

Data Protection & Disaster Recovery

Secured Virtualized Applications with in-built backup and disaster recovery Capabilities



Organizations cannot survive in today's environment without a solid disaster recovery (DR) plan. Guaranteeing data protection is vital for business continuity.

Disasters manifest in various scales and areas of infrastructure. In all situations, businesses react to the impact after the effects ripple through BAU leading to many unforeseen consequences. Referencing a white paper released from the IT Disaster Recovery Preparedness (DRP) Council suggest that one hour of downtime can cost small companies as much as \$8,000, midsize companies up to \$74,000, and large enterprises up to \$700,000.

Challenge

Implementing a solid Disaster Recovery (DR) plan without spending a lot of time and money is a challenge especially when the plan is based on multiple vendor solutions and complex licensing models. A Disaster Recovery plan should be invulnerable and scalable and at the same time flexible.

The key factors to consider while building your Disaster Recovery Plan are:

Business continuity and time to recover

Before making any decisions, the recovery time has to be within the business needs. Even small difference in this number may have a big impact in business continuity.

Skilled staff

Level of expertise needed to achieve full recovery from a disaster depends on the complexity of the solution. Often, companies realize too late that the solutions are complicated to manage. Hiring a DR expert to help with the recovery process will increase the cost and a complex multi-vendor strategy may force you to. A failed recovery would lead to increased downtime, much higher cost and business loss.

Regulatory compliance

Downtime tolerance may vary for regulated industries such as healthcare or financial services. In some cases, it is close to zero and eventually it becomes mandatory for the continuity of the business. Finding a solution that will enable companies to comply with the regulatory issues is not a matter of choice.

Backup

Frequency of backups, automatic backups and the possibility of deleting older versions should be taken into account when choosing any disaster recovery solution.

Maintenance

Different solutions require different levels of maintenance. You should take into consideration the resources required for the maintenance.

Cost

Deploying the most effective Disaster Recovery solution for your business comes at a cost but choosing the right solution will save you much more. Traditional Disaster Recovery solutions require the use of WAN/SAN which means companies need a secondary data center as a backup. While running multiple applications DR needs could be application specific and may add to the cost. A cloud-based solution may look cost effective but will work out more expensive with scale and level. A cloud-based solution could be application specific.

Back-up Options

Backup/Restoration

A full copy of the system and data files at regular intervals. Copy is moved to a remote secondary or backup infrastructure. When the main system goes down, the most recent version stored in the secondary system can be restored and workloads can be shifted to this location until the main one is fixed. Downside is cost, complex management layers and risk of partial data loss.

Mirroring

Two identical systems run all the time and the second one replicates all the activity on the first, close to real time. When the main system is down, the workload shifts to the secondary one. High cost is the downside of this option.

DRaaS (Disaster Recovery as a Service)

Provides data and compute redundancy by utilizing HiveIO Distributed Cloud intelligence. HiveIO replaces archaic architectural methods such as Clustering, Virtual Centers, Database Archives and long lead times in transmitting duplicate data payloads with Message Bus Technology and deduplicated data packets thereby improving Disaster Recovery drastically.


Why Hive Fabric™ for Disaster Recovery

Meeting a tight recovery window requires seamless access to backup data. Consistent with many IT systems, a disaster recovery process fits into the scope of hyperconverged. HiveIO offers a disruptive and comprehensive hyperconverged solution which includes fully automated backup and recovery features coined DRaaS. Utilizing these features eliminates the need for DR experts in the event of a disaster. As a result, the recovery process transacts faster with the added benefit of mitigating human intervention thereby reducing error. Hive Fabric includes Disaster Recovery tools such as Live Migration and a Native DRS that enables instant recovery of virtual machines utilizing non-persistent/ persistent with gold images and shared storage, such as Hive SDS Shared, CIF, CEPH, NFS, Amazon Gateway to S3/ Glacier, across private datacenter WAN/ LAN or public storage.

Hive deploys quickly whilst removing the need for expensive setup, headcount talent and expensive licensing enabling the entire data center and data recovery process. In the event of a disaster recovery, the customer does not suffer a heavy CAPEX and OPEX burden. Hive Fabric transforms a private data center or a cloud environment into a DRaaS and orchestrates it all.

Hive Fabric™

Hive Fabric is a Software-Defined Data Center Solution that delivers commoditized cloud computing consistent with the public cloud on customer private premise. We accomplish this through engineered simplicity and performance at a fraction of the cost associated with both public and private environments. Hive Fabric's hypervisor is based on the industry-leading KVM infrastructure to virtualize and abstract the physical building blocks in the data center. With the lowest overhead of any hypervisor, customers can expect improved density while maintaining highly performant desktops and servers for end-users and applications.



Use Case

MSP / Enterprise

Customer that hydrates Amazon Gateway appliance on Hive Fabric (8 vCPU, 64 GB RAM, 500M Cache Disk) with NFS broadcasting.

Solution

Hyperconverged Fabric infrastructure can bring a quick and cost-effective solution to the event of a disaster. Our Zero-Layer Architecture with High Availability IT infrastructure allows clients to disk back cloud guest deltas and user volumes. In the event of an outage, customers have the option of hydrating to remote data centers or live areas of customer data center. Typically, it servers up to 10T of critical data with a RTO of 3 hours.

Why HiveIO

Complexity can be interesting. It can add richness; it can inject ambiguity and variability. It can lead to redundancy. It can be short-lived. **Simple is Powerful.**

At HiveIO we believe that the real journey should begin immediately, and that the promise can be realized. We do not want you to go from the data center team, to the networking team, to the storage team, to the virtualization team, to the Windows or Unix team. We want you to go to a single team that can see you through your journey.

We do not ask for maintenance agreements, support agreements, or multi-year renewals. We do not require complex hardware configurations to get you to the starting line.



Simple. Powerful. Ready. Realize the promise.