ALTO Transport

Roland Schott
Y. Richard Yang
Danny Lachos
Chunshan Xiong

November 10, 2021

IETF 112, Virtual Meeting

High-Level Objectives of Our Work

- Analyze ALTO transport workload (services, service requirements)
- Evaluate performance and effectiveness of current transport (RFC7285 ALTO base, RFC8895 SSE)
- Evaluate design, benefits/no-benefits of integrating IETF new transport (HTTP/2, HTTP/3)
- Evaluate ALTO transport compared with other network->application information transport (e.g., in-band, 3GPP)

Service	Type (Out if not labeled)	Core Information Structures	Main Size Var	Stability Expectation	Incremental Changes
Information Resource Directory [RFC7285]	application/alto-directory+json	Key-value store; Delegation	#resources	Stable	Add/delete resources
Network Map [RFC7285]	application/alto-networkmap+json	Key-value store: pid -> addrType -> array	#CIDRs	Stable	Add/delete CIDR from pid
Cost Map [RFC7285]	application/alto-costmap+json	Key-value store: srcPID -> dstPID -> value Depend on network map	#SRCPID * #DSTPID	More dynamic than network map	Update cost map entries
Filtered Map Services [RFC7285]	In: application/alto-networkmapfilter+json Out: application/alto-costmapfilter+json	In: selected srcPID, dstPID Key-value store: srcPID -> dstPID -> value Depend on network map	Filtered #SRCPID * #DSTPID	More dynamic than network map	Update cost map entries
Endpoint Property Service [RFC7285]	In: application/alto-endpointpropparams+json Out: application/alto-endpointprop+json	In: addr + prop Key-value store: addr -> prop -> value	#addr * #prop	Depend on property, can be dynamic or stable	Update property value
Endpoint Cost Service [RFC7285]	In: application/alto-endpointcostparams+json Out: application/alto-endpointcost+json	In: srcAddr x dstAddr Key-value store: srcAddr -> dstAddr -> value	#src * #dst	Can be more dynamic than cost map	Update cost value
Cost Calendar [RFC8896]	Extension to CostType	Calendar array	Previous * #num_intervals	Can be dynamic	Calendar window moves
Unified Property	In: application/alto-propmapparams+json; Out: application/alto-propmap+json	In: addr, prop Key-value store: addr -> prop -> value	#addr * #prop	Depend on property, can be dynamic or stable	Update property value
Path Vector	In: see costmap; Out: multipart/related;type=application/alto-costmap+json,	Cost map + unified property map	#src * #dst * #vec size	Depend on metric, can be dynamic or stable	Update path vector
CDNi Cap & Footprint	application/alto-cdni+json	Array; {capability-type: capability-value}	#footprint * #capability	Can be dynamic, bursty	Update capabilities

Performance and Effectiveness of Current Transport

Infrastructures	Basic Workload (ALTO SPEC)	Transport	Collecting metrics

• Benocs is fully open to use its infrastructure as an evaluation environment
 	
• Greater Bay Network is also fully open to use its infrastructure as an evaluation environment
	1. (Filtered) Cost map: distribute inter-site performance metrics and calendar; routing changes + link dynamics => updated metrics
	2. Endpoint/unified property service: endpoint access status query/updates/bwe
	3. CDN node footprint & capability
	4. Flow direction (pointing to CDN nodes) using ECS
	5. Path vector providing available reservable bandwidth
		• HTTP/1.x per rquest full retrieval
		  • keep alive
		  • pipelining
		
		• HTTP/2, HTTP/3 per request, full retrieval
		
		• ALTO/SSE RFC8895 (on HTTP/1.x)
			• ALTO Server processing load
			• ALTO client processing load
			
			• Transport load (bytes)
			• Transport latency
			• Throughput
			
			• Scalability

ALTO Transport using HTTP/2 [RFC7540]

- Goal: ALTO/SSE HTTP/1.x => ALTO/H2 (incremental + push)
- Discussed during ALTO/SSE IESG review; RFC8895 Sec. 3.3

