



Who grasps virtual

There's a fast-changing world of virtualisation out there that's making server and storage consolidation a whole lot more attractive, even for SMEs. Brian Tingham cuts through the hype and misinformation to the real deals

everybody who's anybody is saying it: even in quite small organisations IT managers would be well advised to investigate the potential for savings and efficiencies through consolidation and virtualisation – of their servers, storage, operating systems and whole IT networks. Virtualisation meaning decoupling the physical infrastructure hardware from the software and data it runs by means of an abstraction layer that provides for policy-based management of shared resources.

The business case involves everything from reduced server real estate and complexity to simplified admin, improved utilisation, possibly also resilience and certainly flexibility – including with archiving and back-up. It's about making the infrastructure cheaper to buy, own, run and manage, as well as better able to respond to changing business requirements, even on the fly.

Which is what we want. Manufacturers, like any business, can't be comfortable in the knowledge that their IT departments are working hard just to keep the lights on. We're employing expensive people running expensive kit, so the balance surely needs to shift to that resource focused more on business improvement. Just as important, with average server utilisation often just 10% but some machines running at near 100%,

and primary storage creaking at the seams, the idea of being able to flex all that – hopefully automatically – so it's available when and where it's wanted, sounds ideal.

And the good news is you absolutely can do this – although you're going to come across plenty of people saying you can't, or shouldn't. As always, the devil is in the detail but – primarily because of an unhealthy mix of vested interests, caution and sales people's blind beliefs – sifting the hype from the disinformation, misinformation and finally the reality is a challenge.

Can it be done?

Server virtualisation, for example, is well known and has been being implemented for several years, but you'll hear from many quarters that it's still early days for the same in storage. If by storage virtualisation you understand that all memory, irrespective of type and vendor, can be made available to all applications and users and managed as a cohesive pool, you need to wait, they say. And that's especially the case if you want that 'pool' to include secondary storage – hiving data off expensive, high availability arrays onto archive disk or removable media, even VTLs (virtual tape libraries) as its value and/or immediacy diminishes.

Mark Blowers, for example, senior researcher with

respected analyst Butler Group, puts it like this: “Server virtualisation is well up to the job and a lot of people are adopting it: you can deploy on standard software – just as Unix systems and mainframes have been doing for a long while, but with the likes of VMware for industry standard servers. But storage virtualisation is slightly different. Yes there are SANs [storage area networks] but we’re not quite there yet in terms of creating true virtual pools of storage.”

And Frank Timons, an analyst at asset management firm Robert Baird, adds: “CIOs see virtualisation as important but they’re talking about server virtualisation, not storage virtualisation... That’s not really selling yet. There are pilot studies but end users are saying ‘I don’t believe it and I don’t have time to investigate it.’”

To be fair, neither of them denies the existence of technologies’ that can fully virtualise storage. However, if I was a would-be user worrying about where to bet my IT budget, I wouldn’t rush into a project with that kind of testimony. I’d focus back on the server side.

But I’d be missing a trick. More to the point, so might you. Because there is a lot to go at here, and while some of it might be accused of being relatively

space?

new, there are big brains behind the developments and some big name adopters quietly singing its praises. Which, since you’re going to have to at least centralise your storage to make a good fist of virtualising your servers, has to be worth knowing.

I’m not just talking about virtualising primary storage here. We’ve been doing that for years with SANs (fibre channel and latterly also SCSI – where you want disk volumes to be available anywhere) and NAS boxes (network attached storage – for giving users access to applications across the network). Although user set-up issues have been massively improved, you’re still locked into a particular vendor with these.

[Note: the SCSI and iSCSI (Internet SCSI for IP/Ethernet-based storage) story is important here because it opens the SAN world to manufacturing SMEs that don’t have, or want to buy, specialist fibre channel expertise to implement and maintain these systems – find out more from EqualLogic and HP, with its All-in-One iSCSI networked storage.]

