



Cloud Disaster Recovery: Five Key Steps to Avoid Risk and Protect Your Data

Executive Summary

Hosting applications on the cloud is appealing to IT organizations for many reasons. First, they benefit from the economy of scale and buying power that a large data center can muster. Second, cloud providers generally offer hardened data centers, back-up power sources and other capabilities that only large organizations can afford. Most important, IT organizations can take advantage of these services on a pay-as-you-go basis or guaranteed availability, depending upon organizational requirements.

DR preparedness is not the default configuration for most providers that offer cloud storage infrastructure.

Because many hosting providers maintain multiple data centers, IT managers often assume that disaster recovery (DR) is either inherent in the architecture or that DR is not an issue that warrants concern. However, DR preparedness is not the default configuration for most providers that offer cloud storage infrastructure. One multinational media firm for which Cognizant consulted

was surprised to learn that their cloud-hosted key applications had no DR coverage whatsoever. The 9/11 attacks taught us many things about IT disaster preparedness, including the fact that epic disasters are extremely rare but not unthinkable.

IT managers who have hosted applications through cloud providers, or are thinking of doing so, should perform the same DR due-diligence they would for in-house infrastructure. This includes assessing the risks, laying out the potential solutions and implementing a plan that meets the required service level at the least cost. In most cases, providers will be accommodating for requested DR configurations, but will not have their own reference architecture. IT managers should, therefore, be prepared to work with their cloud providers to help architect and specify an appropriate solution. IT organizations requiring assistance with this process can partner with us on any of the steps, from assessment to planning, architecting, specifying and implementing cloud DR solutions. This paper discusses the storage elements of disaster recovery planning.

Step 1: Assessing the Risks

Disaster risks can be thought of as a continuum of probability, but fall into three major categories:

- **Site disaster:** Fire, broken pipes or long-term power outage that can render the data center (or computer components) unusable for longer than the specified service agreement.
- **Area disaster:** Floods, tornados, hurricanes, major snow storms and even pandemics can render a data center unusable.
- **Regional disaster:** Terrorist attacks, financial failures, train derailments with toxic chemicals, pandemics, etc.

Proper data center design mitigates many of the risks associated with all three of these categories.

Data center location can mitigate likely weather events, train derailments and the like. Hardened data centers will have battery power backup (uninterruptable power supply) for sudden power loss, and power generators for longer term power needs (which could extend for as much as a few days).

If the cost of downtime is greater than the cost of a specific DR strategy, then the strategy is obviously worth the cost.

Even so, an area or regional disaster can render a data center unusable even if it continues to fully exist physically. One of the most likely risks that would require a disaster response is the financial failure of a hosting service.

A common guideline calls for 90 miles of separation between locations.

When the “tech bubble” burst early in the last decade, many hosting providers of the late 1990’s failed financially, leaving customers to scramble for infrastructure alternatives. While such an event does not carry the immediacy of a tornado response, organizations that have a plan in place to re-host applications will do so with less disruption and likely lower cost than those forced to scramble at the last minute. Obviously, a DR plan to move your applications to another site of the same hosting company would not be viable for such a contingency.

The quaint notion of IT personnel boarding a plane with backup tapes underarm headed for a recovery site is a thing of the past. There is no certainty that these sites will be available nor that transportation will be operating. As part of

the risk assessment, IT managers must consider how they would complete a recovery if transportation were severely restricted or if their chosen site could not be accessed.

Step 2: Determining Requirements

After completing an assessment of the risks, IT organizations need to classify their recovery requirements for the applications that are hosted. Requirements should be developed in the context of:

- **Recovery Point Objective (RPO):** RPO is the degree to which data loss can be tolerated. For example, “immediate RPO” indicates zero-tolerance for data loss. A 24-hour RPO would indicate that restoring data as of yesterday’s back-up is sufficient, resulting in a loss of all transactions and data after that time.
- **Recovery Time Objective (RTO):** RTO determines the maximum tolerable time for recovering the data and bringing the application back online. Note: The RTO should include the time to restart systems, databases and applications and re-route communications; simply restoring the data alone is not enough.

Both RPO and RTO requirements are driven by the cost of downtime. Cost of downtime can include actual loss of revenue, loss of employee productivity, loss of customer goodwill and loss of reputation. Tangible financial losses are the easiest to consider and most directly correlate to the cost of mitigation. If the cost of downtime is greater than the cost of a specific DR strategy, then the strategy is obviously worth the cost. Loss of customer goodwill and reputation are less tangible, but just as important. If the cost of acquiring and keeping a customer is high, then a stringent DR plan may be worth it. In any event, it is a business decision weighing risks and consequences against costs.

Step 3: Understanding DR Options

The most basic element of a DR plan is getting the data outside the data center. But how far outside? If the secure facility is less than 10 miles from the data center the distance is insufficient to guard against area disasters and pandemics. It is possible that neither the data center nor the secure storage facility would be accessible at the same time. A common guideline calls for 90 miles of separation between locations to guard against such events as “dirty” bombs. IT organizations can decide for themselves how far is far enough.

Data can be stored offsite asynchronously on tape or disk, or synchronously (disk only).

