CREATING THE LIVING NETWORK. TOGETHER.

Flattening the Protocol Stack of the Internet

New approaches to service provisioning

Dirk Trossen, Senior Principal Engineer

Prepared for publication in joint work with Sebastian Robitzsch (InterDigital), Scott Hergenham (InterDigital), Martin Reed (Essex University) & Janne Riihijarvi (RWTH Aachen)

Background & Drivers

Observed Trends

Move to Cloud-Native Operator Environments

- Micro-service vision with anything-as-a-service
- Efficient service message routing
- Regional data centres with SD-WAN transport (incl. L2 whitebox switching)

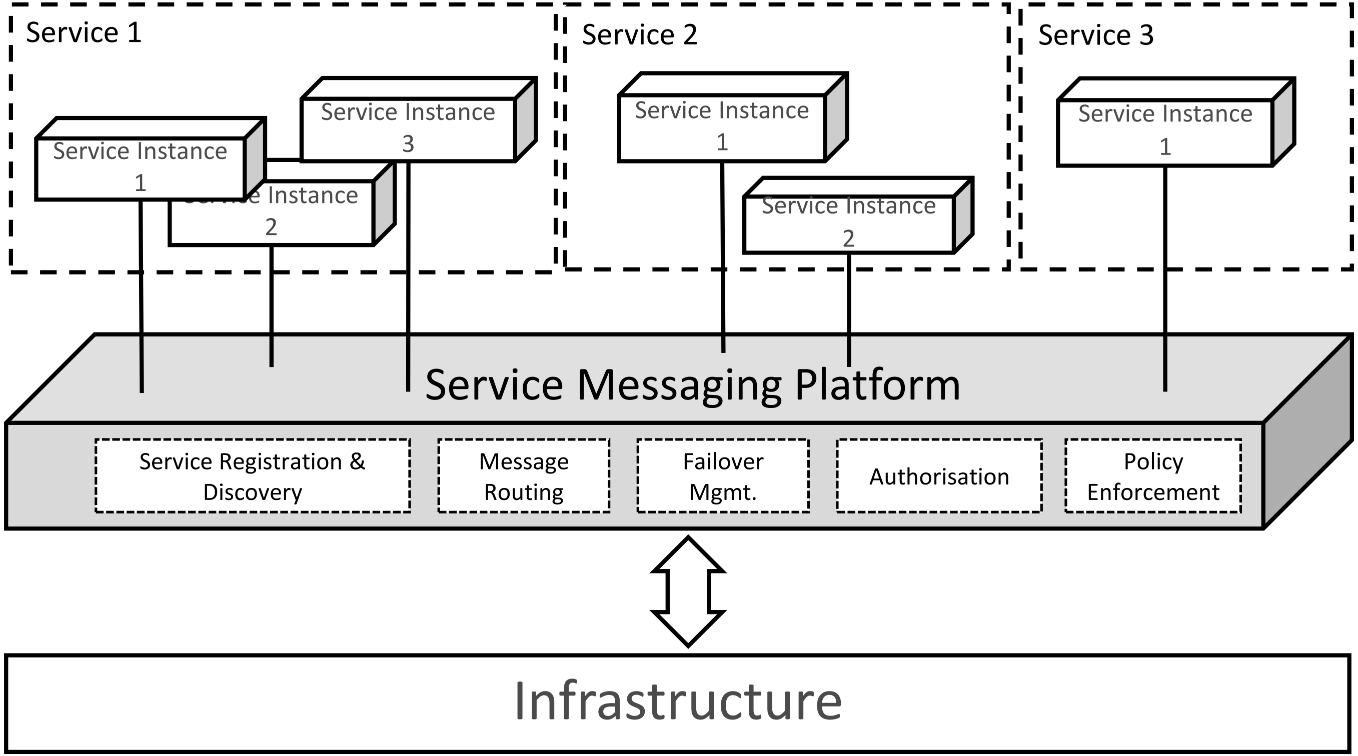

INTERDIGITAL.

Observed Trends

Micro-Services From Far-Edge to Distant Cloud

Anything-as-a-Service (new interactive, immersive experiences, localized where possible)

Service-based architecture across all edge devices and the Internet

Well-proven Internet technology, such as web services, HTTP, IP, … mixed with virtualization technology

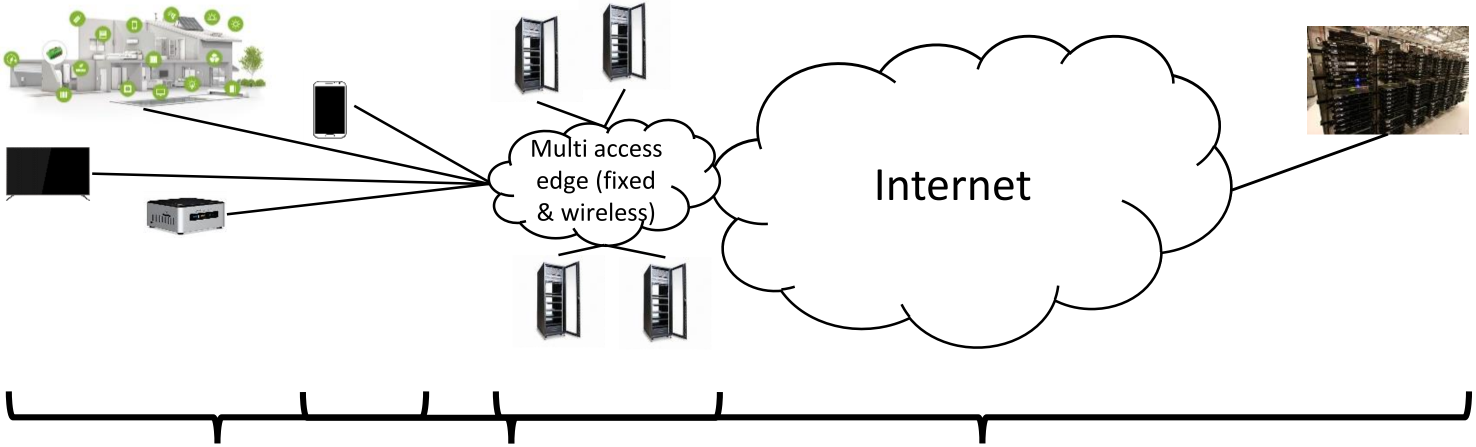

SBA-based far edge

SBA-based 5G cellular

Internet Connectivity

INTERDIGITAL.

Observed Trends

Services over Distributed (Cellular) LANs

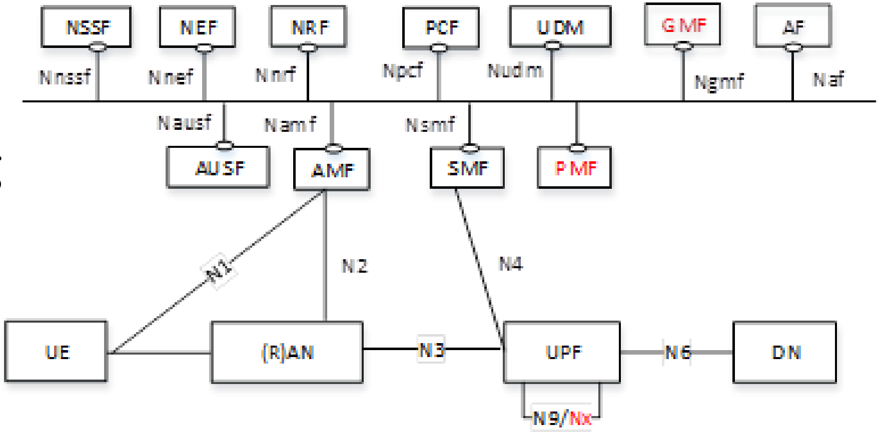

- Idea of every mobile terminal carrying a virtual Ethernet cable
  - Currently being specified in 3GPP Rel 16 FS-VertLAN SI
- Any service being an Intranet service with possible Internet backend connection
- Suitable for scenarios in, e.g.,
  - Industrial IoT or generally site-specific experiences, such as virtual tourist guides
  - ...ultimately any app-centric service experience (demo at upcoming MWC2019)?

