Global Load Balancing (GLB) for AI/ML Fabrics

Kevin Wang
Jeff Haas
Changwang Lin
Jeff Tantsura

Agenda

- Introduction to Global Load Balancing (GLB)
- GLB in 3-Clos Networks
- GLB in 5-Clos Networks
- GLB in General Clos Networks
- IETF Standardization Efforts
- GLB Test Results

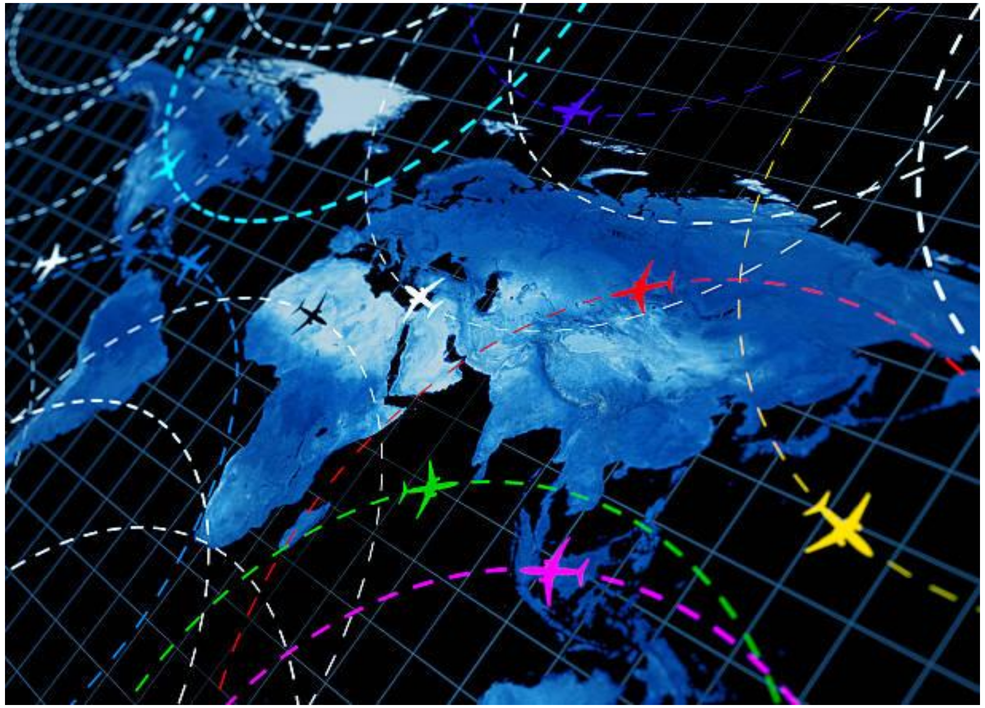

Key Facts about AI/ML Fabrics

- Clos topology
- IP/IPv6 routing
- BGP as the routing protocol
- Packet delay and loss are very costly for AI/ML flows
- High throughput, low entropy
- Congestion avoidance and mitigation is essential

