Proxy Aggregation

Yakov Rekhter
IBM Research Division
T.J. Watson Research Center
yakov@watson.ibm.com

December 12, 1994

Why Aggregate ?

Observations:

- Routing information consumes resources (e.g. memory, CPU, bandwidth, people)
- Resources cost money
- Internet-wide connectivity is provided by a collection of independent, sometimes uncoordinated, quite often competing NSPs

Why Aggregate (cont.) ?

Assumptions:

- Need to match the volume of routing information with the available resources, while providing connectivity service – on a per provider basis

- Need to match the amount of resources with the utility of routing information – on a per provider basis

- Need to constrain routing information growth to less than linear with the size of the Internet:

  - O(# of nets) is not acceptable
  - O(# of sites) is not acceptable

What is Proxy Aggregation ?

- A mechanism to allow aggregation of routing information originated by sites that are BGP-4 incapable
- Being used primarily for site aggregation level
- Useful as a transition tool to a CIDR environment

Proxy Aggregation (cont.)

Current situation:

- Aggregation is controlled by the originator of the routes to be aggregated – destination has full control over the aggregation
- Destination has little (or no) direct benefits from the aggregation – aggregation requires extra work at the originator
- No feedback mechanism to promote aggregation

Observation: disconnect between the entities that directly benefit from aggregation and entities that control the aggregation.

Local Aggregation

Objective: to reconnect the entity that benefits from the aggregation with the entity that performs the aggregation

Local Aggregation (cont.)

- Routes that a domain (e.g. a provider) advertises to adjacent domains (e.g. providers, sites) form a part of the service the domain offers to the adjacent domains
  - The service domain A offers to domain B is controlled by a bilateral agreement between A and B
  - Routing information aggregation is controlled by the bilateral agreement
  - Domain's aggregation policies are established based on the bilateral agreements between the domain and its neighbors – prevents "over-aggregation"

Routing information aggregation performed by a domain becomes local to the domain matter.

Local Aggregation (cont.)

- Aggregation on entry – performed by a Border Router that receives routes from an adjacent domain
  - Action is purely local to the domain (but the reduction of routing information is beneficial to other domains as well)
- Aggregation on exit – performed by a Border router that advertises routes to an adjacent domain
  - Part of the bilateral agreement

From routing information reduction point of view aggregation on exit is better than aggregation on entry, but aggregation on entry provides more autonomy than aggregation on exit.

Potential Candidates for Local Aggregation

- Longer prefix in presence of a shorter prefix
- Adjacent CIDR blocks
- Aggregation over *known* holes (unallocated blocks)

Routing in presence of Local Aggregation

- The autonomous system number and the router that performed the aggregation is identified via BGP (AGGREGATOR attribute)

- A domain should register its aggregation policies in the RRDB

- A domain formulates its aggregation policies to reflect its bilateral agreements with adjacent domains