Publication : CIOL
Date : 23 April 2008
Edition : Online

omnitech
Empowering Enterprises

News

Implementing DR and BCP

Should you go for a hosted DR site or an inhouse one? What's the best DR strategy for maximum system uptime? Which are some of the good tools for implementing DR? We answer all these questions and more in this story

Wednesday, April 23, 2008

After incidents like 9/11, Tsunami and Mumbai Floods-people have understood in no uncertain terms, the catastrophic consequences of disasters, either man made or natural. However, to prevent losses from compounding, CXOs have to figure out ways and means of keeping their core businesses intact in such situations. So, today if you go for any compliance certification, one of the pre-requisites is the kind of disaster recovery measures you have in place for your business. And just because of this reason, most businesses are opting for DR and BCP policies without even completely understanding the complexities involved. And still there are lots of misconceptions and myths surrounding DR and BCP deployments. In this article we try to demystify DR and BCP with some live case studies and implementation scenarios. But before we begin let's try to understand why we really need DR and BCP.

Preparing for disasters

Disasters could be natural, in the form of earthquakes, tornados, floods, etc or they could be manmade such as wars, militant attacks, accidents, etc. Plus, there is one more category of disasters which occur frequently-biological pandemics such as Bird Flu, Chicken Gunia, Plague, etc. Now all these disasters have different characteristics and affect one or more of the four pillars of Business Continuity-Workplace, Infrastructure, Data, and People. But today most of us understand DR as just Data Recovery, ie, if your data is corrupted or lost, you can recover it from some remote storage device. But that's not the only thing in DR. The consequences of a disaster could be more than just data loss. So, whenever disaster strikes, it can take away any or all of the four pillars of your business and your core business could come to a standstill. Nobody would even want to imagine facing such unfortunate events, but that doesn't mean they can't occur.

A DR and BCP deployment is like a medical policy, where you keep investing certain amount of money, and you only see the benefits of the policy when you fall ill. Otherwise, you don't see anything, except money going out of your pocket as premium. As we don't mind spending money for our medical insurance and plan for future, we should do the same for our core business. In a nut shell, a DR and BCP strategy is a medical insurance for your business. We'll now use a few examples to explain how disasters attack the various pillars of your business and how you can take effective precautions. Some of these may sound completely illogical or impossible, but then a disaster warn you before striking. So do take these examples with that in mind.

Let's assume that a large banking company runs its core business from a major city in India. One fine afternoon its network is attacked by cyber terrorists or there's a virus outbreak. In such a situation, the data integrity is lost. The easiest way to maneuver this disaster would be to immediately isolate the cyber attack on the branch and transfer the core job to a DR datacenter hosted at some other location. This would help users to immediately connect to remote DR servers and get back to work.

Take another scenario. One day the same city where the bank was operating from, encounters an epidemic. The Bird Flu virus hits the city, and being an airborne virus, infects anybody walking out in the open. So a city wide red alert is sounded, a curfew is enforced, and nobody can come out in the open. In such a scenario, all your pillars that constitute Business Continuity remain intact except human resources. So your data, equipment and workplace are intact but no one can come to the office and operate from there. So, the strategy to overcome such a problem should be different. Here you must have a DR site with not only data, but also with a backup of employees who can take over the charge of the center and finish the tasks from some other city.
Now let's take another example where an earth quake destroys the entire building, with the data center and all the

equipment. Here, even though peoples' lives might be saved, everything else would get destroyed. In such a situation, a remote DR site is required where you have all the necessary equipment, seating arrangements, data and even a recreation zone, where you can fly in your staff and let them get back to work in as less a time as possible. Such a DR site should not be in the same geographical location as the site in question, so that the calamity does not affect both sites at the same time. On the other hand, it should not be too far away so that it takes a lot of time to fly out people.

Outsourced or In-house?

