
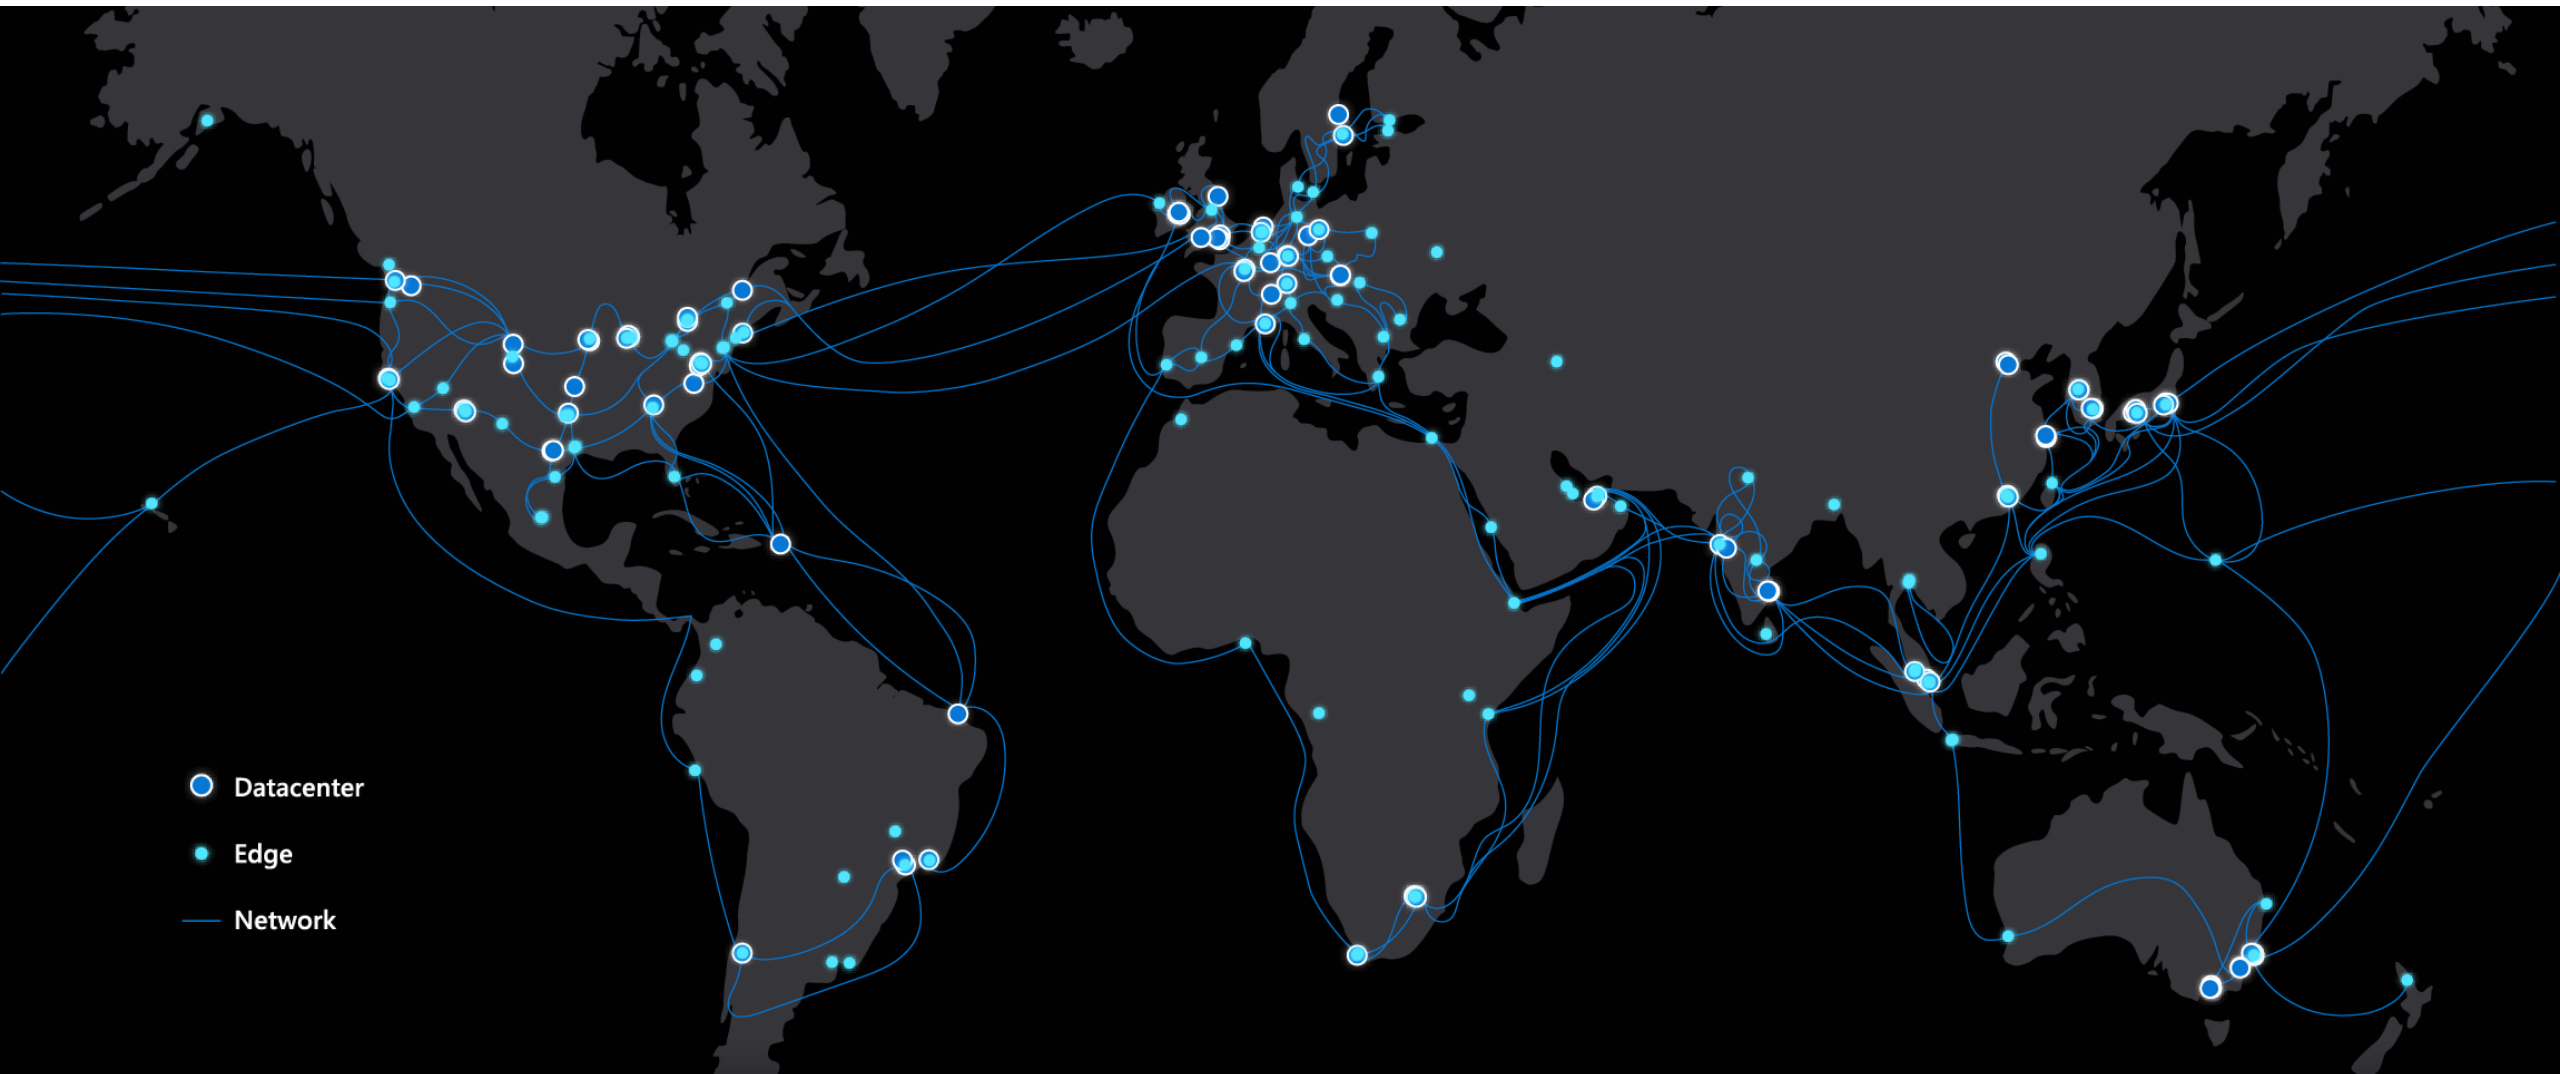


OneWAN is better than two: Unifying a split WAN architecture

Umesh Krishnaswamy*, Rachee Singh, Paul Mattes, Paul-Andre Bissonnette, Nikolaj Bjørner, Zahira Nasrin, Sonal Kothari, Prabhakar Reddy, John Abeln, Srikanth Kandula, Himanshu Raj, Luis Irun-Briz, Jamie Gaudette, Erica Lan



