- Provide for full reachability.
- Don’t break basic value of multiple upstreams/peers.

Basic Proposal

- An alternative to length-based filters
- Sets a bound on the prefix lengths eligible for preference.
- A bound would operate on long prefixes when covering route announcements are available.

Mechanism

- As each prefix is received from an external peer, it is evaluated vs. prefies already received. If two prefixes overlap in space the longer prefix would be marked (to indicate it should not be used for forwarding traffic), and propogated to its internal peers.
- On receiving a marked prefix, the path is compared to other paths towards the same prefix and prefix length.
- If the router is not advertising a bounded path for this prefix as well, it will block advertising of this bounded prefix from its Local-RIB to other BGP speakers within the AS
-

Mechanism (contd.)

- If the router is also advertising a bounded path for this prefix, it will compare router IDs with advertising router; the path originated by the BGP speaker with the lowest router ID will be preferred, so only one BGP speaker will advertise a bounded path into an autonomous system.
- Since the bounded path is not to be inserted into the Local-RIB,it will not be propogated outside the AS.

Failover and ASCII art

- The basic mechanism allows for longer prefixes to re-emerge after a covering aggregate disappears. It's much easier to follow with the ASCII art in the draft, though, and you wanted an excuse to read the draft, didn't you?

Advantages

- An AS reduces the number of prefixes carred internally by the number of longer prefixes that overlap with aggregates. One copy of the prefix is carried through the AS, but not placed in the forwarding table, nor propagated outside the AS.
- Customers still get benefits of multiple upstreams for traffic hand off and redundant inbound links
- Since aggregation is always preserved after 2nd hop, reduces routing table and risk of reachability loss

Costs

- Increases router processing
- May change effect of other traffic engineering mechanisms already in wide-spread use
- Risks transient lack of reachability during re-emergence of longer prefix.
- Risks transient black holes.

Comments, Questions?

- Ted.Hardie@nominum.com
- Ruwhite@cisco.com