ALTO experiences with ISP-CDN collaboration

Hans Seidel

hseidel@benocs.com

BENOCS GmbH
Winterfeldtstrasse 21
10781 Berlin, Germany

July 21st 2016

Problem Statement

Scenario: CDN wants to deliver content to ISP customers

- Only paths from CDN caches towards customers of interest
- CDN caches embedded in foreign ASNs
- How to group prefixes to form PIDs?

Network Details

Large ISP Network to operate from:

- >900 Router
- >760k IPv4 Prefixes
  - >12k IGP Prefixes
  - >750k BGP Prefixes
    - >170k IBGP Prefixes
    - >580k EBGP Prefixes
- >20k IPv4 Ingress Prefixes
  - >950 Ingress Points
  - ~30% of public IPv4 Address Space

We have a live ALTO server providing guidance for CDN

Network Map

- Three different types of PID
  - Internal: AS internal prefixes
  - External: Prefixes belonging to foreign AS's
  - OnNet: Prefixes from directly peered AS's
    - → No third party traffic on peering link
    - → Only for specified ASN (OnNet ASNs)
- External and OnNet prefixes are provided by Ingress Point Detection

Ingress Point Detection

- Tool that detects ingress points of external prefixes
- Processes flow information (e.g. Netflow) for statistical evaluation
- Ingress points are interfaces of border routers connected to neighbored ASs
- Ingress point attributes:
  - Router + Interface
  - List of prefixes
  - Neighbor ASN (Handover ASN)
  - Link utilization

Cost Maps

Three different Cost Maps:

Hop Distance: Number of AS internal hops

Path Weight: Costs according to routing protocols

Custom: Home-brewed metric derived from delay and peering link utilization

Cost Calculation (already presented at IETF 93):

- Costs calculated between PIDs
- Tie breaker solves ECMP
- OnNet PIDs handled like Internal
- Outbound traffic is not considered
  - → No Egress Paths
  - → No Transit Paths

Statistics

- Updates every 5min
- Total Samples: >2000 (7 days)
- Network Map
  - >250k Prefixes
  - >1700 PIDs
    - >750 Internal PIDs
    - >950 External PIDs
    - ∼ 15 OnNet PIDs
  - Average Map Sizes:
    - Map: ∼ 6 MB (∼ 1 MB compressed)
    - SSE Patch: ∼ 1.7 MB (∼ 282 KB compressed)
- Cost Map (Custom)
  - >1.3M PID pairs
  - Average Sizes
    - Map: ∼ 47 MB (∼ 5.6 MB compressed)
    - SSE Patch: ∼ 37.5 MB (∼ 5 MB compressed)

