



OTTER:

Efficient Multi-WAN Transport for 5G

Mary Hogan, **Gerry Wan**, Yiming Qiu, Sharad Agarwal,
Ryan Beckett, Rachee Singh, Paramvir Bahl

Two transformations enabling next-gen apps

Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols



Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols

>10 Gbps
throughput



Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols

>10 Gbps
throughput




<1 ms
latency



Two transformations enabling next-gen apps

5G New Radio


- New frequency bands
- Flexible radio multiplexing
- New RAN protocols



Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols




Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols


Cloudification




Two transformations enabling next-gen apps

5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols




Cloudification



Two transformations enabling next-gen apps

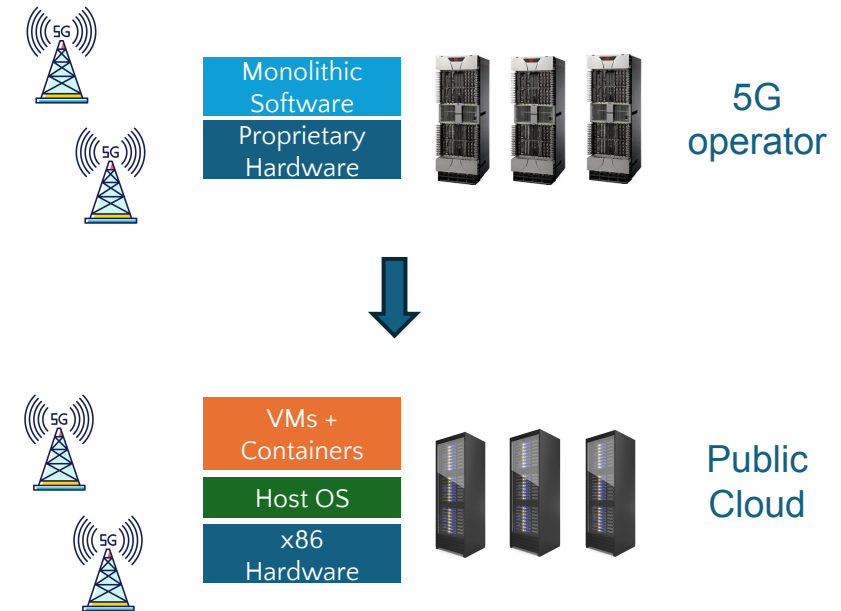
5G New Radio

- New frequency bands
- Flexible radio multiplexing
- New RAN protocols



Cloudification

- Network function virtualization
- Flexible deployment
- Statistical multiplexing



5G NR and cloudification are already underway

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 with 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 Gbps aggregated downlink speed.
- Implementation is based on 5G standalone (SA) New Radio-Dual Connectivity (NR-DC) and carrier

Telecommunications

Rethinking 5G: The cloud imperative

February 19, 2025

Eric Parsons
Head of Emerging Segments, Business Area Cloud
Software and Services, Ericsson

Muninder Sambi
VP, Cloud Networking

Verizon launches the world's first MEC platform with AWS Wavelength.

For the first time, developers can innovate and build new applications at the edge of the Verizon 5G Ultra Wideband network by accessing AWS Wavelength compute and storage services.

5G Core Software as a Service: speed, ease, agility


SaaS 5G SA Mobile Core

The diagram illustrates the SaaS 5G SA Mobile Core architecture. It shows a 5G SA Core connected to AWS via an IPsec tunnel. The 5G SA Core is divided into two main sections: the Control & User Plane and the User Plane. The Control & User Plane includes components like N1, N2, N3, N4, N5, N6, N7, N8, N9, N10, N11, N12, N13, N14, N15, N16, N17, N18, N19, N20, N21, N22, N23, N24, N25, N26, N27, N28, N29, N30, N31, N32, N33, N34, N35, N36, N37, N38, N39, N40, N41, N42, N43, N44, N45, N46, N47, N48, N49, N50, N51, N52, N53, N54, N55, N56, N57, N58, N59, N60, N61, N62, N63, N64, N65, N66, N67, N68, N69, N70, N71, N72, N73, N74, N75, N76, N77, N78, N79, N80, N81, N82, N83, N84, N85, N86, N87, N88, N89, N90, N91, N92, N93, N94, N95, N96, N97, N98, N99, N100. The User Plane includes components like N1, N2, N3, N4, N5, N6, N7, N8, N9, N10, N11, N12, N13, N14, N15, N16, N17, N18, N19, N20, N21, N22, N23, N24, N25, N26, N27, N28, N29, N30, N31, N32, N33, N34, N35, N36, N37, N38, N39, N40, N41, N42, N43, N44, N45, N46, N47, N48, N49, N50, N51, N52, N53, N54, N55, N56, N57, N58, N59, N60, N61, N62, N63, N64, N65, N66, N67, N68, N69, N70, N71, N72, N73, N74, N75, N76, N77, N78, N79, N80, N81, N82, N83, N84, N85, N86, N87, N88, N89, N90, N91, N92, N93, N94, N95, N96, N97, N98, N99, N100. The diagram also shows a connection to the Internet via a Frankfurt location.

VPN to AWS and then out to the internet.