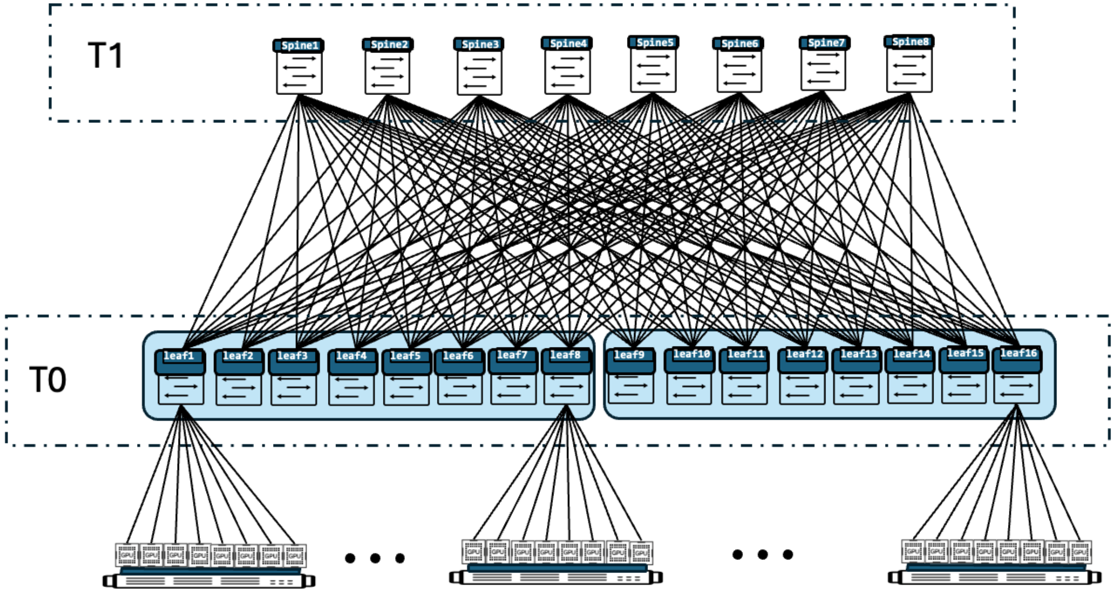

Local Congestion: Dynamic Load Balancing

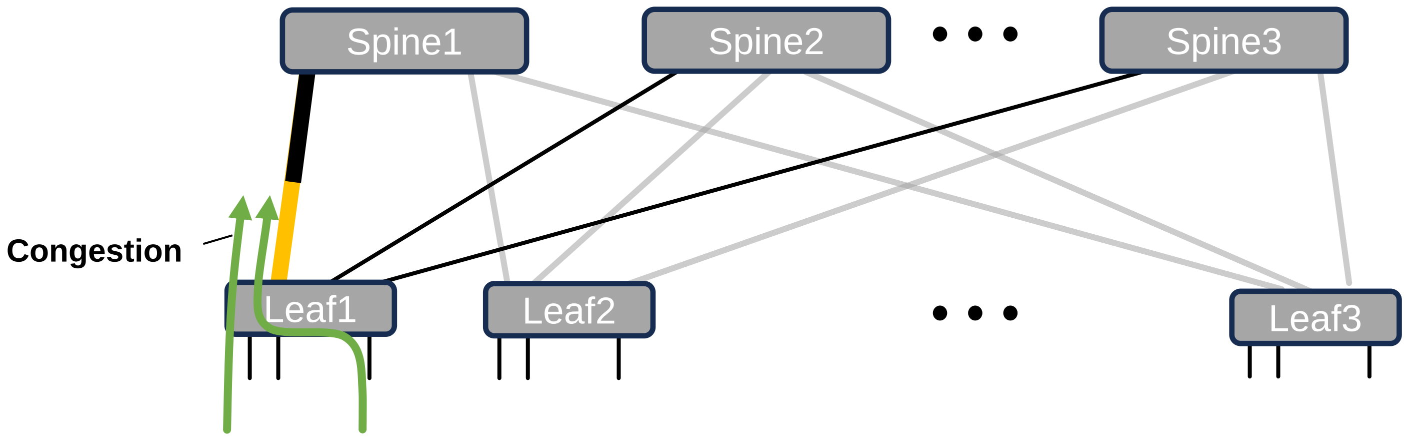

- Random ECMP hashing can’t guarantee each link has the same load
- When congestion is detected locally (Leaf1 to Spine1), Dynamic Load Balancing (DLB) can help to mitigate the congestion