Prefix Changes between Maps

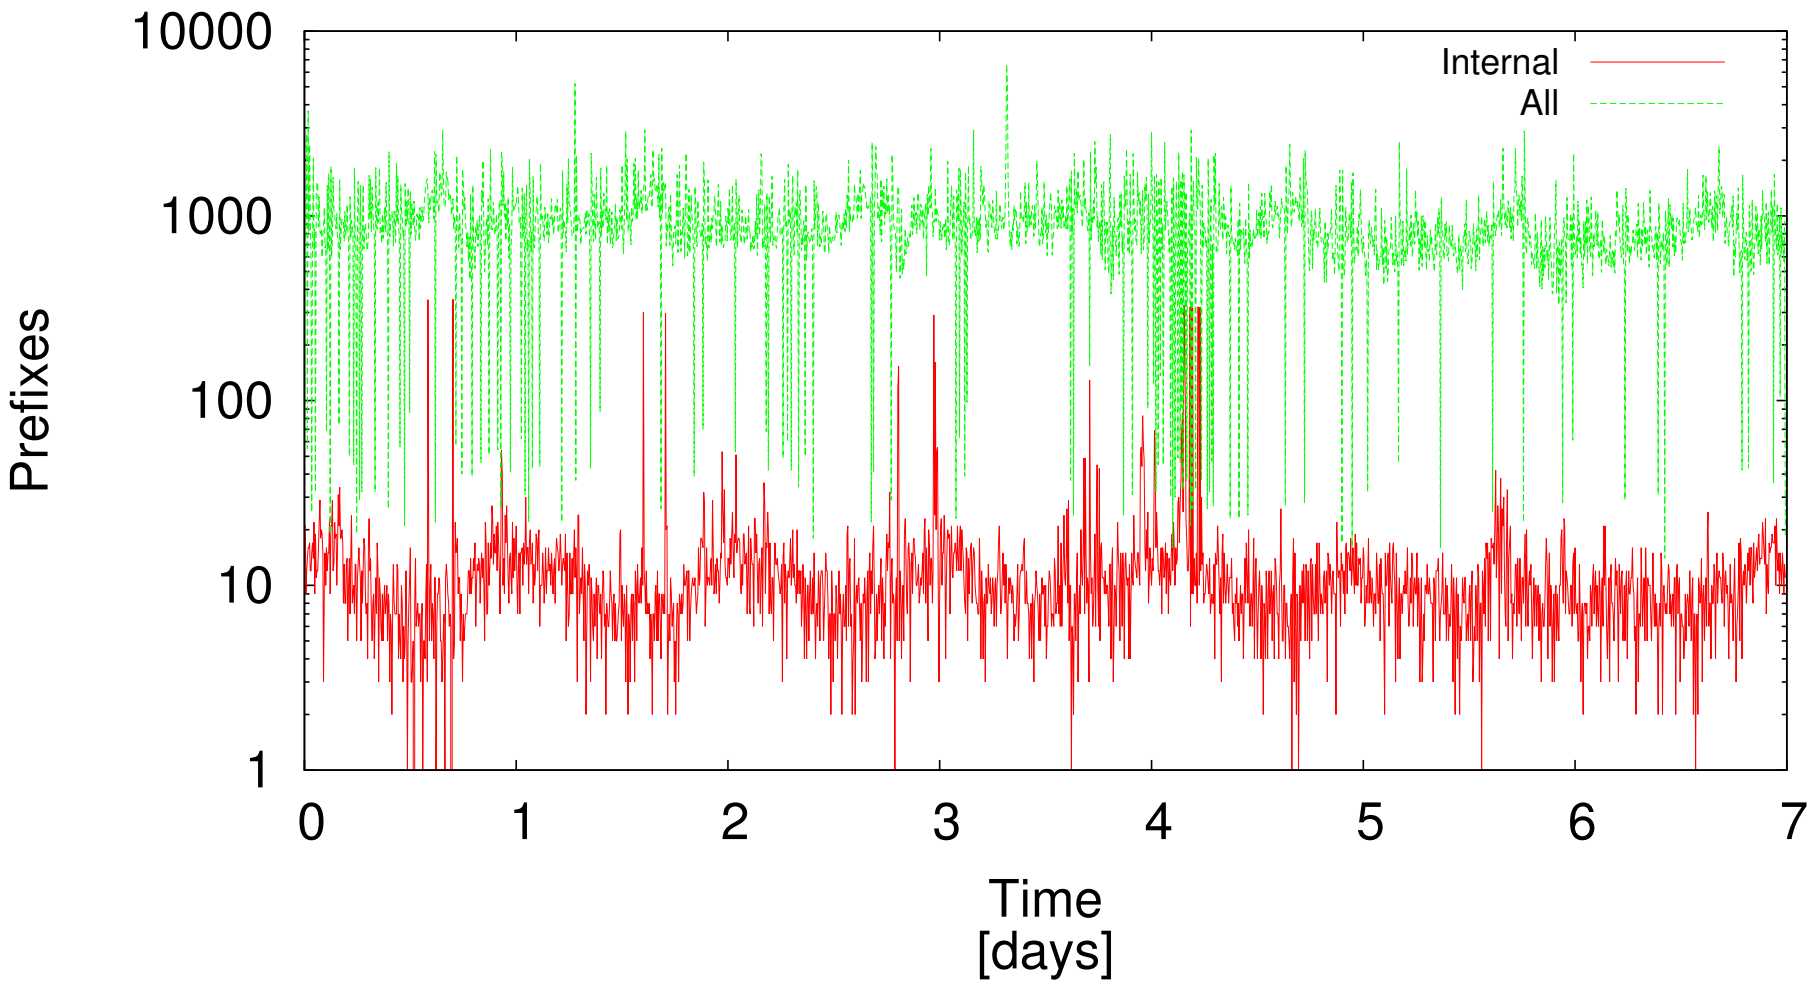

Problems

Long running map calculation process

- Cost Map calculation starts after Network Map is finished
- Network Map newer than current available Cost Map(s)
- Data inconsistency between Network and Cost Map possible

