

Advancing Your Virtualisation Journey

Brought to you compliments of

ARCserve®
More than Backup

Table of Contents

ADVANCING YOUR VIRTUALISATION JOURNEY.....	2
MAKING VIRTUALISATION WORK FOR YOU	2
DATA PROTECTION IN A VIRTUALISATION ENVIRONMENT	2
BUSINESS CONTINUITY AND DISASTER RECOVERY IN A VIRTUAL ENVIRONMENT	3
GROWING THE VIRTUAL ENVIRONMENT.....	4
CONCLUSION	5



Copyright ©2010 CA TECHNOLOGIES. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies. This document is for your informational purposes only. CA TECHNOLOGIES assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA TECHNOLOGIES provides this document “as is” without warranty of any kind, including, without limitation, any implied warranties of merchantability, fitness for a particular purpose, or noninfringement. In no event will CA TECHNOLOGIES be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, business interruption, goodwill, or lost data, even if CA TECHNOLOGIES is expressly advised in advance of the possibility of such damages.

ADVANCING YOUR VIRTUALISATION JOURNEY

Virtualisation enables companies of any size to achieve greater measures of business agility and cost effectiveness. It also offers a range of advantages for transforming an IT department from a source of expense to a tool for enabling greater corporate profit. But for many, the move to a virtualised environment is just the beginning. Although the new virtualised server environment opens many opportunities for savings and increased productivity, organisations soon discover that it also brings a whole new set of operational challenges.

For some, the complexity of navigating further into the virtualisation landscape seems daunting, but the journey can be worthwhile. This white paper offers ideas to help you continue the virtualisation journey and maximise your initial investment.

MAKING VIRTUALISATION WORK FOR YOU

In general, we think of server virtualisation as the process of moving business applications from a large number of physical servers to a large number of virtual servers running on comparatively few physical machines. In many datacenters, virtualisation is driven by server sprawl, which occurs as new applications are switched on and given their own servers to minimise the risk of under-resourcing, allowing them to operate in isolation from other applications and thus reducing the chance of an unplanned outage. Server virtualisation promises to reduce hardware needs by enabling applications to be isolated inside a virtual machine but share physical resources, maximising the use of assets.

Organisations turn to virtualisation to simplify infrastructure, reduce management costs and gain a greater level of flexibility. For many, these aims have been achieved and now IT departments are looking for other ways to exploit virtualisation to realise even greater benefits. For example, disaster recovery is one area of the business that can benefit from the flexibility of a virtualised platform, enabling highly available environments in which business-critical application can run. Such environments can be provided in a much more cost-effective manner with virtualisation than is possible in a completely physical datacenter.

For some organisations, however, the journey has just begun, and they are realising that to fully benefit from server virtualisation, they have to adapt their management practices and expectations. Virtual server sprawl can become a real problem — potentially worse than that in a physical datacenter. With the ease and speed of deployment of virtual servers, it is way too simple to create new platforms for test and development, and even production. IT managers in a virtual datacenter have to manage an ever-changing landscape.

One area that offers a whole set of fresh challenges in the virtual datacenter is data protection, and existing policies need to be reviewed and adapted to secure what for many is a combination of virtual and physical servers.

DATA PROTECTION IN A VIRTUAL ENVIRONMENT

For many years, organisations have been designing and executing data protection policies for their physical datacenters. They may have had to adapt them as applications became business-critical and data grew, but the goal has remained fundamentally the same: to perform sufficient backups to minimise data loss in the event of a failure, while not impacting the performance or availability of applications.

Data protection policies have evolved from backing up to tape to backing up to disk to speed both backup and recovery processes. They have also evolved from capturing live data, to creating a snapshot and using that as the source of the backup data, and from keeping a copy off site, to having multiple copies for local and remote recovery. So when an organisation virtualises the environment, why would the data protection process be any different?

Data protection in a virtual datacenter is designed with the same goal of a physical datacenter: to perform backups without impacting the performance and availability of the application, while ensuring minimal data loss and fast recovery. However, in delivering on these goals, the policy has to be designed to tackle some new challenges.

The ease of deploying virtual machines means that new test, development and production servers arrive frequently. These must be incorporated into the backup schedule, and provisions must be made for the extra resources needed to protect them. The more the policy can capture the arrival of these new machines automatically, the less chance there will be that they will be omitted from backups. If automatic discovery is not possible with the chosen backup solution, then manual procedures need to be defined to update the data protection policies.

