Real-Time Internet Peering for Telephony (RIPT)

Problem Statement and Protocol Summary

Jonathan Rosenberg

CTO Five9

26 March 2020

Deploying SIP/RTP apps in public cloud is much harder than web apps with a growing gap

Hard/Impossible

elastic peering (SIP trunking IP based),
Hitless upgrade on peering (requires call drops)
Call survivability (dropped call on server failure)

Proxy farms, SBCs, custom DNS or IP, Firewall configs, custom HA, custom load balancing, custom monitoring and tracing… your biz logic

VM

Biz logic

Service Mesh

Kubernetes

Cloud Functions

Global Elastic LB + TLS termination + geoDNS

VM

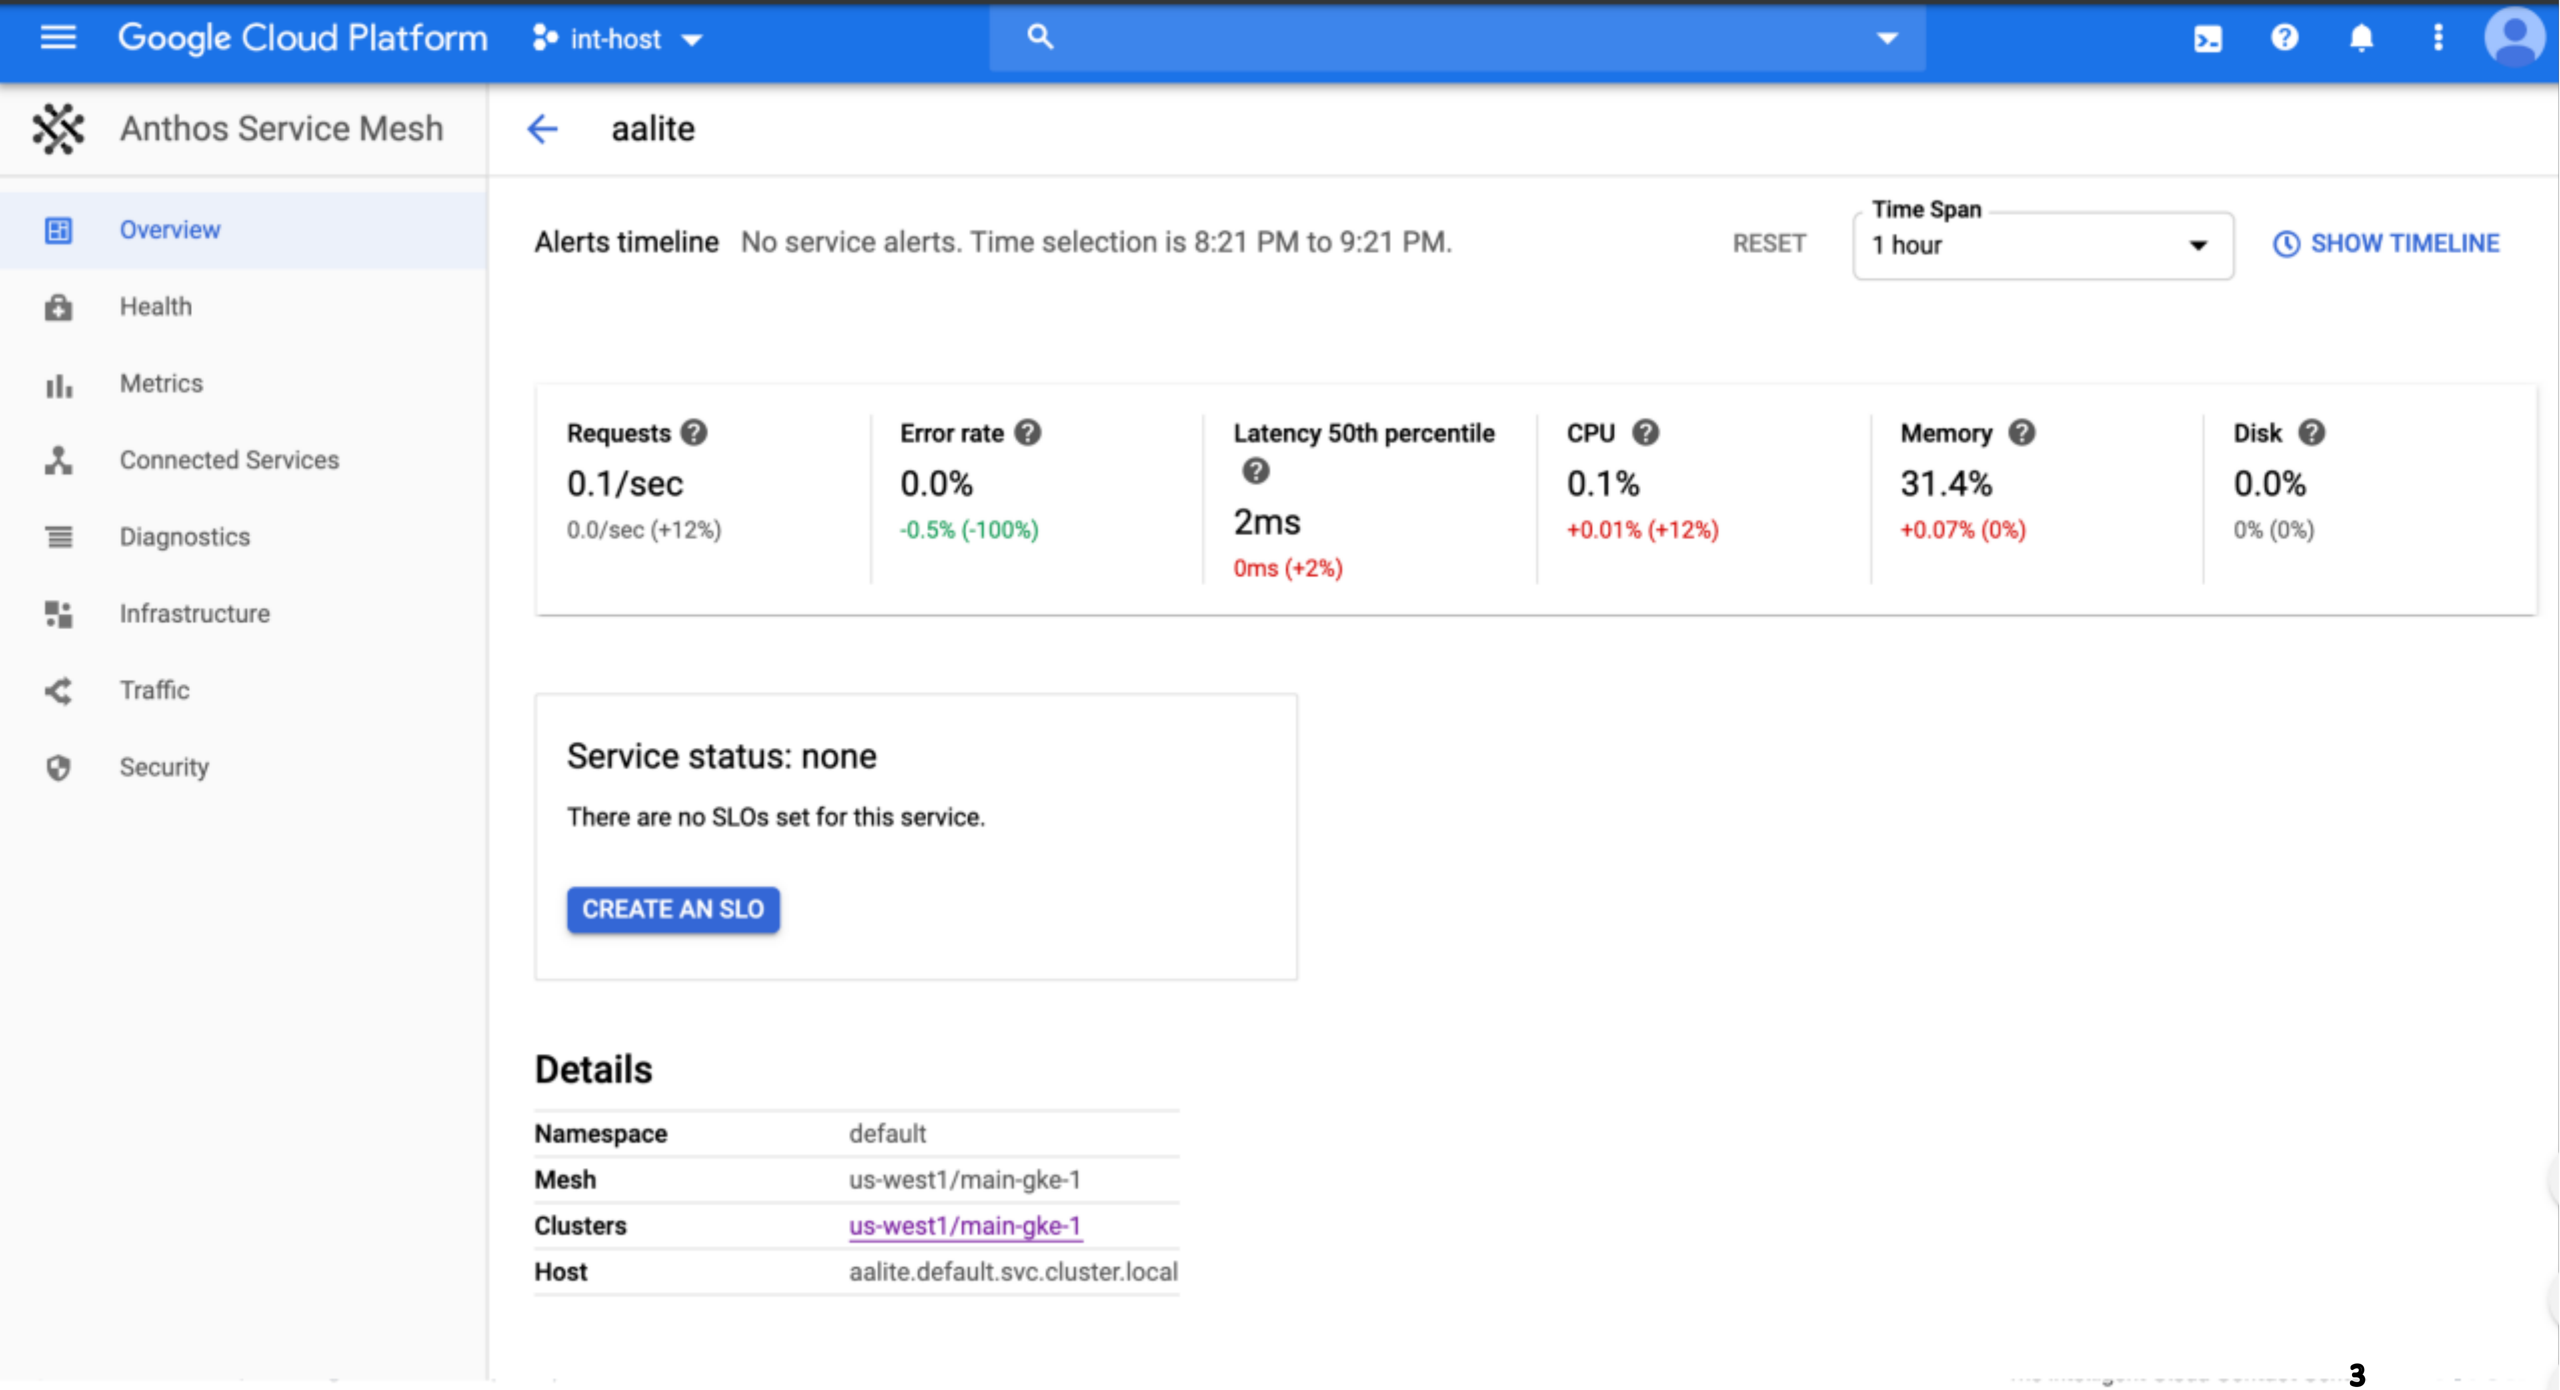
Google Cloud Platform
int-host
Anthos Service Mesh
aalite
Overview
Health
Metrics
Connected Services
Diagnostics
Infrastructure
Traffic
Security
Alerts timeline
No service alerts. Time selection is 8:21 PM to 9:21 PM.
RESET
Time Span
1 hour
SHOW TIMELINE
Requests
0.1/sec
0.0/sec (+12%)
Error rate
0.0%
-0.5% (-100%)
Latency 50th percentile
2ms
0ms (+2%)
CPU
0.1%
+0.01% (+12%)
Memory
31.4%
+0.07% (0%)
Disk
0.0%
0% (0%)
Service status: none
There are no SLOs set for this service.
CREATE AN SLO
Details
Namespace
default
Mesh
us-west1/main-gke-1
Clusters
us-west1/main-gke-1
Host
aalite.default.svc.cluster.local

