

RESILIENCE ISSUES WITH SDN - A MATTER OF TOPOLOGY

ggrammel@juniper.net

JUNIPER
NETWORKS

Engineering
Simplicity

AGENDA

- SDN from the Datacenter to the WAN
- Difference in Network Architecture
- Resilience
- Summary

INTRODUCTION

WHAT SCENARIO ARE WE LOOKING AT?

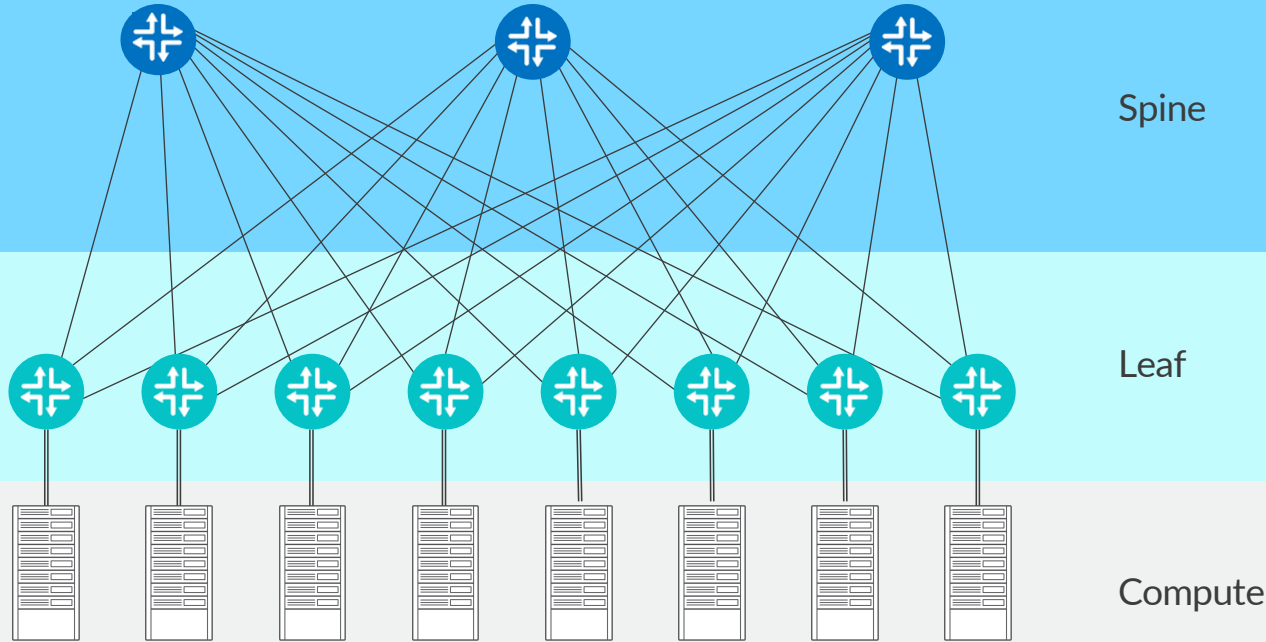
What do we want to achieve?

1. A straightforward and KISS design, where the engineers actually know "what happens when"
2. If customers ask what happened, the Operator will be able to explain

Network Structure

1. A (potentially physically distributed) centralized SDN controller for a Wide Area Network (WAN)
2. "dumb network nodes" without local intelligence and **no pre-provisioned protection mechanisms**
3. **SDN controller as single source of truth**, dealing with failure recovery for a network with ~100 nodes
4. **target recovery time of <2s** for a link failure


DATACENTER SPINE-LEAF STRUCTURE



Data Center Cloud Network

- Spine-Leaf structure
- any to any connectivity between servers
- Handling heavy user traffic


DATACENTER CONTROL STRUCTURE



Control Network

- Provides redundant controller-device connectivity
- is attached to Management ports of devices (limited to 1-2 per device)
- Handling control traffic separate from Cloud network (security)

