



FreeStor Disaster Recovery as a Service (DRaaS) in the Cloud White Paper

FREESTOR DRaaS IN THE CLOUD

White Paper

CONTENTS

Abstract	2
Potential Challenges	3
Implementation	4
Disaster Recovery: FreeStor in the Cloud	6
Use Case: FreeStor with AWS	7
Use Case: FreeStor with Azure	7
Use Case: FreeStor with AWS & Azure	8
Use Case: FreeStor with Standby DR	9
Benefits	10
Summary	11
About FreeStor	12

ABSTRACT

The cloud is the new frontier for data management. Actually, the cloud is not so new; people have been talking about it for a long time. Many say they are doing it right now; some say they would like to do it, but they don't really understand how to do it.

The cloud is sort of like Uber; you only pay when you use it. Instead of buying a car that is expensive to buy and maintain and sits around when you don't need it, you just open your Uber app and order a car when you need to go somewhere. In this case, instead of designing, implementing and maintaining your own computing infrastructure, you can rent compute time and storage space in the cloud. And just as Uber comes in different sizes (i.e. Uber X or XL), you can specify the size and type of server to be created for you in the cloud. In addition to specifying the OS, amount of memory, and CPU count, you can also specify faster storage...or a bigger and faster ride – in the Uber case.

Why is there so much buzz about moving data management to the cloud? Well, the primary reason is the significant cost savings, but there are other benefits as well.

- Cost Savings - The cost savings are substantial. Because cloud services typically offer a "pay-as-you-go" model, it is a bargain when compared to the capital costs and overhead of maintaining traditional IT infrastructures. In some instances, cloud providers offer a more cost-efficient storage platform in terms of dollars per Gigabyte.
- Business Continuity Data Recovery (BCDR) - The cloud is ideal for a disaster recovery solution; Using cloud service providers for disaster recovery can significantly speed up your recovery process so that your organization has minimal disruption to business operations in the event of an outage.
- Agility - Working in the cloud allows organizations to access data from anywhere. Having access to data from everywhere allows companies to be more agile. An agile business can be more responsive to business opportunities, be more customer-centric, and can recover quickly from a disaster.
- Data tiering – The cloud can be used as an additional data tier for cold and seldom used data. As the volume of data grows exponentially in every data center, the need for a tiering solution at different cost points provides flexibility in achieving desired total cost of ownership (TCO) targets.
- Flexibility – While consuming anything as a service can eventually cost more over time, the savings related to capital and people often enables organizations to deploy and move with more flexibility to achieve faster time to revenue. This benefit typically outweighs the longer-term cost factors.

If you have read this far, you are probably interested in making the move to the cloud. Good for you. In a nutshell, it will work something like this:

- A FreeStor Storage Server (FSS), a managed storage server that provides storage virtualization and business continuity services for continuous availability to your data in a virtual or physical environment, will be set up at your local site. The FSS can be a physical device or a virtual device.
- Another FSS – an FSS Virtual Appliance (VA) – is deployed in a public cloud, such as Amazon Web Services or Microsoft Azure.
- The FreeStor Management Server (FMS) is a standalone machine that gathers and consolidates the information coming from your different storage servers into a scalable repository of services, users, and historical data. The FMS includes a web service which allows you to connect to each FSS server for management and monitoring purposes. It is designed to run on any HTML-5-capable browser, including mobile devices.

Refer to the use cases later in this document for more information.

Now that you understand the basic FreeStor cloud solution, we will discuss some potential challenges along with some simple solutions as well as some use case scenarios to help you get started.

This document is intended as a guideline for customers planning to deploy a FreeStor DRaaS cloud solution. It provides a brief overview of FreeStor technology, concepts, and simple disaster recovery cloud solution examples.

POTENTIAL CHALLENGES

One of the most challenging aspects of moving to a cloud deployment is transferring data from your local (on-premise) environment onto the cloud. In order for your data to go from your machine to Amazon EC2, it must be transferred through the Internet onto the cloud-based server.