→ Solution: Mechanism that publishes all maps together when last is ready

Limitation to IP addresses in ECS request

- RFC7285 states that input data for source/destination must be addresses (/32 for IPv4 and /128 for IPv6)
- Difficult requesting ECS for regions

Empty source/destination field in ECS request

- RFC7285 states empty field is replaced with sender address
- Not suitable in ISP-CDN scenario

Modification/Features

- Prefix support in ECS
- Empty source/destination field in ECS filled with sender address
  $\rightarrow$ we let the server choose what to do
  E.g. Adding customer prefixes since the CDN does not know them
- Timestamps and TTLs as meta field in ALTO responses
- OnNet meta field to allow clients to provide OnNet ASNs

Thank you

Questions?

Prefix Changes between Maps

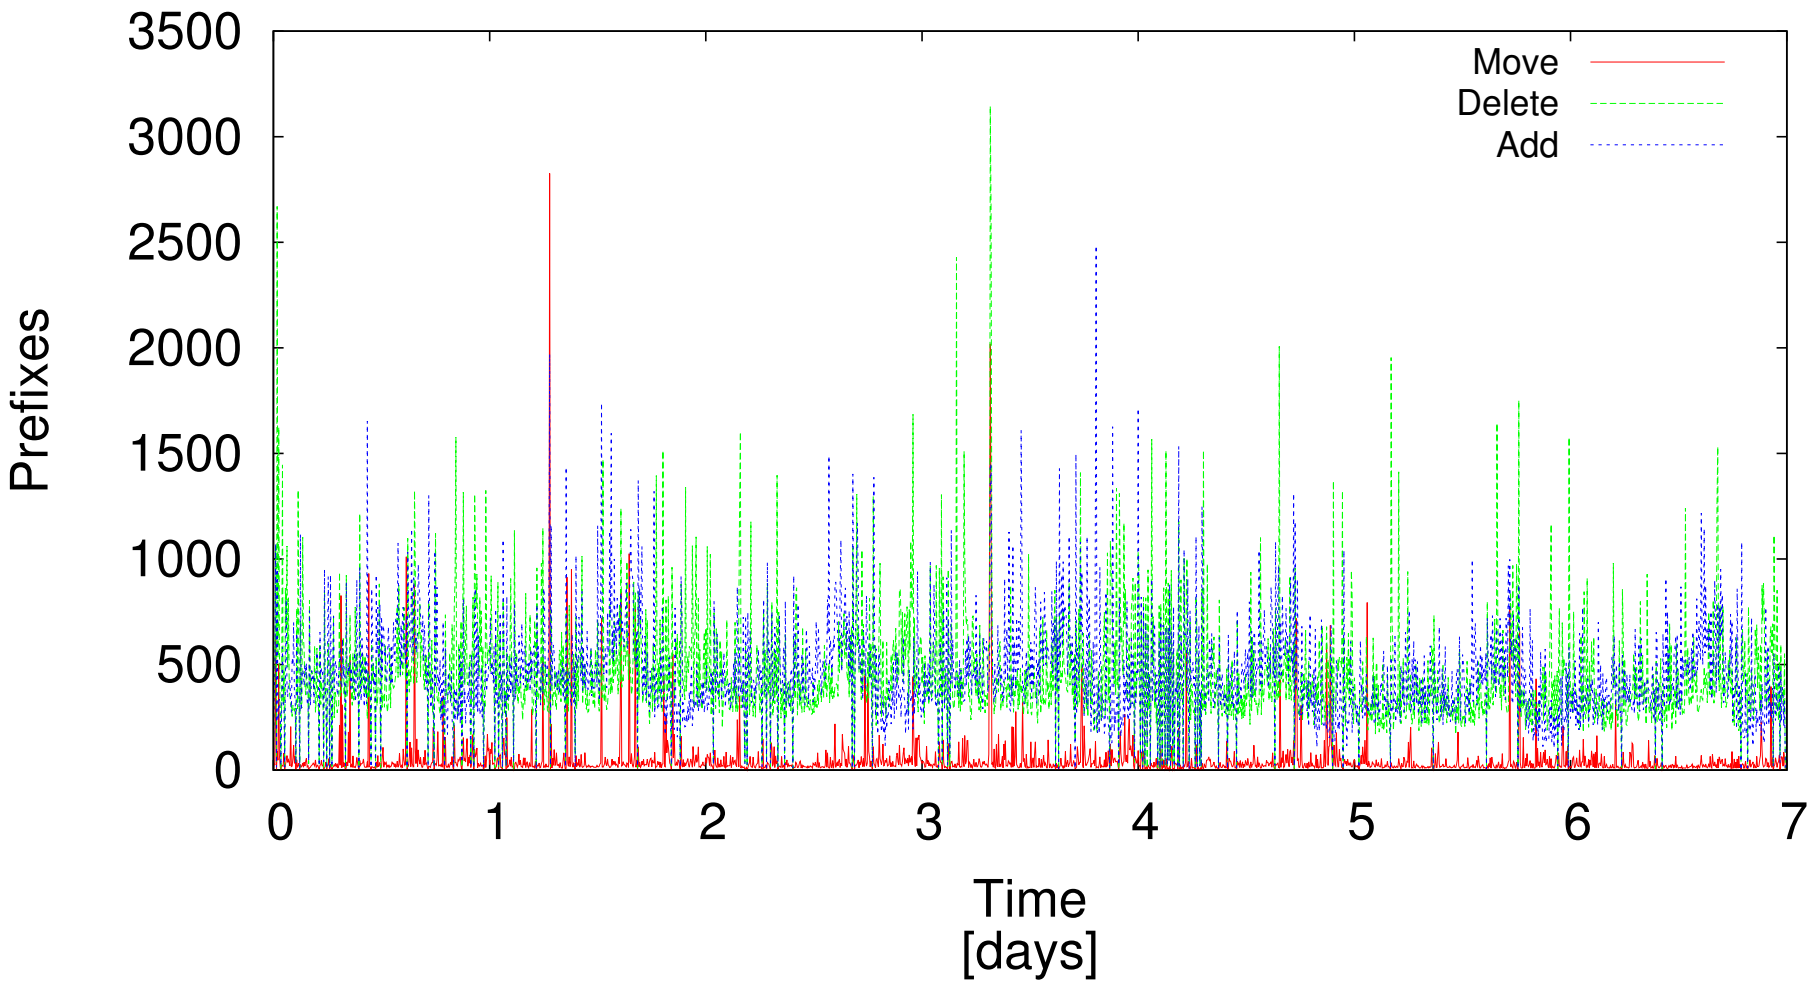

Internal PID

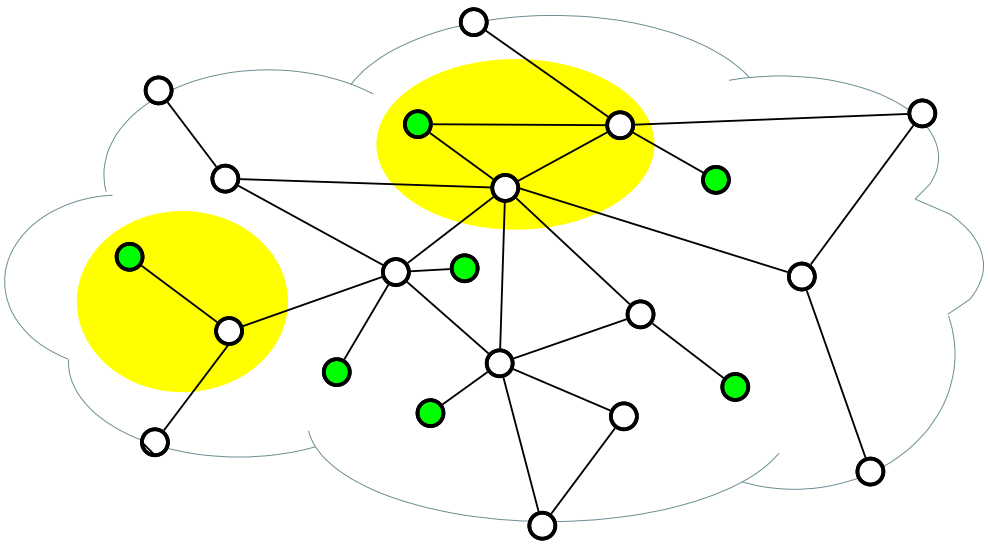

Internal Prefix

- Source: IGP/iBGP
- AS Distance of zero
- Grouped by Attachment Point(s)

