

The five dangers of virtual servers

Vital questions you must ask before you commit your company to virtual servers

First, The Benefits Of Virtual Servers

Virtual servers and cloud computing are changing the way we think about hosting high-performance and high-availability applications.

The technology uses virtualisation techniques to abstract the underlying hardware from the end-product, i.e. the server itself.

In theory this means:

- Higher availability as there is no single point of failure (SPF) associated with any particular hardware;
- Lower costs, as an entire physical node is not dedicated to one server task any more, i.e. multiple virtual servers can run concurrently on one physical node;
- Lower energy consumption per server, for the same reasons as above;
- Quicker lead times on server upgrades as “virtual hardware” can be scaled up or down instantly or near-instantly.

How is a virtual server different from shared hosting?

It may seem that virtual servers are an evolution of low-end shared hosting, where multiple websites live on the same server. However this couldn't be further from the truth. True virtualisation enables each virtual machine to live concurrently on the same physical node, with its' own operating system, allocated memory and disk space, and share of CPU resources. Virtual servers co-existing on the same hardware are completely autonomous and private.

So what are the dangers?

As ever, with this new technology, the devil is in the detail. We've identified five major dangers you need to check before choosing your virtual server hosting service.

1. The Danger of Hardware Failure

One of the benefits of virtualisation should be an improvement of service availability, so it's important that the topology of the cluster is well thought out, such that a failure of a major central component such as a network switch, or even just the physical node (server) that your virtual server is running on, doesn't cause your virtual server to go offline.

Use SAN Storage

The best solution uses centralised SAN storage, so that if the physical node (server) you're running on fails, your virtual server will automatically re-start on a spare node.

An added benefit of this is that your hosting service can take nodes out of service for maintenance without affecting your virtual server. That means you're truly available for the maximum amount of time possible.

Look for an option of fast disks (i.e. 15,000rpm SAS disks) to keep the I/O throughput of your virtual server tip-top.

Use Hardware Resilience

Nodes should be dual-uplinked through two switches, to both the frontend (i.e. the internet) and also the backend (i.e. to the storage) such that the failure of a network switch or uplink cable (or just someone accidentally un-plugging it) does not stop your virtual server from running, or being accessible.

The centralised SAN storage should be built with resilient/failover power supplies and controllers, such that the failure of any component within the SAN does not take offline any virtual server.

A decent service provider will offer the option of offsite backup space to be mounted to your virtual server to give you a self-standing off-site copy of your important data.

2. The Danger of Interference from other Customers' Virtual Servers

There are several ways in which one virtual server can interfere with another, when they're running on the same node. However, when security is properly implemented, this risk is completely mitigated, and the solution is just as secure as a dedicated server.

A responsible host will have thought out these issues and put safeguards in place to ensure that one virtual server cannot impact upon another. Here are some questions to ask your potential virtual server provider:

- **Will I have my own VLAN or are all virtual servers on the same network segment?**
This is important, as virtual servers on the same VLAN are on the same broadcast domain. This means if someone enters your IP address in their network configuration incorrectly, they could take your virtual server offline. It also means you'd receive their broadcast traffic. Finally, it means that there is complete lack of firewalling between your virtual server and others, unless a software firewall is put in-place (see the next point for more on this).
- **Will I have the protection of a hardware firewall?**
With a separate VLAN, it's possible to protect a virtual server in the same way as you would a dedicated server. This means that not only traffic from the outside world, but also to other neighbouring customers, is allowed through only if it meets the firewall rules you've set.
- **Are the kernels shared with the host operating system?**
Slightly more complex this one; basically some virtualisation technologies share the host node's base operating system kernel with the virtual servers, for various reasons, usually so the provider can put more virtual servers on a node. However this means that any kernel vulnerability could mean that a compromised virtual server can effectively take down the whole of the host node, which basically means all the virtual servers running on the node will be taken down too.

