COMMVAULT®

Protect, Monitor and Manage Your Virtualized Data

5 WAYS TO IMPROVE DATA MANAGEMENT ACROSS PHYSICAL, VIRTUAL AND CLOUD ENVIRONMENTS

Rapid growth in virtual infrastructure is introducing complexity which is slowing down IT's ability to deliver value to the business. Left unchecked, it can have a significant impact on IT cost, efficiency, manageability, and productivity. Compounding the challenge are the vast differences that virtual environments can have from each other. These differences can prevent the deployment of different virtualization technologies that have the potential to reduce costs. From their size and architecture, to the workloads they support and their location, each virtual environment requires a unique approach to protect – and recover – the data it contains. Since the attraction of cloud-based (virtualization) platforms is too strong to ignore, most environments will inevitably wind-up with their data spread across multiple virtualization platforms.

To efficiently protect, monitor and manage your virtualized data, regardless of where it may be located, consider these five best practices. They'll help you minimize the complexity, tame virtual machine (VM) sprawl, and work to ensure faster data recovery when you need it.

Infrastructure complexity is the new way of life for most IT organizations, especially as more workloads become virtualized, are moved to the cloud, and unrelenting data growth continues to proliferate. The result is the compounding challenge of complexity, which can increase risk, drive up management and resource costs and slow the adoption of digital transformation strategies.

Gartner reports that through 2019, every dollar that enterprises invest in innovation will require an additional $7 in core execution.[1] Ouch! Of course, the cost of implementing innovation can easily outweigh the benefits of innovation itself. And infrastructure complexity is a large contributor to this cost. So how can you accommodate both: drive innovation, yet enable consistent efficiency? Look beyond the virtualization container and, specifically, at how you protect, monitor and manage your DATA.

DATA MANAGEMENT CHALLENGES IN MIXED ENVIRONMENTS

Today's enterprises have seen an explosion in data types, applications and workloads. More are being deployed on different virtual platforms. They are also leveraging different databases and are deployed in mixed environments - on-premises, in the cloud, or in a hybrid infrastructure. The result is infrastructure cross-virtualization: the creation of virtual servers, storage, operating systems, applications, and other network resources that have been deployed across a mix of local, cloud and hybrid environments. While this provides a powerful way to accommodate massive data growth, agility, and possibly even reduce costs, it also builds layers of complexity.

Among the challenges are:

- **Every application, or workload, has its own unique requirements.** Some of them vastly different from the next. This means that having one or two universal policies will not suffice, requiring customization and therefore adding overhead. The variety of size, location, complexity, and availability requirements change how we must protect – and recover – each of them.

- **VM and data copy sprawl.** A serious issue in its own right, sprawl proliferates when there isn't an efficient process to track VMs across their lifecycle, wasting both storage and resources. Some might say that there are two types of VMs – those you know about and those you don't. It's those hidden VMs that are risky business, as they often contain data that's ungoverned and unprotected. This means they are susceptible to inadvertent duplication, accidental loss, and even exposure to cybersecurity threats.

> "Through 2019, every $1 enterprises invest in innovation will require an additional $7 in core execution."
>
> **- GARTNER**
> *Predicts 2017: IT Services Market Opportunities Expand in the Digital Era, October 20, 2016*

[1] Gartner, Predicts 2017: IT Services Market Opportunities Expand in the Digital Era, October 20, 2016

- **The challenge of rapid operational recovery.** This is particularly important, given that over one-third of servers (physical or virtual) have a downtime tolerance of fifteen minutes or less, and another third of servers (physical or virtual) have a downtime tolerance of less than two hours.[2]

- **Meeting recovery SLAs.** With today's compliance, governance, and security strategies front and center, both disaster recovery (DR) and operational recovery are critical. Without processes in place to enable the fast, easy recovery of virtual data wherever it resides, organizations are unable to achieve their SLAs. And the ability to deliver on-demand access to data without having to wait for a full restore operation to complete is nearly impossible.

5 WAYS TO IMPROVE THE MANAGEMENT OF VIRTUALIZED DATA

To meet these challenges, you must employ a data management strategy and model that enables you to capture, monitor, protect and leverage your data, regardless of how or where you are retaining it. Consider these five requirements. They will ensure that your valued information is not only protected, but compliantly managed and easy to recover, should the need arise.

1. **Natively Support Multiple Virtual Architectures and Hypervisors –** As you evolve and modernize your data protection strategy, ensure that the solutions you rely on support not only the virtual architectures and hypervisors you use today, but any that may be relevant for your environment in the future. There is nothing more limiting than vendor lock-in. Stand-alone virtual data protection products don't easily scale and will limit your protection capabilities to only one virtual environment. You need to know that your data can be easily moved – on-premises to the cloud, cloud-to-cloud, or cloud back to on-premises – in order to support the changing demands of the business. Select solutions with native and heterogeneous support for a complete range of hypervisors, storage, cloud, and hybrid architectures to ensure the mobility, scalability and agility of your business needs, both now and in the future. Look for solutions that not only incorporate the use of native hypervisor APIs, but can also enhance those native capabilities.