Nokia Core TV series #14: Nokia 5G Core Software as a Service in practice

The cloudification of 5G




5G Traffic can now traverse two WANs

The cloudification of 5G




5G Traffic can now traverse two WANs

The cloudification of 5G



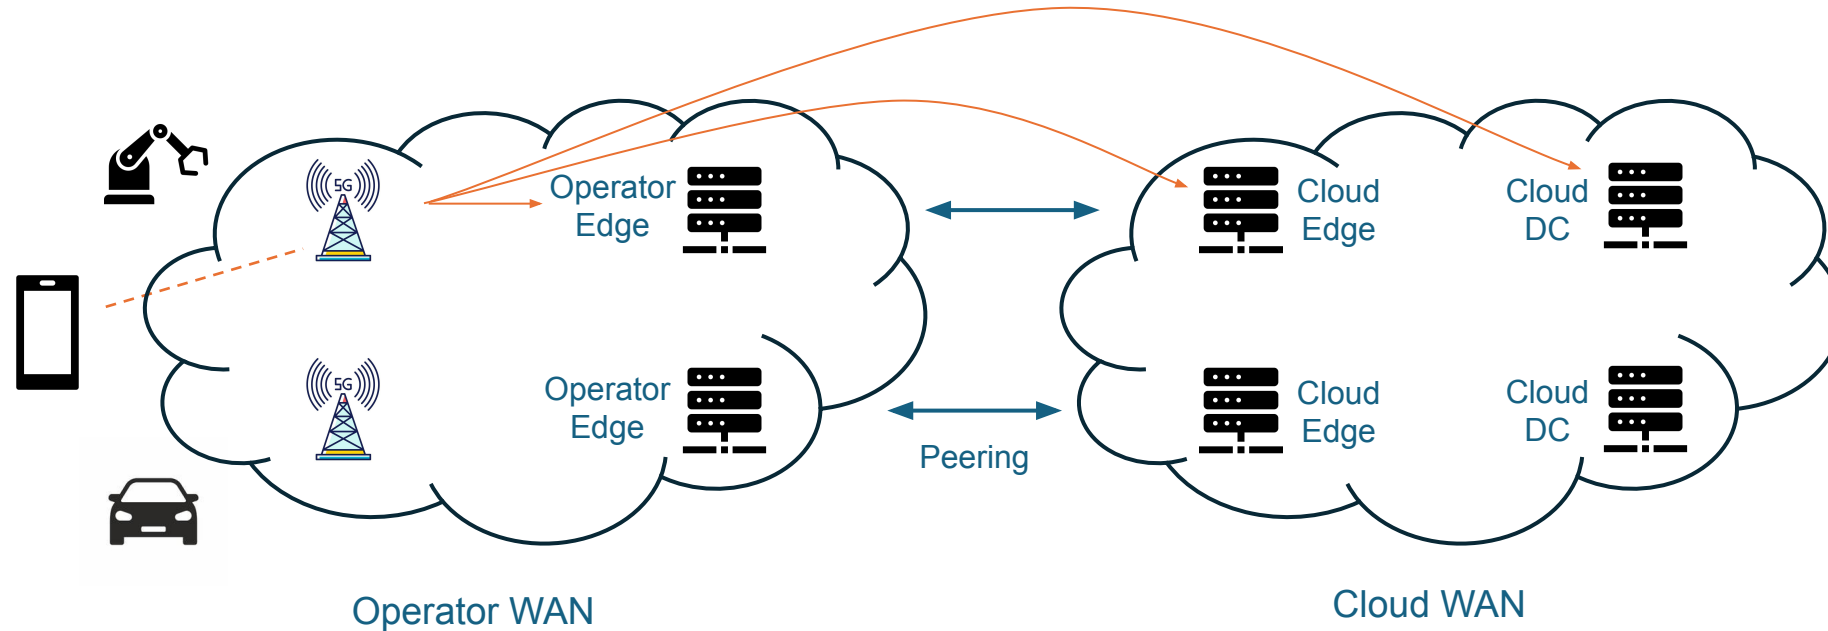
5G Traffic can now traverse two WANs

The cloudification of 5G




5G Traffic can now traverse two WANs

The cloudification of 5G



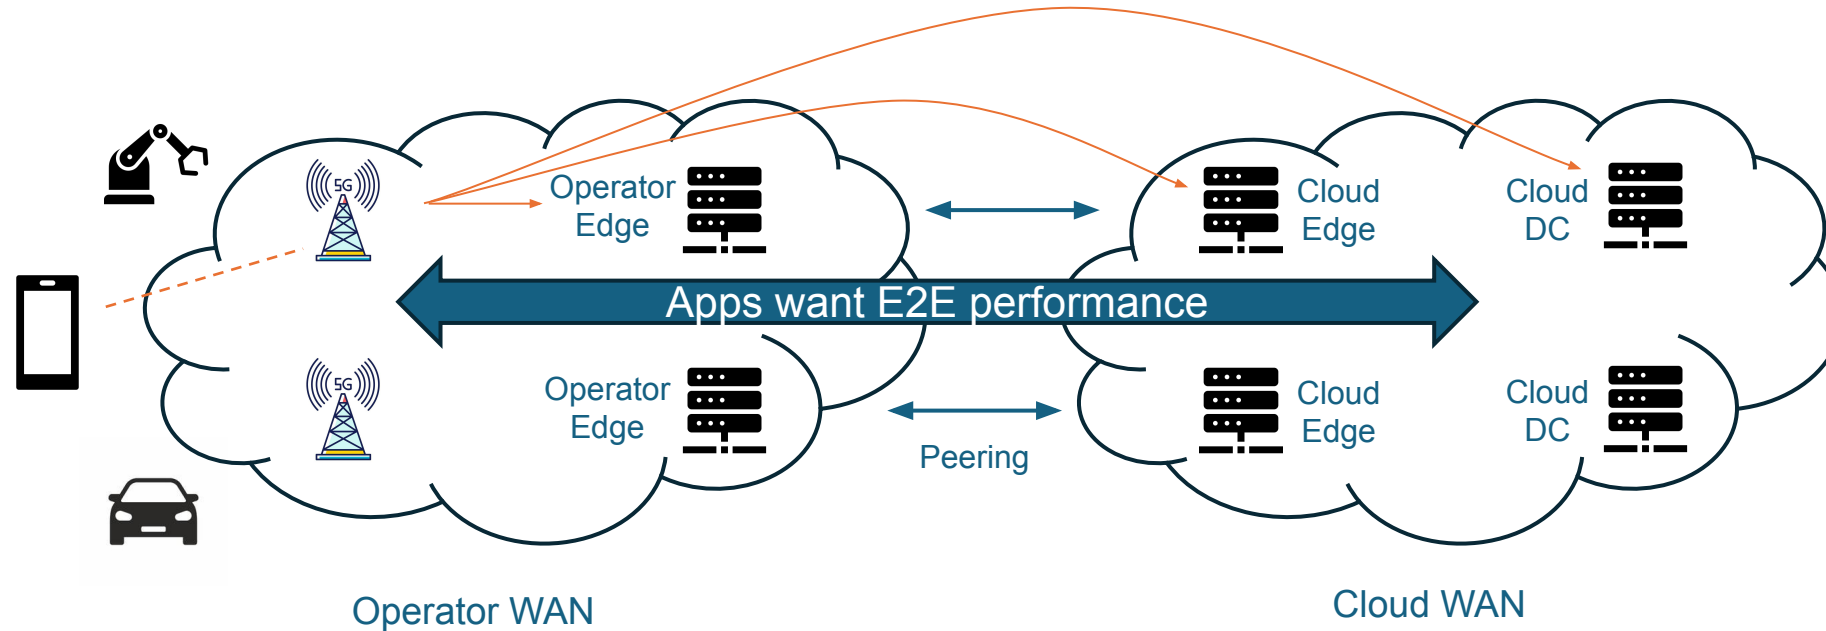
5G Traffic can now traverse two WANs

NR introduces significant pressure on the WANs




- Bandwidth: $O(10\text{Gbps/user})$ hitting cloud services
- Latency: $O(1\text{ms})$ over-the-air compared to WAN latencies & routing inefficiency
- Reliability: going from 99.5 enterprise grade \rightarrow 99.999 carrier grade availability

