Will Traffic Spikes Overwhelm Your Data Center Network?

A Dell Technical White Paper

September 2011

Contents

Figures

One of the most significant data center technologies introduced over the last ten years is virtualization. Both server and storage virtualization are commonly found in today's data centers, bringing more flexibility, greater efficiency and higher availability to IT computing environments. As we try to leverage server and storage virtualization, we need to look at how we deploy network virtualization judiciously so as to implement an optimized data center. As virtualization technology expands beyond the test and development environments where it started to production environments, IT managers are learning the impact of virtualization on the network. With most virtualized servers supporting between 2 and 10 applications, the network traffic profile is changing. Not only is network traffic increasing, but the potential for very large network traffic spikes has increased as well. With each of these virtualized servers, each rack now hosts between 160 and 400 application instances. IT managers must take this into consideration as they deploy more virtualized servers in their data centers.

Data Center Traffic Spikes

In the data center, traffic spikes are caused when a large number of virtual machines experience a burst of network traffic at the same time. In the past, data centers did not experience these network spikes because physical servers were generally inefficient. Servers were heavily underutilized. However, with virtualization and cloud computing, servers are much more highly utilized and more concentrated in fewer locations, and as a result, data center network traffic is changing.

IDC's World Wide Server Forecast 2010 report shows that the rate of virtual server growth is far outpacing physical servers. Prior to virtualization, most servers were supporting a single application. This led to highly inefficient architectures because servers were often sized to support peak application needs. It was common for servers to be only 10 to 20 percent utilized. In today's data centers, as IT managers deploy applications in virtualized environments, servers support multiple applications, dramatically increasing the efficiency of servers to as high as 50 to 60 percent. With virtualized servers supporting 4-10 applications, a rack of 20 servers that in the past would only support 20 applications now supports as many as 200 applications.

> “A new crop of switches is now appearing that offer non-blocking, line-rate performance with several gigabytes of buffering capacity. As a result, buffer capacity is self-regulating and traffic spikes simply cause the use of larger amounts of buffer.”
>
> **Bala Pitchaikani**
> Vice President, Product Line Management
> Dell

To reiterate, using older servers traffic spikes in the network would have happened from tens of application instances in a rack and they would have and still are being addressed as 'network surges.' However, with newer high-density servers, traffic paring mismatch is caused by as many as 200 applications in a rack. When data bound for one of its application instances is lost and it requests a re-transmit of data, more traffic is placed on the network which could impact other application instances.

The occurrence of these major traffic spike events is lower than regular traffic spikes. However, the sad truth is that they are real and when they happen they hit us with no warning. It's the spiraling effect that's created by a combination of multitude of factors that makes its occurrence so unpredictable.

Disastrous Cascading Traffic Spikes

However, the very simple additive changes that are being made today in the data center to meet business objectives is making this a real issue to actively plan for and take the appropriate insurance policy ahead of time. If data centers do not proactively plan for network traffic spikes, disastrous effects can occur. Network traffic spikes last for seconds or even a minute or two, which could result in a hiccup to the system and the data center operation moves on. However, the cascading effect of a network traffic spike impact could last for several hours and/or potentially impact business operations before it's noticed.

The trend towards cloud computing has also been a driver in changing the network traffic profile. With cloud computing, applications are becoming more concentrated in a smaller number of data centers. Hosting and portal companies are managing huge data centers, providing computing services that in the past were supported by internal IT departments. Additionally, internal IT departments are also becoming more concentrated as they shift from having multiple, smaller data centers to having fewer, larger data centers. With fewer, larger data centers, networks are experiencing more traffic. The growing popularity of virtualization and cloud computing are changing network traffic profiles, and as a result, require IT managers to re-architect their data center networks.

Strange Traffic Spikes?

Data center networks supporting today's virtualized data centers are experiencing network traffic that looks different from what it looked like in the past. With a larger number of applications generating more network traffic, not only is the volume of traffic higher, but the profile of network traffic is also different. As the number of applications generating network traffic increases, the aggregated result is mostly a higher level of traffic with fewer peaks and valleys because the network peaks and valleys from each application tend to cancel each other out. However, when a majority of the applications generate a burst of network traffic at the same time, large network spikes occur. These large network spikes in virtualized data centers' increasing server densities can result in dropped packets, and therefore poor network performance.