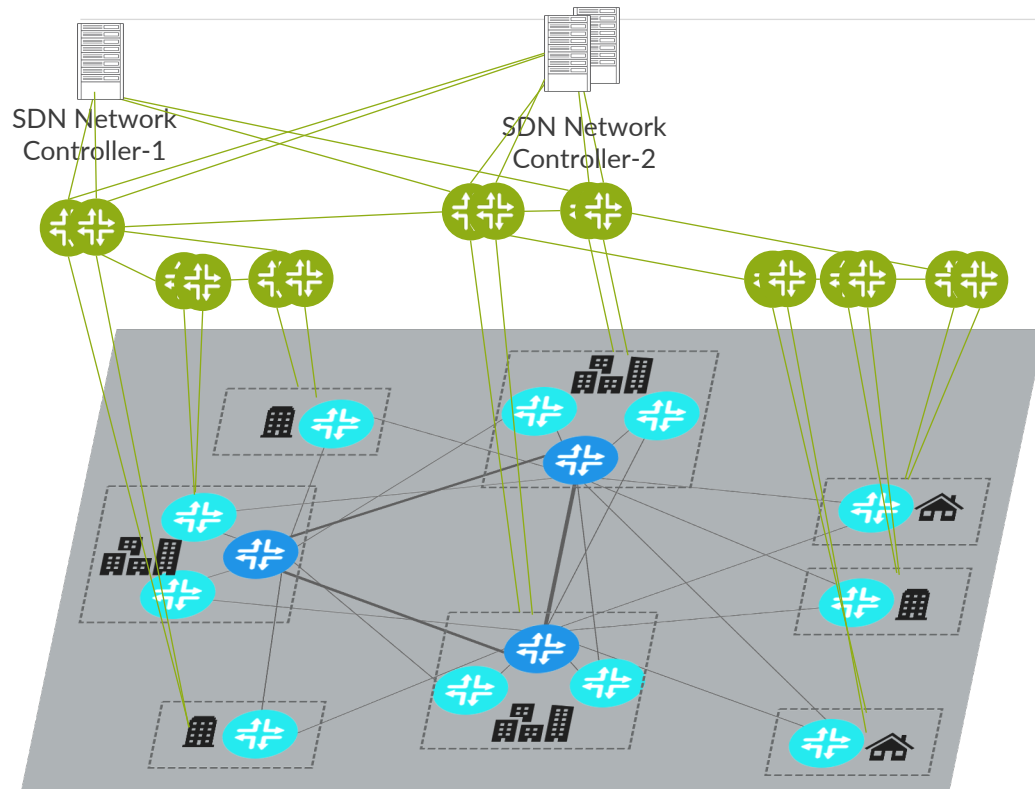
WAN NETWORKING (SIMPLIFIED CONNECTIVITY)



Logical connectivity:

- Routers in POP locations near cities
- Inner core high bandwidth triangle
- Metro routers dual homed to inner core
- Customers connected to metro routers