NR introduces significant pressure on the WANs

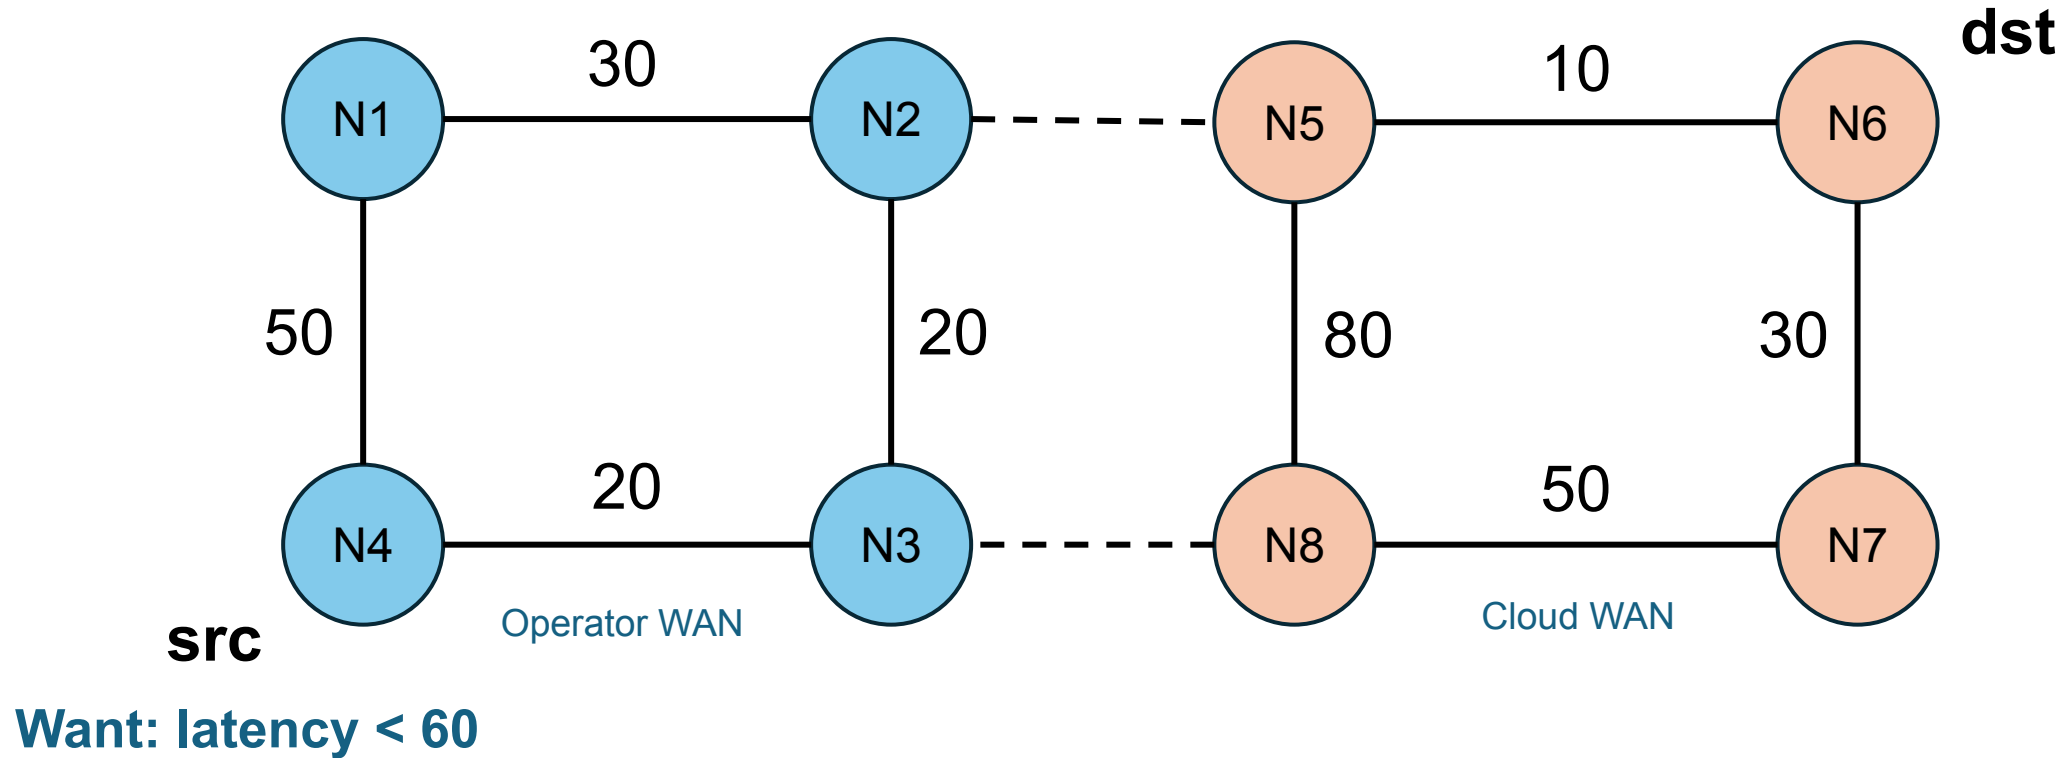


- Shared business incentive: unique opportunity to optimize across two WANs
- Path diversity and differing QoS demands of 5G apps opens door to more intelligent transport decisions

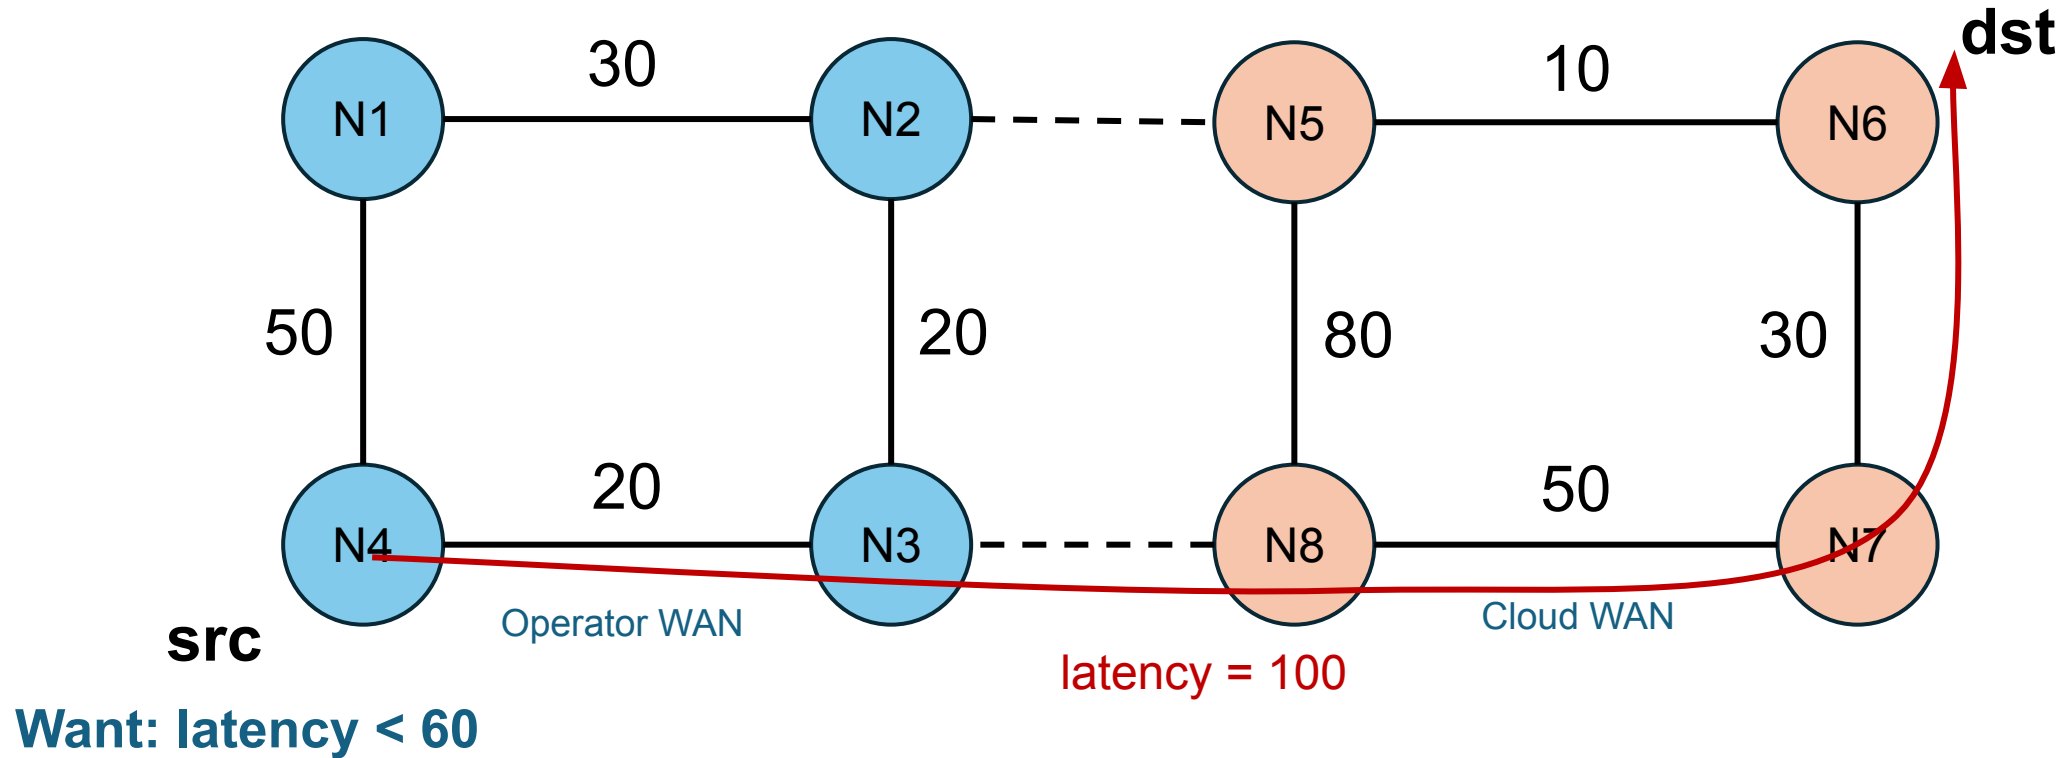
Example: multi-WAN coordination for E2E performance



