IETF-LISP-NEXAGON

IETF 115 Update Sharon Barkai, Nexar

- **H3 to EID LCAF Specification**
- **H3 and BDD Reviews on Lisp-List**

- **AECC KDDI-Oracle-Nexar Public PoC**
- **LISP Application Routing Economics**

H3 LCAF

- Added H3 LCAF to draft in Client-Service Networking and IANA sections

- The specification shows how to encode an H3 ID to an EID + HID type 17

- Updated references from bis drafts to LISP RFCs

The following Lisp Canonical Address Format (LCAF) [RFC8060] is used to encode H3-IDs into IPv6 address:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           AFI = 16387         |     Rsvd1     |     Flags     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Type = 17 HID  |     Rsvd2     |           Length = 8          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       HID (high-order)                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       HID (low-order)                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: LCAF for encoding HIDs in H3ServiceEIDs

Value	LISP LCAF Type Name	Reference
17	H3 ID	Section 4

[RFC9300] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, Ed., "The Locator/ID Separation Protocol (LISP)" , RFC 9300, DOI 10.17487/RFC9300, October 2022, <https://www.rfc-editor.org/info/rfc9300>.

[RFC9301] Farinacci, D., Maino, F., Fuller, V., and A. Cabellos, Ed., "Locator/ID Separation Protocol (LISP) Control Plane", RFC 9301, DOI 10.17487/RFC9301, October 2022, <https://www.rfc-editor.org/info/rfc9301>.

H3Geo + BDD Draft Reviews

- H3Geo Isaac Brodsky, Nick Rabinowitz submitted review to list
  - Verify use of the H3 grid hierarchies for detections and EIDs
  - Major insight regarding H3-EID calculation from detection HID
- BDD Prf. Trevor Darrel and Prf. Fisher Yu submitted review to to list
  - Verify assumptions on detection/localization by moving vehicles
  - Major insight to enumerate speed-limit-signs detections separately

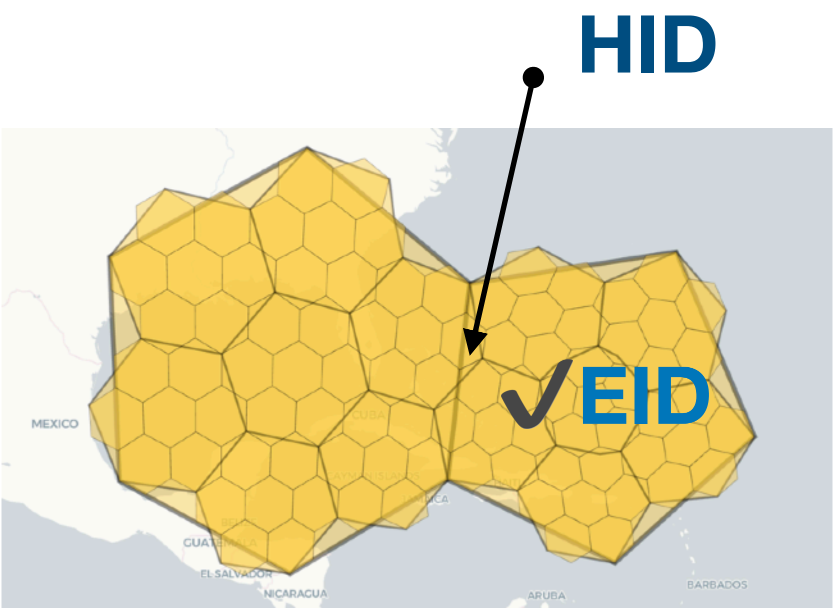

AECC LISP-NEXAGON POC

- RTM: Free Parking PoC
- KDDI SIMs + MEC RTR
- Oracle GPUs + Edge RTR
- Nexar AI-CarCam Stack
- Nexagon Geospatial Twins
- Interoperable Scale out
- EID based Geo-Privacy
- Next Phase: 🚧Blockages, 🔌Power, 💨Emissions, 💵Costs