2. **Seamlessly Extend Your Virtual Environment to the Cloud–** The cloud is quickly changing the economics of business, with lower cost data retention options and elastic scalability. This move will only increase over time. Gartner reports that "by 2020, more than $200 billion in annual IT spending will shift to cloud and cloud-related categories of spending."[3] To take advantage of this new dynamic you need to seamlessly incorporate the cloud into your on-premises infrastructure

[2] Jason Buffington, New Agility Requirements in Data Protection, Enterprise Strategy Group, November 2015
[3] Gartner, Predicts 2017: IT Services Market Opportunities Expand in the Digital Era, October 20, 2016

as a natural extension of your datacenter. There are many current products that can move data to cloud, but do they align with your desired cloud strategy, and do they provide the same capabilities in-the-cloud as they do on-premises in a cloud-cost-optimized manner? To support this, you will want to implement data protection and management solutions that can manage your virtual environment – regardless of where it's located – in a single, unified management console that treats cloud simply as an extension of your virtualized environment. This will streamline your management processes while giving you a holistic view into your entire environment, enabling you to optimize resources, unify policy controls, and ensure SLAs, regardless of where your data is stored.

3 **Enable Global Visibility and Control Under One Umbrella –** To make the most of your virtualized environment, and to avoid the risk and cost of hidden VMs, you need to employ solutions that will improve visibility across the complete lifecycle of your virtualized environment and data. Knowing what exists out there, and where it's located, is the crucial first step in allowing you to act on it. Global visibility into the state of VM usage, as well as control over VM creation and disposition for unused VMs, can also help to mitigate unnecessary risk, improving your compliance posture with active policy control. Solutions for global visibility will also help you control data across your enterprise – from creation to retirement – for improved governance, discovery, and resource planning. Take it even a step further by curbing your VM sprawl. Automating the provisioning of virtual machines and automatically attaching backup, recovery and archiving capabilities immediately as they're created helps to be sure that they're automatically included in your protection strategy, and ensures your data is always recoverable.

4 **Use a Single Index to Eliminate Redundant Data Silos –** With infrastructure complexity and data growth comes the challenge of siloed applications and data. Adding-in new architectures and cloud instances only increases the problem, as redundant data silos proliferate. Stand-alone backup and recovery products are particularly bad when it comes to populating siloed data. Their scale limitation often creates multiple data and management silos. By employing data management solutions that integrate hardware protection capabilities and offer a single data index that spans all your enterprise data, it's easier and more efficient to locate specific data throughout your storage tiers. More importantly, a solution that uses a single virtual repository and deduplication tools will help eliminate redundant data silos while lowering costs and simplifying administration.

The CIO Guide for Virtual Server Protection

Discover the four questions every CIO should ask about virtual server data protection.

READ NOW

PDF

http://bit.ly/1J7DiJ5

5 **Enable Fast Disaster Recovery and Data Access –** Ensuring rapid DR and data access improves the management of your virtualized data. This can be especially challenging within mixed environments. Because both recovery time and the recovery point are critical, you need to implement a comprehensive approach that will span your critical data and workloads. Look for DR options that offer built-in automation, orchestration, replication, alerting, and reporting. More importantly, look for solutions that will allow you the flexibility of deploying multiple and delayed recovery point objectives to protect against new threats such as ransomware. This helps ensure that you can optimize your recoverability for the high degree of protection your workloads demand. To further drive productivity, look for a solution with a self-service interface, to enable greater recovery productivity. This allows you to extend controlled recovery across your organization, lightening the load on your IT recovery management teams.

▶ READY YOUR MIXED ENVIRONMENT FOR THE FUTURE

From here on out, the mixed or hybrid environment is going to be a given, especially as enterprises extend their datacenters to the cloud. But that doesn't mean you need to forgo your data protection and management goals, nor do you have to compromise on how or where you deploy your applications to ensure recovery readiness. Look to solutions that offer native support for a wide range of virtual servers, storage, operating systems, applications and other network resources, no matter where they're deployed. It's *your* data, and regardless of which hypervisor or physical location you choose, ensure that your protection and recovery SLAs move along with your data and don't pose limitations on your platform choices. With the right solution you'll have the confidence you need that your data protection will adapt easily and seamlessly to the changing dynamics of your business.

▶ Protect, monitor and manage your virtualized data. Visit **commvault.com/virtualization**.