Once created, a virtual server could move. During one pass of the backup, it could be running on one physical server, and on the next backup pass, it may have moved to an alternate physical server. The backup process should be able to track this movement and maintain an accurate history of the virtual machine regardless of its location.

What to back up is also an important factor. Is an image of the whole virtual server going to provide the correct recovery capabilities, or is a more granular or application-specific backup needed? This will dictate how the backup solution is deployed. Is it running at the hypervisor level and capturing an image of the whole virtual machine, or does it run within the virtual machine and communicate directly with the applications? The recovery point and recovery time objectives of an organisation should dictate this configuration — it may be that a combination of approaches is needed to deliver on the recovery service-level agreement.

When you back up a lot of virtual machines, the amount of duplicate data is considerable. If you have 10 Windows servers running inside virtual machines, you have 10 copies of the operating system to back up before even getting to the data. Such levels of duplication cost organisations money because more infrastructure is needed to store multiple copies of the multiple machines. Therefore, it is important to consider data deduplication technology as part of your data protection solution. This ensures considerable storage cost savings and reduced operation time.

Finally, it is unlikely that any organisation will be completely virtualised. In fact, many sources report that the level of virtualisation in most companies is still less than 50% of servers. Therefore, data protection solutions that provide a single interface and set of policies for both physical and virtual environments will deliver a more cost-effective solution with less risk to the organisation than purchasing, deploying and managing separate solutions for physical and virtual servers.

Companies such as CA Technologies offer comprehensive data protection solutions, especially when it comes to backing up data in virtualised environments. CA ARCserve® Backup reinforces data protection by enabling high-speed backups from Windows, Linux or Unix servers. It provides backup to disk, to tape, to virtual tape libraries and in disk-to-disk-to-tape backup schemes. From a single user interface, it provides automatic discovery of all virtual servers as well as the necessary management and reporting features to track and document virtual and physical environments. Certified across an extensive range of industry platforms and applications, CA ARCserve provides compatibility and support for heterogeneous environments, including VMware, Microsoft Hyper-V and Citrix XenServer.

Once you solidify your physical and virtual data protection strategy, how can you take it to the next level? Many organisations now realise that periodic backups alone do not meet their service-level agreements or recovery-point objectives. Data lost or damaged between periodic backups or snapshots has to be manually re-entered and, in some cases, may be lost forever. So what can you do? Data replication solutions like CA ARCserve® Replication complement any backup solution by delivering continuous data protection and data rewind capabilities to recover lost or damaged data between periodic backups — perfect for more critical data or demanding service-level agreements. Data can be replicated from one server to another on the LAN or to any remote location to achieve both continuous data protection and disaster recovery goals.

A data replication solution also enables secure, fast and easy copying of backups (even if deduplicated) over the WAN to any remote location as an alternative to risky, costly and time-consuming off-site physical media transport.

BUSINESS CONTINUITY AND DISASTER RECOVERY IN A VIRTUAL ENVIRONMENT

All organisations have one or more business-critical systems that they rely on for day-to-day operations. Interruption to one of these critical systems can have a big impact on revenue, customer service and support, employee productivity and customer loyalty.

New bare metal restore (BMR) technology, like that available in CA ARCserve® D2D, is available to help IT speed server recovery after an unexpected outage or failure. BMR helps eliminate the need to build a server from scratch, including the operating system, system state, application and data. But for some organisations, even recovery using BMR may take too long for critical communications and revenue-generating applications and data. So many organisations have invested in high-availability solutions. This historically has required companies to have not only clusters of servers and replicated disk, but also remote locations with redundant standby servers sitting idle, waiting for a disaster to occur.

In the past, only large enterprise organisations could afford high-availability solutions, as they typically required like hardware and storage configurations at both the production and failover sites. And when those failover servers are physical, they consume a relatively large footprint, which eats into operations budgets. Today, with host-based replication and high-availability software solutions, IT may pick and choose the server (physical and virtual) platform and storage devices (DAS, NAS and SAN) that make sense for the environment. Failover servers may reside locally on the LAN or be deployed at any remote site for both continuous availability and disaster recovery. Now, even small and midsize organisations can attain continuous data protection and high availability for critical applications and data.