ORACLE CLOUD

KDDI

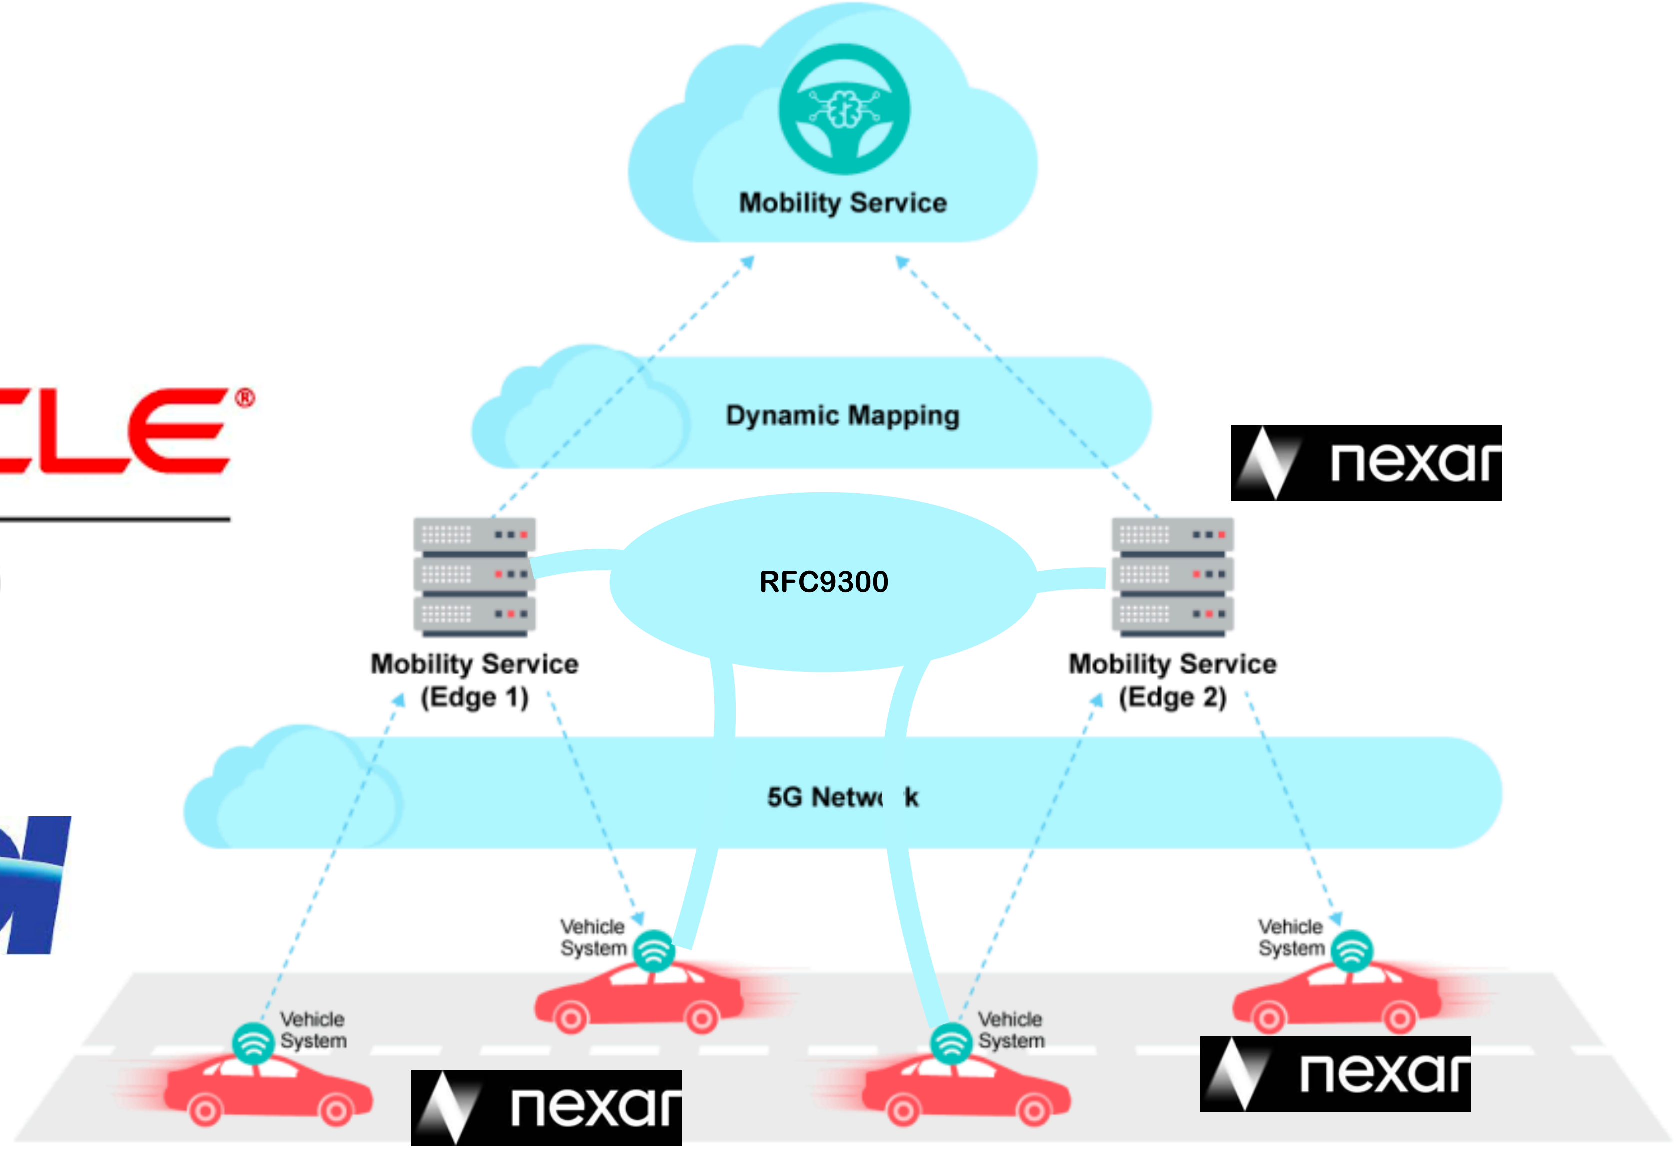

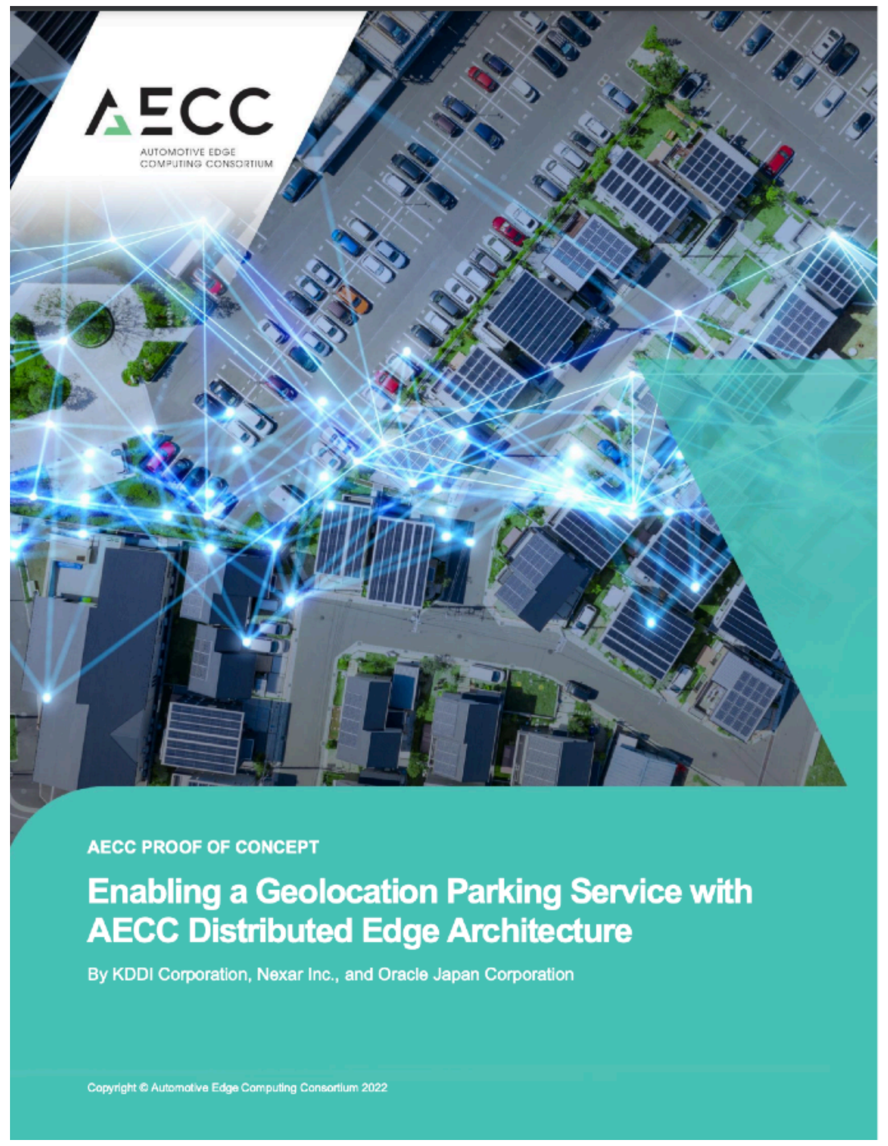

AECC AUTOMOTIVE EDGE COMPUTING CONSORTIUM

Proof of Concept Results

Found free parking spaces. Tests confirmed the ability of the service to accurately detect free parking spaces along street sides. Accuracy was important because the visual presence of space is irrelevant for parking if it is caused by presence of temporary warning cones or space permanently reserved for emergency parking or no parking at all.

