PAPER

Redmond MAGAZINE

UNLOCK HIDDEN RESOURCES IN YOUR HYBRID VIRTUALIZATION INFRASTRUCTURE

As the virtual hybrid infrastructure has evolved, technologies and tactics for backup and recovery have evolved as well.

Quest™

https://www.quest.com

As fundamental IT practices, backup and recovery have not changed in their importance. It remains critical for enterprises to protect their business-critical corporate data stores. The backup and recovery functions have, however, changed significantly with the advent of virtualization and cloud-based environments. Over the past several years, storage infrastructure has evolved from physical servers to virtualized servers, and has now evolved further to accommodate new platforms like containers and hybrid cloud. This has made the entire process of backup and recovery more complex and multi-nuanced.

The very complexion of infrastructure itself has changed with the advent of virtualization and cloud platforms. The relative ease with which organizations can spin up new virtual servers and individual virtual machines throughout their environment has led to some degree of sprawl in most organizations. Scratch the surface at any large organization, and there will invariably be virtual machines and servers that may be unused or simply forgotten about. This can create an unnecessary drain on resources.

This situation immediately leads to challenges simply knowing about your organization's resources. "The biggest challenge is actually having the ability to see everything; having the ability and the tools to help you make sense of what you're seeing and understanding the impact of new techs getting introduced in to your environment," says Chris Jones, product manager for Quest. Jones was speaking during a recent webcast entitled "Unlock Hidden Resources in Your Hybrid Virtualization Infrastructure."

The continued virtual sprawl and even the accompanying physical sprawl can make mapping out an efficient and effective backup and recovery program a challenge. "Environments are now spanning multiple locations and multiple technologies," says Jones.

The process of managing that complexity and ensuring true data protection has had to evolve as the nature of infrastructure and storage itself evolves. "We've seen physical servers being centralized, then decentralized, then centralized again. And now we're getting to a point where we're happy with virtualization," says Adrian Moir, senior consultant and product manager at Quest. "It really changed how we use our infrastructure."

Working in an increasingly virtualized and hybrid cloud-based infrastructure, it becomes critical to monitor and manage

Mathias Rosenthal / Shutterstock.com

> "WE DON'T TALK ABOUT BACKUP ANYMORE. WE TALK ABOUT RECOVERY NOW. WE TALK ABOUT RECOVERY TIME OBJECTIVES AND RECOVERY POINT OBJECTIVES." —ADRIAN MOIR, QUEST

sprawl and complexity to better control costs and ensure sensible investments in technology resources and services. "With the hybrid cloud, you have another change in the way you use our infrastructure. It's more abstracted and when you get to that level of abstraction, you can quickly lose sight of what's going on," says Moir. "When you're delivering on the front end, how much are you noticing what's going on in the back end? When you're in a cloud environment, that's really important because that's where the money is spent."

And while increased use of virtualization and cloud platforms can enable sprawl and add complexity, they can also be leveraged for new and evolving backup and disaster recovery protocols. "One method that's gaining popularity for the simple reason that it's not on your premises is to replicate your data to some bucket with a cloud provider somewhere else," says Moir. That is an effective approach for two reasons. "It's a second copy of your data and it's stored somewhere else. It's a good practice to put it somewhere outside of your environment."

RECOVERY IS CRITICAL

As the infrastructure and storage platforms have evolved, so too has the focus. The general focus is shifting simply from backup to the more practical aspect of recovery. "The important stuff is somewhere else so you can recover. That ability to recover is the bit everyone is interested in," says Moir. "We don't talk about backup anymore. We talk about recovery now. We talk about recovery time objectives and recovery point objectives and all the backup technology is driving how you want to recover."

And again, relying on a hybrid cloud platform can help expedite recovery, but you have to carefully manage and configure the process. "You can leverage the cloud in a couple of ways," says Moir. "The old case is to just have your data sent out there. But if you put your data out there and you want to recover it, how you going to get that data back?"

It's important to consider the size of your data stores and factor in the time it will take to move data back and forth—not only for backup and recovery processes, but also for business activities. Time spent transferring data is time your business is not spending earning revenue.

Applying some level of intelligence and capability to your cloud platforms helps ensure continued access. "It's far more obvious to put a level of capability in that cloud environment in that offsite

"YOU HAVE DIVERGENT TECHNOLOGIES LIKE ONEDRIVE AND SHAREPOINT ONLINE. THEY ARE NOT THE SAME AS TRADITIONAL FILE SHARE, BUT THEY'RE STILL OPEN TO THE SAME TYPE OF RISKS." —CHRIS JONES, QUEST

location to be able to use that data," says Moir. "So you have to have a bit of intelligence in that secondary site in case the first site is not available."

One way to further expedite the recovery process is to place virtual machines in the cloud. "We can put an entire machine set up in a cloud provider and it sits there. It's turned off so it's not consuming as much and not sitting there costing money," says Moir. "That's like the old co-located model, but you have this in the cloud so you don't have to support the infrastructure, but you have these machines ready for disaster recovery and you have your data sets ready."

One potential challenge to this approach is having to deal with an increasing number of data formats. "It is becoming more common to have data mixed," says Jones. "You might say that's not a hybrid environment, but it certainly is. You have divergent technologies like OneDrive and SharePoint Online. They are not the same as traditional file share, but they're still open to the same type of risks. If a user accidentally deletes something, you do need that level of recoverability across all those technologies."

