

# SRv6 for Deterministic Path Placement in AI Backends

draft-filsfils-srv6ops-srv6-e2e-dc-frontend-wan-01

Clarence Filsfils (Cisco)

Chris Martin (Oracle Cloud Infra)

Kiran Sashidran (IBM Cloud)

**Pablo Camarillo (Cisco)**

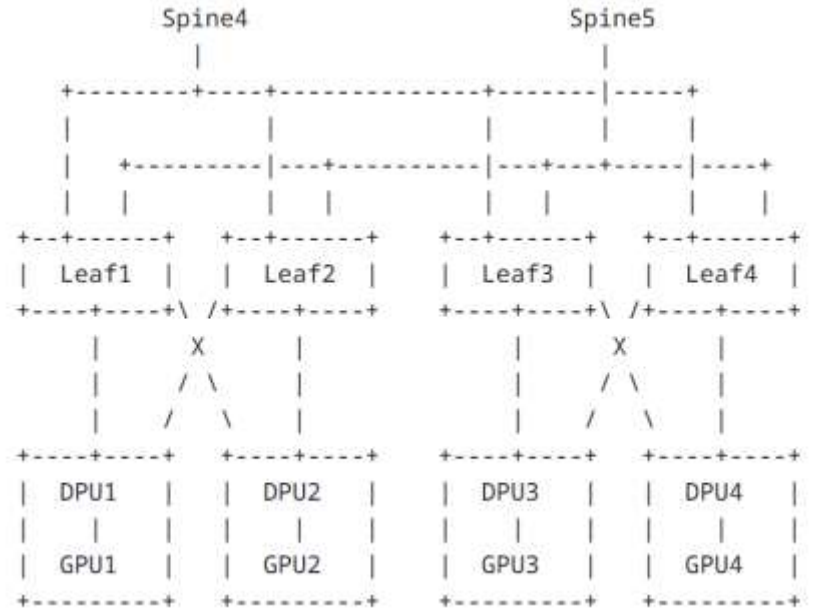
Ahmed Abdelsalam (Cisco)

# AI workloads

- Predictable Elephant Flow
- Synchronized Burst
- Low ECMP entropy
- Resilience

# SRv6 enabled backend

- Day0 provisioning of SRv6 uSIDs on Leaf and Spine devices
- At AI job orchestration, path computation and configuration on NIC.



# SRv6 uSID AI Backend

- Deterministic Path Placement
- Minimum MTU
- Congestion Feedback Loop
- Open Standard
- Ease of operation (no black-box)

# Microsoft

- Implemented in SONiC
- <https://www.segment-routing.net/conferences/Paris25-Microsoft-Rita-Hui/>

## 9. Acknowledgements

The authors would like to recognize the work of Lihua Yuan, Guohan Lu, Rita Hui, and Riff Jiang at Microsoft.

Rita Hui presented this use-case at MPLS & SRv6 World Congress in March 2025. A recording is available here: <https://www.segment-routing.net/conferences/Paris25-Microsoft-Rita-Hui/>

The authors would like to acknowledge the work of the developers who have enabled this use-case in the open-source [SONiC] implementation. In particular: Carmine Scarpitta, Abhishek Dosi, Changrong Wu, Kumares Perumal, Eddie Ruan, and Yuqing Zhao.

# More to come...

- Overlay Tenant Separation (job isolation)
- Cross-DC extension (inter-cluster AI training job)

# Next steps?

- WG Adoption of the document

Thank you