Figure 1. Typical Data Center

Old Switches With New Problems

For traditional network traffic, dropped packets are simply retransmitted, resulting in slower network performance. However, for storage networks, as is the case with iSCSI, dropped packets could result in a re-transmit of the large block of data, increasing the stress on the network. The emerging trend toward consolidating server and storage traffic on the same network using protocols such as iSCSI and Fibre Channel over Ethernet (FCoE) makes it imperative that the network be immune to packet loss. Unfortunately, more complete standards like data center bridging are focused on 10 GbE Ethernet while the majority of data centers are still running 1 GbE (albeit, multiple 1 GbE links) from the servers to the network.

As highly-virtualized data centers begin to show signs of traffic spike activity, IT architects will have to focus on the switching infrastructure to ensure that it is capable of handling significant traffic spikes. It would be costly to model the tipping point when ‘x’ number of virtual application instances running on physical servers with ‘y’ number of cores and ‘z’ number of 1 GbE interfaces making ‘bursty’ network request at p% of the time of file block sizes ‘q’ from ‘r’ number of files or databases at a given point in time.

IT architects have been fortunate not to have to operate on all these variables until recently and have not been aware of the imminent unpredictability of network traffic that’s waiting to overwhelm the data center. The cost-effective answer is to implement switches that have more than enough buffering capacity to handle large traffic spikes.

A typical data center top-of-rack switch has 8-16 megabytes of buffering capacity. When this capacity is overwhelmed, packet loss occurs. Conventional wisdom to avoid packet loss is to (a) look at ways to 'throttle' the traffic by identifying those 'offending' application instances and 'tame' them or move those 'heavy' applications to a separate network/rack or (b) throw more network bandwidth at the problem, say by adding another 1 GbE connection or an even more expensive approach of upgrading the whole infrastructure to 10 GbE or (c) take the Cadillac approach of implementing 10 GbE infrastructure with converged network adapters and FCoE et.al.

Low-Latency Network Future

However, a new crop of switches is now appearing that offer non-blocking, line-rate performance with several gigabytes of buffering capacity. Each port is assigned a given amount of buffer, but if that capacity is exceeded the port can dynamically draw upon a large reserve pool of memory if necessary. As a result, buffer capacity is self-regulating and traffic spikes simply cause the use of larger amounts of buffer.

Figure 2. Dell Force10 Top-of-Rack S60 Switch

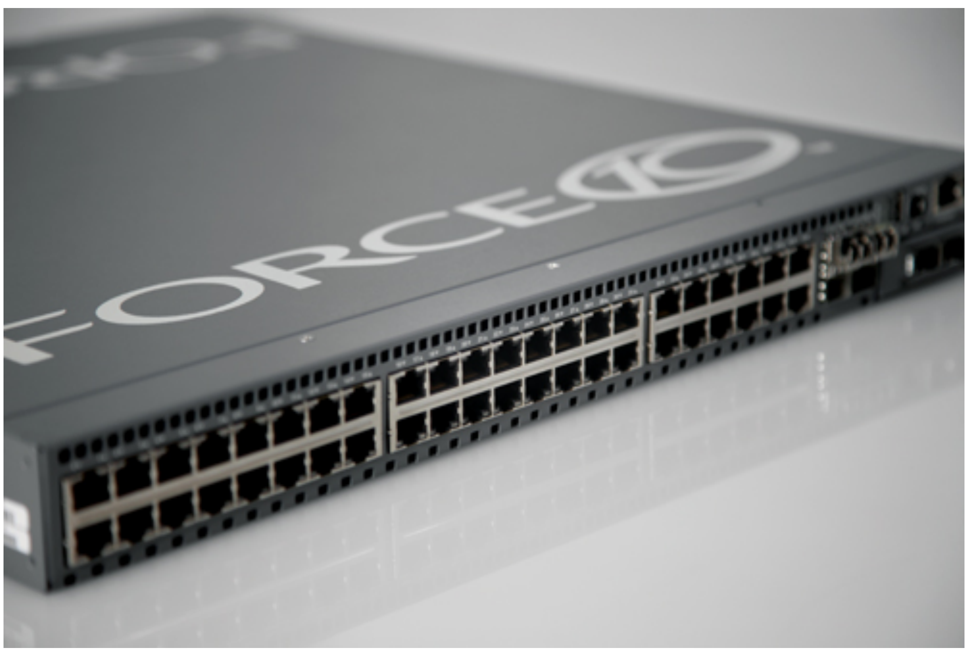

Also, since this new crop of switches is available for 1 GbE non-blocking, line rate performance, IT architects do not have to prematurely upgrade their infrastructure to 10 GbE and protect their data center investment.

As IT architects seek to flatten their networks in order to reduce latency within virtualized environments, they should plan to eliminate legacy switches that can be swamped with traffic spikes, and to implement switches with ample, dynamically-adjustable buffers that can handle whatever the server and storage infrastructure require.