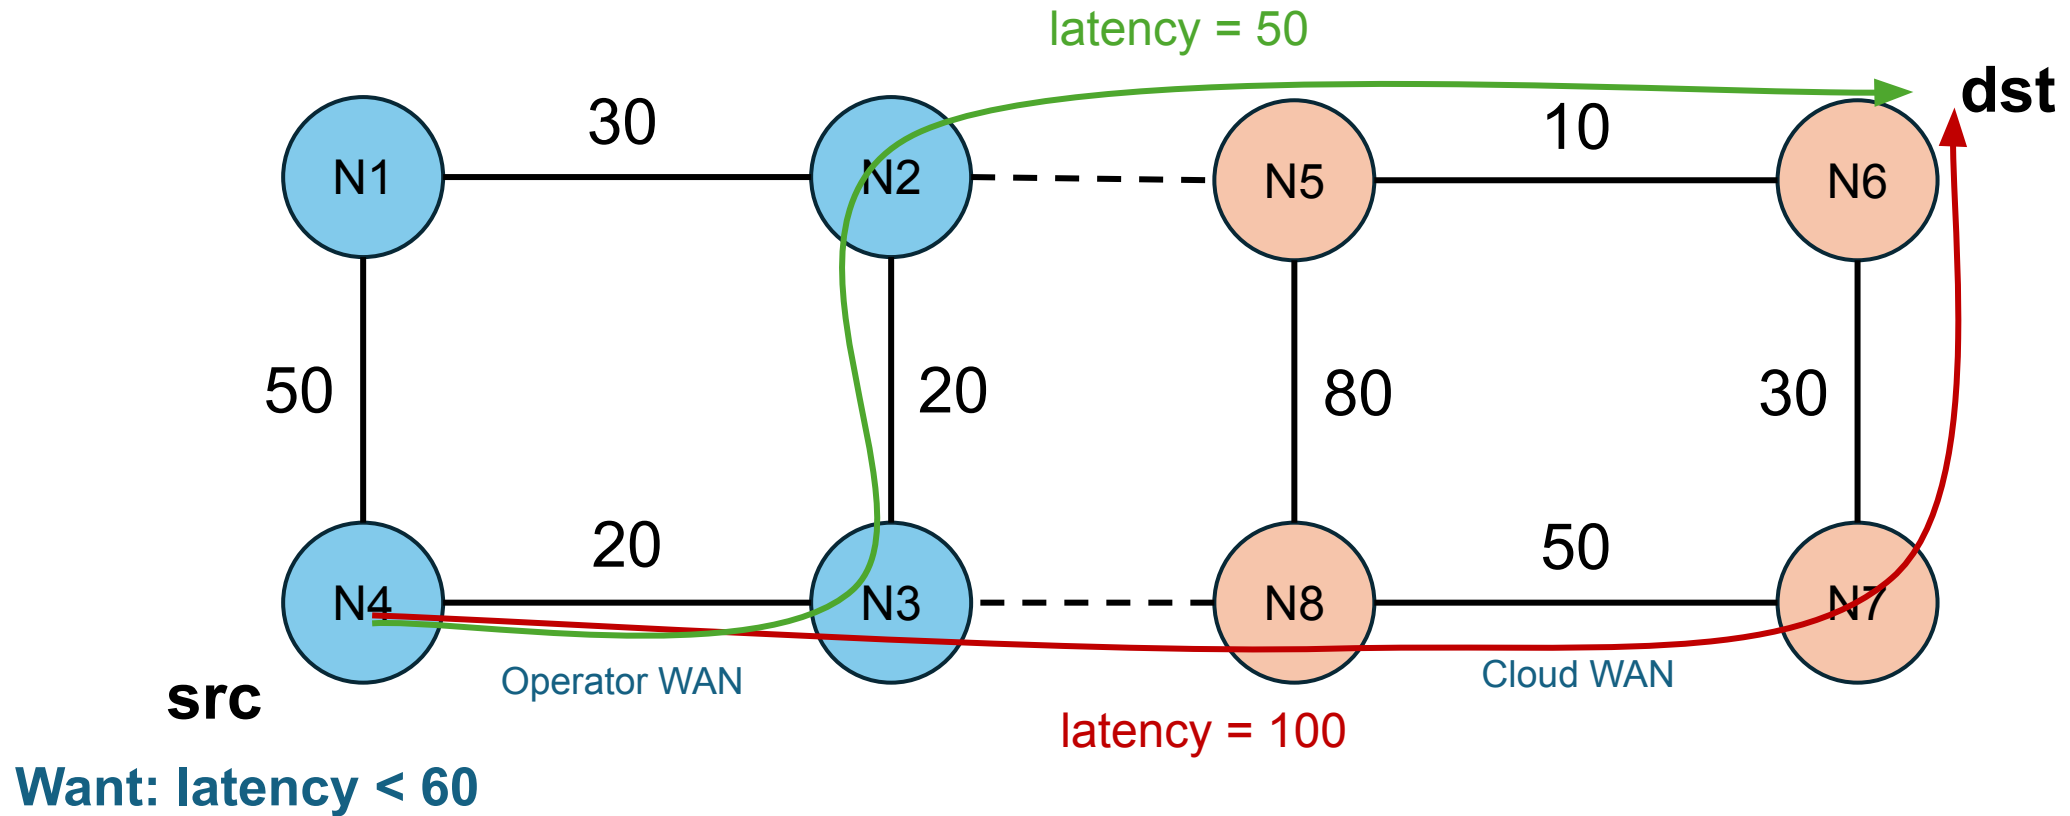
Example: multi-WAN coordination for E2E performance