Remote Congestion: Global Load Balancing

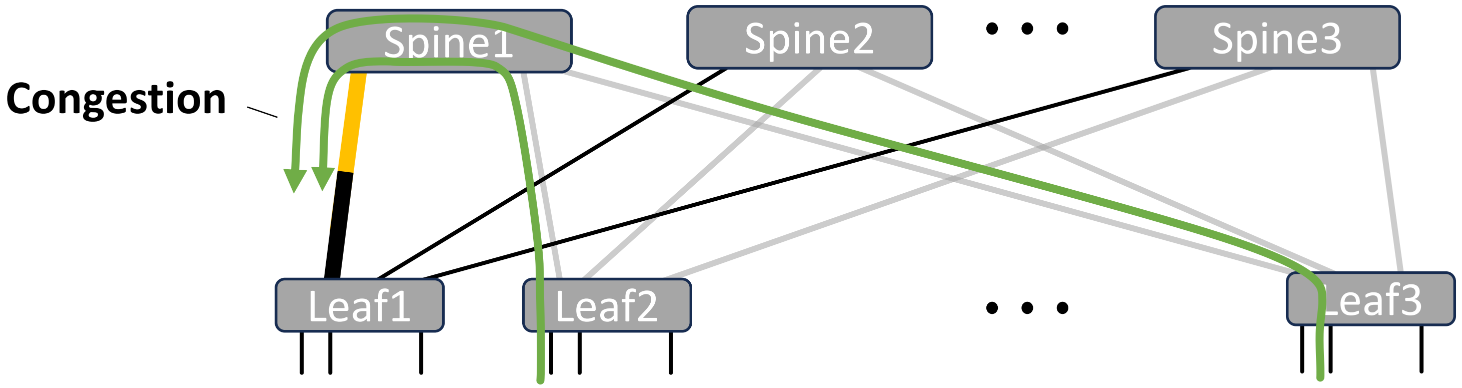

- When Congestion happens on a remote link (Spine1 to Leaf1), DLB on Leaf2 and Leaf3 does not help
- Global Load Balancing (GLB) is needed to mitigate remote congestion

GLB in 3-Clos: Topology Discovery

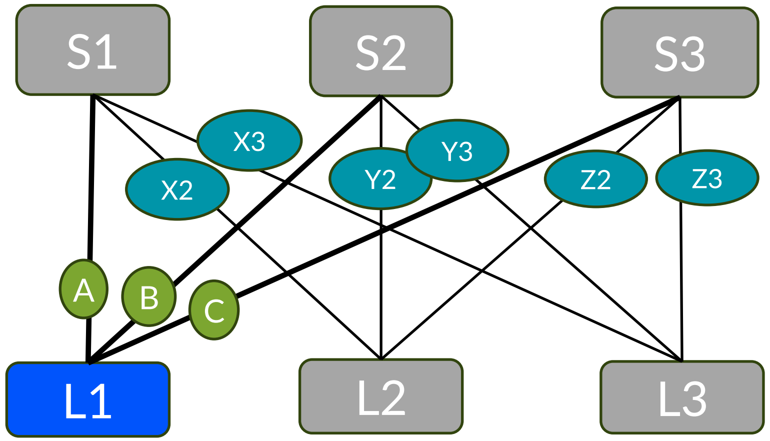

Routes on L1		
Destination	**Next-hop**	**Next-next-hop**
L2	A	X2
	B	Y2
	C	Z2
L3	A	X3
	B	Y3
	C	Z3

- For routes received from L2, L3, L1 keeps track of the <next-hop, next-next-hop> tuple for each path
- BGP protocols is extended to carry the next-next-hop nodes of each next-hop

GLB in 3-Clos: Link Quality Exchange

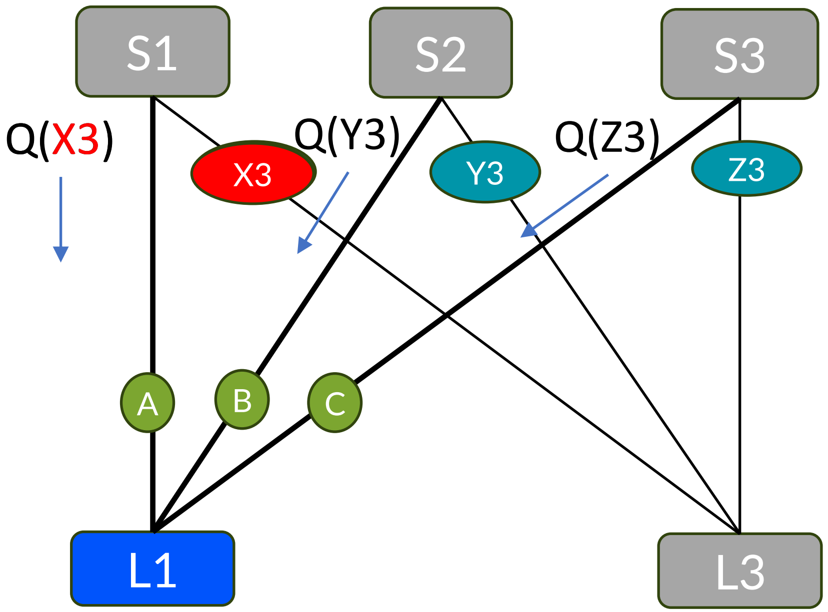

- Each spine broadcast its link qualities to all leaves
- L1 uses the combined link qualities <Q(A), Q(X3)>, <Q(B), Q(Y3)>, <Q(C), Q(Z3)> to make load balancing decisions
- When link X3 is congested, L1 learns it via Q(X3) and reduces traffic to link A to mitigate the congestion on X3
- The Router Information protocol is used to exchange link qualities between neighboring routers within sub-millisecond