After reading all these you must be wondering whether you should have a DR site with redundancy of all for all four pillars of BCP? The answer here should ideally be yes, for at least the first three pillars of BCP. Human resource is the only component for which a 100% redundancy is not feasible. There are ideally two ways in which you can get your DR and BCP site ready. The first and the traditional one is the in-house model where you have your own premises, equipment, and data kept offsite, so that in case a disaster strikes, you get back to that offsite center for DR.

The other and more contemporary way is to outsource Business Continuity to third party DR and BCP sites. This approach has some real benefits and can save you huge money and hassles, but both trends have positives as well as negatives. In the following section we try to discuss the pros and cons of both approaches.

OmniCenter: Managed DR Site Solution

Omnitech is an eighteen year old company based out of Mumbai. Its competencies include Managed Services (IT Infrastructure Management Services/ Remote Management Services, Business Continuity Planning / Disaster Recovery Services), Software Development Lab (Application Management and Maintenance Services) and Independent Software Test Lab (Software Testing Services).

In IT there are four layers where Disaster Recovery (DR) is required to be managed. The first and foremost is data, then comes equipment, site and last is people. DR is a part of Business Continuity Planning (BCP) and Omnitech's Managed DR Center, 'Omnicenter' is a part of BCP.

Most organizations in India do not have a corporate-wide BCM plan in place, and they store entire data backups at onsite locations only. So, why should an organization opt for a third party DR center? The most important reason is that by doing so they can concentrate on their core competency. Moreover, looking at the high rate of attrition in IT, you are assured of support in case an important employee in your IT team decides to leave the organization. Presently, BFSI and ITES sectors are increasingly getting more and more focused on DR/BCP.

Omnicenter is located at Mahape, Mumbai. It is a satellite DR for main hubs like Mumbai. It delivers end-to-end DR services in accordance with global standards and processes like SOX and Basel II. So, whenever disaster strikes at a client site, the team can directly come to the Omnicenter and start their operations. It has a fully operational facility with desktops that contain replica of the software at client side, along with other necessary peripherals such as scanners and printers.

Confidentiality is of utmost importance when it comes to data. Client data is either at their IDC or at Omnicenter's server, which is operated and managed by the client. Thus, the access rights are with the client. BS7799 and BS 25999 are observed at the center. There is high level of security, for eg, a particular location, say location 1 allocated to a customer would not be given access to another customer. There are cards programs where various levels of access rights are defined. Each customer operating in an environment is logically separated. Moving to data equipment, the workplace recovery available could be dedicated or syndicated. One seat could be claimed by more than one customer or the same seat could be dedicated to a particular customer. Dedicated seats ensure that workplace recovery is available to a particular customer throughout the year. They have an in-house product called Omnimonitor, which helps to monitor servers, networking systems, operating systems and databases in the center. It works without an agent and wherever there are threshold parameters set, and if those are achieved, the software would send an alert by mail or an SMS.

Omnitech has also proposed a chain of centers which will serve multiple purposes. One of those would be take care of operations in the vicinity. For eg, if a disaster strikes in Mumbai it does not make sense to have the DR center in Hyderabad; it is more viable to have a DR center in Nashik or Pune. However, such a center would not help if disaster were to strike regionally.

Awareness programs are also held for people, through certified professionals. A dry-run is done regularly. Clients are assisted in doing risk assessment services and business impact analysis, ie to figure which of the businesses are most critical, so that they need to be always available. Based on mission critical applications, the desired plan is defined based on recovery time objectives (RTO) and recovery point objectives (RPO). Based on these analysis, customers are advised to shortlist the services they should outsource. Then you have something called as disaster recovery consulting services. Once the customer gets into the practice of implementing a DR solution they are required to be managed. And once it is managed, they need to be audited. They follow APDIMA (assess, plan, design, implement, manage and audit) methodology.

