

The Extensible Internet (EI)

Based on a CCR editorial with 18 authors:

Hari Balakrishnan, Sujata Banerjee, Israel Cidon, David Culler, Deborah Estrin, Ethan Katz-Bassett, Arvind Krishnamurthy, James McCauley, Nick McKeown, Aurojit Panda, Sylvia Ratnasamy, Jennifer Rexford, Michael Schapira, Scott Shenker, Ion Stoica, David Tennenhouse, Amin Vahdat, Ellen Zegura

Core design team: Arvind Krishnamurthy (UW), James McCauley (MHC), Aurojit Panda (NYU), Scott Shenker

Goal Of This Talk

- Give a very very brief overview of EI
- Just enough to initiate a discussion
- More details in follow-up conversations

Enabling Architectural Change?

- In this talk, “architecture” refers to overall arrangement of dataplane functionality
 - Arch. = Layers and their basic functions
 - Not specific protocols (e.g. IPv4 vs IPv6)
 - Not control plane (e.g., SDN, etc)
- Decades of architectural research...
 - ...but no discernable architectural impact
- The public Internet seems doomed to architectural stagnation

But Not The Hyperscalers!

- Cloud/content providers building private IP-based networks with many PoPs
- These PoPs apply *in-network* services:
 - Flow termination
 - Caching
 - Load balancing
 - ...
- Significant impact on latency/reliability

Good for Apps, Bad for Internet

- Internet increasingly *balkanized*
 - Private nets define own in-network services
 - Public Internet lagging behind
- Rather than these ad hoc and proprietary designs, we need a coherent architecture
- *But isn't Internet impossible to change?*
- *We don't think so; this talk is about why*

Rest of Talk: Four Questions

- Why is the Internet so hard to change?
- How can we overcome this barrier?
- What would this new Internet look like?
- Given the failures, why is now different?

The First Question

- **Why is the Internet so hard to change?**
- How we can we overcome this barrier?
- What would this new Internet look like?
- Given the failures, why is now different?

L3 Has Two Roles in Architecture

1. Allows all L2 networks to interconnect
 - Must be deployed in every router
 - Almost impossible to make significant change
 - Both now, and in the future...
2. Provides the service model to hosts
 - Must support all application requirements
 - Requirements becoming more stringent
 - Which is a reason to change the architecture

This Dual Role Prevents Change

The network must meet additional application requirements, but...

...the only layer that can address those requirements is the only layer that can't be changed!

The Second Question

- Why is the Internet so hard to change?
- **How we can we overcome this barrier?**
- What would this new Internet look like?
- Given the failures, why is now different?

Proposal: Extensible Internet (EI)

- Use current IP protocol/infrastructure
 - Remains essentially unchanged
- Introduce a new network layer above L3
 - Service layer (SL) or L3.5
- SL offers new in-network services to hosts
 - e.g., flow termination, caching, etc.
 - And beyond....(more later)
- *EI is merely the architecturally coherent version of hyperscaler private networks*