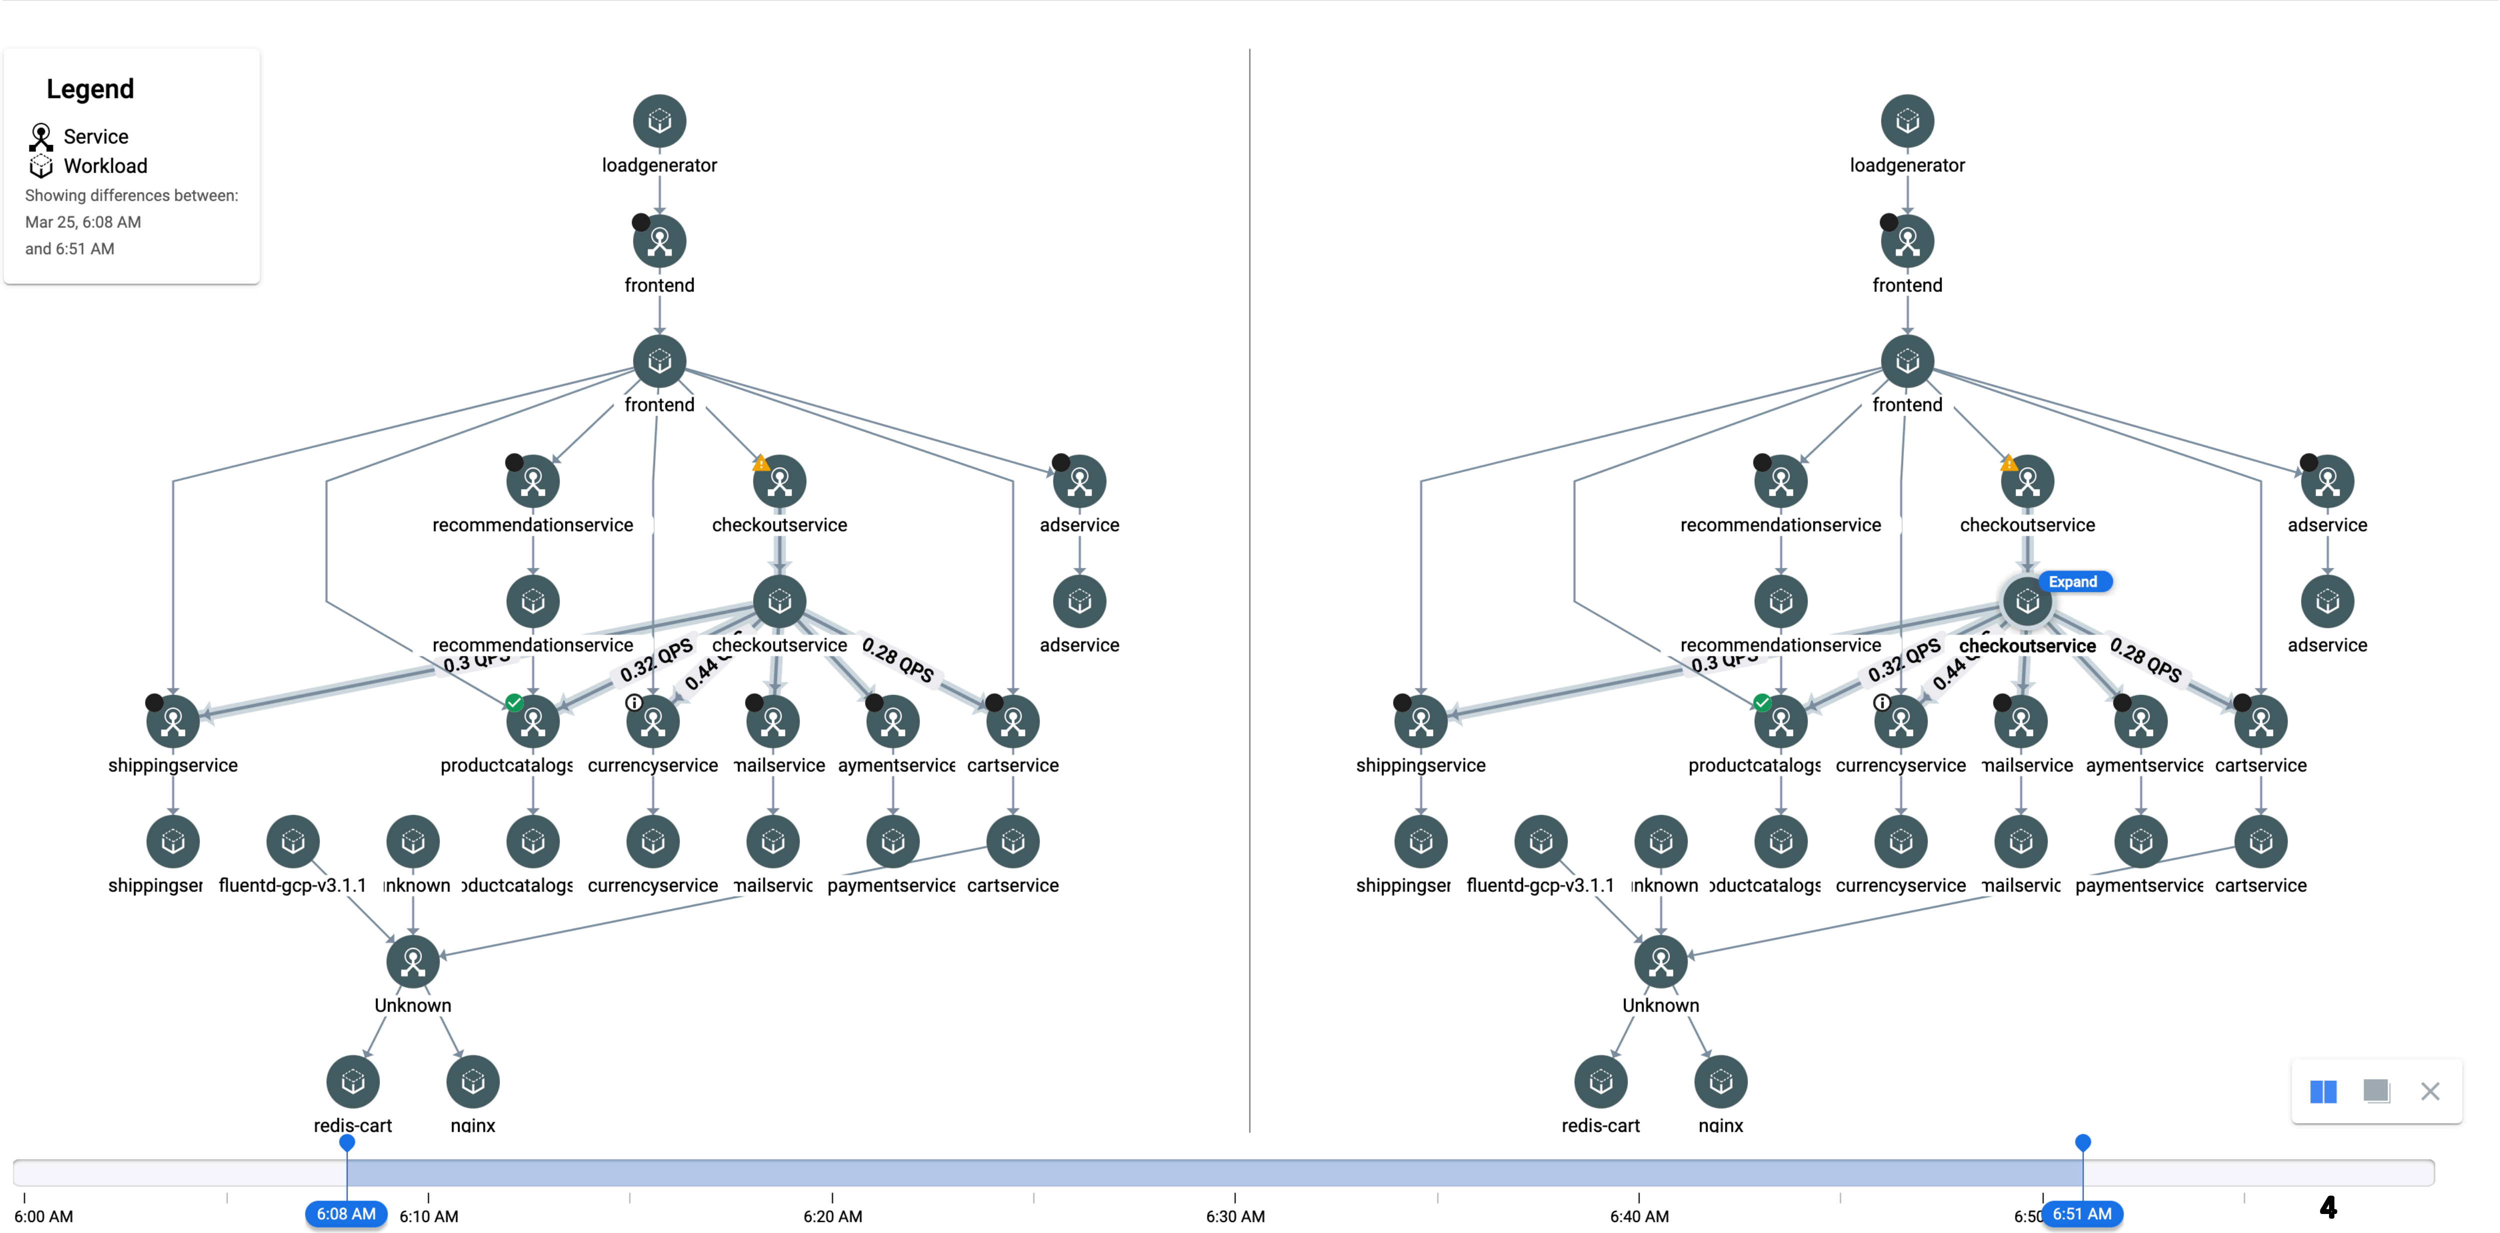

Legend
Service
Workload
Showing differences between:
Mar 25, 6:08 AM
and 6:51 AM
loadgenerator
frontend
frontend
recommendationservice
checkoutservice
adservice
recommendationservice
checkoutservice
adservice
0.3 QPS
0.32 QPS
0.28 QPS
shippingservice
productcatalogs
currencyservice
nailservice
aymentservice
cartservice
shippingser
fluentd-gcp-v3.1.1
ınknown
›ductcatalogs
currencyservice
nailservic
paymentservice
cartservice
Unknown
redis-cart
nginx
Expand
6:00 AM
6:08 AM
6:10 AM
6:20 AM
6:30 AM
6:40 AM
6:50
6:51 AM

What is it about SIP/RTP?

- Its not using HTTP
- IP addresses all over the place breaks load balancers and HA techniques
- RTP in the cloud (bridges, transcoders, etc)
  - Ports for session ID incompatible with kubernetes
  - Usage of IPs for routing incompatible with load balancers, modern HA techniques
  - SIP trunking requires VIP for any form of HA – works very poorly in public cloud

Problem Statement

Enable the deployment of VoIP apps ontop of public cloud platforms such that they can take advantage of the capabilities provided by those platforms for web applications (such as load balancing, HA, hitless upgrades, service mesh, geo-redundancy, etc), minimizing (if not eliminating) the need for customized behaviors unique to VoIP.

Resulting Protocol Requirements

- “just an app” ontop of HTTP (and ideally http3)
- No IP addresses or ports in the protocol
- Server-centric state
- Client-Server
- Signaling and Media Together
- Separation of Call and Connection State
- OAuth

Many Use Cases

Same Provider

	Browser to Cloud App	IP Softphone to IP PBX	App to App (ASR to cloud)	Real-time devices	SIP Trunking alternative
Media Transport (/media)	X	X	X	X	X
Control (start, stop) (/events)		X	X	X	X
Media Negotiation (/handlers)	(API not protocol)	X	X	X	X
Identity Management (/providertgs, /consumertgs)		X			X

