



Using Juniper Networks EX Ethernet Switches and Multiple Virtual Routing-Instances in a Ring Topology to Replace Aging SONET Networks

Introduction

In the last 10 years, the price-point of highly capable Ethernet switches and routers has made it possible in many cases, to replace aging SONET network gear at a fraction of the cost of new SONET gear, and with the same ring-protection based fault-tolerance.

One of the biggest historical problems in using switches to perform this function is in the amount of time that it takes for layer-2 and / or layer-3 protocols (spanning tree, OSPF, EIGRP, etc.) to reconverge around a failure. This reconvergence is on the order of tens of seconds, which can cause a perceptible "downtime" to users and a failure of some automated systems.

An additional problem that has plagued using layer-2 switches is the relative inflexibility of the spanning-tree protocol in link selection and loop-prevention in ring-based networks. Switches such as the Juniper Networks EX-3200 and EX-4200, support a robust set of layer-3 protocols and multiple virtual-routing instances to provide fault tolerance, high performance, and network segmentation for a fraction of the cost of SONET devices. Additionally, with support for 10-gigabit Ethernet, this solution should provide more than enough intra-site bandwidth for years to come.




Figure 1 – Typical EX Ring topology

Layer-3 Ring Topology

Using a ring topology, where every site is connected to two other sites, can provide a high level of fault tolerance. Layer-3 routing protocols can be used to determine the best path to take for any particular traffic that leaves a site. The traditional problem with these layer-3 protocols has been that convergence times are on the order of 45 seconds (for OSPF for example) or longer. This is generally considered an unacceptably long time to recover for most applications. In the event that a device loses physical link, there can be a sub-second recovery, but there are cases where no link is lost, but a routing neighbor is not reachable. In order to combat this unacceptably long lag-time, Juniper Networks EX series switches support a protocol called Bi-directional Forwarding Detection (BFD). BFD is a lightweight “hello” protocol, which can help to detect unreachable routing neighbors in under a second (50-100ms times are usually attainable), which is considered to be an adequate recovery time for most applications.

Multiple Routing Instances

One of the traditional problems with using layer-3 routing between sites is that it can be quite challenging to segment networks within each site. Traditional Layer-2 VLAN's can be used to segment the network at a site, but this would require layer-2 trunks for the site-to-site links, requiring an undesirable loop-prevention protocol like spanning tree running on the uplinks. Juniper Networks EX series switches support multiple routing-instances or virtual routers within a single switch. By separating physical ports on the switch into distinct routing-instances, and by using 802.1Q VLAN tags

on the uplinks (tagged, but layer-3 protocol only) to virtualize the uplink ports into multiple logical ports, distinct routing domains can be created around the ring in order to provide traffic isolation all the way back to a “core” location. Each of these routing domains can be then provided with unique routes from either different upstream devices, or by different virtual-routing instances from a border-firewall cluster (such as the Juniper Networks SRX product line). This can be especially useful if there is a need to transport public-internet traffic from one or more sites (guest access networks) across the infrastructure as private data.

10-Gigabit Ethernet Uplinks

Network uplinks typically run in what is called an “oversubscribed” state, meaning that the aggregate bandwidth of the non-uplink ports at a site is greater than the bandwidth of the uplink. These uplinks then often become the points of network congestion within a network. By utilizing 10-gigabit Ethernet uplinks (where the fiber plant supports it), the overall network bandwidth can be increased to support newer bandwidth-hungry applications such as video-conferencing, IP security camera cameras, offsite backup applications, etc. Additionally, the Juniper EX product line supports advanced class-of-service features so that traffic can be prioritized and treated appropriately to ensure that critical applications receive the bandwidth they require.

Virtual-Chassis and Multiple Routing Engines

Juniper Networks EX-4200 switches support a feature called “Virtual-Chassis”, which is a method of placing up to 10 switches into a single management unit. Although port-density is usually not a design-goal of a transport network, the ability to put two devices into a virtual-chassis becomes useful in this design as these two switches can act as a primary and a backup “routing-engine”. By staggering the uplinks onto these two separate switches, a single switch failure does not take down an entire site as there is always a routing engine with a working uplink at a site.

Conclusion

In conclusion, as Ethernet technologies continue to gain traction in both the LAN and WAN spaces, vendors such as Juniper Networks will provide an increasing array of products and features with which network architects can build reliable networks at a lower price-point than has been traditionally possible. Using the features of the Juniper Networks EX switches, a network can be built at a lower price-point than has been possible, without sacrificing the performance, security isolation, and fault-tolerance of legacy SONET networks.