SuperSizing the Data Center: Big Data and Jumbo Frames

Lori MacVittie, 2012-16-02

#centaur #40GBE **Data center transformation discussions too often overlook the impact on the network – and its necessary transformation.**

For many of the same reasons IPv6 migration is moving slower than perhaps it should given the urgent need for more IP addresses (to support all those cows connecting to the Internet) is the sheer magnitude of such an effort. Without the ability for IPv6-only nodes to talk to IPv4-only nodes, there's a lot of careful planning that has to happen around the globe to ensure success and continued communication between the two incompatible protocols.

In many ways, Jumbo Frames – despite performance advantages – suffer from the same technological incompatibility. Remember that Jumbo Frames – 9000 bytes – are incompatible with regular old sized Ethernet frames (1500 bytes). It makes sense for much the same reasons – you simply can't stuff 9000 bytes into a frame designed to hold 1500. And one of the basic rules of Ethernet is that the smallest MTU (maximum transmission unit) used by any component in a network path determines the maximum MTU for all traffic that flows along that path.

And yet the benefits of Jumbo Frames have been demonstrated many times. It reduces fragmentation overhead (the process of splitting data into chunks small enough to fit into a 1500 byte frame) which translates into lower CPU overhead on hosts. It also allows for more aggressive TCP dynamics, which results in greater throughput and better responses to some kinds of loss. But even though Jumbo Frames can deliver an increase in throughput along with a simultaneous decrease in CPU utilization they are rarely, if ever, used in a data center network.

That, however, is changing.

You might recall some predictions with respect to 10GB adoption in the data center:

> "We expect the Ethernet Switch market to experience two significant years of market growth in 2013 and 2014 from the migration of servers towards 10 Gigabit Ethernet," said Alan Weckel, Senior Director of Dell'Oro Group. "We believe that in 2013, most large enterprises will upgrade to 10 Gigabit Ethernet for server access through a mix of connectivity options ranging from blade servers, SFP+ direct attach and 10G Base-T.
>
> -- Data Center to Drive Ethernet Switch Revenue Growth through 2016, According to Dell'Oro Group Forecast

Historically in the switching market the deployment of 10G in the core networks and the use of Jumbo Frames went pretty much hand-in-hand. Until recently, however, 10GB just wasn't making its way into the data center (costs were too high) and the only place Jumbo Frames were really seen was within storage networks, particularly in conjunction with FCIP implementations.

Aren't jumbo frames bad for multimedia?

Applications sensitive to jitter and burst drops can experience a negative impact from jumbo frames.

But just because you can, doesn't mean you must. The rule is the *smallest MTU* is an upper bound for *every* application. Just because an MTU is set at 9000 (or 1500 for that matter) does not mean an application must fill the payload with data. In fact, the overwhelming majority of HTTP requests have payloads far smaller than even standard sized frames.

Thus, even if jumbo frames are supported, multimedia applications can continue to use smaller frames, as befits their performance requirements.

So what you're saying is watching Pocoyo over Netflix is *not* going to be impacted, right?

For the most part, a lack of support within the data center infrastructure and no real urgency for the efficiency gains that come from Jumbo Frames (and the fact that the Internet is not using Jumbo Frames from end-to-end, which pretty much kills the value proposition) meant enterprise organizations looked at Jumbo Frames with a "someday, but not right now" attitude. But with the increasing adoption of virtualization and movement of 10G networks into datacenters (in part driven by virtualization), Jumbo Frames are becoming more of a reality for a larger population of organizations.

Consider the following support and recommendations for jumbo frames within VMware's documentation:

TCP Segmentation Offload and Jumbo Frames:
Jumbo frames must be enabled at the host level using the command-line interface to configure the MTU size for each vSwitch. TCP Segmentation Offload (TSO) is enabled on the VMkernel interface by default, but must be enabled at the virtual machine level.
-- ESX 4.0 Config Guide, page 57

Optimizing vMotion Performance
Use of Jumbo Frames is recommend for best vMotion performance.
-- Page 188 vSphere 4.0 System Admin Guide

vSphere 4 Performance
Jumbo Frames is one of the suggested means of improving CPU performance with respect to vSphere
-- CPU Performance Enhancement Advice (Table 22-6, page 278)

Add in cloud computing and a desire to more quickly move big data (virtual machines) over the WAN to cloud providers for a variety of business initiatives – a process in which the number of frames sent and low latency is key to success - and Jumbo Frames suddenly start looking a lot more like a requirement than a "Yeah, yeah, we'll get to that eventually. Maybe."

Virtualization and cloud computing are transformative technologies. As some have often – and loudly – reminded us, the *network* is part of the data center, and indeed an integral part of the data center. While we tend to focus on the management and provisioning and automation of the data center and its cultural impact, we should not overlook the impact that these technologies and the changes they bring are having – and will have – on the network.

If cloud and virtualization and consumerization and emerging technologies like HTML5 are going to transform the data center, that's going to necessarily include the network. Ultimately, support for Jumbo Frames will be a requirement – a checkbox item – for every component in the data center.

- Sometimes It Is About the Hardware
- Live Migration versus Pre-Positioning in the Cloud
- F5 and VMware–One Step Closer to the Cloud as a Seamless Data Center Extension
- Cloud is an Exercise in Infrastructure Integration
- Performance in the Cloud: Business Jitter is Bad
- Like Cars on a Highway.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com