Exchanging media: /media

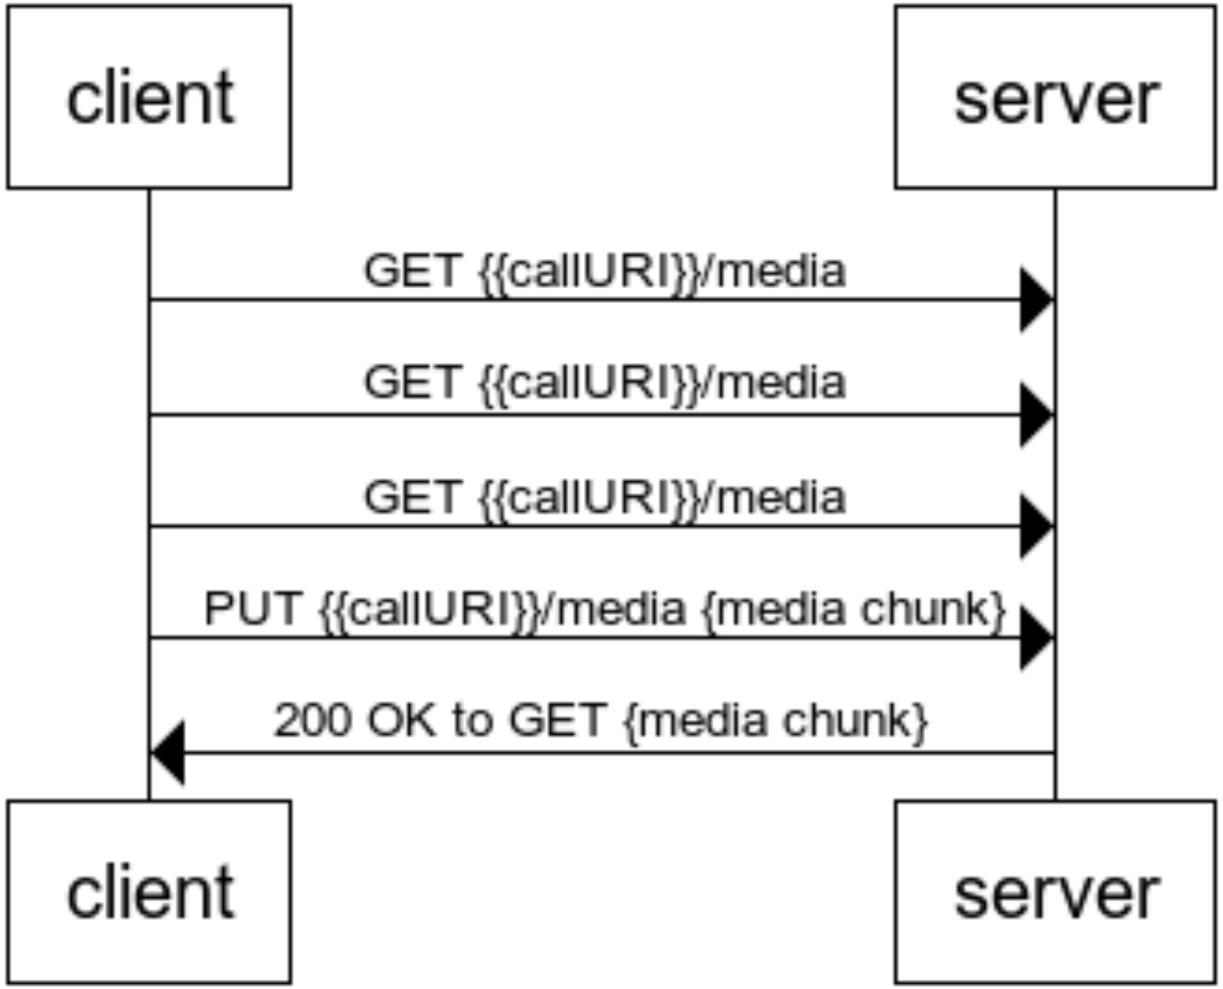

Media Chunk =
payload of RTP packets

Media is always sent to the call URI, which you get….

