Microsoft

Lessons I learned in leveraging AI+ML for 5G/6G systems

Sharad Agarwal

ANRW 2024 Keynote

three lessons

- **L1**: leverage cloud scale to overcome limitations of deployed network protocols & use AI+ML to manage that massive scale
- **L2**: custom learning algorithms can take you far, but are harder to deploy than off-the-shelf AI+ML algorithms
- **L3**: reduce risk from AI hallucinations with careful system design

outline

- cloudification of telecom infra
  - why is it really happening?
  - what are the interesting challenges?

- improving the performance & reliability of cloudified telecom infra
  - high throughput: TIPSY in ACM SIGCOMM 2022
  - low latency: PAINTER in ACM SIGCOMM 2023
  - high reliability: LLexus in ACM SIGOPS OSR 2024

what is happening in the telecom industry?

a confluence of trends behind 5G/6G that is driving a renewed push for enabling new revenue & reducing expenditure

new radios enable new capabilities & revenue models

Home > Press releases > Ericsson and du reach 16.7 Gbps download speed on 5G Standalone with 10 aggregated carriers

Ericsson and du reach 16.7 Gbps download speed on 5G Standalone v 10 aggregated carriers

Available in English 日本語 简体中文 繁體中文 العربية

- Ericsson and du tested 10 carriers per sector on a live 5G network, achieving up to 16.7 G aggregated downlink speed.
- Implementation is based on 5G standalone (SA) New Radio-Dual Connectivity (NR-DC) c aggregation technologies
- The trial opens the doors for differentiated Fixed Wireless Access experiences, and new o for AR/VR and cloud gaming in the United Arab Emirates.

PRESS RELEASE | MAR 25, 2024

5G Advanced Flexible Production Line Features High Reliability and Ultra-Low Latency

November 14, 2023

At Great Wall Motor's factory in Baoding, Hebei, China 5G-Advanced equipment is

evolution of network functions in telecom infra (link)

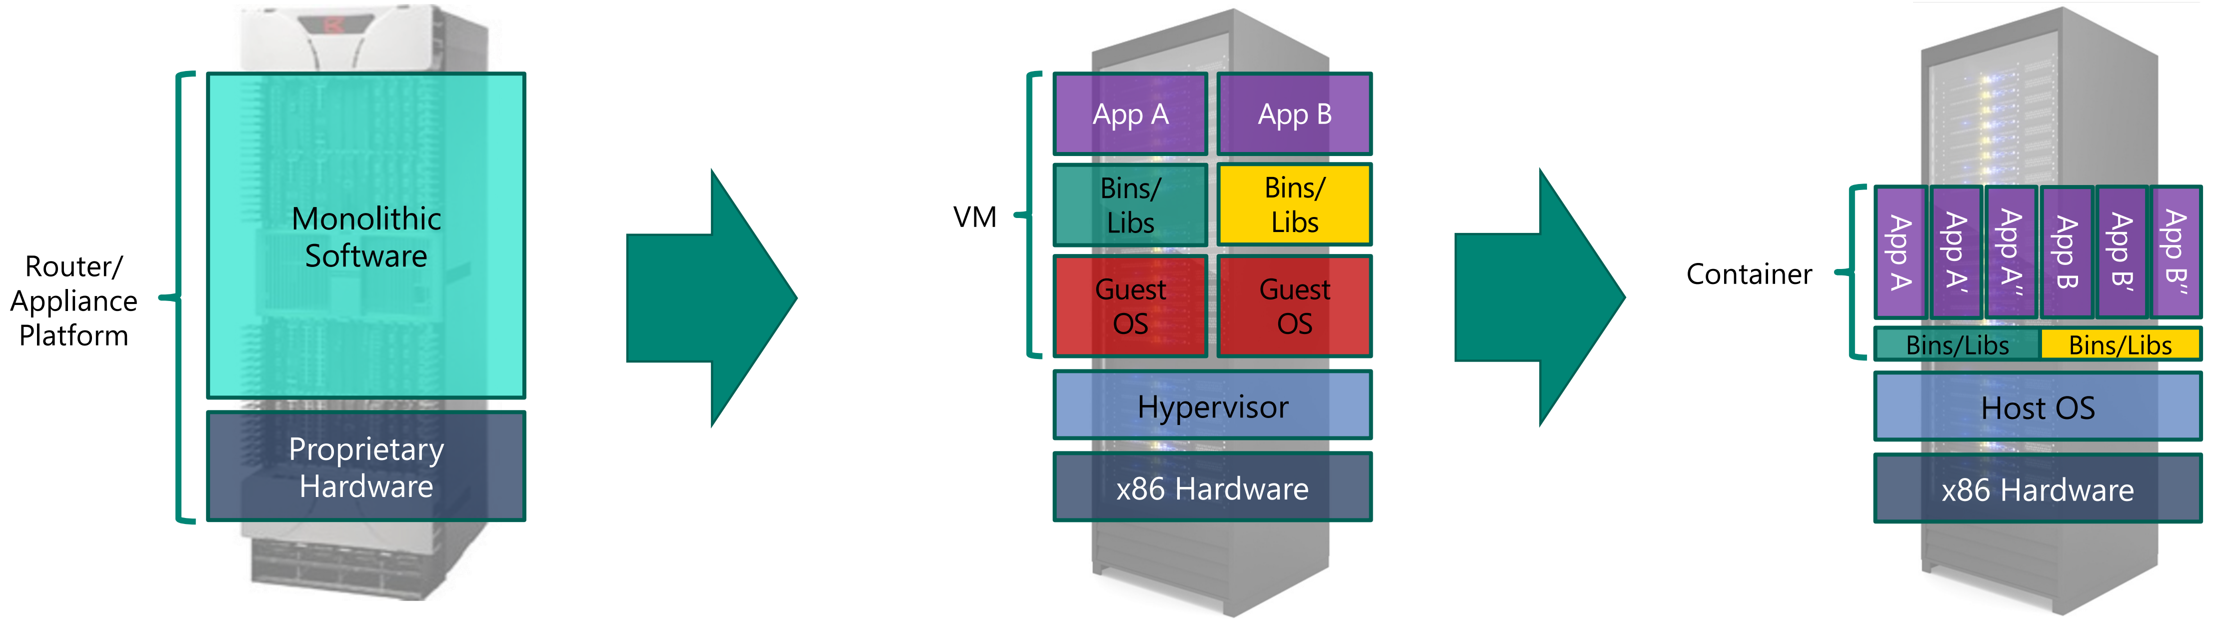

physical network functions

→2012

virtual network functions

2012 → 2018

"Virtualization 1.0"

cloud native functions

2019→

"Virtualization 2.0"

(Tim Irvin (MSFT))

motivation for NF virtualization – ETSI 2012 (link)

- reduced equipment costs
- reduced power consumption
- economies of scale of IT industry
- time-to-market from HW dev to SW dev
- multi-tenancy, scale up/down
- open ecosystems

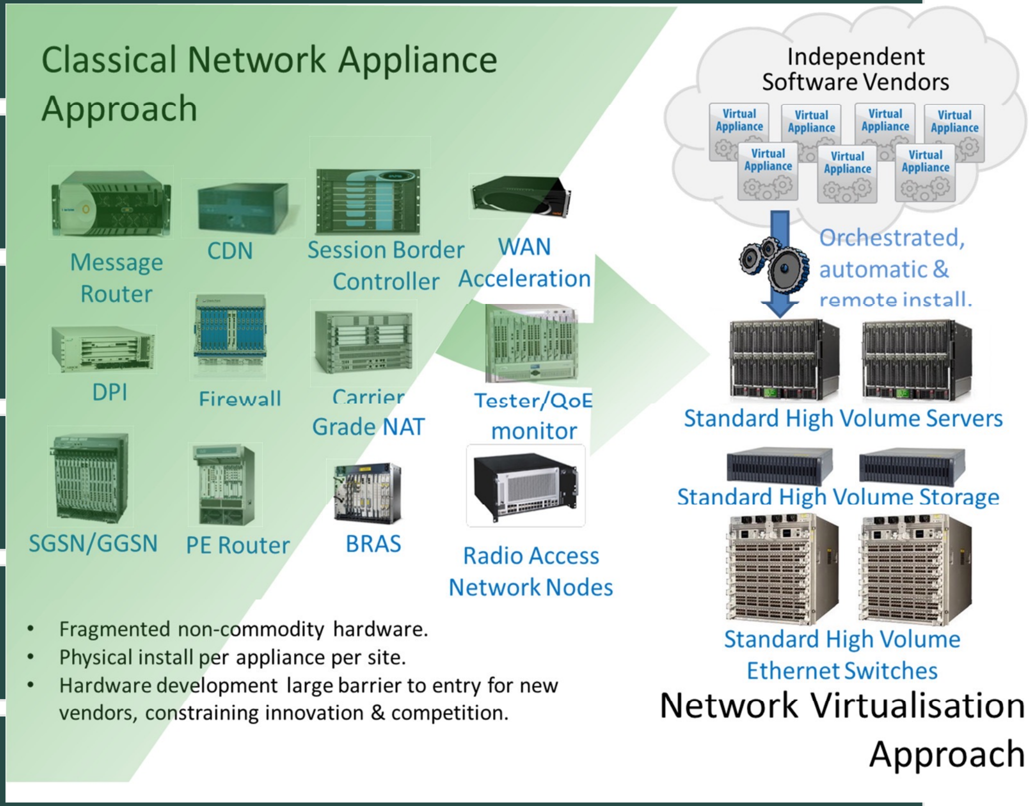

CAPEX heavy → OPEX efficient

closed, locked-in solutions → open-source innovation

summary: what is happening in the telecom industry?

- innovations have been standardized & now available
- new radios, frequencies, protocol improvements
  - high bandwidth (10-20Gbps)
  - low latency (~1ms)
  - private deployments for mission critical apps (e.g., factory automation)
- telco infra has converted from HW to SW
  - cloud-based stack across telco edges & hyperscaler DCs
  - leverage economies of scale in commodity compute servers
  - ease of future upgrades (5G NF → 6G NF and faster servers)
  - promise of AI & ML in cloud for analytics