GLB in 5-Clos: Leaves

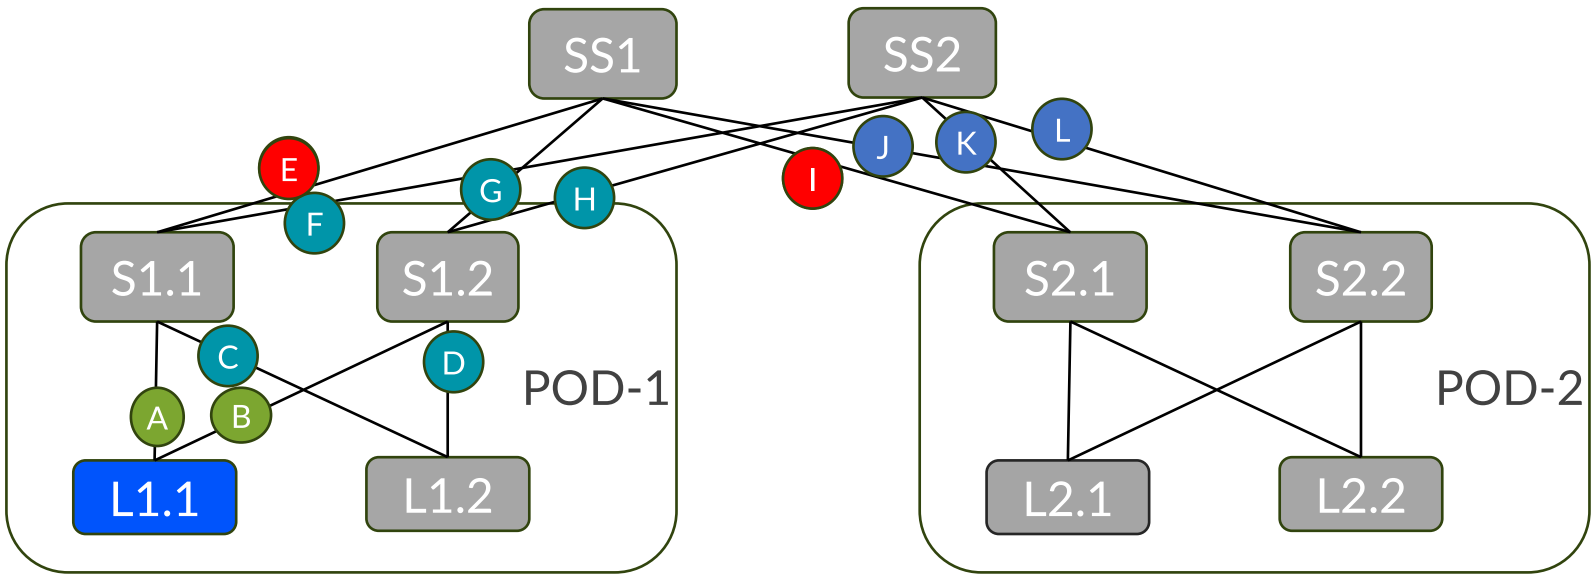

Routes on L1.1		
Destination	**Next-hop**	**Next-next-hop**
L1.2	A	C
	B	D
L2.1	A	E, F
	B	G, H
L2.2	A	E, F
	B	G, H

- Intra-pod GLB is same as 3-Clos GLB
- Inter-pod routes have more than 2 hops, but only <next-hop, next-nex-hop> needs to be kept for each route
- When L1.1 knows link E is congested, it will reduce inter-pod traffic towards link A
- On the other hand, L1.1 does not care about congestion on link I because GLB on S1.1 and S1.2 reacts to it equally. There is no need for L1.1 to prefer S1.1 or S1.2