This blended or varied format is an increasingly common problem in increasingly virtual environments. "When people get into virtualization, the first things they do is go for the easy stuff. They virtualize Active Directory, the print server, and then the file server," says Moir.

Once the file server is virtualized, you have to carefully manage growth in order to preserve your backup capability. "The virtual file server can just grow and grow to the point until snapshot functionality in hypervisor can't deal with it any more. It's just too big. You've got millions of files in there. Then you have to find a different way to back it up," he says.

That typically involves deploying an agent within a virtual machine to do that. That also means at this point you have to think of alternative methods of actually moving those files. "You still have to do it as part of your backup regime," says Moir. "You just have to think of a more intelligent way to do that."

APPLY THE RIGHT SOLUTION

Jones and Moir have been seeing their customers use Quest QoreStor to resolve these challenges. "It does Source ID deduplication and runs anywhere you want to run it," says Jones. "As a virtual machine in remote location, it can replicate data to a primary data center. More importantly, when you're running QoreStor in the cloud or hosted by an

CONTAINERS PROVIDE ANOTHER DEGREE OF TECHNOLOGICAL AGILITY, BUT ALSO POSE SOME UNIQUE IMPLICATIONS FOR DISASTER RECOVERY, BACKUP, AND RESTORE.

MSP (Managed Service Provider) as a service, that's enabling the possibility of Source ID to protect data and send backups directly offsite. Or if you want to backup locally then get that second copy and replicate to the cloud, you can also do that."

QoreStor is a software-defined secondary storage technology. You can install it directly on physical server hardware, in a virtual machine, or a public or private cloud environment. That level of flexibility is more important than ever considering the changing nature of virtualized and cloud-based platforms. And Quest has recently released QoreStor for the Azure marketplace. "You can deploy QoreStor and immediately start getting a copy of data backup or even a primary copy by doing a direct to Azure backup," says Jones.

And almost as importantly, it's available on a subscription basis. "We built cloud-optimized configurations into QoreStor," says Jones. "Our goal was to keep it on a stack at below $200 per month. So, you can protect your remote offices with no CapEx; just pure OpEx. You can support that data protection with a subscription."

That can be an important distinction for most organizations, as it engenders a degree of flexibility that is essential in this business environment. "Everyone from the CTO down has to understand what focus is going to be and what is required for the business to grow," say Moir. "If you can't understand that, then CapEx will be like an anchor. You need to be more agile, so OpEx cloud-based situations are more appropriate. You're spending money on what you're consuming, and can adjust depending on where the business is going and what's it's doing."

MAINTAINING YOUR CONTAINERS

While that organizational and fiscal agility is important, so too is maintaining a similar degree of technological agility. Maintaining an IT infrastructure, complete with the full backup and disaster recovery capabilities, is certainly a different paradigm when looking at on-premises versus cloud platforms. And now virtual containers and containerization is another level of disruption to the composition of your infrastructure. Containers provide another degree of technological agility, but also pose some unique implications for disaster recovery, backup, and restore.

"The backup piece is interesting for containers. Do you back up your containers? What in the container do you actually want to back up?" says Moir.

"Containers provide a service, which is then torn down. During that peak demand, it's an ideal solution. But the most important thing is the database. That's the content you want to protect. You can protect the container images when they're running, but that's it. They're all the same. We used to get asked that a lot with VDI. 'How do you back up desktops?' You should back up the data people are using—not the desktops themselves. We're getting out of the mentality of virtual machines versus desktops."

Container technology has certainly achieved a level of mainstream adoption. "Container evolution has always been interesting. One of the challenges businesses have with traditional apps was understanding how they're performing. Containers complicate that," says Jones.

"Infrastructure teams need to understand where the data is going. 'Are you on-premises or purely in the cloud? And what about scale? Where might that data be persisted? What is the likelihood or rate of change for that data? How do I effectively recover from a full outage?'" says Jones. "You need more consistent storage solutions in containerized environments. Just protecting the container isn't really bringing full value to the business."

Achieving full visibility and ultimately being able to review and mange that via a single pane of glass becomes that much more critical and challenging in this varied and virtual environment. Quest's Foglight has evolved to help meet some of those challenges. "Foglight started as a traditional monitoring tool, and it has evolved over the years. It has become more adept at specializing in databases. The hybrid datacenter has been the challenge, so we're making sure we're supporting them and their environment as it's changing," says Moir. "The next piece we're focusing on is collecting information from the hybrid environment; connecting and pulling down performance information. You can get into a hybrid solution and continue to manage that with Foglight."

Foglight helps organizations gain comprehensive visibility across their hybrid infrastructure. Infrastructure teams can then leverage real-time data, historical data, and full reporting capabilities to inform their decision making and better forecast future capacity needs. Foglight brings full automation, optimization, monitoring and forecasting to all components of a hybrid environment; including physical servers and VMware virtualized servers, as well as Hyper-V, Azure, and AWS platforms.

Faced with an evolving virtual and hybrid infrastructure, organizations need to pay attention to their backup and recovery technologies and tactics. These days, more than ever before, backup and recovery has to be agile, flexible, and capable.

For more information:
https://www.quest.com/foglight/

Quest™