Distributed processing did not confuse vehicles. Test vehicles were in continuous motion, transmitting geospatial data to two different edge networks for local preprocessing, then on to central cloud servers, then back to the vehicles. The system of distributed processing worked smoothly irrespective of which vehicle was crowdsourcing or receiving data while moving within or between either edge network's zone.

No service disruption while protecting privacy. The service operated continuously and protected privacy of each car's data.

Can work at commercial scale. The test proved vision AI technology can scale to process uploads from many vehicles (M) at the same time to enable a reliable detection service for commercial deployment. It proved distributed simultaneous processing of two street segments by two different processes in two different edge networks. By induction, such partitioning that works for two segments will linearly scale to all street segments

Additional Planned Extensions

- Realtime Mapping
- Using Wifi p5g breakouts
- Green edge location distribution
- Cisco, Toyota

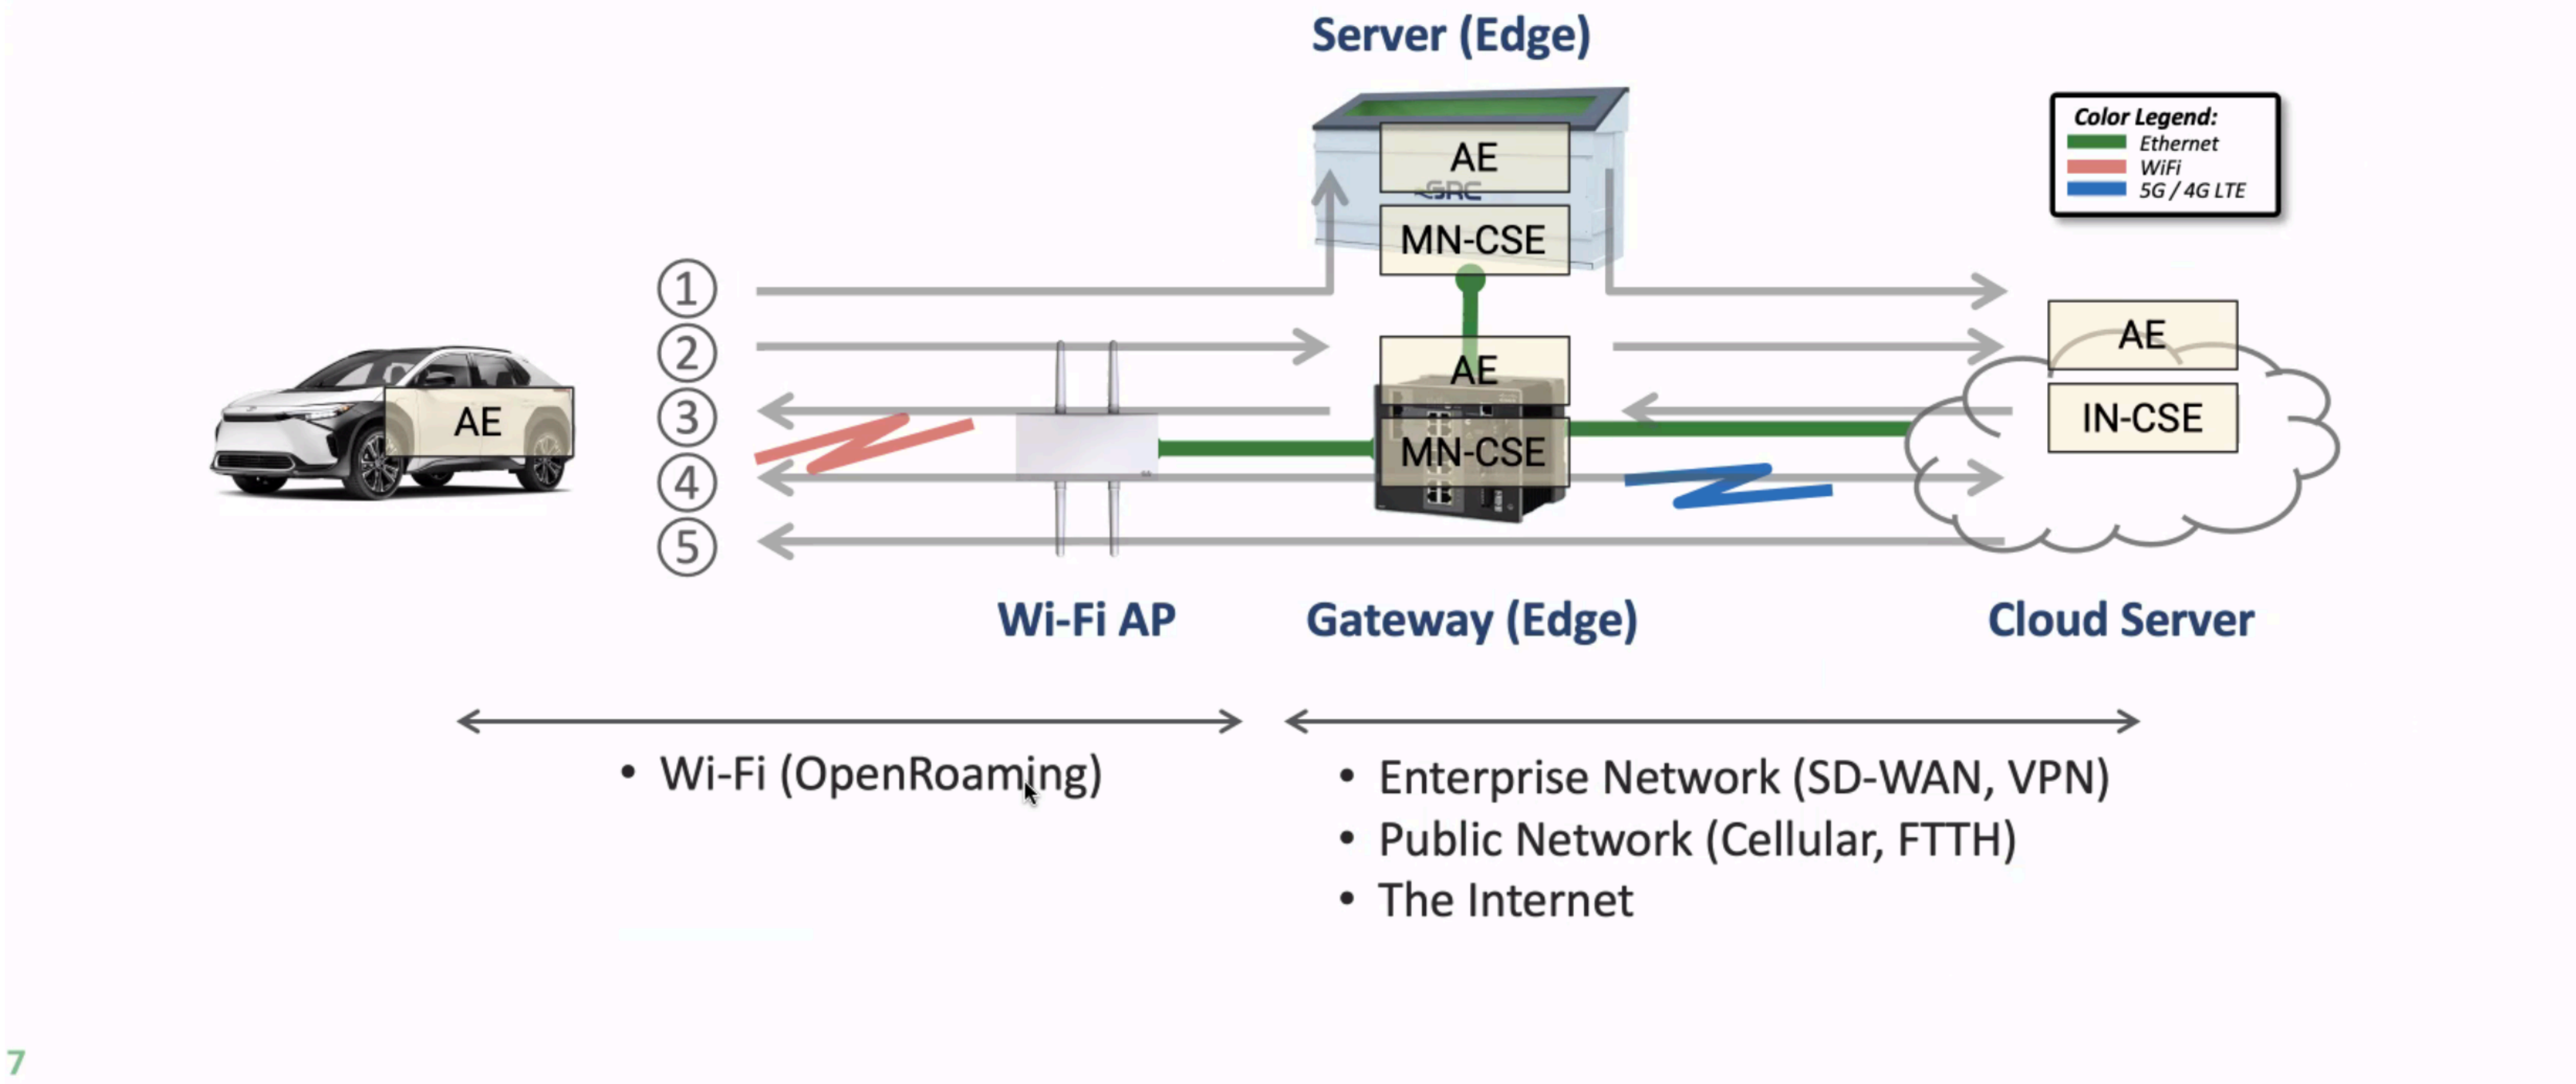

Automotive Digital twin And LISP Economics

From Wikipedia, the free encyclopedia

A **digital twin** is a virtual representation of a real-world physical system or product (a *physical twin*) that serves as the indistinguishable digital counterpart of it for practical purposes, such as system simulation, integration, testing, monitoring, and maintenance. A digital twin can, but must not necessarily, be used in real time and regularly synchronized with the corresponding physical system. A pragmatic litmus test of the efficacy of digital twin is for a system tester to run a robust system verification & validation test suite on both the digital twin and the physical twin.[1] When the system tester cannot reliably distinguish between the digital twin and the physical twin with a high probability, the former is a bona fide digital twin.

In the Actual Field (vs closed tracks and factories):

- *Drastically changing number of sources during the day*
- *Elastic allocation across edge compute locations*

- *Moving sources/destinations and access selection*
- *Sensors (vehicles) change twin (geo) association*