Possible challenges:

- Moving data – The transfer of data to the cloud can be a headache as well as a costly operation. Storage performance (i.e. network latency & bandwidth constraints) can make it challenging to transfer the data in a reasonable amount of time. Also, WAN costs can be some of the biggest costs for a modern Data Center. WAN-optimized replication to and from the cloud is a critical must have.
- Data Synchronization – Just moving bits of data is not enough. Data must be moved in context with the original applications and data sources, and must be synchronized once it is replicated to the new destination. This is critical to ensuring data integrity, avoiding data corruption, and to ensure a successful application recovery in the event of a failover or fallback.
- Vendor lock-in – This means that once you are in the Cloud with a particular cloud service provider (vendor), you may not be able to easily move your data to another vendor if you are unhappy.
- Mobility and visibility – When all of your storage is located in your on-site data center, you can simply walk over and check on everything (i.e. capacity, usage, and throughput) whenever you need. Your data management is local. You can expand a virtual resource if necessary, mirror a device, and set up a tiering system for data seldom used. However, once you move to the cloud, you may have to rely on your cloud service provider to handle these tasks.
- Storage management and data compatibility – It is rare to have common capabilities and management across paradigms. More often than not, the data services and capabilities between on-premise environments and cloud environments don't match. This can present numerous challenges when it comes to data protection and quick recoveries. This disparity may require new or additional tools and software to be purchased, further increasing costs and complexities to implement and manage.
- Data protection – You will need to carefully plan your data protection and disaster recovery strategy as not all data protection plans are the same. Some cloud service providers charge for each snapshot and have limited replication capabilities. You will also need to consider your recovery strategy in relation to your downtime SLA requirements. Even cloud service providers offering 99.9% uptime may be down 40 minutes a month. Depending upon your business, this may or may not be acceptable.

It doesn't matter how fabulous the Cloud is; it means nothing if you can't get there. FreeStor offers you the bridge to cross the gap between on-premise and the Cloud.

IMPLEMENTATION

Once you have made the decision to move to the cloud, you must first identify why you want to move to the cloud and the outcome you wish to achieve. For example, are you looking for a Disaster Recovery (DR) plan to protect your local physical facility in the event of a catastrophic event? Are you looking to lower your backup costs? Or do you want to move to the cloud because you have limited resources (i.e. space, hardware, manpower)? Once you prioritize your objectives, you can determine your plan of action.

The next step is creating a cloud deployment plan, which will include server details along with what exactly you will be moving to the cloud, and for what purpose (i.e. DR). Once your plan is in place and you have selected a cloud service provider, it is time to move your data.

Two of the biggest cloud service providers are Amazon Web Services (AWS) and Microsoft Azure. Therefore, the use cases described in this document will primarily discuss these two cloud service providers.

ABOUT AMAZON WEB SERVICES (AWS)

Amazon Web Services is an Infrastructure-as-a-Service (IaaS); it is a virtually limitless data center. AWS uses Elastic Block Storage (EBS) for block level storage volumes for use with EC2 instances in the AWS Cloud.

For object storage, AWS uses Simple Storage Service (S3). Amazon S3 stores data as objects within "buckets", allowing users to store as many objects as they want within each bucket resource.

ABOUT MICROSOFT AZURE

Microsoft Azure is a cloud computing platform and infrastructure created by Microsoft for building, deploying, and managing applications and services.

Azure uses object storage, also called "Blob storage" to store unstructured data in the cloud as objects/blobs. Blob storage can store any type of text or binary data, such as a document, media file, or application installer.

Now that you have had a brief introduction to some cloud vendors, let's address the challenges mentioned in the previous section.

Moving data

As discussed earlier, one of the most challenging aspects of moving to a cloud deployment is transferring data from your local (on-premise) environment onto the cloud. While moving data is not a new concept, moving it across a WAN can be one of your most expensive costs for the entire solution. It is strongly recommended that your cloud solution take advantage of WAN optimization; it is critical that you move the data in context to the workload for DR as is ensuring proper data synchronization. If the data is not synchronized in context to workloads, corruption can occur and successful DR may not be possible.

To address both of these key success factors, FreeStor integrates with Amazon Web Services (AWS) as well as Microsoft Azure. The cloud solutions discussed here describe a DR solution using FreeStor Virtual Appliances with cloud-based services from Amazon AWS and Microsoft Azure. While Cloud deployments are not limited to Amazon or Azure, they are the cloud solutions being used here as relevant examples of how FreeStor can deliver cloud-based backup, DR and Cloud Hopping services. What is *Cloud Hopping* you ask? That is a great question, which leads us to the challenge of vendor lock-in, which will prevent you from cloud hopping.

