draft-janapath-opsawg-flowoam-req-00

Richard Groves
Microsoft

Janardhanan Pathangi
Balaji Venkat Venkataswami
Dell

Peter Hoose
Facebook

Problem Statement

- This document specifies OAM requirements to improve traditional OAM tools such as Ping and Traceroute.

- Ping and Traceroute do not provide enough information for troubleshooting, performance, and network planning in large scale datacenter environments

- Gathered from operators of large data center networks
  (Microsoft, Facebook, etc)

The Traditional Network

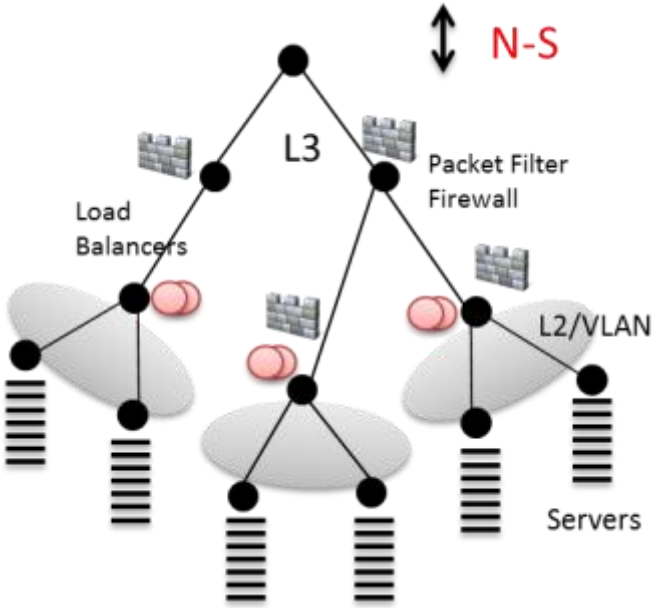

Hierarchical Tree Structure – Optimized for N-S traffic

- hierarchical tree optimized for north/south traffic
- firewalls, load balancers, and WAN optimizers
- not much cross datacenter traffic
- lots of traffic localized in the top of rack

Background

In the traditional data center network……

- Operators manage IP networks with classic OAM tools like Ping and Traceroute.
- Use Ping for end to end connectivity checks
- Use Traceroute for hop by hop path information and failure isolation

