



ALTO Southbound Based on IETF Mechanisms

Supporting Base ALTO Base Services

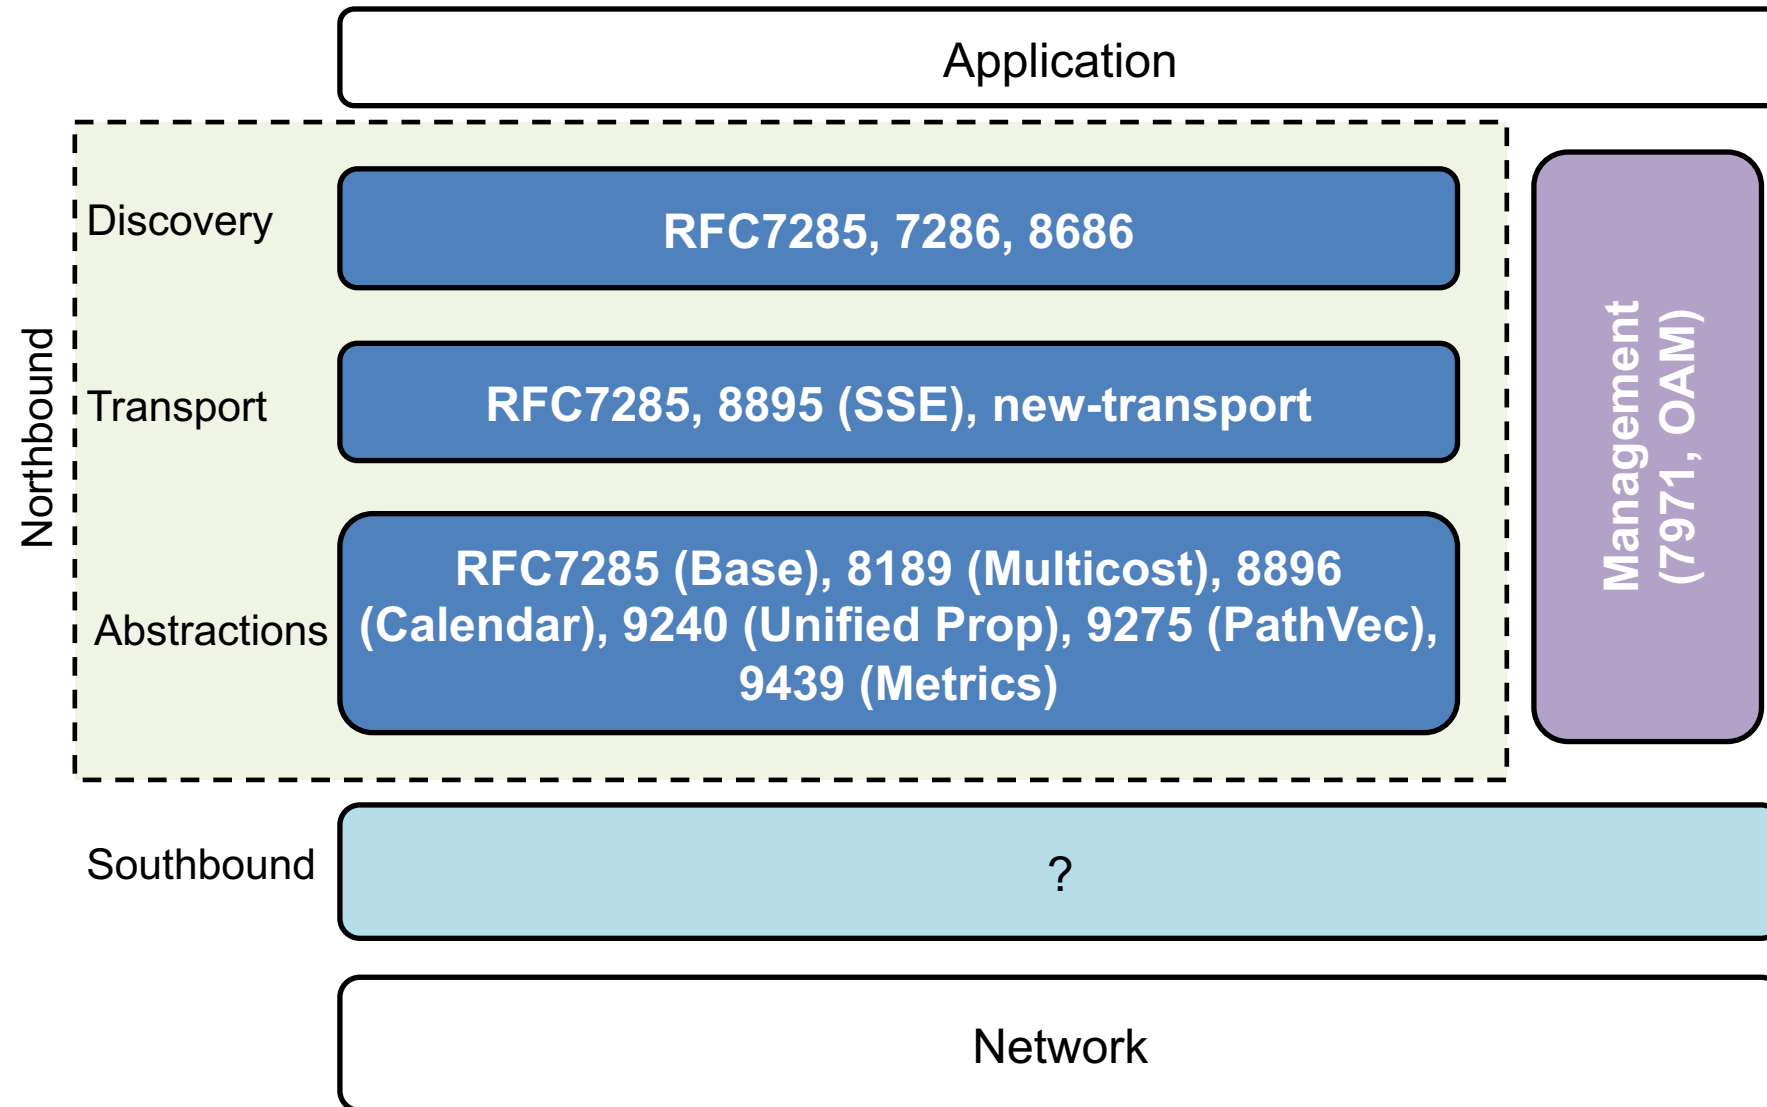
(Deployment Catalyst)

