

Amazon Web Services (AWS) cloud computing resources, networking facilities, and multi-user applications reduce the amount of financial and business resources spent maintaining businesses' IT infrastructure. Off-site storage in the cloud further protects organizations' data against damage to their facilities. This whitepaper will explain best practices for backup and disaster recovery on the AWS Cloud.

AWS has 33 centers in 12 geographic regions – with 5 more regions and 12 centers coming online over the next year – to ensure the security and availability of your data. The "pay-as-you-go" pricing model charges only for the resources used. The AWS Cloud is highly scalable, allowing you to increase workloads from photo storage to support video hosting at a glance, then scale down to a text repository. Almost any kind of machine or server can be virtualized, giving cloud systems an infinite potential.

This whitepaper will explain best practices for backup and disaster recovery on the AWS Cloud.

# **Backup Scenarios**

There are two parameters to evaluate the level of data protection in a network: Recovery Point Objective (RPO) and Recovery Time Objective (RTO). RPO shows the amount of data a company can afford to lose after an outage and consequent recovery. RTO defines the time between an outage and the full recovery of the infrastructure, which corresponds with the time taken for data from the latest backup to be retrieved.

Naturally, businesses want to get their information back from storage as quickly as possible and avoid the loss of any information. There are two ways to do so:

- Reduce RTO: for example, by retrieving data from storage faster.
- Reduce RPO: make backups more often.

However, each company faces technical limitations. For example, it takes about 24.5 hours to transfer a 10 TB copy through a 1 Gbps channel. Thus, the ability to reduce RTO is limited by network bandwidth. As for RPO, more frequent backups increase the network load, making upload and download more difficult.

# 3-2-1 Backup Strategy

The simplest backup strategy entails making a backup for network storage, keeping a few copies of backups to be consequently overwritten one by one. It is easy to retrieve these backups, and they are protected from software failure. A storage crash can make the simplest scenario dangerous. This dilemma can be resolved simply by utilizing a different storage device, in congruence with your existing hardware. For example, a tape can be simply lost or get corrupted, but any repository is prone to physical damage, therefore additional safety measures are required.





To ensure your data is safe and up to date, we recommend the **3-2-1 backup strategy:** 

- Have at least 3 copies of your data.
- Keep two copies of your data on two different types of media.
- Store one copy of your data offsite.

Three copies— the original data, and two clones — provide protection from human error, such as accidental deletion of data. Keeping data on at least two different kinds of storage devices makes it less likely for data to be lost due to a hardware fault. For example, a Hard Disc Drives (HDDs) can crash individually within a short period; this is most likely to occur if they were purchased together and installed at the same time. Offsite storage maintains data outside your city or country, keeping it safe in case of disaster that could destroy both hard disks and external storage.

Offsite storage is obviously extremely important; thus it should be:

- Reliable.
- Able to store any kind of data.
- Located as far as possible from your current location.
- Able to be accessed immediately.

AWS cloud facilities satisfy all these demands. Building your own data center is less secure and takes years of planning. Renting a rack in a commercial data center is less scalable and is not immune to disasters. Both options are much more expensive than the AWS Cloud.

Nowadays, enterprises need automated and customizable solutions – enabling terabytes of data to be backed up and restored faster. Let's see what backup capabilities exist and how you can utilize them.

# **File-Level Backup**

The standard copying of select files and folders is the easiest way to back up. It doesn't require much disk space, compared to other methods, and fits any kind of storage. File-level backup is mostly applied to working user data and documents. It's also possible to perform quick rollback to the previous version, whereas deduplication services prevent backup size overflow by uploading only new or changed data.

The best solution for keeping working data safe is file-level backup. To maintain the server state and be ready for disaster recovery, it's better to use another backup option.





### **Image-Based Level**

Image-based backup creates a copy of the operating system (OS), and the data associated with it, for its respective computer or virtual machine (VM). In case of failure, users can leverage the copies to retrieve data. All the machine's data from working files to system configurations are stored in a single file. This strategy requires more space on a storage system, but ensures availability of all server data.