5G is just the Starting Point Albeit a Necessary One!

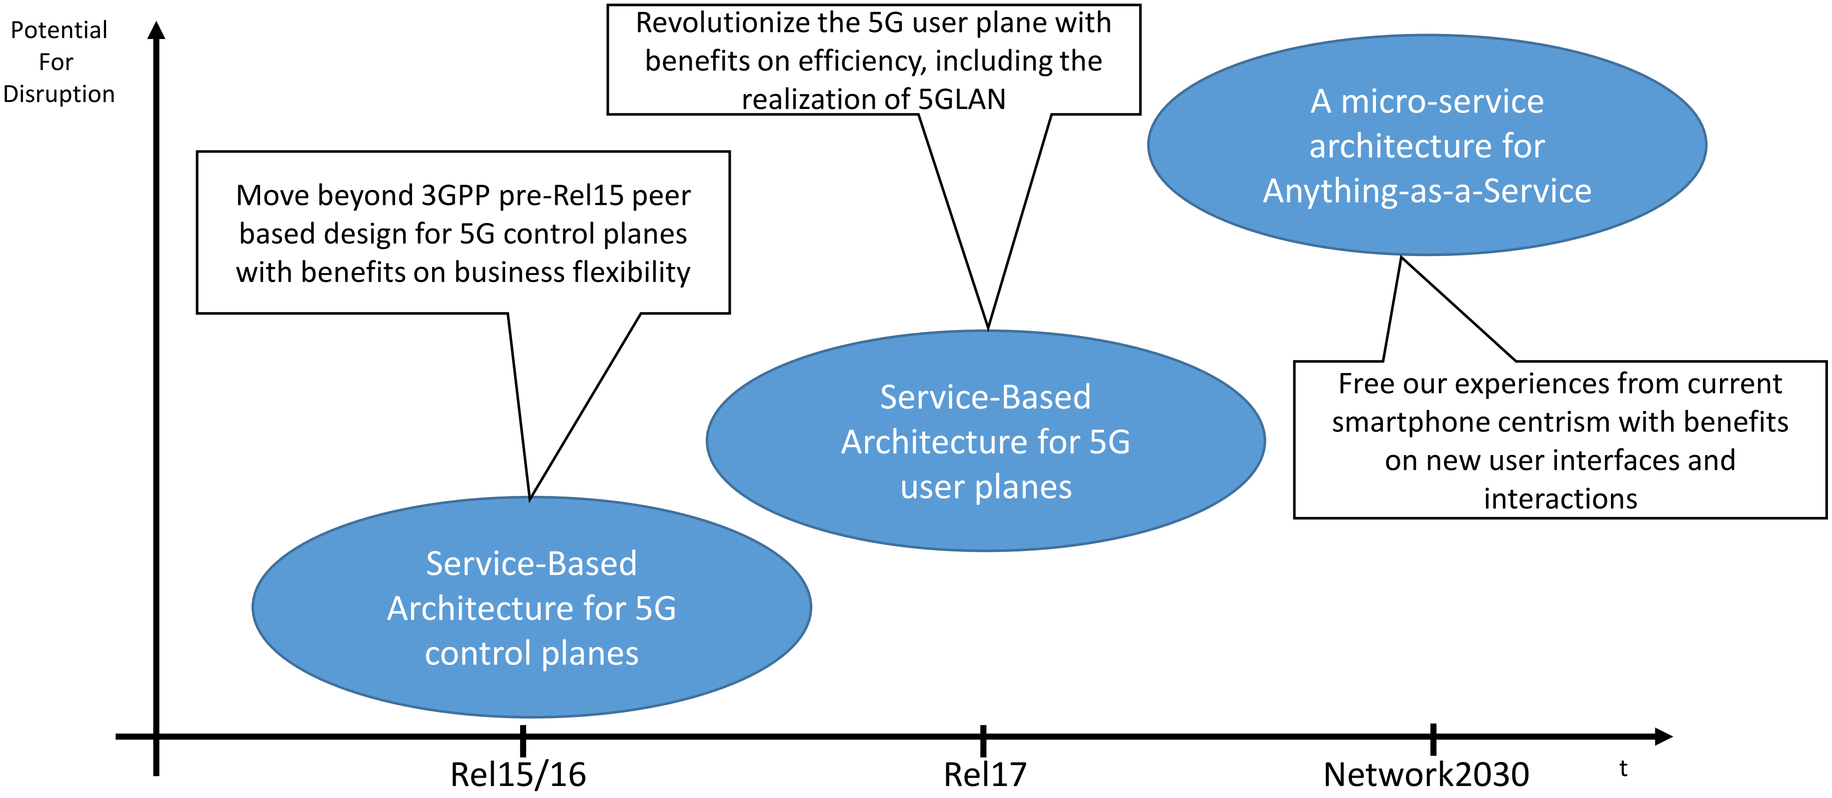

INTERDIGITAL.

A micro-service design for 5G (and beyond)

What is SBA in a Nutshell?

Service-Based Architecture for 5G means integrating mobile networks into the existing large-scale cloud-native Internet service architecture, i.e. apply the principles that made major Internet players successful in the design of the cellular (and fixed) sub-system

-> real winners will be **operators** with increased flexibility in going after new business models with a single cloud-native architecture

-> possibility for **new entrants** in the SW-driven services market!

-> possibility for **new applications** beyond pure client-based ones, e.g., true mobile edge computing, distributed immersive experiences, …

-> possibility for **new devices**, ending the reign of the smartphone!

What are the Features of an SBA?

- Interpret anything as a named service
  - Network function becomes a (named) control/user plane service
- Decompose services
  - A control/user plane service can consist of several (sub-)services, composed from several vendors
- Route service requests to specific service instance
  - There can be several instances in existence, thanks to virtualization
- Decouple service invocation from delivery
  - There are many ways to get bits across the wire

Efficient service routing is key to realizing the SBA micro-service vision!

System Architecture for efficient Service Routing

Main Idea

Directly map individual IP protocols onto name-based routing, integrated into emerging and existing Layer 2 transport networks

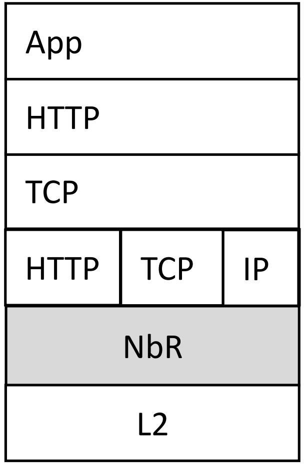

- **Backward-compatible for protocols**, i.e., works with any IP-based protocol
- **Backward-compatible** for any IP-based device, i.e., works with current IP-based devices
- **Forward-compatible** by interpreting any IP-based protocol as a **name-based exchange** (over L2)
- Reliably secure any such exchange via a novel **name-based routing** over an efficient Layer2 **path-based forwarding**