GLB in 5-Clos: Spines

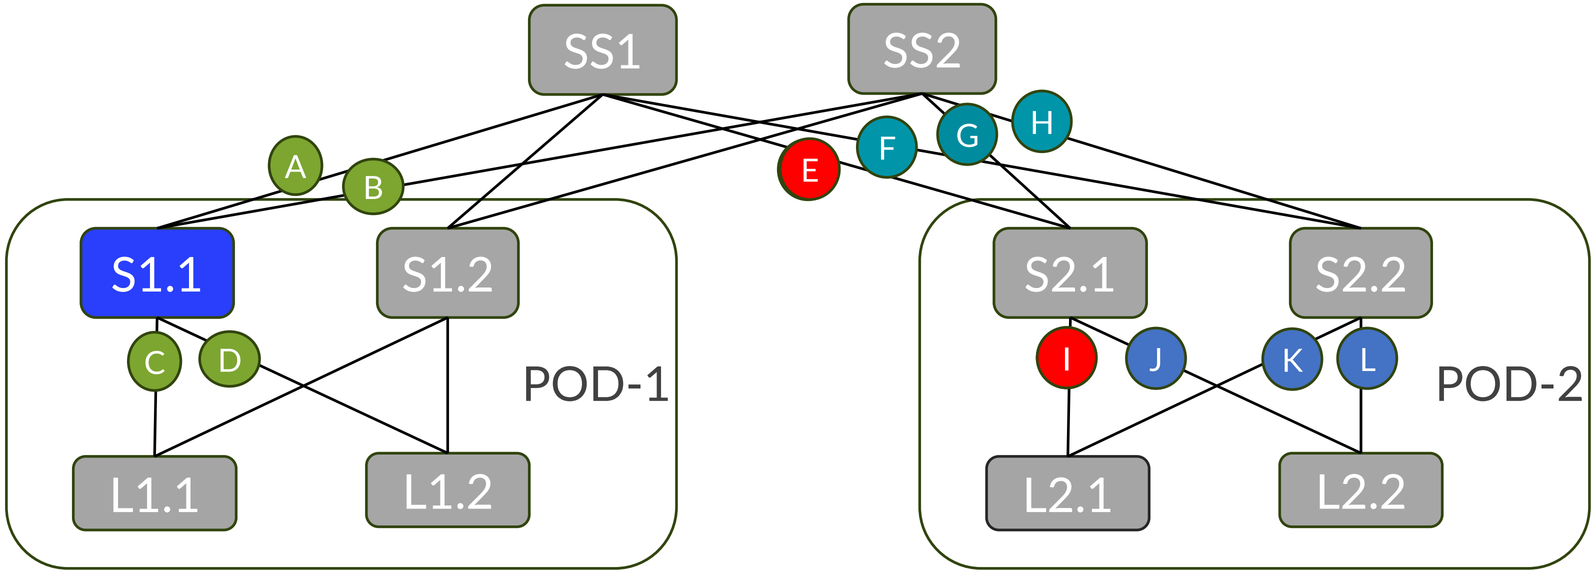

Routes on S1.1		
Destination	**Next-hop**	**Next-next-hop**
L2.1	A	E, F
	B	G, H
L2.2	A	E, F
	B	G, H
L1.1	C	-
L1.2	D	-

- Spines also need to run GLB for inter-pod traffic
- When S1.1 knows link E is congested, it reduces traffic towards link A
- On the other hand, S1.1 does not care about the congestion on link I because GLB on SS1 and SS2 reacts to it equally. There is no need for S1.1 to prefer SS1 or SS2

GLB in General Clos Networks

- The symmetry and rich redundancy of Clos networks ensure that only next-hop or next-next-hop link congestion impacts load balancing decisions
- Each node tracks only the <next-hop, next-next-hop> for each route
- Congestion beyond the next-next-hop has no impact on GLB
- This allows non-link-state protocols like BGP to determine topology for GLB
- It also allows link quality exchange to be limited to direct neighbors
- The design works for any levels of Clos networks

IETF Standardization: IDR

GLB needs to discover next-hop to next-next-hop mapping via BGP

- https://datatracker.ietf.org/doc/draft-ietf-idr-next-next-hop-nodes/ defines NNHN TLV for the NHC attribute:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Capability Code = 2       |  Capability Length (variable) |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Next-hop BGP ID                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
~                   Next-next-hop BGP IDs (variable)            ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

IETF Standardization: RTGWG

GLB also needs neighboring nodes to exchange link quality information

- https://www.ietf.org/archive/id/draft-zzhang-rtgwg-router-info-05.txt defines a UDP based protocols for neighboring routers to exchange router information TLVs
- The Neighbor Path Information TLV is used for GLB

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   (Type) 1    |               Length          |S| Refresh Rate|
+-------------------------------+-------------------------------+
~                   repeated per-Neighbor Records               ~
+---------------------------------------------------------------+
```

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Neighbor ID                          |
+---------------+---------------+-------------------------------+
| encoded info  |
+---------------+
```

Test Results for 3-Clos (100G Links)

