

# ALTO Re-charter Discussions

IETF 108, Virtual Event

July 27, 2020

# Acknowledgement

- The technical slides are the outcome of recent meetings, discussions, and email exchanges with many contributors:
  - China Mobile: Franck Li
  - Telefonica: Luis Miguel Contreras Murillo
  - T-Mobile: Farni Boten, Lyle Bertz
  - Tencent: Yunfei Zhang Chunshan Xiong, Yixue Lei, Wei Huang, Yunbo Han
  - Benocs: Ingmar Poesse
  - The SENSE project: Harvey Newman (CalTech), Chin Guok, John McAulay, Tom Lehman (esnet), Justas Balcas (CERN)
  - The qzcloud project: Shu Yang (SZU), Zhongxing Ming (SZU), Xiaonan Xie (qcloud)
  - Nokia: Sabine Randriamasy, Anwar Walid
  - Ericsson: Borje Ohlman
  - SCU: Kai Gao
  - Tongji: Jensen Zhang
  - UNICAMP: Danny Perez, Christian Rothenberg
  - Univ. Minnesota: Zhi-Li Zhang
  - Yale: Qiao Xiang

# Area Director (Martin Duke) Guidance

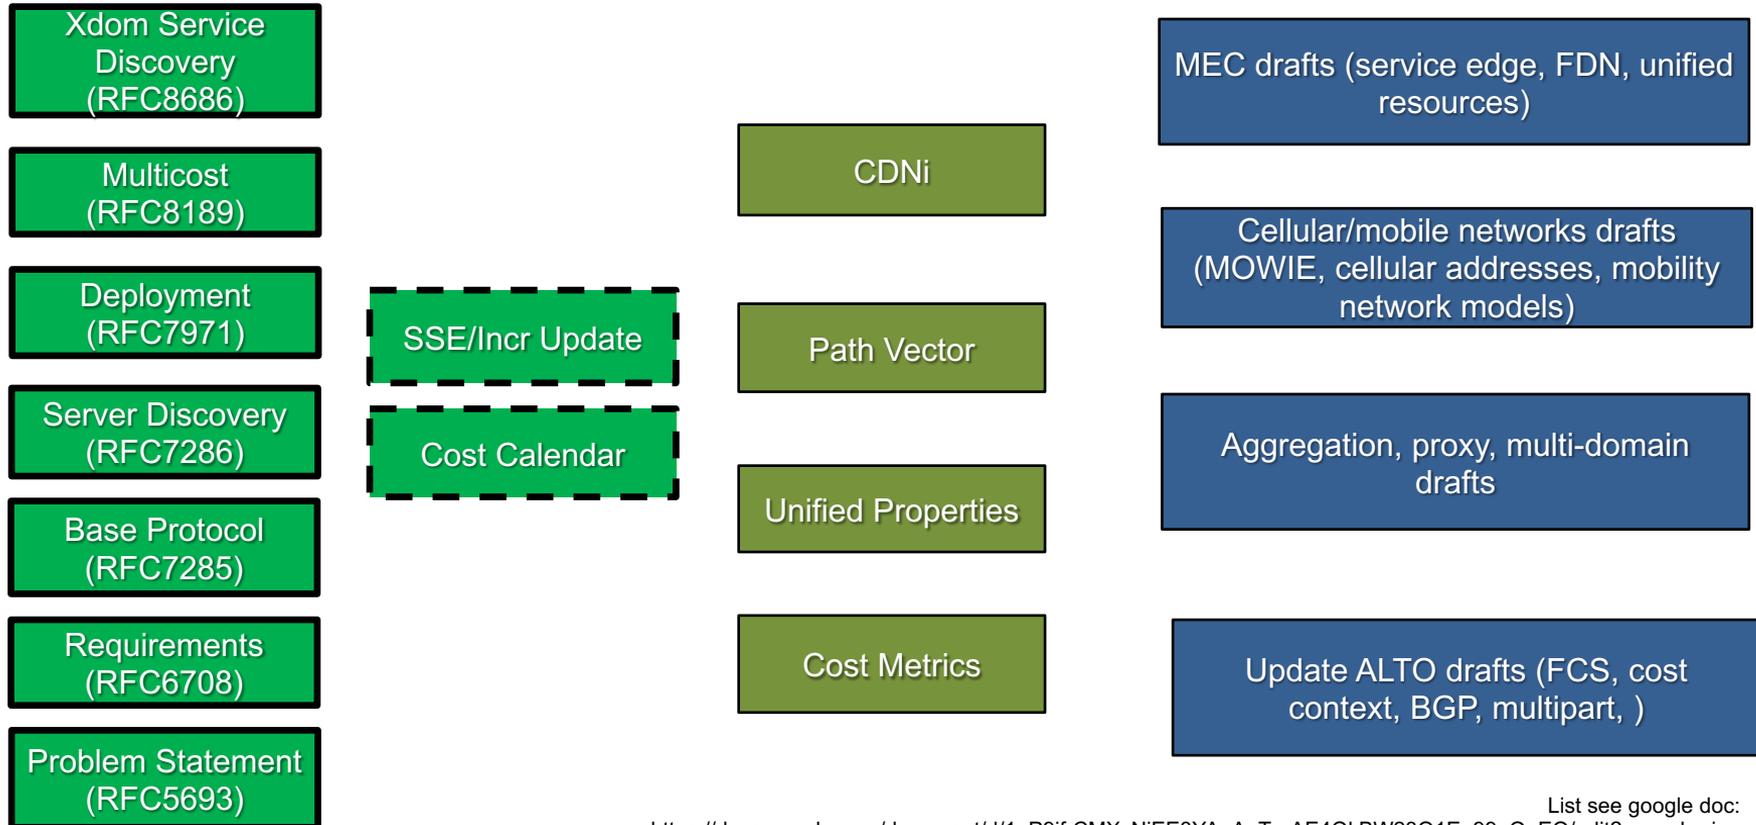
- “There are willing working group chairs and a charter that we can agree on
- **Relevance**
  - are we solving a real problem and are people willing to implement \*and deploy\* the output of the WG?
- **Review**
  - Does the WG have critical mass to provide quality review of documents?
- **Feasibility**
  - Is the problem statement and solution space tightly bounded so that we can set milestones for the next 1-2 years? Is the architecture well understood to be feasible, or do we have to solve several research problems to succeed?
- Looking for hidden problems from the transport perspective (and security, etc to a lesser extent)
- Goal: not have a finished charter at the end of 108
  - at best, we'll have **broad consensus on the kind of problem to solve next**, and can spend the time between now and 109 refining the text
  - at the very least, we'll have an idea of what the possibilities are”

