

IETF 110 [Virtual] ALTO Working Group

Chairs:

Vijay Gurbani vijay.gurbani@gmail.com

Jan Seedorf ietf@j-f-s.de

Qin Wu bill.wu@Huawei.com

- Online Agenda and Slides: <https://datatracker.ietf.org/meeting/110/session/alto>
- Data tracker: <http://datatracker.ietf.org/wg/alto/>
- Tools: <http://tools.ietf.org/wg/alto>
- Notes: <https://codimd.ietf.org/notes-ietf-109-alto>

Note Well

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (<https://www.ietf.org/contact/ombudsteam/>) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- BCP 9 (Internet Standards Process)
- BCP 25 (Working Group processes)
- BCP 25 (Anti-Harassment Procedures)
- BCP 54 (Code of Conduct)
- BCP 78 (Copyright)
- BCP 79 (Patents, Participation)
- <https://www.ietf.org/privacy-policy/> (Privacy Policy)

Also see: <http://www.ietf.org/about/note-well.html>:

Administrative

- Meetecho
 - Using meetecho queue control
 - Chat/Jabber available for use
 - Bluesheets
- Note taking
 - <https://codimd.ietf.org/notes-ietf-110-alto>
 - Please help with minute taking
(only discussion needs to be captured)
- Online Agenda and Slides at:
<https://datatracker.ietf.org/meeting/110/session/alto>
- Data tracker: <http://datatracker.ietf.org/wg/alto/>
- Jabber: <xmpp:alto@jabber.ietf.org?join>
- Tools: <http://tools.ietf.org/wg/alto>



Agenda

- **Introduction**

Session Intro & WG Status Chairs (5 minutes)

- **Recharter Discussion:**

Update of ALTO Implementation Collection Chairs (5 min)

- **ALTO Draft Recharter Discussion (30 min)**

- * Charter Goal and Scope Chairs (5 min)

<https://mailarchive.ietf.org/arch/msg/alto/FxRoRoesddhOhgkYhzQOyHjghPc/>

- * Proposed Charter item Use Case discussion All proponents (25 min)

- o Generic Protocol Extension for Policy Attribute Definition Sabine&Gang (5 min)

- o Protocol Extension for Pub Sub Mechanism Chunshan&Gang (5 min)

- o ALTO deployment for Operation Automation Luis (5 min)

- o Operation Automation Data Model Dhruv (5 min)

- o Multi-Domain Setting Richard&Ingmar (5 min)

- * Open Discussion on Charter: All (15 minutes)

- **Wrap up: (5 minutes)**

Document Updates

- Post WGLC

Document	↕ Date	↕ Status	↕ IPR	↕ AD / Shepherd
Active Internet-Drafts (4 hits)				
draft-ietf-alto-cdni-request-routing-alto-16	2021-01-12	I-D Exists		Vijay Gurbani
Content Delivery Network Interconnection (CDNI) Request Routing: CDNI Footprint and Capabilities Advertisement using ALTO	41 pages	WG Consensus: Waiting for Write-Up <i>Jul 2020</i>		
draft-ietf-alto-path-vector-14	2021-02-22	I-D Exists		Vijay Gurbani
ALTO Extension: Path Vector	51 pages	Held by WG <i>Jul 2020</i>		
draft-ietf-alto-performance-metrics-15	2021-02-04	I-D Exists		Jan Seedorf
ALTO Performance Cost Metrics	33 pages	Held by WG Reviews: tsvart <i>Jul 2020</i>		
draft-ietf-alto-unified-props-new-16	2021-02-22	I-D Exists		Martin Duke
ALTO extension: Entity Property Maps	57 pages	Held by WG <i>Jul 2020</i>		Vijay Gurbani

Status of drafts

- **draft-ietf-alto-unified-props-new** and **draft-ietf-alto-path-vector**
 - Chair review done by vkg.
 - Shepherd writeup done (shepherd: vkg).
 - Shepherd to move the drafts to IESG after reviewing changes resulting from chair review.
 - Thanks to authors for the revisions.

Status of drafts

- **draft-ietf-alto-performance-metrics**
 - -15 released on 2021-02-04, addressing comments and issues previously identified from WGLC.
 - IPR check done.
 - Shepherded writeup needed (Jan).

