



Options for Prospective Hosting Providers

From the Desk of:
Sid Herron
Dir. of Channel Management
VirtualQube
(425) 939-2700
www.VirtualQube.com

If you're a Value-Added Reseller (VAR) or Managed Services Provider (MSP), you've been hearing voices from all quarters telling you that you've got to get into cloud services:

- The [451 Research Group has estimated](#) that, by 2015, the market for all kinds of "virtual desktops" will be as large as \$5.6 Billion. IDC estimates that the portion of these virtual desktops sourced solely from the cloud could be over \$600 Million by 2016, and growing at a more than 84% annually.
- John Ross, technology consultant and former CTO of GreenPages Technology solutions was quoted in [a crn.com article](#) as saying, "This is the last time we are going to see hardware purchases through resellers for many, many years." He predicts that 50% of the current crop of resellers will either be gone or have changed to a service provider model by 2018.
- The same article cited research by UBM Tech Channel (the parent company of CRN) which indicated that "vintage VARs" that stay with the current on-premises model will have to add at least 50% more customers in the next few years to derive the same amount of sales, which will require them to increase their marketing budgets by an order of magnitude.
- Dave Rice, co-founder and CTO of TrueCloud in Tempe, AZ, predicted in the same article that fewer than 20% of the current crop of solution providers will be able to make the transition to the cloud model. He compares the shift to cloud computing to the kind of transformation change that took place when PCs were first introduced to the enterprise back in the 1980s.

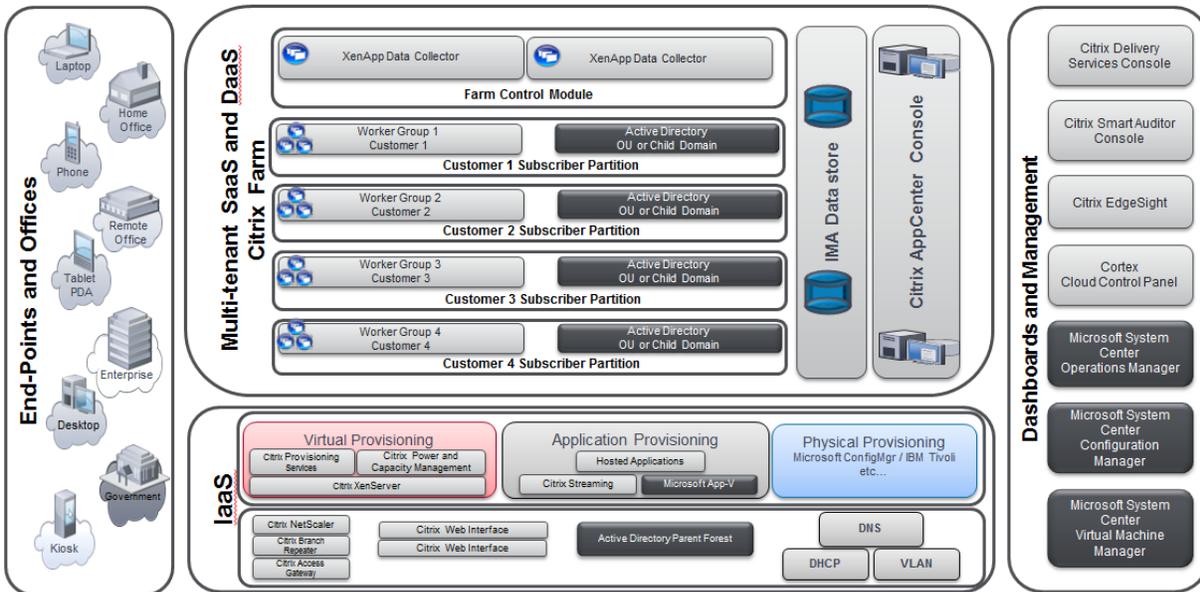
If you place any credence at all in these predictions, it's pretty clear that you need to develop a cloud strategy. But how do you do it?