System Overview

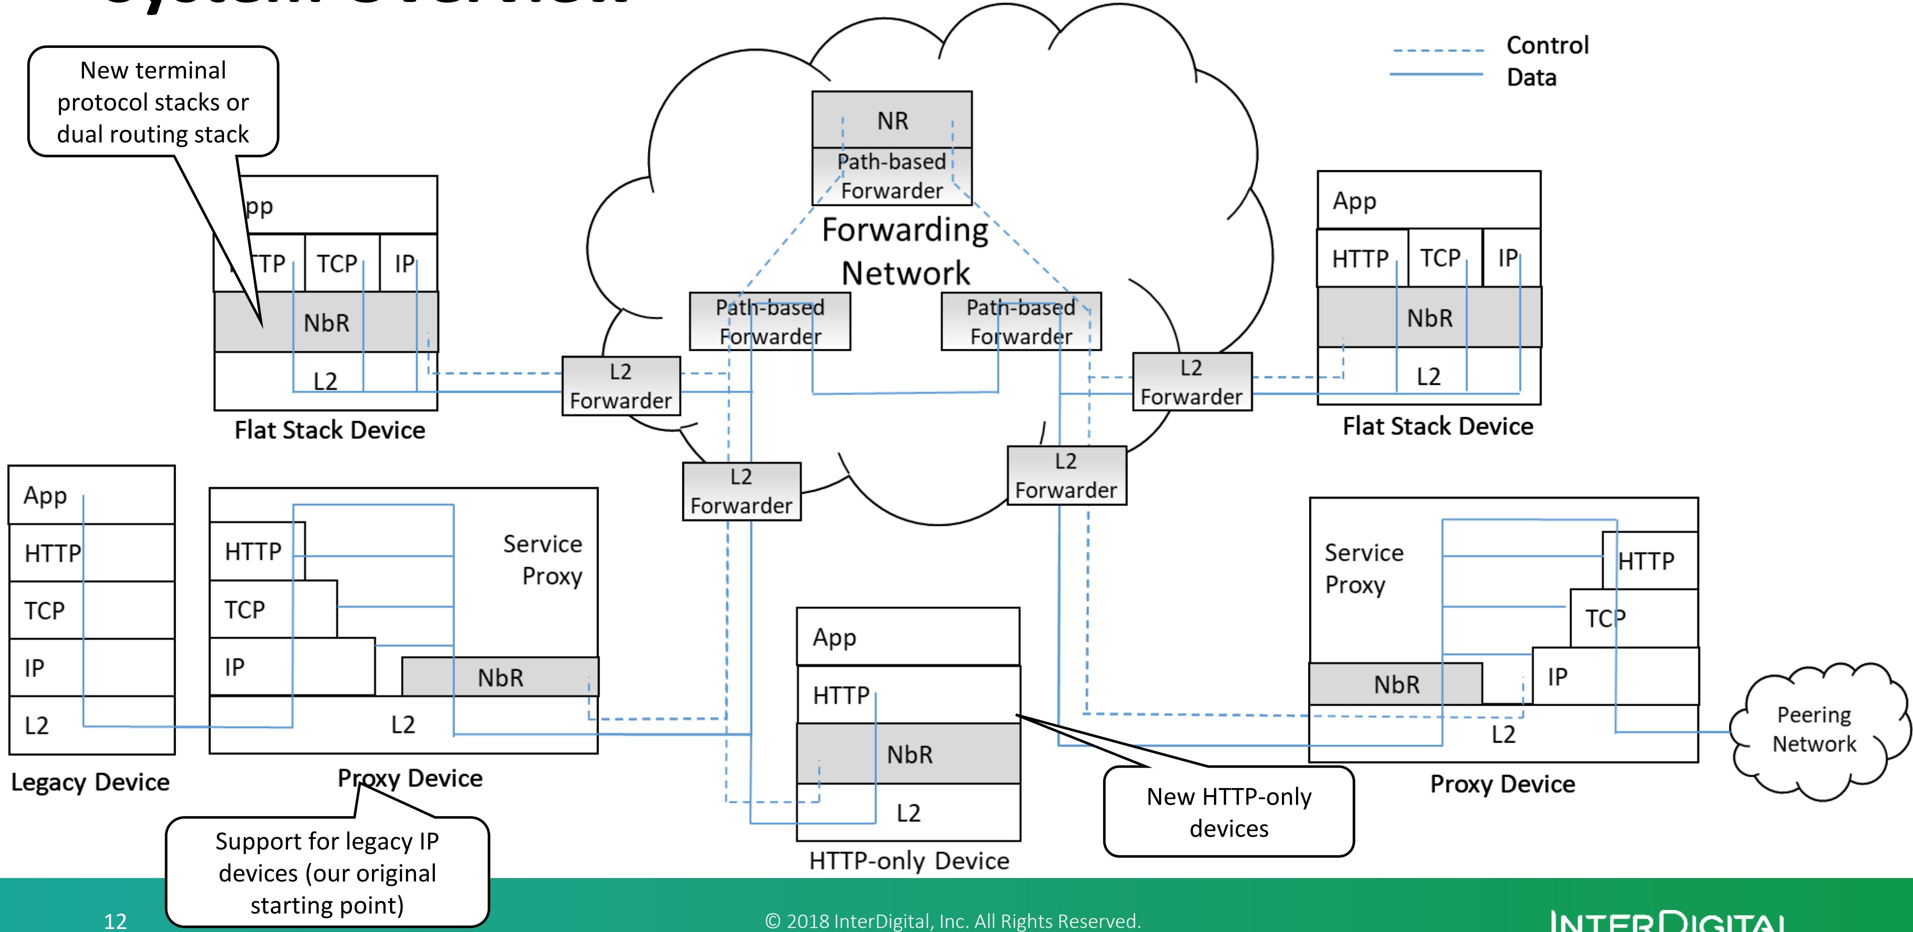

Path Forwarding through Bitfields

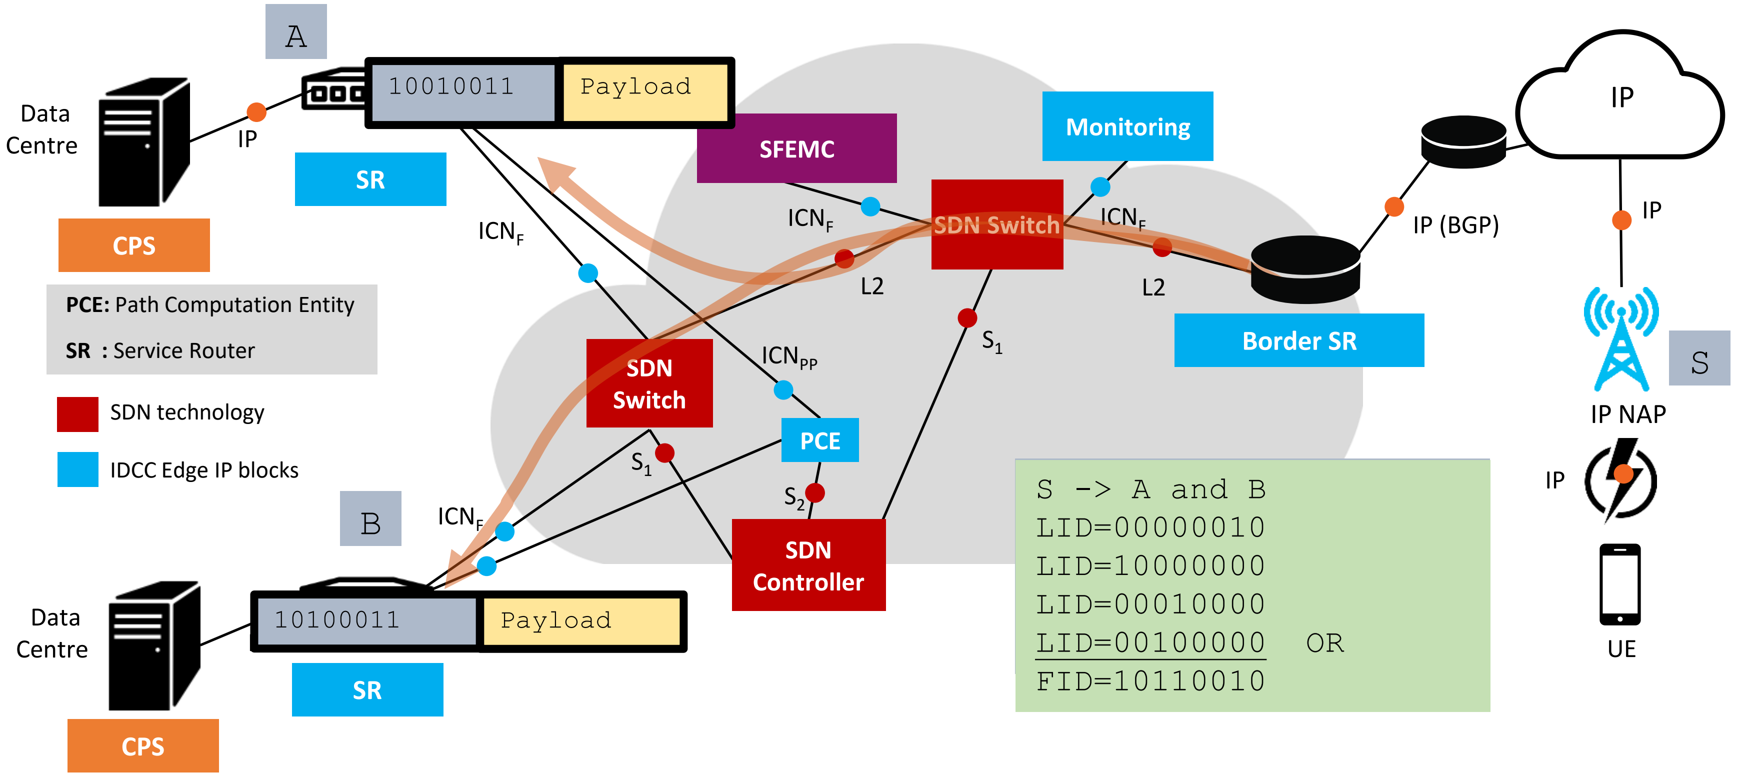

Advantages of Path-based Forwarding

- Only proactive insertion required
  - Rule changes/additions only when inventory changes!
- Can be deployed with **low TCAM requirements** in SDN environments
  - TCAM sizes are important
  - in practice quite limited (thousands of entries)
  - TCAM is the most power-hungry
  - Will co-exist with existing protocols – plenty of TCAM left for IP, MPLS, L2switch …
- Solution can provide native multicast **with no additional TCAM entries!**
  - Existing technologies (or proposed solutions) either require high-state churn (IP-multicast) or large amount of state (various MPLS multicast proposals).
- Compatible with SDN, P4, BIER (for overlay multicast networks)

Generic Packet Structure

Src MAC	Dst MAC	pathID	NAME_ID	Payload

- We assume Ethernet-level Layer2 abstraction, i.e. any device is addressed in a single LAN -> integrates with vision of 5GLAN as a cellular-based distributed LAN!
- `pathID` is a static bitfield of given size (use IPv6 for SDN environments)
- `NAME_ID` depends on IP-level protocol being realized
  - Subject to standardization
- `Payload` is the raw IP protocol level packet, e.g., the HTTP request (but without any TCP/IP wrapping)

NbR Generic API

Simple API for name-based transactions with natively supported return result association

- `conn = send(name, payload)`
  - E.g., send(FQDN, HTTP request)
- `send(conn,payload)`
  - E.g., subsequent HTTP requests
- `conn=receive(name, &payload)`
  - E.g., receive(FQDN, &payload)
- `receive(conn, payload)`
  - E.g., for responses send to initial FQDN-based request
- Note that conn can be the path identifier or combination of (MAC, pathID)

Realizing Name-based Ad-hoc Multicast
Example: HTTP

(*) Note that multicast for conn = (MAC, pathID) is a combination of bitwise OR over all pathID and setting MAC to Ethernet broadcast/multicast! The requestID is then used for L2-level multicast of responses