3. Is there a proper SLA for the virtual server; and what is a "proper" SLA?

It's one thing to say that a service is reliable, but quite another to put a money-back guarantee or Service Level Agreement (SLA) on it.

Obviously what you want is a working service, not a money-back guarantee, but it certainly puts an emphasis on the provider having a standard to work towards.

Many providers say SLAs are a waste of time, and marketing hype. If they truly believed this, surely they'd offer an SLA anyway as they have nothing to lose?

Some virtual server products come without an SLA; some even charge you extra for an SLA. In short, any decent high-availability virtual server offering should have an SLA on both the network connectivity, but more importantly the availability of the "virtual hardware" of your virtual server.

A decent SLA should provide a realistic and achievable service level (100% SLAs are usually vanity only) and a penalty for the host if the SLA is not achieved. It should also be clear what the SLA covers. For example, a lot of network/connectivity SLAs cover only the provider's own network in their data centre, and not their upstream providers. That means your server may be inaccessible but if it was due to their provider and not their own internal network, you're not covered.

4. How contended are the servers, and why does that matter, are there any performance guarantees?

There's a potential for over-selling with virtual servers, as the provider may not tell you how many virtual machines they intend running on one physical node, whether memory and CPU time are contended, or how fast the physical node is uplinked to the network.

The whole point is to put limitations in-place to stop one virtual server from hogging the resources at a performance cost to neighbouring virtual servers.

A responsible provider uses a virtualisation technology that does not allow memory or disk space to be contended, and has fixed parameters in place to ensure CPU and network resources are fairly shared out. This includes setting an upper limit to the number of virtual machines that can run on a physical node, and also ensuring the physical nodes are uplinked to the 'net at a suitable speed such that every server gets a decent sized connection. For example a gigabit connection shared between 30 virtual machines gives approximately 30Mbps to each virtual machine.

5. Does the customer get console access and remote reboot facility?

You would generally want KVM over IP and remote reboot facilities on a dedicated server, to give you "sat in front of the machine" access in the case of a major OS failure, or to correct network settings when the machine is otherwise inaccessible.

The same should apply to a virtual server. This technology can be a life-saver when you need to work on a virtual server at 2am without waiting for your provider to respond to help you out. It gives you complete self-sufficiency.

Conclusion

There's plenty of scope for virtual server providers to cut corners. That said, if you research your provider well, using the above questions as part of your decision-making process, you're likely to find a service that's a high-availability and high-performance alternative to dedicated servers.

To sum up make sure that any virtual server provider can meet the following criteria and you've done your utmost to mitigate the dangers we've described:

1. Uses SAN Storage
2. SAN and host servers have redundant critical components
3. Nodes dual-uplinked through two switches (front-end and backend)
4. You have your own VLAN
5. Provides a hardware firewall
6. No kernel sharing between host node and virtual machines
7. A reasonable SLA
8. Console access with reboot facility

And now a little about Melbourne...

Started in 2000, Melbourne now owns and runs two data centres and is building a third to cope with client demand. Overall the centres host over 2,000 servers and are used by such companies as The Chartered Institute of Tax, JJB Sports PLC, Travel Counsellors and Freedom Finance PLC.

Obviously web sites and ecommerce applications must be available 24x7x365 year on year so Melbourne provides redundancy on practically every part of the server, power supply and Internet access provision.

The data centres all operate with multiple internet connectivity providers to ensure Internet redundancy. All servers have dual power supplies and with the firm intention to stay in play even if there's a prolonged power outage Melbourne has a regularly tested UPS and generator set on standby that can power the entire facility for 7 days by itself.

Melbourne's Virtual Dedicated Server offering, called UltraCloud™, was launched in early 2009, providing robust reliability with an affordable price.

If you've any questions about this report or want to have a tour of our Manchester data centres please phone us on 0800 915 8771 or visit our website at www.melbourne.co.uk

Melbourne Network Solutions Ltd | Registered office: Turing House, Archway, Manchester, M15 5RL