By control /events – when a call is established

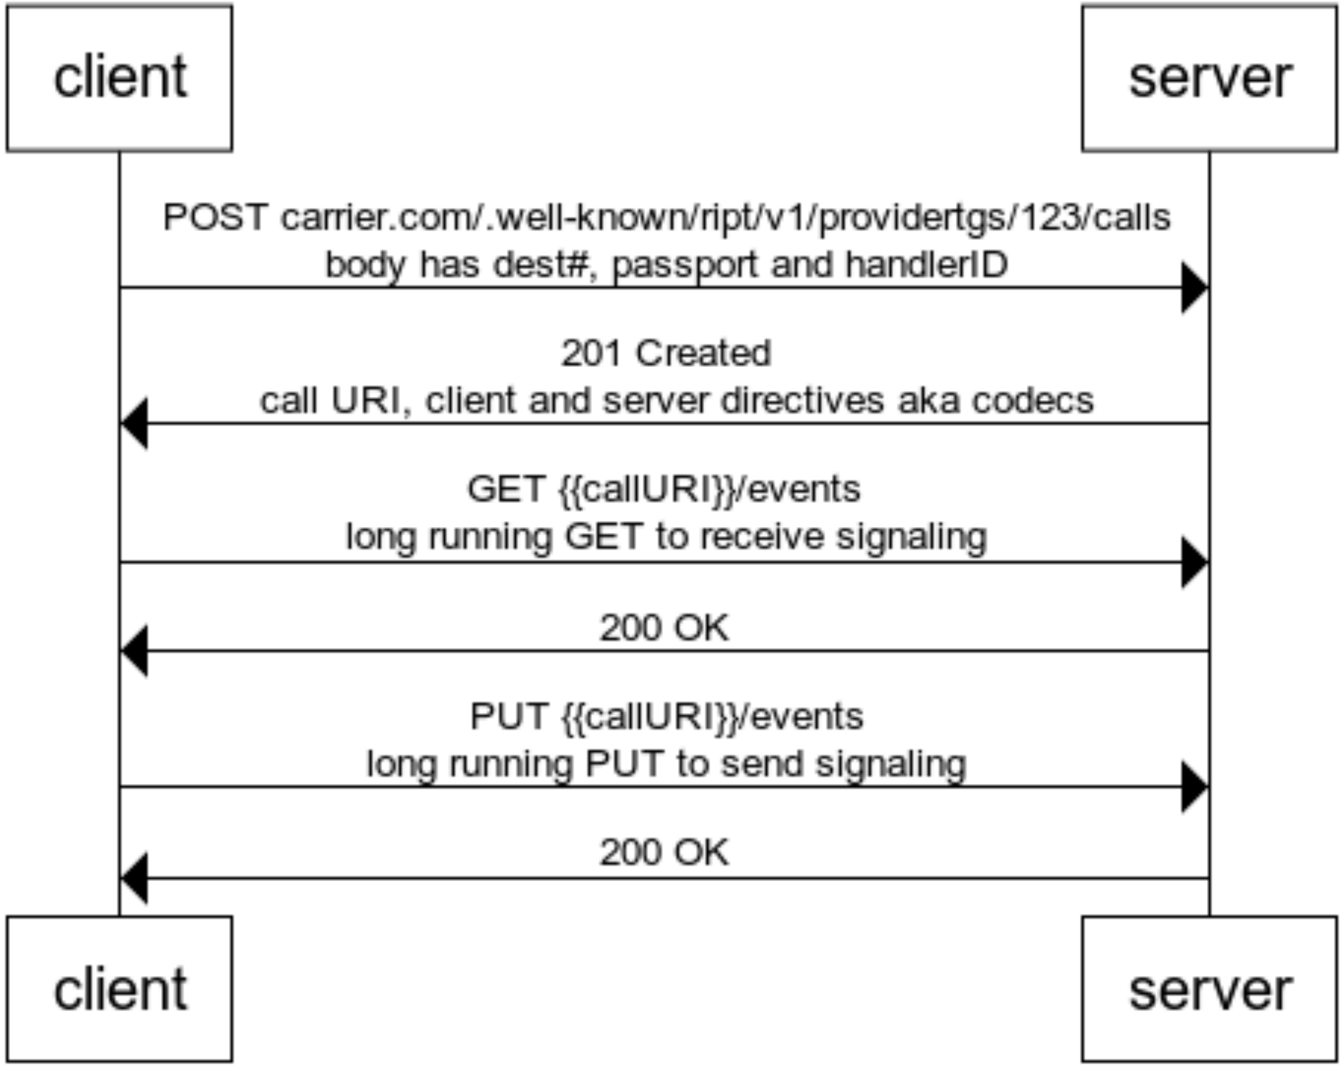

Events:
Proceeding
Alerting
Answered
Declined
Ended
Migrate
Moved
Ping
pong

Handler describes media capabilities, which was shared…

By posting to /handlers
one time event when new client comes online

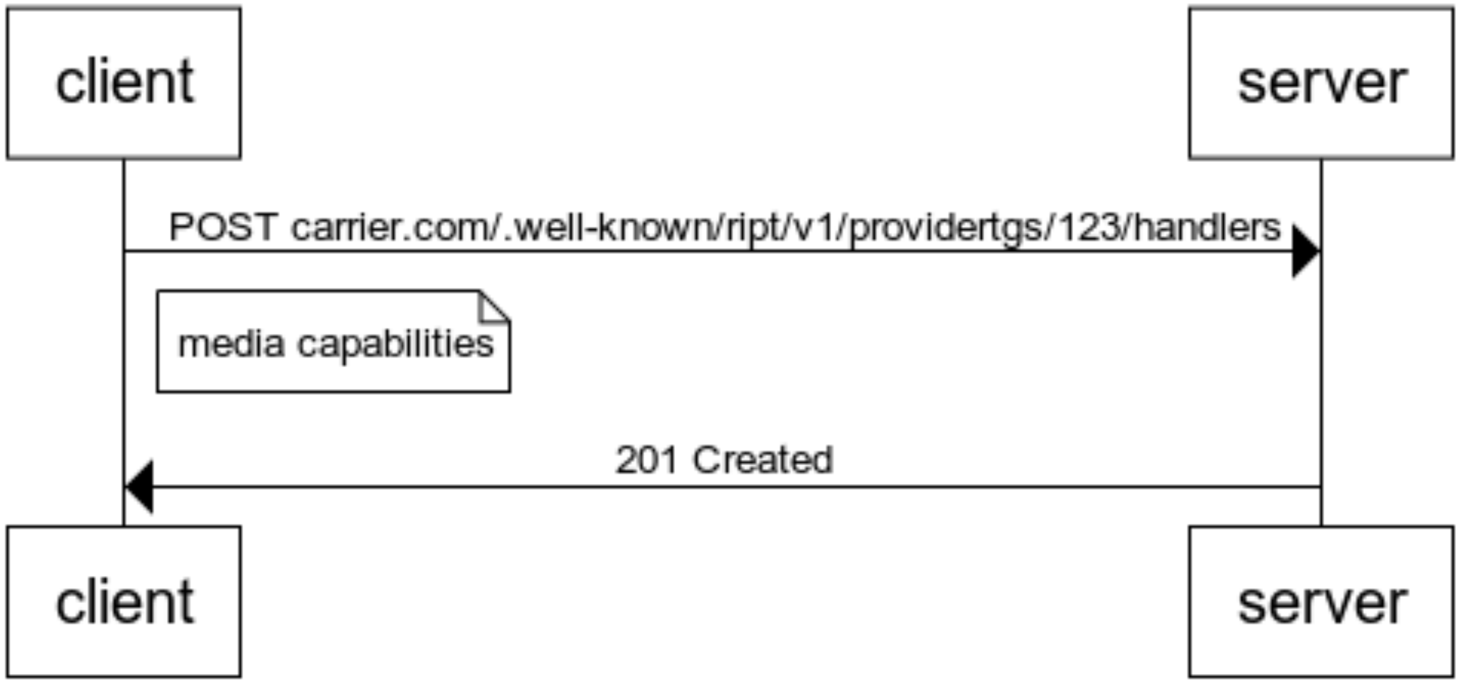

Which was sent to a “tg” – a bag for policy and identity management, which was gotten…

By GETting to /providertgs
And for receiving inbound calls – POSTING to /consumertgs

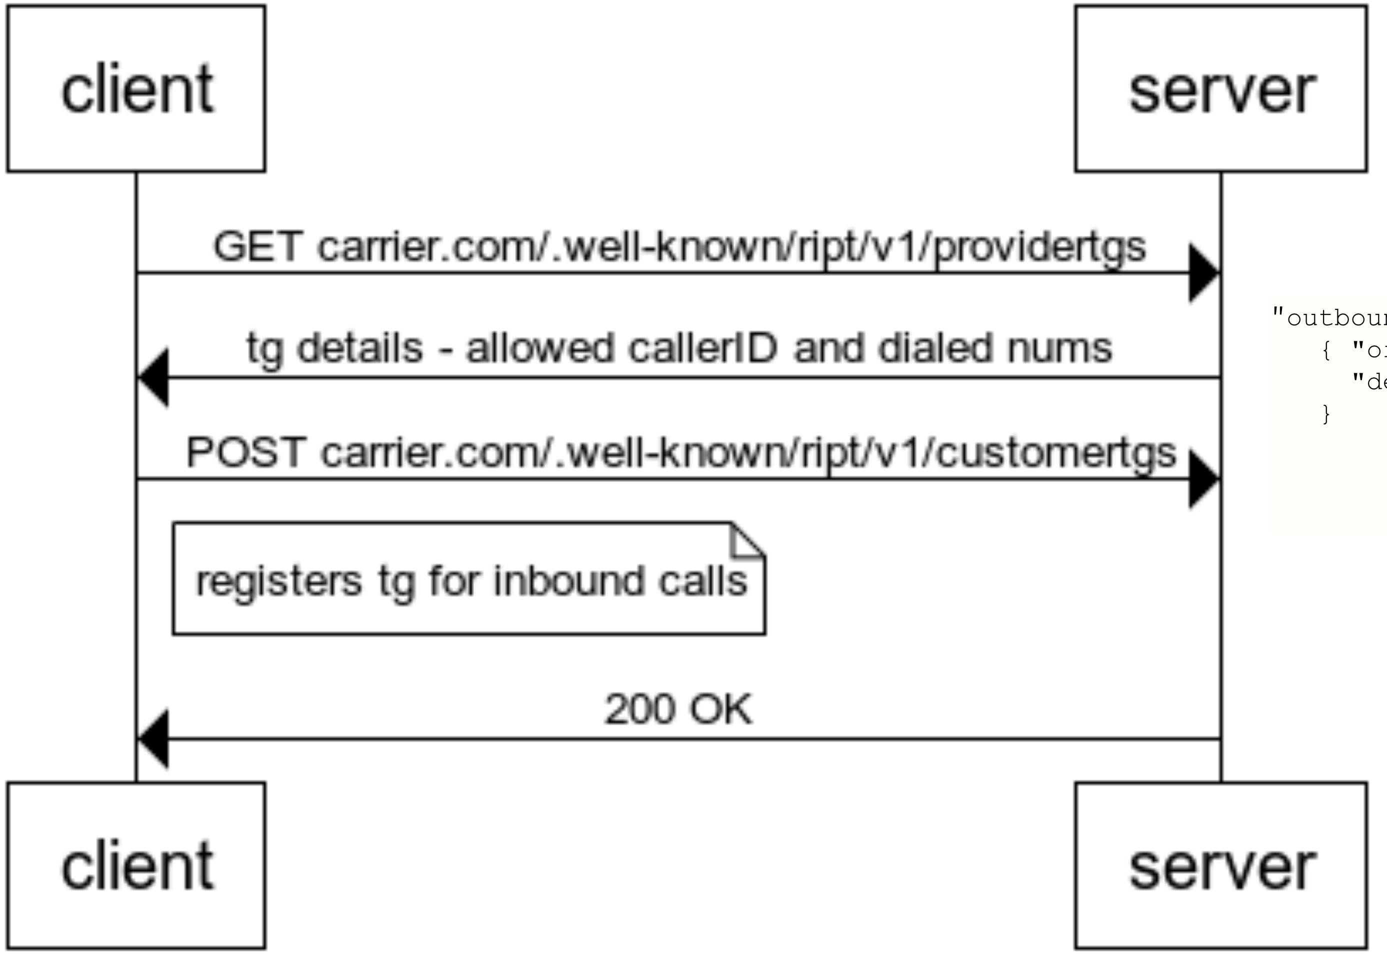


```
"outbound":
   { "origins" : "RFC 8226 cert with +14085551000 and +14085551002",
     "destinations" : "+1*"
   }
```