

Virtual servers; real problems

By Chris Procter.

Chris examines how businesses can realise the benefits of virtualization without risking downtime due to hardware and application failures... by consolidating your servers you have also consolidated any issues that may affect them, increasing the scope of a failure from a single machine to all the virtual machines hosted on the same machine.

Probably the biggest of the latest crop of buzzwords is virtualization. The big win of virtualization technologies is the ability to replace many servers with a single box running multiple virtual machines. Most machines have far more resources (CPU, RAM etc.) than they normally need, this is because they are specified by what they are expected to need at their peak use. Careful consolidation of servers onto virtual machines can allow better use to be made of unused capacity which is otherwise only available for SETI research.

Consolidating multiple real servers into a single box has other advantages too, it reduces power consumption and cooling requirements, it simplifies network infrastructure and other such ancillary services (less machines means fewer ports on the switch, fewer KVM ports, fewer racks etc). Management is also simplified, by tying everything together on a single box, resources can be assigned as needed.

However, by moving all your machines onto a single machine you are proverbially putting all your eggs in one basket. Whereas previously a hardware failure would only affect a single machine, now a hardware failure on the physical server will cause a cataclysmic failure as the entire virtual machine dies. So by consolidating your servers you have also increased the scope of a failure from a single machine to all the virtual machines hosted on the same machine.

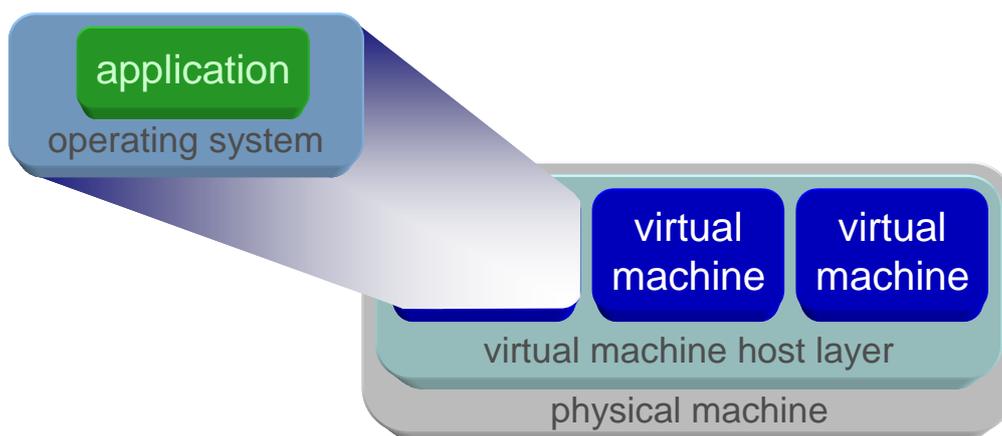


Fig 1: Virtual infrastructure with multiple o/s and applications on a single machine. Provides consolidation, but increases impact of hardware failure.

The truth is that virtualization reduces the availability of the services that run on the virtual machines.

Effects of virtualization:

- increased scope of system failures;
- powering down (eg to add new hardware) will involve shutting down all virtual machines;
- one virtual machine can hog shared resources (eg CPU, memory) and cause all servers to run more slowly.

These problems can be reduced by careful planning of the virtual environment and choosing the right virtualization technology for the task. Some solutions allow the administrator to limit the physical CPU usage of each virtual machine for example. But there is a fundamental single point of failure in the physical machine, the more functionality you build on top of this single foundation the more significant the consequences of losing it, and no additional functionality that also relies upon it will ever solve that problem.

So how do you gain the benefits of virtualization without opening your business up to failures and outages?

High availability clustering can be used to run an application on one machine (normally known as the primary or active machine) with one or more machines (known as the secondaries or backups) acting as warm backups. Cluster management software monitors the members of the cluster and if the application dies, due to either the application hanging or the machine itself dying, it brings it into service on the backup machine, allowing clients to continue as if nothing had happened.

High availability clustering and virtualization can complement each other in a number of ways.

The most obvious combination of the two technologies is to use two servers, each with one or more virtual machines, one server running the application and the second providing the warm backup, much like a traditional HA cluster. As each of the virtual machines is running on a separate physical machine, this removes the single point of failure. Running multiple virtual machines on the primary each clustered with a secondary virtual machine on the second physical machine, reduces an entire server infrastructure to two physical boxes while increasing application availability.

Of course, there's nothing that requires all of the virtual machines to use the same physical machines as their primary. It's quite possible to have three physical machines each hosting three virtual machines, backed up by a single physical machine running a total of nine virtual machines (one for each of the virtual machines on the primaries).