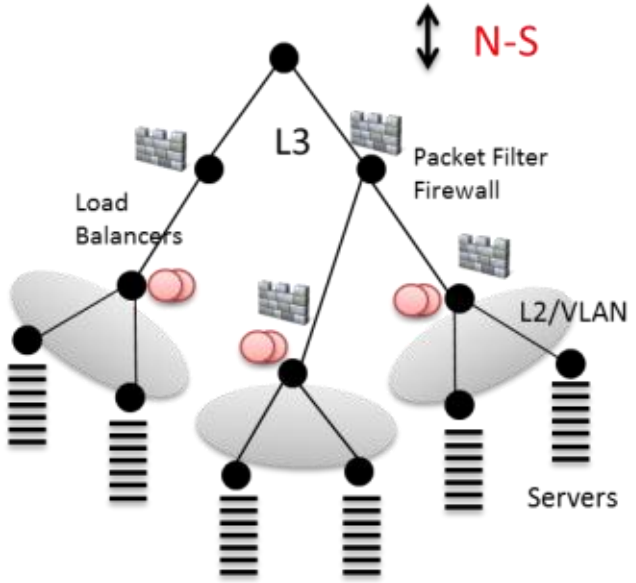

Large Scale Datacenter Network

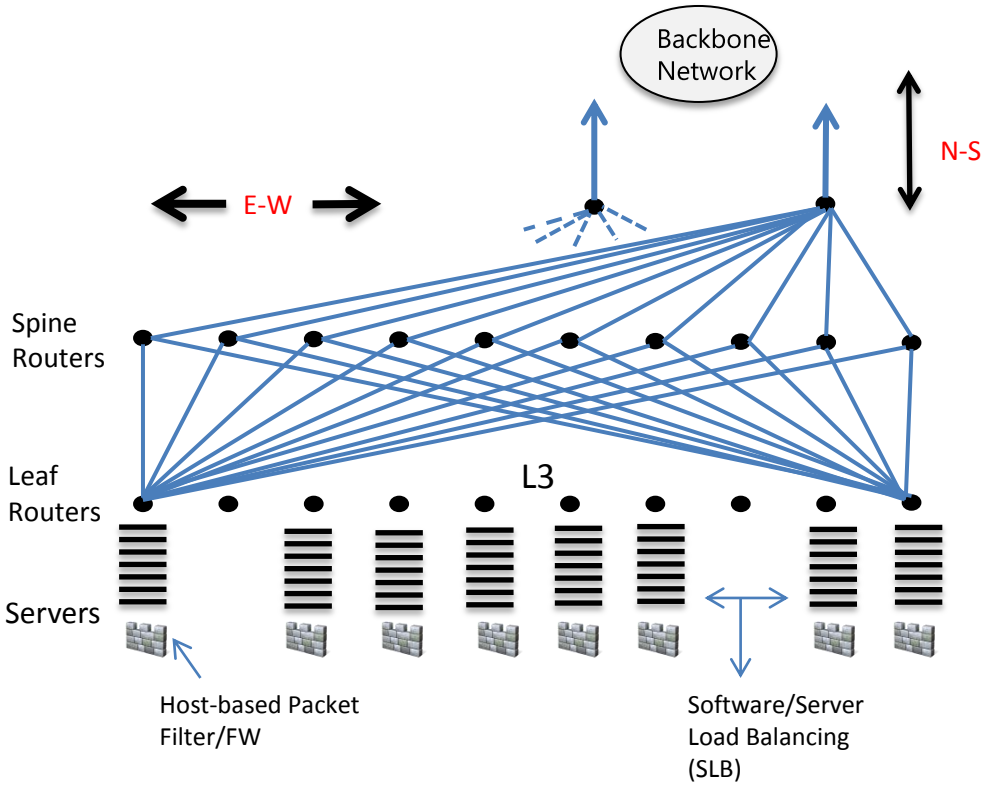

- optimized for east to west traffic
- large scale with lots of ECMP

  draft-lapukhov-bgp-routing-large-dc-01
- load-balancing and packet filtering moved to servers
- network based virtualization

  NVGRE etc
- commodity hardware

Background

However increasingly Datacenters……

- Large amount of ECMP
- existing OAM tools are unable to identify flow specific problems
- also do not provide sufficient information on the various paths which includes performance characteristics.
- unable to provide sufficient information about the performance aspects of a path

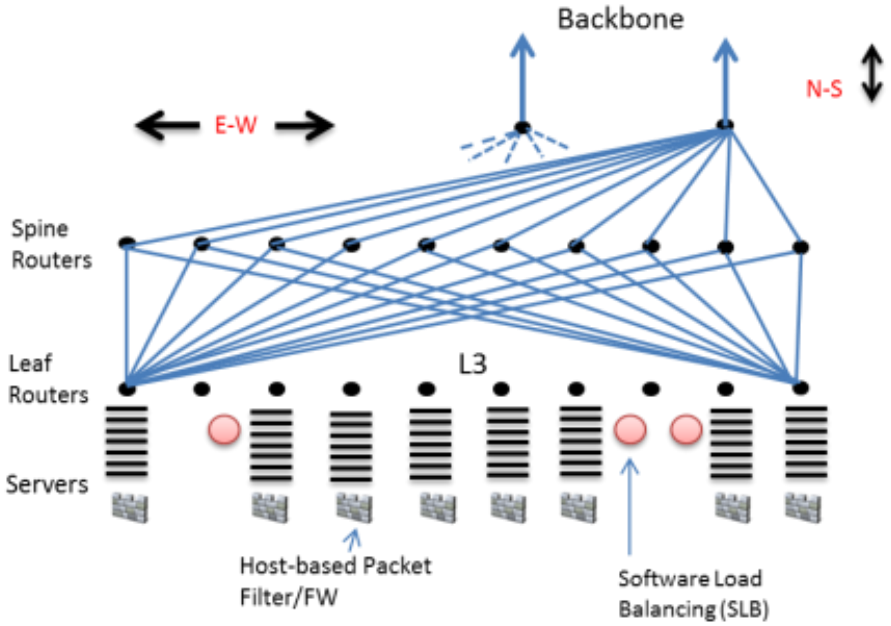

What is needed is a set of tools that will perform the OAM functions based on header fields of actual user traffic.