Status of drafts

- **draft-ietf-alto-cdni-request-routing-alto**
 - Waiting for shepherd writeup.
 - Token with vkg.

Update of ALTO implementations

• ALTO Implementation/deployments:

- **Benocs:** ALTO implementation in Telekom
 - Key component: Flow Director
 - Objective: Optimize CDN content delivery
 - Deployment Details:
 - Number of machines 2
 - Cores/threads/hyper-threads 16/32/64 (Xeon E5)
 - RAM > 700 GB
 - Number of IPv4/IPv6 routes ~850 k / ~680 k
 - NetFlow records > 45 billion per day
 - NetFlow rate (peak) > 1.2 Gbps
 - BGP peers > 600
 - CoNEXT 2019 Best Paper Award; IETF/IRTF 2020 Applied Networking Research Prize
- **China Mobile/Tencent:** Mobile and Wireless Information Exposure (MOWIE)
 - Objective: Improve User experience for Cloud Gaming Application
 - NAA QoE POC Test have been conducted.
- **Sprint/T-Mobile:** ??
- **Telefonica (Telco CDN Provider):** integrate ALTO with the Telefonica CDN
 - Objective: Keep track of the network topology change and Surrogate selection in the CDN network
 - Plan to move to production in 2021 H2.
- **Unicorn:** Multi-Domain, Multi-Controller Resource Orchestration
 - Key component: multi-domain orchestration
 - Objective: Discovery service, resource, topology info from multiple domain involved in the federation
 - Implemented in Opendaylight platform, demonstrated in SC 18, 19 in small federation networks connecting Dallas, Texas, Los Angeles, California, and Denver, Colorado
- **Cisco:** Network Position System information Exposure
 - Objective: Optimize SP resources utilization while improving service delivered to CDNs, applications and OTT overlays.
 - Key component: Network Data Accumulation (Policy data, Geo-location data, Topology data, etc.)
 - Source: <https://meetings.apnic.net/31/pdf/NPS-sprevidi-201102-APRICOT.pdf>
- **Juniper:** Content/Service request routing using ALTO
 - Objective: Locate best copy of content for the end user
 - Source: <https://archive.nanog.org/meetings/nanog53/presentations/Monday/Ward.pdf>


Update of ReCharter Discussion

- Recharter discussion started from IETF 108 and continue in IETF 109
 - 5 proposed work items were discussed
 - Generic Protocol Extension
 - ALTO Cellular Network information Service
 - New Transport
 - Multi-domain ALTO
 - Operation automation for ALTO
 - It was suggested to address overlapping among new work items and existing work item
- Recharter scope and goal, problem space documenting
 - Overlapping issue among new work items was revisited
 - The first version of Recharter text was documented and sent to AD and co-chair in December 30 for feedback collection
- ALTO meeting is planned for third round charter review.
 - Syncup meeting with AD and co-chairs was scheduled to set the goal for the meeting.
 - Circulate the draft charter on this list by 22 Feb
 - Collect ALTO implementation/deployments information
 - "consensus in the room" on the charter items
- ALTO Charter proposal has been posted on the list for public review.
 - 10 cases have been investigated.

ReCharter Goal and Scope

Basic Goal:

- Provide network information that an application may not easily get by itself
- Provide abstraction to simplify complex network information/internals



Recharter Goal:

1. To support wireless network info exposure, property and cost extension to approach real time indication is needed.

2. To support Rate adaptation for Cloud Gaming and wireless network info exposure, fine granularity pub sub mechanism is needed, cost metrics and entity property extension is needed.