# ALTO High-Level Goals and Bigger Context

- (Current) ALTO high-level goals:
  - Provide network information that an application may not easily get by itself
  - Provide abstraction to simplify complex network information/internals (goal: as simple as it can be, but not simpler)
- Bigger context: Network-application integration
  - Welcome to attend SIGCOMM'20 Network-Application Integration/CoDesign Workshop on August 14, 2020

# ALTO Now

# Current ALTO RFCs/WG Docs/Drafts



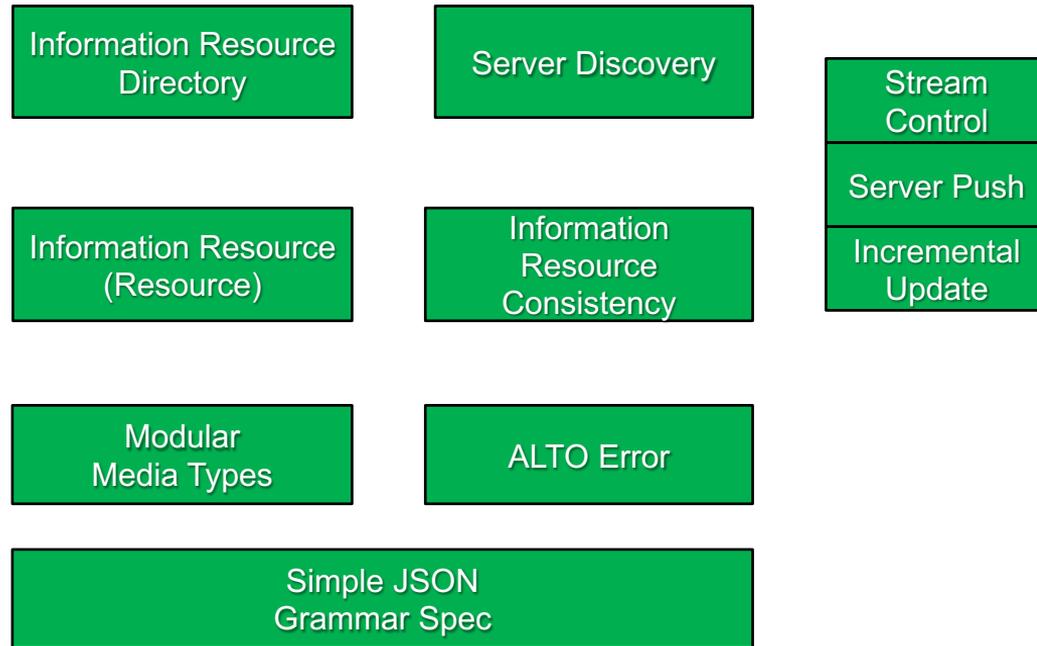
List see google doc:  
<https://docs.google.com/document/d/1qP9jf-CMXvNiEE3YAnApTczAE4QkBW23Q1Eg99uOaEQ/edit?usp=sharing>