- **Backup to Tape and Offsite Storage:** Tape remains the cheapest method for moving data to a second site or archiving it. There are some “gotchas” IT managers should consider with tape backup:
 - Tape format: The format between the source and the target must be compatible. For example, LTO tapes cannot be read by DLT tape drives. Beware of different generations of the same technology as well.
 - Back-up and recovery software (B/R): B/R software applications write the data in proprietary formats (unless explicitly specified to write in TAR or CPIO formats). For example, a tape created by Symantec NetBackup cannot be restored by CommVault Galaxy.
- **Asynchronous Site-to-Site Data Replication:** This method moves data offsite, but to magnetic disk drives vs. magnetic tape. Disk backups can significantly reduce recovery time in the event of a disaster.
 - Most B/R applications can back-up to disk using compression and/or de-duplication technology. Thus, the back-up image is much smaller than the actual data image. With change-only back-up methods, transmission bandwidth and data storage requirements are minimized.
 - Virtual tape libraries (VTLs) are specialized storage devices (disk only or disk-to-tape) that can further automate the DR process.
- **Synchronous Site-to-Site Data Replication:** Synchronous data replication ensures every piece of data entered or changed is simultaneously replicated. While synchronous replication is typically the most expensive offsite replication option, for some critical applications the cost may be justified. Synchronization delivers an immediate RPO and an RTO limited to the time that it takes to declare a disaster, restart the application from the second site and re-establish communication.

Synchronous DR does not eliminate the need for a back-up and recovery solution. Even where

the cost of synchronous data replication is justified, synchronous backup will typically be combined with asynchronous methods for point-in-time restore capabilities supporting rollback to previous application or file versions. Whatever the backup methodology, it is important to ensure that the system image (i.e., the operating system) be included with the back-up set so that the entire environment can be recreated if necessary.

Step 4: Auditing Cloud Providers

Cloud providers should be willing to provide users with documentation regarding their data center protection strategies and, in fact, many have published literature describing these features. IT managers should compare this against their own list of requirements, just as they would for their own data center. They should examine it for location and should not assume that the hosting company considered nearby rail facilities or manufacturing operations that use toxic chemicals when choosing the site.

IT organizations should also understand the range of data protection solutions offered by the provider. Most offer daily back-up to disk capabilities and some supplement that with periodic tape backup (e.g., weekly). However, few organizations realize that off-site tape transfer is rarely included in the base service. So, while on-site backups can help recovery from data corruption, and inadvertent data deletion and allow point-in-time restores, they obviously do little to protect organizations from any of the disaster scenarios described previously.

Elements that should be considered in a Cloud Provider audit include:

- Location
- Possible events
- Power grid/communications considerations and contingencies
- Proximity to potential terrorist targets (e.g., airports, seaports, national landmarks)
- Relationship to recovery destinations
- Data center hardening features
- Vendor’s DR contingencies

Synchronous DR does not eliminate the need for a back-up and recovery solution.

Step 5: Implementing and Managing Your Cloud DR Solution

Over time, your organization will evolve. Applications will be added, enhanced or abandoned; offices will open and close; new technologies and formats will emerge. Moving forward, many cloud providers will likely be on the buying or selling side of a data center acquisition and integration. As important as it is to evaluate and select the right provider and the right DR solution, it is equally important to review your DR requirements and solutions over time. Make sure your cloud provider has instituted a process for simulating and testing your DR solution and ensuring that all systems are performing as promised. A rolling quarterly test of a subset of applications may be sufficient, as long as most or all of your systems are eventually tested on an annual basis. Internal requirements should also be reviewed quarterly. Plan on a comprehensive audit of your

Plan on a thorough audit of your cloud provider requirements and solution annually, as a DR recovery solution that was appropriate a year ago may become non-operative over time.

cloud provider requirements and solution annually, as a DR recovery solution that was appropriate a year ago may become non-operative over time.

Proper Planning Prevents Poor Performance

IT managers should not passively assume that their cloud hosting provider has the disaster recovery contingency addressed. In most cases, it is a capability that must be explicitly requested. Moreover, IT managers should be prepared to drive the conversation and deliver specifications to the provider. Most providers are accommodating but often lack the expertise to guide customers toward an appropriate solution.

Just as with an in-house solution, IT managers should assess the potential risks for disasters and the impact of a protracted recovery. Any solution must consider the necessary RPO and RTO required by the application. It is not necessary to have the ultimate recovery for every application, so IT managers should balance the cost of providing a specific RPO/RTO against the cost of downtime for the application.

About the Author

Phil Goodwin is a Senior Manager and Principal Architect in Cognizant's IT Infrastructure Services group, where he assists clients in the development of adaptive storage architectures, storage management best practices, back-up and recovery, disaster recovery and data archiving. Phil holds a BSBA in Marketing and a Master of Technology Management from the University of Denver. He can be reached at Phil.Goodwin@cognizant.com.

About Cognizant

Cognizant (NASDAQ: CTSH) is a leading provider of information technology, consulting, and business process outsourcing services. Cognizant's single-minded passion is to dedicate our global technology and innovation know-how, our industry expertise and worldwide resources to working together with clients to make their businesses stronger. With over 50 global delivery centers and more than 100,000 employees as of December 1, 2010, we combine a unique global delivery model infused with a distinct culture of customer satisfaction. A member of the NASDAQ-100 Index and S&P 500 Index, Cognizant is a Forbes Global 2000 company and a member of the Fortune 1000 and is ranked among the top information technology companies in BusinessWeek's Hot Growth and Top 50 Performers listings.

Visit us online at www.cognizant.com for more information.



World Headquarters

500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277
Email: inquiry@cognizant.com

European Headquarters

Haymarket House
28-29 Haymarket
London SW1Y 4SP UK
Phone: +44 (0) 20 7321 4888
Fax: +44 (0) 20 7321 4890
Email: infouk@cognizant.com

India Operations Headquarters

#5/535, Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060
Email: inquiryindia@cognizant.com