Microsoft global network



60+ Azure regions

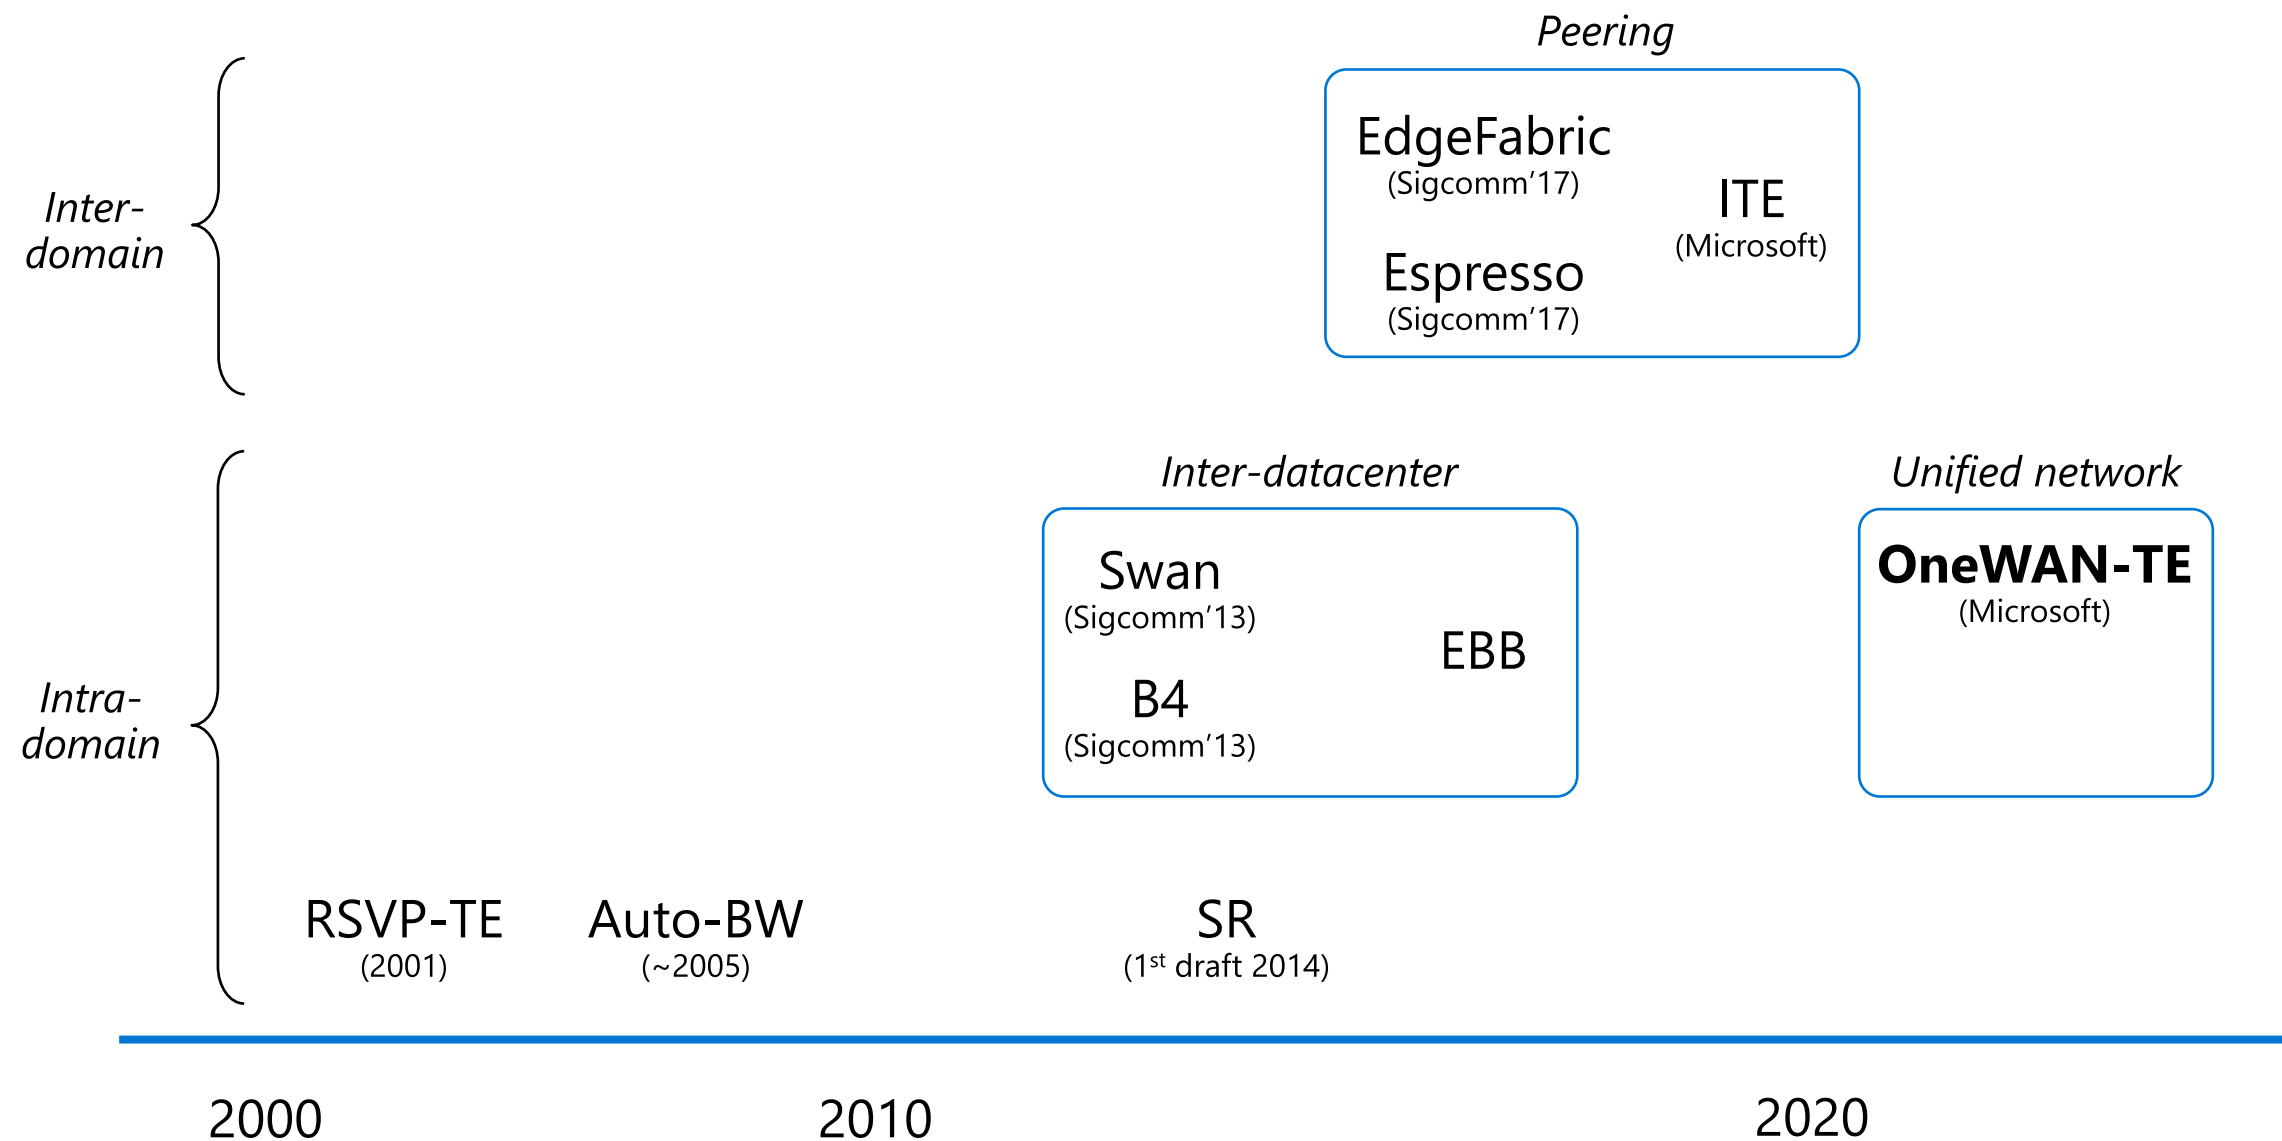
175k+ miles of fiber + subsea cables

190+ Network Edge PoPs


200+ Express Route partners

20k+ peering connections

Traffic Engineering



Two-WAN architecture



Two-WAN positives

Global TE, operate Swan at higher utilization, and time-shifting discretionary traffic lowered capex.

Core carried Internet flows with higher SLO. Swan evolved to serve critical customer workloads.

BlastShield (NSDI'22) slices enabled global TE without global blast radius.

| Traffic type | SLO |
|---|---------|
| Customer (e.g. Azure, Office, Teams) | 99.999% |
| Discretionary (e.g replication, backup) | 99.9% |





Two-WAN headwinds

Internet traffic growing faster since 2020.

RSVP-TE reaching scale limits in Core. Full mesh involves network-wide change.

Higher device and power costs inside regional gateway due to two WANs.

Capacity planning harder with two WANs. One network can be over-utilized while other is under-utilized.



New challenges with unified WAN

10x increase in route scale: $O(10^5) \rightarrow O(10^6)$.

10x increase in number of TE devices: $O(100) \rightarrow O(1000)$.

Match RSVP-TE fast reroute repair times.

Use existing hardware.

Hitless transition in live network.