CloudBerry directly connects to the cloud and uploads new images in real-time and analyzes images to identify the difference between previous images and only uploads modified data blocks so it doesn't use up additional disc space.

Folders and separate files can be easily restored from an image too. An image-based backup is the main tool for server and cloud migration, as well as disaster recovery. Alongside with simple restoration, images also allow businesses to:

- Restore as a virtual machine in the cloud (Amazon EC2).
- Restore with USB Flash directly from the cloud.
- Restore to dissimilar hardware.
- Restore to Hyper-V or VMware.

CloudBerry Backup empowers you to not only to create, restore, and transform images on the fly, but deploy VMs on AWS from your backups. An image-based backup is typically combined with the file level. Images are created to deal with system or hardware failures and disaster recovery, whereas file backups are good for daily routine losses and errors.

# SQL Database Backup

While SQL server can be backed up on the image level, the database itself is often the most valuable thing on the server. It is possible to only protect the database, with no extra storage expenses and efforts.

There are two main database backup strategies:

- 1 Full Backup: Save all data and logs commonly used for periodical service or initial data seeding to storage.
- 2 Differential Backup: Only update modified data blocks basic maintenance strategy.

The best strategy for data base backups is making a full backup as the initial seed, then updating it with differential backups as often as possible.





Additionally, CloudBerry Backup supports SQL server clusters and transaction log backups. AWS lets you deploy your database from your backups as a virtual machine on the cloud, or add it to your existing database. Amazon offers three database platforms:

#### Amazon Relational Database Service (RDS)

offers a wide choice of relational database engines, compute and storage options, Multi-AZ availability, etc. The main feature here is management simplicity.

#### **Amazon DynamoDB**

provides a fast, scalable, and cost-effective NoSQL database, where data is automatically replicated among data centers.

#### **Amazon RedShift**

is a tool for fast, scalable big data storage management. Its primary function is data warehouse maintenance.

### Microsoft Exchange Backup

There are two primary backup cases on Microsoft Exchange:

- Enterprise Database (EDB) files the database by itself
- Log files attached to EDBs.

These two items depend on each other, therefore, optimizing their backup requires caution and accuracy.

Block-level maintenance is extremely useful for databases, where backups can easily overflow storage. CloudBerry provides a special Purge feature that maximizes the benefits of block-level backup without risking the loss of data.

The best strategy for Exchange maintenance is to upload a full backup once a week, and update it by purged block-level backups without overloading the network and system while keeping everything up to date.

### **NAS Backup**

Network Attached Storage (NAS) can be used for backups or maintaining working user data. CloudBerry Backup can be installed on supported NASs to manage cloud data backup.

Another CloudBerry feature useful for network storage services is its consistency check. If the disk was removed from NAS or the backup was corrupted, a consistency check will help to solve all issues. To guarantee accuracy, CloudBerry uses timestamps to track changes made and compares them to see if there were any modifications.





## Mac & Linux Backup

Backup rules are the same for any operating system. CloudBerry and AWS enhance tools and cloud support for Mac and Linux. A variety of backup instruments, with both graphical and command line interfaces, give you the opportunity to match any kind of machines and storage types. AWS allows you to store and virtualize most Linux and OS X versions.

# EC2 Backup

Amazon Elastic Compute Cloud (EC2), as a VM, supports all backups. EC2 instances are deployed from preset software packs using Amazon Machine Images (AMIs) and simple images. You can focus on protecting configuration and stateful data using simple file or app-level backup types. This makes it possible to create backups more often, resulting in recovery with minimum data losses.

# **Big Backups Upload and Initial Seeding**

All backup strategies, beginning at the image-level, can scale. Data is easily transferred on the local level via fast-speed inner networks, but sending big data to offsite storage requires significantly more work. Transferring backups to cloud storage via the Internet can take time and incur additional costs. Amazon developed solutions to facilitate this transfer.

#### Amazon S3 Transfer Acceleration:

a new feature that is built-in to Amazon S3. When enabled, it speeds up data exchange with selected S3 buckets up to 6 times. Increased speed is realized by selecting Amazon transfer routes with a higher bandwidth rate, giving your data upload priority.

#### AWS Import/Export Disk:

a 16 TB hardware data transfer tool that lets you send data from your own device to an Amazon data center.

#### AWS Snowball:

allows up to 80 TB of data to be stored and transferred to the cloud. It is self-encrypted, armored,

Note: the delivery speed of hardware data devices is determined only by postal services performance.

CloudBerry Backup supports online and offline upload services. You can access acceleration tools and hardware storage delivery pages directly from the GUI. CloudBerry products track the entire lifecycle of Amazon Snowball transfers.

The best strategy for backup upload depends on the data size, urgency, and the transfer facilities available. For example, if a full enterprise backup comprises 15 TB of data and the Internet connection bandwidth is 100Mbps, it will take approximately 18 days to upload the initial seed with 80% network utilization. Amazon Snowball operates faster. It takes approximately two days for a Snowball device to be delivered to customers, upon which data can be uploaded instantly. Amazon S3 Transfer







Acceleration is best used for uploading weekly and monthly backups. Disk Export helps with backing up new elements introduced into the IT-structure.

When choosing the tool for your backup transfer, take the peculiarities of your region into account; including the quality of Internet connection and data upload destination. This is further explored, in detail, in the **Data Transfer** section of this document.

# **Storage Facilities**

Amazon Web Services offers different classes of storage for various usage scenarios. This allows organizations to reduce storage costs for backups which are not accessed often. Classes have a high level of reliability and support SSL data encryption during transmission, but differ in cost. AWS offers the following for data maintenance:

## Amazon S3 Standard

Amazon S3 Standard is designed for high-usage and has the following features:

- High capacity and low latency.
- 99.99999999% reliability (risk losing one object for every one-hundred billion).
- 99.99% availability (one hour of unavailability for every ten thousand hours).
- Use of storage is covered by Amazon S3 Service Level Agreement, which considers compensation if the level of uninterrupted operations is lower than it was declared.

Standard storage is suitable for file-level backup of working files and documents. These may be rolled back, changed, and recovered dozens of times per day. S3 is also the first place where Snowball transferred data is uploaded. It is common practice to use S3 intermediate storage for image-level and database backups.

# Amazon S3 RRS

**Amazon S3 Reduced Redundancy Storage (RRS)** reduces storage costs for replicable, and non-critical data. Amazon RRS is intended to sustain the loss of data for single facilities. This can be achieved by reducing the amount of data replicated across multipole devices and facilities. The main difference between RRS and S3 Standard is reliability (99.99%).

This solution is perfect for non-critical, or easily replicable, data of applications. It doesn't suit maintenance for crucial data, though it can be used as a data buffer if you are dealing with large backups being transferred to multiple storage systems.





### Amazon S3 Standard Infrequent Access

Amazon Standard Infrequent Access (S3-IA) is designed for data which require less frequent access than Standard class. Low delays combined with high capacity and reliability (99.9999999999) ensure the safety of objects for a long periods of time. Amazon S3-IA differs from Standard in the following ways:

- 99.9% availability (e.g., slightly greater chance of a request error, compared to standard storage).
- Charges for data retrieval.

The minimum storage period is 30 days, and the minimum size of an object is 128 KB. This tier is recommended for long-term storage of user files, disaster recovery data and backups. In CloudBerry Backup, S3 Standard IA class can be attached.

## Amazon Glacier

Amazon Glacier is for the long-term storage and archiving of backups that don't require instant access. The service allows storing large volumes of data at a low price. Amazon Glacier differs from S3 Standard in the following ways:

- Extremely low cost.
- Uninterrupted operation is not guaranteed by Amazon S3 Service Level Agreement.
- Minimum period of storage is 90 days.
- There is a charge for data retrieval of more than 5% from the average monthly volume. You can access your data in four hours after the first request.

The service is optimized for infrequently accessed data, with a retrieval time of several hours. It's beneficial for storing items such as old backups and outdated database records.

AWS does not save objects directly in Glacier. S3 archives data in accordance with the Lifecycle Policy. CloudBerry products can manage this policy and circumnavigate S3 as intermediate storage by transferring files directly to Glacier.

# Lifecycle Policy

All Amazon S3 classes are supported by the lifecycle policy, meaning you can optimize storage costs for objects by setting rules for automatic transfer to cheaper storage. It is also possible to set up the lifecycle termination policy so that the files are automatically removed after a certain period. This is useful for recovery data maintenance, as hot backups will always be available and old backups will go to the archive, resulting in reduced expenses.

For example, you can save a backup using Amazon S3 Standard, transfer it to Standard IA storage, and finally to Glacier. Later the backup can be removed, or placed into archive storage.





# Recovery

CloudBerry Backup can make the recovery process easier with built-in consistency verification and by setting up multiple backup schedules. Make sure that the lifecycle policies on your local and cloud storage are correctly configured.

# **Image-Level Recovery**

If your storage is accessible, the process of full system restoration from an image can be initiated by a couple of clicks via the CloudBerry Backup GUI. Image recovery is a difficult process. To avoid the loss of data, it is important to follow best practices to fully execute a successful recovery. Take all precautions to avoid losing data that has been added or changed since the last image update. In general, these are settings and working files; the most important things in production. Here are a few actions to take before image-level recovery:

- Make a fresh backup of working data, system and application settings.
- Make sure that the databases are maintained separately, and their last backup is up to date.
- Carry out file-level recovery after image restoration to bring the machine into the state of readiness.
- Be aware of your application recovery peculiarities for example, Microsoft Exchange and Active Directory restoration may require to carry out additional adjustments after being unpacked from the image.

How do you go about an image-level restoration in case of disaster recovery? The steps are nearly identical, but you may not have a chance to make a differential file backup. A well-planned maintenance schedule and consequent image and file-level backups are key to IT-infrastructure safety. Data transfers should be also taken into consideration. Downloading images from 20 desktops and 3 servers can take a while. You must be ready for extensive downtime while backups are transferred, or additional costs for quick data transfer. AWS is faster and more convenient than offsite tape or disk storage services, as there are additional measures that can be taken to simplify recovery.

# **Cloud to Virtual Machine Recovery**

Amazon Web Services enables you to decrease your systems' downtime with virtual machines located in the same cloud as your backups. Recovery to Amazon EC2 can be made directly within CloudBerry Backup's Recovery Wizard. Here you can configure the virtual machine's type, connect it to the subnet, and adjust all other settings. The steps to achieving a full recovery is the same as that of an image-level recovery: start with the latest full image, then deploy the latest version of working data.





While deploying Amazon EC2 instances from CloudBerry, you can choose to create AMIs. This option can be initiated from the AWS console anytime and helps start your machine with a personalized IP-address right after finishing Recovery Wizard. Otherwise, IP-addresses should be configured within Amazon Elastic IP service. Newly created instances should be launched from Amazon EC2 Management Console.

# **Pilot Light Recovery**

The term *pilot light* is often used to describe a disaster recovery scenario in which a version of an environment is always running in a cloud. The term "pilot light" comes from a gas heater system that always has a small flame burning, as an ignition source for when the entire heater is turned on and more gas comes in. This principle can be used in recovery by deploying an AWS virtual machine which runs the most important core element of your IT-structure. If your local structure fails, you can deploy all other elements, such as databases and file storage systems, around the pilot light core. To provision the rest of your structure, you should preconfigure the servers or other multi-user networking machines and save them as AMIs, which are ready to deploy on demand. When starting recovery, EC2 instances come up quickly from these AMIs with their predefined role (for example, Web or App servers). CloudBerry Backup converts your backups into operable virtual machine images, facilitating pilot light strategy preparation.

The auto-scaling feature of AWS ensures the virtually recovered infrastructure has enough computing resources to cope with the existing workloads, but you still have to restore working data to get the system to operating condition.

There are a few main recovery scenarios, which may be used at your discretion. Regardless of the strategy, ensure that your backups are fully functional, and recovery goes well. This can be done by doing:

- Regular backup consistency checks to make sure that files are not corrupted.
- Regular recovery tests to ensure that restored systems and applications are functional.
- Regularly create additional copies to have a plan B if something goes wrong (remember 3-2-1 rule).