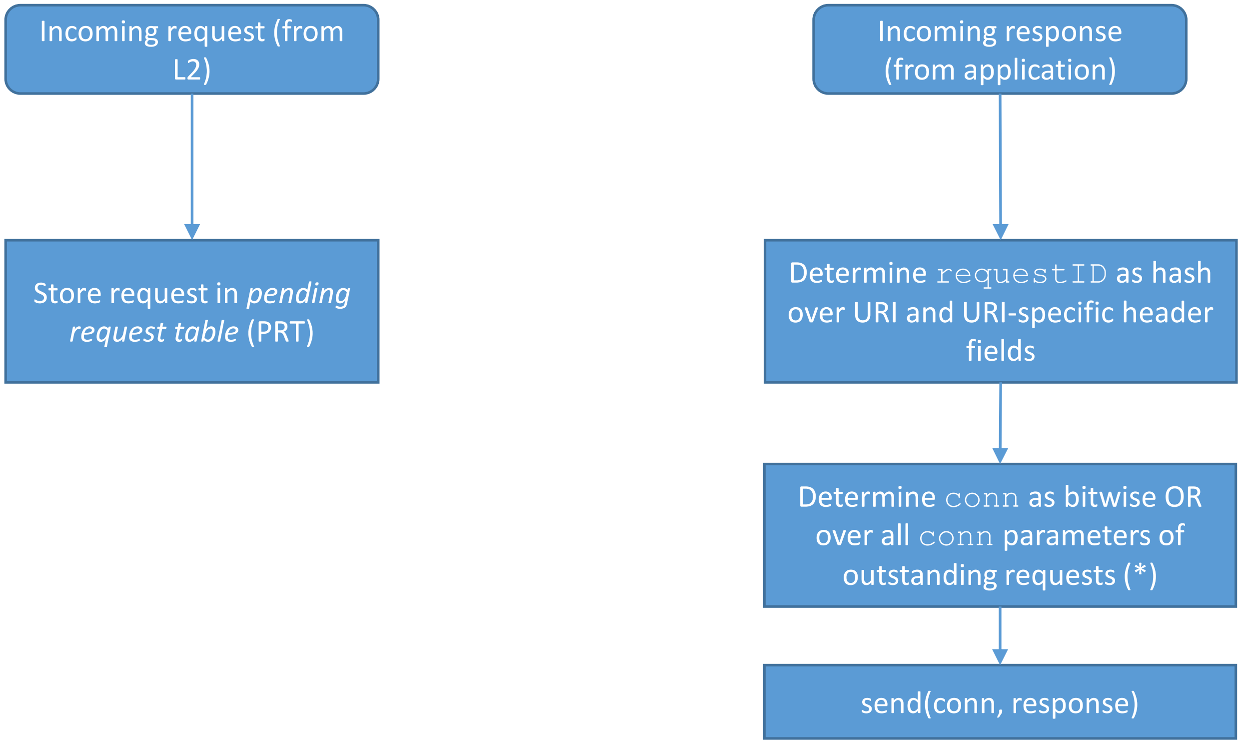

Flow Management

Mapping onto Name-based Transactions

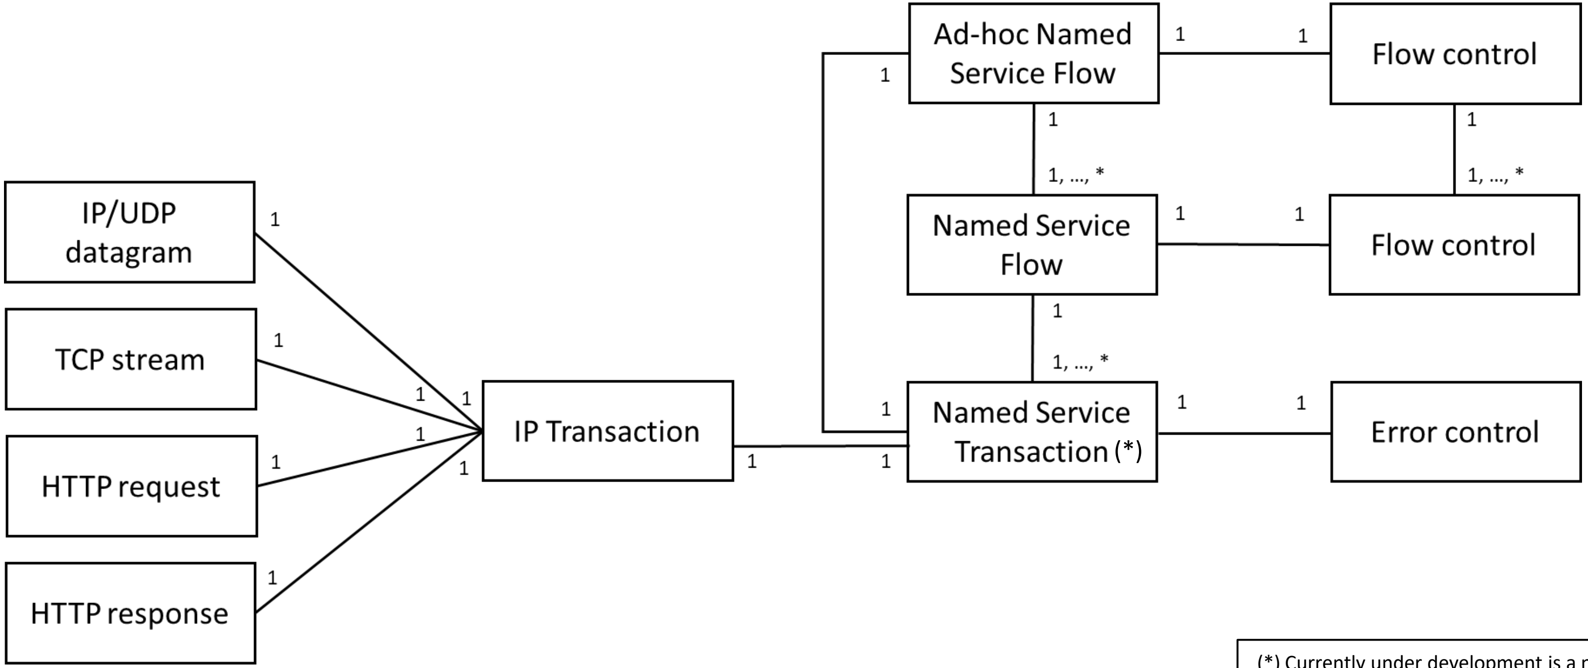

(*) Currently under development is a named service transaction protocol, supporting ad-hoc multicast based on random linear network coding to minimize receiver feedback

Evaluation Plans

Identifying the opportunities to test against

InterDigital

Potential for Layer 2 Multicast

- Utilize simulation driven approach with popularity models for HLS video transmission and operator-scale topologies
- Complement with insights from use case driven insights with site-local deployments and multicast down to the radio links

- NOTE: we have shown multicast gain as early as MWC 2016 with scenarios for
  - Statistically quasi simultaneous OTT video viewing
  - VR videos where multicast is achieved when viewing angles align
  - VR videos where time alignment is part of the use case (e.g., VR tourish guide) – to be shown in Jan 2019 in UK trial

INTERDIGITAL.

Ensure Flow Fairness

- Use reference topologies (dumbbell) with bottleneck link for intra-service flow fairness as well as cross-traffic scenario for inter-service flow fairness.
  - Also include scenarios such as multi/many-TCP P2P download to show removing impact of such scenarios on user-level fairness

- Use setups with legacy as well as new devices

INTERDIGITAL.

Reduce Flow Setup Latency

Objective: evaluate flow setup latency

- Qualitative evaluation of resource management regimes
  - End-to-edge and edge-to-edge for legacy devices
  - End-to-end for new devices

  Expectation: flapping limited to client-to-edge

- Baseline scenario executing HTTP/1.0 (i.e. no TCP re-use)
  - Vary the NbR flow length proportion compared to end-to-edge

- Extended scenario to compare against HTTP-over-QUIC

INTERDIGITAL.

Deployment Insights

INTERDIGITAL.

InterDigital SBA Platform

Realize our Vision, Aligned to Emerging Standards

5G compliant service delivery platform that has proven MEC services & capabilities can be delivered in minutes or less over managed cloud-native operator infrastructures

- **We have also showcased** an array of capabilities is possible through this technological approach:
    1. Linear cost increases of HTTP based streaming can be capped through easily enabled L2 multicast method
    2. E2E latency can be reduced significantly by dynamic end point selected nearest to end users in less than 20ms
    3. Recovery from service & network failures can be reduced to <1sec compared to DNS-based failovers in minutes
    4. Device batter performance can be increased 50% by offloading device functions in real time to edge resources
- **Experimental** solution, deployed in 5G UK test bed in Bristol & Bath (in UK) by end of September 2018
- **Trials** planned in Bristol, Bath and Barcelona
- **Showcased** container realizations in multi-DC deployment for CP & UP services with significant performance gains along the above capabilities