In-house DR

The biggest problem with an in-House model is that if you plan to build a DR site with 100% redundancy of equipment, data and seats, then it's like having investing double the capital for the same task. And in case you don't encounter any disaster for years at a stretch, the ROI for the investment comes out to be zero. And then, it's not only about building a DR site initially and forgetting about it completely. You have to regularly monitor, manage and test equipment, and data kept at the site, so that in case of an emergency you know everything is working fine. Such a kind of monitoring and management and the associated drills require a huge amount of investment. So in a nut shell, it's like buying your own hospital instead of getting a medical insurance done.
There are some good things about an in-house DR site though. For instance, if security and data theft is a major concern, then you might not even think of opting for an outsourced solution as you might have reservations in sending your data regularly to places that are not under your complete control. For businesses such as stock exchanges, BFSI, etc an in-house model alone can work.

Outsourced DR

On the other hand, in a third party managed model you just have to pay for the number of seats you hire from the managed service provider (MSP). And you don't even have to bother about the management and monitoring of the equipment. All that is taken care by the MSP. So, you just paying for the number of seats and the investment is pretty less. Now depending on your requirements, you can even acquire shared or dedicated seats. In case of dedicated sets, you purchase a particular number of seats and those seats are isolated for ever unless and until you meet with a disaster and come to the MSP's DR site to use those seats.

However, in case of a shared model, the same seats are sold to more than one organization. As the number of organizations that hire seats grow, the investments decrease. If you get, let's say n number of seats for a 1:3 ratio, this would mean that the same seats would be sold to three organizations. And obviously, the investment for acquiring those seats would be around one third as compared to that for a 1:1 ratio. But there's a catch. Let's say three organizations from Mumbai acquired n number of seats for a 1:3 ratio, at a nearby third party managed DR site. In case of a disaster, the whole city would be affected all three companies would get affected. All three companies would try to reach the same DR site and acquire the n seats. This means they would require thrice the number of seats actually available. In such a situation all companies would have to make do with only one third of the resources allocated to them. This is the biggest disadvantage of a shared model. However, if you want to keep your investments low then with proper planning and selection of third party DR sites, you can solve most of these problems and have a fully functional DR site.

Now as we have already talked about what are the reasons and different ways of having DR and BCP policies. Let's now go a step further and make our hands dirty to give some of the DR and BCP related tools a try.

Virtualization and BCP

Virtualization has made its way in most of the enterprises, in one way or the other. Virtualization is also starting to change the way companies do Business continuity planning. Now, when a server crashes it doesn't necessarily means that we have to recover whole of the data to a physical server. You can easily recover a virtual server and continue all your business processes. Many enterprises are also opting for strategies where they have running replica VMs of the servers running their critical business applications; so in case of a server failure, users are automatically switched to the replica VM.

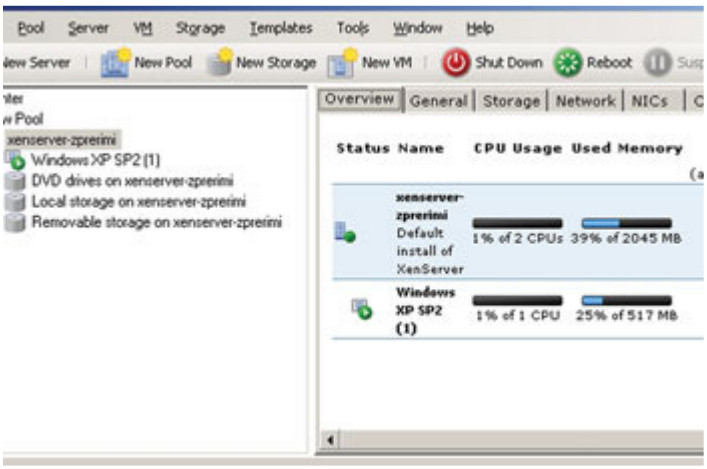

In Citrix XenCenter you can view status of all XenServers running in the resource pool as well as VMs through a single console

There are many virtualization solutions available which have gone a few steps further to provide BCP in their unique ways. One such solution is Citrix delivery center which delivers applications right on your desktop anywhere in the world. So, the applications can be delivered in India while being run on some other part of the world. The advantage here is, in case a disaster happens in the region/country where application was delivered and the whole site gets wiped out, the users can still access their business applications from their homes and another office. This also means that you need to create a mission critical DR site only at the place where you have the application server running. While at the local site you only worry about the workplace DR. This means you just need to have separate office building with dumb terminals which can connect to the main site in case of a disaster at local site. This not only makes implementation of DR site easy, but also saves good amount of money as you don't have to take care of data and equipment redundancy at the local site.