1. Traffic steering
2. Local repair
3. TE optimization
4. Traffic matrices
5. Traffic migration

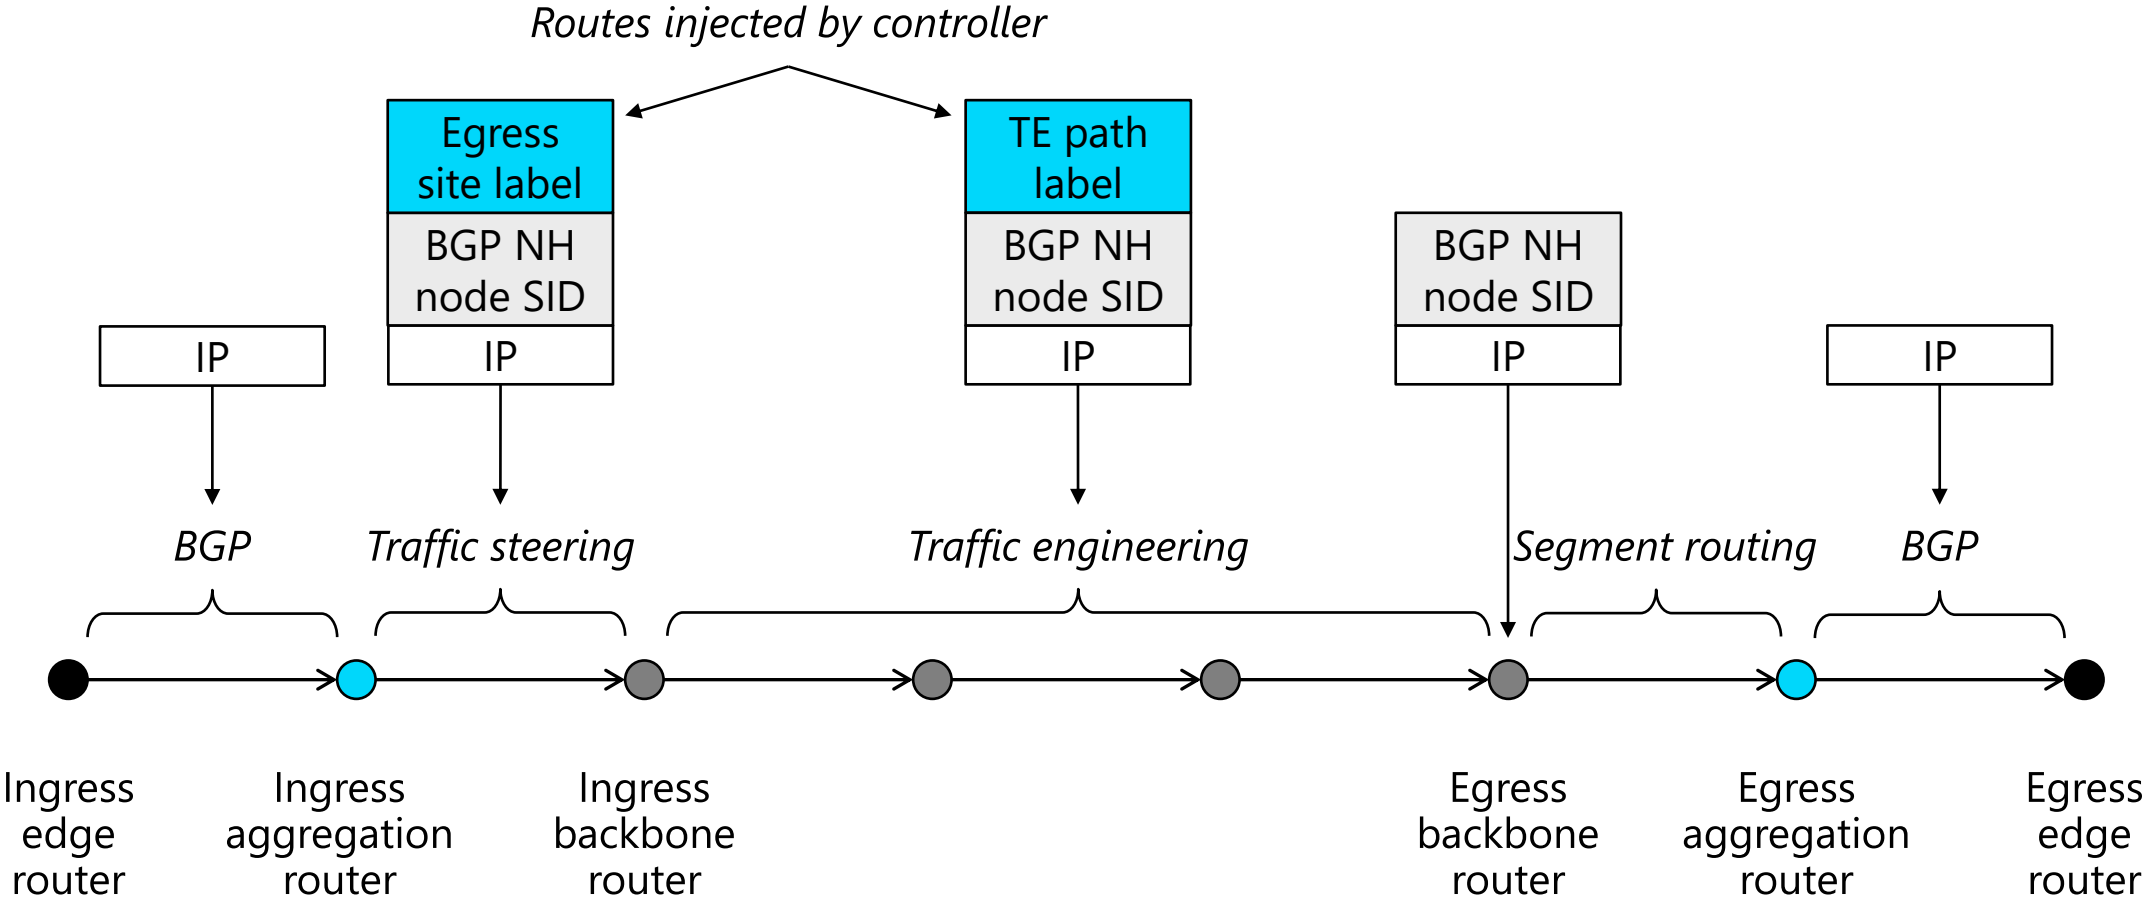
Why traffic steering?

The controller does not need to deal with route scale.

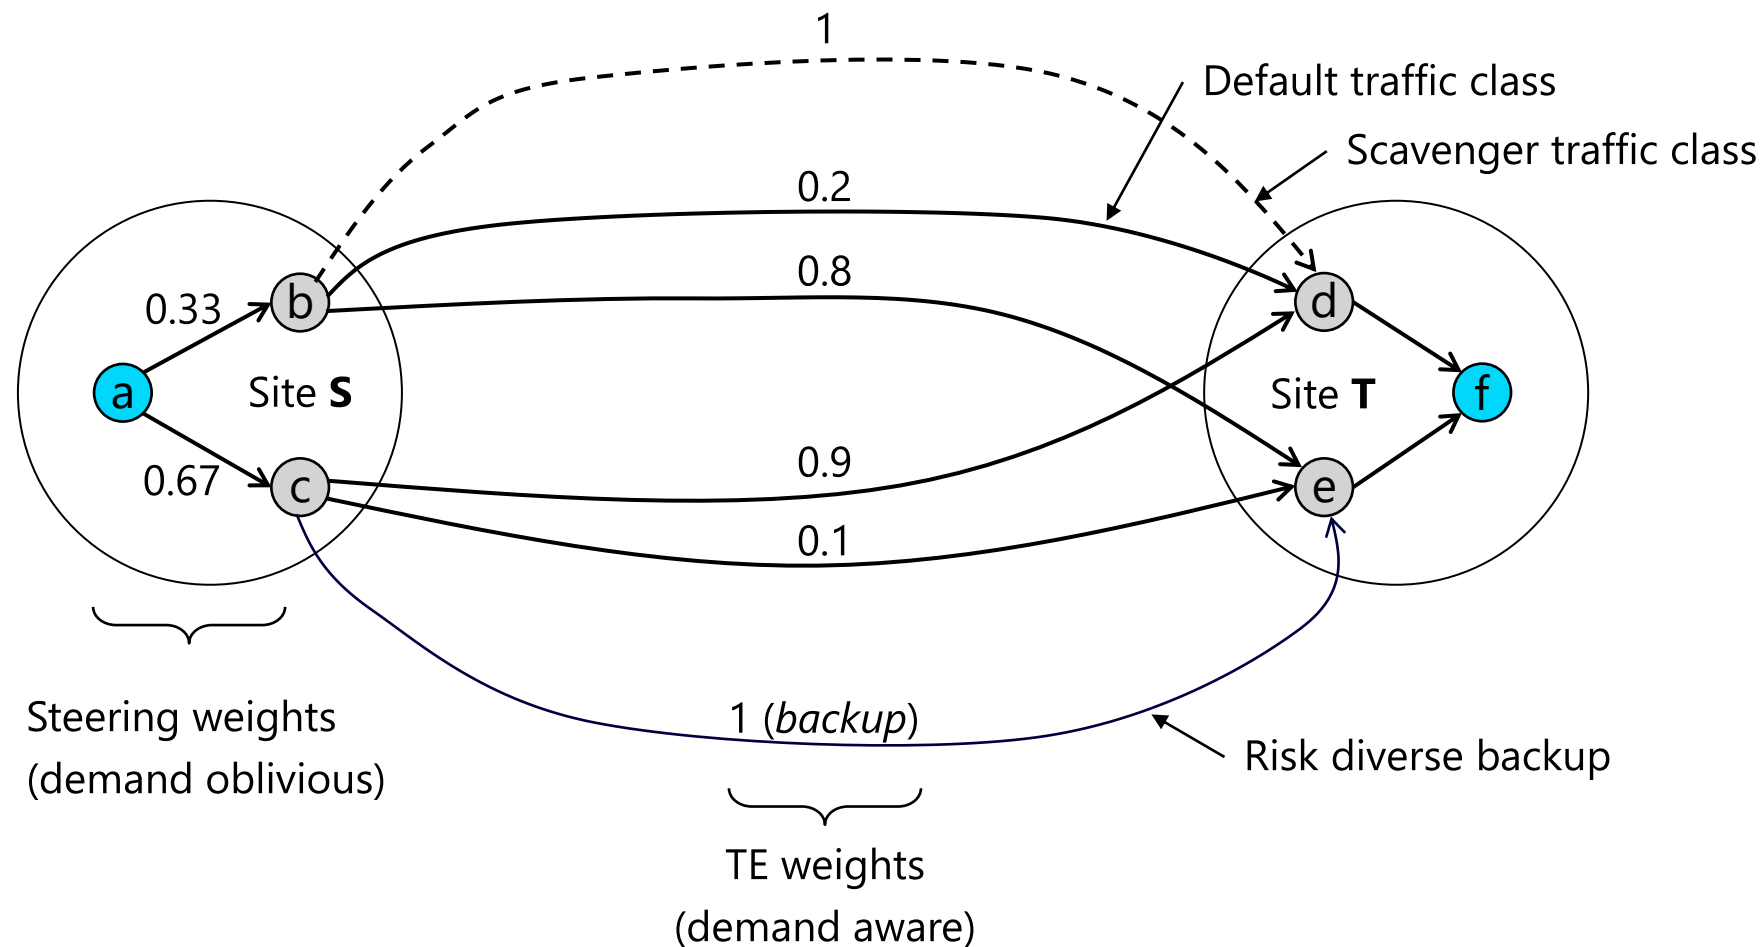
Use standard BGP to resolve prefixes to controller routes.

Only aggregation routers hold full Internet routes.

OneWAN routing in three parts



OneWAN-TE traffic splitting



1. Traffic steering
2. **Local repair**
3. TE optimization
4. Traffic matrices
5. Traffic migration


Why local repair?

Interactive applications need fast route convergence.


Agent can make pre-programmed route changes faster than a controller.

Local repair



Primary default




Backup default



Primary default after local repair




Unidirectional tunnel probing in Swan




Probes return to tunnel ingress using IS-IS routes.