ietf116-interim-data-sources, draft-yang-alto-multidomain, ietf116-cascading-alto

IETF 117; July 27, 2023


Presenter: Sabine Randriamasy, Y. Richard Yang
on behalf of team

High-Level: ALTO Specifications



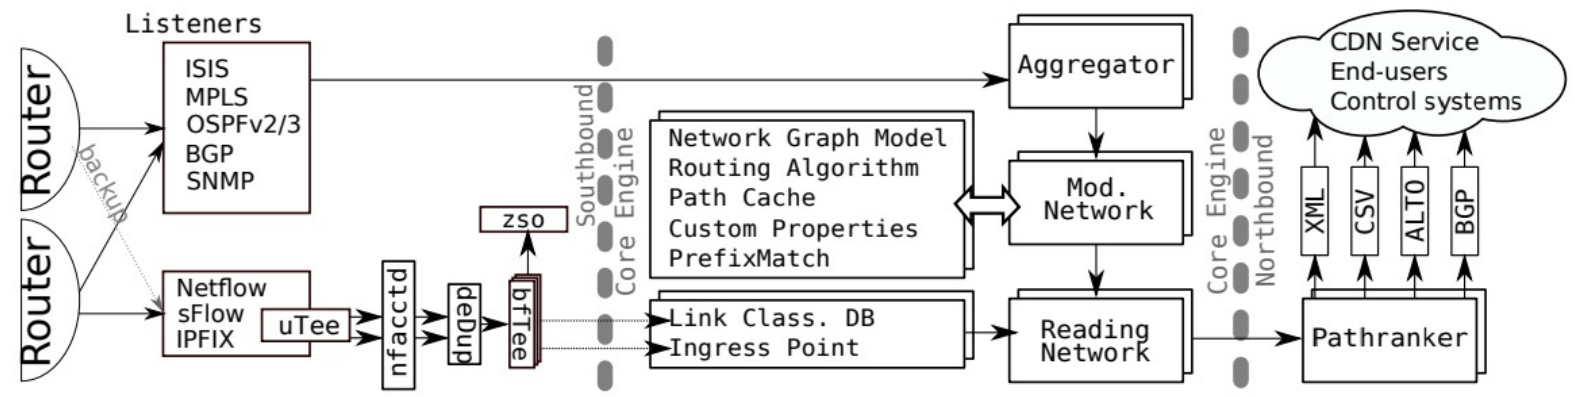
- Northbound: Relatively well specified
 - What (info) and how (to discover and transport)
- OAM: Good progress lately
- Southbound (how to obtain ALTO info)
 - **Some ALTO info (e.g., fancy cost maps) can be network secret sauce, but basic ALTO info (e.g., standard, e2e perf metrics) has a standard answer but is hard to compute, impeding ALTO deployment**
 - **The difficulty has become clearer in deployment efforts**