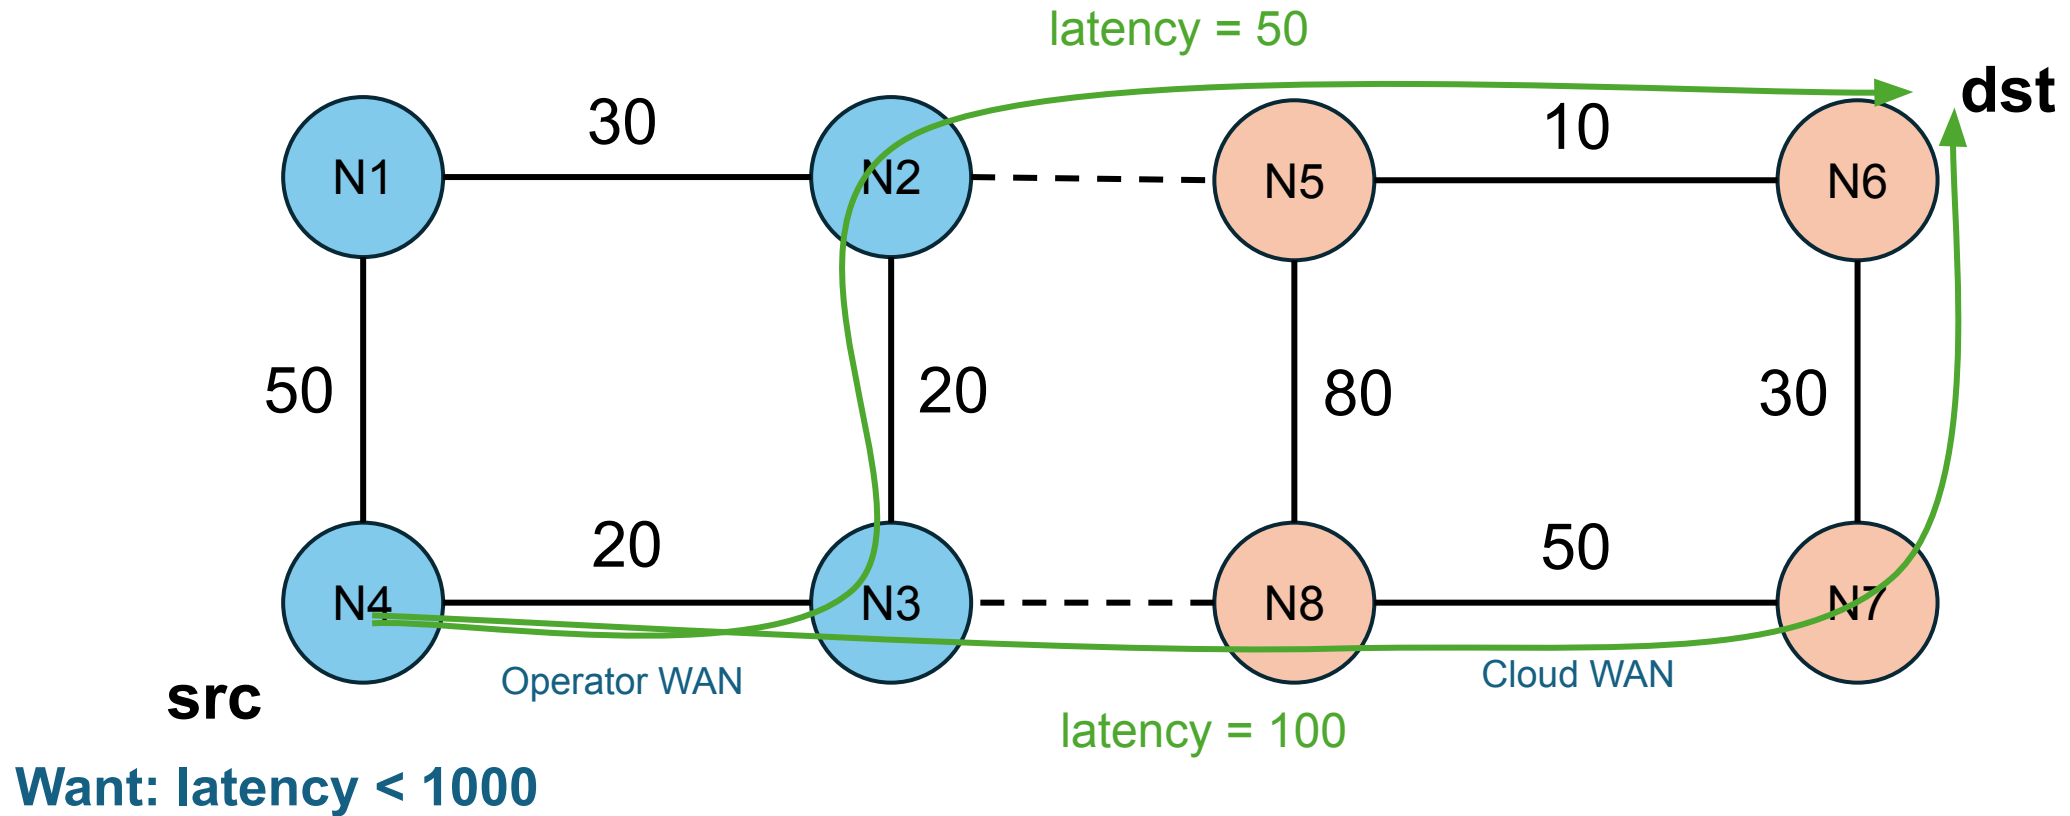
Example: multi-WAN coordination for E2E performance



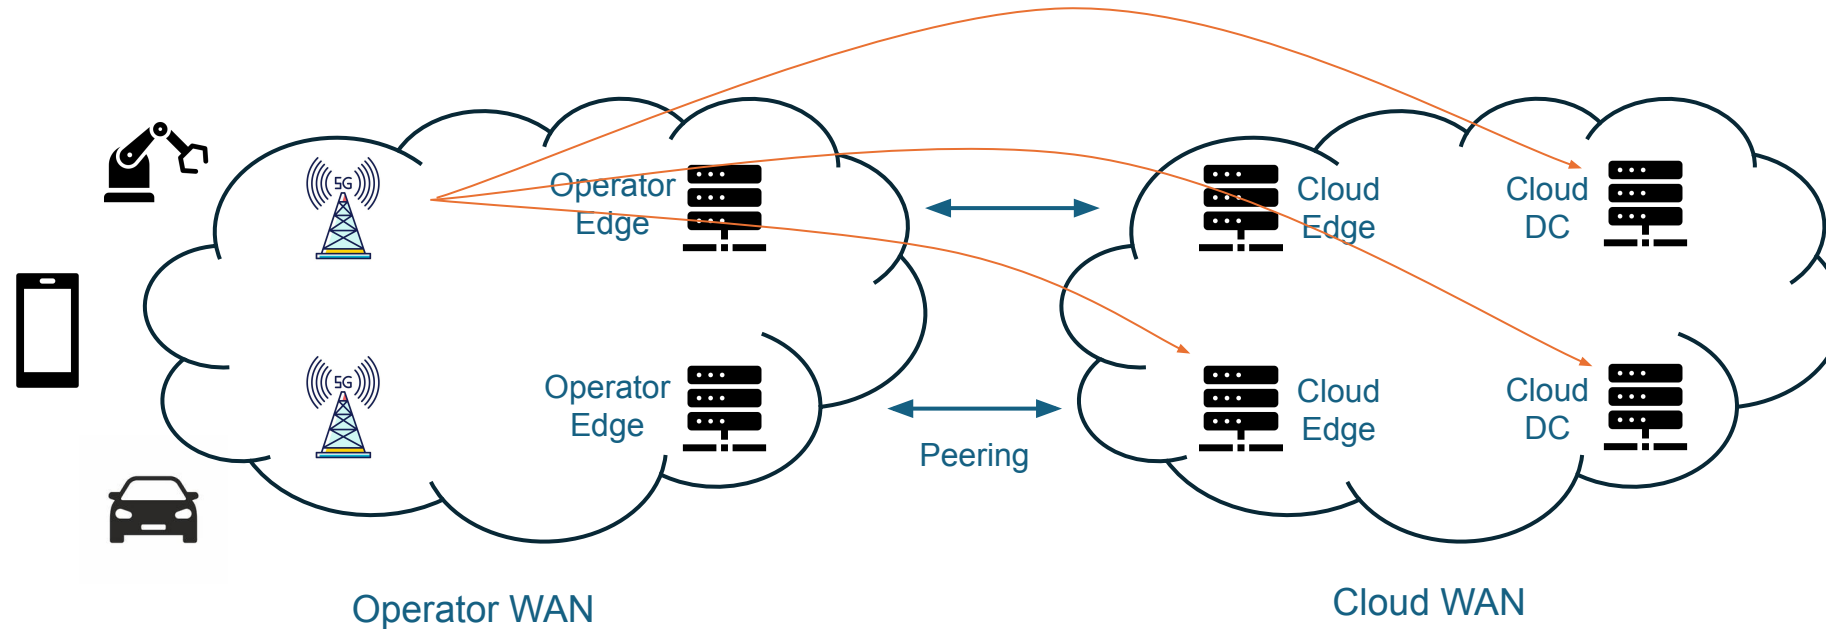
Example: multi-WAN coordination for E2E performance



Example: multi-WAN coordination for E2E performance



Not just paths, but destinations too




- Which edge site?
- Which cloud DC?
- Differentiated performance can come from both choice of path and destination