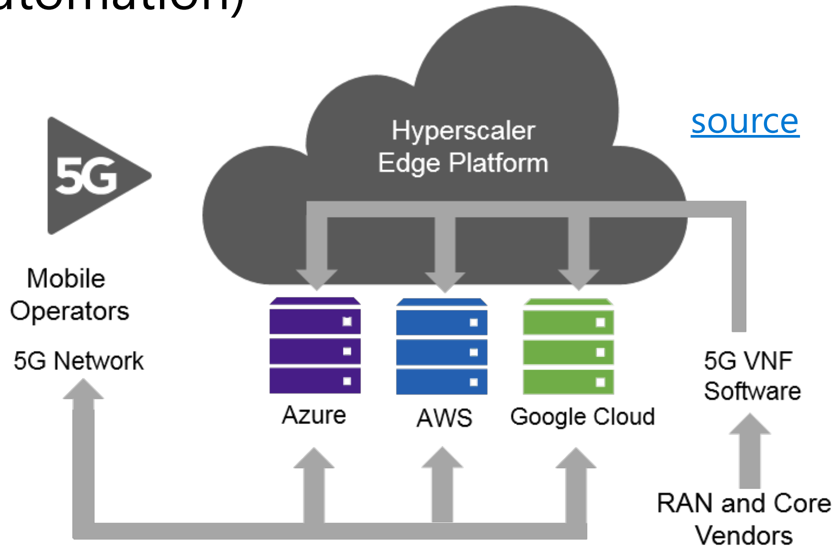

source

what a telecom network on a hyperscaler looks like

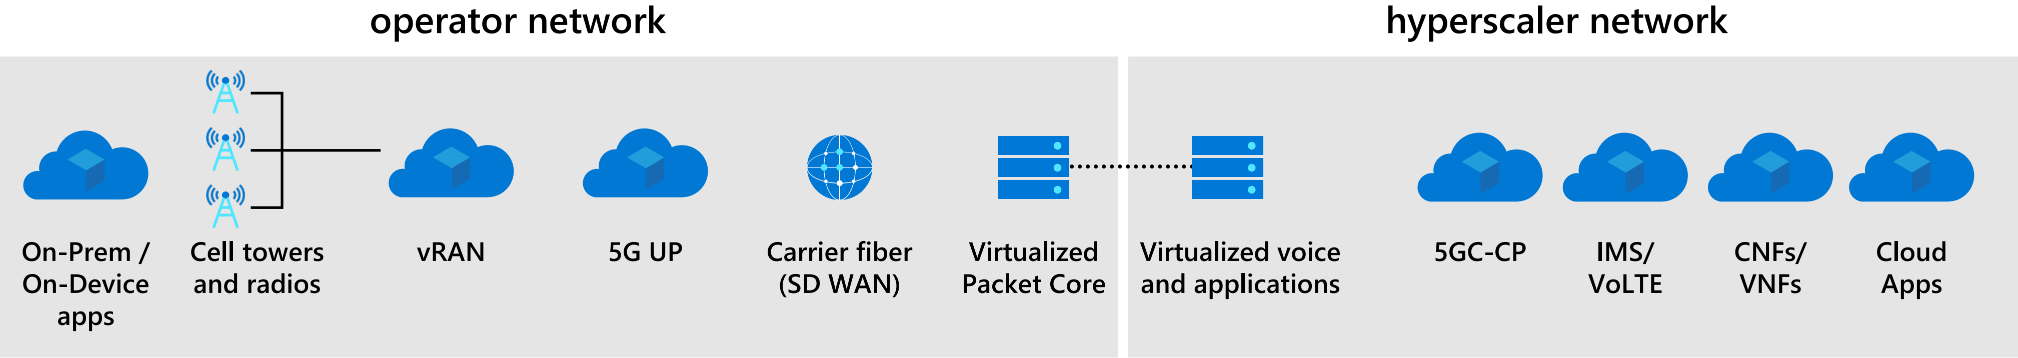

AT&T to run its mobili
on Microsoft's Azure fo
Operators cloud, delive
efficient 5G services at

Microsoft to acquire AT&T's Network Cloud technolog
operators increase competitive advantage through st
service differentiation

NEWS PROVIDED BY
Microsoft Corp. →
Jun 30, 2021, 10:35 ET

DALLAS and REDMOND, Wash., June 30, 2021 /PRNewswire
strategic alliance provides a path for all of AT&T's mobile net

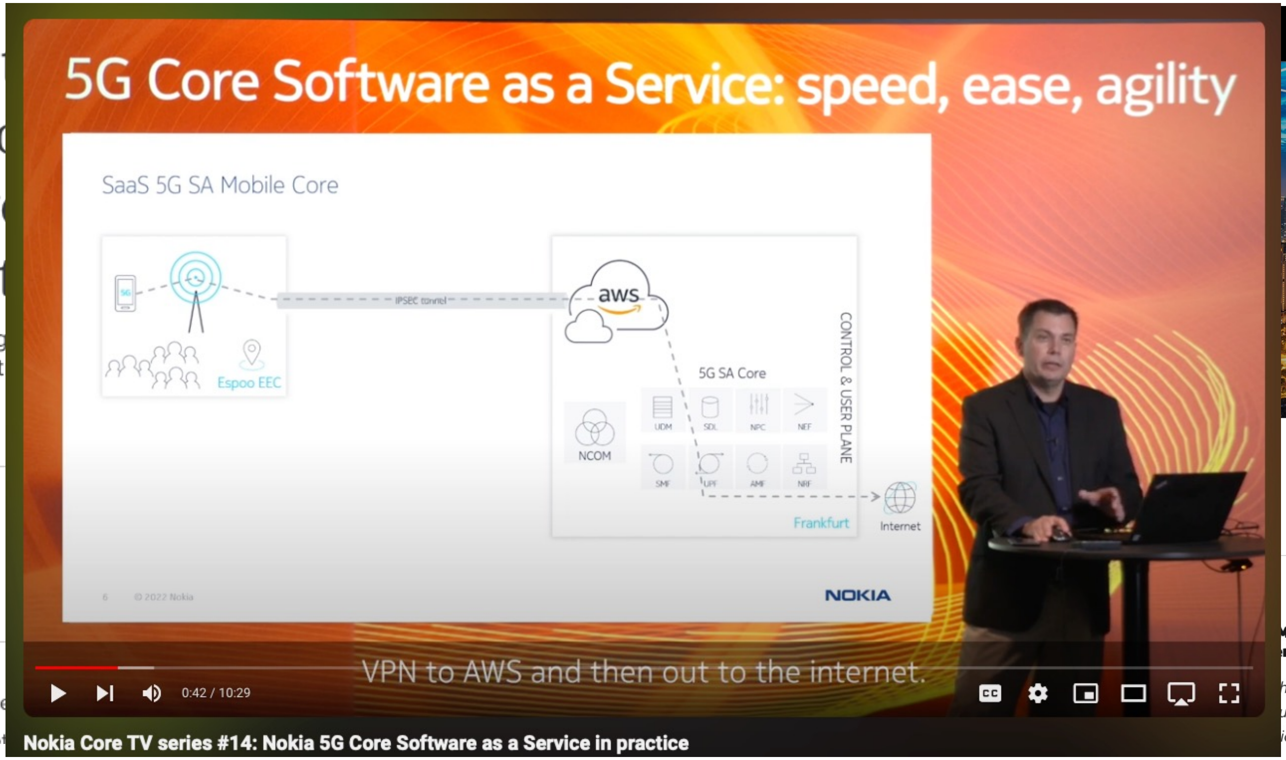

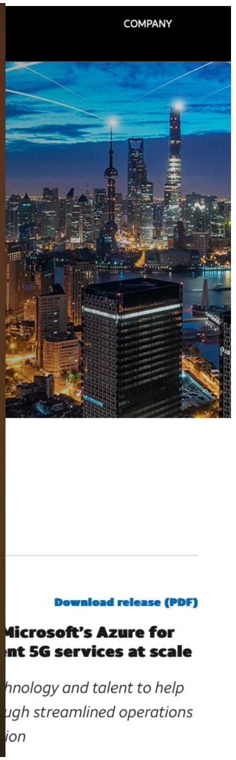

what a telecom network on a hyperscaler looks like

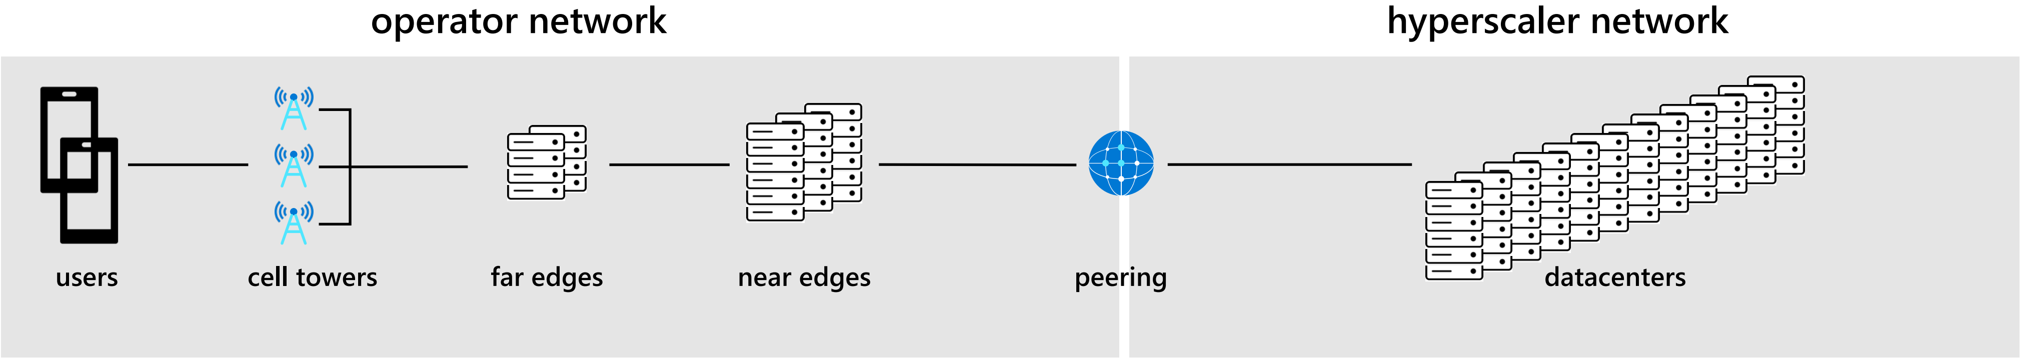

- introduces significant pressure on existing cloud metrics
  - bandwidth: O(10Gbps/user) hitting cloud services
  - latency: O(1ms) over-the-air compared to WAN latencies & routing inefficiency
  - reliability: going from 99.5 enterprise grade → 99.999 carrier grade availability

outline

- ~~cloudification of telecom infra~~
  - ~~why is it really happening?~~
  - ~~what are the interesting challenges?~~

- improving the performance & reliability of cloudified telecom infra
  - high throughput: TIPSY in ACM SIGCOMM 2022
  - low latency: PAINTER in ACM SIGCOMM 2023
  - high reliability: LLexus in ACM SIGOPS OSR 2024

Microsoft

TIPSY: Predicting where traffic will ingress a WAN

ACM SIGCOMM 2022

Michael Markovitch
Sharad Agarwal
Rodrigo Fonseca
Ryan Beckett
Chuanji Zhang
Irena Atov
Somesh Chaturmohta

users with high BW overwhelm cloud peering links

- biggest peering links are 100G or 400G
  - at 10G/user, only 10-40 users can overwhelm a peering link
- many examples from post-mortem analysis
  - e.g., image & video uploads from certain phones overwhelmed 100G ingress in Europe
- core problems
  - small numbers of users cause traffic drops for many users
  - BGP is not capacity & congestion aware
  - no control & visibility over how others select our routes, which makes traffic ingress appear non-deterministic
  - egress control such as Espresso not widely used
- opportunity
  - spare capacity at other peering links

