



SWIFTSTACK WHITEPAPER

Technology Use Cases:

# How to Improve Disaster Recovery with Object Storage

November 2014

Object storage is gaining widespread interest from IT professionals because it provides an answer to the challenge of cost-effectively scaling storage to meet the unrelenting growth of unstructured data. However, as more organizations deploy object storage they are discovering yet another truth: Object storage is a breakthrough technology that can reduce the costs and simplify the processes involved in disaster recovery.

With object storage, all storage nodes, regardless of location, are managed as one single pool, with no file system hierarchy. When an object is stored, it gets a unique identifier: Then, when it is updated or changed, it is stored as a new object.

The architecture of an object storage platform is inherently fault tolerant, meaning that organizations can use commodity hardware and sustain dramatic amounts of hardware failure without experiencing data loss. For example, if a particular server or storage device goes down, it can have no impact on operations because the data it contains is already stored on multiple devices at different physical locations within the infrastructure.

This can eliminate the need to set up separate sites specifically for disaster recovery, which means less costs, less management, less maintenance and, importantly, far greater speed, accuracy and simplicity in recovery. In certain use cases, there is in fact no recovery—even if there is a failure, no applications ever go down and no user ever knows that there was a problem. In this white paper, we examine two types of use cases where object storage is changing the dynamics of disaster recovery.

### **Use Case No. 1: Backup and Recovery to a Storage Target**

In doing backup and recovery, you typically need a storage target. The best storage target is one that is going to be available across multiple locations. With an object storage solution, you are creating a single storage environment that is deployed across multiple data centers. So when data is put into the object storage system, it ultimately exists in all of those locations. Therefore, if one of the sites is inaccessible, you can still get to that backup and recovery data from a number of other locations.

This type of disaster recovery deployment has been utilized successfully by Budd Van Lines, a New Jersey-based nationwide moving company. Budd has 150 workers in its office staff and more than 300 workers in the field. The IT department consists of just five people, delivering 24/7 compute services.

Because trucks need to keep moving around the clock and across the country, Budd required a storage infrastructure that would continue working even if a natural disaster occurred on one coast. In addition, the wide distribution between regional offices and field workers made it imperative that Budd deploy high-performing storage distributed across and replicated between multiple sites.

To address its challenges in storage and disaster recovery, Budd turned to the SwiftStack Object Storage Platform and Filesystem Gateway, which was brought in to replace SAN-based archiving and disaster recovery backups for virtual server and virtual desktop snapshots. While the snapshots were relatively small at 10-15 terabytes (TB), SwiftStack's ability to cost effectively replicate them between sites was critical to meeting disaster recovery goals.

Budd started with three SwiftStack nodes in each of its New Jersey and California offices. As the deployment expands, Budd has been adding at least one SwiftStack node in four other regional service centers, ensuring that dispatch order documents are accessible even if the site is completely cut off from the other offices.

In addition to disaster recovery, Budd has been able to use the SwiftStack object storage to reduce the costs of virtual machine backups, while also simplifying management. Budd is using SwiftStack to capture and replicate virtual machine (VM) backups produced by Veeam Backup & Replication tools. These backups drive cost savings by eliminating the need to have alternate copies on different media other than the network-attached storage (NAS). In addition, because of SwiftStack's replication among regional offices, there is no need to keep off-site backups of the VMs up to date and accessible.

Another important point to consider in thinking about disaster recovery is that an object storage platform is much more fault-tolerant than any storage platform you will have experienced before. You can sustain significant amounts of hardware failure without data loss, which gives you the ability to achieve significant cost savings through the use of commodity hardware throughout the storage infrastructure.

For example, when Budd Van Lines did a full total-cost-of-ownership analysis, it determined that the SwiftStack Object Storage Platform was far less expensive than alternate solutions when measured on a cost-per-terabyte basis. The SwiftStack solution came in at \$100 per TB for the first year, and only \$35 per TB for the second and third years. Comparable cloud-based and appliance-based solutions were orders of magnitude more expensive.

## **Use Case No. 2: Building Applications Around the Storage Environment**

For organizations that build their applications around the storage environment, there is actually no disaster recovery—because failures don't affect the applications, and there is nothing to recover. The data is already there because it has been automatically replicated across other physical sites and is always available.

Here's how disaster recovery works in this use case:

When you build the application with object storage, the data is stored in multiple locations from the beginning. In the event of a failure at one of the data centers, the DNS will route users to an available location. Users don't experience any downtime or even performance gaps in their applications and, in fact, they will likely be completely unaware that a failure has even taken place.

An example of this would be in a file sync-and-share application. If a file (object) is housed in an object storage environment, and a user makes a request to retrieve the file, the request will be routed through the DNS to determine where to get that file. If the file doesn't exist at one site, the request will be rerouted to another site where it does exist.

If a storage node is down, the DNS will automatically and transparently reroute to another node at another location. Users continue to be productive without even knowing an event has taken place. It's like a built-in disaster recovery plan, only there is no need for manual intervention to recover.

This level of resiliency is one of the clear advantages of an object storage platform, and it is particularly useful for applications in which users are dealing with large amounts of unstructured data and can't afford any disruptions, in either the availability of data for their applications, or in performance. Among the obvious use cases where this can be critical are video streaming, multimedia content creation, graphics and computer-aided design, file sync and share, and a wide range of high-performance computing and life sciences/research and development activities, such as genomics.

## Conclusion

In evaluating object storage platforms that are capable of delivering these types of advantages for disaster recovery, there are certain factors you will want to consider. First, it is essential that you deploy a software-defined object storage platform. Software-defined storage (SDS) is the next generation of storage deployment and management and eventually will be the model for how all storage is deployed. With object storage, an SDS model is necessary so that you can pool all storage resources and abstract the management functions to take advantage of commodity hardware, centralized management and the myriad benefits of object storage.

A second important consideration is to deploy a software-defined object storage solution that is built on open source OpenStack Swift object storage, which is a multi-tenant, highly available and durable object storage system designed to store large amounts of data at low cost via a RESTful HTTP API. Swift is already one of the industry standards in object storage, powering many of the world's leading public cloud object storage solutions.

Finally, you also want to make sure that you deploy a platform that provides seamless integration with your existing NAS solutions to enable existing applications to integrate with your object storage platform. When you evaluate vendors, make sure they deliver a file system gateway solution so that you are not creating yet another storage silo and are, instead, creating a unified and centralized storage platform that can be managed from a single pane of glass.

One of the first steps you should take toward deploying object storage for disaster recovery is to contact SwiftStack. SwiftStack was formed to deliver software-defined object storage solutions, and thus has no legacy environment to protect. The company's expertise in object storage is widely acknowledged, and it plays a prominent role in the open source Swift community. Furthermore, SwiftStack's solutions will enable your organization to put in the right foundation for object storage, to not only address your disaster recovery requirements, but to also reduce costs, simplify manageability and enable massive scale for unstructured data across your entire storage footprint. Ready to take the next step? **Here's where to start.**