Requirements

- exact path tracing of a flow while obtaining all relevant info about the links along that path

Requirement 1

Exact Path or Flow Tracing

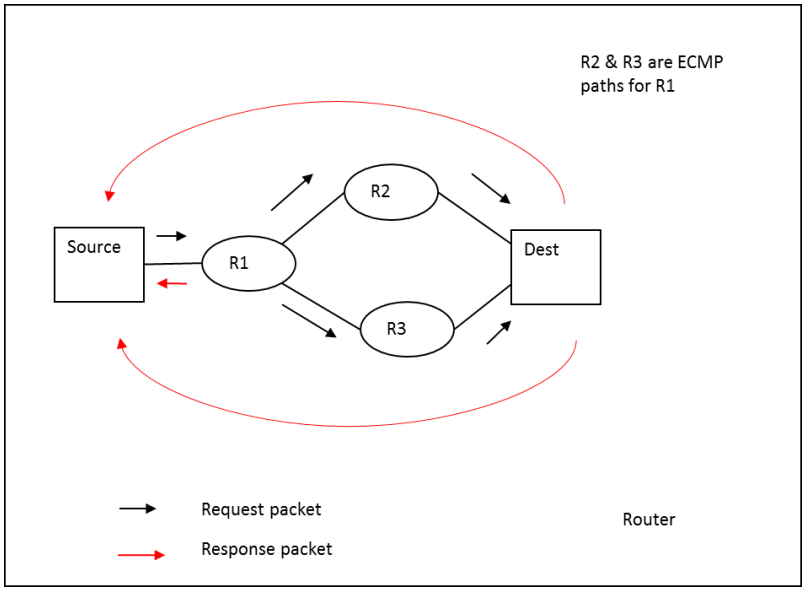

- Tracing through any multi-path configuration.
  (L3, TRILL, SPB, LAGs between adjacent routers)
- trace packet should undergo the same processing at each node as would the actual application flow that is being traced.

- It should be possible to initiate flow monitoring on one or all of the intermediate devices, and should have the following capabilities
- Understanding the hash algorithm is important with respect to which ECMP or LAG member is chosen to forward the packet on.
- It would be useful to know which fields play a part in the computation of this choice.