InterDigital 3GPP Rel 16 Compliant SBA Development Platform

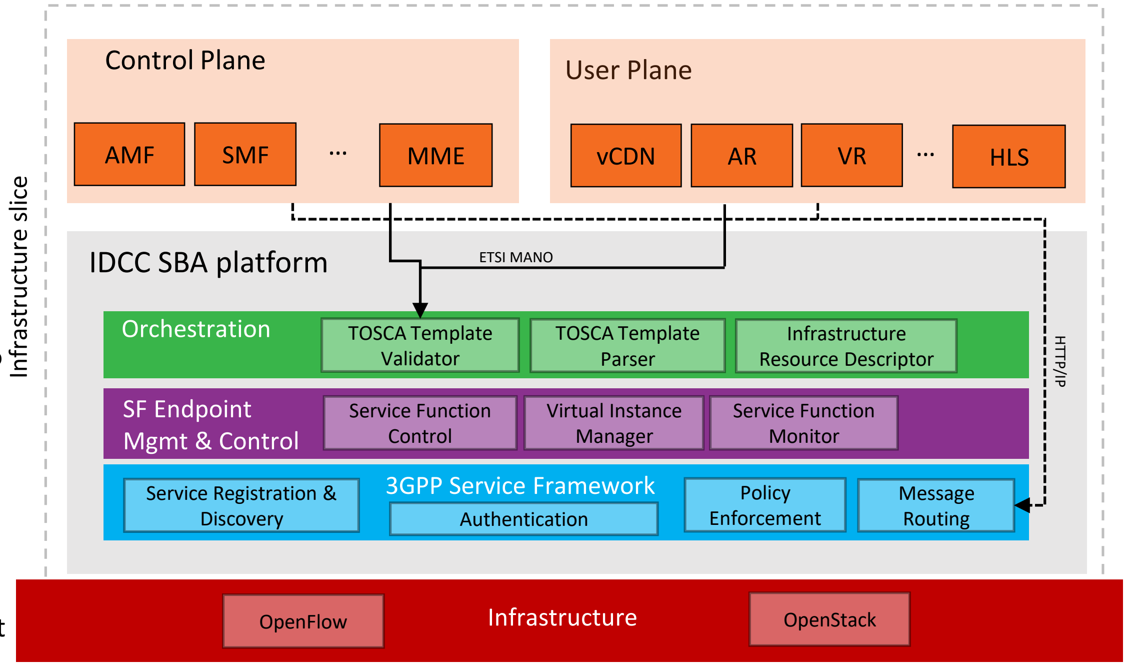

INTERDIGITAL

System Deployment

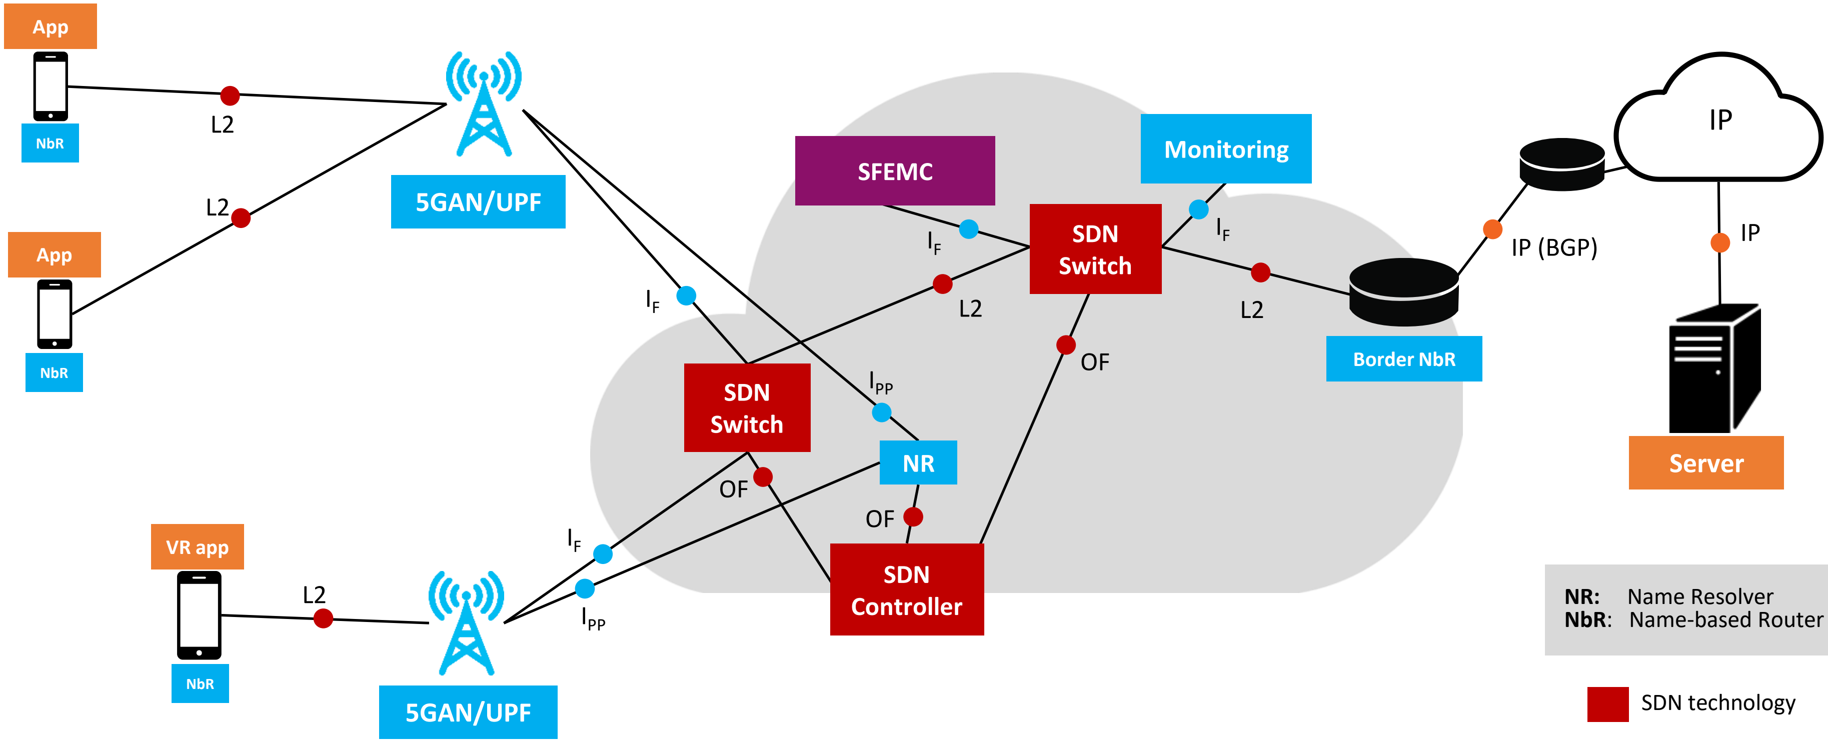

NR: Name Resolver
NbR: Name-based Router

SDN technology

IDCC IP blocks

INTERDIGITAL.

Android Realization as User Application

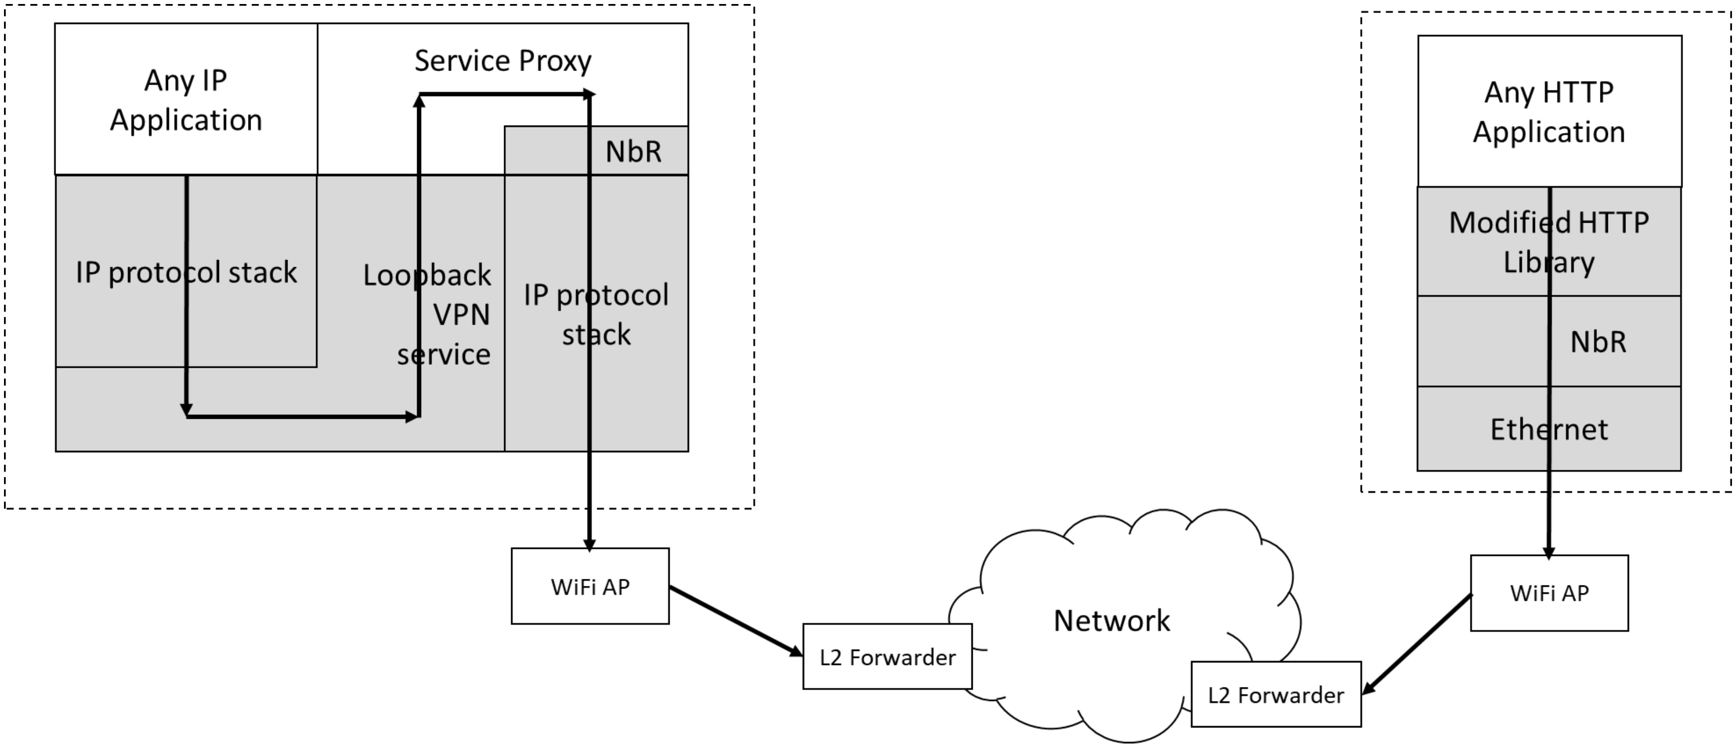

INTERDIGITAL.