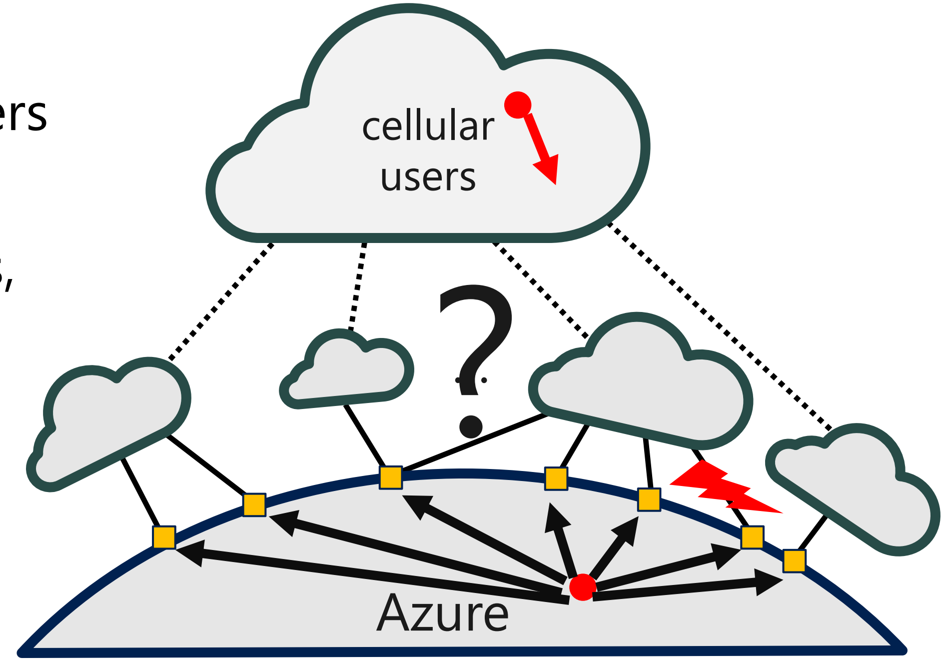

Microsoft global network

60+
Azure regions

300+
Datacenters

4000+
AS peers

190
PoPs

175K
miles lit fiber

TIPSY: shift ingress traffic with predictive withdrawals

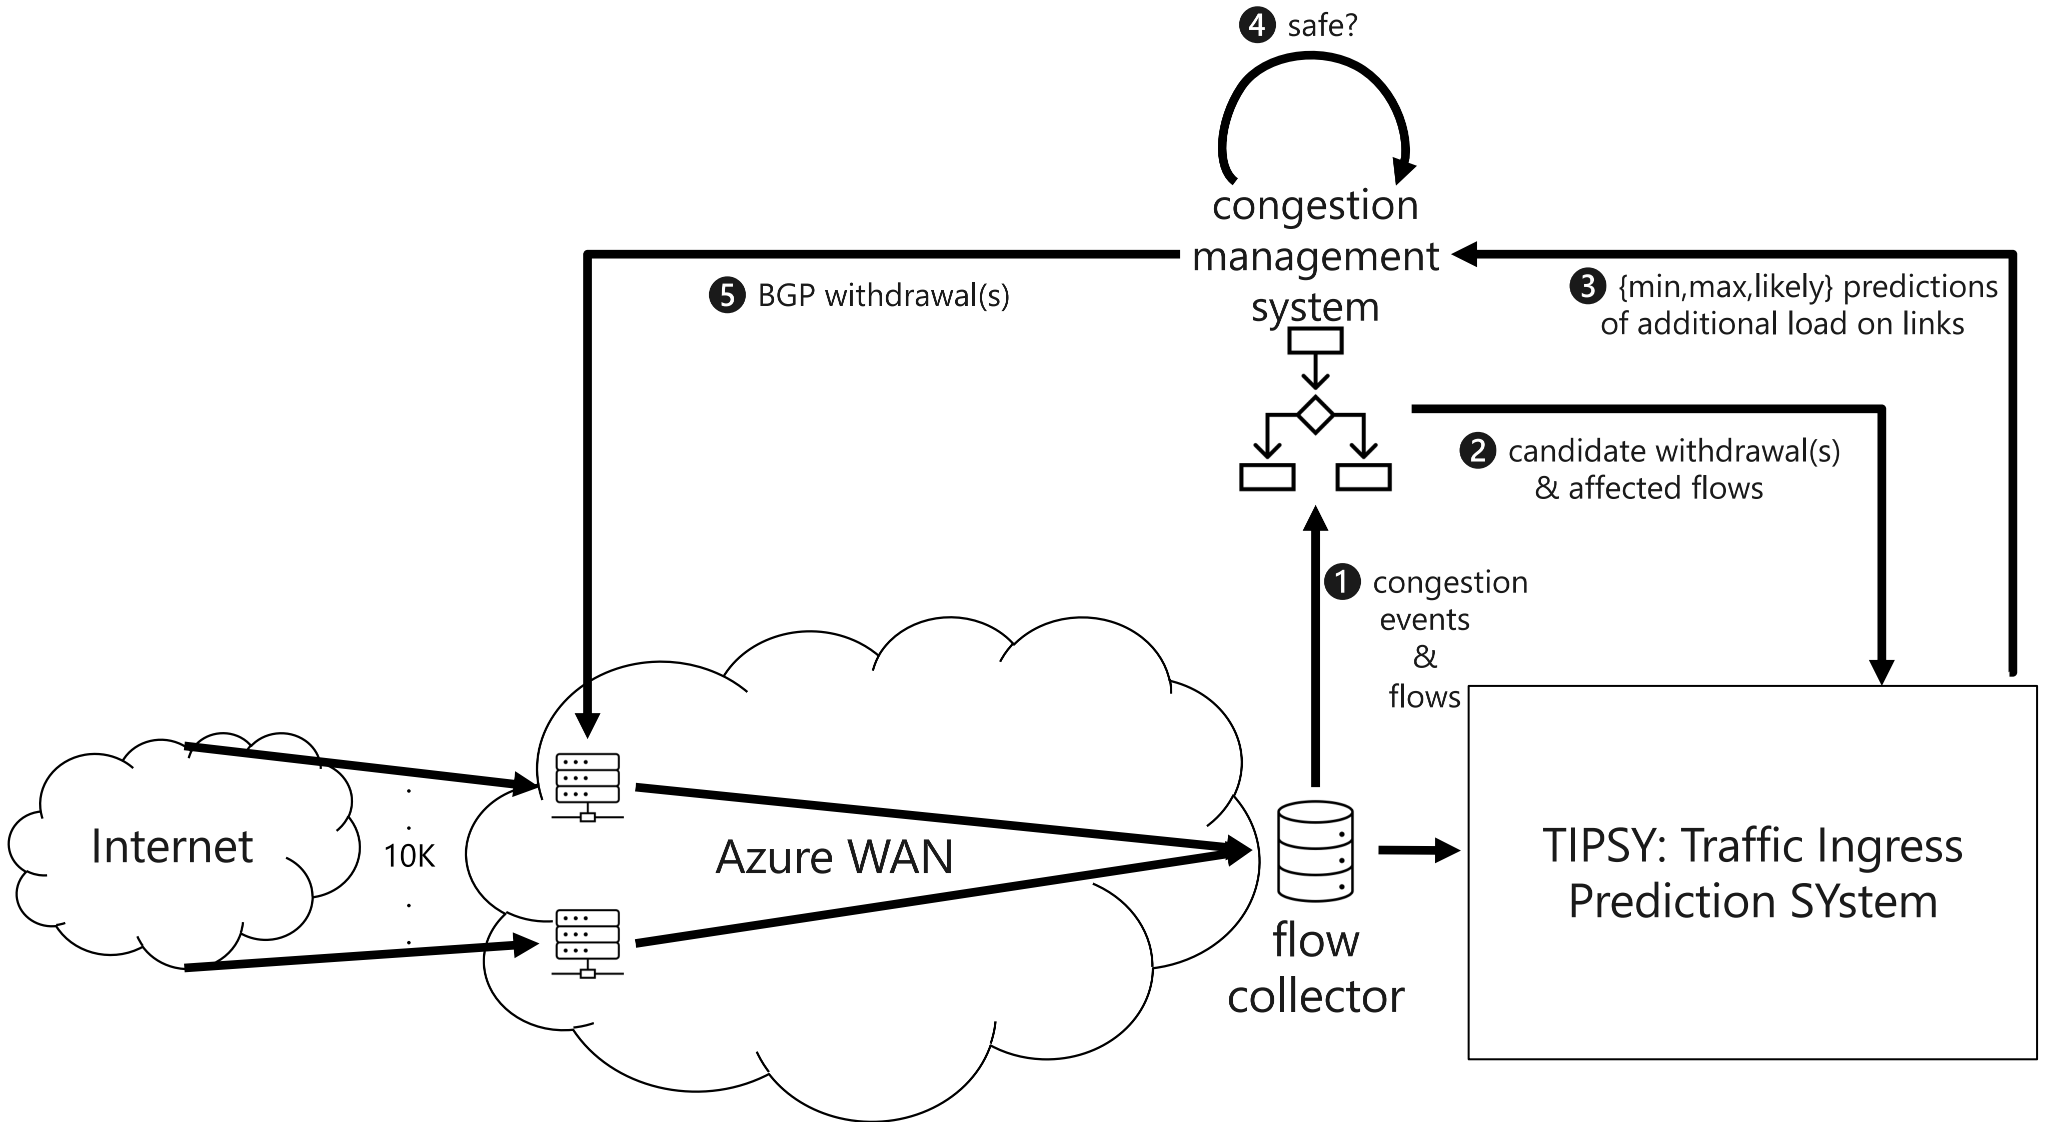

ensemble of statistical classification models

- predictive algorithm
  - given O(10K) peering links for a WAN
  - for a given traffic flow that enters the WAN
  - produce the top k predicted links most likely to receive flow
- input to train model
  - sampled flow data + annotations (source AS & location, destination location & type)
  - list of peering interfaces & metadata
- supervised learning with custom models
  - historical model: weighted average cache of prior traffic
  - geographic model: Haversine formula for distance
- combined into ensemble, with different feature sets

L1: leverage cloud scale to overcome limitations of deployed network protocols & use AI+ML to manage that massive scale

L2: custom learning algorithms can take you far, but are harder to deploy than off-the-shelf AI+ML algorithms

summary: we learn a way out of BW bottlenecks

- the core of the Internet has relied on statistical multiplexing
  - but this assumption is at risk: new radio promises 10-20 Gbps
  - typical peering links on the Internet tend to be 10G/100G, with max of 400G
- we built TIPSY, a learning system for Azure WAN
  - accuracy is 76.4%-97.9% & demonstrated on incidents with mobile users & telco traffic
  - more details in ACM SIGCOMM 2022 paper
- take-aways
  - leveraged cloud scale to avoid having to change BGP & wait eons for deployment
  - but cloud scale necessitated a learning algorithm to predict 10K classes from 1PB of data
  - biggest barrier to adoption was product engineering to own this custom algorithm
  - $¼M/year in compute was not a fundamental barrier

Microsoft

PAINTER: Ingress traffic engineering & routing for enterprise cloud networks

ACM SIGCOMM 2023

Thomas Koch
Shuyue Yu
Sharad Agarwal
Ethan Katz-Bassett
Ryan Beckett

private 5G/6G for e

pwc

PwC Global › Industries and sectors › Techno

Smart Manu

The modern factory is already
machines and robots are equ
connected to high-powered a
performance, manage produc
orchestrate all the activities o

By eliminating the need for w
high-speed manufacturing en
flexibility. And the sheer richn
have the capacity to maintain
either wired or previous wirele
just about anything.

verizon business

Factory of the future: How 5G and MEC can help transform factory ope

Factory of the future: How 5 and MEC can transform fac operations

Author: Keith Shaw

Global supply chain disruptions—whether caused by hea
new technology innovations to transform their operations
benefit from deploying a combination of private 5G and M

The smart factory of the fu

Many manufacturers already deploy automation within th
Industry 4.0 includes a wealth of new innovations to help

5G Factory of the Future

Who We Are | Our Vision | Industrial Applications | News | Contact

5G for Manufacturing: How companies are deploying 5G in their factories

News | 5G for Manufacturing: How companies are deploying ...