# Large Recovery Download

Nowadays businesses generates huge amounts of data, making recovery to new hardware and local storage systems more complex. While ordinary downloads from the Internet are possible, alternative tools are very useful with large recoveries.

For recovery, download tools that are also used for uploading backups. Consider the peculiarities of your region, the quality of Internet connection, and data upload destination.





# Data Transfer

Corporate information is distinguished in terms of urgency. A company may need a set of assets on a daily, weekly or yearly basis. Therefore, backup and recovery operations can meet one of the following scenarios:

- Accessing archived information that is infrequently accessed
- Immediate recovery for information that is crucial to business operations, where backup occurs daily.

The second scenario is common for big companies. The bigger a business's infrastructure, the more prone it is to data loss.

A standard AWS storage scheme copies user data to a geographically distant AWS server. Distance increases latency so that creating several big backups are more challenging. Amazon has developed two features to avoid reduce time spent on backup: S3 Transfer Accelerator and AWS Snowball.

# **Amazon S3 Acceleration**

Amazon S3 Transfer Acceleration is a new built-in feature implemented in Amazon S3. When enabled for a bucket, it speeds up data exchange with this bucket up to 6 times. Users can enable data transfer acceleration for a single S3 bucket, enabling a special Accelerated mode. After that, an alternative name is generated for this bucket. Specifying this name in the backup plan enables faster uploading and downloading.

Amazon uses its own networking resources to ensure accelerated data transfer. For each session, an optimal route based on AWS own backbone network is selected. This route is built dynamically: channels with bigger bandwidth and less traffic load are found for each separate upload/download session.

Speeding up data transfer does not result in security level reduction. All data uploaded in the accelerated mode, as well as standard mode, are encrypted. Data is not stored in transit AWS edge locations, eliminating the risk of their leakage.

# **AWS Snowball**

AWS Import/Export Snowball is a hardware solution which transports up to 80 TB of data directly from the user's location to AWS network. Snowball devices are shipped to AWS customers on demand. A request is sent from the standard AWS console and takes approximately two days to ship. Upon receiving an Amazon Snowball, the user must connect it to their network, install AWS S3 client on it, transfer the file or system image copies, then ship it back to AWS. The data transferred to a Snowball device are automatically encrypted. After the device is shipped back, the user receives notifications on the progress of data extraction and transfer to the user's S3 buckets. In the USA, the overall cost of creating a large backup using Snowball will be several times lower than via high-speed Internet connection.







Total cost of retrieving 50 TB of data: comparison between S3 Acceleration and S3 Snowball			
	S3 Acceleration	S3 Snowball (80TB)	
Cost	\$0.04 (acc. fee) + 0.07 (std. fee)) * 50 * 1024GB (data retrieved) = \$5,632	\$250 (device) + \$0.03 (fee) * 50 * 1024 GB (data retrieved) = \$1,786	
Time	5 days	2 days	
Additional Fees	None	Shipping fees	
Other issues	Depends on your network	Depends on your physical location	

Physically transporting large volumes of data this way helps to save networking resources. Imagine a situation where the user needs to make a 25 TB backup to an AWS S3 bucket but the network can supply only 4 Gbps. It will take about 25 hours to complete, using the full bandwidth. In fact, the actual time of backup creation will be longer as other tasks generate additional loads in the network. The retrieval time will also vary depending on the type of data retrieved and changes in traffic load. A Snowball device will carry all necessary data itself, making the user's network fully available for other tasks.

CloudBerry and Amazon collaborate closely to provide the best experience for customers leveraging AWS storage. CloudBerry Backup and CloudBerry Explorer support Amazon S3 Acceleration and Amazon Snowball features. All options for your S3 bucket can be accessed directly from CloudBerry user interface.

# **Billing Calculations**

Now, when you are aware of cloud-based backup and recovery, it's time to estimate its cost. Let us imagine a common company office with the following structure:

- 20 desktop computers with Mac and Windows aboard.
- 1 Microsoft SQL Database server.
  - 1 Microsoft Exchange server.