Android Realization as System Application

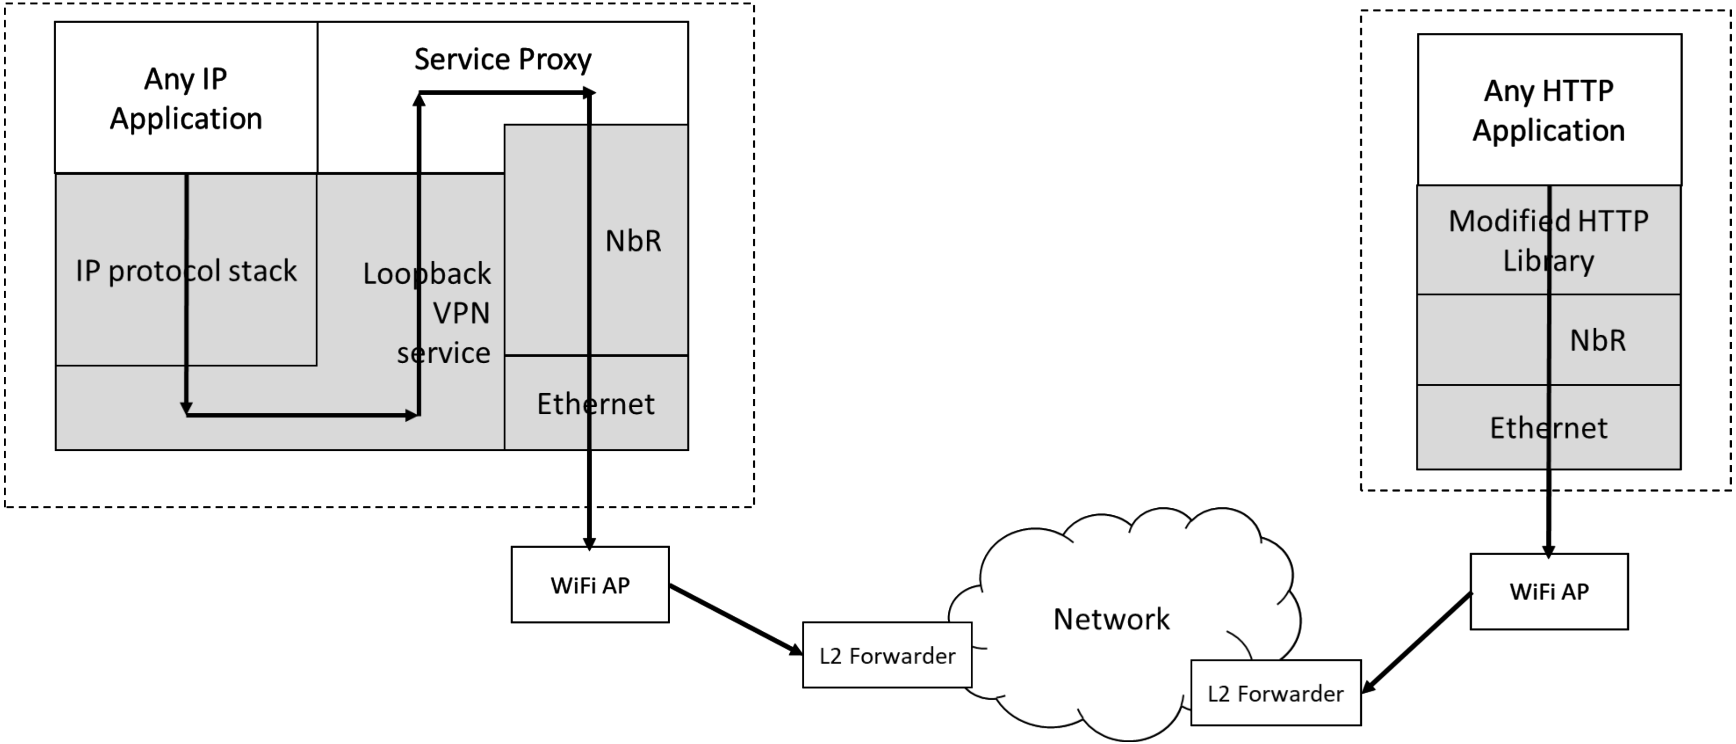

Validation through Urban Scale Trials & Experiments

- Validate platform capabilities by trials conducted by ecosystem partners
  - 5 operator infrastructures
  - 25+ customer trials
- New media formats (AR, VR, 360) and distribution channels
- Engagement with media service providers, content providers, infrastructure operators and beyond
- Trials will be conducted in 3 waves from January 2019 to June 2020
- Public funding available through H2020 FLAME project

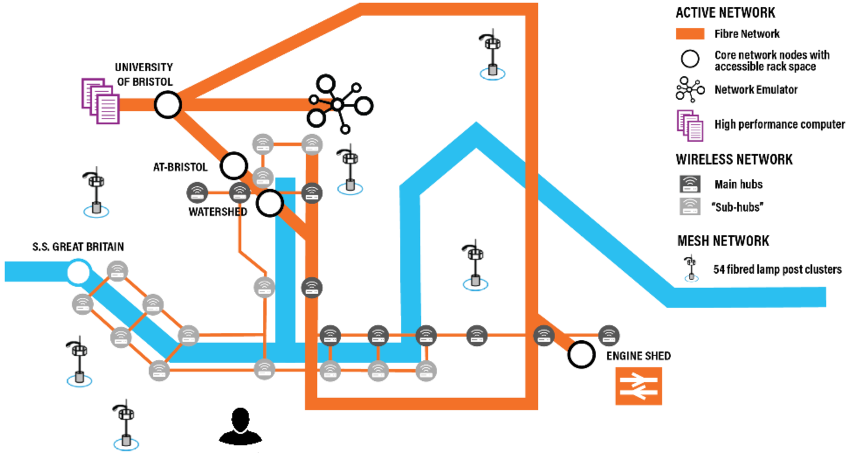

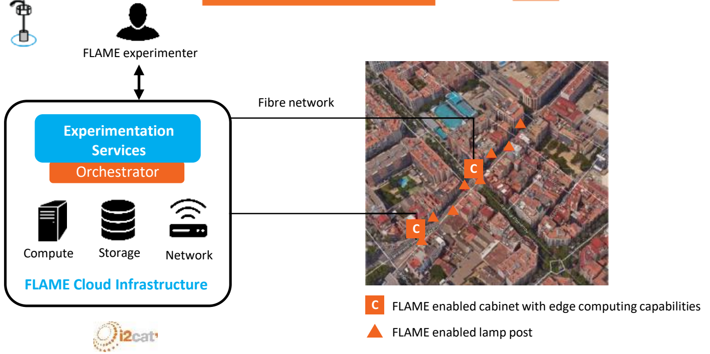

INTERDIGITAL.

Deployed Across Two UK Cities for Large-Scale Trials in 2019

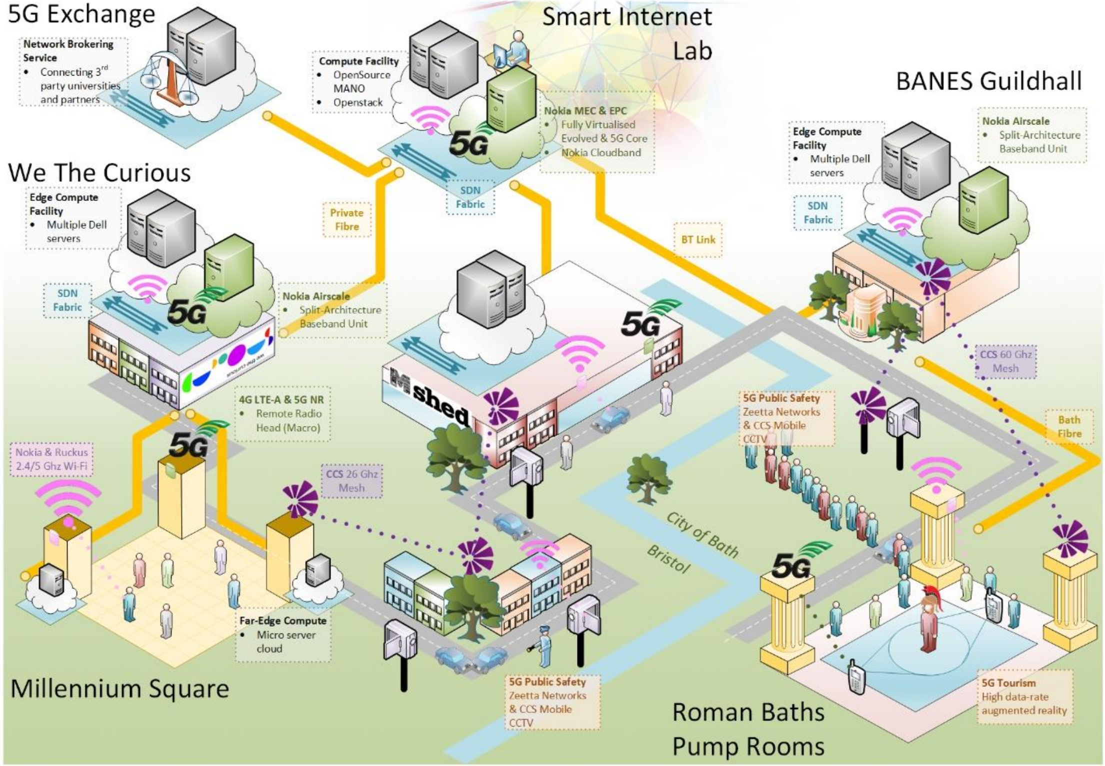

Technology Highlights

- Multi-RAT
  - 5G NR & 5G mmW
  - LTE-A
  - Wi-Fi
- Micro-data centers in
  - Roman Baths & Guildhall
  - We The Curious
  - Smart Internet Lab

  providing MEC services
- Use cases in guided VR tours, AI-assisted image recognition & public safety

INTERDIGITAL.

Working with Partners Across Horizon 2020

	Partners (38)
Platform providers	InterDigital, ATOS
Vendors	Huawei, NEC, Intracom, Thales
Content	Disney Research, VRT Belgium
Operators	Deutsche Telekom, Orange, Telenor Guifi.net, Avanti, Primetel
SMEs	CTVC, Ell.i, Martel Consulting, Ubitech, B-COM, Eurescom, Nextworks
Municipalities	Barcelona, Bristol-is-Open
Academia & Research Institutes	Athens University of Economics & Business, Aalto University, RWTH Aachen, TU Munich, Cambridge University, University of Essex, TU Kaiserslautern, i2CAT, iMinds, King's College London, Fraunhofer Fokus, IT Aveiro, IT Innovation, University of Bristol, ETH Zurich

POINT	EUH2020 funded under grant #643990 http://www.point-h2020.eu
RIFE	EUH2020 funded under grant #644663 http://rife-project.eu
FLAME	EUH2020 funded under grant #731677 http://www.ict-flame.eu
5G CONFIG	Unfunded international research and innovation project http://www.5g-control-plane.eu

13.5MIL EURO
Combined H2020 Funding

Showcased as **ETSI MEC PoC**
Successful **5G network trial in 2017**
Won **CSI Award 2018**

INTERDIGITAL.

From Research to Standards

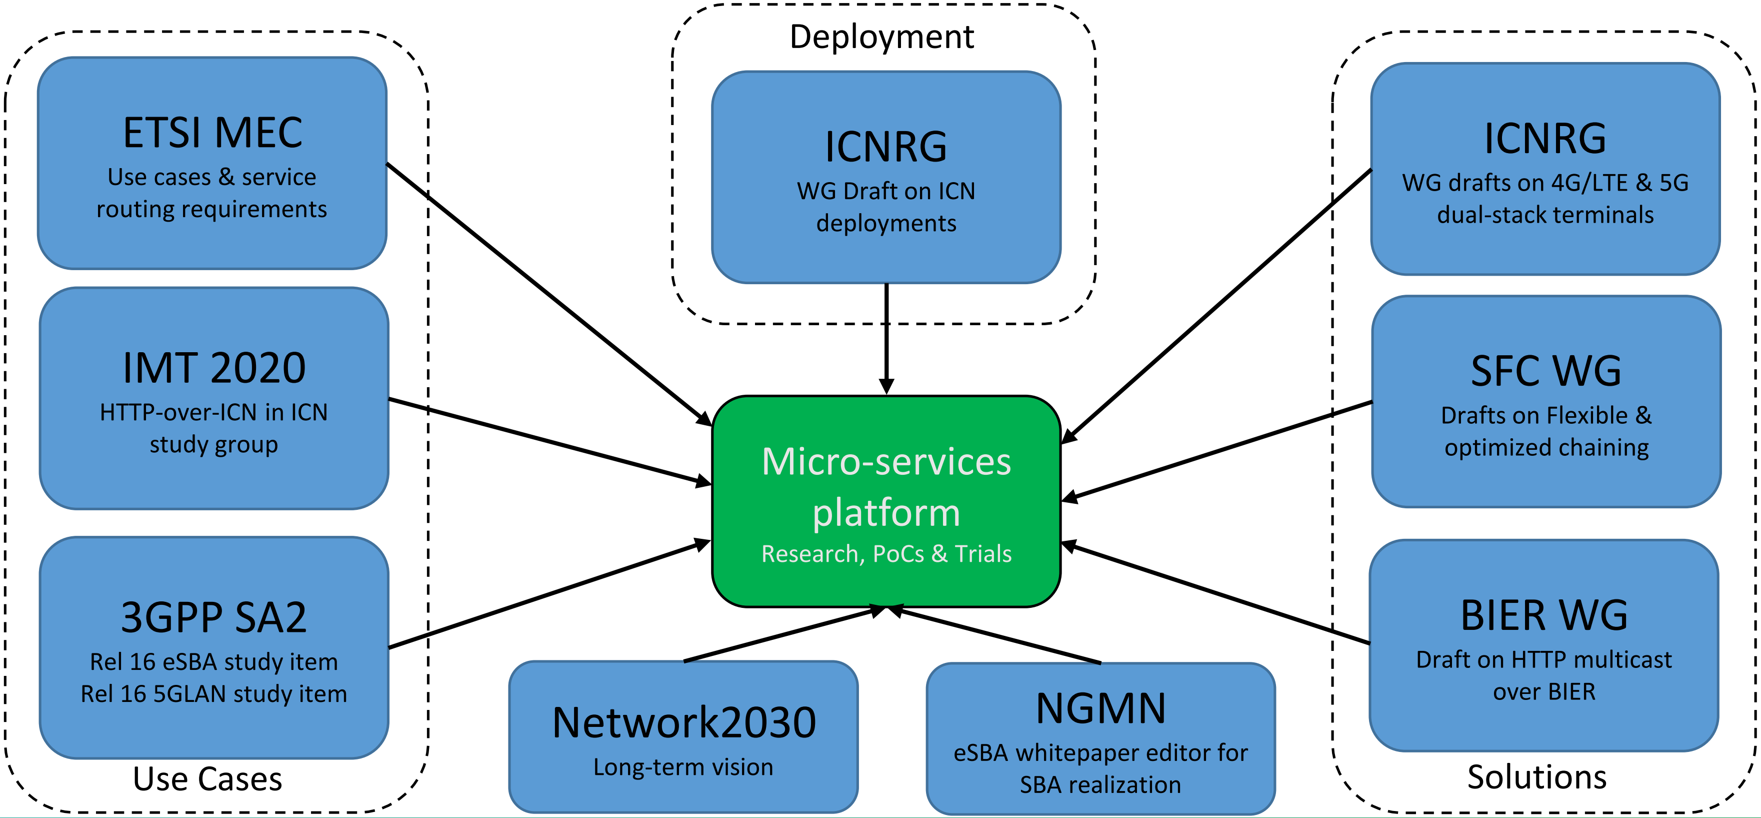

INTERDIGITAL.

Conclusions

- Micro-service vision is all about dynamic service routing, often localized and at the very far edge of the Internet
  - Recognized in 3GPP and ETSI through mobile edge computing and SBA efforts!
- Presented our approach to flexible service routing
  - Name-based routing has long-standing provenance in previous research efforts
  - Prototyping and even early trialing is ongoing
- Still lots of open research, such as (and among others)
  - Chaining services towards an expected overall task/experience
    - Dynamic chaining based on changing constraints -> dynamic programming problem?
  - Move from L2/L3 chaining to name-based chaining
  - Ensure correctness of resource scheduling

Much of it is about redefining the line between management and control!

INTERDIGITAL.