October 2013

Meeting Business Data Protection Needs in a Mixed Physical / Virtual Environment

arcserve®

Virtualization has without doubt enabled IT organizations to obtain maximum efficiency from their physical hardware, but virtualization has also increased the complexity these organizations have to manage, and that includes data protection. This paper discusses some of the hidden complexities virtualization introduces to backup and recovery requirements, and the challenges created when backup involves a combination of physical servers and virtual machines. It also discusses ways that backup and recovery processes can be simplified by avoiding point solutions.

The Myth of Total Virtualization

To read the technology press, it would seem that virtualization is technology's version of a done deal - more or less the only way computing gets done these days. This simply isn't true. It is true that organizations ranging in size from the U.S. government to small consulting firms with half a dozen employees are making heavy use of virtualization technology. In fact, many companies now have "virtualization first" policies, which means IT has to justify *not* running a new application on a virtual machine (VM). It is also true that adoption rates for virtual servers are high. According to one recent survey conducted by the independent Spiceworks IT community, 72 percent of small-to-medium sized businesses (SMBs) have already adopted server virtualization technology. Numerous vendor-sponsored surveys put that number somewhat lower, in the mid 40 percent range, but virtual server technology is clearly well beyond the early adopter stage.

What is also clear, however, is that the one-application-one-physical-server model is not going away. *eWeek* recently reported that a mere 4 percent of SMBs run all their critical applications in a virtual environment. Forty-four percent of the sample used a combination of physical and virtual servers for these applications, and 25 percent ran the majority of their critical applications exclusively on physical servers. Even VMware evangelist David Davis recently stated, "I don't think those [physical storage arrays associated with physical servers] are going away any time soon."

Virtualization has created numerous storage headaches. The most obvious has to do with the CPU capacity available during backup windows. With numerous VMs running on a single physical server, there is much less headroom to accommodate back-up operations, and careful scheduling is required to avoid degrading application performance and falling out of compliance with an SLA commitment.

In private clouds with load balancing among multiple physical servers, the situation is worse, because allocation of resources is often automatic, rendering careful backup schedules meaningless. Further complicating the situation is the fact that in private clouds individual users can create their own VMs via self-service portals, leading to the virtual equivalent of physical "server sprawl."

Finally, when multiple VMs are running on one physical server, there are complications because the different applications/VMs may have different back-up needs, e.g. they may have several different Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs), and IT teams may want different levels of granularity for restoration purposes. It's about recover-ability.

There are, of course, point solutions specifically designed to manage the backup of VMs and deliver the precise level of granularity required. But point solutions have their own disadvantages. These include:

- Acquisition and installation costs
- Extra personnel training
- Allocation of over-stretched resources for day-to-day administration
- Harmonizing two sets of storage data so that business and regulatory requirements can be met

In short, point solutions introduce even more complexity into an already complex situation.

A third option is backup to a public cloud. This option has the advantage of shifting costs from CAPEX to OPEX, but there are significant disadvantages as well. For many IT managers, the most significant is loss of control. There is a legitimate question of whether or not it is wise to entrust data that is truly mission critical to a third party. In some cases, there are compliance ramifications because some agencies require that companies have visibility into the physical location where sensitive data is being processed or stored, and in a public cloud, this may be literally impossible. Also, when data is physically located outside the U.S., legal disputes fall under the jurisdiction of the country where the data is physically stored. This can cause major problems.

There's no doubt that all these challenges can be overcome, but they highlight the fact that the public cloud strategy for storage, which is supposed to simplify operations, has a major component of hidden complexity.

The Business of IT Is Business

From a high-level perspective, the best way to deal with the new complexity introduced by virtualization is to look at backup and recovery first and foremost as *business* issues. Senior IT managers in SMBs and mid-sized organizations typically see their primary job responsibility as meeting SLAs and, more broadly speaking, "keeping the lights on." When the "lights" go off - when a system fails - they are likely to hear about it via angry phone calls within minutes. In a recovery situation there are two questions in the minds of users and senior management alike:

- How soon will we get back online?
- How much data have we lost?

In other words, they are concerned with RTO and RPO, even if they haven't heard these terms or aren't aware of their practical implications. Given the new complexity introduced by virtualization, IT managers face a daunting challenge dealing with these two issues while meeting the demands of senior management. Ideally, the RTO for any given application should be as short as possible and the RPO should be as close to the present moment as possible. In reality, achieving this ideal is more important for some applications than others. For an e-commerce site, both RTO and RPO are extremely important. For an HR department, there is less urgency (with the important exception of payroll). The key point here is that decisions about how to protect the data in various applications are *business* decisions. The factors to consider are somewhat different for the two recovery objectives.

RTO: What's at Stake?

Generally speaking, applications can be deemed mission critical if, when they fail:

- The company can't sell.
- Employees can't work.
- Processes, such as manufacturing, are shut down.
- People can't communicate.
- Customers receive poor service, leading to customer dissatisfaction.