Problem summary

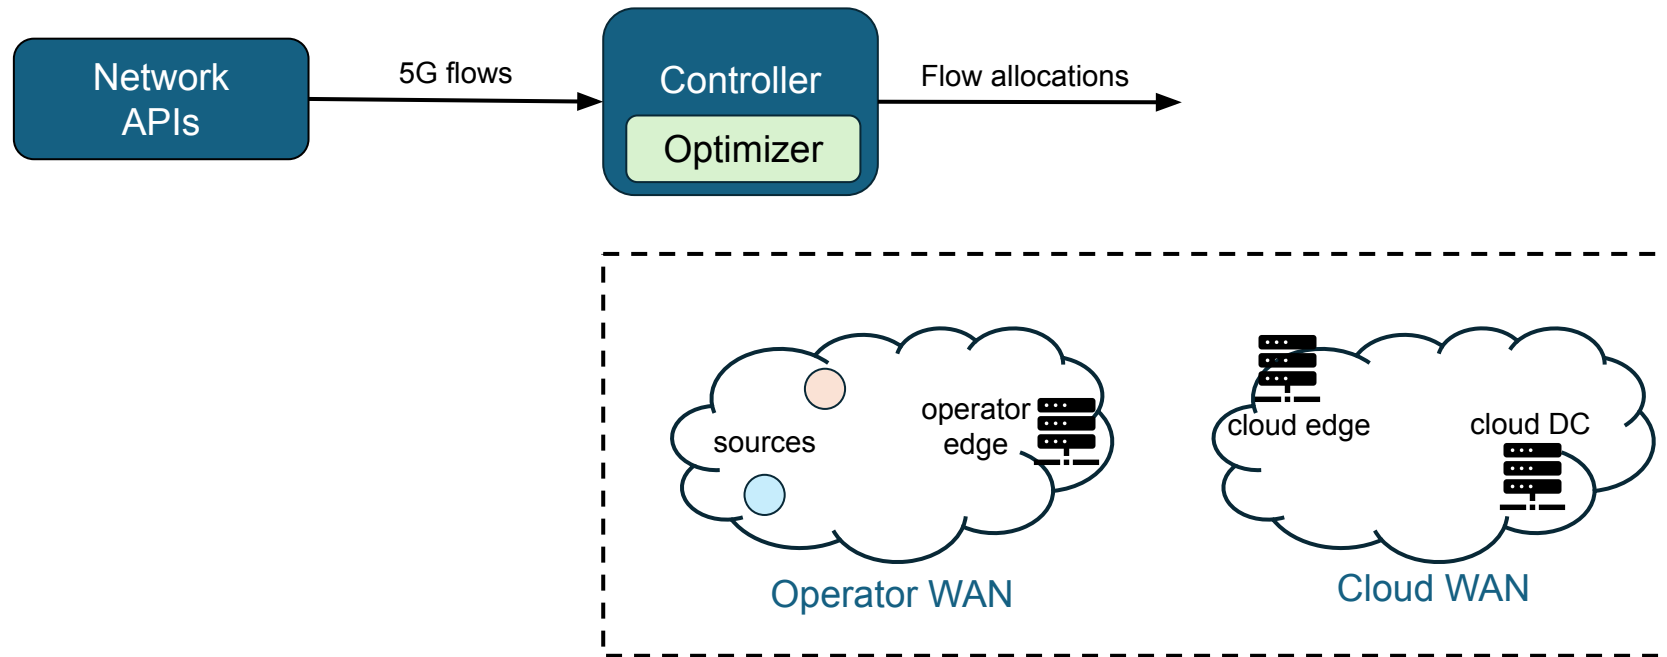
- Next-gen apps enabled by 5G NR and cloudification need E2E QoS
- Transport may include two WANs – operator & cloud
- WAN inefficiencies can negate benefits of NR
- Default routing doesn't support diverse, fine-grained service objectives

OTTER: Overlay Traffic Transport and Efficient Resource allocation

- System for efficient, on-demand transport across both operator and cloud WANs for 5G flows with diverse QoS needs



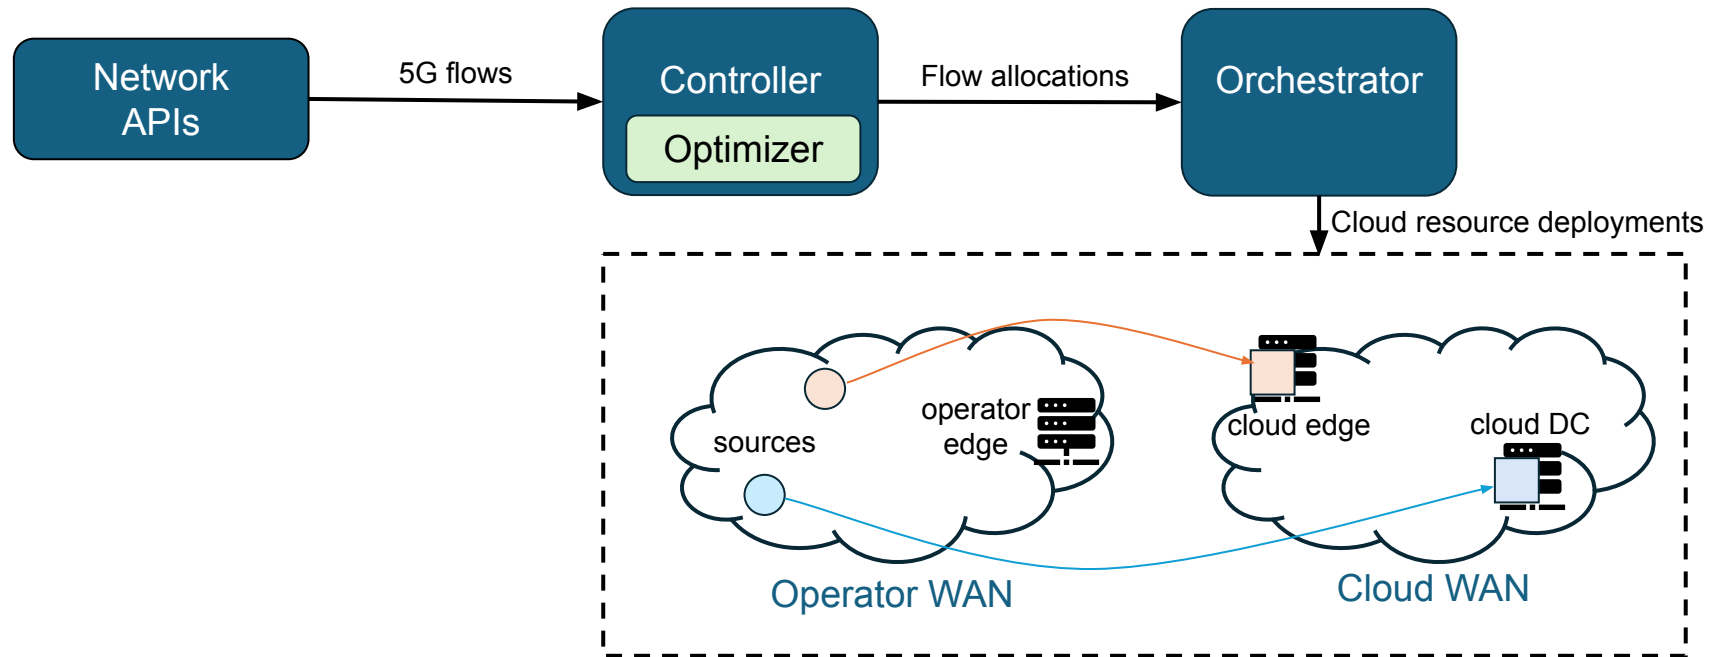
OTTER: Overlay Traffic Transport and Efficient Resource allocation



Three main technical components:

- **Controller** to place 5G flows w/ QoS demands on paths
- Scalable multi-WAN overlay that supports path **orchestration**
- **Measurements** on overlay to identify performance of paths and dsts