Citrix XenServer 4.1 beta

Citrix recently released beta version of Citrix XenServer 4.1, which is a part of Citrix delivery center. It is a server virtualization platform which uses Xen hypervisor to enable servers to host virtual machines simultaneously. XenServer allows users to create a Resource Pool by combining multiple Xen-enabled servers. Memory, storage, CPU and other resources of this storage pool can be dynamically controlled and allocated to either running or new virtual machines. With XenServer administrators can create multiple clusters of Resource Pools and manage them through a single console. This creates a virtualized data center environment which not only reduces the complexity of managing server, but also ensures that you use each server to its optimum level.

Another interesting feature of XenServer is live relocation capability called XenMotion. When a XenServer host in the pool requires maintenance, VMs running on that host can be easily relocated to other servers in the pool without disturbing the state of VM. Similarly, in case of hardware failure, VMs can automatically move to new resources ensuring that Business critical Applications are running smoothly at all times. If you are planning to run VMs with Windows you will require an AMD-V or Intel VT x86-based server with one or up to 32 CPUs. Recommended RAM is 2 GB and disk space of 60 GB. When XenServer is installed, it creates two 4 GB partitions on the machine; rest of the space can be used for virtual machines.

Deploying XenServer

For running XenCenter you need a Windows 2003,XP or Vista with .NET framework 2.0 or above and min 1 GB of RAM. To prepare a resource pool and enable XenMotion, you will need a shared storage on the network. Installing XenServer host is simple; it is made of stripped down Xen-enabled Linux OS, VM templates, a local storage repository for VMs and a management agent. To install just boot the server with XenServer installation CD, the installer will automatically detect all hardware present and will ask for the basic information depending on the configuration of the server. You have to provide things such as setting root password, DHCP, hostname etc. Once installed the server will reboot and is ready to be managed through XenCenter.

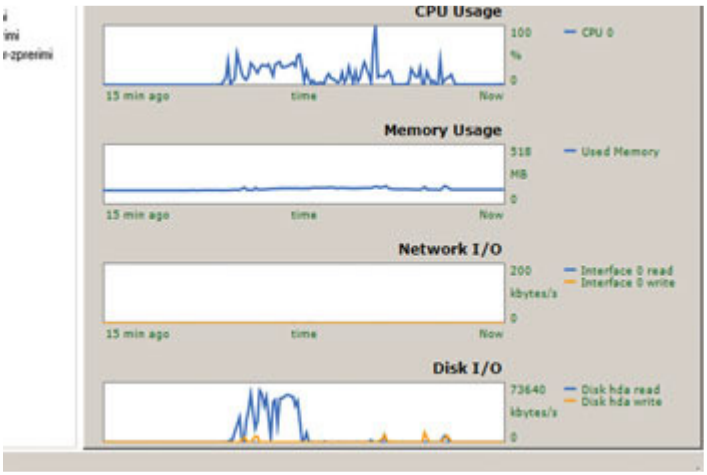

Citrix XenServer comes with Xentools, through which you can contantly monitor resources used by VMs running on the pool

After this, you need to install XenCenter on a Windows machine on the same network. Install XenCenter and then launch it from the Programs menu.