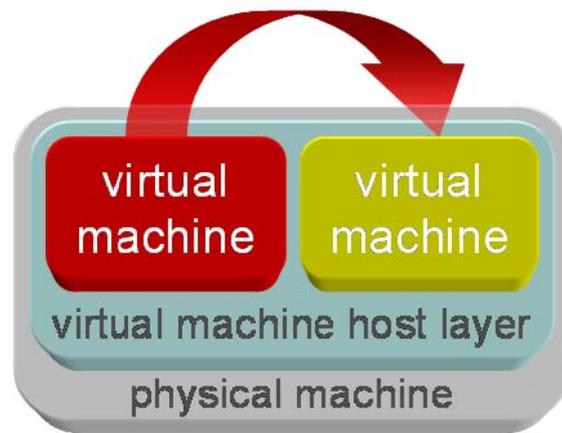


Fig 2: Failover between virtual machines on different physical machines. This can provide protection against hardware and application failure.

While if all nine virtual machines on the backup were running applications at the same time performance could be adversely affected. This is a very unlikely scenario (requiring all three primary physical machines to die at the same time) and the cost savings of consolidating nine servers down to four while increasing availability more than compensates. HA clustering takes away the weaknesses of a virtual environment by providing a more resilient solution to cater for failures.

A refinement on this solution comes with the realisation that, coupled with real time data replication, there is no requirement for the backup server to be in the same physical location as the primary. This opens up the possibility of running the three primary servers from the previous example on the company's site while having the backup hosted in a data centre in another building, another town, even another country. In a purely physical infrastructure hosting nine servers locally and nine in a third party data centre is an expensive proposition, a single off site box running nine virtual machines however is much cheaper and can provide not only resilience against an application, operating system or hardware failure, but also protection against losing the company building to a fire or other disaster. In these situations the ability to get back up and running immediately with all the critical services, each with an up to date copy of the data can make the difference between recovery and disaster.

Of course, there's nothing to say that you need to use virtual machines. If there is a pre-existing server infrastructure, it may prove simpler to cluster them with off-site virtual machines running on a single box to provide disaster recovery protection. This gives the advantages of a high availability cluster in terms of application uptime and disaster recovery protection, without the disruption and expense of having to migrate existing services and data to new, virtual servers, or having to purchase and host an identical physical server.

The final method of high availability clustering virtual machines comes from the realisation that virtual machines are just applications running on an underlying host operating system. This means that with the right clustering software it is possible to fail not just an application running within a virtual environment over onto a backup box but to fail the virtual machine itself over. In the event of a failure of the primary physical machine the entire virtual machine is started up on the backup, the

advantage of this is that because it is the entire virtual machine that is started, the configuration is guaranteed to be the same.

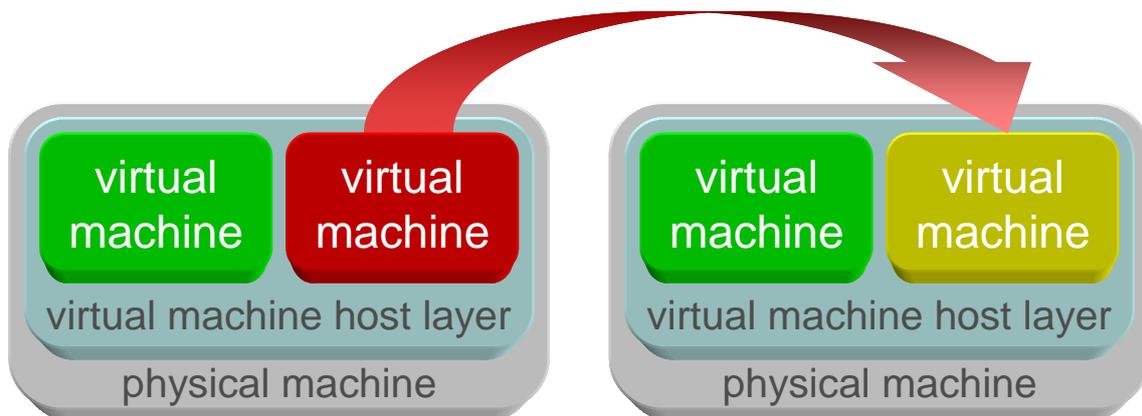


Fig 3: Failover from virtual machines on multiple physical machines to a single physical backup. Multiple physical or virtual machines can failover to virtual machines on a backup physical machine.

The disadvantage of monitoring the virtual machine rather than the application is that if the application crashes or hangs it will not be detected. This makes this solution ideal for those situations where many changes are made on a machine that is not confined to a single application.

Virtualization is trendy, and by improving operation of expensive computer resources, it's also very useful. Its real downside is that by reducing the hardware/services ratio it also reduces the resilience of those services.

This disadvantage means that virtual technology should not be adopted without first fully considering consequences and devising a recovery process. However, by clever use of the existing and well-tried methods of high availability clustering these disadvantages can be nullified and even turned into advantages by providing cheap and effective disaster resilience.

Chris Procter is solution architect for Open Minds, SteelEye Technology's Competence and Support Centre in the UK (www.steeleye.com)