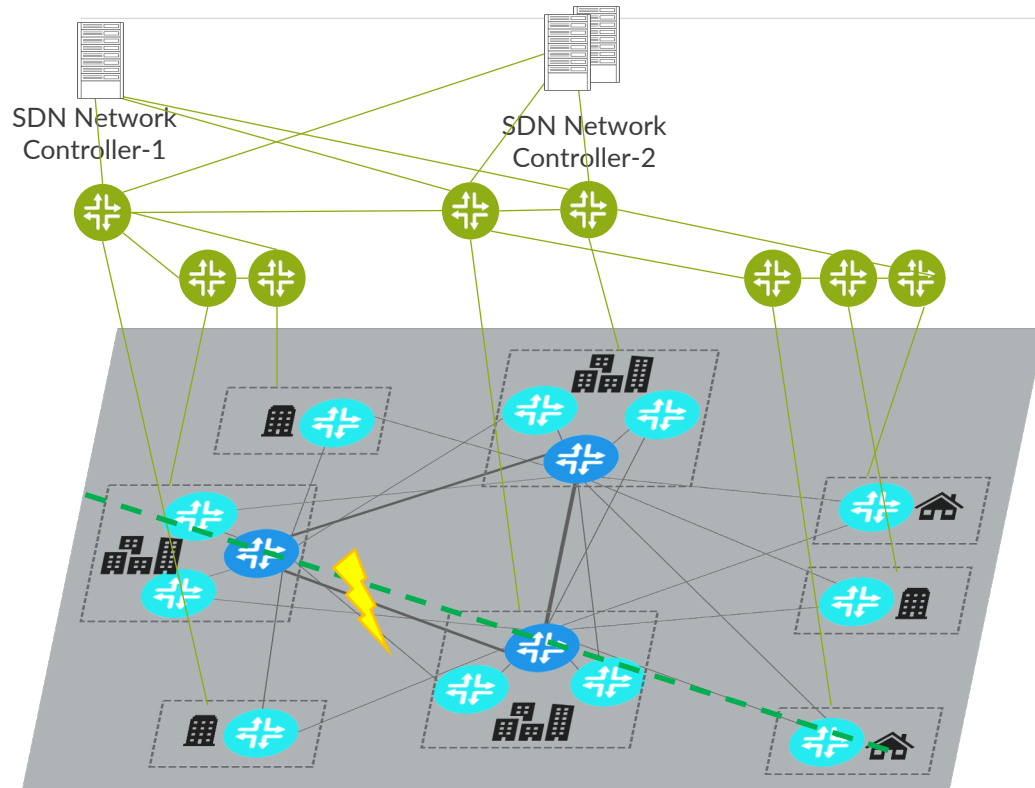
WAN NETWORKING CONTROL STRUCTURE



DCN Logical connectivity:

- Provides redundant controller-device connectivity
- is attached to Management ports of devices (limited to 1-2 per device)
- Redundant DCN routers in POP locations near cities
- Used for control traffic only

