internet.com®

Five Reasons Classic Ethernet Switches Won't Support the Cloud

By Drew Robb

When a company starts looking at providing cloud services to its clients, or consuming cloud services itself, it often becomes quickly apparent that the network needs to be upgraded to support this new service delivery model.

"External and hybrid cloud models mean that computing resources are remote, and access/ communications with those workloads will have to traverse a wide area or Internet network link," says Jim Frey Managing Research Director of analyst firm Enterprise Management Associates. "This introduces substantially greater latency (and bandwidth constraints) than are normal for traditional campus/data center-based computing."

Before implementing cloud computing, therefore, it is critical to examine the strength of the network to make sure that any weak points are located and eliminated. This helps to ensure that users receive the same QoS they are used to when using local servers. In addition, IT has to be made aware that traditional Ethernet network architectures may not be adequate in achieving the desired service levels. Here are five reasons why this is the case.

1. Classic Ethernet Networks are Hierarchical

Ethernet typically relies on a hierarchy of three or more tiers. Moving traffic between different server racks, for example, requires transitioning up and down this logical tree. The important point here is that each step along the path adds to the total latency. To avoid bridge loops between different network segments, IEEE 802.1D Spanning Tree Protocol (STP) allows only one active path at a time between two switches. However, this leads to congestion on the Inter-Switch Link since the bandwidth is limited to that single link. There are, of course, variations of that STP protocol that seek to address this issue. Rapid Spanning Tree Protocol (RSTP) is a method defined in IEEE 802.1w that reduces the time it takes to respond to changes in network topology. Another approach is Multiple Spanning Tree Protocol (MSTP), IEEE 802.1s, which extends RSTP to include a separate spanning tree for each Virtual Local Area Network (VLAN) group and blocks all but one possible passage for that VLAN within the spanning tree.

2. Classic Ethernet Networks are Unintelligent

Classic switches are designed to connect physical, not virtual, resources. Each port is individually configured with the QoS, security, VLAN traffic and other network policies required by that particular server. In a virtual world, however, VMs move from one box to another depending on traffic loads or maintenance requirements. Unless the port configuration follows the VM to its new physical location, the connection fails. What this means to IT is either manually reconfiguring the ports with each move, or implementing a more intelligent class of switches that have all the applicable policies available at all ports, and can utilize those policies that are applicable to the service that is connected at the time.

3. Classic Ethernet Networks Don't Scale

Since Classic Ethernet networks are designed to allow only a single connection between switches, this limits the amount of traffic they can carry.

"The scalability of traditional Ethernet is problematic, especially if an organization assumes, as it nearly always will, that cloud infrastructures will grow," says Charles King, Principal Analyst at Pund-IT, Inc. "That said, some vendors are delivering or proposing cloud solutions based on higher throughput 10Gb and 40Gb Ethernet."

To reduce the impact of this bottleneck, link aggregation groups (LAGs) permit multiple physical links to share the ISL. But this requires manually configuring all the ports in the LAG, a task that becomes increasingly complex as more switches are added.

4. Classic Ethernet Networks Are Inefficient

Due to the limitations of STP, only one link or LAG can be active at a time, lowering the utilization rate. Adding or removing a new link causes the entire network to pause while the network reconfigures a new loop-free tree. This process can take from a few seconds up to several minutes, disrupting traffic and sometimes causing server crashes.

5. Classic Ethernet Networks Are Complex

With classic Ethernet networks, each port on each switch needs to be configured with the applicable VLANs, QoS, security, network policies and protocols. As servers are added, more switches are needed and they can't share a common configuration.

Switching to Fabric

To eliminate the various bottlenecks caused by traditional Ethernet architectures, switch vendors have begun to adopt fabric architecture into their Ethernet switches. This eliminates or mitigates each of the problems listed above.

Fabric switches flatten the network, reducing latency by eliminating unnecessary steps and preventing bottlenecks that can occur at the ISL.

Since they are more intelligent than their traditional counterparts, all ports on all switches within the fabric can share common configuration parameters. When VMs migrate from one location to another, for instance, ports don't require reconfiguration. Further, they can scale much higher because a logical trunk is automatically formed when more than one connection is required instead of allowing a single connection between switches.

Ethernet fabric switches are more efficient because they use link state routing with multipath routes. This allows traffic to take the shortest path through the network, rather than following the hierarchical path imposed by STP. Finally, Ethernet fabrics are simpler to manage since configuration information is shared with all ports in the network.

Classic or Fabric?

When is a classic Ethernet network adequate and when should one change to an Ethernet fabric? It depends on the applications. Some applications require large amounts of raw throughput; transaction systems are sensitive to latency.

“It comes back to what is the workload one is trying to project up into the cloud,” says Dan Kusnetzky, analyst and founder of the Kusnetzky Group. “Until we know more about the application and its requirements, it is hard to say whether or not a typical router will satisfy the requirements or not.”

By analyzing one’s current network and any proposed new cloud services, it will become clear whether the existing infrastructure will work.

“If you are going to Fibre Channel over Ethernet (FCoE) as part of a converged and cloud environment, you will want to invest in new technologies to take full advantage of those capabilities,” says Greg Schulz of the Server and StorageIO Group. “However, if your network and switches can handle the local as well as wide area traffic, cloud service access and security, they might be able to support cloud access and the question should be what upgrades are needed for wide area access and optimization. Ultimately, though, some network upgrades and enhancements will be in your future.”