Route convergence impact






Bidirectional tunnel probing in OneWAN




1. Traffic steering
2. Local repair
3. TE optimization
4. Traffic matrices
5. Traffic migration

Online path computation

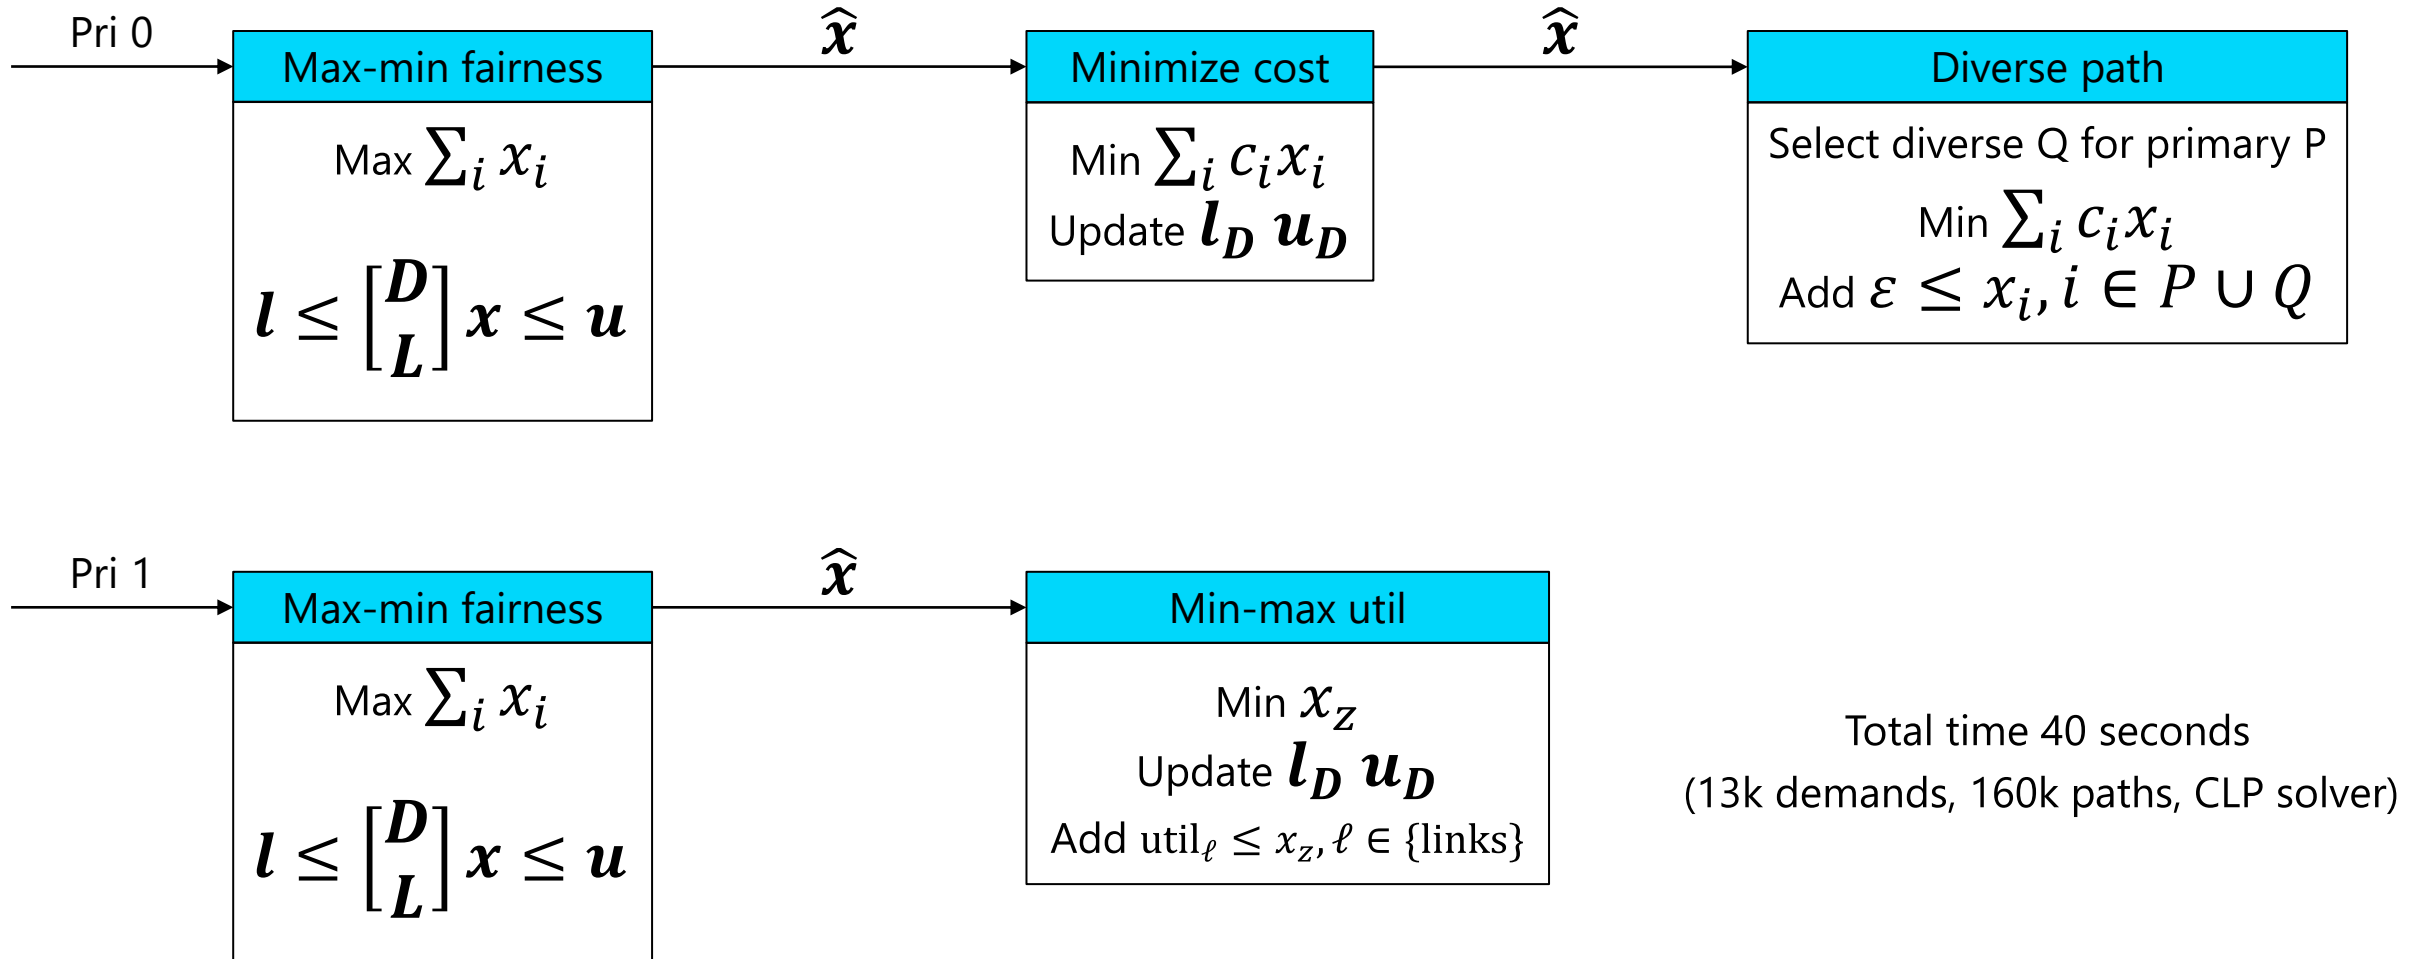


Why we do not use k-shortest paths




k -shortest path finds $s - u - t$ after 668 shorter paths.
Single commodity max-flow path finder is more efficient.