Three in four manufacturers intend to adopt private 5G networks by 2024.

So says an international study by network management company Accedian. By comparison, 92% of manufacturing facilities today use Wi-Fi for local networks.

Explaining why interest in private 5G is at an all-time high, Accedian's Jay Stewart says that manufacturers clearly understand the impact it can have on their businesses. "Private 5G supports a wide variety of existing manufacturing applications while enabling new ones that aren't practical with Wi-Fi, Ethernet, and other technologies," he adds.

The research, which canvassed the connectivity ambitions of manufacturers in the UK, Germany, the US and Japan, identified five key factors influencing 5G deployment model decisions:

- 63% – Network Security
- 49% – Network Performance
- 49% – Speed / Simplicity of Deployment
- 45% – Application Performance
- 43% – Data Sovereignty / Privacy

actual example of customer complaint

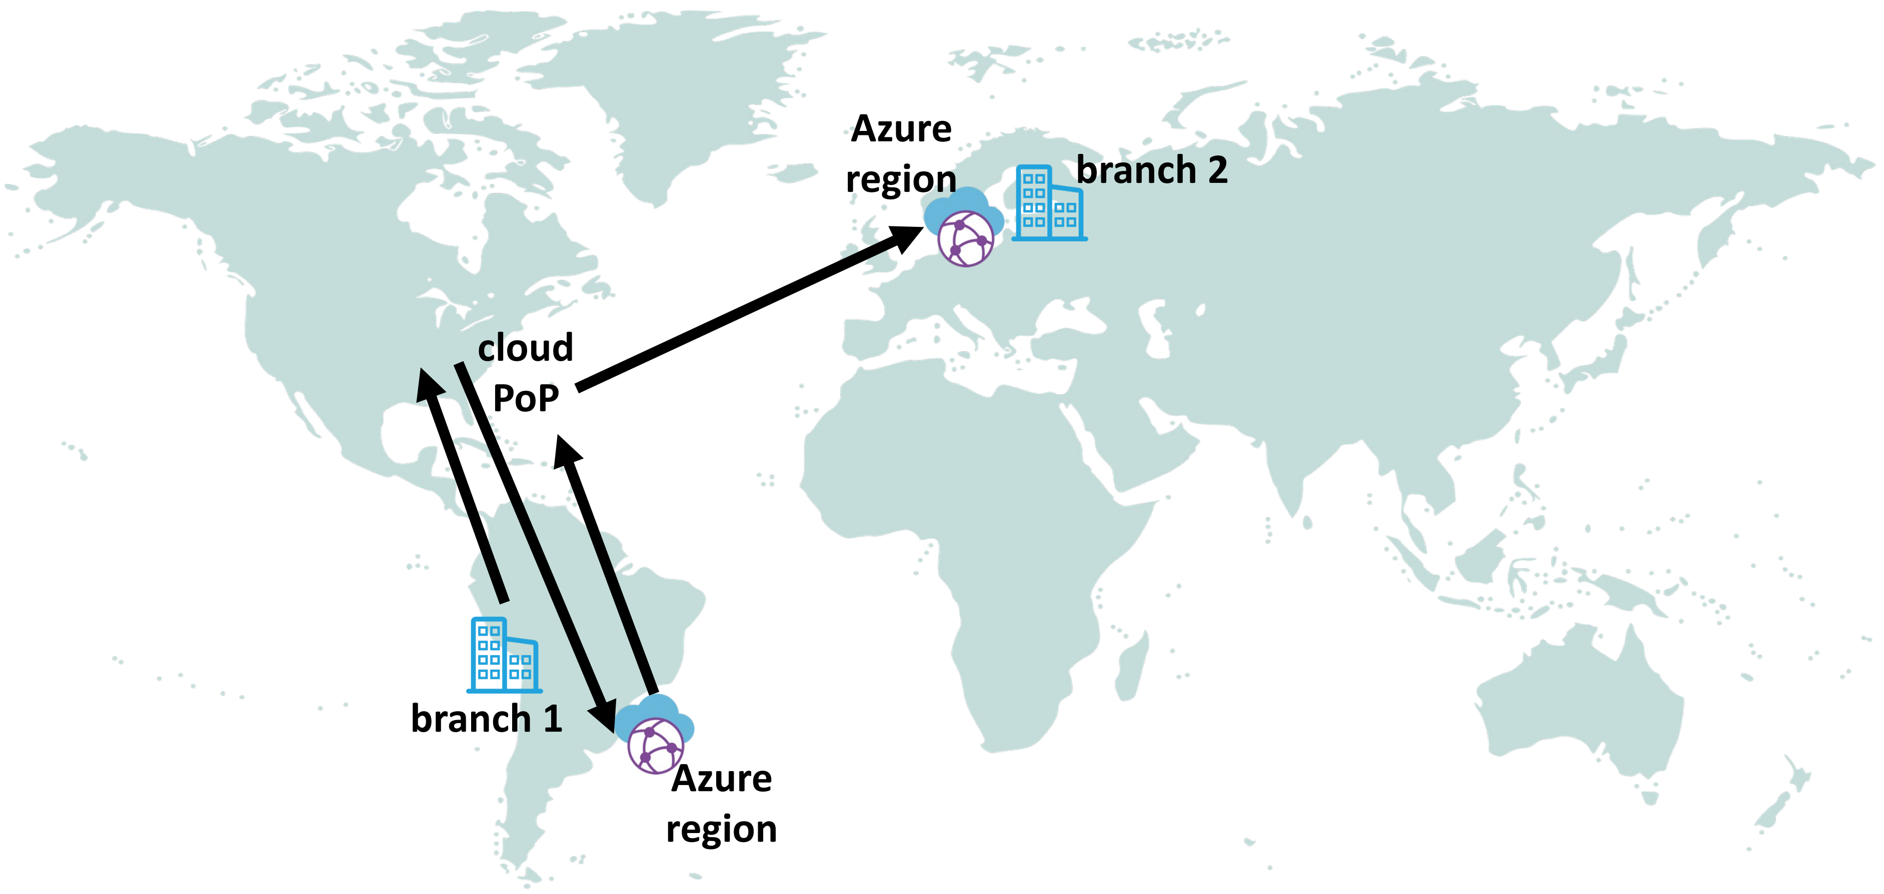

problem: need for resiliency & low latency

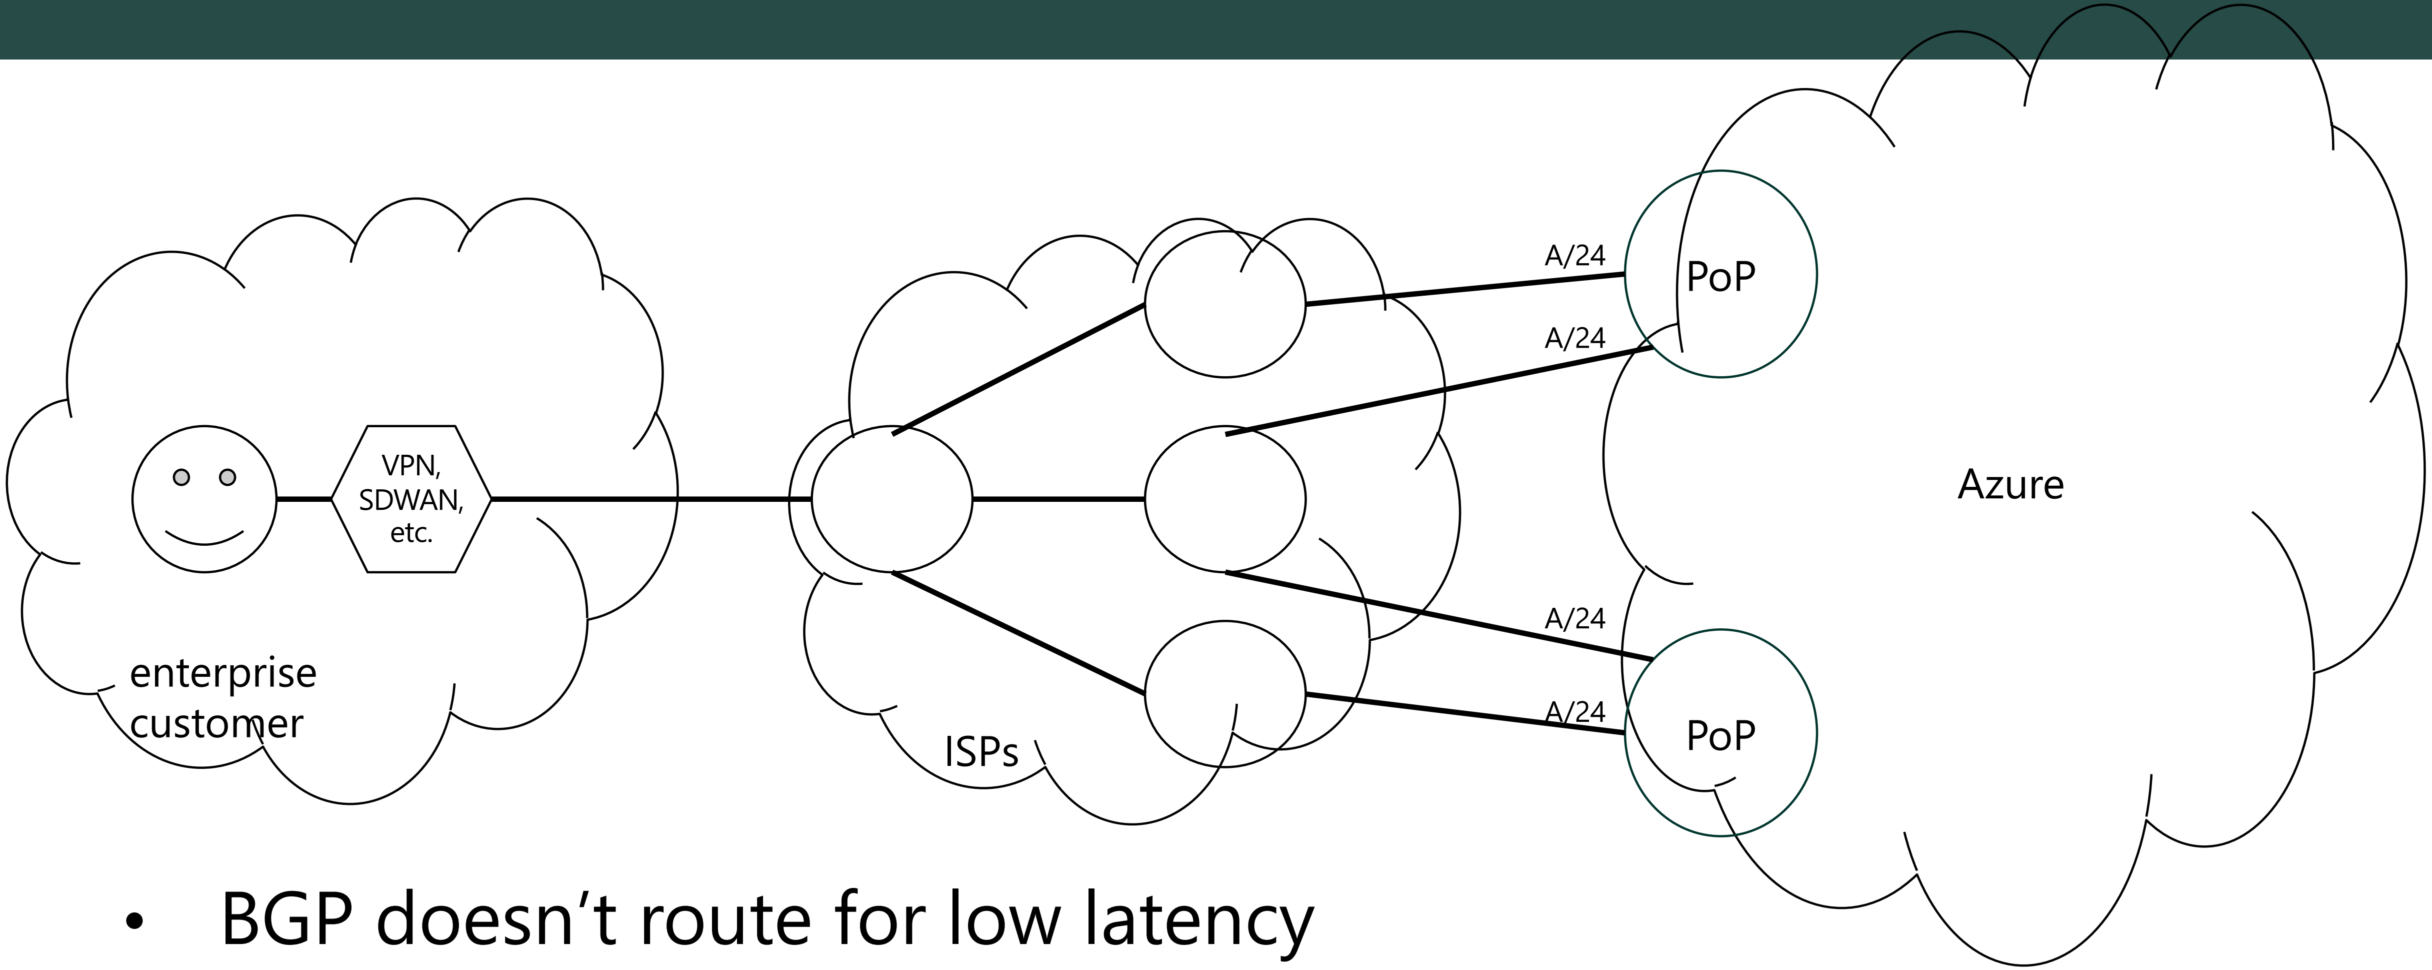

- BGP doesn’t route for low latency
- hence, most clouds peer in many locations and use anycast
- tends to work, but not in all cases

PAINTER: learn the best direct paths to cloud

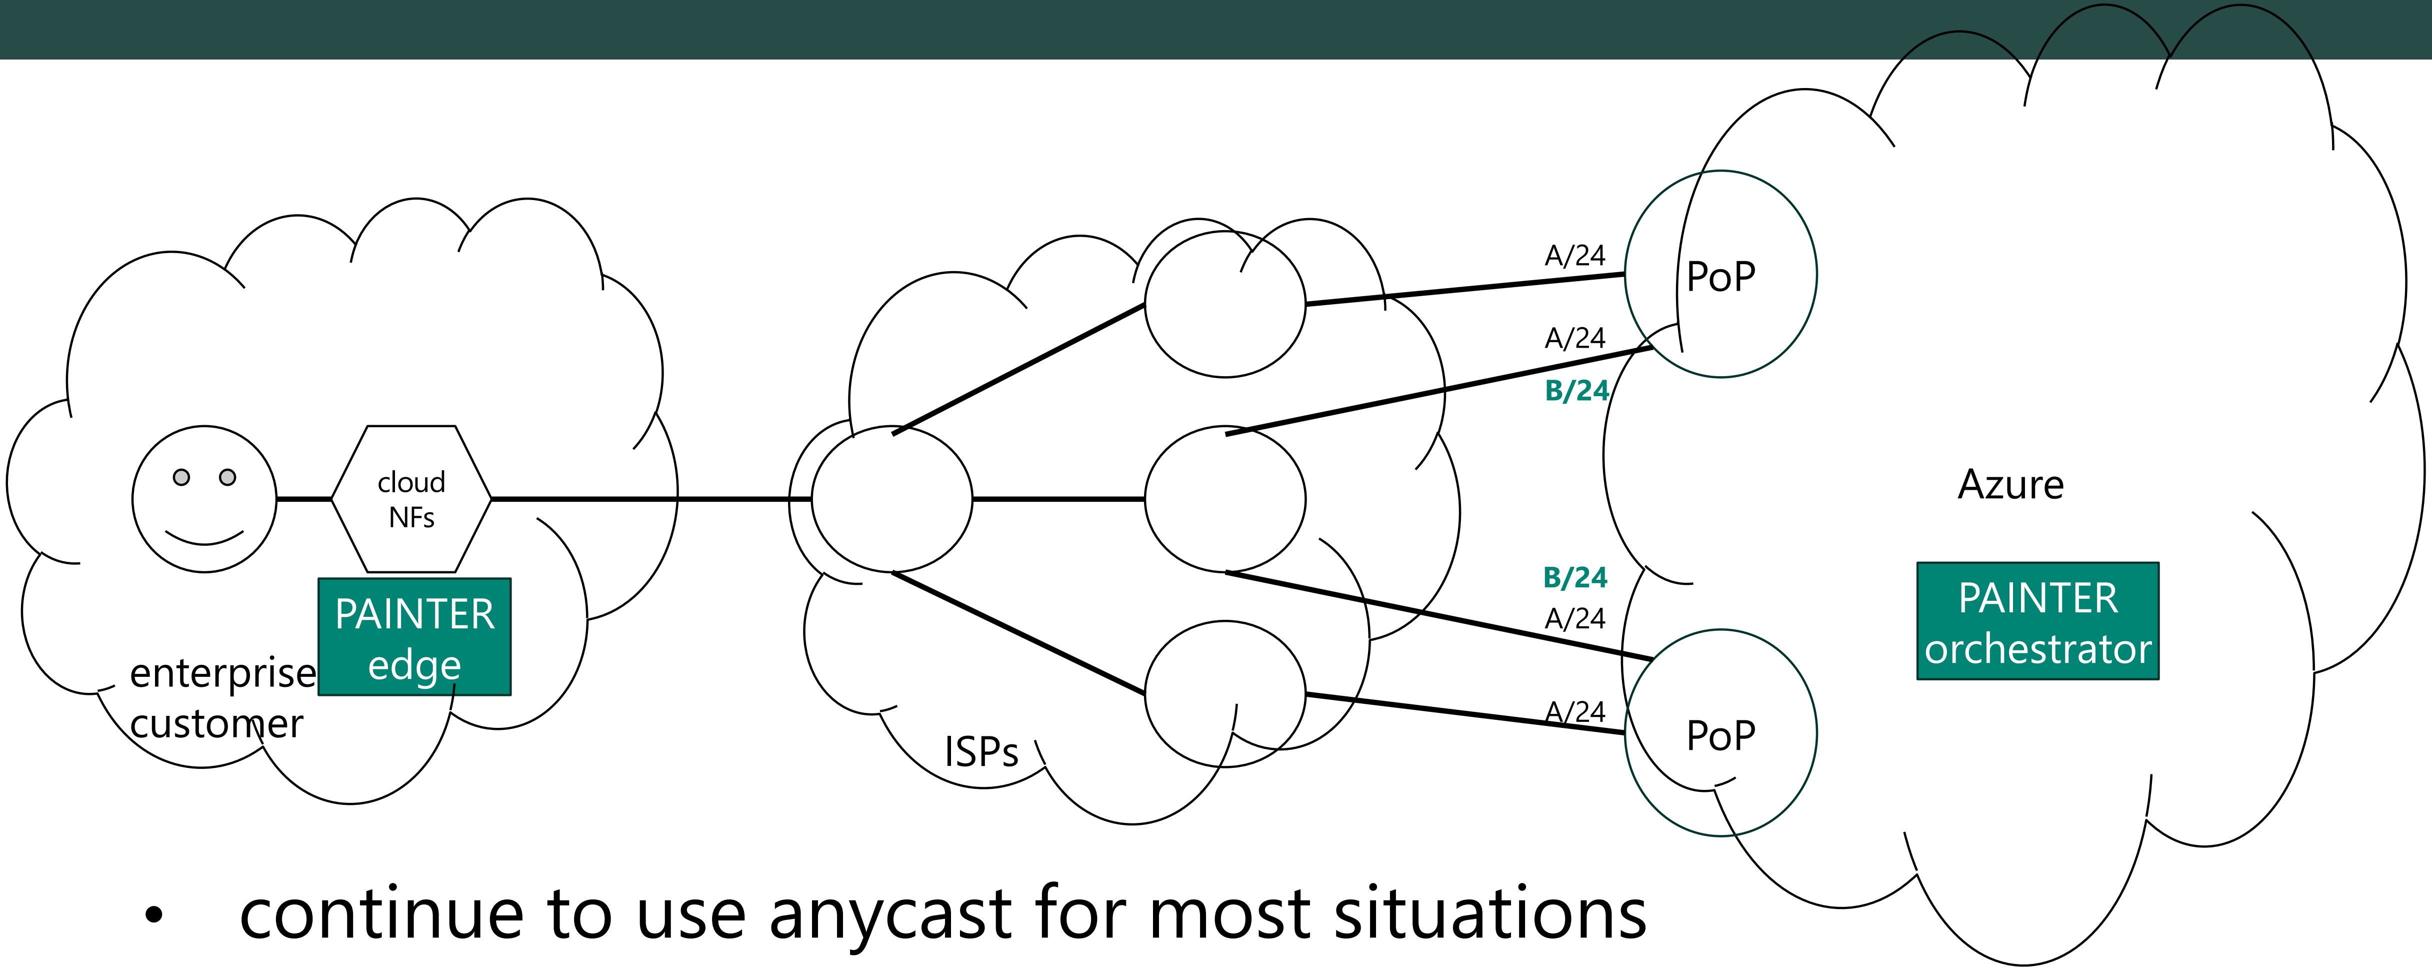

- continue to use anycast for most situations
- exploit large peering surface area to open up unicast paths
- learn & adjust to pick best path & limit prefix costs

challenge: scaling to O(10,000) peering connections

- O(10K) prefixes is $$ and bloats routing tables → prefix reuse
- how will an advertisement reach a customer?
- which advertisements overlap?
- which path is lower latency?