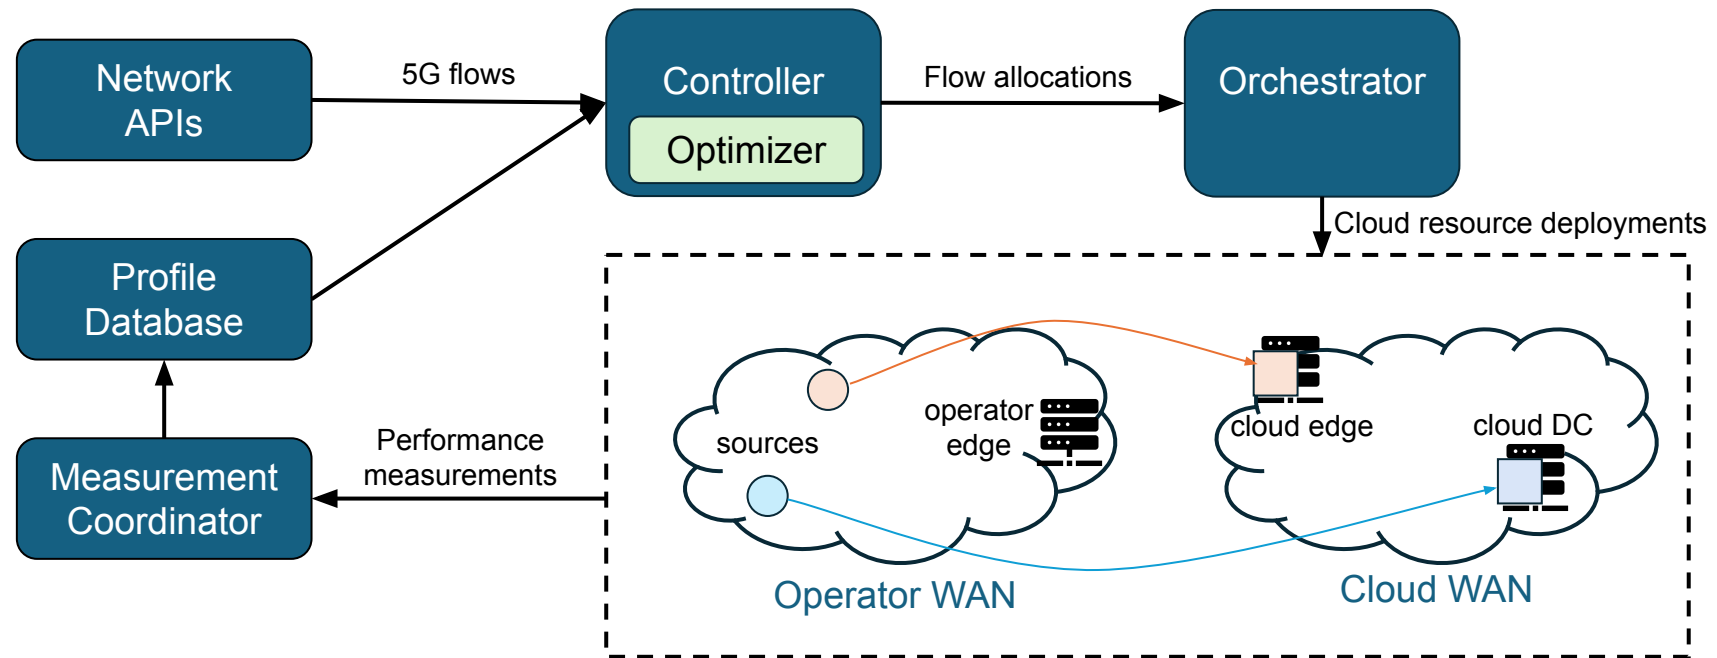
OTTER: Overlay Traffic Transport and Efficient Resource allocation



Three main technical components:

- **Controller** to place 5G flows w/ QoS demands on paths
- Scalable multi-WAN overlay that supports path **orchestration**
- **Measurements** on overlay to identify performance of paths and dsts

OTTER: Overlay Traffic Transport and Efficient Resource allocation



Three main technical components:

- **Controller** to place 5G flows w/ QoS demands on paths
- Scalable multi-WAN overlay that supports path **orchestration**
- **Measurements** on overlay to identify performance of paths and dsts

OTTER Controller

How to optimally place 5G flows with different QoS demands across two WANs?

OTTER Controller

How to optimally place 5G flows with different QoS demands across two WANs?

Optimizer inputs:

- Set of flows F (s , d , bw)
- Set of paths P (s , d , $links$)

OTTER Controller

How to optimally place 5G flows with different QoS demands across two WANs?

Optimizer inputs:

- Set of flows F (s , d , bw)
- Set of paths P (s , d , $links$)

Standard multi-commodity flow problem

OTTER Controller

How to optimally place 5G flows with different QoS demands across two WANs?

Optimizer inputs:

- Set of flows F (s, d, bw)
- Set of paths P ($s, d, links$)



- ($s, D, bw, \text{demand fns, resource req}$)
- ($s, d, links, \text{path metrics}$)

Standard multi-commodity flow problem

OTTER multi-WAN flow placement

Objective Function

$$\operatorname{argmax} \sum_{f,p} \left(x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m) \right) / bw_f$$

Objective Function

$$\operatorname{argmax} \sum_{f,p} (x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m)) / bw_f$$

Amount of bandwidth of
flow **f** assigned to path **p**

Objective Function

$$\operatorname{argmax} \sum_{f,p} \left(x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m) \right) / bw_f$$

Degree of service
demand satisfaction

Objective Function

$$\operatorname{argmax} \sum_{f,p} \left(x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m) \right) / \boxed{bw_f}$$

Bandwidth demand

Objective Function

$$\operatorname{argmax} \sum_{f,p} \left(x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m) \right) / bw_f$$

subject to destination resource constraints

Objective Function

$$\operatorname{argmax} \sum_{f,p} \left(x_{f,p} \cdot \sum_{m \in \{rtt, jit, loss, bw\}} \mathcal{D}_f^m(\sigma_p^m) \right) / bw_f$$

subject to destination resource constraints

- Allocate as much bandwidth as possible while satisfying QoS demands

On-Demand Flow Placement

- 5G traffic is more dynamic compared to typical WAN workloads
- For scalability:
 - Greedily allocate new flows to paths with best demand satisfaction
 - Periodically run global optimization