Nor am I just talking about the similar systems for virtualising secondary storage for replication, mirroring, disaster recovery, archiving and the rest. [Again there are improvements worth noting, mostly around enhancing the protection of secondary storage and re-thinking its potential for productive use, for example in business continuity by auto-copying data to multiple storage types and locations in the virtualised pool – find out more from Bridgehead Software.] And I’m also not just banging the drum for virtualised clustered file systems, like those from Polyserve, that enable many

servers to mount the same, variable size file system and vice versa – flexible, redundant utility computing.

No, you may be surprised to learn that actually there are now perfectly good systems for fully virtualising all your storage – and that the technologies have been around for years. So you can in fact create that holy grail of a flexible pool of storage despite what some of the big boys fail to tell you. What’s more you won’t find it expensive or difficult: quite the reverse. Today, you can match virtually any business requirement and IT preference with storage virtualisation, quickly building on and repurposing your existing investment without prejudicing your IT staff’s skill sets.

Powerful new twists

In fairness again, this is relatively new at the enterprise level – meaning good enough for serious business. As IT consultancy Morse Group’s principal consultant Andy Holpin says: “Virtualisation has been around in storage for longer than in the server market. Think about storage management products like Veritas Volume Manager that present a number of server DAS [direct attached storage] disks as one storage device. And think about virtualisation on storage arrays like SANs and NAS boxes. But this is at the storage fabric [the switches and host bus adaptors] level – and that’s what makes it suitable now for use in business-critical data centres.”

A couple of names coming slowly to prominence here with software-based virtualisation products that demonstrate quite different approaches are StorAge (recently acquired by disk giant LSI) and DataCore. StorAge provides what’s termed ‘split path’ or ‘out-band’ virtualisation, meaning that it creates a virtual storage volume but read/writes don’t pass through the virtualisation engine. Meanwhile, DataCore’s is an ‘in-band’ approach with the software running on a commodity server connected into the SAN such that all I/O goes through the virtualisation server – much as IBM’s own SAN Volume Controller (SVC) launched in 2003 but again, hardware independent.

Opinion as to which is best waxes and wanes around single points of failure and outright performance, but there’s more similarity than dissimilarity. The real point is best expressed by Ziya Aral, formidable chairman and CTO of DataCore, who made his name in the Unix and mainframe worlds back in the ‘80s and ‘90s. To those that say the technology for total storage virtualisation isn’t ready, he says: “Bullshit... It’s been a reality for a long time, and it’s so, so simple.” And to those who question its use in secondary storage: “If people tell you that with virtualisation you can replace back-up with CDP [continuous data protection], that’s absolutely true: virtualisation enables CDP.”

How so? Because his company has created what you can think of as a couple of multi-vendor, multi-type disk array controllers, each with a GUI (graphical user interface) front end: SANsymphony for big and relatively complex SANs, and SANmelody for SMEs with more modest requirements, and each in the familiar SCSI, Ethernet environment. What’s special about them is



▶ that they'll talk to just about any make and type of storage device – fibre channel, SCSI, SAS (serial attached SCSI), SATA (serial ATA) even internal IDE disk – on any operating system (Novell, SUSE, Red Hat, HPUX...), yet the platform is the ubiquitous Microsoft Windows.

Says Nick Broadbent, DataCore's European operations director: "Think of any disk type out there; if it can work with Windows it can work with us." And then the GUI: "You can create tiered storage pools just by dragging and dropping storage of your choice into whichever pool. For your Tier One, for example – your business critical storage, Oracle database, whatever – you can drag and drop your fibre channel fast disk into it, maybe RAID Five, or RAID Five plus one if you want to mirror as well. And the same applies to the other tiers."

Broadbent sees DataCore as struggling to get out from the shadows of the storage giants that largely control popular opinion. After all, no hardware vendor is going to be chuffed about some upstart neutralising proprietary value by abstracting all that into a software layer that leaves customers free to buy low cost disk for mirroring instead of more of the same expensive stuff.

However, he believes software-based virtualisation is about to have its day as IT managers are forced to think more carefully about ILM (information lifecycle management) and seek new solutions to save money and enable their infrastructures and people to cope with spiralling performance requirements, increasing data volumes and archive times that threaten to exceed the number of hours in a day.