Algorithm 1 Algorithm for se

```
Input Prefix Budget PB, n
RM ← []
while learning do                                          ▷ Ter
    CC ← []
    for p in range(PB) do
        while True do
            peering_improvements                           [state]
            ranked_peerings ← sort
            found_peering ← False
            for next_best_peering in ranked_peerings do    ▷ Greedy search
                NP ← (p, next_best_peering)                ▷ Proposed new prefix,peering
                if B(NP; CC) > 0 then                      [reward]  ▷ Require positive benefit.
                    found_peering ← True                   ▷ Choose this one
                    break
                end if
            end for
            if found_peering then
                CC.append(NP)
            else
                break
            end if
        end while
    end for
    RM ← execute_advertisement                             [action]
end while
return CC
```

L1: leverage cloud scale to overcome limitations of deployed network protocols & use AI+ML to manage that massive scale

L2: custom learning algorithms can take you far, but are harder to deploy than off-the-shelf AI+ML algorithms

summary: we predict a way out of high latency paths

- WAN latencies eat up any benefit of new radio latencies
  - major new era of last mile no longer being the bottleneck
  - deployed WAN protocols haven't evolved
- we work around them by exploiting cloud scale of O(10K) peerings
  - massive scale requires learning
  - ~50ms saving in many cases, with tail savings of ~175ms
- we added complexity for efficiency
  - small number of learning iterations gets us to optimal
  - worsens understandability, maintainability, deployability
  - could have been simpler but less optimal to use off-the-shelf ε-greedy RL

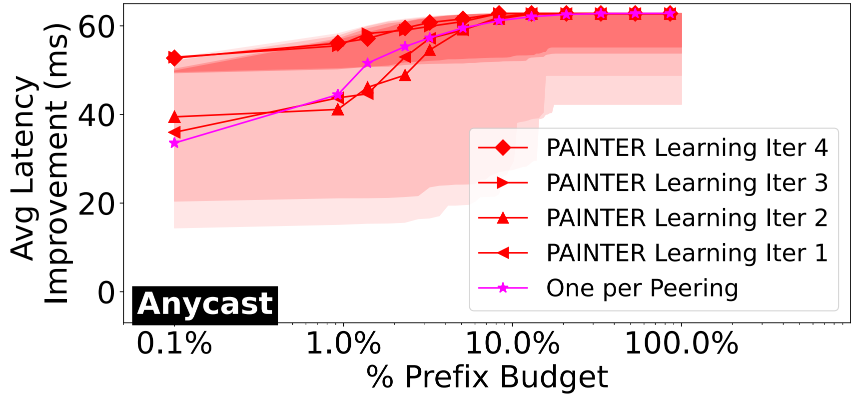

Microsoft

LLexus: an AI agent system for incident management

ACM SIGOPS OSR 2024

Pedro Las-Casas
Alok Gautum Kumbhare
Rodrigo Fonseca
Sharad Agarwal

3GPP Overall Architecture and Specifications

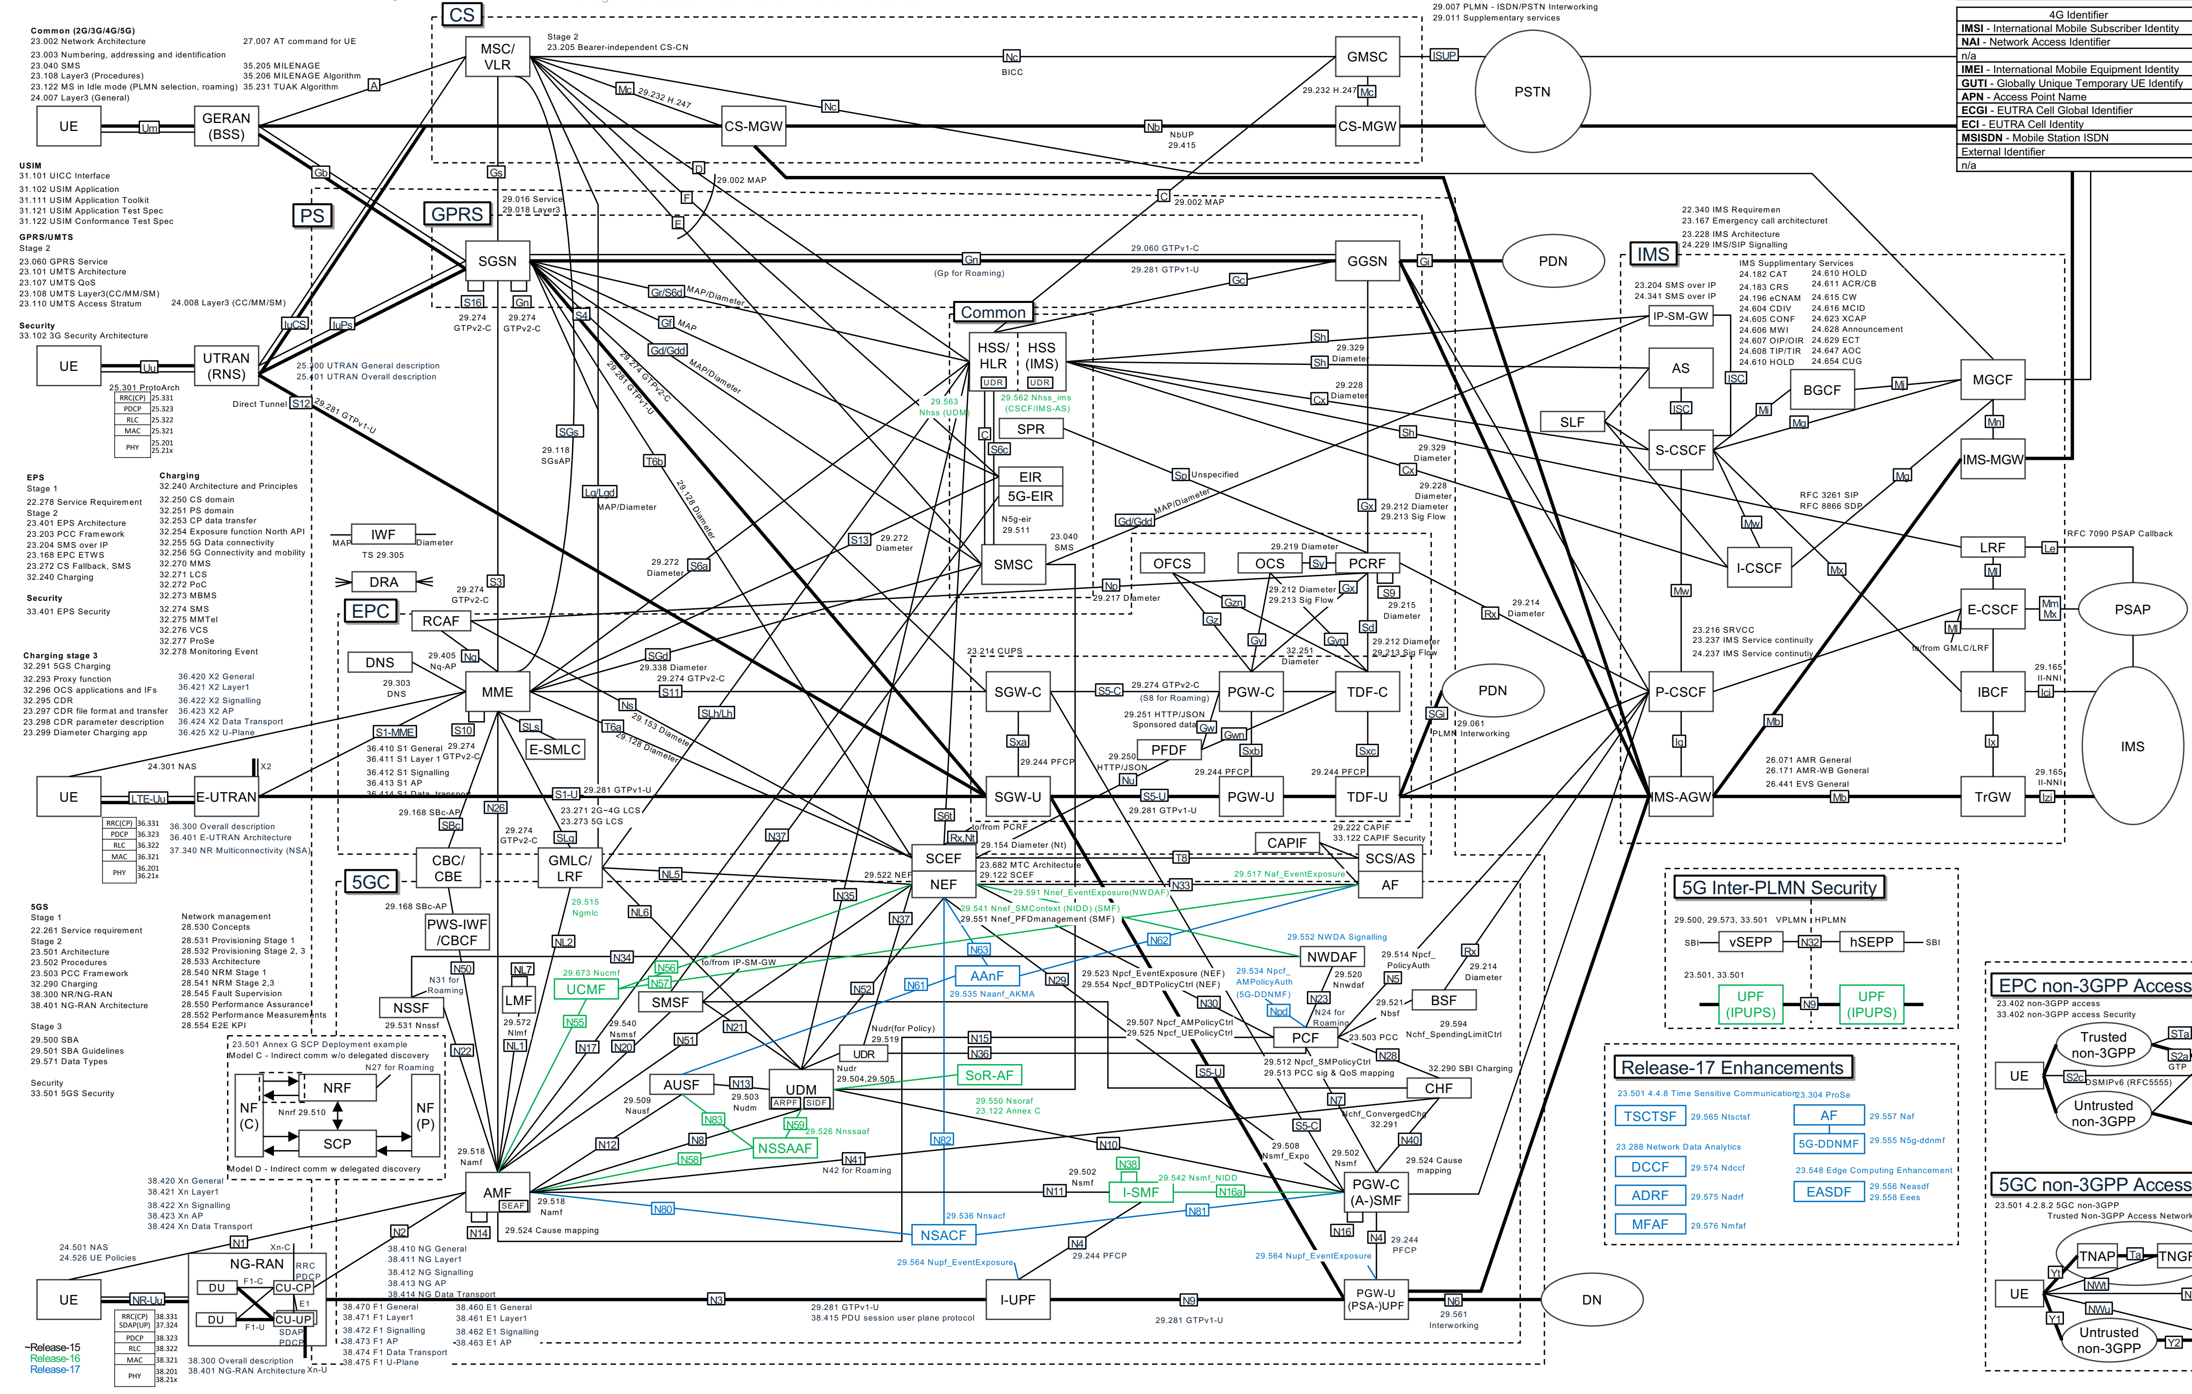

4G and 5G Identifier mapping

4G Identifier	5G Identifier
IMSI - International Mobile Subscriber Identity	**SUPI** - Subscription Permanent Identifier
NAI - Network Access Identifier	**SUPI** - Subscription Permanent Identifier
n/a	**SUCI** - Subscription Concealed Identifier
IMEI - International Mobile Equipment Identity	**PEI** - Permanent Equipment Identifier
GUTI - Globally Unique Temporary UE Identify	**5G-GUTI** - 5G Globally Unique Temporary UE Identify
APN - Access Point Name	**DNN** - Data Network Name
ECGI - EUTRA Cell Global Identifier	**NCGI** - NR Cell Global Identity
ECI - EUTRA Cell Identity	**NCI** - NR Cell Identity
MSISDN - Mobile Station ISDN	**GPSI** - Generic Public Subscription Identifier
External Identifier	**GPSI** - Generic Public Subscription Identifier
n/a	**S-NSSAI** - Single-Network Slice Selection Assistance Information

