

WebSocket-based server-to-client notifications for the Application-Layer Traffic Optimization (ALTO) Protocol

draft-marocco-alto-ws-01

Enrico Marocco, Jan Seedorf

IETF 84, Vancouver

ALTO WG

July 31st, 2012

Background

Discussions on enhancements/extensions to the current ALTO protocol

- I2AEX BoF and ALTO WG session at IETF83
 - Check presentations and see also `draft-marocco-alto-next`
- Possible useful extensions
 - Incremental updates (see `draft-schwan-alto-incr-updates`)
 - Server-initiated Notifications ← **focus of this draft/presentation**
 - ALTO Information Extensions
- Use Cases that would benefit from such ALTO enhancements:
 - ALTO for CDNs and CDNI
 - ALTO for SDN
 - ALTO Server-to-server Communications
 - Virtualized Applications in Datacenters

Goals

Goals of draft-marocco-alto-ws-01

- Propose WebSockets (RFC 6455) as a possible solution for server-initiated ALTO information update notifications
- Highlight the benefits of this approach
- Provide high level example
- Start discussion of design tradeoffs in the ALTO WG

Why WebSockets?

- WebSockets are only one option for server-initiated ALTO update notifications
 - Other options include e.g. XMPP, BGP and SNMP
- Advantages of WebSockets
 - Explicitly intended to provide bi-directionality to HTTP, i.e. the transport the ALTO protocol is based on
 - Simple identification of related ALTO resources in a WebSocket URI
 - Easy advertisement of such WebSocket URIs in the Information Resource Directory (IRD)
 - Unambiguous definition of resources through MIME type
 - Reuse of HTTP authentication

Example

```
{
  "resources" : [
    .
    .
    .
    {
      "uri" : "http://alto.example.com/networkmap",
      "media-types" : [ "application/alto-networkmap+json" ],
      "updates" : "ws://alto.example.com/networkmap"
    }, {
      "uri" : "http://alto.example.com/costmap/num/routingcost",
      "media-types" : [ "application/alto-costmap+json" ],
      "capabilities" : {
        "cost-modes" : [ "numerical" ],
        "cost-types" : [ "routingcost" ]
      },
      "updates" : "ws://alto.example.com/costmap/num/routingcost"
    }
  ]
}
```

Information Resource Directory returned by an ALTO Server

- IRD contains the WebSocket URI of the update notification service associated to the specific resource
- ALTO Server provides both a network map and a cost map with corresponding update notification services

Conclusion and Discussion

- Design tradeoffs for server-initiated ALTO notifications
 - WebSockets implies keeping a connection alive to the server for updates
 - useful in cases where there are frequent notifications
 - However, there may be overhead in cases where an ALTO server only very rarely updates ALTO information
 - What is the main use case (with respect to frequency of updates) we are targeting with update notifications?
 - Possibility of having only one WS connection for all the update notifications of the resources the client is interested in
 - Pro: one connection only
 - Con: additional complexity due to the need of some kind of subscription protocol

BACKUP SLIDES

Need for Server-initiated ALTO notifications: Example of ALTO-Guidance within CDN Request Routing (DNS Example)