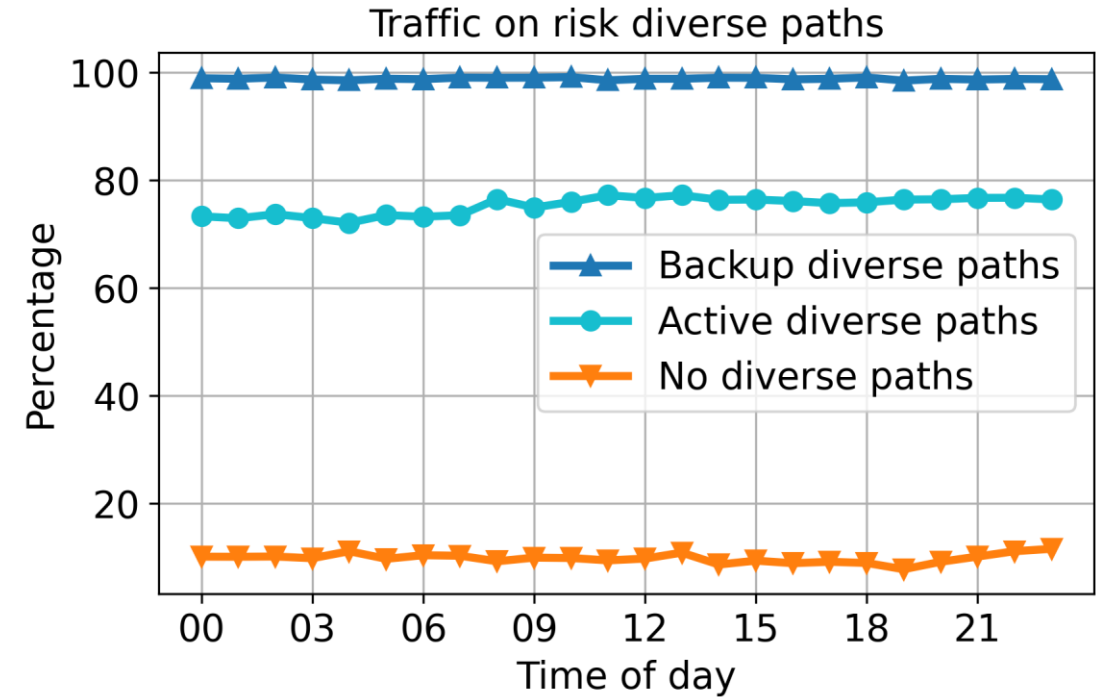
Priority fairness solver chaining



Diverse path solver



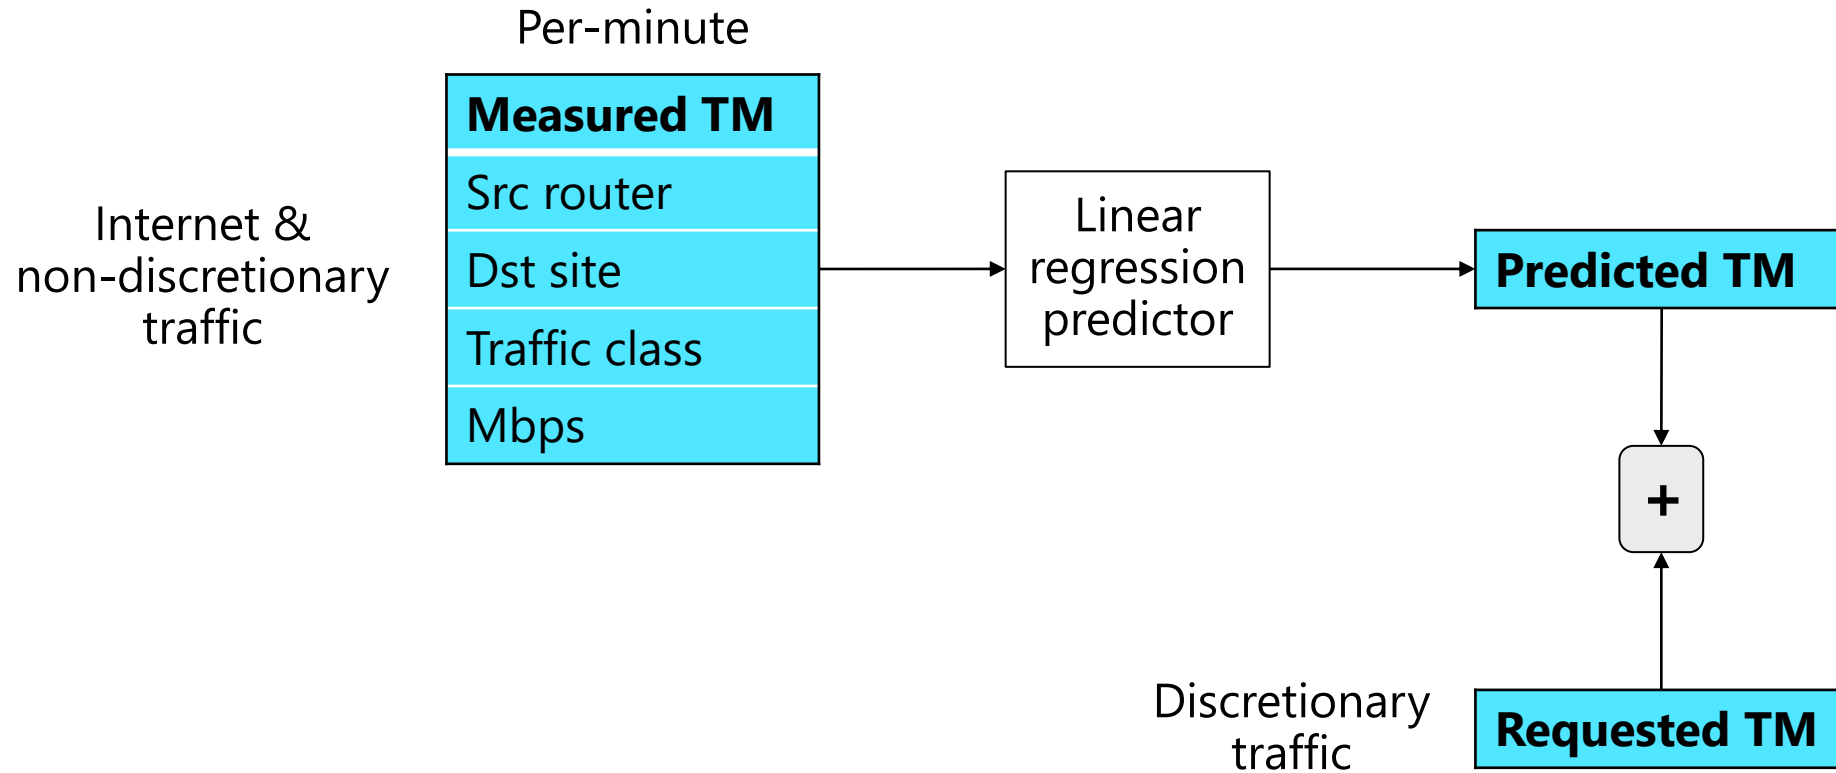
Path $s - u - t$ only protects risk $\{b\}$.
Penalizing path finder explores risk diverse paths.