Vendor lock-in

Cloud hopping is when a user wishes to change cloud providers. Cloud services, capabilities, and formats are often proprietary to that specific cloud vendor, exacerbating the problem of vendor lock-in. This lack of compatibility can make it difficult or impossible to move the data from one cloud vendor to another (due to cost or proprietary formats). However, by implementing a heterogeneous FreeStor Storage Server (FSS) as a VA, you can eliminate vendor lock-in; FSSVA ensures data can be replicated in a patented WAN optimized fashion while ensuring synchronization and continuity to eliminate risk. FreeStor can even do a non-disruptive test to ensure data is ready in context to the application and that failover/replication is successful.

Lack of mobility and visibility

When you have all of your storage in your on-site data center, you manage your data locally; it is always there and always in view. If you need to expand a resource, you do it. Once you move to the cloud, you will typically need to work with your cloud service provider (i.e. AWS) and manage your data via their interface.

With the FreeStor solution, the FMS provides you with real-time and historical information on capacity, bandwidth, IO, latency, etc. From the underlying storage resources all the way out to the client application or Operating System (OS), IT administrators now have an end-to-end view across a heterogeneous storage infrastructure from a single pane of glass.

In addition, FreeStor provides native mobile apps which make monitoring across your entire storage infrastructure a breeze from anywhere. The mobile app allows you to view and manage your data from anywhere. FreeStor mobile provides a very simple workflow and at-a-glance view for monitoring and reporting of all storage servers, giving you greater insight as to what is happening with your data.

Storage Management

With FreeStor, different types of environments are no problem. FreeStor Intelligent Abstraction decouples storage hardware, networks, and protocols to create a free flow of data and common services across heterogeneous storage, regardless of environment or location.

Only FreeStor offers analytics and insight across heterogeneous storage environments, allowing users to take action, both proactively and reactively, as needed. This approach allows you to better manage capacity, performance, and availability.

Through our centralized storage management interface, we simplify data management across the entire storage infrastructure.

Data Protection

FreeStor offers flexible DR solutions that automate complex recovery processes, ensuring all your data can be recovered before business operations are impacted.

DISASTER RECOVERY: FREESTOR IN THE CLOUD

FreeStor in the cloud provides data mobility, as well as business continuity. FreeStor can move any workload – Windows server, virtual machine, etc. - from an on-premise environment to a public cloud infrastructure.

While there are other technologies that move virtual machines to the Cloud, many have limitations, such as requiring the same hypervisor and cloud orchestration. FreeStor allows you to move both virtual and physical machines bi-directionally - without hypervisor or cloud barriers.

With FreeStor's Disaster Recovery as a Service (DRaaS), you can virtualize your data and keep a current copy of your data in the cloud – on whichever cloud service provider you choose. Your data is replicated and remains in the cloud until you need it. If disaster strikes, your replicated cloud environment is available for use.

DRaaS is a smart solution when you have a limited budget. In this case, instead of struggling for 50% of your production budget for disaster recovery (DR), you can eliminate the need to set up an expensive remote site by moving your DR site to the cloud. Since you typically only pay for what you use, if/when a disaster occurs; you can have a permanent data center in the cloud running at 100% capacity.

Your FreeStor solution will work something like this:

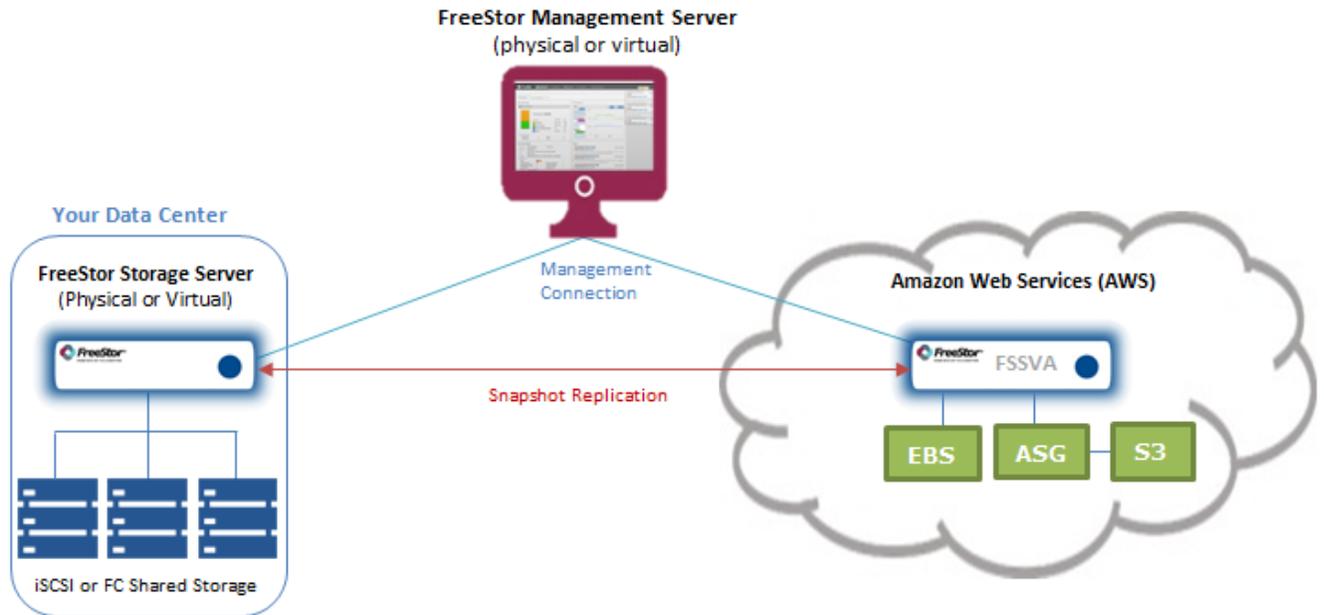
- A FreeStor Storage Server (FSS) sits at your local site. This can be a physical device or a virtual device.
- Another FSS – an FSSVA – is deployed in a public cloud, such as Amazon Web Services or Microsoft Azure.
- The FreeStor Management Server (FMS) allows you to monitor, report, and analyze your storage. The FMS includes a web service as well as a mobile application. You can install the FreeStor mobile application from the Apple® AppStore and Google® playStore.

Some basic FreeStor in the Cloud use case scenarios are described below.

USE CASE: FREESTOR WITH AWS

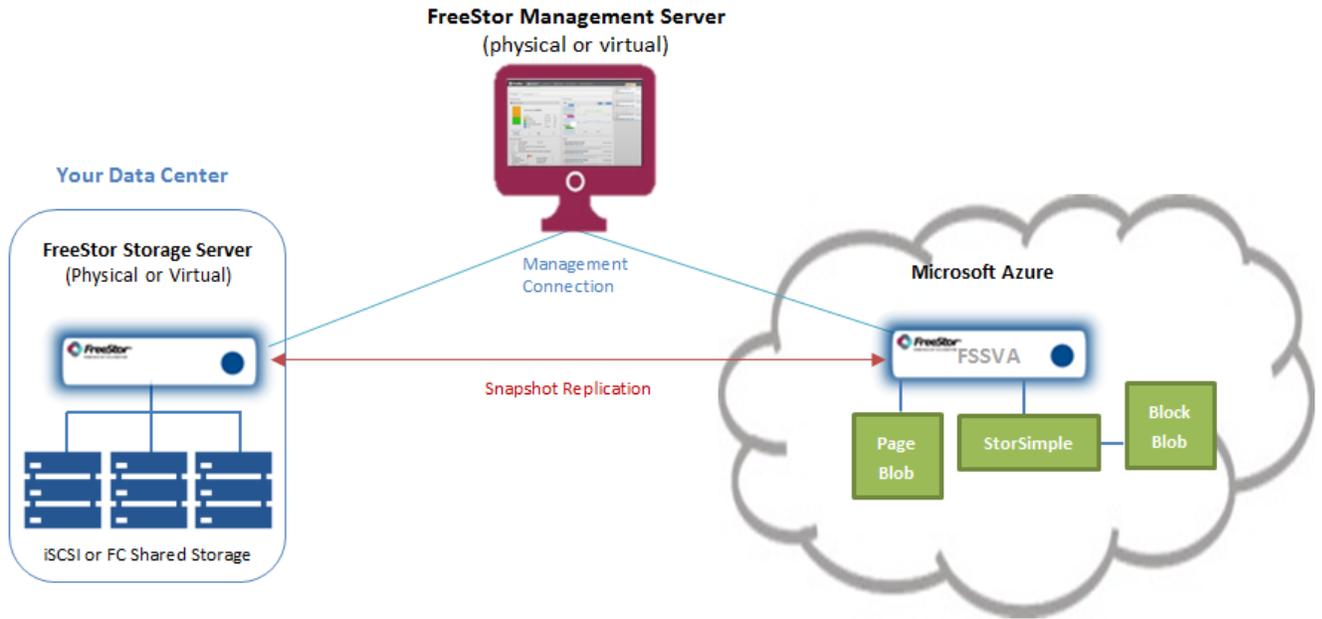
The diagram below illustrates a simple disaster recovery (DR) solution using FreeStor with AWS. In this case, the FSS server is replicating snapshots to the FSSVA server in the Amazon AWS cloud.