<a href="#">draft-contreras-alto-bgp-communities-00</a>	Extending ALTO by using BGP Communities	7/13/204 pages New
<a href="#">draft-contreras-alto-service-edge-01</a>	Use of ALTO for Determining Service Edge	7/13/208 pages New
<a href="#">draft-huang-alto-mowie-for-network-aware-app-01</a>	MoWIE for Network Aware Application	7/13/2022 pages New
<a href="#">draft-lachos-alto-md-info-exposure-00</a>	Multi-domainn Information Exposure using ALTO	7/13/2019 pages New
<a href="#">draft-lachos-alto-multi-domain-use-cases-01</a>	Supporting Multi-domain Use Cases with ALTO	7/13/2026 pages New
<a href="#">draft-lachosrothenberg-alto-brokermdo-04</a>	ALTO-based Broker-assisted Multi-domain Orchestration	7/13/2022 pages New
<a href="#">draft-lachosrothenberg-alto-md-e2e-ns-02</a>	Multi-domain E2E Network Services	7/13/2015 pages New
<a href="#">draft-xiang-alto-multidomain-analytics-05</a>	Resource Orchestration for Multi-Domain, Exascale, Geo-Distributed Data Analytics	7/13/2023 pages New
<a href="#">draft-randriamasy-alto-cellular-adresses-03</a>	ALTO cellular addresses	3/9/208 pages
<a href="#">draft-randriamasy-alto-cost-context-03</a>	ALTO Contextual Cost Values	3/9/2018 pages
<a href="#">draft-xiang-alto-unified-representation-03</a>	ALTO Extension: Unified Resource Representation	7/13/2015 pages New
<a href="#">draft-yang-alto-deliver-functions-over-networks-01</a>	Delivering Functions over Networks: Traffic and Performance Optimization for Edge Computing using ALTO	7/13/2013 pages New
<a href="#">draft-zhang-alto-bgp-ls-01</a>	Considerations of Deploying ALTO using BGP - Link State (BGP-LS) Advertisement	7/13/2017 pages New
<a href="#">draft-gao-alto-fcs-07</a>	ALTO Extension: Flow-based Cost Query	3/16/2028 pages
<a href="#">draft-zhang-alto-multipart-04</a>	Multiple ALTO Resources Query Using Multipart Message	7/13/2026 pages New

# Current ALTO Protocol Framework

- Network information **divided** into (network) information resources
  - Explicit division allows modularity (**media types**), flexibility, scalability
  - **Dependency (consistency)** among information resources can be specified
- List of available information resources provided by **Information Resource Directory (IRD)**
- Bootstrap server provided by **server discovery** [7286, 8686]
  - Server discovery, xdom discovery
- Each individual information resource is provided as a **RESTful** service
  - Has a simple, but so far working well **grammar**
- A generic, SSE-based framework [SSE] to **stream-control, push, incrementally update** information resources

