
Maximizing Data Center Uptime with Business Continuity Planning

Next to ensuring the safety of your employees, the most important business continuity task is resuming business critical operations. Having a co-location site for interruptions will help ensure your ability to provide the same level of service regardless of the circumstance.



Introduction

As with any successful Business Continuity Plan (BCP), executive-level support is crucial. But once the company has decided move forward with a comprehensive plan it is often left to the BCP team members to fill in the details. The support of the company and development of a mission statement that reinforces a commitment to the BCP can come in handy later on should the initial risk assessment return a higher cost than anticipated. Those costs can, and often do, include creating a co-location facility in order to resume business operations and provide the same level of service to both employees and customers in any circumstance.

The typical business continuity objective is to resume business operation to a functional level. This means providing only the necessary assets to keep the company operations moving forward. For some businesses this is perfectly acceptable, especially for smaller operations where there may be only a handful of servers to bring online (and not multinational locations and thousands of employees). But for Fortune 500 companies, functional operation may not be acceptable. Gartner Group reports that larger companies can lose more than \$42,000 per hour when operations are down, while other studies show the cost of lost data is nearly \$10,000 per megabyte¹. If operations are down longer than 48 hours, there is a high probability that the company may not be able to recover, and certainly won't be competitive. In fact, the U.S. National Fire Protection Agency reports that 43% of companies never resume business following a fire. Another 35% are out of business in three years. Herein lies the benefits and cost justification for a data center co-location facility. Duplicating the most business-critical systems will not only help reduce the possibility of a single point of failure in a disaster, it also eases the responsibilities of the data center manager by allowing them to be prepared for any outage - planned or unplanned.

1. Smith, David M. "The Cost of Lost Data." Graziadio Business Report Vol. 6 No. 3 Using Double-Take Availability to Create a Co-location Facility

Finding a Co-location Facility

So, where to begin? Let's assume you have a business continuity plan with executive approval, and a budget for the project. Now how do you find a co-location facility designed not only for disaster recovery but also for maximizing uptime for business-critical systems? When looking for a co-location facility there are a few factors to consider such as location, accessibility and how you will duplicate your entire data center configuration.

The co-location should be within driving distance, as your data center team will likely need to make frequent trips to set up hardware or infrastructure. (This is what distinguishes a co-location facility from a disaster recovery facility, as disaster recovery facilities are usually much farther away). Also, typically a disaster recovery facility does not have the infrastructure in place to begin resuming operations right away.

Depending on service level agreements, you may need to have hardware drop shipped, before you can begin the process of restoring images from tape or other means. The co-location facility is really a hot site, completely duplicated and ready to go at a moments notice. This ability provides true benefit to the company, not only from a disaster recovery stand point but also in the ability to better manage the data center day-to-day operations.

Once the hardware and infrastructure are in place the next step in the co-location setup is to transfer the systems. Using a continuous data protection product such as Double-Take Availability enables you to begin replicating data from one or many locations and allows you to decide when you are ready to switch over (and eliminate dependence on the last backup completed). The next task is to determine when you want to bring servers online, test or configure them. Most likely the data center manager's goal when creating a co-location will be to move a very complex web of data and applications (including Exchange Server 2003, BlackBerry® Enterprise Server, OpenText, Elite®, PeopleSoft®, Windows File® Sharing and multiple custom applications) without users experiencing any downtime.

Protect Your Business Critical Servers First

Using a continuous data protection product such as Double-Take Availability enables you to begin replicating data from one or many locations and allows you to decide when you are ready to switch over (and eliminate dependence on the last backup completed).

Starting with a phased approach will help you protect your most business-critical servers first and steadily get the co-location created as you roll out less critical servers. When the servers are in place at the co-location facility you can install Double-Take Availability to begin the replication process to transfer the data. Depending on the amount of data you have and bandwidth available, the length of this process will vary. However, because you aren't tied to an outage window and there isn't any downtime or user interruption, you no longer need to worry about how long the transfer process takes. Because you are using world-class replication technology you are assured to capture whatever changes are made to your production data.

Consolidating Exchange for Centralized Management

One of the options you may want to consider with your business continuity team when building the co-location is whether you should centralize your operations from your other sites. This provides several benefits:

Centralized Management – consolidating and centralizing remote Exchange servers reduces the need to have a separate set of Exchange experts at each location for regular maintenance. It also makes it more efficient for a centralized management team to maintain your primary communication and maximize uptime.

Clustering – Consolidating multiple single standalone Exchange servers also provides the ability to implement a cost-effective high availability solution. Microsoft Cluster Services is a good option, as all the Exchange servers could be placed on a multi-node Exchange cluster to help maximize uptime for your most business critical servers.

Bandwidth – Bandwidth between the two facilities is important, especially if you are implementing across a WAN or split node cluster. MSCS requires a 500ms round trip heartbeat to maintain the integrity of the communication between the nodes of the cluster.