- *Geoprivacy*

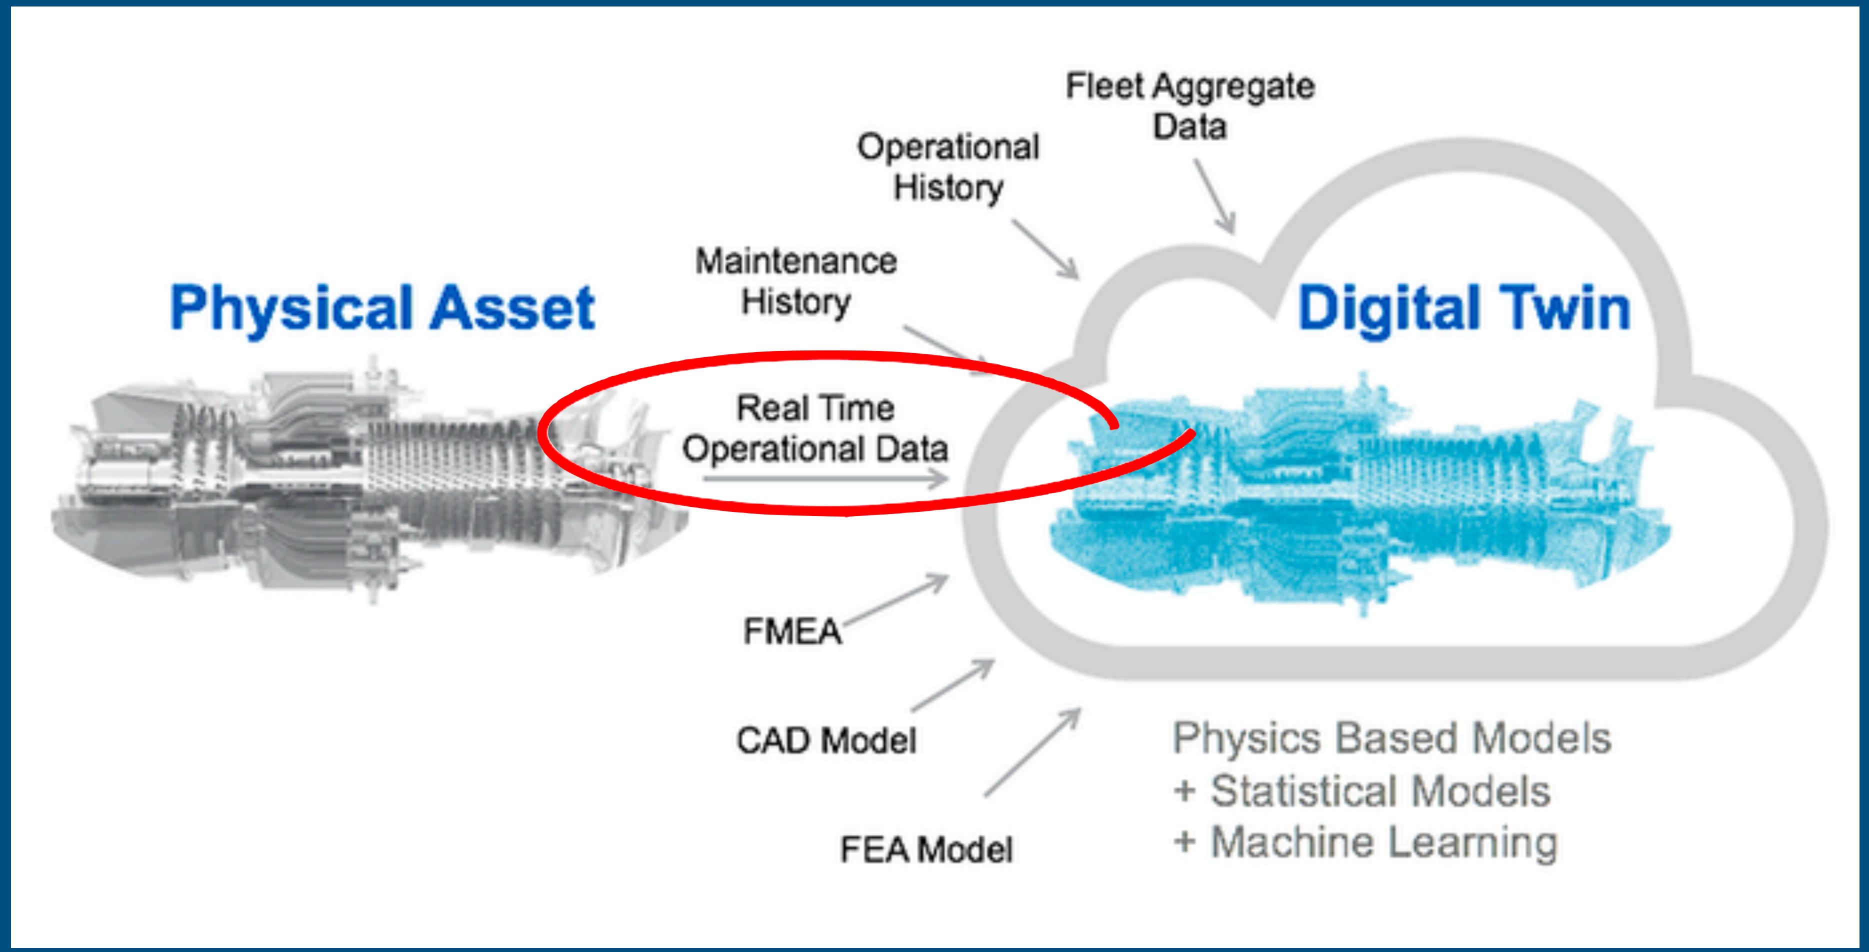

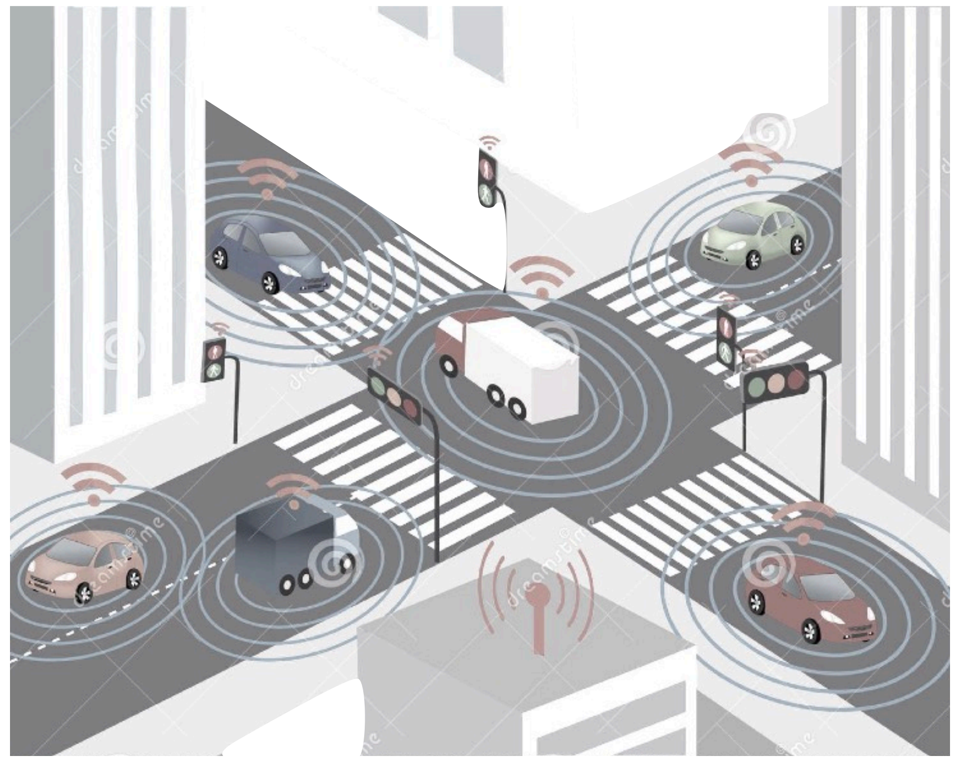

Edge Unlimited Distributed Capacity ↔ Resource Fragmentation ↔ LISP RFCs

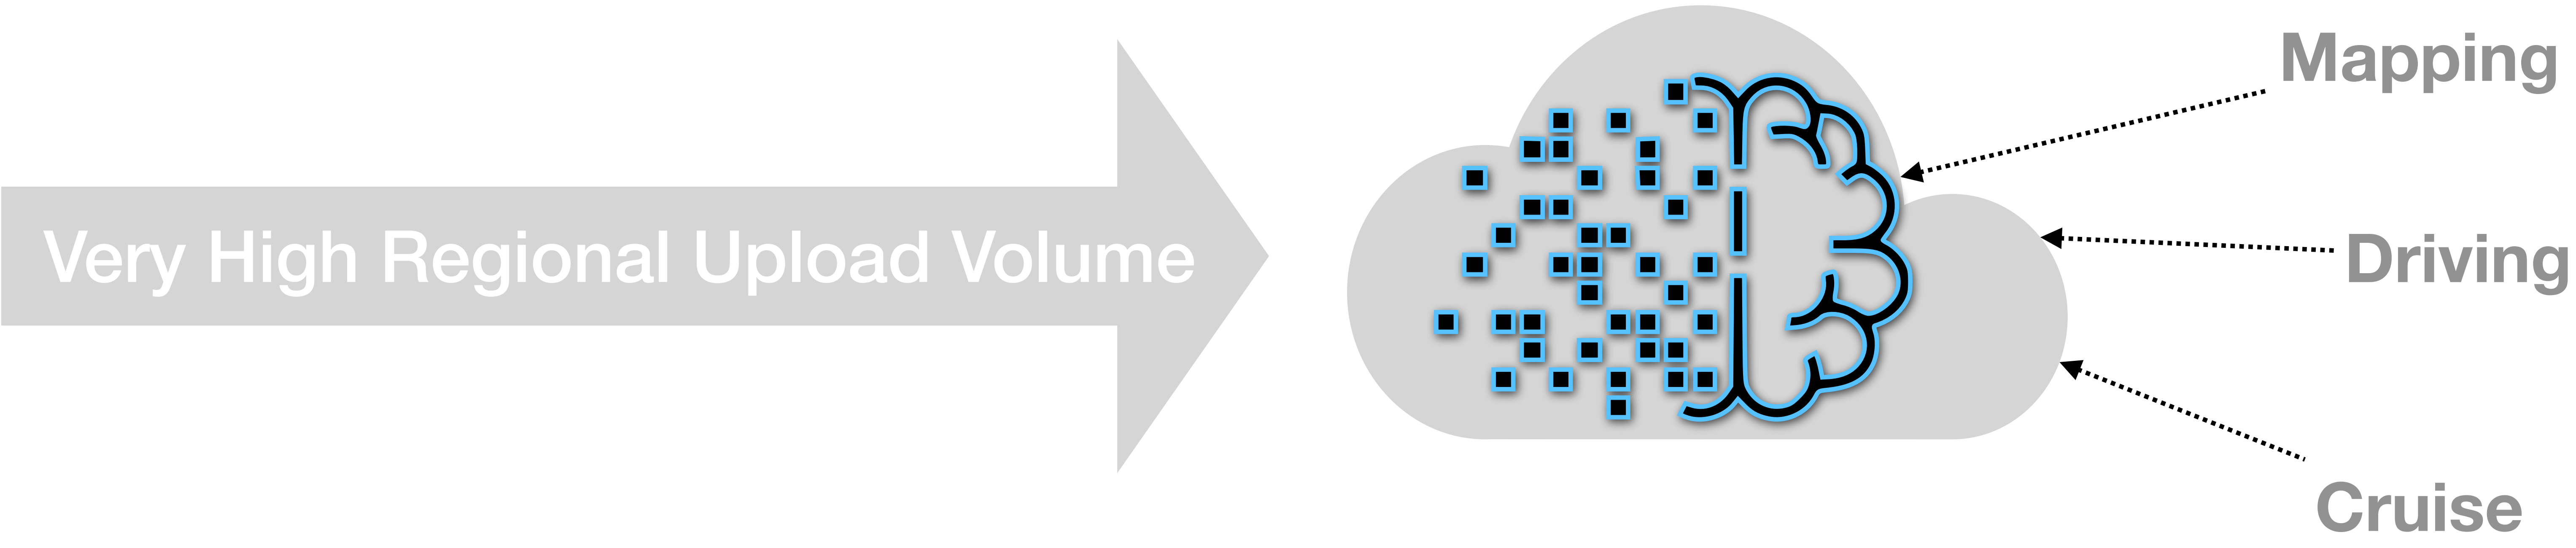

Routed Geospatial Twins - Interoperable Flexibility - Cost & Capacity

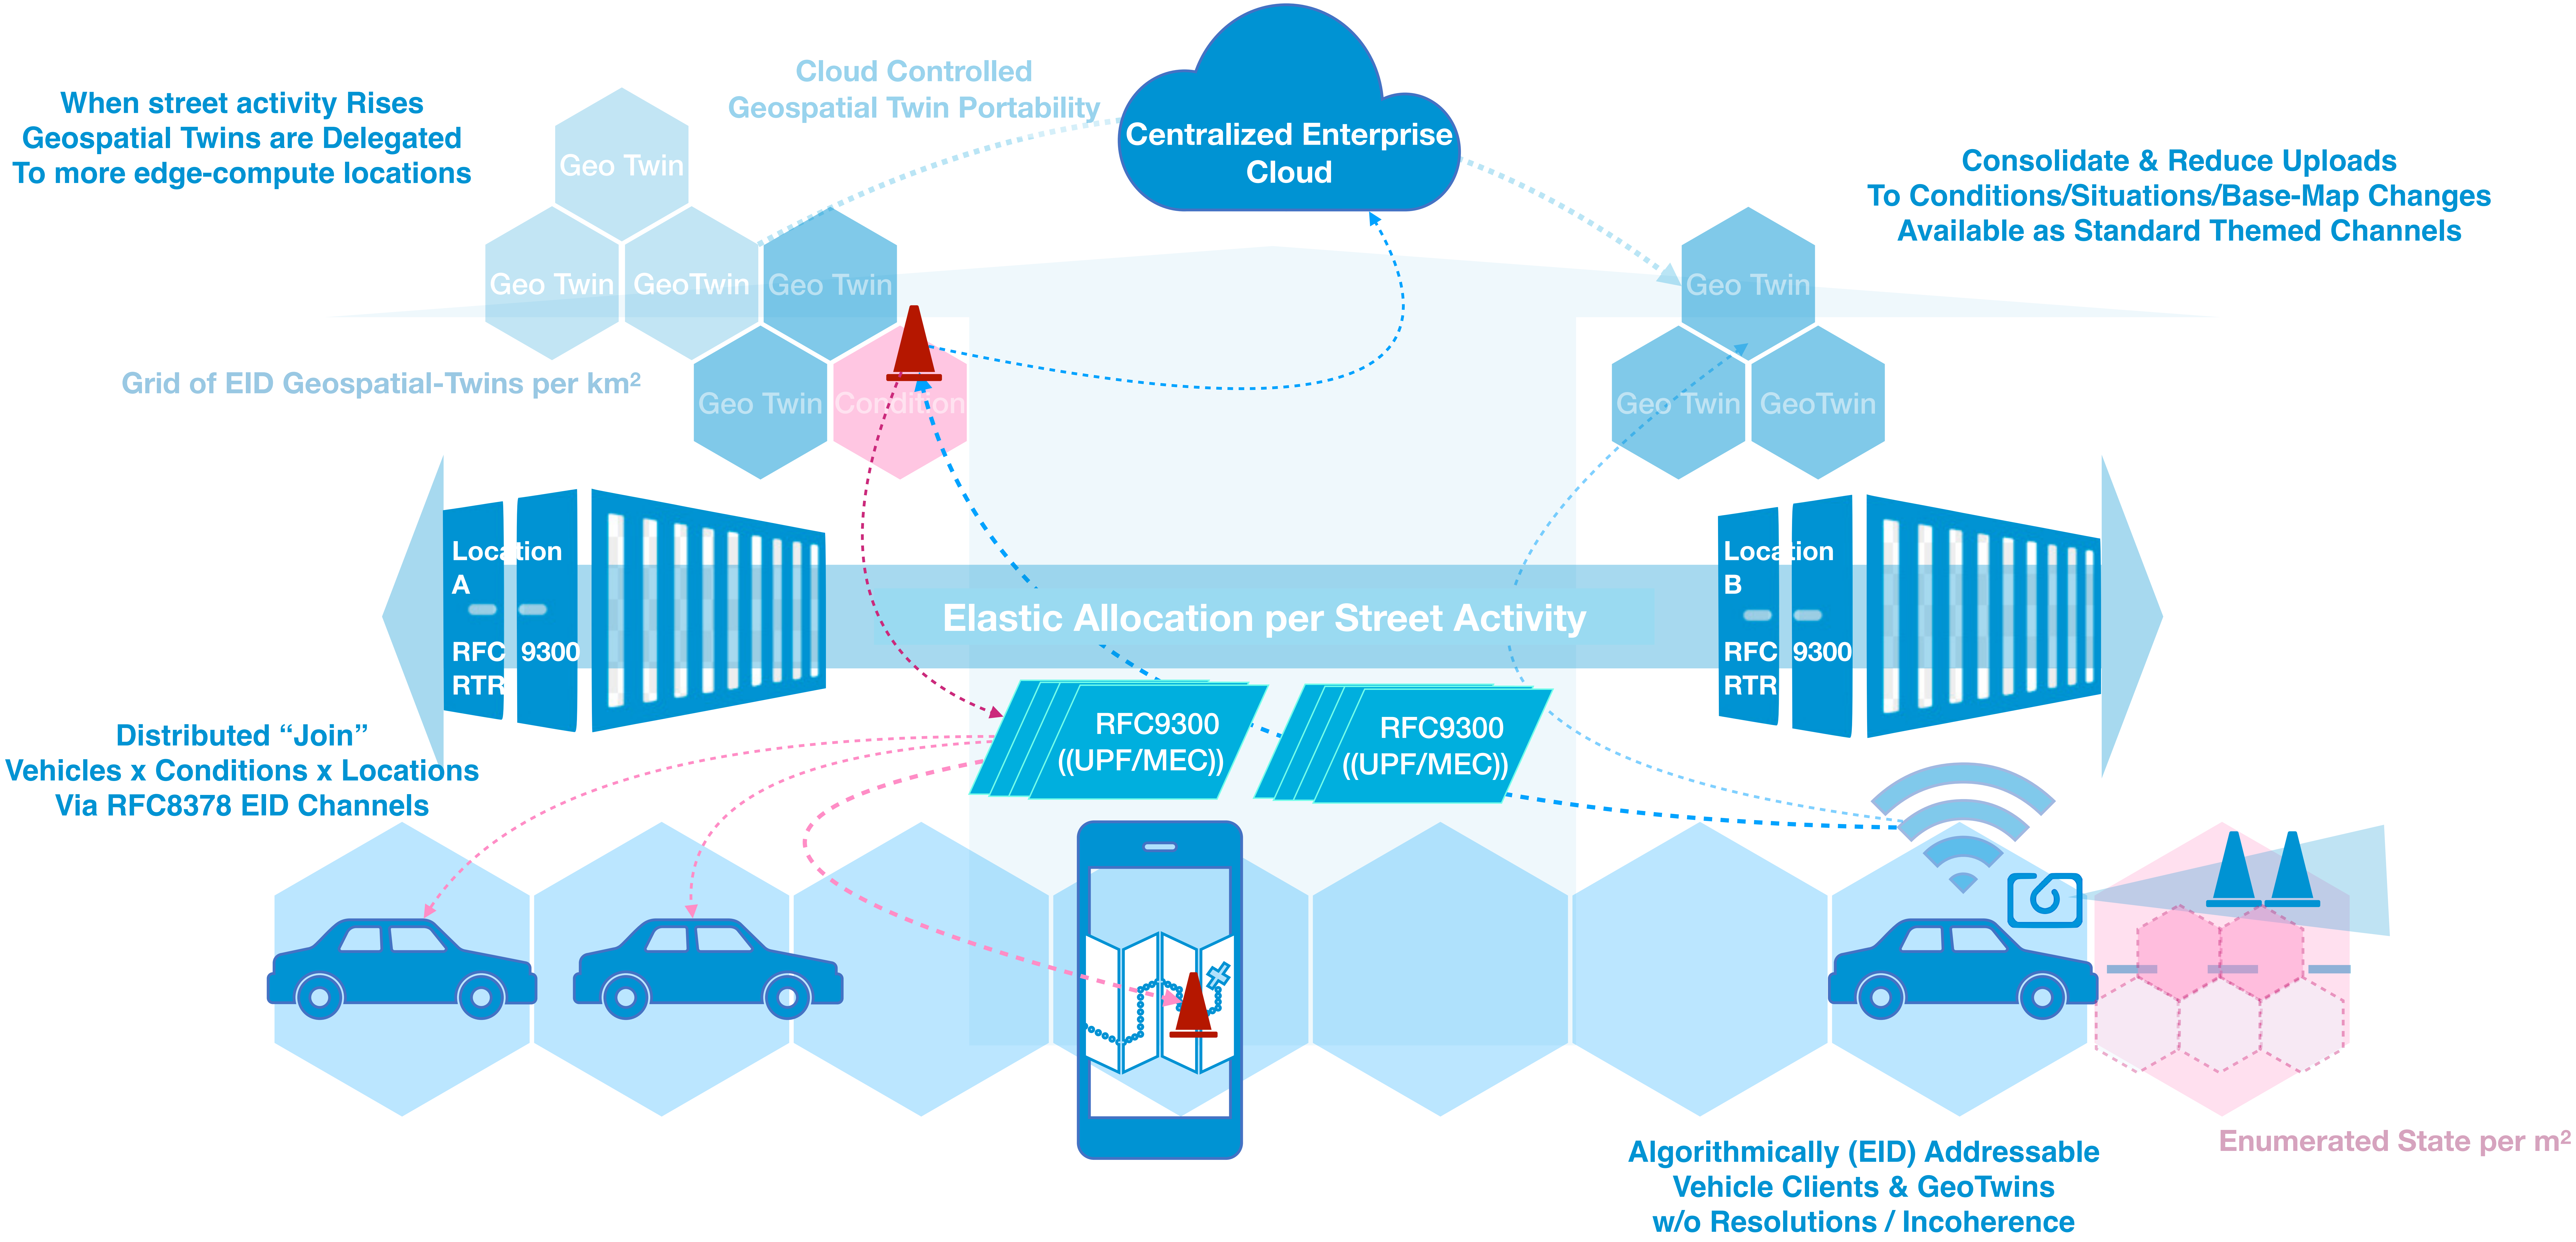

LISP Ties Fragmentation: Low-Cost High-Capacity

< $100k *MRE, 1k FPS per Automotive Metro Area Edge Cloud