# Current ALTO Protocol Framework



# Current Provided Abstractions

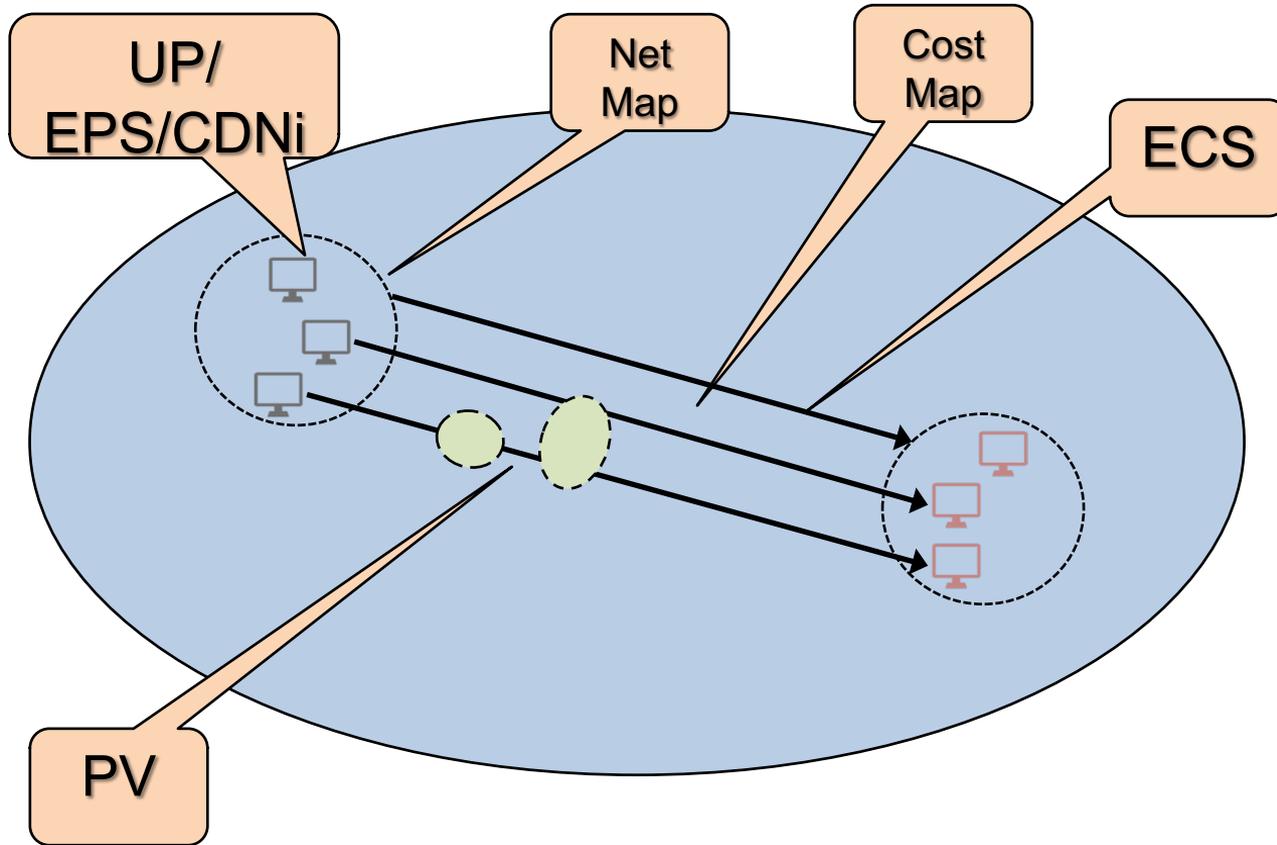
- Entity: A network consists of a set of entities [UP]
  - Endpoints, aggregation of endpoints [PID defined in 7285], abstract network elements [PV]
  - Each entity can have a set of properties
- Path: a network-traversal path from (some types) of a src entity to a dst entity
  - Path has properties called cost metrics [7285, PerfMetrics, MultiCost, CostCalendar]
  - ECS queries endpoint to endpoint, cost map queries aggregation to aggregation
- PathSet (co-flow): path vector
- Entity as destination [FCI]
- A generic framework supporting **entity properties, inheritance**
- Information resources can be **filtered**

# ALTO Network Abstractions

A network consists of nodes and paths

- A nodes can be an
  - endpoint
  - aggregations of endpoints (PID)
  - abstract network element
- Endpoints, aggregations of endpoints, abstract network elements are called **entities**
- Entities have **properties** that can be **inherited**
- Entities can have **capabilities**
- A path has **path properties**:
  - cost metrics, multcost, calendars
  - vector of abstract network elements
- A set of paths can form a co-flow, with:
  - shared abstract network elements cross the co-flows

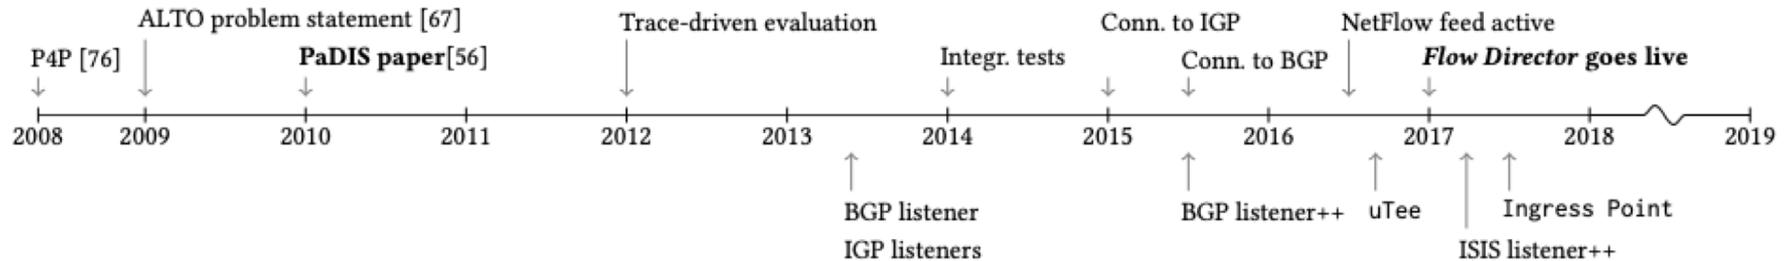
# ALTO Network Abstractions and Services



