

Forget Performance IN the Cloud, What About Performance TO the Cloud?



Don MacVittie, 2011-15-09

In an N-Tiered architecture, the network connection between tiers becomes a truly important part of the overall application performance equation. This is a fact we have known for a couple of decades now. If your network performance is down for some reason (from mis-wiring to hardware mis-configuration to over utilization), your application performance will, by definition, suffer. I ran a test once while writing for Network Computing where the last person to use the lab had programmed the ports I was using to shunt bandwidth over threshold X onto a different VLAN. It took weeks and the help of one of the vendors under test – Juniper – to help me figure out exactly what was going wrong. Only after they observed that it was topping out and then dropping packets were we able to track down the one changed configuration setting. The funny thing is that this particular change would not have impacted the vast majority of testing that we did in the lab, I was just unlucky enough to pick the wrong ports for a test that purposefully overloaded the network. Until the Juniper crew said “no, that’s not right, we’ve done this test before”, I was blaming the products for the failures under high load... A symptom of the fact that when network performance degrades, systems appear to degrade. Thankfully, we caught the problem and I did not go to print with misinformation.

But it does highlight the need for network performance to be top-notch if your application performance is to be top-notch. And we’re starting to head into uncharted territory where this fact is concerned. The problem is that enterprises don’t generally throw a whole bunch of data over the WAN, and if they do, they have specially provisioned ways to do so – because they’re going to a remote datacenter or partner network that is known and data volumes can be calculated for.



But as we approach cloud computing we need to be more aware of the network performance aspects than we would be either in the datacenter or transferring between two datacenters. The reasons for this are manifold. First, you don’t have control of the connection to the cloud provider. You own one end, and with the right tools (skipping plug for [F5 WOM](#) here, look it up if you have a need) you can “control” both ends of the connection, but you don’t really have control of what is in-between. The volume of traffic, the size of the pipes, etc. are all dictated by outside forces. For some DC-to-DC connections this is true also, but unlike cloud, DC-to-DC is point-to-point. If you are dealing with one of the major cloud vendors, you can’t even be certain what *country* your app resides in at the moment, let alone the route to that application.

Some of this concern is automatically mitigated. If the platform provider is large enough to have datacenters in multiple countries, they have pipes bigger than New York sewers connected to those datacenters. Some of it can be mitigated by those same “right tools” that help at the ends of the connection, because they can create tunnels or p2p connections to handle communications, and offer bandwidth reduction capabilities to make certain you are only sending smaller amounts of data, reducing the impact of the network by having less round trips across it.

In cloud storage, you have a bigger issue because the whole point is to send massive amounts of data across the WAN. While the right products will reduce the footprint of that data with compression (skipping plug for [F5 ARX](#), look it up if you have a need), you still are sending a lot, or you wouldn’t have to go to a cloud platform, you’d just store it locally. So the question becomes how to make certain that performance to the cloud storage vendor is optimal when you don’t own both ends of the connection? That’s a tricky question, because it is not just a question for today, it is a question forever. You see, the more you store in a cloud storage providers’ space, the less likely you are to want to change providers. But the more business a cloud storage provider receives, the more companies are cramming huge volumes of data in and out of their DC. Which could cause performance problems for you... Unless you have SLAs that are rock-solid and you’re willing to enforce.

The long and the short of this post is some advice. Get tools to reduce the amount you’re sending over the wire, and make certain you have an SLA that will cover you if your usage (or the other people using the providers’ usage) jumps. And yeah, if your vendor charges by the megabit or megabyte, check out some products that might reduce your throughput. I might have mentioned a couple here, but there are more out there.

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

©2016 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS04-00015 0113