OTTER Orchestrator

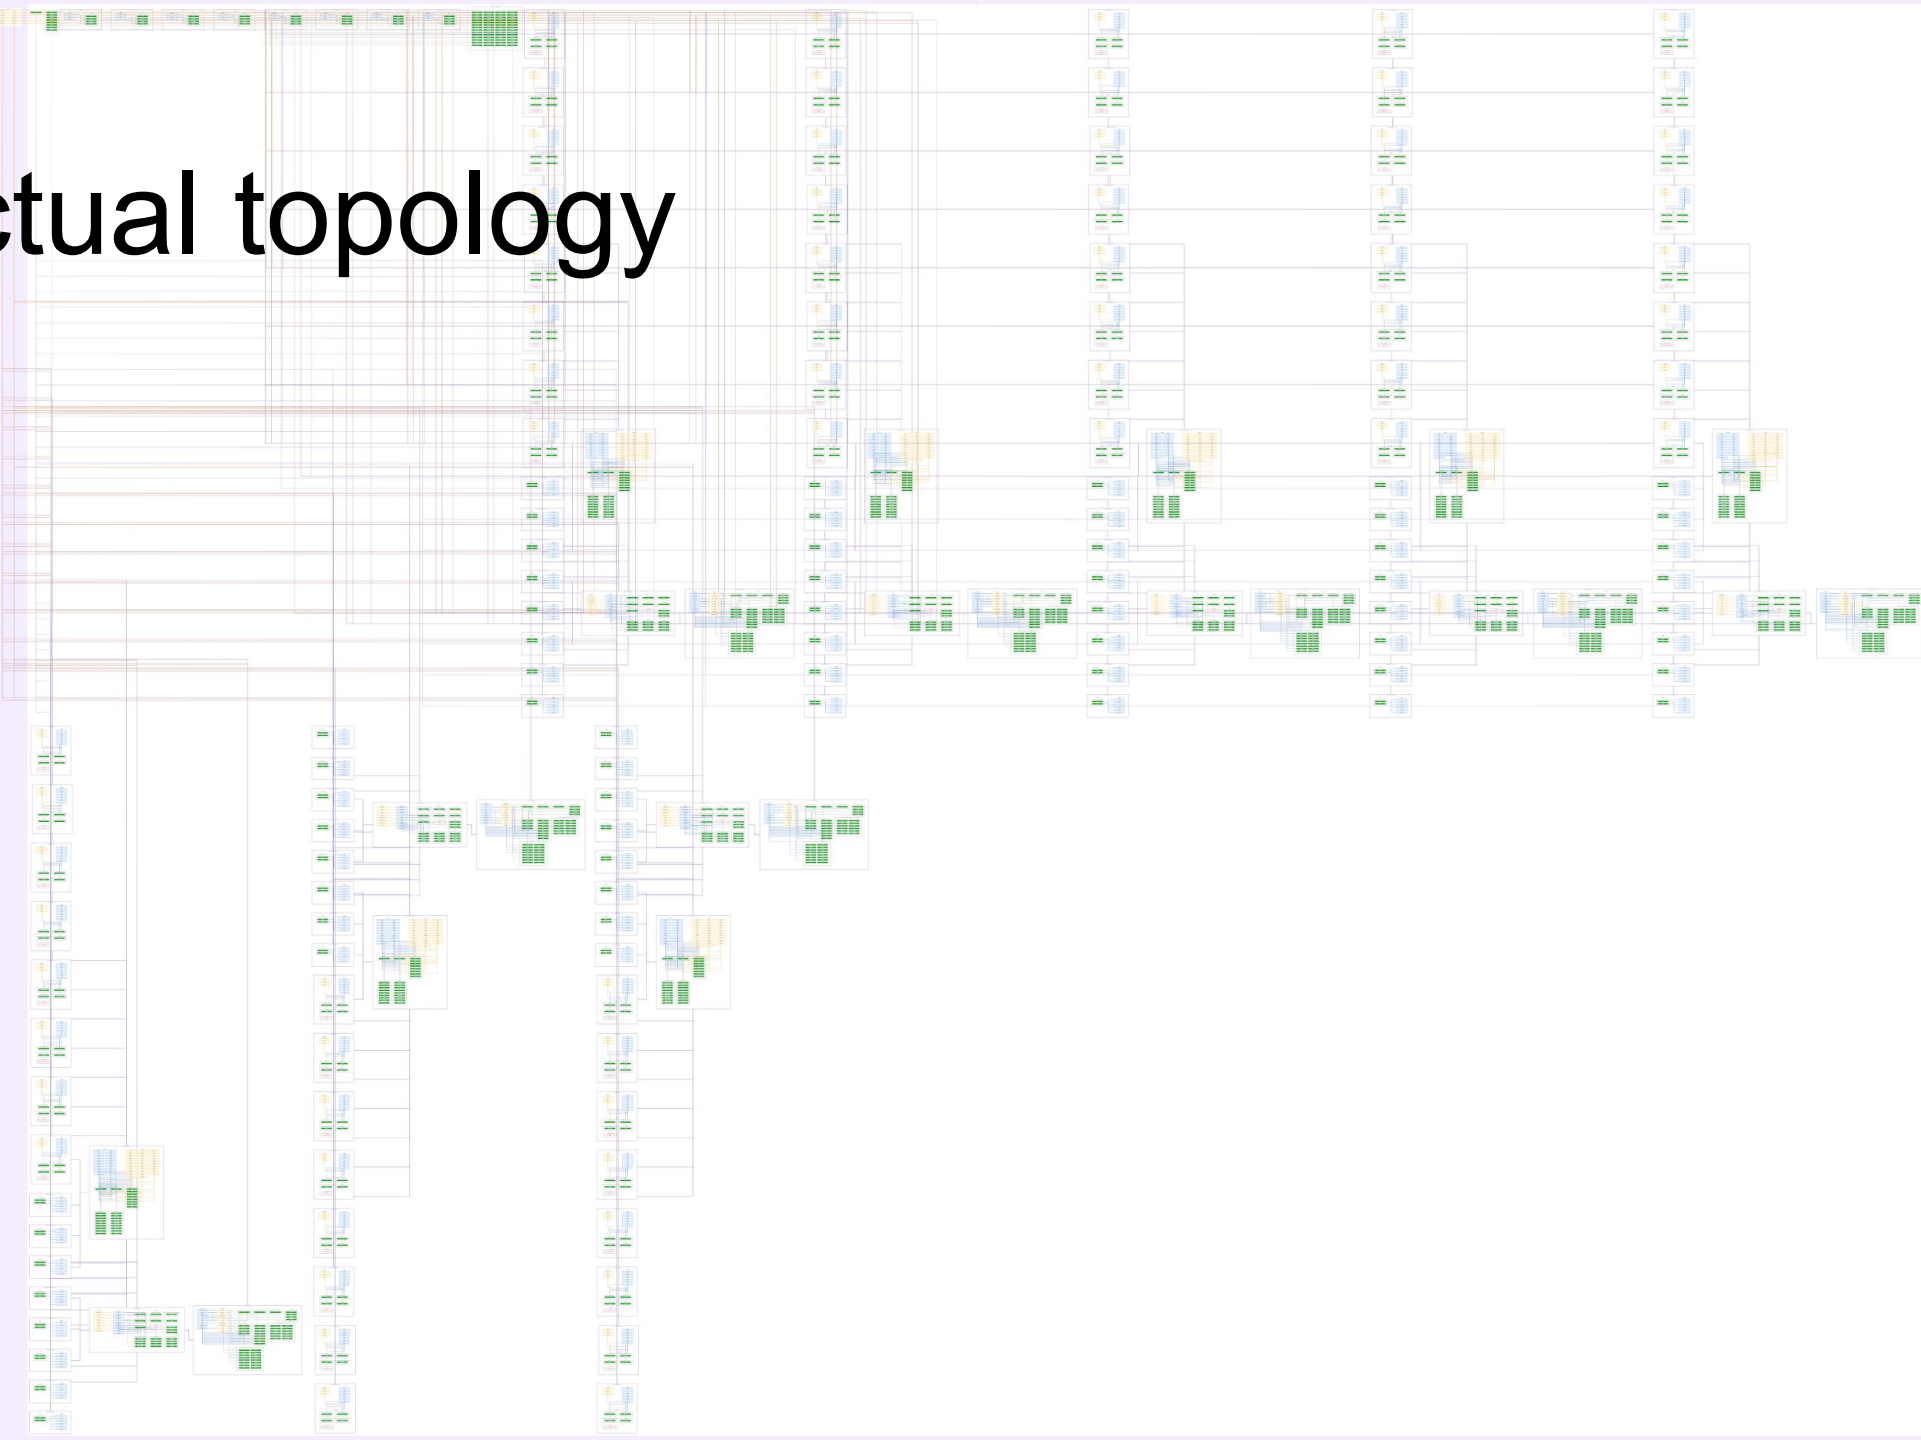
How to steer traffic along the routes computed by the Controller?

OTTER Orchestrator


How to steer traffic along the routes computed by the Controller?

- Solution: a multi-WAN overlay using native cloud functionality
 - No need for explicit cooperation between WANs
 - No need for significant engineering effort to exploit path diversity
 - Can easily scale-up and scale-out


Actual topology




Example QoS choices



Example QoS choices






Example QoS choices



- OTTER offers a choice among paths with different performance characteristics to satisfy QoS demands

QoS Improvement Results

- Throughput: +13% average, +136% best case (6-10Gbps higher)
- Latency: -15% average, -56% best case (21ms lower each way)
- Jitter: -45% average, -99% best case (5ms lower)
- Loss: -0.06% average, -0.4% best case (after removing outliers)
- 25-46% more bytes allocated with higher service demand satisfaction vs. greedy baselines



Summary

- 5G NR and cloudification are already here
- E2E performance is being bottlenecked on WAN paths
- **OTTER**
 - Co-optimizes paths and destinations for flows with diverse QoS demands
 - Deploys a scalable, multi-WAN overlay that unlocks inter-domain routing flexibility