Problem Statement

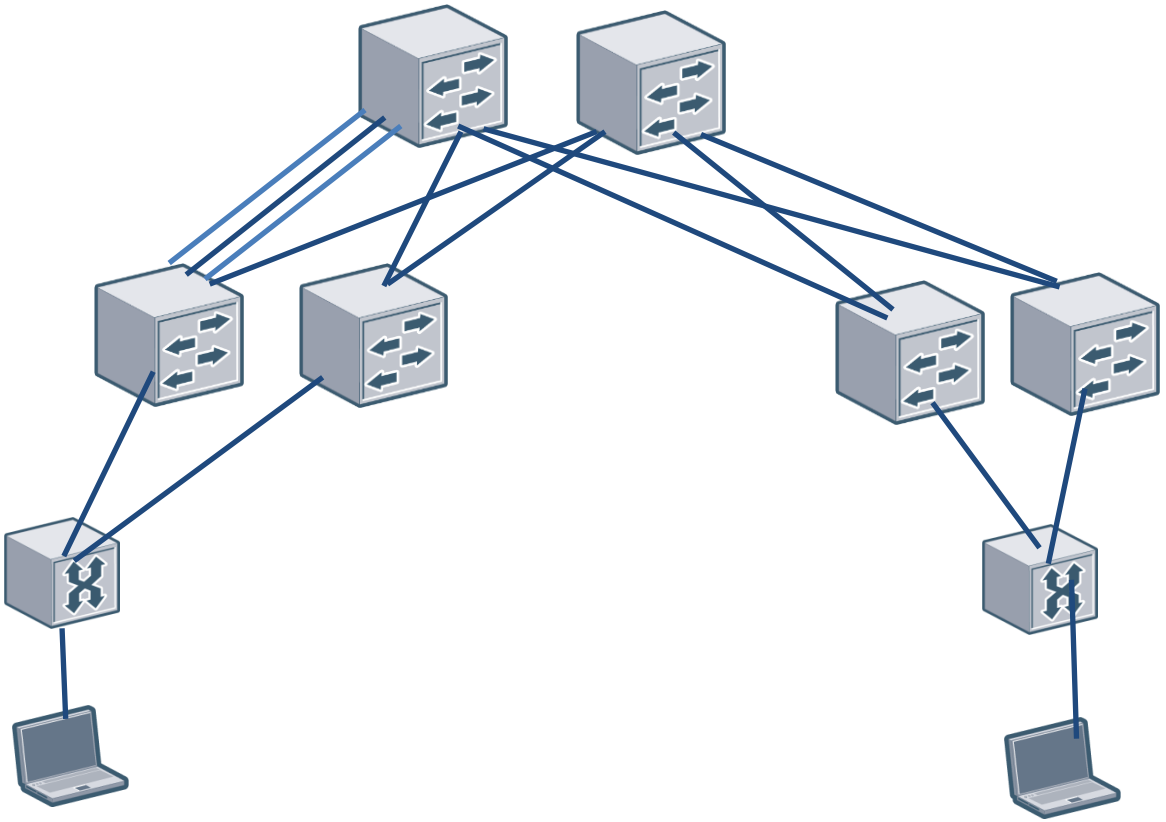

Host A

Host B

- Multiple paths between Host A and Host B

- Network would be using ECMP / LAG for data flow between A and B

- Tracing the path for a particular flow is problem

- Ping, Traceroute give reachability between nodes A and B

- This is not necessarily the path that a particular flow from A to B would take

- Tool needed to trace the exact path a particular flow would take

- Tool needed to identify all paths that exist between nodes A and B

Statistics

- There should be a way to collect the utilization of links along the path in addition to the fan-out information
- The tool should provide an extensible mechanism by which the monitoring station can ask for monitoring of certain parameters for the flow like input rate, packet drops, etc at a given network node.
- Enables a server to detect link polarization and select source ephemeral ports in such a way as to avoid over-utilized links.
- Enables the operator to tweak hashing functions to better match their needs for load distribution

Requirements

- exact path tracing of a flow through the network while obtaining all relevant info about the links along that path
- probes should share fate with the actual flow while not affecting real production flows

Requirement 2
Fate Sharing and Flow Interference

- OAM probes must share the same fate as the real application
- Probes should not affect the real application in progress at the time of troubleshooting.
- The OAM request should not interfere with the real application at the target host
- The OAM response should not go back to the real application at the originator of the OAM query.

Requirements

- exact path tracing of a flow through the network while obtaining all relevant info about the links along that path
- The OAM probes should share fate with the actual flow while not affecting real production flows
- Configurable trace start and stop

Requirement 3
Start/Stop Behavior

- should have the capability to terminate the trace at a specific hop (address/hops)
- The packet should carry a time period and frequency of sampling which capable devices will honor
- A mechanism to start a trace and then monitor until a new request to turn it off is seen should be included
- The device honors request based on its policy, authentication and available resources on the device. It should be able to indicate back in the response if and what parts of the monitoring are activated.
- A local policy override for any of the above should be included.

Requirements

- exact path tracing of a flow through the network while obtaining all relevant info about the links along that path
- The OAM probes should share fate with the actual flow while not affecting real production flows
- Configurable trace start and stop
- Packet Drops and their Reasons

Requirement 4
Packet Drops and Reasons

- The OAM mechanism should be able to indicate the actual reasons for a packet drop.

- The response OAM packet should indicate the error code appropriately.

- Packet drops and their reasons such as Access list based drops, Administratively disabled and Routing Failures

Requirements

- exact path tracing of a flow through the network while obtaining all relevant info about the links along that path
- The OAM probes should share fate with the actual flow while not affecting real production flows
- Configurable trace start and stop
- Packet Drops and their Reasons
- Loop Detection

Requirement 5
Loop Detection

- Should be able to detect when OAM packets are being looped
- If this happens the operation should be aborted.
- Appropriate heuristics may be considered while implementing this feature.

Requirements

- exact path tracing of a flow through the network while obtaining all relevant info about the links along that path
- The OAM probes should share fate with the actual flow while not affecting real production flows
- Configurable trace start and stop
- Packet Drops and their Reasons
- Loop Detection
- Secure Considerations

Requirement 6
Security Considerations

- Securing Requests and Responses
- Information Hiding
- Rate limiting and obviating attack vectors.

Asks from the Working Group

- Feedback on this draft

- Moving it to WG draft status - feedback