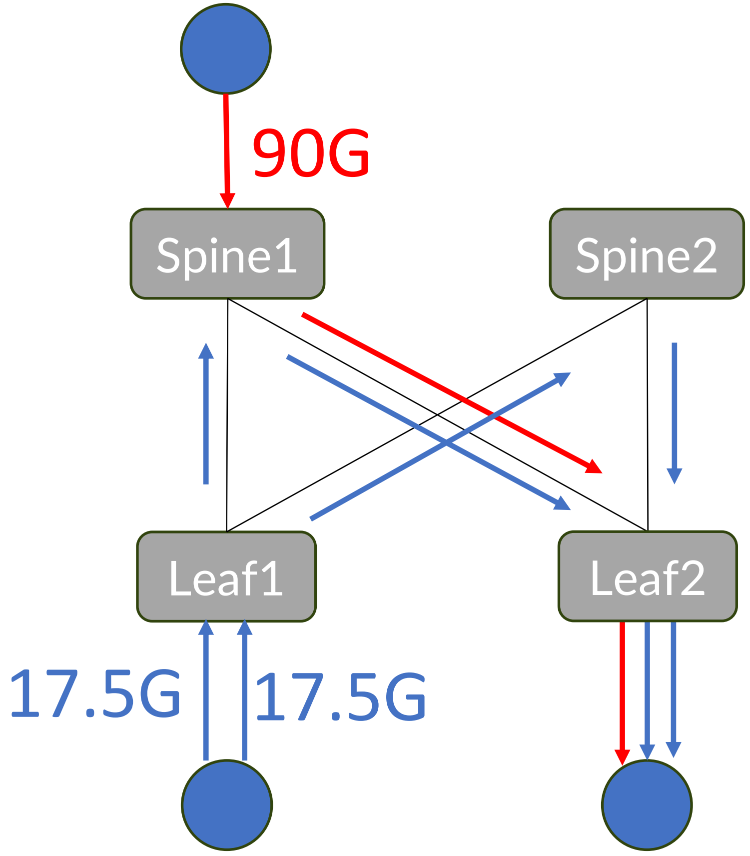

Without GLB: drops happened on Spine1 => Leaf2

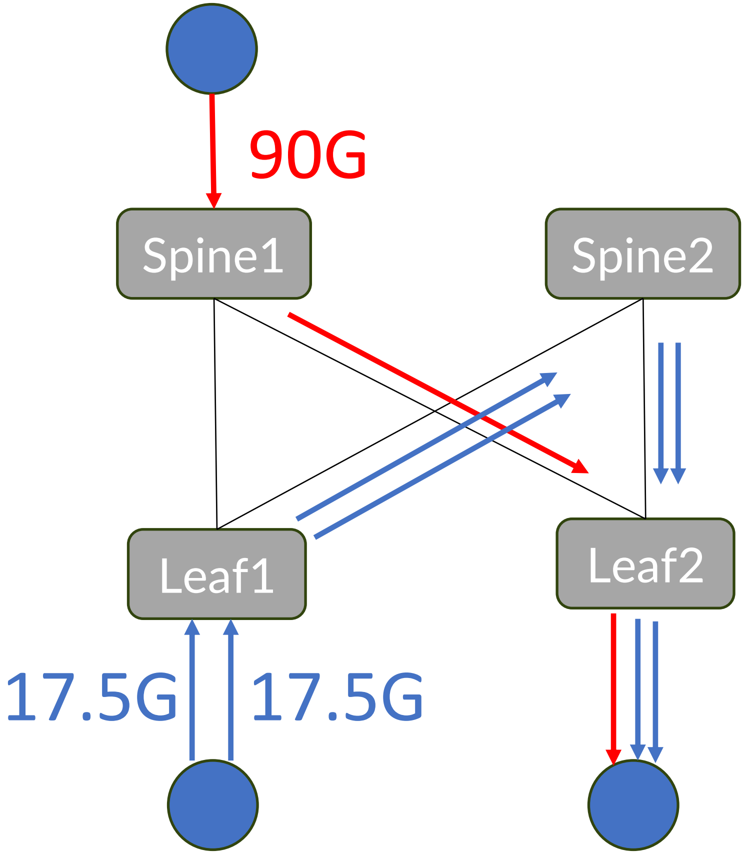

With GLB: no drop

Summary

AI/ML flows are prone to congestion but are less tolerant of it

GLB is designed to prevent and mitigate congestions for AI/ML flows

GLB keeps the simplicity of the existing AI/ML fabrics

- Clos networks
- IP/IPv6 routing
- BGP as the routing protocol

The scalable design enables sub-millisecond level reaction to potential congestions

Thank you!