The ALTO Path Vector Extension

draft-ietf-alto-path-vector-09

K. Gao[2] Y. Lee[3] S. Randriamasy[1] Y. R. Yang[5] J. Zhang[4,5]

[1]Nokia Bell Labs [2]Sichuan University [3]Sungkyunkwan University
[4]Tongji University [5]Yale University

Oct 20, 2019 @ IETF 106

Quick Summary

What new feature does this extension provide?

- Reveal **internal structures** and **detailed property information** from the ISP's point of view for **end-to-end application-layer communications**
  - e.g., bottleneck links, 5G UPF, MEC, service edge

Why is this extension essential?

- Such information is **useful** in many networking scenarios
  - deriving resource correlations of flows, e.g., shared risk resource groups and "co-flow" scheduling[1]
  - context-aware service selection and optimization, e.g., 5G UPF, MEC for cloud gaming, video streaming
- Such information is **fundamentally new** in ALTO
  - ALTO only has "cost" for (src, dst) pairs

[1] Chowdhury, M. and Stoica, I. 2012. Coflow: A Networking Abstraction for Cluster Applications. Proceedings of the 11th ACM Workshop on Hot Topics in Networks (New York, NY, USA, 2012), 31–36.

Quick Summary (2)

How does this extension provide such information?

- Internal structures: abstract network elements (ANE)
- Detailed property information: unified property map[2] for the ANEs
- End-to-end: ALTO cost map and endpoint cost services

What are the potential technical problems and how to address them?

- Representation issue: how to represent the internal structures?
  - ~~physical~~ v.s. abstract
  - persistent v.s. temporary
  - Decision: **abstract network element both persistent and temporary**
- Practical considerations
  - Scalability & consistency: one-round communication v.s. two-round communication
  - Complexity: design a new message format v.s. reuse ALTO message format
  - Decision: **one-round communication with a multipart response**

[2] Unified Properties for the ALTO Protocol, draft-ietf-alto-unified-props-new-09

Finalize the specification for cost type

- cost mode: array, cost metric: ane-path

Clarify the property negotiation process

- Available properties are announced in an IRD entry capability
- Selected properties are submitted in a query

Introduce persistent-entities property as an initial registry entry

- An array of entity identifiers that are persistent in the scope of an ALTO server

Clarify Part Resource ID (integration with SSE)

- Sync'd with SSE draft -16 (draft-ietf-alto-incr-update-sse-16)
- ResourceID of each part = Client ID + '.' + Part Resource ID

Propose solutions for cost calendar compatibility

- Flows only interfere in the same time interval
  $\Rightarrow$ The calendar results can be inferred from the PV of each interval
- Both correlations and properties may change over time
  $\Rightarrow$ Only make the PV part calendared (enough to represent both changes)

In -09 (a minor revision)

- We emphasize that ANE by design is dynamic to the query in multiple places in the document (in introduction, terminology, specification, etc.)
- We also highlight the benefits of on-demand dynamic ANEs
  - It reduces the information leaked to multiple queries
  - The ALTO server can use property-specific optimizations to compute ANEs

Remaining issues

Dependency on the UP draft

- Terminology from the UP draft (e.g., Entity, Entity Domain, etc., Sec 3)
- The property map part reuses the response data format from UP (in Sec 7.1.6 and 7.2.6)
- One property domain and two properties are registered using the UP registration procedure (which may lead to an IANA dependency, Sec 12)
- Sync'd with UP -08

Dependency on the SSE draft

- Sync'd with SSE -16
- SSE -17 includes multipart handling so the related part can now be removed from PV (to be done in the next submission)
- Terminology inconsistency (**part resource Id** in PV and **content Id** in SSE)

Writing

- Fix the dependency issues
- Improve the quality of writing
  - Need feedback from the WG

Heterogeneous ANE?

- Why
  - The Internet infrastructure has heterogeneous components already
  - Side meeting talks (e.g., cloud gaming) and some other IETF work (e.g., CFN) show that capability discovery is useful in network-aware end-to-end communication
  - ALTO PV can be used as a mechanism to expose capabilities for end-to-end communication
  - This strengthens the power of ALTO extensions and extends the scope of ALTO
- How
  - Define the entity type hierarchy for ANEs
  - The capabilities announced in IRD reuses the UP capabilities
- What follows
  - Identify ANE types (maybe work with other WG) and register the entity type, properties and their bindings to UP

Conclusion

- Current status
  - The motivations and potential problems are relatively clear
  - Most part of the specifications are relatively complete and stable
  - New inputs are received during IETF 106
- Great thanks to the coauthors and the WG for the feedback and guidance
- Next steps:
  - Make a revision
  - Set a milestone for WGLC? (Maybe IETF 107)
  - Call for reviews

Q & A

Join the Discussion at `alto@ietf.org`!

Questions and Comments are Welcome!