EI Decouples L3's Two Roles

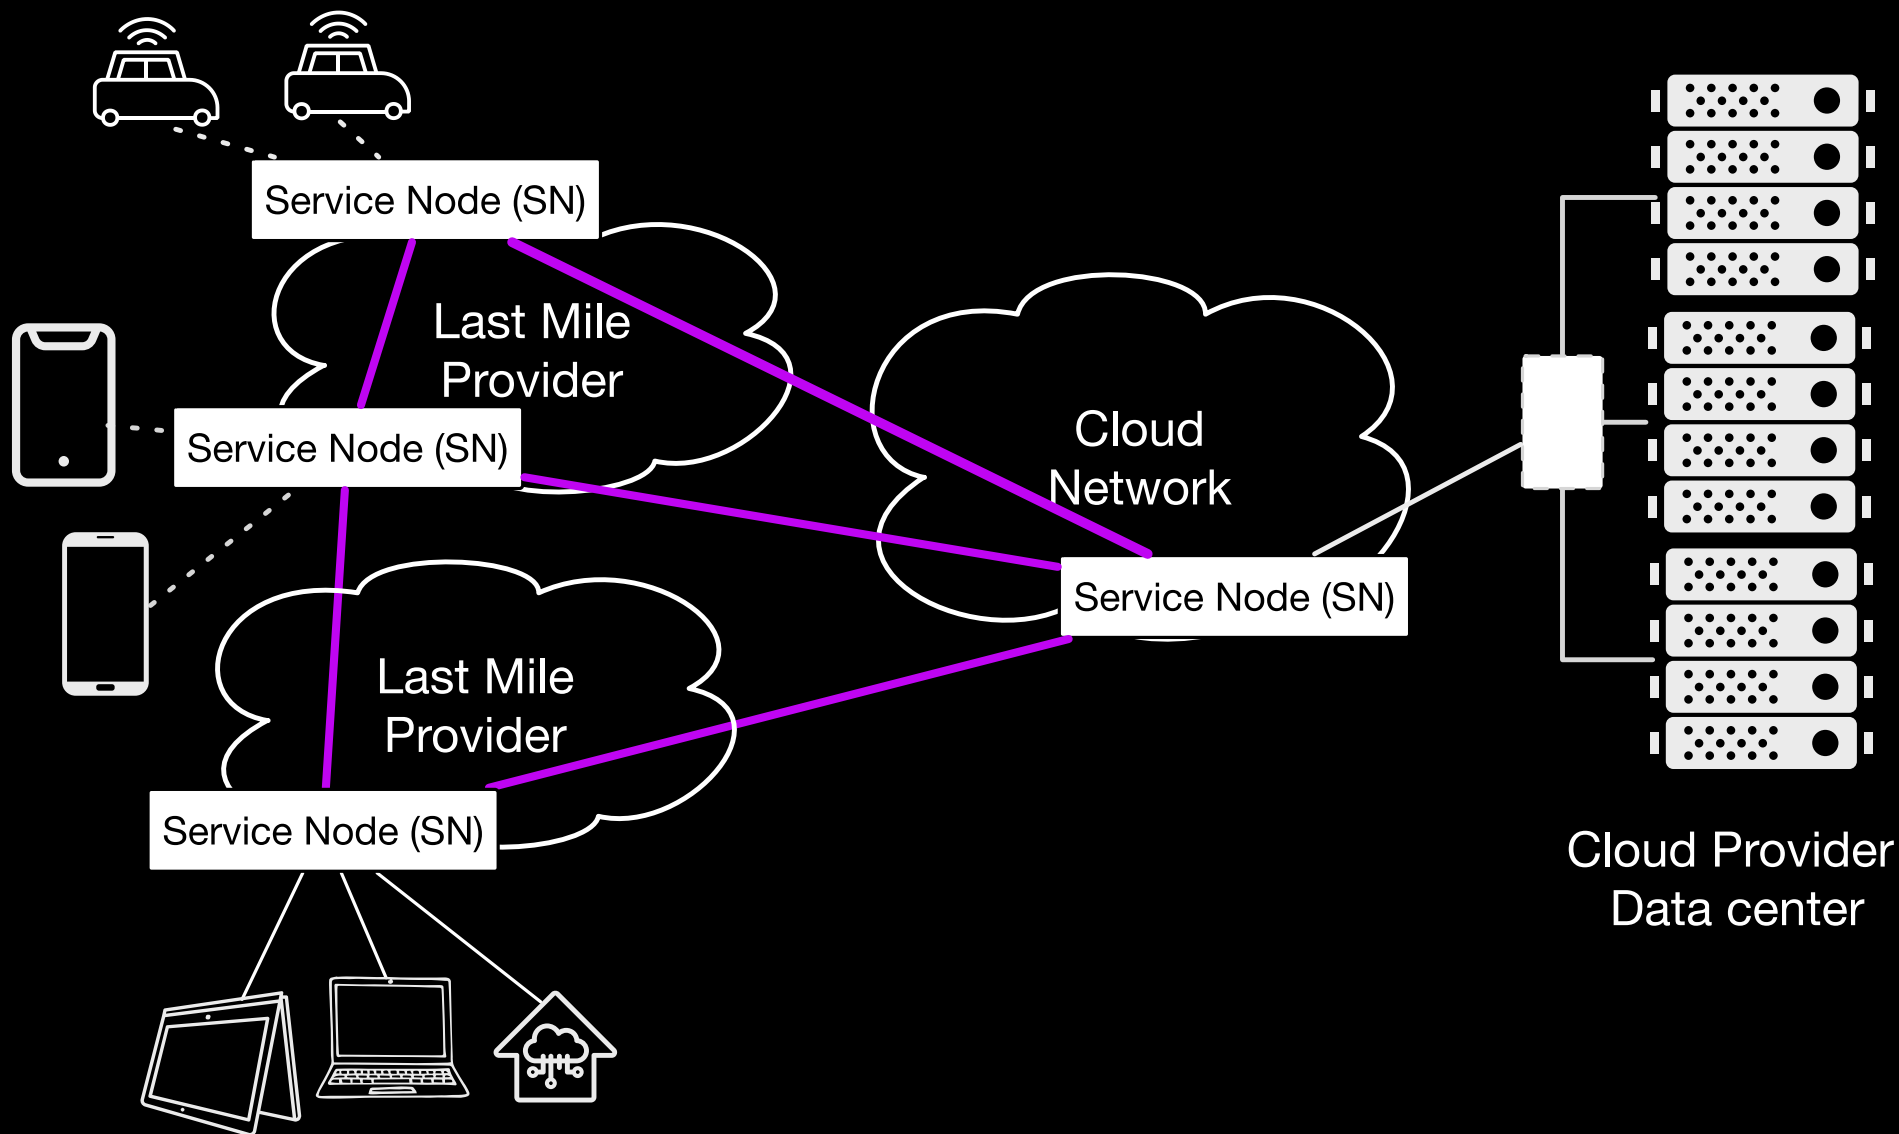
- Current L3 is the interface to both
 - L2 networks: so it must be narrow waist
 - L4 on hosts: defining Internet's service model
- In EI
 - L3 is still the interface to L2 networks
 - But L3.5 is the interface to L4
- No reason for L3.5 to be in every router
 - So it can be more easily deployed/changed
- No reason for L3.5 to be “narrow”
 - So it can meet a wide range of needs

