

Against the Grain. Private Clouds as Cost Containment



Don MacVittie, 2010-24-08

My father was an antique dealer that specialized in furniture refinishing. All of us children spent some amount of time down at the shop getting instruction in how to handle antiques from dishes to weapons to furniture. But each of us got special instruction in how to treat a piece of furniture. The man looked at a piece of broken down furniture with a critical eye, and then caressed it like it was special, he could recover some of the most horrifically damaged furniture with nothing but experience and trial-and-error. The one lesson all of us received over and over, sometimes violently, was *never go against (across) the grain*. Not scraping furniture stripper and the toxic goo it created off with a putty knife, not with steel wool, and certainly never with sandpaper. Doing so damages the patina, the layer of wood at the top that contains the original finish and makes the grain stand out.



I've found in my career in Computer Science that this advice, so sound in woodworking, is just not as sound in a field where things are constantly changing and many marketing type folks are selling the latest, greatest thing. Sometimes you have to look at what is commonly believed – or someone has paid big money to make you think is commonly believed, and consider what's best in real life.

So goes the “public/private” cloud debate, which has been declared dead by those with a vesting interest in one side or the other. For those of you who are out there trying to make this stuff work, the “debate” really isn't dead. Instead it is something you're considering on a daily basis to determine what is best for you and your organization. Increasingly, IT staff are doing what many of us expected – implementing cloud internally, at least for some solutions, and watching cloud externally, but not diving in other than specific use like testing with scrubbed data.

We here at [F5](#) have no real stake in this consideration. With [iRules](#), [iControl](#), [TMSH](#), and other [ADC](#) functionality, we're a good fit in either scenario. You, on the other hand, have a huge stake in this decision. Your very business and the adaptability of your IT department depend upon the choices you make. And private cloud answers the two biggest issues in cloud computing – security and availability. It also answers what I am certain will become the third big question in cloud computing... Cost Containment.

The pricing models for cloud providers make the pricing for CDNs look simple. Throughput, servers, connections, phase of the moon... You name it, you're getting billed for it. As I've said before, the problem is not that it will cost too much, it is that the costs cannot be reliably predicted. While it sounds great to say that your costs will grow at the rate of your usage, it is rarely a linear progression, and the costs are wildly variable. But with internal cloud, you really do grow at the rate of usage. The significant difference is that you have to own the hardware, but the fact is that if you want to control the data, you *need* to own the hardware it resides on. There are some applications for which you do not need 100 percent control of the data, and those will do just peachy keen on the public cloud, but there are some for which the data leaking is a liability – either a competitive liability, a PR liability, or worst of all, a legal liability. Those apps are the ones that internal cloud is best suited to.

With [F5](#) and [VMWare](#) products, you can get a long way to automated internal IT. Since some universities are already doing automated internal on-demand self service IT – albeit with a more technical audience than most enterprises – the tools are there, you have but to decide how far you want to go and what amount of hardware you want to invest in. Of course all of the caveats about thin provisioning and virtual sprawl and automation weaknesses apply, but they're surmountable problems.

There are costs of course, you have to move hardware to a cloud model before you can transition applications, which implies you have more hardware on hand than you need to handle traffic during the switch over. You'll have to have two sets of SOPs while you're making the transfer, one for the old and one for the new. And if you've already virtualized a lot of systems, then you'll need three – one for hardware, one for virtualization, one for cloud. But this is nothing new, we have commonly worked under different sets of procedures – like one for Windows XP and one for Apple OS, for example. It's just a question of figuring out the process for a user to request and receive an instance of an application. And how it grows... With a device like [BIG-IP LTM](#), you can know when the system is near capacity and another instance is required, you can add the new instance to the pool of available instances once it is brought up, and remove it from the pool of available instances when it is decommissioned. With [VMWare vCenter AppSpeed](#), you have an even more complete look at the application and its performance. But all of this does require some development work, we're not yet to the point where it just automatically grows and shrinks on demand. We are so close that you can implement it in your private cloud with a little bit of work though.

If, while you're setting up your private cloud, you segregate the cloud from the rest of the DC, you can start to look into how you're going to manage the overall conglomeration of IT with cloud services, start to understand costs and bottlenecks, and prepare for the day when much more of your application infrastructure might be on the public cloud. You can consider security implications by looking at the weak points in your internal "private" cloud and determining steps to remedy application and data storage security weaknesses. All while becoming more agile.

That's against the grain for many, there is this push to convince you that your entire world just changed, but it didn't. You know it, I know it, the goals and needs of IT are very much the same, you just have another tool at your disposal for delivering the IT services that the business needs on a timeline even more suited to the business. But you'll still have apps, databases, mail servers that consume a seemingly endless stream of disk – no matter where that disk sits – a LAN and a larger list of projects than you can hope to accomplish. And that's okay, it's a job because it requires work, not because it just became "so simple a cave man can do it".



Related Articles and Blogs

- [Building a private Cloud](#)
- [Virtualization, Cloud, Automation, but where are we going?](#)
- [Cloud Players Unite for Enterprise Private Cloud Initiative](#)
- [Building a Private Cloud](#)
- [So You Put An Application in the Cloud, Now What?](#)

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