5G Network Function Abbreviations

Release-15
- **5G-EIR** - 5G-Equipment Identity Register
- **AAnF** - AKMA (Authentication and Key Management for Applications) Anchor Function
- **AF** - Application Function
- **AMF** - Access and Mobility Management Function
- **AUSF** - Authentication Server Function
- **ARPF** - Authentication credential Repository and Processing Function
- **BSF** - Binding Support Function
- **CAPIF** - Common API Framework for 3GPP northbound APIs
- **CHF** - Charging Function
- **I-UPF** - Intermediate UPF
- **LMF** - Location Management Function
- **LRF** - Location Retrieval Function
- **N3IWF** - Non-3GPP InterWorking Function
- **NEF** - Network Exposure Function
- **NRF** - Network Repository Function
- **NSSF** - Network Slice Selection Function
- **NWDAF** - Network Data Analytics Function
- **PCF** - Policy Control Function
- **SCP** - Service Communication Proxy
- **SEAF** - SEcurity Anchor Function.
- **SEPP** - Security Edge Protection Proxy
- **SIDF** - Subscription Identifier De-concealing Function
- **SMF** - Session Management Function
- **SMSF** - Short Message Service Function
- **TNAP** - Trusted Non-3GPP Access Point
- **TNGF** - Trusted Non-3GPP Gateway Function
- **TWIF** - Trusted WLAN Interworking Function
- **UDM** - Unified Data Management
- **UDR** - Unified Data Repository
- **UDSF** - Unstructured Data Storage Function
- **UPF** - User Plane Function

Release-16
- **IPUPS** - Inter PLMN UP Security
- **I-SMF** - Intermediate SMF
- **NSSAAF** - Network Slice-specific and SNPN Authentication and Authorization Function
- **UCMF** - UE radio Capability Management Function
- **SoR-AF** - Steering of Roaming Application Function

Release-17
- **5G DDNMF** - 5G Direct Discovery Name Management Function
- **ADRF** - Analytics Data Repository Function
- **EASDF** - Edge Application Server Discovery Function
- **MFAF** - Messaging Framework Adaptor Function
- **DCCF** - Data Collection Coordination Function
- **NSACF** - Network Slice Admission Control Function
- **TSCTSF** - Time Sensitive Communication and Time Synchronization Function

EPC non-3GPP Access

23.402 non-3GPP access
33.402 non-3GPP access Security

5GC non-3GPP Access

23.501 4.2.8.2 5GC non-3GPP
Trusted Non-3GPP Access Network (TNAN)

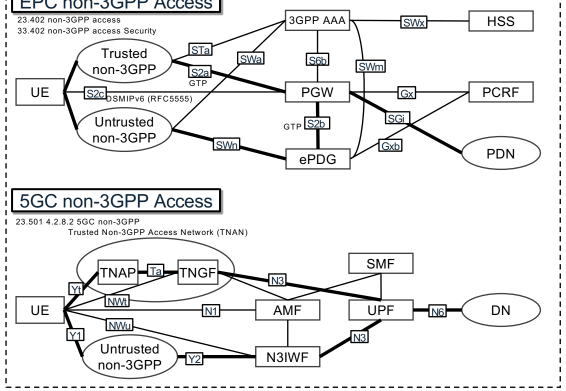

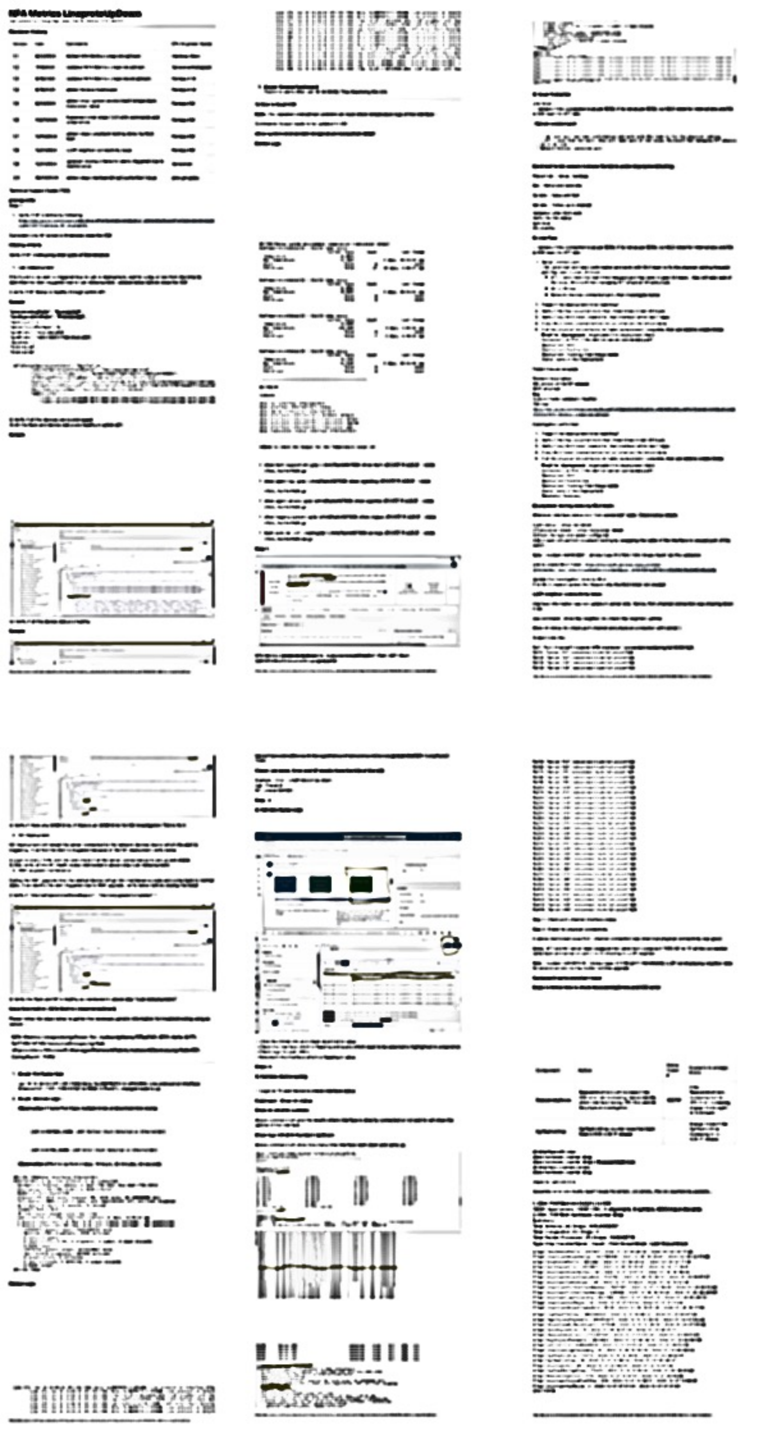

what happens when a 5G NF has a failure?

- SRE (site reliability engineer) follows a TSG (troubleshooting guide)
- TSGs are long, with many laborious steps & actions
  - e.g., link went down
  - switch down? server down? server load too high? physical port too hot? fan RPM? syslogs? metrics? dirty power? is link flapping?
- TSGs are continually evolving as product evolves
  - and new causes and/or behaviors are discovered & documented
- immense pressure to maintain 99.999% availability

LLexus automates TSG execution using AI agents

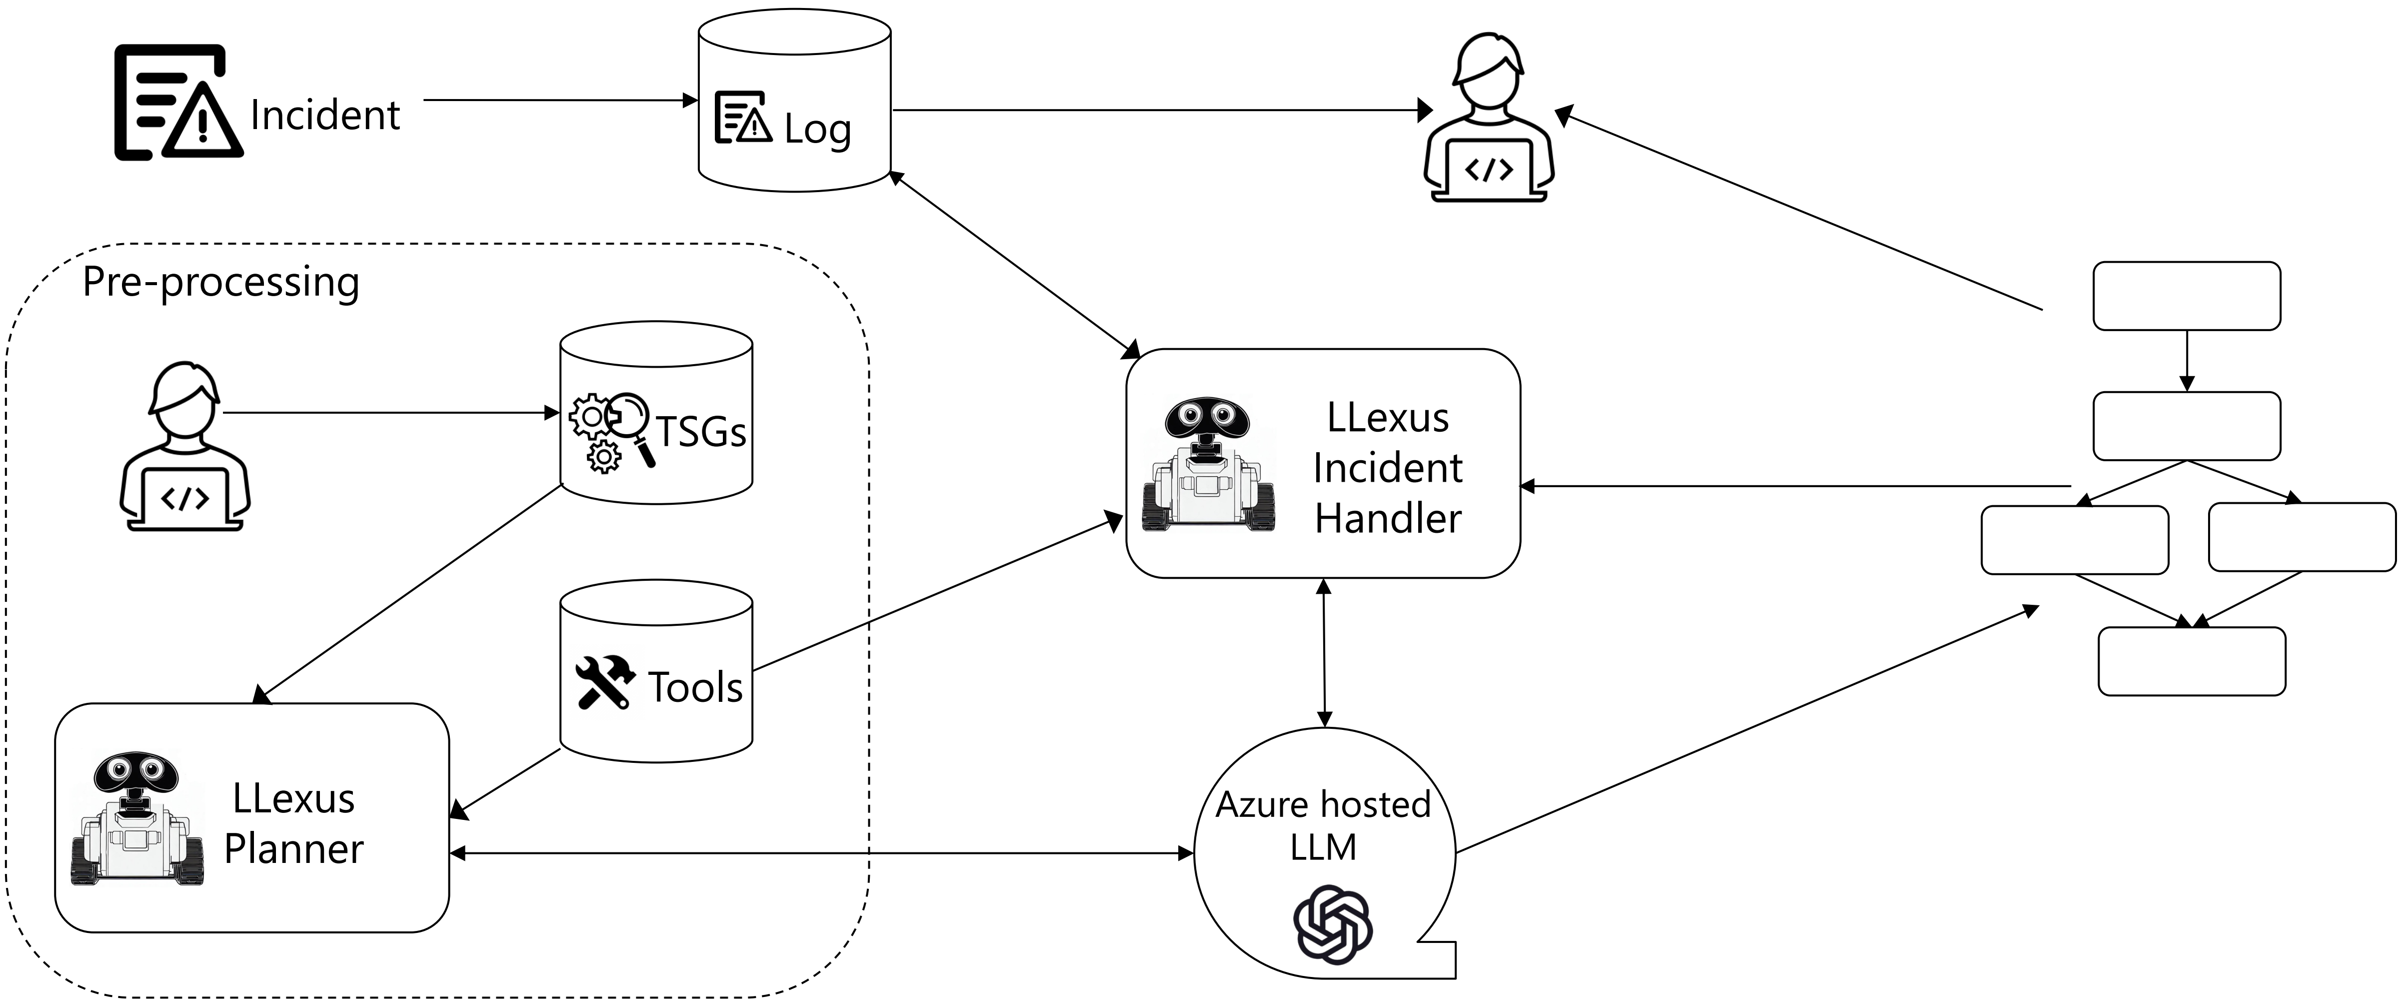

design to limit hallucinations

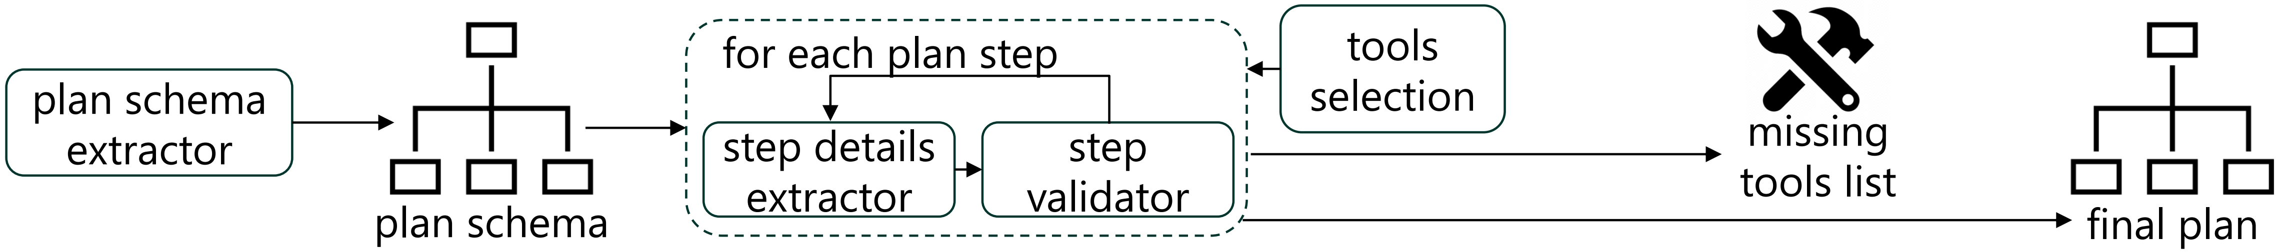

- iterative plan generation (high level plan, then focus on each step)
  - each call to LLM is a smaller task, with more limited scope for hallucinations
  - validation rules to check output of each LLM call
- pre-generate plan, instead of doing it at incident occurrence
  - allow human audits to catch any hallucinations when TSG is updated
- use existing tools (e.g., tool for "show interface status")
  - far safer than giving an LLM unfettered access to ssh on a network device
- and more …

L3: reduce risk from AI hallucinations with careful system design

summary: we use AI agents for better network uptime

- LLexus targets automatic mitigation of live site incidents
  - surprisingly, it can make sense of long, complex, technical documents
  - it can create methodical plans with many steps, branches, etc.

- value-prop: improve TSGs, reduce MTTR & SRE cost, scale support

- careful systems design mitigates many hallucinations
  - use multiple calls to iteratively refine complex tasks
  - use tools with specific scope to limit what AI agents can do
  - use human audits at key but hopefully rare points as safeguards

Microsoft

in conclusion...

edge + cloud compute has finally arrived at scale

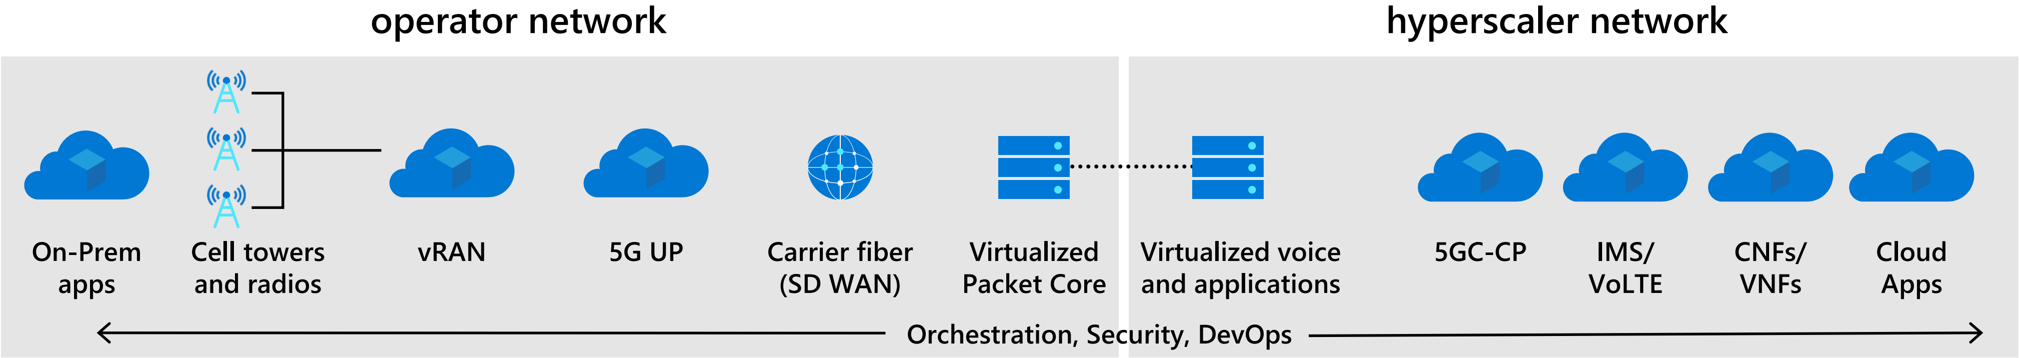

- exciting new capabilities are behind cloudification of telecom infra
  - new radios, NF virtualization & use of commodity compute at edge & DCs
- why build on hyperscalers?
  - new business models tend to be cloud centric
  - leverage economies of scale in compute
  - ease of future upgrades (5G NF → 6G NF and faster servers)
  - promise of AI & ML in cloud for analytics

there are several avenues for impactful research

- carrier-grade reliability
  - challenging target for cloud, edge compute, and the network in between
  - many new NFs & other components that are being deployed at scale
- end-to-end QoS
  - new radio promises ~1ms latency, 10-20 Gbps throughput, 99.999% reliability
  - only useful if the end-to-end path can also provide those guarantees
- security & privacy
  - control of SW & HW is spread across multiple administrative domains
  - NFs & other components are from multiple vendors
- cost (e.g., packets per core)
  - compute at edges is limited due to space & power constraints
- and more…

Microsoft

Thank you!