# Deployment Took Time

Steering Hyper-Giants' Traffic at Scale

CoNEXT '19 December 9–12, 2019 Orlando, FL, USA



**Figure 13: Timeline: From a research idea to a fully operational CDN-ISP collaboration. Top: Project management and infrastructure roll-out events. Bottom: FD's development milestones and main overhauls (++)**

source: <http://people.csail.mit.edu/gsmaragd/publications/CoNEXT2019/CoNEXT2019.pdf>  
CoNEXT 2019 Best Paper Award; IETF/IRTF 2020 Applied Networking Research Prize

# Real Deployment Need to Handle Complexity

Steering Hyper-Giants' Traffic at Scale

CoNEXT '19 December 9–12, 2019 Orlando, FL, USA

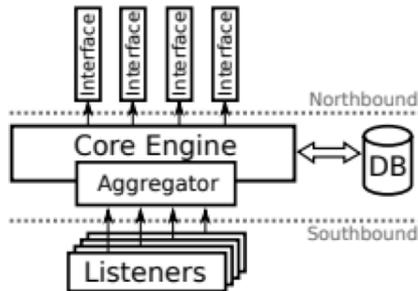


Figure 9: *Flow Director*: High-level system architecture.

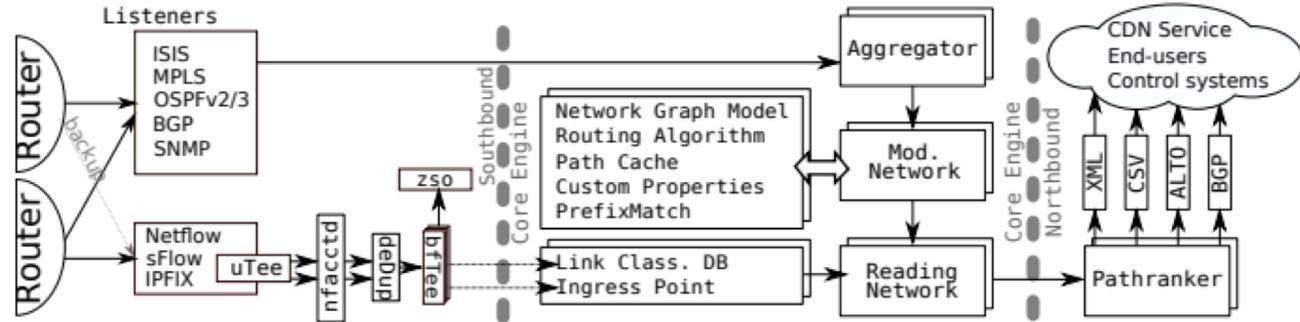


Figure 10: *Flow Director* processing pipeline.

source: <http://people.csail.mit.edu/gsmaragd/publications/CoNEXT2019/CoNEXT2019.pdf>  
CoNEXT 2019 Best Paper Award; IETF/IRTF 2020 Applied Networking Research Prize

# ALTO Re-charter: Relevance, Review, Feasibility

# Collaboration Essential for Relevance/Review/Feasibility

- Industry
  - Network operators, vendors, application providers
- Within IETF/IRTF, e.g.,
  - NMRG, PANRG
- Broader SDO, e.g.,
  - ETSI ZSM, 3GPP, SENSE/AutoGOLE/GNA-G PNWG
- Academia
  - NAI, SOSR, ...