Now, before creating resource pool remember that for a resource pool each CPU has to be from the same vendor i.e. if you have one AMD-V CPU and second one as Intel VT then both of them cannot be used together. Resource Pool can be created on XenServer Host through CLI as well as through XenCenter management console running on any other host. To create a Resource Pool launch XenCenter from the programs menu and click on 'Connect to New Server' option. In the window which pops up, provide Hostname of the XenServer with username and password. Once connected to the XenServer, click on New Pool option on the menu bar, this will launch a New Pool wizard where you've to provide a name for the pool and select the Master XenServer and the slave servers. Once you finish this, the Resource Pool has been created and you are ready to host virtual machines on this pool. You can also add Storage Repository to this pool by right-clicking on the pool you just created and selecting the option New Storage Repository.

To create a Virtual Machine on the pool, from the menu click on New VM button. This will start New VM wizard, wherein you will found 26 templates for VMs; if you are planning to run any OS from those templates select that otherwise select Other Install media option. Further you'll be asked to provide name, location of source files, and how much CPU and memory should be used for the VM. Once the wizard gets finished, it will automatically start the VM and look for installation media on the specified location. After the VM has been created, you have to install XenServer tools on it to be able to measure its performance. To install these tools, simply click on the VM name and select the option Install XenServer tools. Now an installation wizard will start inside the VM. Just follow the on screen instructions to install the tools. Once installed, VM will reboot, and now you can monitor performance of this VM from anywhere in the network through XenCenter.

CA XOsoft Replication File Server

CA XOsoft Replication File Server is a real-time data replication solution for file servers and applications like MS Exchange, IIS, SQL, Oracle etc. It transfers changes to files instantly to the replica site which can be hosted locally or over WAN. This solution supports Automatic synchronization i.e. if replica or the master server reboots, the soln automatically resynchronizes master and the replica server. It comes with a single management console for performing all the operations and monitoring the replica as well as master servers. In case of a disaster, users can either be diverted to replica site or data can be easily restored to the master through easy wizards. Using the CA XOsoft Assured Recovery feature of this soln one can perform automated disaster recovery testing.

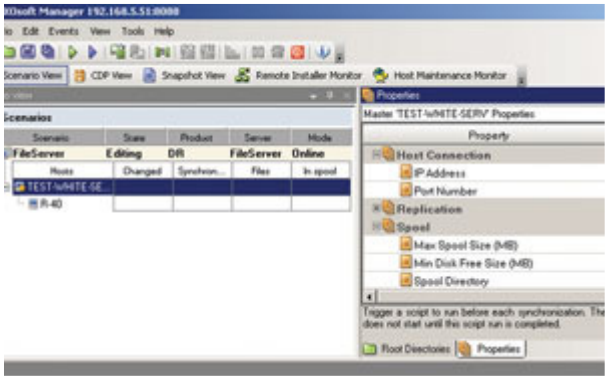

In CAXosoft you can easily create scenarios for replication of data from man site to DR site

Deploying CA XOsoft

The CA XOsoft has five main components namely CA XOsoft Control Service, Engine, Management Center, PowerShell and CA XOsoft CDP Repository. Deployment of these components can vary according to the architecture of enterprise networks and DR needs. When deploying this solution for high-availability purposes, its control services should not be deployed on master or replica server. Through this module administrator can control operations of CA XOsoft; it communicates with engines as well as the managers. It accepts requests from the Manager module and converts them into commands which are then passed to engines to perform a particular task. All critical scenario files are maintained on the server running control services. The soln is also responsible for authentication of users. Installing part of CA XOsoft is simple; first you need to install XOsoft Control Service, followed by XOsoft core Engine and XOsoft CDP Web Server.

