Does Cloud Solve or Increase the 'Four Pillars' Problem?

Don MacVittie, 2012-11-10

It has long been said – often by this author – that there are four pillars to application performance:

1. Memory
2. CPU
3. Network
4. Storage

As soon as you resolve one in response to application response times, another becomes the bottleneck, even if you are not hitting that bottleneck yet.

For a bit more detail, they are

- “memory consumption” – because this impacts swapping in modern Operating Systems.
- “CPU utilization” – because regardless of OS, there is a magic line after which performance degrades radically.
- “Network throughput” – because applications have to communicate over the network, and blocking or not (almost all coding for networks today is), the information requested over the network is necessary and *will* eventually block code from continuing to execute.
- “Storage” – because IOPS matter when writing/reading to/from disk (or the OS swaps memory out/back in).

These four have long been relatively easy to track. The relationship is pretty easy to spot, when you resolve one problem, one of the others becomes the “most dangerous” to application performance. But historically, you’ve always had access to the hardware. Even in highly virtualized environments, these items could be considered both at the Host and Guest level – because both individual VMs and the entire system matter.

When moving to the cloud, the four pillars become much less manageable. The amount “much less” implies depends a lot upon your cloud provider, and how you define “cloud”.

Put in simple terms, if you are suddenly struck blind, that does not change what’s in front of you, only your ability to perceive it.

In the PaaS world, you have only the tools the provider offers to measure these things, and are urged not to think of the impact that host machines may have on your app. But they do have an impact. In an IaaS world you have somewhat more insight, but as others have pointed out, less control than in your datacenter.

Picture Courtesy of Stanley Rabinowitz, Math Pro Press.

In the SaaS world, assuming you include that in “cloud”, you have zero control and very little insight. If you app is not performing, you’ll have to talk to the vendors’ staff to (hopefully) get them to resolve issues.

But is the problem any worse in the cloud than in the datacenter? I would have to argue no. Your ability to touch and feel the bits is reduced, but the actual problems are not. In a pureplay public cloud deployment, the performance of an application is heavily dependent upon your vendor, but the top-tier vendors (Amazon springs to mind) can spin up copies as needed to reduce workload. This is not a far cry from one common performance trick used in highly virtualized environments – bring up another VM on another server and add them to load balancing. If the app is poorly designed, the net result is not that you’re buying servers to host instances, it is instead that you’re buying instances directly.

This has implications for IT. The reduced up-front cost of using an inefficient app – no matter which of the four pillars it is inefficient in – means that IT shops are more likely to tolerate inefficiency, even though in the long run the cost of paying monthly may be far more than the cost of purchasing a new server was, simply because the budget pain is reduced.

There are a lot of companies out there offering information about cloud deployments that can help you to see if you feel blind.

Fair disclosure, F5 is one of them, I work for F5. That's all you're going to hear on that topic in this blog.

While knowing does not always directly correlate to taking action, and there is some information that only the cloud provider could offer you, knowing where performance bottlenecks are does at least give some level of decision-making back to IT staff. If an application is performing poorly, looking into what appears to be happening (you can tell network bandwidth, VM CPU usage, VM IOPS, etc, but not what's happening on the physical hardware) can inform decision-making about how to contain the OpEx costs of cloud.

Internal cloud is a much easier play, you still have access to all the information you had before cloud came along, and generally the investigation is similar to that used in a highly virtualized environment. From a troubleshooting performance problems perspective, it's much the same. The key with both virtualization and internal (private) clouds is that you're aiming for maximum utilization of resources, so you will have to watch for the bottlenecks more closely – you're "closer to the edge" of performance problems, because you designed it that way.

A comprehensive logging and monitoring environment can go a long way in all cloud and virtualization environments to keeping on top of issues that crop up – particularly in a large datacenter with many apps running.

And developer education on how not to be a resource hog is helpful for internally developed apps. For externally developed apps the best you can do is ask for sizing information and then test their assumptions before buying.

Sometimes, cloud simply is the right choice. If network bandwidth is the prime limiting factor, and your organization can accept the perceived security/compliance risks, for example, the cloud is an easy solution – bandwidth in the cloud is either not limited, or limited by your willingness to write a monthly check to cover usage. Either way, it's not an Internet connection upgrade, which can be dastardly expensive not just at install, but month after month.

Keep rocking it. Get the visibility you need, don't worry about what you don't need.

Related Articles and Blogs:

- Don MacVittie - Load Balancing For Developers
- Advanced Load Balancing For Developers. The Network Dev Tool
- Load Balancers for Developers – ADCs Wan Optimization ...
- Intro to Load Balancing for Developers – How they work
- Intro to Load Balancing for Developers – The Gotchas
- Intro to Load Balancing for Developers – The Algorithms
- Load Balancing For Developers: Security and TCP Optimizations
- Advanced Load Balancers for Developers: ADCs - The Code
- Advanced Load Balancing For Developers: Virtual Benefits
- Don MacVittie - ADCs for Developers
- Devops Proverb: Process Practice Makes Perfect
- Devops is Not All About Automation
- 1024 Words: Why Devops is Hard
- Will DevOps Fork?
- DevOps. It's in the Culture, Not Tech.
- Lori MacVittie - Development and General
- Devops: Controlling Application Release Cycles to Avoid the ...
- An Aristotlean Approach to Devops and Infrastructure Integration
- How to Build a Silo Faster: Not Enough Ops in your Devops

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