However, replicating data using Double-Take Availability is less restrictive because of the built-in compression algorithm - reducing SQL® Server and Exchange transactions by as much as 80%, and file servers anywhere from 30-50% (depending on the type of data).

Multi-node Exchange clusters are more and more common and provide an array of benefits that range from being highly available systems to the ability to better manager servers without requiring downtime. However, MSCS still has vulnerabilities because of the single point of failure on the share storage. Should there be an issue with the disk array to which the clusters are attached - there isn't much you can do besides fix the disk as fast as possible. Striping the disks with a RAID 10 configuration will help prevent a disk failure, but not if it is a more severe issue such as a controller card or power supply. However, a co-location would provide another multi-node cluster that is setup and ready to stand in.

Using Double-Take Availability on each of the Exchange cluster nodes allows you to avoid a single point of failure. Double-Take real time replication resides on each node and monitors the transactional activity of Exchange mailboxes then replicates those changes to the appropriate node at the co-location facility. Because Double-Take is configured with its own cluster resource, if you have a less severe or node failure, the original node would roll to the available passive node, reconnect and continue replicating the data to the appropriate nodes at the co-location facility.

The Exchange resources remain offline at the co-location facility so the replicated changes of the source server transactions can be applied when they come across.

In the event all of the active nodes or the shared storage for those nodes fail, Double-Take Availability provides the ability to seamlessly failover all of the nodes to the standby Exchange cluster at the co-location. Failover monitoring is configured with defined parameters that provide a threshold for when to initiate failover (or determine when a failover condition is met). Should there be an intermittent outage the Double-Take target server will attempt to ping the source server for responses. Should the source server or cluster node

not reply within the predetermined number of attempts, a failover condition will be identified. When providing high availability via a WAN it is recommended that you be notified by the Double-Take server that a failover condition has been met, giving you the opportunity to investigate the issue versus automatically proceeding with the failover process. Double-Take can then provide one-click failover to the co-location facility. Double-Take can also provide automatic failover capabilities, however WAN availability and uptime requirements should be considered when deciding whether to use automatic or one-click failover. Double-Take's native SMTP can monitor Double-Take events, send e-mail communication and tie into existing management tools such as SCOM (Microsoft Systems Center Operations Manager).

When choosing one-click failover, Double-Take Availability will provide a message in the management console that a failover condition has been met and prompt you to initiate the failover process. Once you determine that the source cluster failure is severe enough to warrant executing the failover process, simply click "Yes".

With both data centers being in the same AD domain, you have the option of configuring the target Exchange cluster nodes to use the same named EVS (Exchange Virtual Server). One of the benefits of this option is that it doesn't require any Active Directory mailbox redirection (which Double-Take Availability automates for non-clustered Exchange servers). Because the Exchange files are installed with the same virtual name, the only actions that Double-Take Availability will initiate upon failover are starting the EVS on the target cluster and updating DNS if the two data centers are on different IP subnets - allowing Outlook users to automatically re-connect to the Exchange cluster at the co-location facility.

Configuring the Multi-Node Target Cluster

Double-Take Availability provides real-time replication between data centers - giving options to the data center managers when more maintenance is required. No longer do managers need to disappear for an entire weekend or upgrade systems in the middle of the night.

To begin the configuration of the Exchange multi-node target cluster install or configure the Microsoft Cluster service on the target cluster to create a unique cluster name (different than the source) so that it can participate on the network as a separate cluster.

1. Create an Exchange Group on the target cluster that matches the name of the Exchange group on the source.
2. Configure the disk resource on the target cluster to match the same disk configuration of the source. If the source cluster uses a D, E and F volume for Exchange data then the same disk resources are required on the target cluster.
3. Using Windows Add/Remove programs, install SMTP, NNTP, ASP.net and any other tools required by Exchange to be installed.
4. Install Microsoft Exchange on the target cluster into the same install path as the source. Apply any necessary hot fixes or service packs to Exchange to ensure that the versions of Exchange are identical.

Configuring a Like Named Exchange Virtual Server (EVS)

In order to configure the liked name EVS, the source Exchange group should be taken offline and planned for accordingly. Then, take the source Exchange Group offline to enable the creation of the same name on the target cluster. On the target cluster in the Exchange Group, create an IP Address resource and then the identical network name of the source exchange group. It is essential that the Network Name resource be identical as this is the key behind the 'like name' configuration. Under the resource properties, check 'DNS Registration Must Succeed' and 'Enable Kerberos'. In the Exchange Group on the target cluster a Microsoft Exchange System Attendant resource is created. This will configure the EVS on the target cluster using the same configuration as the source EVS. After the resources have been successfully created, bring all of the new Exchange resources online to verify integrity. Once the target EVS configuration is complete, take the target resources offline and bring the source cluster resource back online.