The Third Question

- Why is the Internet so hard to change?
- How we can we overcome this barrier?
- **What would this new Internet look like?**
- Given the failures, why is now different?

The Service Layer (SL)

- Implemented at service nodes (SNs)
 - SNs are servers deployed at network “edge”
 - Every host associated with a service node
- Typical communication pattern:
 - Source → SN → SN → Destination
 - Service layer communication tunneled over IP
 - Source specifies which service to invoke at SN
 - So services need not be backwards-compatible



Key: All Services In Software

- Standards: Open-source code, not specs
- 3 necessary software components on SNs
 - Standardized service modules
 - Standardized execution environment (WORA)
 - Open-source runtime/orchestration available
- Services involve limited computation
 - Packet forwarding and payload processing
 - Simple functions like caching
- SNs could have enclaves, hw accelerators

Choosing The Services

- IETF-like governance determines the set of *public* services and implementations
 - Public services are deployed on all SNs
- EI's biggest change: deployment model
 - Services are approved software modules
 - No per-vendor/per-domain decision process

Some Potential Services

- Basics:
 - Flow termination, caching, load balancing...
- Support for other delivery models:
 - Multicast, pub/sub, redirection, QoS...
- Support for security and privacy
 - DDoS, attestation, oDNS-like, Tor-like,...
- Incorporate other frameworks
 - Istio, OPA, telemetry,...
- Support for new architectures
 - ICN designs, other radical redesigns,...

The Last Question

- Why is the Internet so hard to change?
- How we can we overcome this barrier?
- What would this new Internet look like?
- **Given the failures, why is now different?**

#1 Backwards Compatibility

- Can continue using traditional Internet
 - Nothing changes for hosts/domains if they choose not to use EI
- Addition of EI only requires:
 - Deploy SNs using EC or existing PoPs

#2 Fear Is A Great Motivator

- Internet architecture has resisted change
 - Most have given up on this...
- But there has never been an alternative
- There is now: the private cloud networks
 - The clouds could easily reduce the public Internet to primarily the last mile
- The Internet's choice: *change* or *shrivel*
- Extensible Internet is one possible change
 - That will preserve the Internet's role

#3 EI Based On Simple Conjecture

- That in-network support for applications:
 - Is important for many current/emerging apps
 - Can be done at SNs, rather than each router
 - Can be done in software, not ASICs
- **Evidence for conjecture: private networks**
 - *EI is merely the architecturally coherent version of these private networks*
- That's why we think EI might succeed...

Current Status and Future Plans

- Prototype of SN done (McCauley)
- Next steps:
 - Finish development
 - Deploy on FABRIC and other testbeds
- Next steps: Community
 - Engage research community (via FABRIC)
 - Continue discussions with industry
- ***We would love to have your participation!***

Thank You For Your Attention!

Questions?