"Costs and ROI are extremely important, but IT managers have felt they had to use high availability, high end, expensive kit for Tier One storage," he explains. "We're saying you can get availability and ease of use

without spending that money. It can still be high end, or we can provide high availability on fast and slow disks in a virtualised array."

And on secondary storage for test, development, data replication, back-up and archiving, he adds: "For many companies the back-up window is disappearing and impacting on production. So we say, take that off the LAN, put it onto virtualised SAN and you can do your backups without interrupting production – to tape, fast or slow disk, SATA disk if you want to switch to snapshots for back-up, whatever. And it's safer than RAID arrays because it needn't all be in one place."

Proof of concept

DataCore was instrumental in the Telegraph's move from Canary Wharf to Victoria, involving shifting an HP SAN, switches and the rest while enabling the group to get better performance and flexibility. "We implemented SANSymphony and they only had half an hour of downtime to reconfigure their storage to be in the path and rediscover the disks, then migrate and mirror the data across to the new site. Now they can choose any storage, and get faster performance because of caching... And the cost of that was two SANSymphony nodes at about £30–40k on standard HP servers."

So where does that leave you? All the consultants I speak to advise reviewing the business case, and if you're looking for the kinds of operational improvements and savings discussed, and especially if server virtualisation is part of that, then think seriously about virtualising your storage. Best advice is to go for a business-based pilot, if only to mitigate apparent risk.

As Computer Associates storage primo Jason Phippen says: "No-one is suggesting virtualising your whole environment as a starting point. It's still an emerging concept and there's a perception of risk – maybe if you lose the virtualisation engine you could lose access to the whole environment.

"But if you have a heterogeneous environment and you're using volume managers, that's expensive and difficult, so virtualisation is a powerful solution. Similarly, if you already have expensive Tier One storage – like IBM Shark ESS, Hitachi Data Systems Lightning, or EMC Symmetrix or Clariion – and you're looking for disaster recovery on a second data centre but can't justify the cost of the same again, virtualisation will let you use Second Tier LSI or HP EVA technology for example."

And he adds: "The cost is now quite low: you're talking about an entry point for the first few Tb of around \$10,000, and the ROI can be pretty quick. In fact, once customers bed in a storage virtualisation solution and really see the benefits and stability, they tend to rapidly expand: it almost become viral."

Finally, ponder this: why did LSI – the great hardware builder that badges disk for giants like IBM, Sun and SGI – buy a storage virtualisation developer? Could it be that the company knows it's only a matter of time before the market shifts wholesale this way? ■

YOU CAN'T HAVE ONE WITHOUT THE OTHER

If you're spending money on virtualising your servers to grab the benefits of consolidation and flexibility you're going to have a tough time if you don't centralise your storage. So it makes a lot of sense to consider virtualising that too, ideally as a single project, to reap the full rewards.

Few will go the whole hog because of perceived risk but as Ziya Aral, chairman and CTO of Datacore, says: "Any data centre is about one third network, one third servers and one third storage, so if you only virtualise your servers that doesn't help you deal with the long tail of physical storage detail – reconfiguring the wires, proxying, mapping and the rest. Storage is very physical so you've got to point to real disks, and with different systems, the tools are different, the methods are different and the procedures are different."

He concedes it's not always essential. Some application servers are very CPU-intensive so they don't use a lot of storage. "You may have a limited data set covering, say, one to 10 disks and they're DAS. The application may even be so minimalist you can use a NAS box," says Aral. "But if you have an Oracle database or an ERP system with a bunch of disks, you just have to virtualise the storage."

Incidentally, Aral insists this could have been possible years ago, but for the traditional storage hardware vendors fearing losing market. "The limitations of a SAN are the disk controller at the back end. That's it. IBM's SVC [SAN Volume Controller], for example, runs on AIX and that operating system will run with any [storage] hardware if they emphasise the communications to the back end. So you could do true storage virtualisation right there – but IBM doesn't sell it that way."

Enter 340 at www.mcsolutions.co.uk/enquiry