10x $10/h p5G RU
100 floating UEs Each
RFC9300 UPF

₵35/kWh x 14kW +$10/h dep.
Or 16 x $1/h GPU
RFC9300 EdgePods/Zones RTR

EC2
S3

100k FPS

1k FPS

$100/h

<< x24 x30 >>

$15-16/h

10 FPS

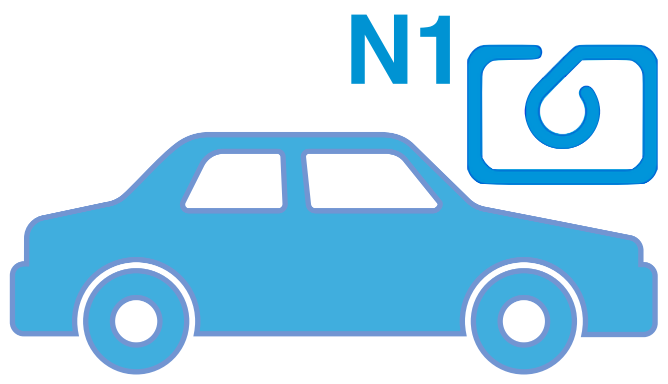

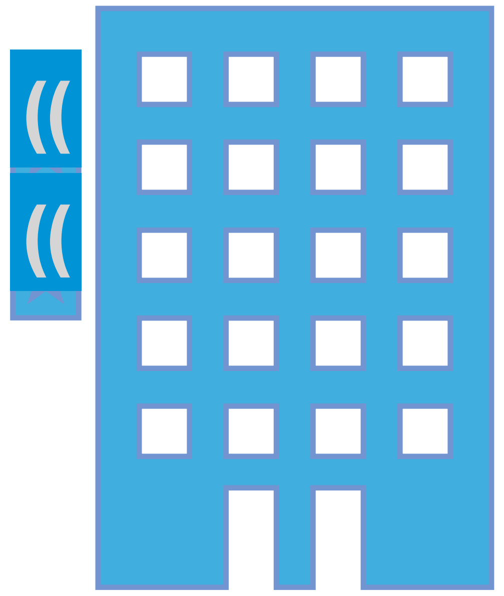

RFC9300
MAN

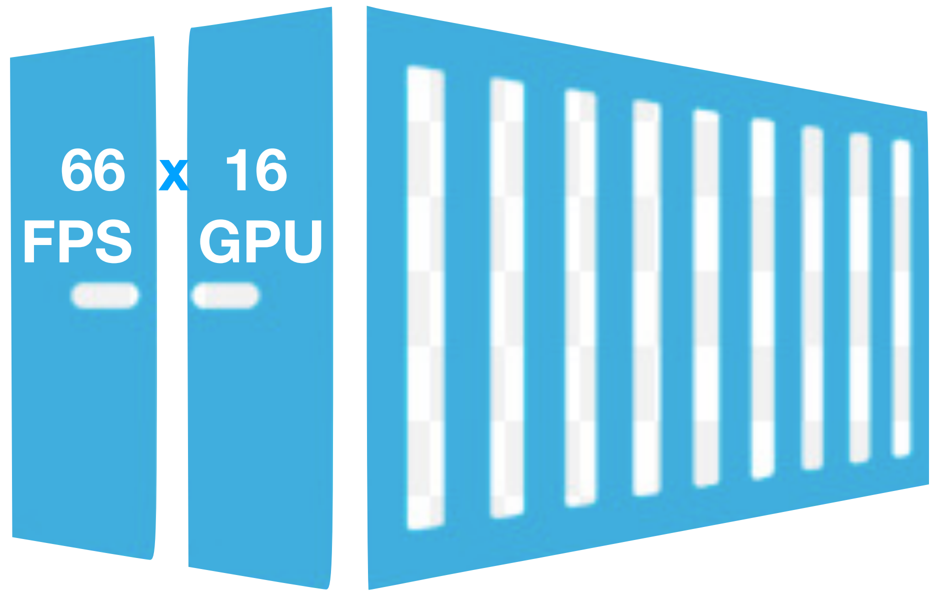

RFC9300
WAN

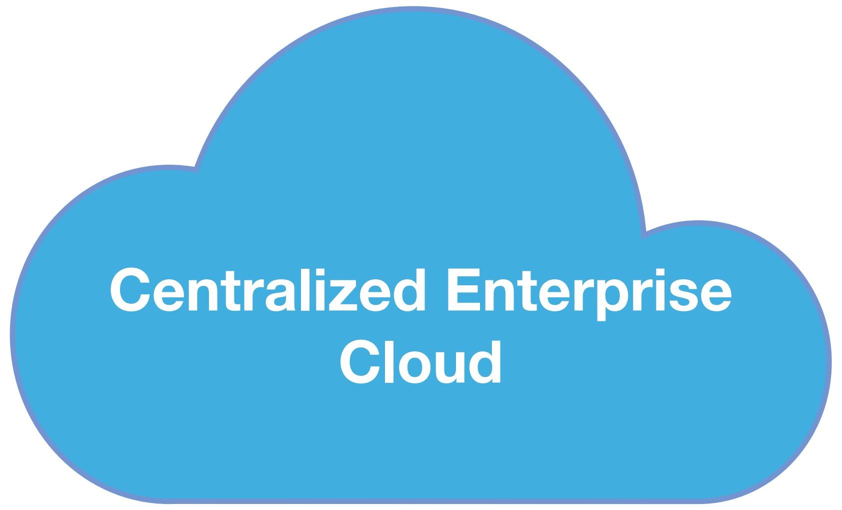

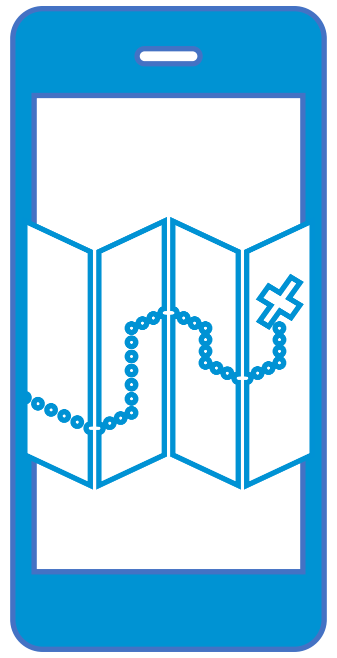

$1/Month
Per Themed Channel Feed
RFC9300 Clients <<

Standard Monthly Plans

Geospatial Twins <<
IETF-LISP-Nexagon RFC
<< RFC9300 Channels >>

<< control Mobility Service Provider
(Nexar)
>> Data & Evidence Backend

LISP Mobility Edge Economics

Each street detection frame objects + meta-data = 100KB
Benchmarking 1k FPS 2.6b frames per month => 260TB
60-66 FPS/GPU Edge 16 GPUs Vs $100/M Frames Cloud

- $4/GB Continues Spectrum Upload: 260k x 1GB x $4 = $1m per month
- $100/h 1kUE 5G Hotspots Upload: $100 x (24 x 30) = $72k per month
- $100/M frames Cloud GPU = 2.6b frames / m x $100 = $260k per month
- $0.35kWh Edge GPU x 14kw + $10/h dep. x (24 x 30) = $10k per month

16 GPU DGX power and hourly depreciation

LISP Application Routing Based Mobility/Elasticity => 15x cost reduction
=> Viable vs Not Viable for Mass Consumer Scale Reality+

Thank You