In the event that recovery is required, the user can spin up a new VA and instantly mount and boot the most recent snapshot in the cloud. Recovery time is in minutes, not hours or days.



USE CASE: FREESTOR WITH AZURE

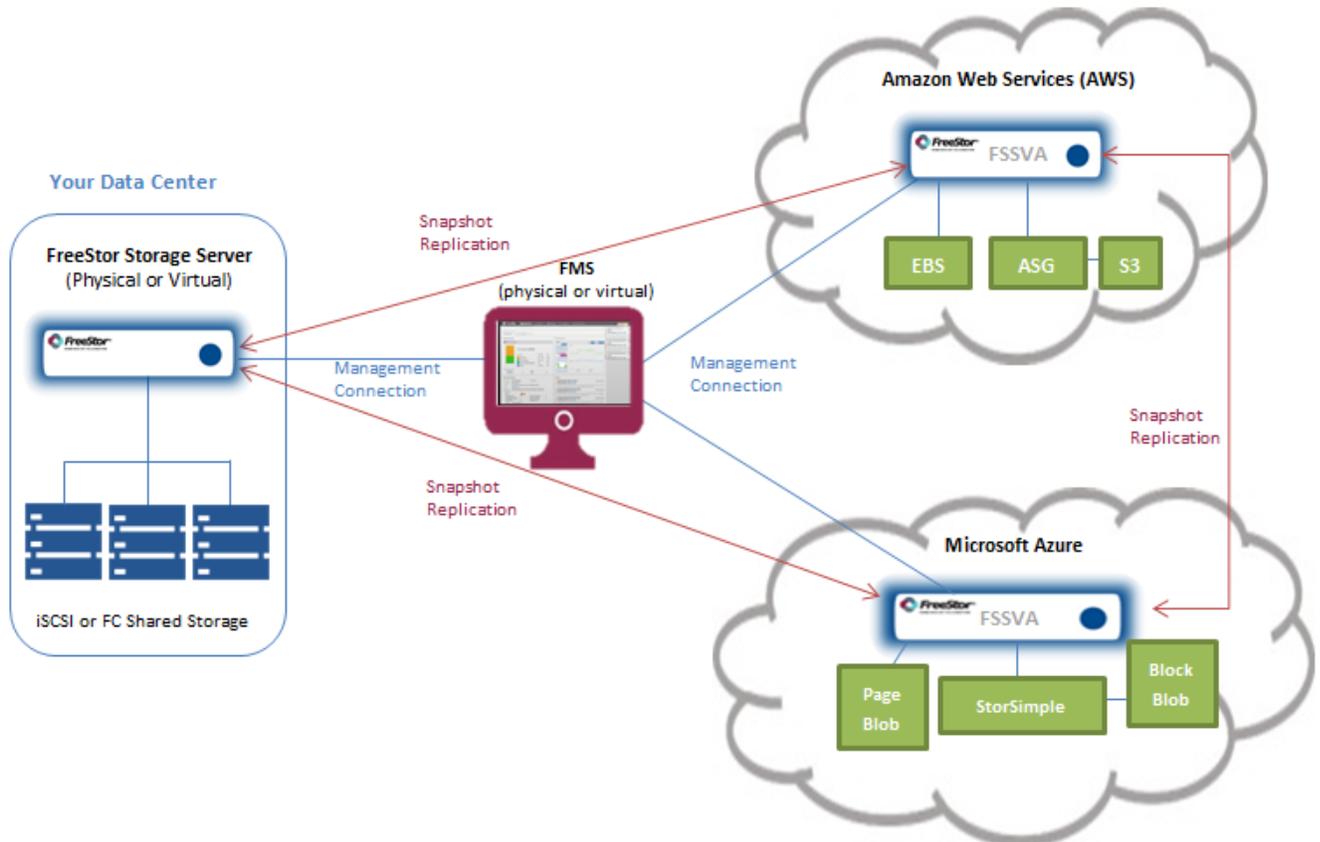
The diagram below illustrates a simple disaster recovery (DR) solution using FreeStor with the Microsoft Azure cloud. In this DR case, the FSS server is replicating snapshots to the FSSVA server in the Microsoft Azure cloud.