# Essential Collaboration

- The technical slides are the outcome of recent meetings, discussions, and email exchanges with many contributors:
  - China Mobile: Franck Li
  - Telefonica: Luis Miguel Contreras Murillo
  - T-Mobile: Farni Boten, Lyle Bertz
  
  - Benocs: Ingmar Poesse
  - Tencent: Yunfei Zhang Chunshan Xiong, Yixue Lei, Wei Huang, Yunbo Han
  - The SENSE project: Harvey Newman (CalTech), Chin Guok, John McAulay, Tom Lehman (esnet), Justas Balcas (CERN)
  - The qzcloud project: Shu Yang (SZU), Zhongxing Ming (SZU), Xiaonan Xie (qcloud)
  
  - Nokia: Sabine Randriamasy, Anwar Walid
  - Ericsson: Borje Ohlman
  
  - SCU: Kai Gao
  - Tongji: Jensen Zhang
  - UNICAMP: Danny Perez, Christian Rothenberg
  - Univ. Minnesota: Zhi-Li Zhang
  - Yale: Qiao Xiang

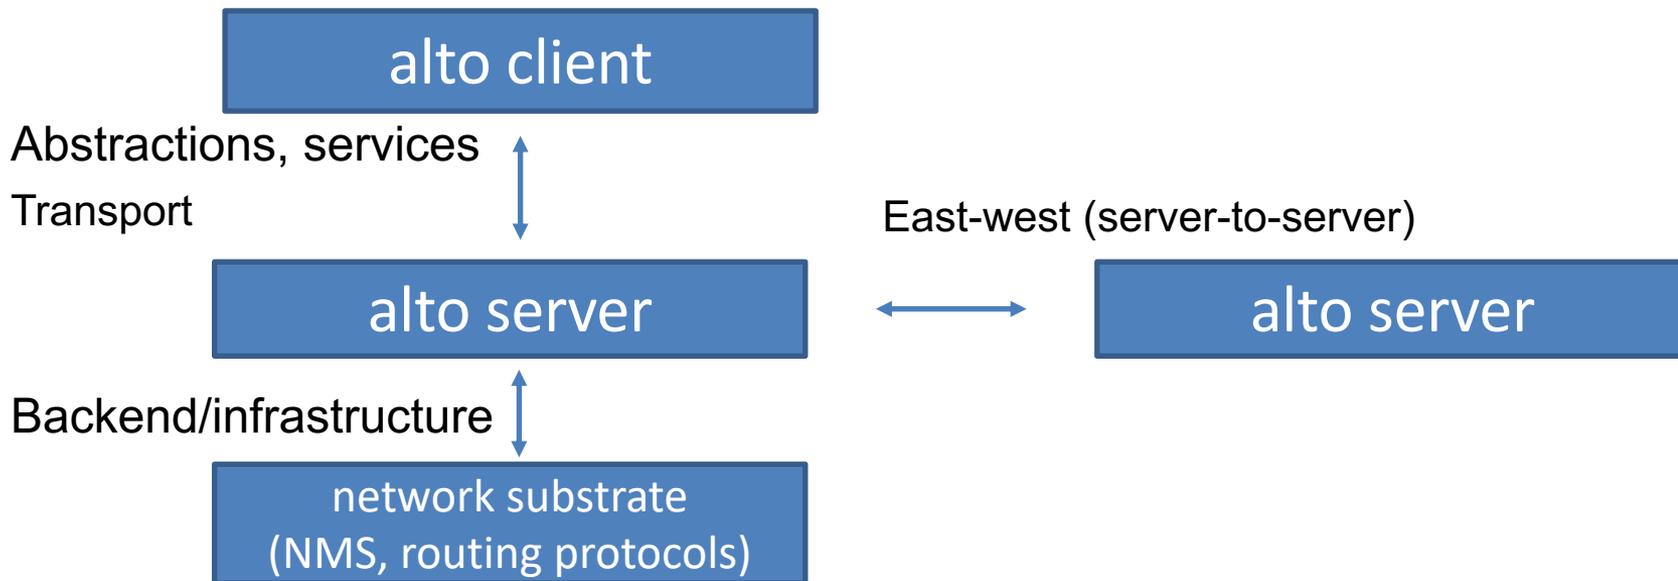
# ALTO Re-charter Technical Discussion

# WG Chairs/AD Discussion Guidance

- “everyone who wants to contribute a presentation that will have impact on the rechartering discussion should produce 1 slide (and 1 slide only, no title slide, just a single slide that includes architecture, results, and discussions on their specific work item for rechartering).
- the creator of the slide may come up to the mic and speak to their slide for 1 minute (and 1 minute only). This limit will be enforced by the chairs (apologies in advance).”