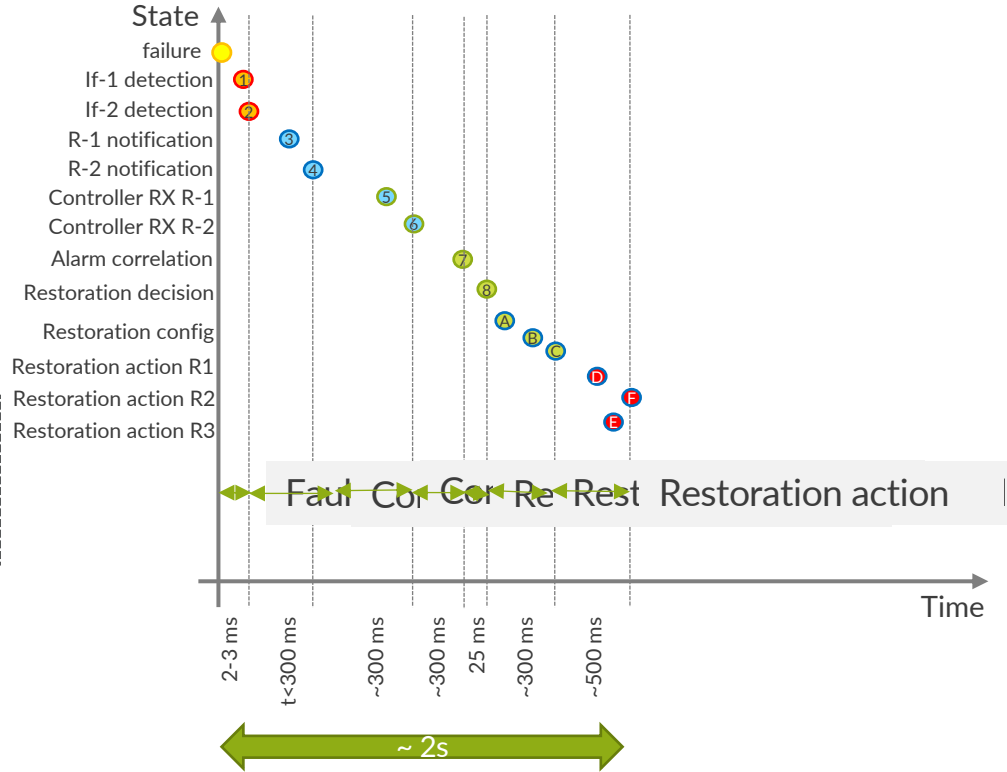

WAN NETWORKING CONTROL STRUCTURE



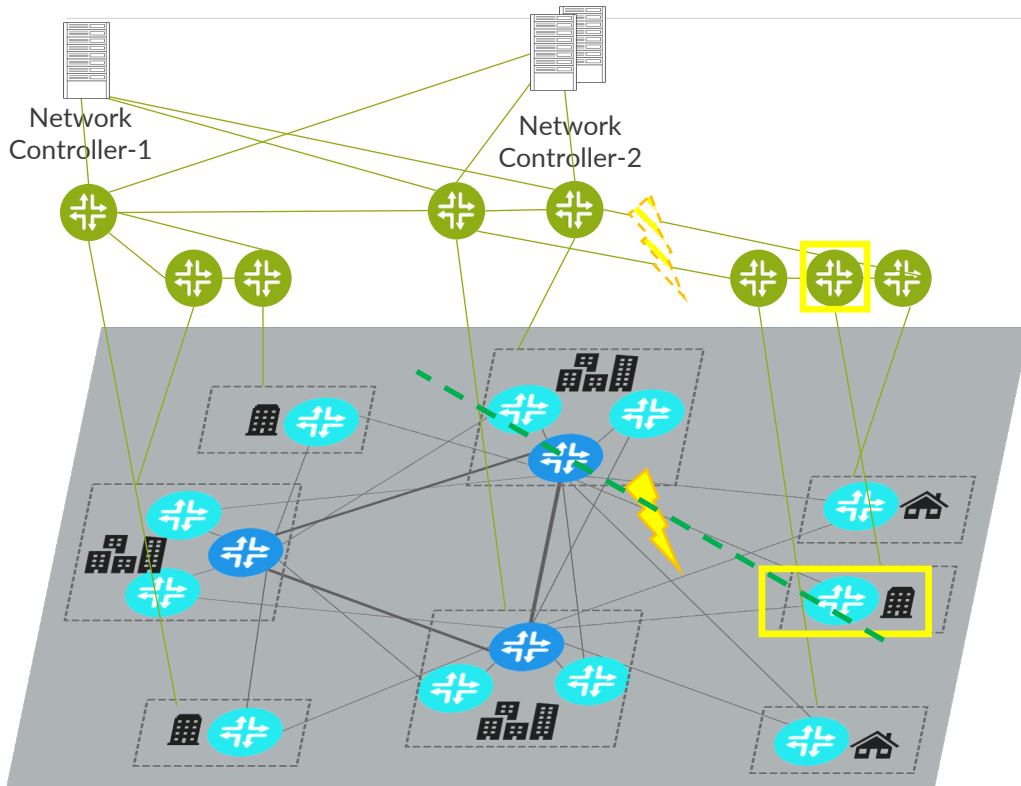
Use case:

1. SDN controller in charge to dynamically restore traffic in case of failures
2. WAN devices are fully under SDN control and routes configured
3. No pre-provisioned protection mechanisms in WAN devices

