



Enabling WCMP in SONiC using PINS and ONOS



Niloofar Toorchi



Daniele Moro



Don Newton



Brian O'Connor



Pier Luigi Ventre

PINS = SONiC + SDN


- Software-Defined Network (SDN) brings many features and capabilities to the network which is difficult to achieve using traditional embedded control planes
- P4 Integrated Network Stack (PINS) adds SDN capabilities to SONiC, which is a widely deployed network OS
- New SDN use cases: **Weighted Cost MultiPath (WCMP)**



Weighted Cost MultiPath (WCMP)

WCMP distributes traffic flows on multiple links proportionally to the assigned weights

- Helps optimally distribute traffic in unbalanced networks
- Watch Google's presentation on WCMP for more information here:



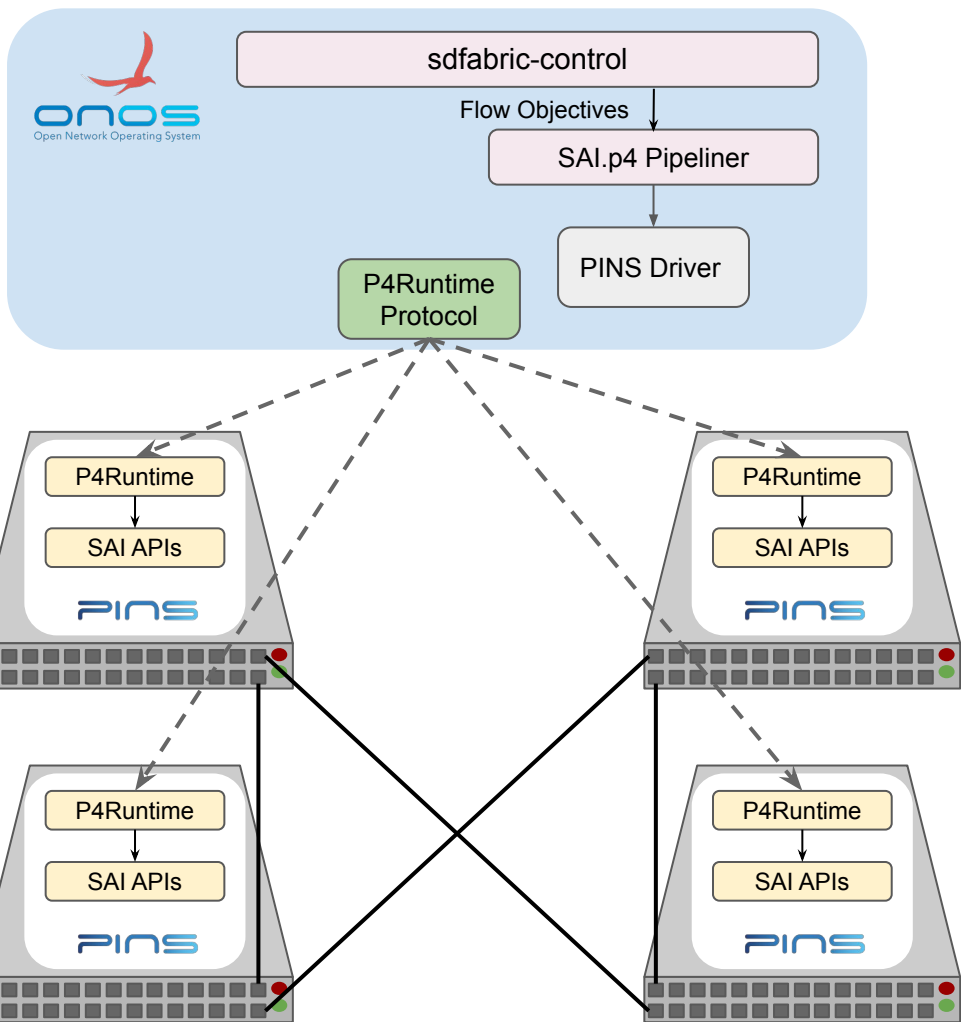
We implement WCMP using the open source **ONOS** and **SD-Fabric** platforms.