OnNet PID

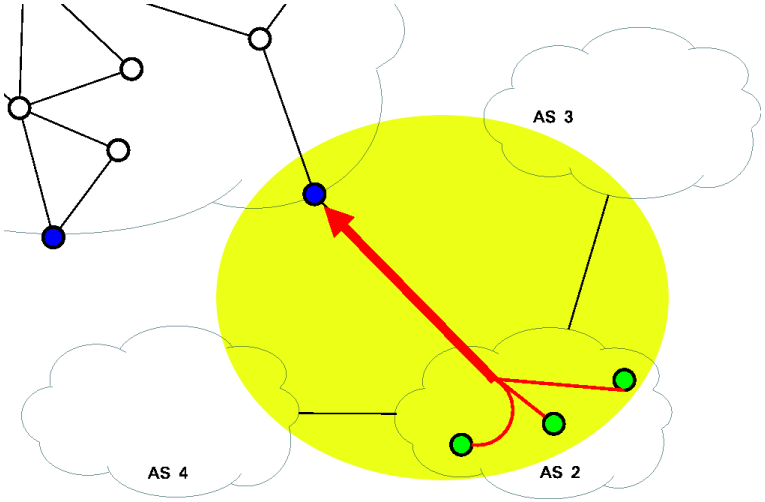

OnNet Prefix

- Source: Ingress Point Detection
- AS distance == 1
- Origin ASN == OnNet ASN
- Handover ASN == OnNet ASN
- Grouping by Ingress Router and ASN

External PID

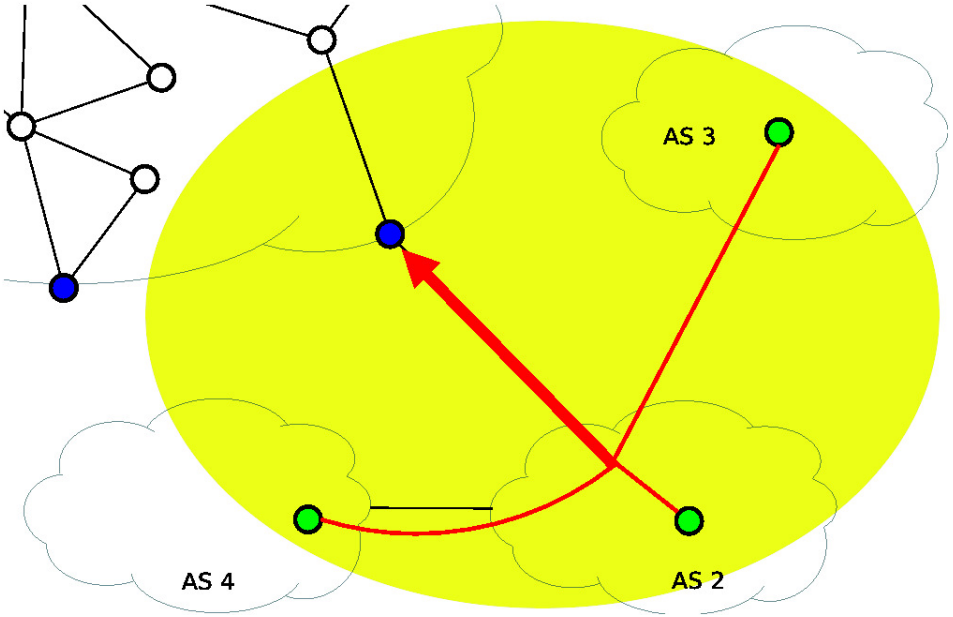

External Prefix

- Source: Ingress Point Detection
- AS distance >0
- Not OnNet
- Grouping by Ingress Point