USE CASE: FREESTOR WITH AWS & AZURE

The diagram below illustrates a simple disaster recovery (DR) solution using FreeStor with both AWS and Microsoft Azure cloud services. This diagram shows a multi-site DR scenario across sites or clouds.

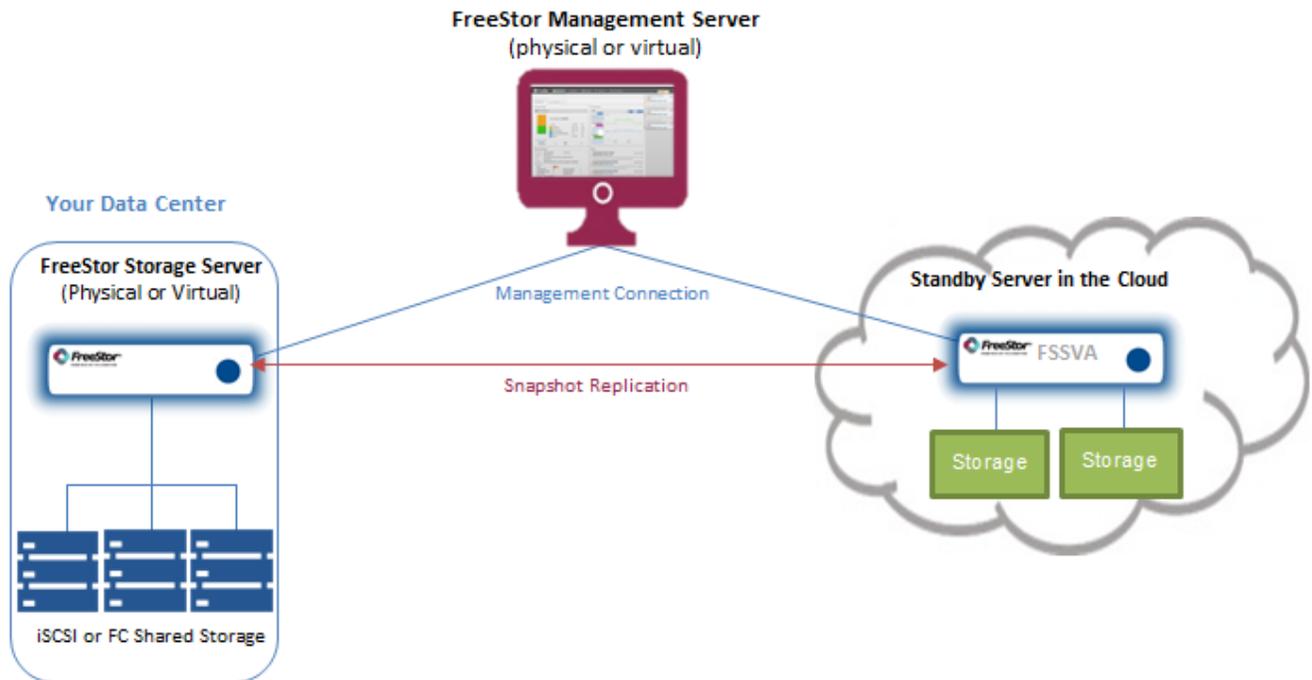
This case also illustrates “cloud hopping”, where you can change cloud vendors whenever you want. The ability to cloud hop is important for users not wanting to be locked in by proprietary cloud formats and services.



USE CASE: FREESTOR WITH STANDBY DR

The diagram below illustrates an on-demand disaster recovery (DR) scenario with a standby recovery server. In this case, your data is in the cloud being hosted by FSSVA. The data stays fresh and current because it is being replicated on a regular basis. If/when you need it; you can just spin it up, assign it and you are back in business. This standby DR scenario is a popular, fast, and cost-efficient method of on-demand disaster recovery.