Southbound Complexity in Realizing Base ALTO Cost Service



- Basic southbound task deploying ALTO server for Net N: each link has a link metric m , compute aggregated e2e metric m for the src->dst path in N
- Basic subtasks:
 - Where is the ingress point?
 - What is the set of links traversed?
 - What is the metric of each link (if not totally static)?
- Even though Net N owns all devices, they can come from multiple vendors as closed source systems---need ways to extract the info

ALTO Southbound using IETF Mechanisms: Routing System Listener/Sampling Based Deployment



Flow Director Software Structure CoNEXT'19; PaDIS ACM SIGCOMM Software Systems Nomination

Subtask

Complexity and Issues

Obtain routing state

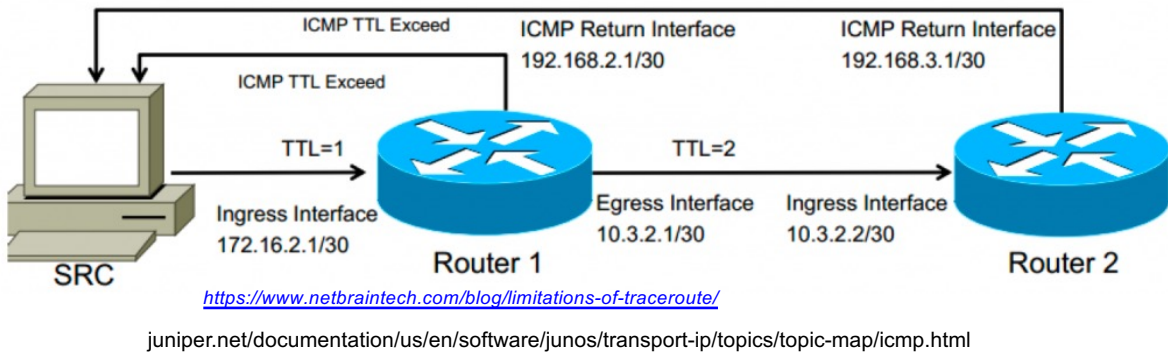
IGP passive peering (BGP-LS retrieval); handling network structures properly (e.g., areas), and replicating behaviors from topology data

Determine ingress


xFlow/IPFIX sampling/report; must have traffic, scalable concern

ALTO Southbound using IETF Mechanisms: Path Measurement Based Deployment

- Base ALTO server at CERN/WLCG, which uses PerfSonar data
 - Based on IETF-defined ICMP (RFC792), OWAMP (RFC4656) mechanisms
 - Many issues: e.g.,
 - ICMP has issues including node aliasing, (multi-)path aliasing;
 - PS servers are different from target endpoint locations (complex, limited anchoring algorithms)



Node aliasing



Path aliasing

Work Item Proposal (Related w/ This Presentation) and Participants

- Work item
 - Evaluate, guide, and propose solutions to ALTO southbound implementation and deployment barriers
 - Focus on computing ALTO base services providing ALTO **standard** metrics, using IETF standard mechanisms
- Participants/champions
 - CERN/WLCG ALTO server as driver: Y. Richard Yang, Sabine Randriamasy, Dong Guo, Jordi Ros Giralt, Jensen Zhang and Kai Gao
 - Benocs systems as driver: Danny Lachos and Ingmar Poese
 - Telefonica BGP-LS based ALTO server as driver: Luis M. Contreras and team

Some Details

- IETF standard mechanisms [to be advised by WG/AD]
 - standard routing systems (e.g., IGP) passive peering; BMP (RFC7854);
BGP-LS based topology export (RFC7752 and TE extensions draft-ietf-idr-te-lsp-distribution, rfc9085, RFC9351);
YANG topology models RFC8345/8795;
IETF-standard driven measurements integration such as IPPM.
- Ongoing, related, non-charter southbound implementations
 - Investigate robust, efficient, standard-based mechanisms (e.g., cascading ALTO) to compute ALTO cost metrics in a network for cross-domain traffic
 - Implement and deploy ALTO southbound based on hybrid topology (BGP-LS) and path retrieval (retrieve path segment)