First of all, let's be clear that, in our opinion, selling Office 365 to your customers is not a cloud strategy. Office 365 may be a great fit for some customers, but it still assumes that most computing will be done on a PC (or laptop) at the client endpoint, and your customer will still, in most cases, have at least one server to manage, backup, and repair when it breaks. Moreover, [you are giving up a great deal of account control, and account "stickiness," when you sell Office 365.](#)

In our opinion, a cloud strategy should include the ability to make your customers' servers go away entirely, move all of their data and applications into the cloud, and provide them with a Windows desktop, delivered from the cloud, that the user can access any time, from any location where Internet access is available. (Full disclosure: That's precisely what we do here at VirtualQube, so we have an obvious bias in that direction.) There's a pretty good argument to be made that [if your data is in the cloud, your applications should be there too](#), and vice versa.

The best infrastructure for such a hosting environment (in the opinion of a lot of hosting providers, VirtualQube included) is a Microsoft/Citrix-powered environment. Currently, the most commonly deployed infrastructure is Windows Server 2008 R2 with Citrix XenApp v6.5. Microsoft and Citrix both have Service Provider License Agreements available so you can pay them monthly as your user count goes up. However, once you've signed those agreements, you're still going to need some kind of hosting infrastructure.

Citrix can help you there as well. Once you've signed up with them, you can access their recommended "best practice" reference architecture for Citrix Service Providers. That architecture looks something like this:



When you've become familiar enough with the architectural model to jump into the deep end of the pool and start building servers, your next task is to find some servers to build. Broadly speaking, your choices are:

1. Buy several tens of thousands of dollars (at least) of server hardware, storage systems, switches, etc., secure some space in a co-location facility, rack up the equipment, and start building servers. Repeat in a second location, if geo-redundancy is desired. Then sweat bullets hoping that you can sign enough customers to not only pay for the equipment you bought, but make enough profit that you can afford to refresh that hardware in three or four years.
2. Rent hardware from someone like Rackspace, and build on their platform. Again, if you want geo-redundancy, you're going to need to pay for hardware in at least two separate Rackspace facilities to insure that you have something to fail over to if you ever need to fail over.
3. Rent VMs from someone like Amazon or Azure. Citrix has been talking a lot about this lately, and has even produced some helpful pricing tools that will allow you to estimate your cost/user/month on these platforms.
4. Partner with someone who has already built it, so you can start small and "pay as you grow."

Now, in all fairness, the reference architecture above is what you would build if you wanted to scale your hosting service to several thousand customers. A wiser approach for a typical VAR or MSP would be to start much smaller. Still, you will need at least two beefy virtualization hosts – preferably three so if you lose one, your infrastructure is still redundant – a SAN with redundancy built in, a switching infrastructure, a perimeter firewall, and something like a Citrix NetScaler (or NetScaler VPX) for SSL connections into your cloud.

Both VMware and Hyper-V require server-based management tools (vCenter and System Center, respectively), so if you've chosen one of those products as your virtualization platform, don't forget to allocate resources for the management servers. Also if you're running Hyper-V, you will need at least one physical Domain Controller (for technical reasons that are beyond the scope of this article). Depending on how much storage you want to provision, and whose SAN you choose, you're

conservatively looking at \$80,000 - \$100,000. Again, if geo-redundancy is desired, double the numbers, and don't forget to factor in the cost of one or more co-location facilities.

Finally, you should assume at least 120 – 150 hours of work effort (per facility) to get everything put together and tested before you should even think of putting a paying customer on that infrastructure.

If you're not put off by the prospect of purchasing the necessary equipment, securing the co-lo space, and putting in the required work effort to build the infrastructure, you should also begin planning the work required to successfully sell your services: Creating marketing materials, training materials, and contracts will take considerable work, and creating repeatable onboarding and customer data migration processes will be critical to creating a manageable and scalable solution. If, on the other hand, this doesn't strike you as a good way to invest your time and money, let's move on to other options.

Once you've created your list of equipment for option #1, it's easy enough to take that list to someone like Rackspace and obtain a quote for renting it so you can get a feeling for option #2, so we won't spend a lot of time and space here on that option. Instead, let's move on to option #3.

Citrix has spent a lot of time lately talking about Amazon as a platform for Citrix hosting providers. There was an entire breakout session devoted to this subject at the 2014 Citrix Summit conference in Orlando. So it may be productive to spend a little time looking at the cost model that Citrix has provided for building a hosting infrastructure on AWS. At first glance, the costs may look fairly attractive. But there are a number of assumptions built into the cost model that should make any aspiring service provider think twice:

- It assumes that you've got enough users lined up that you can get the economies of scale from building an infrastructure for several hundred, if not thousands, of users.
- It assumes that you've got enough free cash to pay up front for 3-year reserved instances of all the servers you'll be provisioning.
- It assumes that, on average, your servers will need to run only 14 hours per day. If your customers expect to be able to work when they want to work, day or night, this will be a problem.
- It assumes that you will be able to support an average of 150 concurrent users on a XenApp server that's running on a "Cluster Compute Eight Extra Large" instance. Anyone who has worked with XenApp knows that these assumptions must be taken with a very large grain of salt, as the number of concurrent users you can support on a XenApp server is highly dependent on the application set, and doesn't necessarily scale linearly as you throw more processors at it.

If all of these assumptions are correct, the Citrix-provided spreadsheet says that you can build an infrastructure at AWS that will support 1,000 concurrent users (assuming 10 customers with 100 users each for the multi-tenancy calculation) for an average cost/user/month of \$45.94 over a three year period. But that number is misleading, because you have to come up with \$377,730 up front to reserve your AWS instances for three years. So your first-year cost is not \$551,270, but \$803,081 – that's actually \$66.92/user/month for the first year, and then it drops to \$35.45/user/month in years two and three, then back to \$66.92/user/month in the fourth year, because you'll have to come up with the reservation fees again at the beginning of year four.

There are a couple of other things about this model that are troublesome:

1. It assumes only a single file server for 1,000 users, meaning that you would administer security strictly via AD permissions. It also means that if anything happens to that file server, all of your tenants are impacted. If we instead provision ten file servers, so that each of the ten tenants has a dedicated file server, it bumps the average cost by roughly \$5/user/month.
2. If your user count is 100 users per tenant, but you're expecting to support 150 users per XenApp server, you'll obviously have users from multiple tenant organizations running concurrently on the same XenApp server. This, in turn, means that if a user from one tenant organization does something that impacts XenApp performance – e.g., launches the Production Planning Spreadsheet from Hell that pegs the processor for five minutes recalculating the entire spreadsheet whenever a single cell is changed – it will affect more than just that tenant organization. If we assume only 100 users per XenApp server, so that we can dedicate one XenApp server to each tenant organization, it bumps the average cost by roughly another \$1.50/user/month.

“But wait,” you might say, “not many VARs/MSPs will want to – or be able to – build an infrastructure for 1,000 users right off the bat.” And you would be correct. So let's scale it back a bit. Let's look at an infrastructure that's built for 250 users, and let's assume that breaks down into five tenants, with 50 users each. Let's further assume, for reasons touched on above, that each customer will get a dedicated file server, and one dedicated XenApp server. We'll dial those XenApp servers back to “High CPU Extra Large” instances, which have 4 vCPUs and 7.5 Gb of vRAM each. Your average cost over three years, still assuming 3-year reserved instances, jumps to \$168.28/user/month, and you must still be prepared to write a check for just over \$350,000 for the 3-year reservation fees. Why the big jump? Primarily because there is a minimum amount of “overhead” in the server resources required simply to manage the Citrix infrastructure, the multi-tenant Active Directory and Exchange infrastructure, etc., and that overhead is now spread across fewer users.