3. To support multi-domain and operation automation, data model for ALTO server and client is needed.

4. ALTO deployment in RFC7971 is outdated and needs to be enhanced with latest protocol extension to support emerging new applications

5. To support multi-domain setting, Path vector needs to be extended for multi-domain application (e.g., SD-WAN, distributed computing)

Proposed Chartered items and Use Cases

Proposed Chartered item	Use Case(s)	Proposed extension and mechanism
Item 1: Generic Protocol extension for policy attribute	<ol style="list-style-type: none"> 1. Rate Adaptation for Cloud Gaming 2. Cellular Network Information Exposure 	<ol style="list-style-type: none"> 1. Extensions to the Property and Cost service 2. New properties and metrics on QoS profile
Item 2: Protocol extension for Pub Sub Mechanism	Rate Adaptation for Cloud Gaming	<ol style="list-style-type: none"> 1. Fine granularity pub sub 2. Pub sub extension to support HTTP 2.0, HTTP 3.0
Item 3: ALTO Data model	3 ALTO system management and operation	Data model for ALTO client and server configuration And ALTO system state monitoring
Item 4: Multi-Domain Setting	<ol style="list-style-type: none"> 4. Multi-Domain Resource Orchestration 5. Hierarchical service function chaining in the Distribute DC Network [RFC8459] 6. Muilti-Domain SD-WAN Edge Discovery 	<ol style="list-style-type: none"> 1. Multi-Domain Server Discovery 2. Path Vector extension for Multi-Domain Connectivity discovery
Item 5: Operation automation	<ol style="list-style-type: none"> 7. ATLO Deployment enhancement [RFC7971] 8. Compute Network Edge Service Discovery 9. ALTO Aggregation 10. CDN overlay/underlay integration 	<ol style="list-style-type: none"> 1. Automatic IRD generation 2. Network and Computing information Aggregation 3. Aggregation from LSPDB/TEDB/RIB/FIB

Related work: Alternative QoS profiles and Notification control in 3GPP

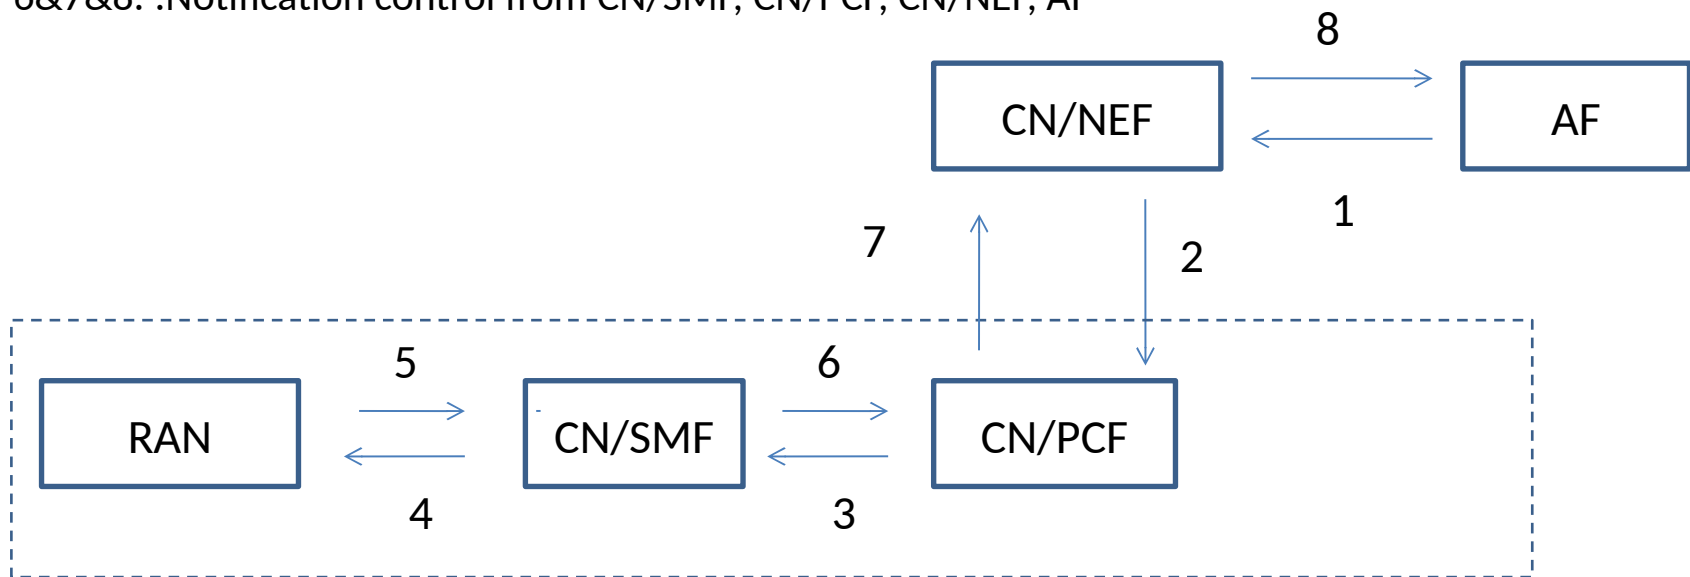
1&2 : QoS reference, (optional) Alternative Service Requirements (containing one or more QoS reference parameters in a prioritized order)