The total amount of enterprise data is 10 TB; 70% of which is hot data (documents, working files, backups of crucial importance etc.), database archives, and old system images. The business is growing quickly, with its data volume increasing 20% up per year. The company chose the Amazon Northern Virginia region for its offsite storage. How might it go about backing up this infrastructure, and how much does it cost per year?





### Software Licensing Cost

For the office depicted above, the cost of CloudBerry Backup (CBB) would be the following:

• 20 CBB Desktop Edition copies will cost \$21.99 per computer: \$439.80 total, regardless of Windows and Mac versions ratio.

*Note: The basic price is \$29.99 per copy. The wholesale discount for 20 computers is 27%.* 

- CBB for MS SQL costs \$149.99 per machine.
- CBB for MS Exchange costs \$229.99 per machine.
- CBB Server Edition for the file server costs \$119.99 per machine.

*Note: With server CBB tools you pay only for the required features. Universal CBB Ultimate Edition costs \$299.99 per copy.* 

• You might also want to have CloudBerry Explorer on the administrator desktop, which allows you to manage storage and data transfer. It costs \$39.99 in Pro Edition.

The total price for CloudBerry licensing is \$939.77 (without Cloud Explorer). A license is infinite per machine, and you don't need to relicense recovered machines. If you have moved to different hardware or want to use CloudBerry with recovered servers, the license can be easily moved from the old infrastructure to the new one.

An additional service available is CloudBerry Maintenance, which includes annual technical support and update subscription. The yearly fee for the configuration will be \$219.95.

# Storage Cost

For a 10 TB stack, comprised of 70% hot data and 30% cold data, the billing would be the following:

- 7 TB of frequent access data on Amazon S3 Standard costs \$2,541.84 per year.
- 3 TB of old archive data on Amazon S3 Standard Infrequent Access with 1TB per month retrieval possibility costs \$583.68 per year.
- Since the data grows, 3 TB of archives will be superseded to Amazon Glacier Storage, costing \$258.12 per year.

#### Note: AWS technical support is provided for free for all kinds of storage systems.

The annual fee for AWS will be \$3,383.64. This total cost would be much lower if there was less hot data.





## **Initial Seed Transfer Cost**

Let's estimate the fee for sending a full enterprise data block to a cloud:

- Using the standard upload method, it will be free to send data, though time-consuming with 100 Mbps it takes 233 hours to transfer 10 TB of data.
- With Amazon S3 Transfer Acceleration it will cost \$400 (\$0.04 per GB) when powered by the US, Europe and Japan Edge locations, or \$800 (\$0.08 per GB) if the bandwidth is accelerated by other locations.
- Transfer via AWS Snowball is priced in a couple of stages. First, you pay a service fee for Snowball job, which will be \$200 for a 50 TB device. Then shipping expenses are paid, and after that a Snowball device can be located for 10 days free onsite. For every day of expiration, a \$15 daily fee is charged. Uploading data to an Amazon data center is free.
- Sending data with AWS Import/Export Disk, you pay \$80 for every device handled. Alongside with shipping, data loading time is paid \$2.49 per hour of uploading, partial hours are billed as full ones. With a fast 500 MB/sec drive, data upload will cost \$14.94 and \$94.49 with a device handling in the upshot.

# **Disaster Recovery Cost**

The first thing to do is to initialize recovery data download. The charge for it is as follows:

- With standard Internet transfer, you pay \$921.52 to get 10 TB data from S3 Standard Storage, and \$30.72 for access to 3 TB of S3 Infrequent Access data, \$952.24 in total.
- Using Amazon S3 Transfer Acceleration, you will be charged \$400 for getting all data out of the cloud. Together with the standard transfer fee, it will cost \$1321.52.
- AWS Snowball will cost \$200 per device and \$300 for downloading 10 TB of data from the cloud, \$500 in total. Additional \$15 will be charged every day upon exceeding a 10-days limit of onsite usage.
- AWS Disk expenses are the same as with backup upload \$80 for device handling, \$14.94 for upload, \$94.49 in total. