SGI NAS
Data Disaster Recovery & Business Continuity

Business-Critical Data: Use It, Never Lose It

Data Recovery Issues

Data loss disasters can come in many forms, ranging from human errors to disastrous acts driven by external circumstances (fires, floods, computer viruses, power failures) that cannot, for the most part, be controlled. What would happen to your business if you had a major data loss?

Data protection solutions need to be an integral part of any business continuity plan and, for many businesses, are mandated compliance requirements.

The Challenge

Implementing a data recovery plan is one thing, but less than optimal solutions can be costly and the source of huge failure points. With the typical data warehouse doubling in size every six to nine months, companies need reliable and accurate replication of business data, along with dependable failover, as critical requirements for data recovery success.

In most cases, regular backups of pure data are not considered to be adequate for a disaster recovery plan for business-critical data. What is needed is business continuity planning and disaster recovery procedures that incorporate high-availability of all IT services and IT systems.

Solution Overview

SGI NAS combines unlimited full or incremental snapshot capabilities with unlimited cloning. Software RAID solutions can be implemented by most operating systems as a component of the file system (ZFS), thereby providing a single virtual device. Mirroring is the replication of logical disk volumes onto separate physical hard disks in real time to ensure continuous availability—in short, extra safeguards via data replication. In addition, SGI NAS provides clustering capabilities, which enable an HA cluster to connect geographically separated nodes, providing greater data protection, along with cost reductions, when compared with legacy systems.

SGI NAS also includes **Auto-CDP (Automated Continuous Data Protection)**, which provides block-based replication that optimizes recovery point objectives (RPO), recovery time objectives (RTO), and backups. Auto-CDP promotes remote mirroring capabilities (synchronous).

Auto-Sync replicates files between two different SGI NAS appliances, asynchronously. It re-creates source snapshots (data and metadata) at the destination.

Auto-Tier has a capability to snapshot its (tiering) destination upon each successful execution, as opposed to auto-sync, which actually transfers existing snapshots from the source to its (syncing) destination.

Both Auto-Sync and Auto-Tier are designed from the ground up to use a variety of transports (a.k.a. protocols), which the required flexibility to execute over Internet and Intranet, from behind a firewall and in the environment that requires extra security.

Auto-Sync and Auto-Tier support incremental intelligent replication, whereby, once the initial replication is done, only accumulated differences (the delta) between the source and the destination is transferred during all subsequent service runs.

Both Auto-Sync and Auto-Tier support any schedule. You can schedule the services to run every minute, every hour at a given minute of the hour, every few hours, every day or at a certain day of every second month. Whatever is required by business needs.

Data Disaster Recovery Features

- SGI NAS data recovery and business continuity solutions deliver unparalleled flexibility in supported hardware, software, and applications.
- **Data Consistency:** Identify data corruption and self-heal those errors.
- **Data Redundancy:** Physical disks are grouped to redundant logical volumes (RAID sets).
- **Data Replication:** RPO and RTO values are reduced by the use of data replication techniques.
- **Snapshots and Cloning:** Automatic archiving and synchronizing of local and remote data to one or multiple destinations (asynchronously).
- **Clustering:** Provides high-availability of services, applications, and business processes. Clustering also provides load-balancing and horizontal scalability of services.

Target Values

Recovery Point Objective (RPO) and Recovery Time Objective (RTO):

RPO and RTO values are reduced significantly by the use of data replication techniques, especially in comparison to traditional backup and restore procedures. Data replication enables you to provide business process continuity even in case of an entire site failing.

RPO defines the interval between two backups. It also defines the maximum amount of data and number of transactions permitted to be lost.

RTO defines the time for data service recovery. Ideally, both values should be near zero.

During the development of disaster recovery procedures, RPO and RTO values are defined for the business processes and matched to the capabilities of the underlying IT infrastructure.

With replication and deduplication, data is scrubbed before transfer, thereby creating only one transfer and lowering, by significant amounts, the total volume of data sent.

Data Consistency

ZFS is a transactional, object-based, 128-bit file system supporting unlimited storage capacity. As a unique feature, ZFS calculates the checksums of metadata, user data, and additional copies of data-blocks (so called ditto-blocks) per data set.

Utilizing this information, ZFS is capable of identifying data corruption and self-healing the errors.

ZFS supports software-RAID mirror (RAID1), RAIDZ (similar to RAID5), and RAIDZ2 (similar to RAID6).

ZFS combines data management functionalities, storage and disk management, volume management, and file system management into one. The use of a RAID controller and battery-backed NVRAM becomes redundant with the use of ZFS, so potential data loss caused by possible failure of those components is avoided, which adds reliability to the SGI NAS solution.

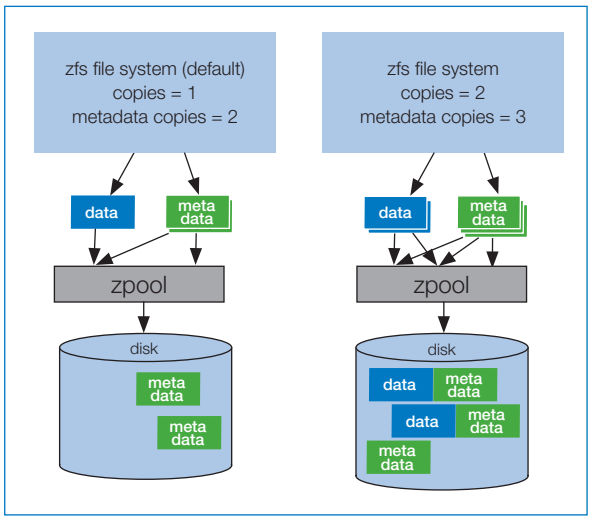

Snapshots and Cloning

Automatic archiving and synchronizing of local and remote data to one or multiple destinations asynchronously is provided by the SGI NAS Auto-Sync Plug-in.

SGI NAS, utilizes the implementation of ZFS as a copy-on-write (COW) file system, creating a new block for each operation instead of overwriting the original data block like traditional file systems.

Due to COW, snapshots are inherent in ZFS and created immediately and online without any further impact on the system.

The RPO value for data backup with snapshots is almost zero. User can access snapshots to restore data, as well as single files by themselves, thereby bringing the RTO value also to almost zero.

Continuous Data Protection

Synchronous real-time block-level data replication is provided by the SGI NAS Auto-CDP Plug-in.

Volumes can be mirrored seamlessly between geographically distant appliances providing security against site failures.

Clustering

Whereas data replication supports high-availability of data, clustering provides high-availability of services, applications, and business processes.

In addition, clustering may provide load balancing and horizontal scalability of services. Special cluster configurations (stretched cluster) protect against service loss in case all systems at one site fail (for example, a power outage).

Active / Passive Cluster

Active / Passive clusters are failover solutions and easy to administer. In case the primary system fails, the passive system automatically takes over the primary role and all its services. The RPO value in this configuration is zero. The RTO value depends on the configuration and amount of services that need to failover.

In case of a standby cluster configuration, the RTO values increase with the time needed for the system boot.

Active / Active Cluster

SGI recommends the use of Active / Active Clusters. In this configuration, all servers are active and services, as well as workloads, are distributed (load balanced) across all systems.

In case of failure of one system, failed services automatically are taken over by the remaining system.

Compared to the Active / Passive configuration, the RTO value in this configuration is reduced by 50% because half of the services already are running on the second node.

Stretched Cluster

SGI NAS supports cluster configurations where the cluster nodes are located in two independent, geographically distant, data centers, miles away from each other.

In case of a disaster and a complete loss of one data center, all services are taken over by the second data center so that organizations are able to continue to conduct business.

Shared Storage

In general, storage access in Active / Active Cluster configurations is divided into “shared nothing” and “shared-all” configurations. In “shared-nothing”, all cluster nodes access their own dedicated data partitions.

In case of a “shared-all” configuration, all cluster nodes access the data concurrently. While providing load balancing and scalability, this configuration supports full resilience by providing immediate access to all common data to the second node in case of a failover.

In case one storage system fails, SGI NAS supports spreading of ZFS RAID sets across two or more storage systems. This way each JBOD represents one disk of a ZFS RAID set.

In Conclusion

No business is immune to disaster or downtime. And every business, regardless of size, is obligated to have some level of disaster recovery in place to ensure ongoing business commitments and competitiveness within the marketplace.

SGI NAS offers a full range of enterprise-class features to protect your data and to provide high-availability and business continuity for your data, as well as for your applications and services.

Many of SGI NAS’s features are inherent in ZFS, such as snapshots, clones, disk management, volume management, the functionality to create redundant disk groups, and file system management with self-healing capabilities.

Additional plug-ins protect your business against system failure and data loss, even in extreme cases such as entire data centers failing.

SGI NAS delivers the flexibility and control that is expected by both IT and end users. SGI NAS simplifies IT’s task by extending the benefits of OpenStorage to open backup and disaster recovery.

Global Sales and Support: sgi.com/global