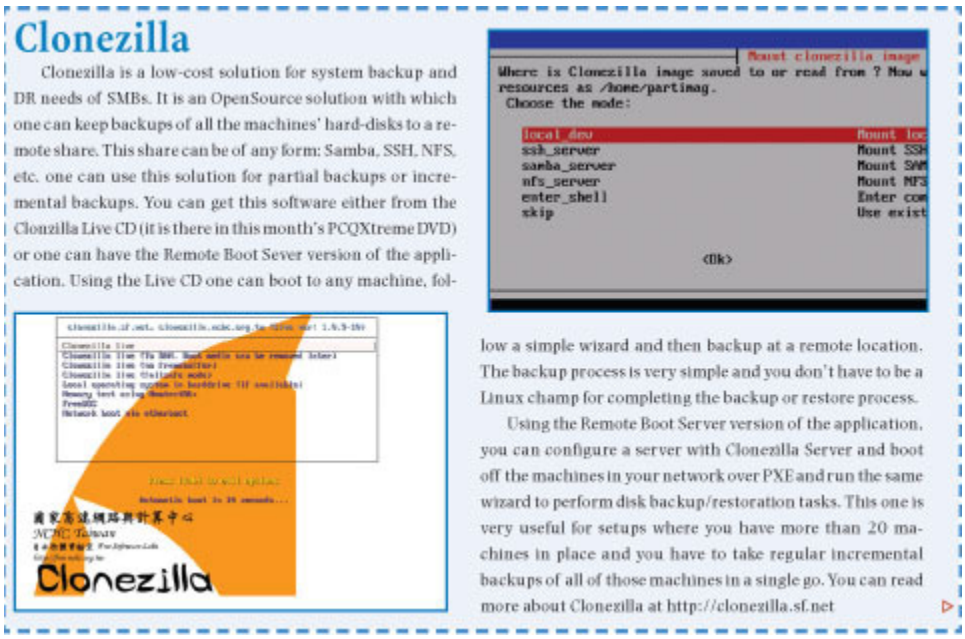

After installation, open a browser and type http://host_name:port_no/start_ page.aspx to open CA XOsoft Manager. It will ask for the username and password; provide the credentials you chose during installation. Now from the Quick Start pane, click on Scenario Management option. This will start CA XOsoft Manager. To create a new scenario go to Scenario menu and click on the New option, which will open Scenario Wizard. The wizard will then ask you to choose the server type for which you want to create the scenario, and then specify the master as well as replica hosts. Now, the wizard will verify the connectivity with replica and master hosts specified and whether the CA XOsoft engine is running on these hosts or not. If the engine is not detected, you will be asked for the credentials of the remote machine so that it can install Engine service on the host. Next choose the files on the master which Replica host has to replicate and also the location on replica host where data must be backed up. Configure properties of replica and master and click on 'Run Now' to start the data synchronization process. After this process finished, a final report will displaying time taken, files modified, directories created etc. This scenario can be run any time later on by clicking on Run button and could also be modified later on as per the needs.

To recover data from a replica, first stop the scenario. Then from the scenario folder select replica host and from the tools menu select Restore Data. This will start Recovery Method wizard, provide the information required and data will get restored on the master side. By default after successful synchronization XASoft will show a report.

5 Steps to an Effective Disaster Recovery Strategy

The following five strategies can help enterprise IT organizations implement robust high- availability and disaster strategies that maximize system availability for day-to-day operations.

Strategy One: Solve problems faster

Traditionally, one of the key challenges in executing timely disaster recovery was a delay in alerting IT staff to an outage, and subsequent problem diagnosis.

Advanced clustering technology notification and reporting capabilities can pinpoint when an outage occurs, and immediately notifies administrators about the problem.

Clustering technology then takes immediate action by starting up applications at the secondary data centers and connecting users to the new data center.

Administrators can then use configuration management tools to diagnose the cause of the downtime, such as identifying a change that may have been made by another administrator. The tools can display the nature and time of the change, speeding problem identification, and resolution. When the change is reversed, the normal operating environment can then be restored. With configuration management tools, data center administrators can be confident that their systems can prevent similar outages in the future.

Strategy Two: Automate recovery processes

For many organizations, system recovery is a manual process. It often requires time-consuming trouble- shooting to identify and solve the problem, and then administrators must rebuild the infrastructure step by step, including restarting servers, installing software, mounting data, starting up and configuring the software, and reconnecting users to the secondary site. Pressure builds on administrators as time, revenue, and customer loyalty slip away, and the potential for human error rises.