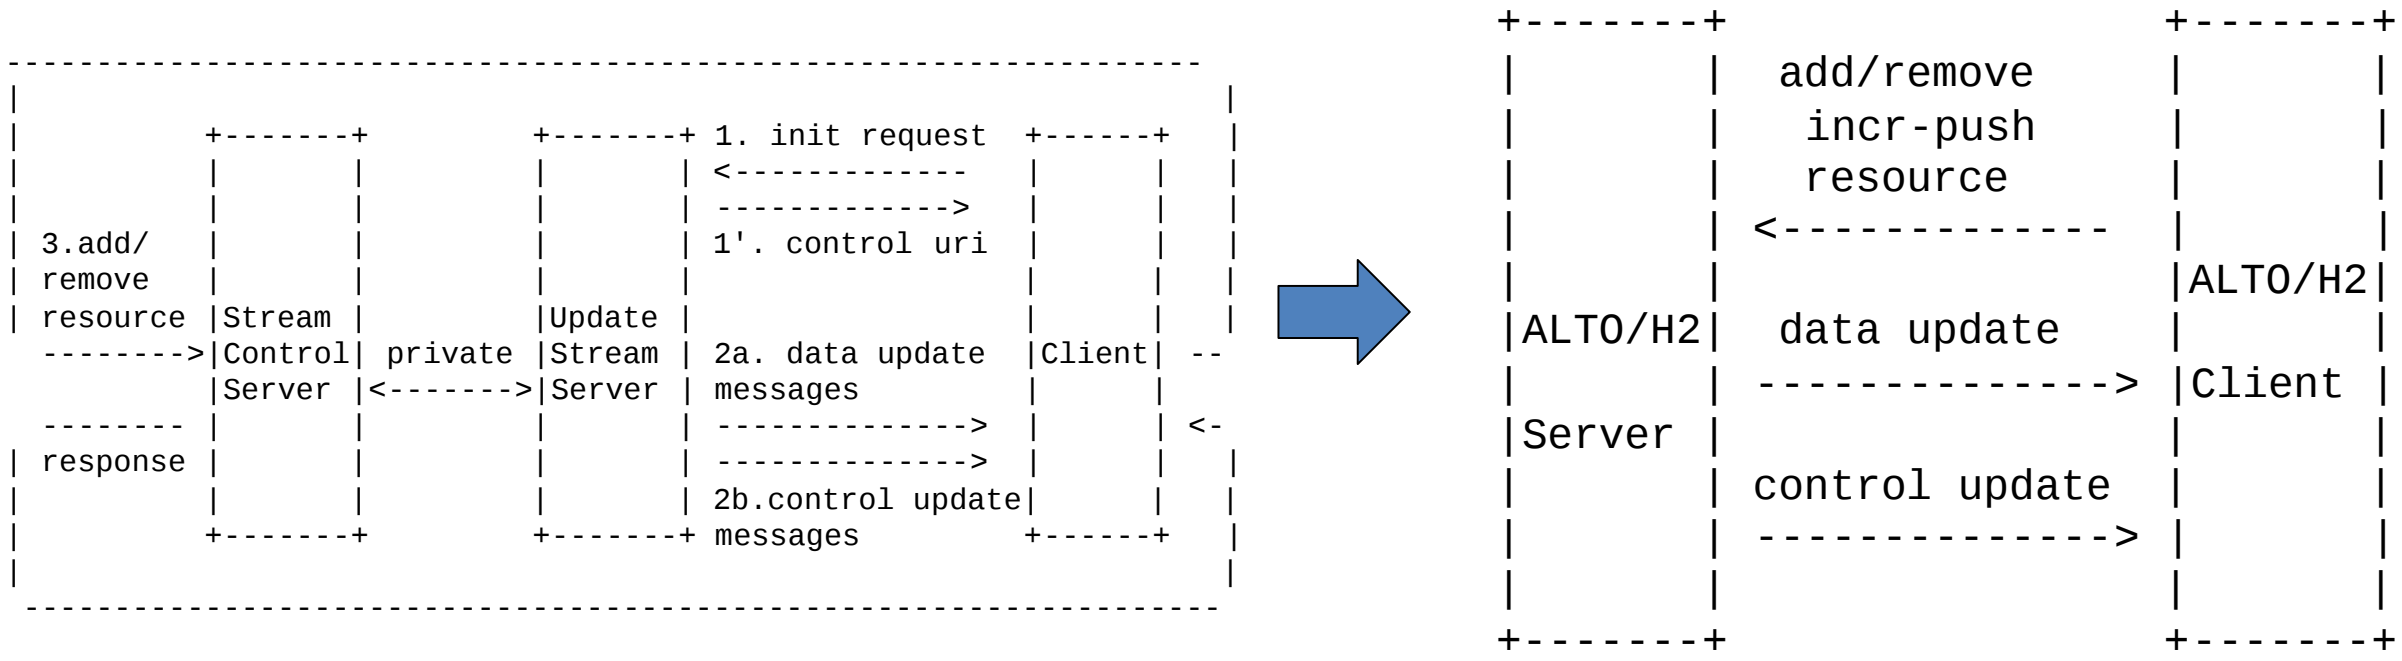

ALTO/H2 Requirements

- R0: Client can request any ALTO resource using the connection, just as using ALTO base protocol using HTTP/1.x
- R1: Client can request the addition (start) of incremental updates to a resource
- R2: Client can request the deletion (stop) of incremental updates to a resource
- R3: Server can signal to the client the start or stop of incremental updates to a resource
- R4: Server can choose the type of each incremental update encoding, as long as the type is indicated to be acceptable by the client
- R5: Design follows basic principle of HTTP---Representational State Transfer
- R6: Design takes advantage of HTTP/2 design features such as parallel transfer and respects HTTP/2 semantics

ALTO/H2 Design Issue

- ALTO incremental updates support multiple content types for incremental updates
  - Original resource type (for complete information)
  - JSON Patch
  - JSON Merge Patch
- => Server push needs to indicate, for each incremental update,
  - the content type of the incremental update
  - which initial request
- ALTO/SSE solution:
  - <content-type>+<stream-id> in each update event, e.g.,
    - event: application/alto-networkmap+json,my-network-map
    - data: …

ALTO/H2 Design Issue Space

- ALTO/H2 design options
  - Incremental-update stream <-> HTTP/2 stream
    - Need content indication layer,
      - Design 1: Exactly ALTO/SSE
      - Design 2: Content-type: <type>
  - Incremental-update stream -> multiple HTTP/2 streams
    - A design point is each update is a new HTTP/stream [RFC8895 Sec. 3.3]
      - Realize using PUSH_PROMISE, but semantics may not match
  - HTTP/2 extension, e.g., add a new Frame type, to indicate structure in stream, support generic incremental push notifications

General Space of Network->App Information Transport

Protocol/Base Reference	What information	How transported	Frequency
ALTO/RFC7285, RFC8895	Network map, cost map, unified property, …	HTTP/1.x client/server	Request/response; ALTO/SSE incremental push
ECN/RFC3168	Congestion notification	2 bits in IP Traffic Class header; ECN-Echo/CWR flags in TCP header	Per packet marking
NEF, SCEF<->AS/3GPP, …	Network capabilities/events -> AF	HTTP	Request/response
…			

Next Step Planning IETF 113

- Benchmarking deployment and initial reports [Richard, Danny]
  - openalto.org
  - Open platforms, all are welcome to join
- Initial design and analysis of ALTO/H2 [Roland, Richard]
  - Design team
- Initial draft surveying of Net->App information and transport in major SDOs [Roland, Chunshan, Richard]

Backup Slides

ALTO Transport using ALTO/SSE [RFC8895]

