Grid Computing Now!
Knowledge Transfer Network

Comment Article

quocirca

# Grid Computing Now! – Grid: Virtualisation

By Clive Longbottom, Service Director, Quocirca Ltd

*This is the third in a series of articles written for the Grid Computing Now! Knowledge Transfer Network*

The whole concept of grid computing is predicated on the idea of virtualisation – and yet we at Quocirca keep finding that the term means completely different things to different people.

Are we talking here about network virtualisation, or about hardware, operating system or application virtualisation? How about storage virtualisation, or database virtualisation?

And here lies the rub; grid computing is something that falls under that wonderful term of a "holistic" solution – it includes everything. And yet, internal IT departments tend not to be set up that way –we have network professionals, storage experts, application developers, hardware specialists and so on.

It's like the discussion that goes on for 10 minutes between a network specialist and a banker – both think they know what the conversation is about until one or the other really figures out that the ATM they are talking about is not a common concept, and that the banker's ATM (Automatic Teller Machine) may well be connected to the rest of the system by ATM (Asynchronous Transfer Mode), but that's as far as any commonality goes.

Even for those techies that understand that virtualisation is a many-faceted term, there still seems to be some confusion over the two high level aspects of virtualisation – federation and partitioning. By federating, we bring multiple different resources together and make them appear to be a single resource. Once we have this single resource, we can then partition it at will to appear as completely different logical items, such that two application servers can appear to be 5, or that a 1Tb and a 4Tb storage system can appear as a 2Tb and a 3Tb system, for example.

If a bunch of techies get confused about the term "virtualisation", then there's little hope for grid computing as a mainstream concept – so here goes with an attempt to clear the decks with a high level, yet definitive, virtual tour.

**Network virtualisation**

So, today's no-brainer: the vast majority of us are already on a virtualised network – we use TCP/IP, a packet based system that is not dependent on point-to-point direct connections, but virtualises all possible connections to provide the network that we are all used to seeing as a cloud within network diagrams. However, for the sake of security, we keep attempting to undermine this concept by utilising point-to-point tunnelling technologies for virtual private networks – but that's a different story.

**Hardware virtualisation**

A tougher nut to crack. At the server level, we have a highly dynamic environment, where we may be looking at silicon-level virtualisation (a la IBM's Power5 chip, which has the inherent capability to virtualise and then partition itself, and of the new capabilities within Intel and AMDs CPUs), or we have the capability to layer on a hypervisor, such as EMC's VMware or Microsoft's Virtual Server, such that we can layer multiple operating systems on to the same piece of hardware. However, what we really need is the way to create a single resource pool of compute capability, where all CPUs and all memory are available to all services – and this needs careful managing through a real grid control mechanism. Newer server blade architectures are essentially predicated on having a suitable virtualisation capability, and it's here where we see the main focus from vendors on creating suitable tooling.

**Storage virtualisation**

Taking all available storage resources and making them appear as a single pool has been talked about for many years. However, the continuing proprietary nature of storage hardware has mitigated against this becoming a reality, but with newer storage area networks

(SANs) and network attached storage (NAS) architectures, it is getting closer. Again, blade architectures are evolving to remove disk storage as a separate entity from the basic CPU blades, so creating a fully poolable storage resource.

**Database virtualisation**

Information is the lifeblood of an organisation, and if it is held as a set of disparate silos, it is difficult to gain full value from it. Virtualising the databases, such that all data appears to be within a single repository, and so making the information held available as a service to the organisation, means that the value can be more easily gained. Also, database virtualisation means that a more complete picture is available of all the information that is required to make a decision – and the quality of the decisions is therefore improved.

**Operating system virtualisation**

Using hypervisor technologies (again, such as VMWare or Virtual Server), it is possible to install multiple instances of an operating system on to a single piece of hardware – or on to logical units of hardware that have previously been virtualised. Also, through paravirtualisation, it is possible to create a similar end result using hypervisors such as Xen and L4.

**Application or service virtualisation**

Once we have an underpinning of virtualised platform, such as the network, the servers, the storage and the operating system, we need to apply the desired functionality in such a way that it makes the most of the resources available to it. This can either be done by abstraction, where the underlying assets are made to seem "real" to an existing application, or we can look to writing applications specifically to make the most of virtualisation, with fully dynamic capabilities for provisioning and for resource allocation as required. The best way to do this is to go for a service-based approach using a service oriented architecture (SOA).

Historically, grid computing has been more focused on providing a shared resource capability for a single application. The move is towards supporting a more varied set of workloads within a grid architecture, such that commercial usage of grids becomes the rule, rather than the exception. When we look at the growth of grid computing in the commercial sector, we see the need for a highly virtualised environment where services can be provisioned at will, the resources that each service requires can be called and created dynamically, new hardware can be introduced transparently without the need for downtime, and where virtualisation provides built in business continuity through fully redundant capabilities in the total stack.

But, we need a common vocabulary, and just talking about "virtualisation" may be misleading and counterproductive to where we need to go. Best to get everyone concerned into a room and agree that virtualisation applies to all aspects of the IT infrastructure – and that those involved need to be more aware of how their piece of the jigsaw fits in to the desired end result.

The lack of one virtualised layer within the infrastructure will heavily impact the capabilities and performance of a grid aimed at providing a single infrastructure for the running of multiple composite applications within an organisation, but each virtualisation capability needs to be fully managed, and its impact on other environments known and controlled. A further article in this series will look at how this management should be carried out.

 

## About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, Microsoft, IBM, Dell, T-Mobile, Vodafone, EMC, Symantec and Cisco, along with other large and medium sized vendors, service providers and more specialist firms.

Details of Quocirca's work and the services it offers can be found at
http://www.quocirca.com