On-demand disaster recovery is a common configuration because it is cost-efficient and offers quick recovery times. It is cost-efficient because you're not using it until you need it. And it is quick because it remains up-to-date and ready to be called upon and assigned to duty.



BENEFITS

- Moving to the cloud can free up resources, such as staff, hardware, and capital, making your business more agile and giving you more opportunity to deploy other revenue-generating projects.
- Data Services speak your language – whatever that language may be. Unified data services across heterogeneous storage – whether it is physical, virtual or hybrid – FreeStor in the Cloud eliminates vendor lock-in and platform incompatibilities.
- The FreeStor solution implementation process is fairly straightforward; you can move storage to the cloud with minimal risk and without interruption to the production environment.
- The FreeStor solution offers robust data protection and recovery capabilities for storage resources, ensuring that data can be recovered before business operations are impacted. Lower risk and improved uptime means you can be back in business in minutes, instead of hours or days, resulting in a higher level of business continuity and a peace of mind for IT administrators.
- FreeStor in the cloud utilizes its horizontal, heterogeneous software-defined storage platform to move and optimize your data. This enables you to eliminate silos and modernize your existing storage infrastructure – without expensive upgrades.
- Costs are lower because of the reduced number of software solutions needed and reduced complexity – fewer things to manage equals higher productivity and lower costs.

SUMMARY

FreeStor is not your ordinary data protection and disaster recovery cloud solution. It is a robust, yet flexible solution designed for organizations that are looking for a dependable business continuity disaster recovery solution while cutting costs, complexity, and recovery time.

FreeStor in the cloud gives you the power to seamlessly migrate, recover, and protect data in the cloud without tying your business to specific hardware, networks, or protocols.

To learn more about FreeStor or request a demo, visit www.falconstor.com.

ABOUT FREESTOR

FalconStor® FreeStor® is a horizontal, heterogeneous software-defined storage platform that helps IT organizations realize more economic value out of existing environments as well as future storage investments with maximum flexibility and operational efficiency. FreeStor customers have the power to seamlessly migrate, protect, and recover data – on or off the cloud – without being tied to specific hardware, networks, or protocols.

The core components of FreeStor are:

- FreeStor Management Server (FMS) - The FMS is a standalone machine that gathers and consolidates information coming from different storage servers into a scalable repository of FalconStor services, users, and historical data. The FMS includes a web service that provides management and monitoring.
- FreeStor Portal - The web-based portal provides centralized management and monitoring of multiple FSS servers via a web browser or mobile device.
- FreeStor Storage Server (FSS) - Each FSS is a managed storage server that provides storage virtualization and business continuity services for continuous availability to business data in virtual and physical environments.

FreeStor – software-defined storage with no surprises

Today's challenges require a new, software-defined approach that eliminates vendor lock-in, proprietary platform silos, increased complexity, and lack of hardware and software compatibility. Only FreeStor addresses those challenges head on and delivers real value to help organizations reduce costs, eliminate silo's while providing real flexibility and freedom.

- Intelligent Abstraction – Virtualizes existing AND future storage for a more efficient and available data pool
- Predictive Analytics – Provides real-time and historical analytics across the heterogeneous storage pool to better manage capacity, performance, and availability
- Intelligent Action – Enables users to take action based on real-time information to optimize and maintain their storage environment from a single-pane-of-glass regardless of storage vendor or location

For more information, visit www.falconstor.com or contact your local FalconStor representative.

Corporate Headquarters

2 Huntington Quadrangle, Suite 2S01
Melville, NY 11747
Tel: +1.631.777.5188
salesinfo@falconstor.com

Europe Headquarters

Landsberger Str. 312
80687 Munich, Germany
Tel: +49. (0) 89.41615321.10
salesemea@falconstor.com

Asia Headquarters

PICC Office Tower No. 2
Room 1901
2 Jian Guo Men Wai Avenue
Chaoyang District
Beijing 100022
China
Tel: +86.10.6530.9505
salesasia@falconstor.com

Information in this document is provided "AS IS" without warranty of any kind, and is subject to change without notice by FalconStor, which assumes no responsibility for any errors or claims herein. Copyright © 2016 FalconStor Software. All rights reserved. FalconStor Software, FalconStor, FreeStor, and IPStor are trademarks or registered trademarks of FalconStor Software, Inc. in the United States and other countries. All other brand and product names are trademarks or registered trademarks of their respective owners. 063016