# Organization

- Use case driven
  - Cellular, MEC, huge data, automation, interdomain, ALTO extension
- Structure driven (structure is not rigid, just help)



# Organization

		Northbound	Transport	Backend (Southbound)	East-west (server-to-server)
Use case/Setting	MoWIE	Cellular info and benefist (Yunfei/Franck/Chunshan/Anwar)	in-band+out-band, fast (Chunshan); ftp/http/2/alto-see;in-band		
	MEC	Footprint for MEC + generic capabilities (Shu/Xiaonan/Zhongxing: qzcloud/SZU); Luis/Danny			
	HugeData	Generic query language to cover metrics such as time-bandwidth window (Harvey); predictive TCP tput (Jensen/Kai); FCS ?			
	Automation			Automatic measurements; automatic derivation from southbound (Luis/Jensen/Kai); Lyle; Farni (ZSM)	
	Interdomain	ALTO Interdomain (Danny/Qiao)			ALTO Interdomain (Danny/Qiao)
	ALTO extension	Generic query language (Qiao), Cost context (Sabine)	HTTP/2/3 (Richard); Multipart (Jensen)		

# List of 1-Pagers

- ALTO Services Extension for Cellular Networks with MOWIE
  - **Franck Li/China Mobile;**  
w/ Yunfei Zhang (Tencent), Chunshan Xiong, Yixue Lei, Wei Huang, Yunbo Han (Tencent), Anwar Walid (Nokia), Zhi-Li Zhang (Univ. Minnesota)
- ALTO Extension to Support ZSM
  - **Farni Boten/T-Mobile**
- ALTO Service Extension for MEC
  - **Luis Miguel Contreras Murillo/Telefonica,**  
w/ Danny Perez, Christian Rothenberg/UNICAMP
- ALTO Services Extension for Functional Delivery Networks
  - **Shu Yang/SZU**  
w/ Zhongxing Ming/SZU, Xiaonan Xie/qzlcoud
- ALTO Service Extension for Huge Data
  - **Harvey Newman/CalTech**  
w/ Chin Guok/esnet, Tom Lehman, John McAulay, Justas Balcas, Jensen Zhang

# List of 1-Pagers (Cont')

- ALTO Service Extension for Interdomain
  - **Danny Perez/UNICAMP**  
w/ Christian Rothenberg/UNICAMP, Luis Miguel Contreras Murillo/Telefonica, Borje Ohlman/Ericsson, Sabine Randriamasy/Nokia
- ALTO Service Extension for Predictive Network Performance
  - **Jensen Zhang/Tongji**  
w/ Kai Gao/SCU
- ALTO Flow-Based Information Services
  - **Kai Gao/SCU**  
w/ Jensen Zhang
- ALTO Services Extension Supporting Cost Context
  - **Sabine Randriamasy/Nokia**
- Generic Query Language Extension for ALTO
  - **Qiao Xiang/Yale**

# List of 1-Pagers (Cont')

- MOWIE Transport
  - **Chunshan Xiong (Tencent)**  
w/ Franck Li/China Mobile; w/ Yunfei Zhang (Tencent), Yixue Lei, Wei Huang, Yunbo Han (Tencent), Anwar Walid (Nokia), Zhi-Li Zhang (Univ. Minnesota)
- ALTO Transport Extension supporting HTTP/2
  - **Richard Yang/Yale**
- ALTO Transport Extension with Multipart
  - **Jensen Zhang/Tongji**
- Automatic Derivation of ALTO Information from Southbound
  - **Luis Miguel Contreras Murillo/Telefonica**  
Jensen Zhang, Kai Gao, A. Escibano/Alten, P. Cano/UST Global, F. Cano/Telefonica
- ALTO Automation Extension
  - **Lyle Bertz/T-Mobile**

# Backup Slides

# New Abstractions and Services

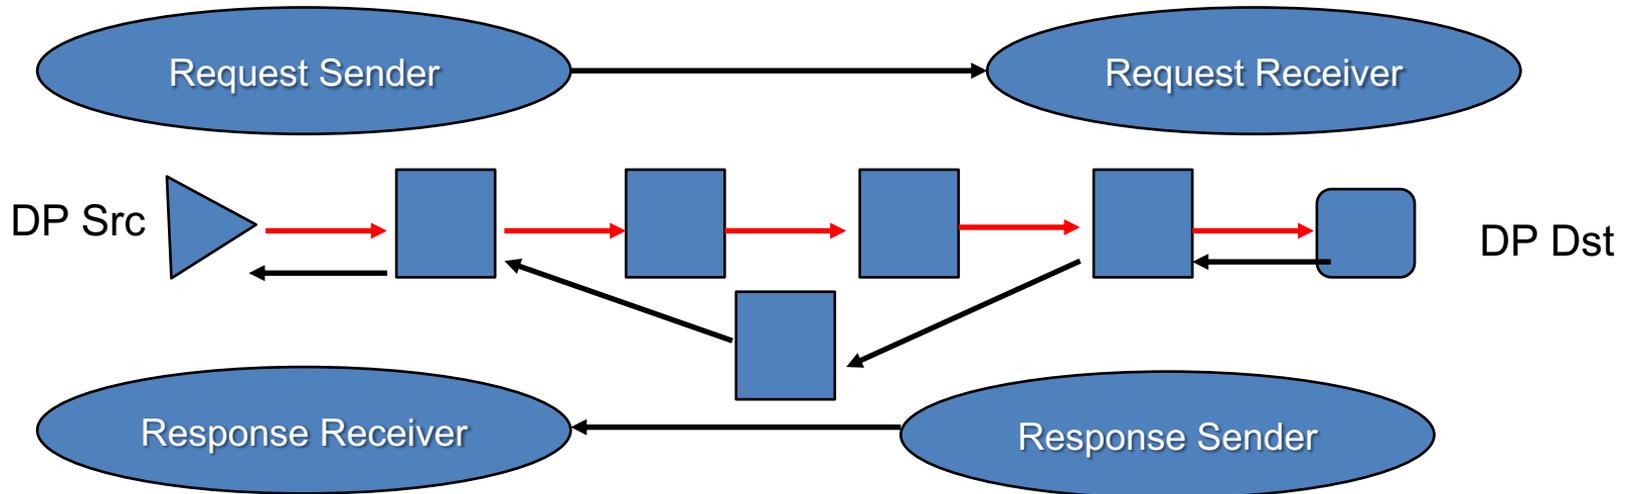
- Cost context
- Cellular info, cellular info
- Multiple resources, capabilities [CPU, ..., edge]
- SENSE type resource discovery
- Predictive network information (ToN/INFOCOM) reactive flow
- FCS
- Interdomain/multi-domain [guidance for a set of networks]
- SMT/query language

# New Transport

- In-band+out-band
  - Chunshan
- HTTP/2
  - Richard
- Multipart

# Network Information (NI) Exposure Transport: IB+OB

- Request Sender: who issues request for info
- Request Receiver: who receives request for info
- Response Sender(s): who send/update network info
- Response Receiver(s): who receive network info
- Network Info Source(s): source of network information



# Backend (Infrastructure)

## (Deployment, Operation, Implementation)

- How to acquire the information from lower layer to realize the services exposed to applications/query from clients
  - Automatic measurements
  - Automatic aggregation
  - Server-to-server
- Sources of raw information
  - Measurements, routing protocols, prefix
- How to acquire the information from lower layer to realize the services exposed to applications/query from clients
  - Automatic measurements
  - Automatic aggregation
  - Server-to-server





# Potential Charter

- Flexible resource model and query
  - Model: Resource query language [supports flow/co-flow level, cost context, SENSE, predictive, ...]
- Functions/capabilities
- Interdomain resource abstraction
- Transport: HTTP/2, Multipart, in+out band
- Backend/operations: Automatic measurements, operations, ...

# Application-Aware (AA) Networks Can Have Diverse AA Capabilities, Requiring Different Support

Example Capability	Possible Support & Assumption
<b><u>Treat each packet the same (aka not AA)</u></b>	
<b><u>Aware at app-level granularity</u></b>	
Create different networks/slices (e.g., voice vs data networks)	IP, access; SDN; scheduling
Identify packets by ports (e.g., ACL)	Packet header port; scheduling
<b><u>Aware of sub-app granularity</u></b>	
Scheduling each packet according to app-level deadline (e.g., fastpass'14)	Custom packet header; scheduling
Distinguish application-level structures (e.g., I frame vs P frame)	
Fancy: Co-flow scheduling (e.g., VARYS'14, AALO'15)	Network state; scheduling
<b><u>Aware of cross-app/protocol dependency</u></b>	
Fancy: identify full dependency (e.g., application-level dependency such as DNS->handshake->...)	Network state; scheduling
...	

# Network-Aware (NA) Applications

## Can Have Diverse NA Capabilities, Requiring Different Network Information/Support

Example Capability	Support & Assumption
<u>Transfer time selection</u>	Network state in time; can delay
<u>Server direction</u>	Path properties from client to potential servers; has multiple servers
<u>Rate adaptation</u>	
CC, reacting to loss/delay/ECN bit/INT (e.g., HPCC'19)	None/None/ECN/INT
Adaptive streaming	
Lower-than-best-effort (e.g., LEDBAT)	
Multi-path TCP	
...	

# Basic Challenge: Architecture

- Applications and networks can be designed with different objectives
  - Application: optimizes application's utility
  - Network: optimizes network's utility, enforces fairness, ...
- The end-to-end principle which mostly argues for the minimization of AA-networking