RESTORATION IDEAL CASE



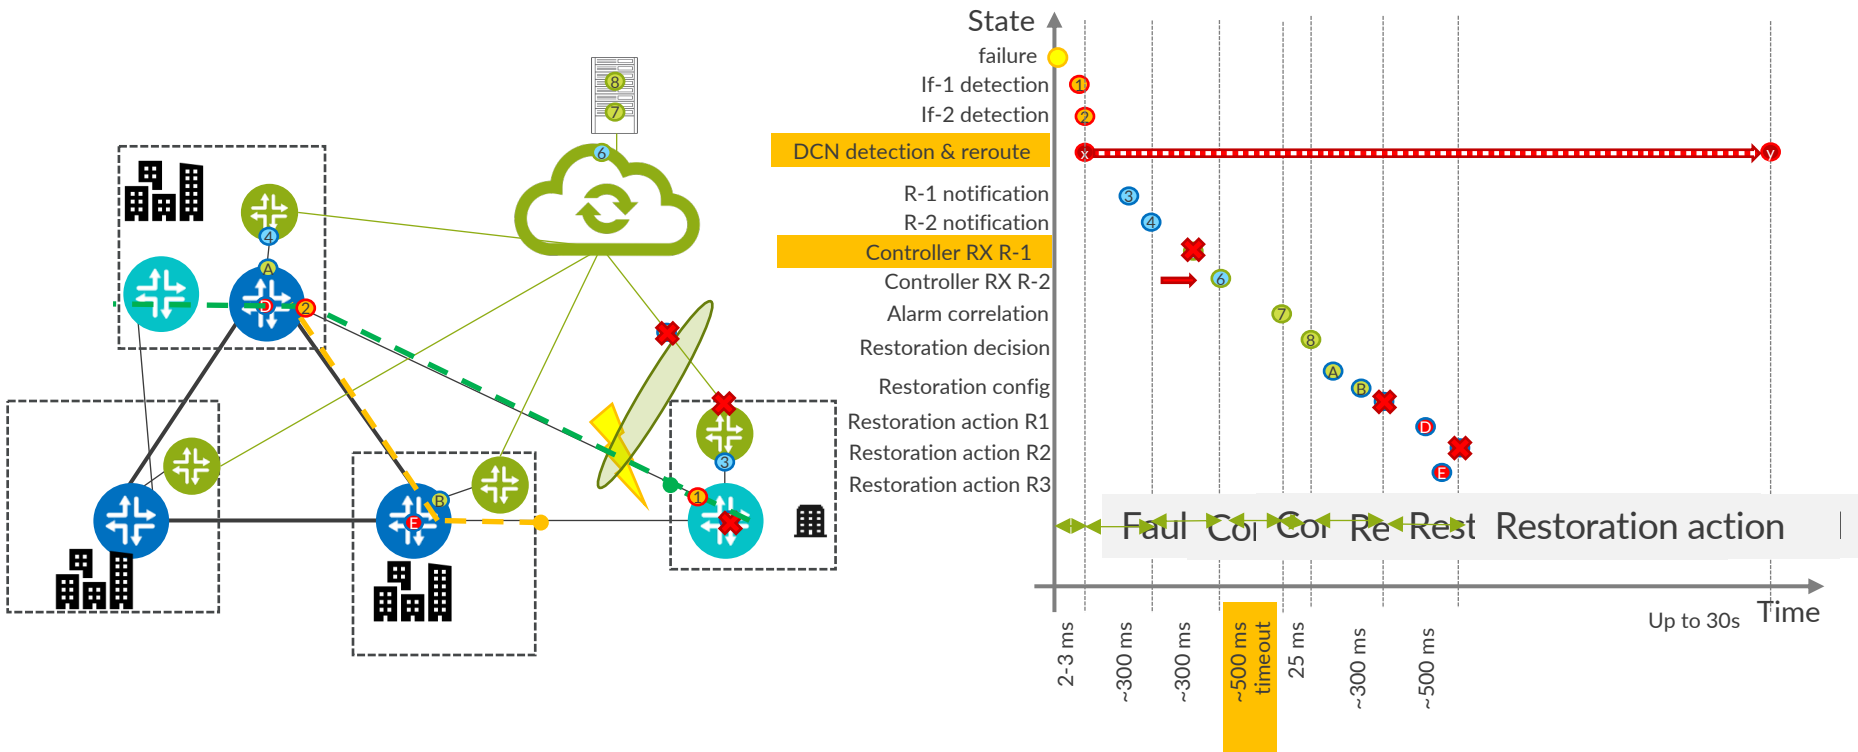
WAN NETWORKING REAL CASE



DCN Physical connectivity:

- Controllers tend to be located at the busiest location, just where most traffic is processed
- DCN is using the same physical ducts and cables as normal traffic, it's just minimally separated: different wavelength or fiber
- Both: user traffic and DCN traffic tends to follow the shortest path
- There is a substantial likelihood that traffic and DCN is co-routed and fails together

RESTORATION BAD CASE: PARALLEL CO-ROUTING



EXACERBATING EFFECTS OF PARALLEL CO-ROUTING

1. Nodes participate in re-routing at the same time as they should send notifications
2. Failure notifications hit the DCN when it is busy with re-routing
3. Only part of the failure information arrives at the controller in the correlation time window
4. Failure correlation needs to deal with partial information
5. Restoration decision may be influenced by partial information but head- & tail nodes may not be reachable
6. If restoration is not completed after tbd-timeout, circuits need to be cleaned up (orphan control)

COUNTER-MEASURES IN THE DCN

1. Traffic Engineer DCN:
 1. Requires detailed knowledge of SRLG for Data and DCN
 2. Needs TE extensions in DCN
2. Provision Fast protection of DCN links i.e. FRR or LFA using BFD:
 1. Protects the DCN link before notification (3), (4) goes out
 2. Requires additional OPEX to set-up and maintain
3. Route control traffic between SDN controller devices →
 1. If controller device SDN-X can't reach node-1, maybe SDN-Y could?
 2. Inter SDN connections need fast protection too
4. If a tail application is virtual, the VM could be shifted to a location with connectivity.
5. Force co-routing of DCN and data traffic and combine with fast DCN failover