Exchange Maintenance Made Easy - Failover Isn't Just for Disasters Anymore

Now that Exchange clusters exist at both data centers, Double-Take Availability can be used to maximize Exchange uptime during maintenance periods. Because of the service level commitment of the co-location, the data center managers now have the ability to maintain any and all of their servers without interrupting any company operations. Failover isn't just for disasters anymore; Double-Take Availability provides real-time replication between data centers - giving options to the data center managers when more maintenance is required. No longer do managers need to disappear for an entire weekend or upgrade systems in the middle of the night. Failover can be executed from almost anywhere – including home. Because the systems are online at the co-location, once the failover has occurred the IT staff can take as long as needed to upgrade the systems or perform maintenance. More and more companies are failing over their business - critical servers to co-location facilities to perform routine server maintenance. Because most vendors are releasing hot fixes and or service patches once a month, the potential downtime alone for applying these fixes will interrupt your business operations. Multiply that by the number of applications you are running, by the number of operating systems, hardware, networking, communication and you too may wonder how a data center manager keeps any uptime at all.

Guaranteeing Setup and Reducing Duplicated Effort using Full-Server Failover

When a data center has hundreds of servers how do you guarantee that the co-location setup will function exactly as the production servers? The likelihood of having all of the installation instructions and configuration parameters for all of your production servers is slim. What if one of your vendors is no longer in business or the engineer who wrote the application it is no longer with the company? If that is the case, where do you begin to get that application setup at your co-location facility? Even if you had all the documentation

for every application in the data center, the amount of time it would take to just restore those applications to working order could be insurmountable. With Double-Take Full Server Failover and or Server Recovery Option the process is streamlined. Double-Take Availability will replicate the system state with all the program files and apply them to the target server when necessary. The only requirement is that the co-location server has the same base operating system installed that the source production server is running. Using these products will not only reduce the amount of time it takes to create all the applications, it guarantees that the configuration is exactly the same as your production source.

Full-Server Failover

Server consolidation and data center moves can deliver significant benefits – including cost savings, enhanced business continuity, optimized service management, and improved regulatory compliance. With the flexibility of Double-Take Availability and the Full-Server Failover feature's ability to work on all different applications, data center managers will be able to use Double-Take for a complete data center move with almost zero down time.

Replication, Consolidation and Data Center Migrations

The argument for consolidating servers from remote branches into a primary location could go either way. Some members of your team may say that distributing your environment is the best way to plan for a regional disaster. The others may say the more distributed you are the more you are at risk for interruptions. Whichever side you are on, you can hardly disagree that consolidating helps you manage your environment more efficiently. Especially if during the consolidation process you are virtualizing machines to help reduce overall data center requirements and logistics. And, managing you environment more efficiently can only help you resume business operations more quickly. Duplicating your data center cost can be reduced dramatically via virtualization and replication technology. By virtualizing, you can cut the cost of duplicating a data center in half by reducing the amount of hardware, space and power needed to resume business operations.

Summary

A Business Continuity Plan is certainly an important part of the justification for a co-location facility. But the plan and co-location alone won't maximize data center uptime. Flexibility in the types of controls that are implemented will be critical to your ability to provide server migration, consolidation, high availability and overall ease of management. Vision Solutions' Double-Take brand of server and operational recovery products not only will provide the technology necessary to deploy, move and migrate a duplicate data center, it also provides ease of management and the peace of mind to know that critical systems will be able to fully protect even the smallest interruption, whether planned or unplanned. And - long after your co-location is deployed Double-Take will continue to provide fast, efficient and effective business continuity.

Easy. Affordable. Innovative. Vision Solutions.

With over 25,000 customers globally, Vision Solutions is one of the industry's largest providers of high availability, disaster recovery and data management solutions for Windows, IBM i (i5/OS), AIX, Linux and Cloud environments. Vision's MIMIX, Double-Take and iTERA brands keep business-critical information continuously protected and available. With an emphasis on affordability and ease-of-use, Vision products and services help customers achieve their IT protection and recovery goals, which in-turn improves profitability, productivity, regulation compliance, customer satisfaction and quality of life. Vision Solutions oversees a global partner network that includes IBM, HP, Microsoft, VMware, Dell and hundreds of resellers and system integrators. Privately held by Thoma Bravo, Inc., Vision Solutions is headquartered in Irvine, California with development, support and sales offices worldwide.

For more information call 1 (800) 957-4511 or visit visionsolutions.com.



15300 Barranca Parkway
Irvine, CA 92618
800-957-4511
801-799-0300
visionsolutions.com

iTERA

MIMIX

Double-Take® AVAILABILITY™

© Copyright 2011, Vision Solutions, Inc. All rights reserved. IBM and Power Systems are trademarks of International Business Machines Corporation. Windows is a registered trademark of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds.