

A Practical Guide to Cost-effective Disaster Recovery Planning

Contents

Measuring Total Cost of Ownership.....	3
Measure Performance	4
Assess Your Risk through Regular Testing	4
Changing the Game through Virtualization	4

The global recession, increased competition requiring just-in-time processes, squeezed information technology (IT) budgets, explosive data growth, and new regulatory requirements have all increased the importance of disaster recovery (DR). As a result, organizations are now under pressure to create, re-evaluate and update their disaster recovery plans.

Looking across the DR industry today, there are hundreds of solutions whose value should be measured in three ways: cost, performance and risk.

Measuring Total Cost of Ownership

Cost is typically the most important concern for customers evaluating DR solutions, particularly during these tough economic times. Ideally, disasters are problems that never surface, so budget allocation is often difficult to secure for what should be unlikely use. On the other hand, when talking about the data center, any server in production is a server worth protecting, as it must have some level of business value. With this in mind, allocating budget appropriately and balancing protection costs with the business importance of the components of the data center is critical.

Over time, the market for server protection has evolved into two major categories or approaches which have further complicated the concerns around cost. The first approach consists of solutions built around the concept of infrastructure mirroring and redundancy: By mirroring the entire server environment, organizations are able to achieve the greatest degree of protection. The second approach to protection is simply to back-up or archive all of the data within the data center. Let's look at each in turn.

An infrastructure mirroring approach offers a fully redundant infrastructure, which provides tremendous performance for both Recovery Time Objective (RTO, or total time to go from an outage back to running in production) and Recovery Point Objective (RPO, or tolerance for data loss).

The problem with this approach is the total cost of ownership. Inherently duplicating anything already doubles initial cost, not to mention the additional costs of the tertiary components, along with the soft costs associated with implementation and maintenance.

While organizations can easily justify the expense of duplicating business-critical server workloads such as customer-facing applications (e.g. Web servers and online order processing), it is harder to find sufficient funds to protect workloads deemed less critical—such as email servers, internal Web servers or batch reporting applications—in this redundant manner.

In comparison, back-up or archive-based solutions leverage everything from inexpensive tape to increasingly economical disk. As a whole, these solutions tend to be very cost effective. The downside to this data-focused approach is the recovery performance.

If we look at the RTO performance of archive-based solutions, most of them tend to be quite poor. This is primarily due to the process of taking back-up data from a tape or disk, and rebuilding it back to a usable workload, which can be lengthy and complex.

The choice that customers are left with is expensive, redundant infrastructure versus cheap data back-up, and this creates a difficult budget allocation problem for customers. Statistics show that ultimately companies end up using 80 percent of their budget to protect only 20 percent of their servers.

Measure Performance

When looking at the performance associated with DR, the conversation becomes clearer if you break DR into three phases: replication, failover and failback.

In most solutions, an organization's focus has traditionally been placed on the first phase of the cycle: Replication. To date, data back-up solutions have been primarily very focused on the technologies and processes associated with keeping that data current, and those solutions now range from simple daily tape back-ups to sophisticated, synchronous Storage Area Network (SAN) based replication.

The reality is that in most cases, equal if not greater importance should be placed on the other two remaining phases: Failover and failback.

The solutions offering the best performance when looking at failover again trend towards complex and expensive redundancy-based approaches. However, customers who are already facing budget constraints, and have implemented more cost-effective back-up solutions, will find their failover processes to be lengthy and error prone, missing the mark on performance. The problem with these back-ups is the processes involved in converting raw data into a useable server workload state. Again, the choice between cost and performance has become an issue.

The final phase of the DR lifecycle, failback, is usually the most overlooked part of DR planning. With many solutions, especially the more cost-effective data back-up solutions, only a one way trip is considered. Once you land on your recovery site, there is no plan in place to get "back to normal". Obviously this thinking can lead to unexpected or unnecessary headaches, as you try to return to 'business as usual.'

Assess Your Risk through Regular Testing

Typical DR plans will include an annual test period or DR Event. With the speed of business and technology today, a full year can be a very long period of time, and the amount of change that can occur across a data center and in business process over a 12-month period can be tremendous. One reason for this inefficient testing standard is due once again to the fact that most solutions are focused on the front end of the DR Lifecycle. Testing process will often mirror the recovery process, and bring with them all the problems and complexities already mentioned with Failover.

Taking a practical view, only a thoroughly tested DR plan is a reliable DR plan. All too often testing is not adequate, and issues with the plan are not identified until restore procedures are executed, which by then is too late.

Another issue that arises is that DR infrastructure that is only touched during once-a-year testing periods inherently loses its value. If businesses were able to test more often, in an easier way, not only would their plans be safer and more reliable, but the infrastructure associated with DR would gain more day-to-day value.

Changing the Game through Virtualization

The emergence of virtualization has allowed businesses to change the dynamic by which IT looks at DR. However, whether it's due to virtualization technology still being in its early stage, or costs, or regulation and policies associated with the data center, virtualization still has not been adopted on a large scale in production environments.

That said, remote or DR sites can present organizations with a relatively safe environment to deploy virtualization (similar to how the first Virtual Machines (VMs) were deployed in the test lab). The cost savings achieved with virtualization become even more compelling at these secondary sites as the biggest issues with DR, as mentioned above, always revolve around cost.

Virtualization effectively allows IT to take the concept of infrastructure mirroring, and be able to achieve this with the lowest overall infrastructure investment possible—a virtual infrastructure. This drastically reduces the costs previously associated with this strategy. One simple, small footprint virtual server environment can now be used to mirror or protect a large footprint of physical production servers. Virtual machines can provide the same recovery workflow as traditional data archiving, with the flexibility and performance of a “boot in place” machine. Leveraging virtualization has already improved RTO performance by creating an archive that does not have to be built from the ground up like a traditional back-up.

From a testing standpoint, by leveraging VM snapshots, virtualization offers the ability to create a snapshot copy of an archive that can be booted in place and tested not only very easily and quickly, but with absolutely no impact on production.

Next generation solutions leveraging virtualization can now allow customers to effectively bridge the gap between mirroring and back-up, when protecting the physical servers in their data center. These emerging technology solutions enable users to efficiently protect more for less, alleviating the budget concerns of yesterday’s DR solutions.