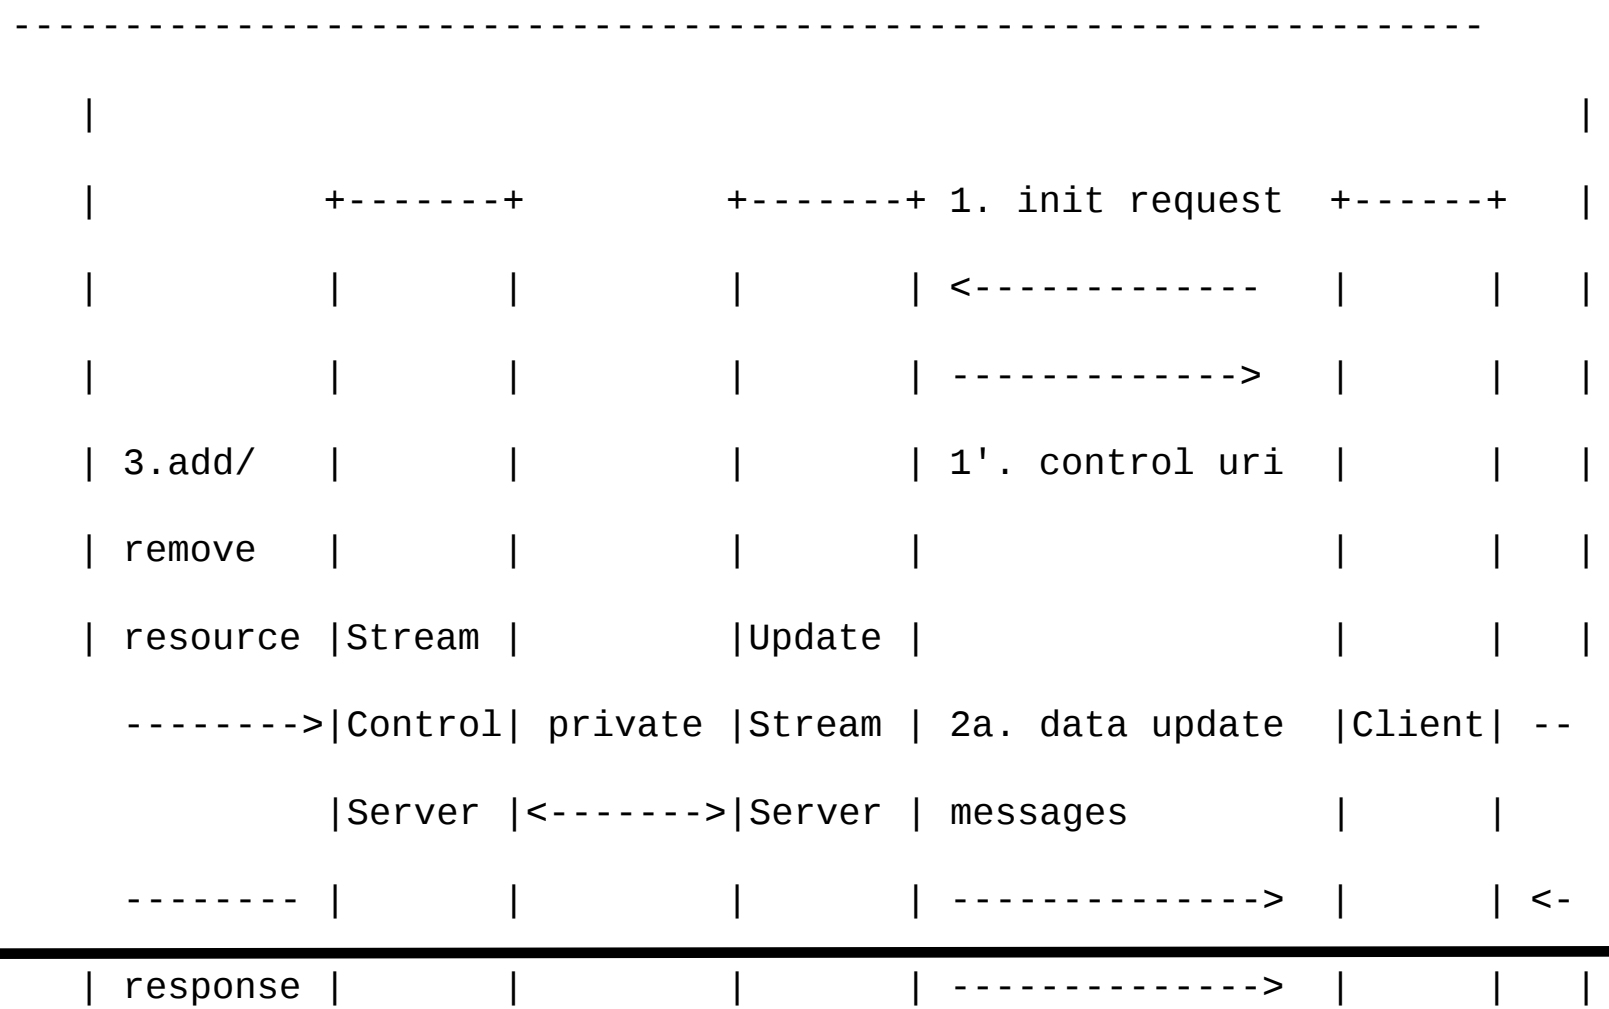

ALTO/H2

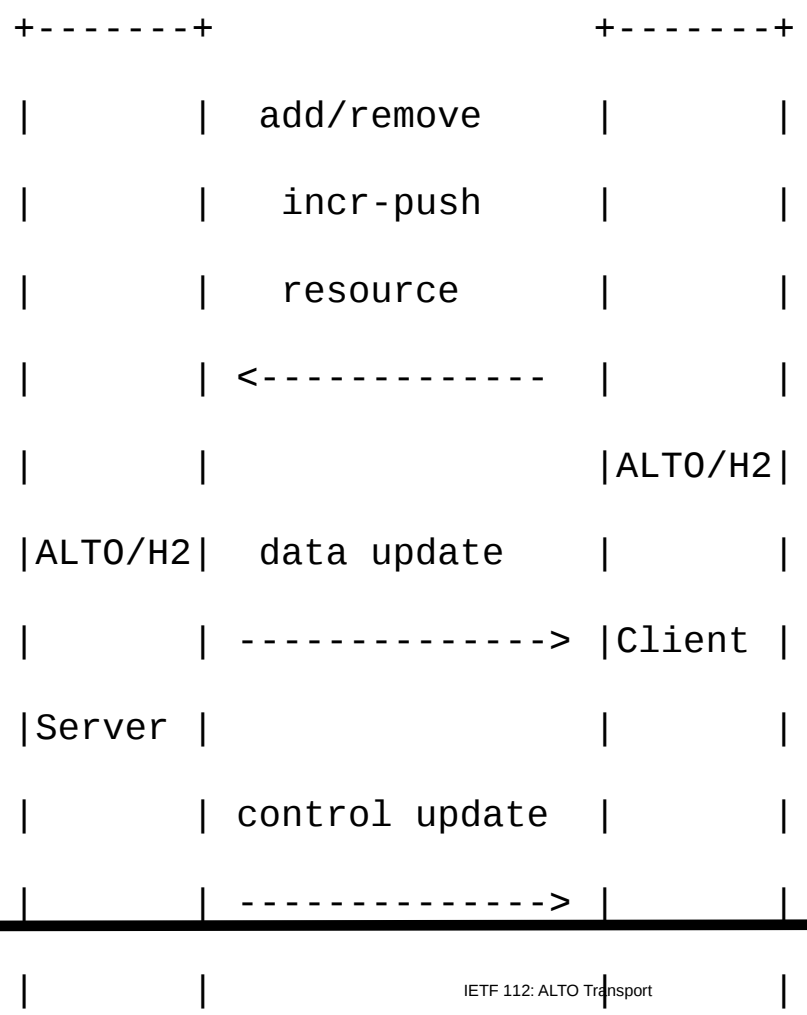
add/remove
incr-push
resource
data update
control update
ALTO/H2
ALTO/H2
Server
Client