3&4: QoS profile and any alternative QoS Profile in a prioritized order

5: Notification control from NG-RAN to CN (GFBR guaranteed, not guaranteed)

- The NG-RAN should always try to fulfil the QoS profile and any Alternative QoS Profile that has higher priority than the currently fulfilled situation
- In order to avoid a too frequent signalling to the SMF, it is assumed that NG-RAN implementation can apply hysteresis (e.g., via a configurable time interval), PCF has ensured that the QoS values within the different Alternative QoS Profile(s) are not too close to each other

6&7&8: :Notification control from CN/SMF, CN/PCF, CN/NEF, AF



Legend:

RAN:Radio Access Network; CN :Core Network SMF :Session Management Function GFBR: Guaranteed Flow Bitrate
NEF: Network Exposure Function AF : Application Function PCF: Policy Control Function


Use Case 1: Rate Adaptation for Cloud Gaming

1. [ALTO] Multiple QoS requirements based on ABR in a prioritized order

2&3&4&5&6&7: refer to 3GPP


8. [ALTO]Event based notification for each QoS requirement (guaranteed, not guaranteed), **which will be used for cloud gaming server to adapt bitrates**

- Note: significantly reduce the signaling overhead, maybe > minutes level



7

2



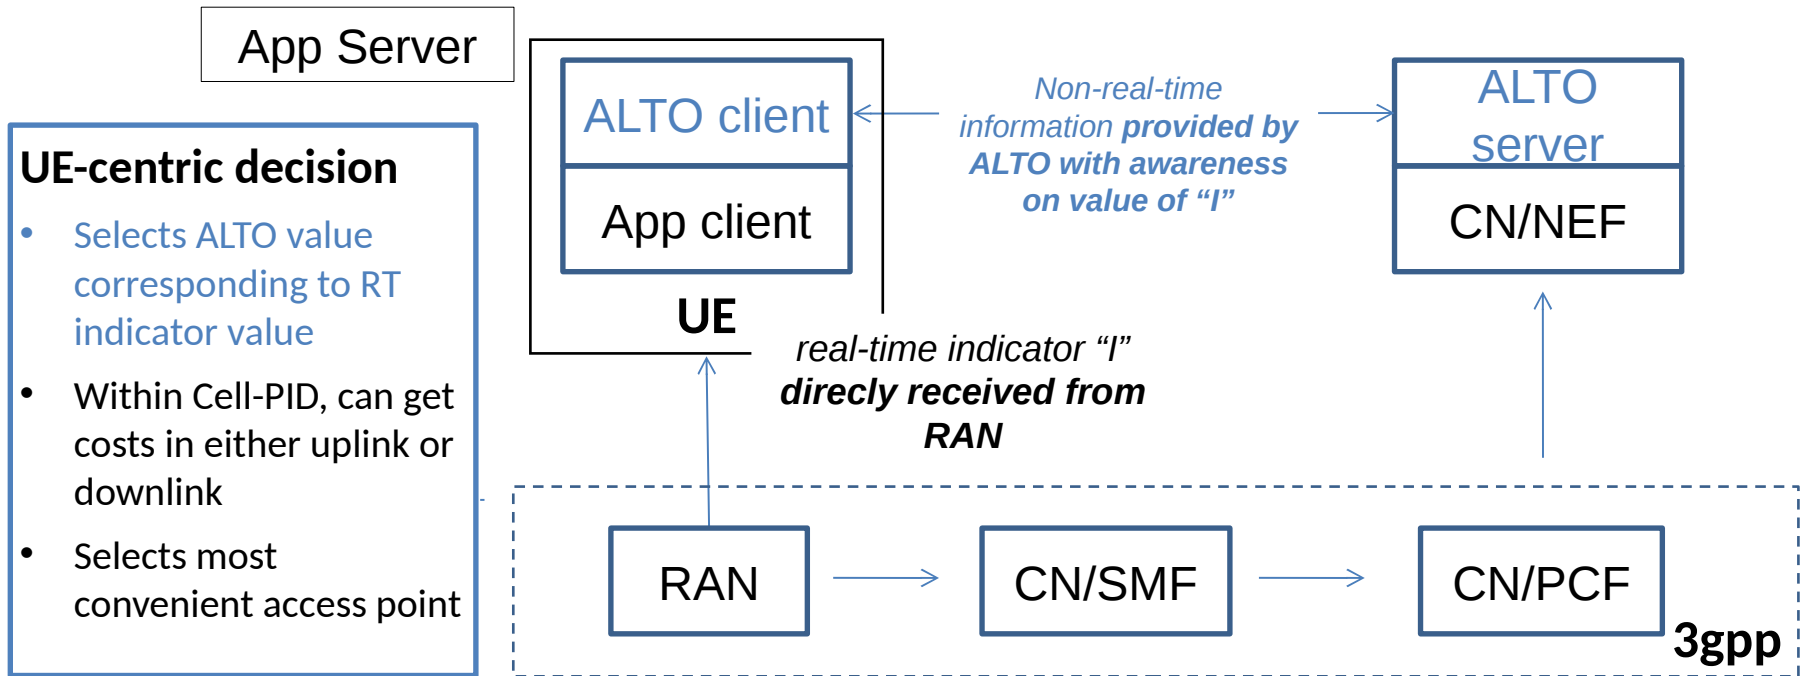
- Corresponding work items:

- Work item 1: Generic Protocol Ext
- Work item 2: Pub Sub Mechanism

- Proposed Protocol Extension:

- Define ALTO extension for pub sub mechanism and allow subscribe to QoS profile change (e.g. for 1080P, 720P) and receive notification for QoS requirements changes
- Define new policy attribute in the ALTO interface and allow the client to express QoS requirements such as performance objectives (throughput, latency, priority, error rate etc.), golden, silver, bronze standard profile, real time data indication or not)

Use Case2: cellular network information exposure

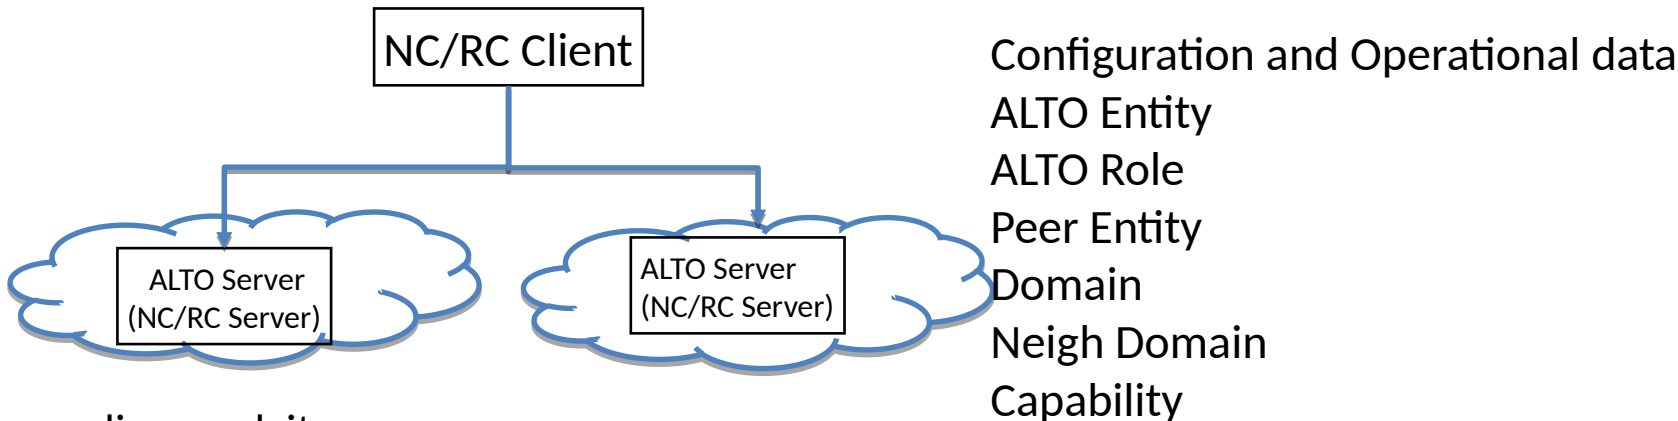


- Corresponding work item: General Protocol Extensions
- Proposed Protocol Extensions – **ALTO scope in blue**
 - Define policy attributes for ALTO information queries to indicate dependency of ALTO path-costs & properties on e.g. real-time (RT) network indicators, link direction,
 - Example RT indicators: udi, others correlated to cell load, ...
 - Example attribute: uplink, downlink
 - **Can apply to Server centric decision making as well**

Use Case 3: ALTO Data Model

In case ALTO client and servers are enabled on devices that already uses YANG based mechanism for configuration, status and monitoring or multiple domain server to server communication is enabled, **then providing a standard YANG data model for ALTO is needed.**