End-to-End Architecture

SD-Fabric is a P4 programmable network fabric that includes:

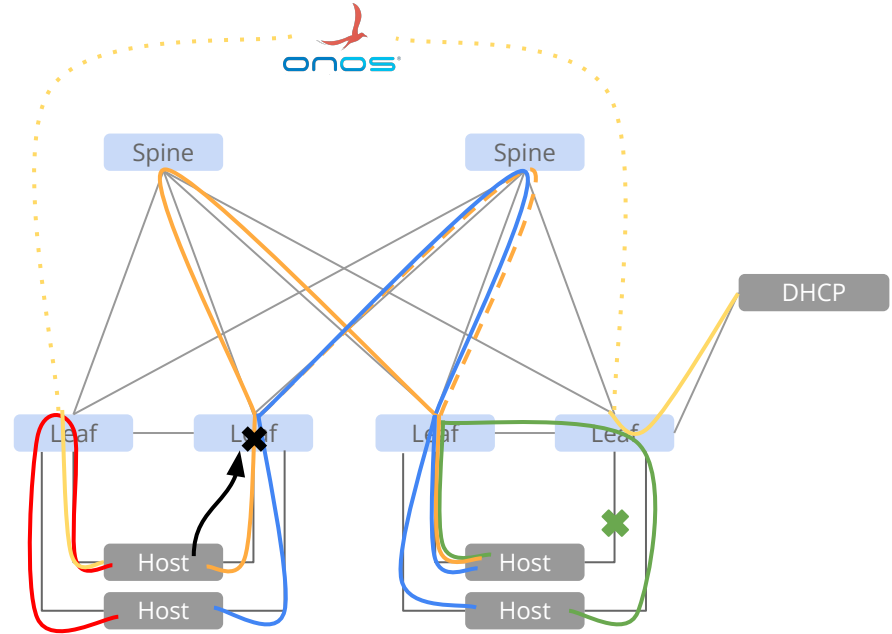
- SDN Controller
 - ONOS
 - sdfabric-control application
- Leaf-spine fabric of programmable switches
 - P4 Integrated Network Stack (PINS)
 - sai.p4
 - P4 entities to SAI APIs call
 - Remote Packet I/O



SD-Fabric Overview

Traditional Fabric Features

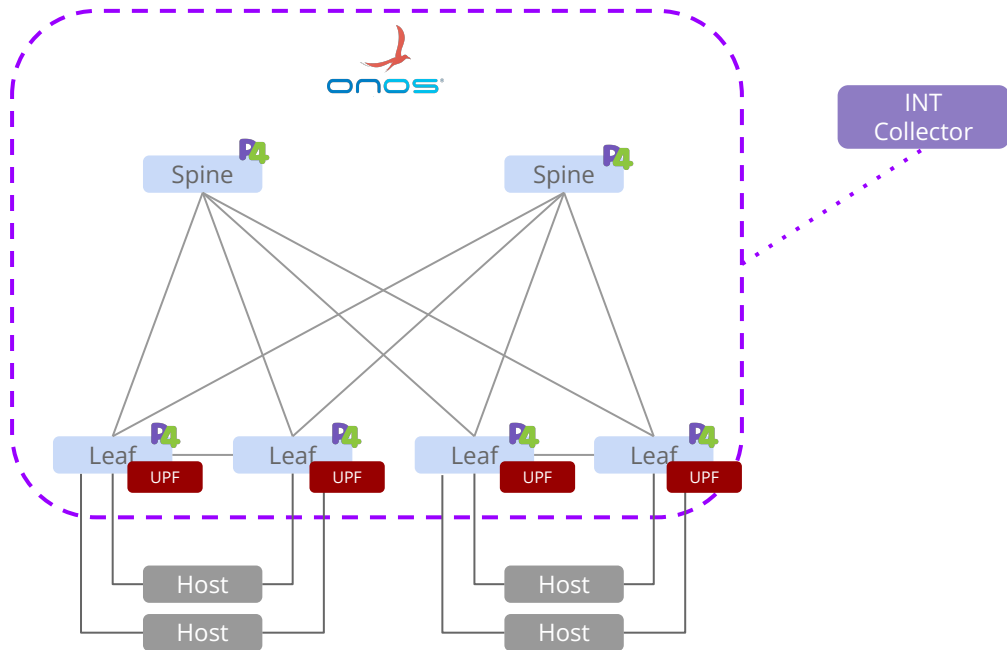
- Bridging (L2) at the Leaf
- Equal Cost MultiPath (ECMP) routing (L3) across the fabric
- IPv4/IPv6 unicast/multicast
- Dual-homing
- DHCP relay
- Access Control List (ACL) to drop or redirect traffic



