

IETF
Berlinq
19th July 2016

FedEx – Extensive Federation

<http://www.cl.cam.ac.uk/~jac22>

Team Jon

Assume an edge cloud world where we embed significant processing, wireless base station, and storage in pretty much every building, vehicle and in your pocket/wearables.....what scale of federation are we looking at?

Outline

1. Smart City like london
 - 10M people with 1000 devices=10Billion
 - Lucky we already did IPv6
 - No clue how to do 10^{10} Ass
- And dynamics
 - assume 5M people move per hour
 - Device set changes –
 - From home, to car/bus to work/school/shop
 - Need to group re-key associate 5Billion potentially
 - If we want perfect fwd/backward secrecy/privacy
 - Which I assume we do.

The baseline 21 years ago for IPng requirements assumed this figures – I think they stand the test of time well. Of course the model was that most devices would do pairwise communication, RPC like. And that the traffic matrix would be client/server – we now see this is sort of true with client device talking to cloud services so massive aggregation of services in a small number of data centers, at least for smart phones/tablets and some IoT systems – there's no reason for this, and an edge cloud system serving local IoT devices and sourcing personal data (mail/messaging/social media) in a P2P way seems perfectly feasible, and aligned with the large deployment of base stations necessary to meet 5G capacity requirements, plus resilience, latency needs of users and applications and so on

Simplifying assumptions

- Most devices group around a principle/person
 - Or possibly vehicle, or place
- Each principle (person) has a small social net
 - 150 mean friends (dunbar)
 - May be subsetted by context
 - (home/kin, work/colleague)
- Life complicated by multipath&replicas
 - Needed for high availability (6-fold)
 - Requires scalable consensus algorithms

People, society and devices will socially cluster – this gives us good locality for many services – for high availability, though, to meet similar Mean Time Between Failure of cloud service, in a P2P/Edge Cloud world, we need approximately 6-fold replication (empirical studies show mean downtime of home broadband and cellular data service coverage...)

More simplifications

- Most data is append only
 - So consistency easy
 - Can certainly use block chain storage too
 - E.g. Irmin or other append only storage
- Notify services have to scale too
 - Not just for IoT adverts
 - E.g. for emergencies
 - But also for device software updates

If you need to run paxos across 6 copies in the lan, metro and wan net, forget it – but you don't need to as the update rate is low and only on a very small fraction of the actual flows –so lazy synchronisation/ replication will work well.....for much data – and serving some stale iot or social data to friends or smart meter reader apps is not necessarily a problem....just needs factoring in to the analytics algorithms

So future city = current internet

- Except has harsher requirements if fed
 - For privacy, integrity etc
 - For notification
 - For resilience
 - For consistency

Smart homes putting the total digital footprint of all occupants into a centralised cloud service would represent a massive privacy problem and possibly security threat – better to keep decentralisation – reduces the attack surface n -fold for n people 😊

Questions?

- I'm happy to respond to followups.
jon.crowcroft@cl.cam.ac.uk