An automated approach, such as high-availability clustering, eliminates vast amounts of downtime compared to the traditional manual recovery process. If a system fails in the primary data center, the software can restart the application automatically on another server.

The administrator may be notified by a text message or via an email, and has visibility into problems at all times, but the series of activities required for maintaining business continuity is handled by the software; with limited action required by IT employees.

If a disaster threatens to cripple an entire data center, an automated approach can eliminate human error and reduce downtime by triggering failover of the critical applications to the secondary site.

The failover solution should determine which replicated data the application needs to continue operations. Then a single-click starts an automated procedure that restarts the application and connects the users to the secondary site.

Automated failover also addresses a common weakness in many disaster recovery plans ? the assumption that key employees will be available to physically enter the data center and manually restart applications. If the employees are unavailable, business continuity suffers. Automation helps reduce this potential point of system failure.

Strategy Three: Test your DR plan

Recent studies have shown that few companies test their DR plans on a regular basis, and as a result, most companies have little faith that their DR plans will work when needed.

Companies have been reluctant to conduct DR testing because testing often involves bringing down production systems, mobilizing a large segment of the work force, thus taking them off of more urgent projects, and forcing employees to work during inconvenient hours such as weekends or nights.

With automated failover capabilities, IT organizations can test recovery procedures using a copy of the production data ? without interrupting production, corrupting the data, or risking problems upon restarting a production application.

This capability means that tests can be run during business hours instead of over the weekend, hence reducing staff overtime. As an added benefit, automated tests run during peak production periods can re-create and approximate the conditions that would occur during a true failover situation.

Configuration management tools can also give more confidence to IT managers that their DR plans will work by ensuring that servers at DR sites are consistent with those in production sites. Server builds change over time as patches are implemented or as application dependencies change.

This can prevent clustered servers from working properly, as stand-by servers may have not received the latest patch or configuration updates. The latest configuration management tools can run consistency checks that will alert administrators that servers have drifted from the standard build. Action can then be taken to make the appropriate changes and ensure that HA/DR technology will work when called upon.

Anand Naik
Director, Systems Engineering, Symantec India

Strategy Four: Extract value from secondary sites

For most enterprise IT organizations, secondary sites are viewed strictly as cost centers, sitting idle much of the time. New advances in server provisioning software allow more value to be extracted from secondary sites, enabling them to be used for test development, quality assurance, or even less critical applications.

If a disaster strikes and the primary data center goes down, administrators can use provisioning software to automatically re-provision server resources to match the production environment.

Advanced clustering software also reduces high cost of the traditional condition that applications must be failed over to the identical hardware that the production applications run on. The most sophisticated clustering software permits failovers between different storage and server hardware within a data center or at remote sites.

With the flexibility to dynamically reconfigure and reallocate resources, the secondary site becomes a resource that can be used for multiple purposes the majority of the time, but can be quickly reverted to its backup designation when needed.

This underscores the value a secondary data center can deliver, making it more accessible to more companies.

Strategy Five: Achieve High Availability and Disaster Recovery in Virtual Environments

Server virtualization has become mainstream technology in today's server-centric data center. Server virtualization employs virtual machine technology that allows multiple operating systems to be run on a single server, each functioning independently of the others with its own operating system.

Restarting virtual servers at secondary sites has traditionally been a manual process, requiring personnel who may not be available during an actual disaster.

New clustering software allows companies to deploy server virtualization technology and receive the same automated disaster recovery benefits they can expect in their physical server environments. Furthermore, new high availability and disaster recovery tools are available that reduce the complexity of protecting and managing both physical and virtual server environments.

With clustering software, administrators can fail over applications from physical servers to virtual servers, and manage physical and virtual resources from a single graphical user interface. The result is that, through effective management of physical and virtual servers, hardware costs at secondary sites can be significantly reduced.