Backup diverse paths do not carry traffic in non-fault cases.
Greater freedom to select diverse paths when backup.

1. Traffic steering
2. Local repair
3. TE optimization
- 4. Traffic matrices**
5. Traffic migration

Traffic matrix measured by flow sampling




Measuring anycast traffic


| |
|------------------------|
| Switch : a |
| Egress interface |
| BGP nexthop : b |
| ... |

Flow records


| |
|------------------------|
| Switch : c |
| Egress interface |
| BGP nexthop : d |
| ... |



Traffic matrix error correction


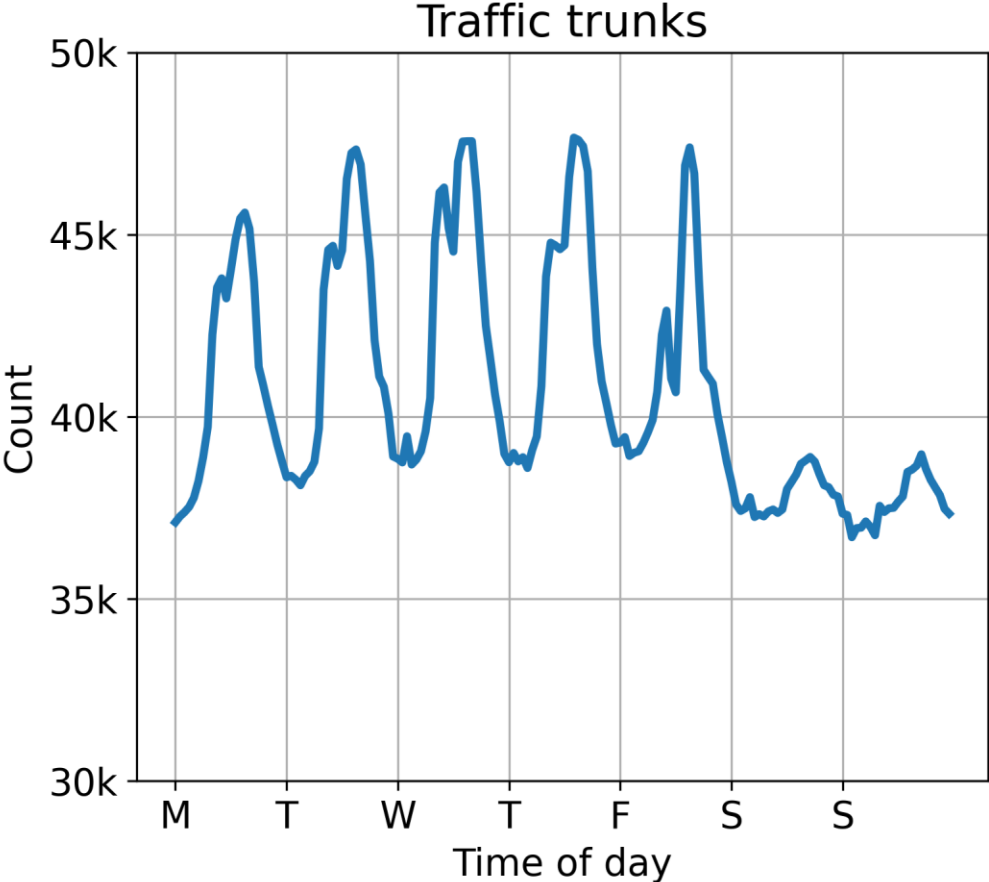


Interface counters correct for underprediction from data loss



TM operations correct for firmware limitations

OneWAN traffic matrix characteristics

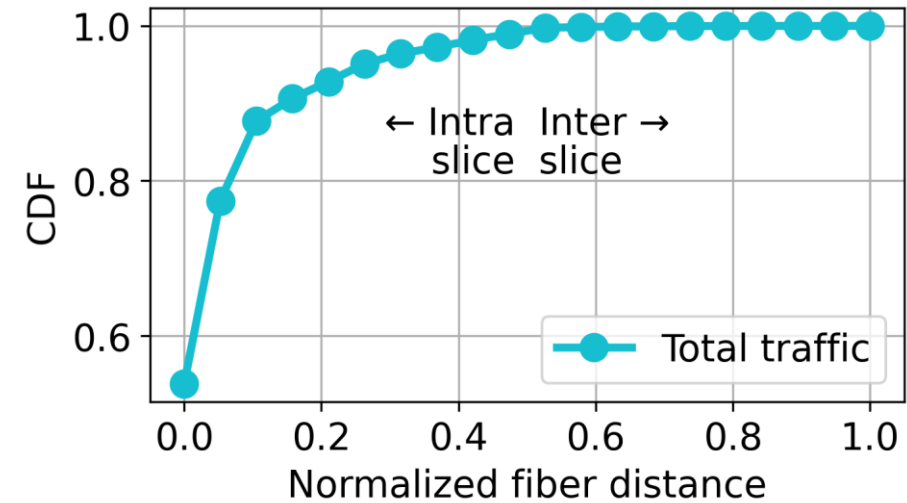


1. Traffic steering
2. Local repair
3. TE optimization
4. Traffic matrices
5. Traffic migration

Traffic migration




New BlastShield slices for OneWAN controllers.




Migrate intra-slice before inter-slice.

Traffic migration



Step down RSVP bandwidth reservation till zero.



Traffic migration from RSVP to OneWAN.

Summary