Coupled with advances in network bandwidth, software replication, system monitoring and failover, server virtualisation enables organisations to deploy highly available business continuity and disaster recovery environments in a much more flexible and cost-effective manner.

Since new servers can be created and deployed quickly in a virtualised environment, sophisticated business continuity and disaster recovery configurations can be developed without a lot of the complex infrastructure needed in a physical environment. High-availability solutions that incorporate real-time replication and failover help solve the server downtime dilemma. Real-time replication enables the transfer of data to a failover server while the production server and applications remain online and available to users. If a server, application or network failure occurs, application workloads and end users are automatically redirected to the failover server with minimal to no interruption. When the failover servers are virtualised, cost is significantly reduced.

GROWING THE VIRTUAL ENVIRONMENT

The server virtualisation market is evolving quickly. In a relatively short period of time, the dominance of a single vendor has been eroded by others entering the market. What this often means for an organisation that has successfully started the migration from physical to virtual (P2V) is a re-evaluation of the hypervisor platform and a possible transition to a heterogeneous virtual environment.

When this occurs, new migration and data protection challenges surface. What was often viewed as a one-time P2V migration has changed into potentially a multistage virtual-to-virtual (V2V) movement of servers. This makes the choice of data protection, disaster recovery and migration tools even more important. Now there is a need for a single view of physical and heterogeneous virtual environments, and the ability to move data between different platforms and recover data to very dissimilar architectures.

Administrators and those charged with facilitating the move from physical to virtual servers also need to keep in mind that as the company grows, virtual-to-physical migrations are often necessary. For example, an application that has been virtualised, such as a SQL database, might need to be moved back to a physical environment due to the increase in database size or the need for more CPU cycles.

Migration can be facilitated using one of two methods: cold or hot. A cold migration is simply a migration that occurs while the virtual machine is in suspended mode or powered off. This type of capability is facilitated by a backup application such as CA ARCserve D2D. It can use its hardware-independent BMR capabilities to migrate whole servers, applications and data from a physical environment to a virtualised one, virtual to virtual and, of course, virtual to physical.

CA ARCserve D2D enables IT to apply incremental snapshots of data located on a physical server and to quickly restore that data to a virtualised one. Not only will the storage space, network traffic and load on production servers be reduced, you will be prepared for a range of possible outages, from application failures to server malfunctions or other unforeseen natural disasters. CA ARCserve D2D also alleviates the time-consuming nature of physical data restores due to new hardware acquisitions or the task of moving data to an entirely new data center.

In the case of hot migration, server consolidation occurs without taking the entire system offline, thus assuring high availability. One server can continue to operate while the destination server is being created and populated. The process can be performed transparently, with no discernible interruption to users or applications, which is crucial for 24/7 business operations and user availability.

CA ARCserve® High Availability enables this type of real-time system and data migration, solving both 24/7 availability issues and P2V migration requirements. Because uninterrupted workloads can be immediately transferred from the production server (physical or virtual) to another virtual server, users maintain continuous access to applications, even during P2V migration. Real-time replication, server and application monitoring and automated failover and failback help ensure that there is virtually no downtime for users or applications. And CA ARCserve High Availability can even ease V2V system, application and data migration for organisations looking to migrate from VMware to Hyper-V virtual servers.

CONCLUSION

Conclusion The variety of heterogeneous products available today can create hurdles in your virtualisation journey. This paper examines both the potential successes of a virtualised environment and the challenges. In order to maintain control of a virtualised data center, administrators need to approach the process from a holistic perspective. It may seem like a significant challenge to coordinate different approaches. However, the CA ARCserve® Family of Products provides a number of tools and approaches that will enable you to achieve comprehensive data migration and protection across your company's virtualised infrastructure.

To learn more about the ARCserve Family of Products visit www.arcserve.com/gb

ABOUT CA RECOVERY MANAGEMENT

The Recovery Management and Data Modeling Business Unit of CA Technologies (NASDAQ: CA) delivers the acclaimed CA ARCserve and CA ERwin products. Providing much more than backup, the CA ARCserve Family of Products gives customers control over their changing business by delivering total protection, recovery and availability for systems, applications and data — across physical and virtual environments. These award-winning products come together under the Business Unit's commitment to a 100% channel business model driven by more than 10,000 partners worldwide.