SD-Fabric Overview

Advanced Features

- P4 Forwarding Pipeline (fabric-tna.p4)
- In-Band Network Telemetry (INT) support
- Network Function Embedding
 - bng.p4, upf.p4



SD-Fabric for PINS with WCMP

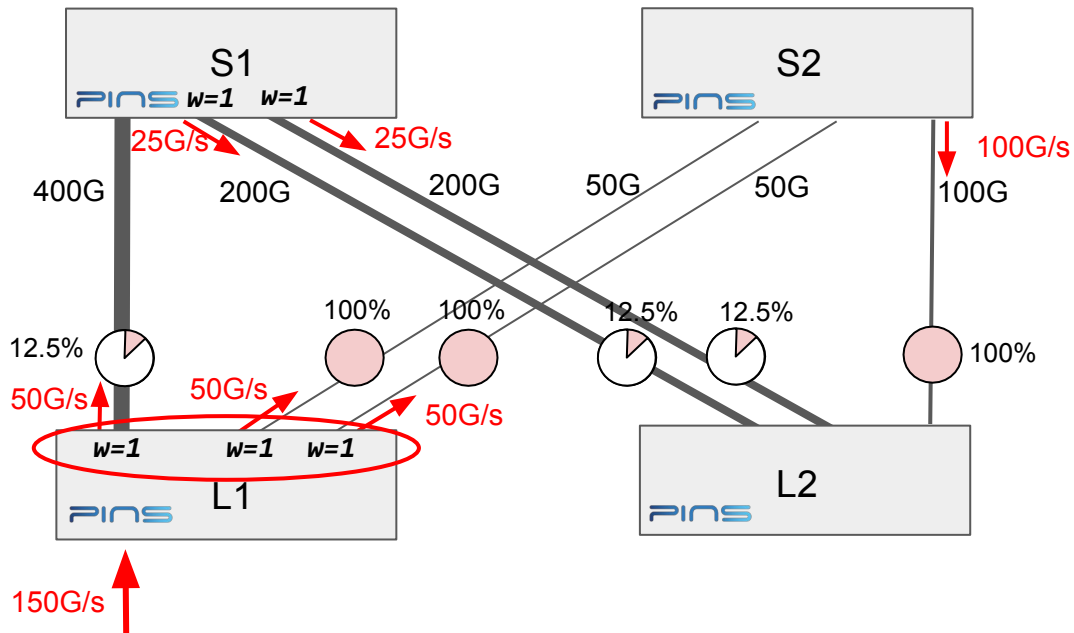
- ONOS Pipeliner adapted to work with SAI.p4
 - Use pure L3 routing also on the spines instead of MultiProtocol Label Switching (MPLS)
 - Considering MPLS and Segment Routing v6 (SRv6) for the future
- Added a WCMP handler to SD-Fabric
 - New module, periodically calculates and adjusts weights of the hash groups (next hops) based on the current topology
 - Trivial algorithm: weight proportional to the link speed
 - The algorithm can be extended to find the bucket weights according to different metrics such as Quality of Service (QoS) metrics, packet drop rate, etc.
- Embed Network Function (*future*)
 - Considering SAI generic extension mechanism to extend pipeline functionality



WCMP vs ECMP

ECMP Weight Distribution Example

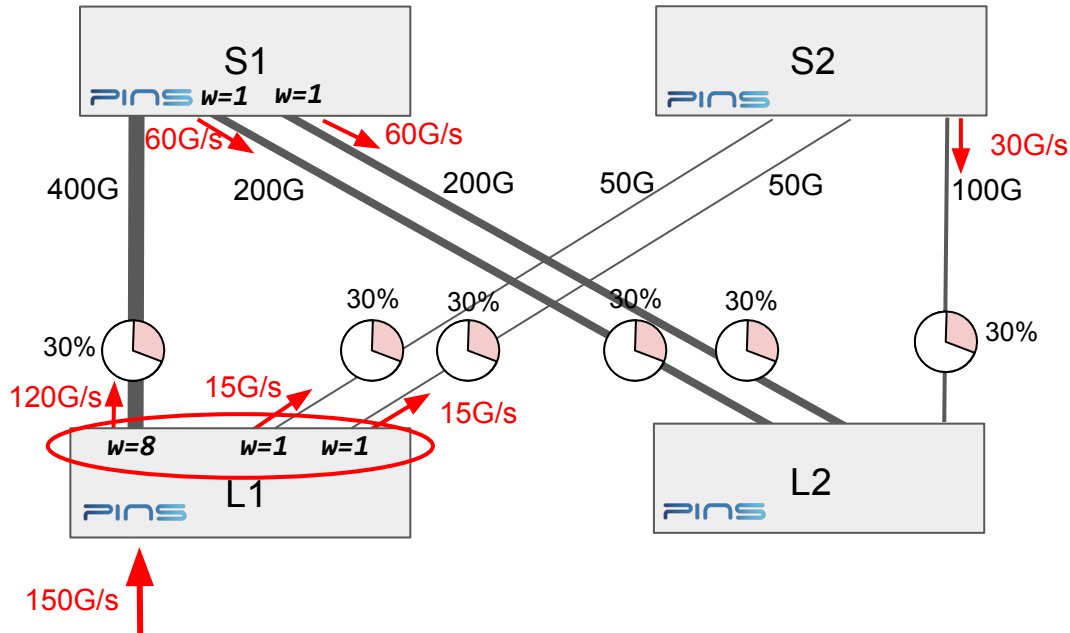
- ECMP equally distributes the traffic through the links
- Not optimal with unbalanced networks