- This data model should support configuration of the ALTO client as well as ALTO server.
 - Similar to draft-ietf-netconf-http-client-server-06
- The data model needs to also support the operational status as well as monitoring.
- This is not to be used as a replacement of ALTO protocol.

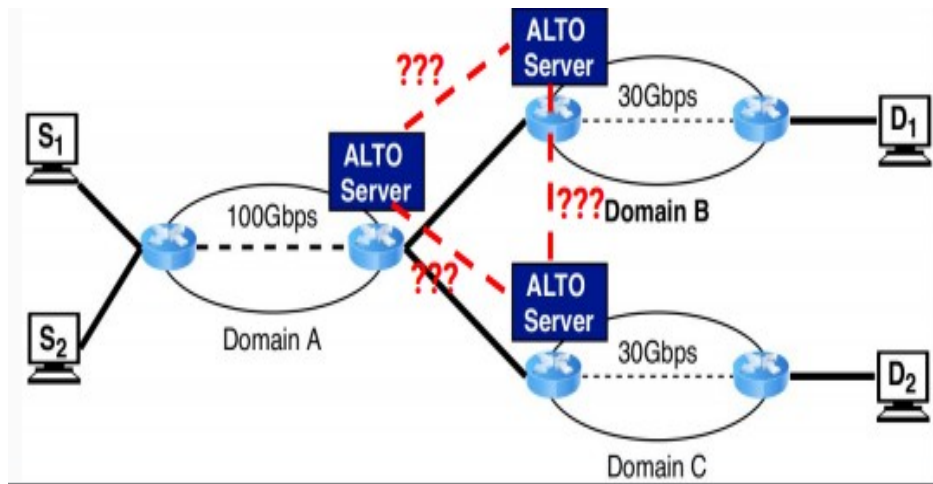



- Corresponding work items:
 - Work item 3: Data Model
- Proposed protocol extension or building block
 - YANG data model definition

Use Case 4,5,6: Multi-Domain Setting

- Server-to-Client ALTO communication is not enough.
- It is necessary multi-ALTO server communication to allow exchanging network information from multiple domains.
- The ALTO protocol specification states [RFC7285]:
 - "It may also be possible for an ALTO server to exchange network information with other ALTO servers (either within the same administrative domain or another administrative domain with the consent of both parties) ...".
 - However, such a protocol is outside the scope of the specification

- ALTO may consider either a **hierarchical** or **mesh** architectural deployment design [INTER-ALTO][MERCATOR][SFC-MD]
 - **Hierarchical design**, ALTO servers in domain partitions gather local information and send it to central server.
 - **Mesh deployment**, ALTO servers may be set up in each domain independently, and gathering the network information from other connected domains.