Now consider that all of the prices we've been looking at so far cover only the compute and storage resources. We haven't begun to factor in the monthly cost of Citrix or Microsoft Service Provider licensing. In round numbers, that will add another \$25/user/month or so to your cost, including MS Office. Nor have we accounted for the possibility that some of your users may need additional SPLA applications, such as Visio or Project, or that some tenants may require a SQL server or some other additional application server. Nor have we accounted for the possibility that some of your tenants may require access to the infrastructure on a 24x7 basis, meaning that their servers have to run 24 hours per day, not just 14.

This is why, at the aforementioned session at the recent Citrix Summit conference in Orlando, the numbers presented in the session were challenged by several people during the ensuing Q&A, the general feedback being that they simply didn't work in the real world.

Our approach is a little different.

First, we believe that VARs and MSPs need a platform that will do an excellent job for their smaller customers – particular those who do not have a large staff of IT professionals, or those who are using what AMI Partners, in a study they did on behalf of Microsoft, called an “Involuntary IT Manager” (IITM). These are the people who end up managing their organizations' IT infrastructures because they have an interest in technology, or perhaps because they just happen to be better at it than anyone else in the organization, but who have other job responsibilities unrelated to IT. Often these individuals are senior

managers, partners, or owners, and in nearly all cases could bring more value to the organization if they could spend 100% of their time on their primary job. Getting rid of on-site servers and moving data and applications to a private hosted cloud will allow these people to regain that lost productivity.

Second, we believe that most of these customers are going to need access to their cloud infrastructure on a 24/7 basis. Smaller companies tend to be headed by entrepreneurial people who don't work traditional hours, and who tend to hire managers who also don't work traditional hours. Turning their systems off for 10 hours per day to save on run-time costs simply isn't going to be acceptable.

Third, we believe that the best mix of security and cost-effectiveness for most customers is to have a multi-tenant Active Directory, Exchange, and SharePoint infrastructure, but to dedicate one or more XenApp server(s) to each customer, along with a file server and whatever other application servers they may require (e.g., SQL Server, accounting server, etc.). This is done not only for security reasons, but to avoid "noisy neighbor" problems from poorly behaved applications (or users).

In VirtualQube's multi-tenant hosting infrastructure, each customer is a separate Organizational Unit (OU) in our Active Directory. Each customer's servers are in a separate OU, and are isolated on a customer-specific vLAN. Access from the public Internet is secured with a common Watchguard perimeter firewall and Citrix NetScaler SSL/VPN appliances. Multi-tenant customers who need a permanent VPN connection to one or more office locations can have their own Internet port and their own firewall.

We also learned early on that some customers prefer not to participate in any kind of multi-tenant infrastructure, and others are prevented from doing so by security and compliance regulations. To accommodate these customers, we provision completely isolated environments with their own Domain Controllers, Exchange Servers, etc. A customer that does not participate in our multi-tenant infrastructure always gets a customer-specific firewall and NetScaler, and customer-specific Domain Controllers. At their option, they can still use our multi-tenant Exchange Server, or have their own.

Finally, we believe that many VARs and MSPs will benefit from prescriptive guidance for not just how to build a hosting infrastructure, but how to sell it. That's why our partners have access to a document template library that covers how to do the necessary discovery to properly scope a cloud project, how to determine what cloud resources will be required and how to price out a customized private hosted cloud environment, how to position the solution to the customer, how to write the final proposal, how to handle customer data migration, and much, much more.

We would respectfully submit that, for most VARs/MSPs, partnering with someone who has already built the infrastructure makes the most sense, and we would love to have them partner with us. We believe that partnering with VirtualQube makes sense for VARs and MSPs because that's the world we came from. Our hosting platform was built by a VAR/MSP for VARs/MSPs, and we used every bit of the experience we gained from twenty years of working with Citrix technology. We invite prospective partners to visit us on the Web at <http://www.VirtualQube.com/partners>.