WCMP vs ECMP

WCMP Weight Distribution Example


- WCMP distributes the traffic through the links according to their **weight**
- Can reach optimality with unbalanced networks



WCMP vs ECMP

WCMP Weight Distribution Example

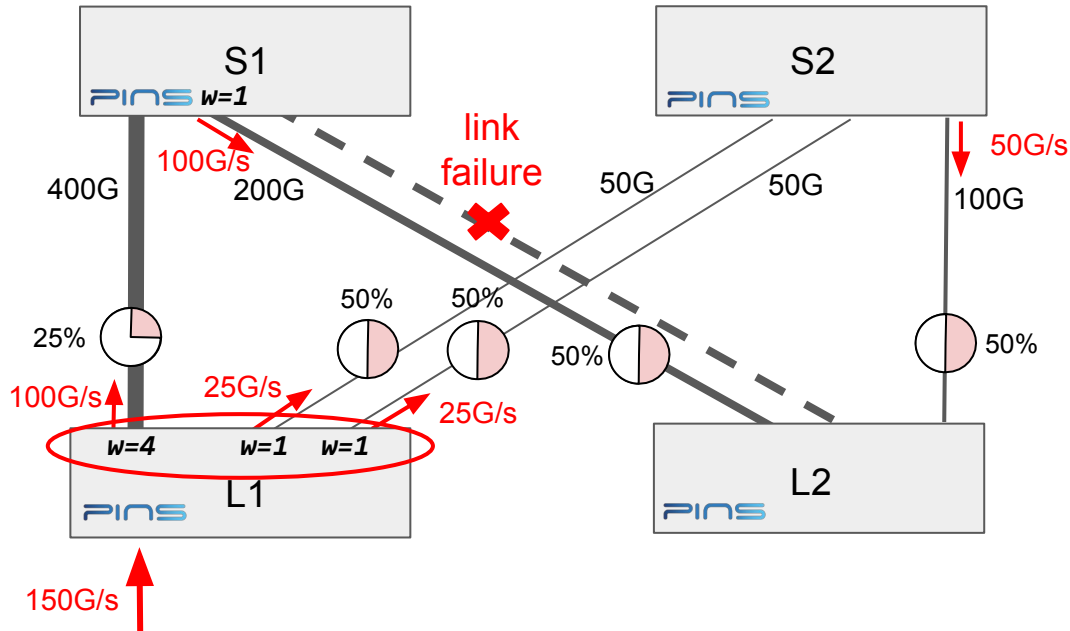
- WCMP distributes the traffic through the links according to their **weight**
- Can reach optimality with unbalanced networks



WCMP vs ECMP

WCMP Weight Distribution Example

- WCMP distributes the traffic through the links according to their **weight**
- Can reach optimality with unbalanced networks



DEMO

In Summary

- PINS brings SDN capabilities to SONiC
- SD-Fabric and ONOS are used to program routing tables with WCMP
- WCMP helps in optimally exploit unbalanced networks
- The SDN controller can dynamically adapts weights based on network conditions

Notices & Disclaimers

- Intel technologies may require enabled hardware, software or service activation.
- No product or component can be absolutely secure.
- Your costs and results may vary.
- © Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.



Thank You!

https://github.com/sonic-net/SONiC/blob/master/doc/pins/pins_hld.md

<https://github.com/Azure/sonic-pins>

<https://github.com/pins/tutorials>