Beyond these basic criteria, deeper analysis is possible, at least for the first three points. In fact, the criticality of applications can often be quantified, putting decision-making on a firm basis. If a company's revenue comes from online sales and its e-commerce web site is down, the cost of the lost sales per minute or hour can be calculated based on historical data. If design

engineers can't work because a CAD system is down, the cost is the value of their lost time: the number of individuals who are "out of work" times their hourly cost to the company (wages, benefits and amortized overhead). If an assembly line is shut down because the ERP system that coordinates operations is down, the cost is the lost labor (as with the CAD example) plus the amortized cost of each machine on the line. The same approach for valuing lost time can be applied to "non-physical" processes like claims adjudication in an insurance company.

The cost of lost communication capabilities is harder to quantify, although there is a general consensus that email is mission critical to any business. In some companies collaborative applications like Microsoft SharePoint may play a similar mission critical role. Also, the inability to communicate externally, particularly with customers, can tarnish a company's reputation. Any time the customer base or public knows about a failure, there is a chance of a steep reputational cost.

RPO: What's at Risk

As with RTO, RPO decisions should be driven first and foremost by business factors, and secondarily by technical factors. How often the data in a database changes is important in determining RPO, but not as important as the business value of that data (or the costs and risks associated with losing it). Here are some business factors to consider:

- Value of lost intellectual property. Whether it's an engineering team designing a valve or a financial consultant working on a spreadsheet, if three hours of their work is lost, that work has a value that can be calculated.
- Value of lost goods. If industrial quality control monitors collect inline data to ensure that a product meets certain specifications and that data is subsequently lost due to a system crash, the batches involved often have to be destroyed or, at minimum, retested. Again, those costs can be calculated.
- Cost of re-creation. According to one estimate, the cost of retyping 20 megabytes of sales data is $17,000, while the same amount of accounting data costs $19,000. In a paperless office, such data could simply be lost forever - a true catastrophe from which some businesses couldn't recover.
- Fines. In some industries, the failure to produce data for regulatory agencies can result in fines.
- Accounting inaccuracies affecting business reporting/decisions.
- Inventory inaccuracies affecting sales.

The main point is that the approach to recovery should be based on business value, and in many cases that value can be quantified.

Technical Decisions

Once business value has been determined, the decisions about appropriate backup and recovery options can be made. Depending on the criticality of the application, there are four basic and often complementary choices:

- Traditional tape backup
- Image-based backup to disk, which can dramatically reduce recovery times compared to tape
- Continuous data replication for demanding RPO
- Solutions that monitor systems and applications and provides automatic failover to another server or VM to prevent an unplanned outage for high availability

The technical details of how these solutions can be implemented lie beyond the scope of this paper. The point, however, is that recovery solutions can and should be matched to various applications. One size does not fit all. Further, it makes sense to deploy a single, integrated solution that can provide all of these options for both physical and virtual environments, rather than rely on a collection of point solutions that can't deliver an overall picture of what's happening.

Measure and Test

Obviously, the most carefully thought out backup plans have no value if they don't work in a crisis. To ensure that they work, organizations must monitor and test them.

In the complex environment created by the coexistence of conventional physical backup and virtual backup, monitoring the backup and recovery process is more important than ever because, quite simply, more can go wrong. Managers need to have visibility into both the virtual and the physical devices involved in the backup and recovery processes (including the physical servers that host the VMs). They need alerts so that they can take immediate action when necessary. Finally, they need to have visibility into the whole process so they can locate bottlenecks or potential capacity problems before they occur. Without these measures, it's impossible to know if backup-related SLAs are being met.

Actually testing the ability to recover from a catastrophic failure is as important as monitoring the backup process, and this is an area where mid-sized businesses are remarkably lax. According to one recent study of businesses with fewer than 150 employees conducted by

SearchStorage, almost one third (32 percent) had *never* tested their backup plan. This is a grave error, because companies whose backup plans fail suffer extreme consequences. The statistics are grim:

- 60% of companies that lost their data shut down within 6 months of the disaster
- 93% of companies that lost their data center for 10 days or more due to a disaster filed for bankruptcy within one year of the disaster.
- 50% of businesses that found themselves without data management for this same time period filed for bankruptcy immediately

The lesson is that RTO and RPO need to be tested and measured. The objectives need to become actuals with proven metrics. Furthermore, this needs to be achieved in ways that are not disruptive to ongoing business processes.

Summary

Whether backup and recovery are viewed from a pure IT perspective or a business perspective, the rules have changed.

- VMs are playing an increasing role.
- Physical servers are not going away any time soon.
- Whether virtual or physical, servers have different backup and recovery requirements based on their business value.
- Using multiple point solutions to solve all these problems increases complexity. What's needed is a solution with a single pane of glass.

CA arcserve® addresses these new realities, and provides a comprehensive solution for virtual and physical environments, for backup and recovery of applications and data, for disaster recovery and replication of complete systems or applications, and for ensuring high availability of mission-critical services.

For more information on the CA arcserve Family of Products, and to download a free trial, please visit www.arcserve.com.