

Redirection to the rescue – I/O tools boost virtualization projects

By: Beth Pariseau,

Publication: SearchServerVirtualization:

Date: May 31, 2011

Data storage and backup are often the most expensive and complex elements of a server virtualization project, but some IT pros say the bleeding edge is a good place to look for these tools.

Case in point: A new generation of companies making their way into virtualization that use I/O redirection and optimization technology to solve vexing performance and backup problems.

IO Turbine, for example, uses I/O redirection to help IT managers take advantage of super-fast solid-state drives (SSDs) and improve virtual machine (VM) density. Actifio Inc.'s I/O technologies, meanwhile, speed up backup and replication times, helping to make users more comfortable virtualizing key workloads.

SSD caching: “We don’t use arrays and we never will”

When it comes to pushing virtualization into the world of enterprise Tier 1 apps, some users find that moving swap files out of traditional disk arrays using host-based caching software and SSDs can improve virtualization performance and density.

NextCloud LLC, a hosted private cloud service provider, uses beta software from [IO Turbine](#) called Accelio, plus Flash cache made of [SanDisk](#) (formerly Pliant) SSDs, to eliminate the need for swap files that live on external disk. The software automatically redirects I/O for “hot” workloads in VMware virtual environments to a Flash cache attached directly to the host, rather than allowing memory to page to external disk, which NextCloud said was more effective than traditional disk arrays for boosting the I/O operations per second (IOPS) performance of virtual servers. “We don’t use arrays and we never will,” said Gary Lamb, NextCloud co-founder and CTO.

Even the fastest spinning disks attached to a host via a network are no match for SSDs directly attached to the host bus. Internal testing has shown each 150 GB Pliant drive can produce up to 250,000 read IOPS and up to 125,000 sustained write IOPS. Lamb estimated the company would have to buy about 150 15,000 RPM Fibre Channel or Serial Attached SCSI drives for similar IOPS performance on performance-intensive data such as swap files.

NextCloud configures each of its blade servers with 1 GB of RAM and 2 GB of SSD cache dedicated to each VM. But because IO Turbine is not yet out of beta, Accelio isn’t in production in most of NextCloud’s client-facing environments.

Still, Lamb said the company plans to put the product into production as part of its second-generation infrastructure design, and one client that has been willing to work with NextCloud on testing it has been happy with the results. The client, who was running a smart building monitoring application drawing from 144,000 sensors around a hospital campus, was able to eliminate about 10 GB of RAM from the host and still achieve satisfactory performance. Ultimately, Lamb estimates that NextCloud will be able to double its VM density from about 30 VMs per host to about 60 or 70 performance-intensive workloads per physical server.

Getting users on board with new virtual infrastructure

Meanwhile, improving backup performance can help IT managers convince end users to accept virtualized workloads, while also simplifying operations.

Paradise Tomato Kitchens Inc., a sauce manufacturer for chain restaurants based in Louisville, K.Y., used a combination of tools from Actifio and Kroll Ontrack Software to replace its virtual and physical backup and replication tools. In the process, it freed up space on production storage that had previously been used to store backup data, which in turn persuaded end users at Paradise to embrace a new Intranet environment based on virtualized SharePoint, according to Paradise IT director Ryan Swain.

Users generally “don’t want change. So the whole scenario of the shared storage, [disaster recovery] and quick recovery, is added incentive for those guys to move off of their legacy systems and embrace the new one. When they see the performance gains that’s what really brings them over,” Swain said.

Actifio’s Virtual Data Pipeline (VDP) appliances sit between servers and data storage systems, and manage data copies, which can otherwise sprawl in virtual environments. Once data reaches the storage array managed by the Actifio appliance, it’s deduplicated for capacity efficiency and can be migrated or replicated between any two storage systems, using built-in WAN optimization.

Swain said the decision to go with Actifio to manage backup copies first arose from frustration with his environment’s old backup products, particularly Symantec Backup Exec 10d. The Actifio appliance handles data distribution through virtual copies, rather than requiring admins to switch between different consoles depending on whether data has been backed up to local disk, a secondary disaster recovery site, or is to be cloned for testing or development purposes.

Because Actifio’s VDP presents virtual copies, restoring a multi-gigabyte virtual machine takes, in Swain’s estimation, about 20 seconds.

“That’s just unheard of in the tape and software backup world,” Swain said. Also, “the other beauty of it...is that because [Actifio] interface[s] with everything at the [storage area network] level, it doesn’t impact the server or the network.”

Finally, some 24 TB of space in Swain’s environment has been freed up where backup data had been stored on a Hewlett-Packard Co. EVA disk array, in favor of a cheaper commodity disk array fronted by the Actifio appliance. Swain estimated the commodity array came at half the cost of the same amount of EVA capacity.

Actifio’s product can restore a virtual server image quickly, but one thing it doesn’t do yet is allow the restoration of a single object such as a file or mailbox directly from a backup. This is where Swain brought in a tool called [PowerControls](#) from Kroll Ontrack. PowerControls allows the restoration of single items to Exchange and SharePoint servers without requiring separate recovery space.