- Corresponding work items:
 - Work item 4: multi-domain setting
- Proposed protocol extension
 - Multi-domain ALTO server discovery: e.g., BGP extension or RFC8686
 - ALTO server communication for network information and compute information exchange
 - Multi-Domain Connectivity discovery:
 - Using BGP-LS export traffic engineering information with external domains using the BGP routing protocol (Not in the scope).
 - Define mechanisms where a ALTO entity cooperates either with other ALTO entities in adjacent domains or with a parent ALTO entity.

Use Case 7,8 : Operation Automation

Goal: Enhanced / enabled automation of network operation by integrating ALTO information with different purposes, assisting Network control entities to perform decisions and trigger informed actions

Corresponding Work Items -> Work Item 5

Case 1

New ALTO protocol extensions (e.g., the incremental update mechanism, cost calendar, unified properties, and path vector) require new architectural and deployment considerations.

Investigate best practices (extending RFC7971) including:

- Automatic IRD generation
- Aggregation of different information sources

Case 2

ALTO deployment enhancement with new protocol extension

- Aggregation of network information and compute information
- ALTO protocol extension for edge **selection combining network and compute info**

Common Network Function Virtualisation Infrastructure Telecom Taskforce (CNTT)

The screenshot displays the CNTT Compute Flavours interface. It features two main profiles: 'Basic Profile' and 'Network Intensive Profile'. Each profile has 'Interfaces Options' and 'Compute Flavours' sections. The 'Basic Profile' includes options like 'tiny', 'small', and 'medium'. The 'Network Intensive Profile' includes options like 'large', '2xlarge', and 'Axlarge'. A blue arrow points from the interface to the table on the right.

Flavor Name	Type of instance (T)	Interface Option (I)	Compute flavor (F) (CPU, RAM, disk and bandwidth)	S.	A.
Small-1	Basic	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1,512 MB,1 GB,1 GBps}
Small-2	Network Intensive	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1,512 MB,1 GB,1 GBps}
Medium-1	Network Intensive	{25, 50, 75, 100, 125, 150 Gbps}	{2,4 GB,40 GB,1 GBps}
Large-1	Compute Intensive	{50, 100, 150, 200, 250, 300 Gbps}	{4,8 GB,80 GB,1 GBps}
Large-2	Compute Intensive	{100, 200, 300, 400, 500, 600 Gbps}	{8,16 GB,160 GB,1 GBps}

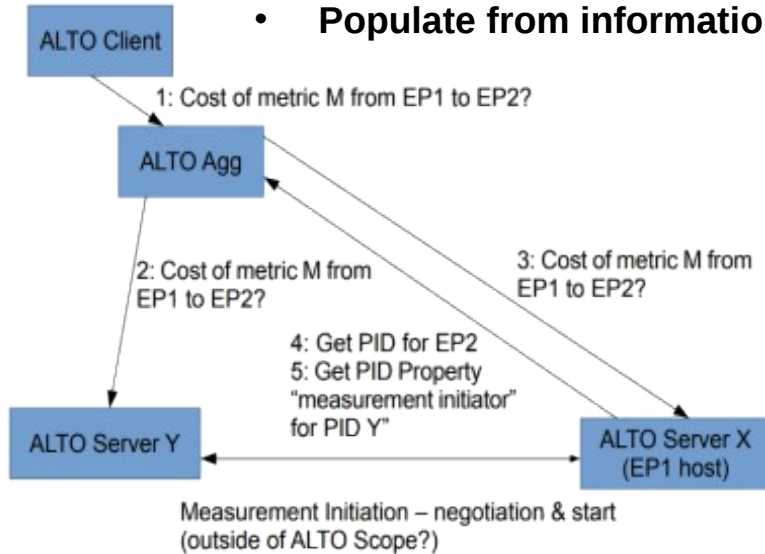
- Multiple data centers of different sizes across the network of distinct sizes (CPUs, memory, storage, bandwidth, etc)
- Identify the proper DC for a given application considering both compute and transport substrates

Use Case 9,10 : Operation Automation

Case 3

ALTO deployment enhancement with new protocol extension

- Aggregation of measurement information from the underlying network
- Reactive ALTO information resources update (trigger in time measurement)
- **Populate from information collected from TEDB, LSPDB, etc.**



UC 1 : Data is present in underlying Server BUT not visible in ALTO Server (implies a Filter) => ALTO MUST support Server side filtering (On Demand Measurements is proposed)


UC 2: Data is not present server at all => ALTO MUST support some form of Measurement Initiation

UC 3: Metric not currently supported by ALTO Server => Dynamic loading of measurement data

Case 4

Overlay / underlay integration

- Provide and maintain proper insights of underlay capabilities and status to overlay constructs in order to optimized service delivery
- Example: CDN delivery



Charter Proposal

- The ALTO working group was established in 2008 to devise a request/response protocol to allow a host to benefit from a server that is more cognizant of the network infrastructure than the host is.
- The working group has developed an HTTP-based protocol and recent work has reported large-scale deployment of ALTO based solutions supporting applications such as content distribution networks (CDN).
- ALTO is now proposed as a component for cloud-based interactive applications, large-scale data analytics, multi-cloud SD-WAN deployment, and distributed computing. In all these cases, exposing network information such as abstract topologies and network function deployment location helps applications.
- To support these emerging uses, extensions are needed, and additional functional and architectural features need to be considered as follows:

Work item 1: Generic Protocol Extension

- Protocol extensions to support a richer and extensible set of policy attributes in ALTO information update request and response. Such policy attributes may indicate information dependency (e.g., ALTO path-cost/[QoS](#) properties with dependency on real-time network indications), optimization criteria (e.g., lowest latency/[throughput](#) network performance objective), and constraints (e.g., relaxation bound of optimization criteria, domain or network node to be traversed, diversity of paths).

Work item 2: Pub Sub Mechanism

- Protocol extensions for facilitating operational automation tasks and improving transport efficiency. In particular, extensions to provide "pub/sub" mechanisms to allow the client to request and receive a [diverse types \(such as event-triggered/sporadic, continuous\)](#), customized feed of publisher-generated information. Efforts developed in other working groups such as MQTT Publish / Subscribe Architecture, WebSub, Subscription to YANG Notifications will be considered, and issues such as scalability (e.g., using unicast or broadcast/multicast, and periodicity of object updates) should be considered.

Charter Proposal-Continue

Work item 3: ALTO Data model

- The working group will investigate the configuration, management, and operation of ALTO systems and may develop suitable data models.

Work item 4: Multi-Domain Setting

- Extensions to ALTO services to support multi-domain settings. ALTO is currently specified for a single ALTO server in a single administrative domain, but a network may consist of multiple domains and the potential information sources may not be limited to a certain domain. The working group will investigate extending the ALTO framework to (1) specify multi-ALTO-server protocol flow and usage guidelines when an ALTO service involves network paths spanning multiple domains with multiple ALTO servers, and (2) extend or introduce ALTO services allowing east-west interfaces for multiple ALTO server integration and collaboration. The specifications and extensions should use existing services whenever possible. The specifications and extensions should consider realistic complexities including incremental deployment, dynamicity, and security issues such as access control, authorization (e.g., an ALTO server provides information for a network that the server has no authorization), and privacy protection in multi-domain settings.

- The working group will investigate recent protocol extensions (e.g., cost calendar, unified properties, and path vector) and new extensions that the WG develops. New considerations will include decisions about the set of information resources (e.g., what metrics to use), notification of changes either in proactive or reactive mode (e.g., pull the backend, or trigger just-in-time measurements), aggregation/processing of the collected information (e.g., compute information and network information) according to the clients' requests, and integration with new transport mechanisms (e.g., HTTP/2 and HTTP/3).

Charter Proposal-Continue

- When the WG considers standardizing information that the ALTO server could provide, the following criteria are important to ensure real feasibility:
 - - Can the ALTO server realistically provide (measure or derive) that information?
 - - Is it information that the ALTO client cannot find easily some other way?
 - - Is the distribution of the information allowed by the operator of the network? Does the exposure of the information introduce privacy and information leakage concerns?
- Issues related to the specific content exchanged in systems that make use of ALTO are excluded from the WG's scope, as is the issue of dealing with enforcing the legality of the content. The WG will also not propose standards on how congestion is signaled, remediated, or avoided.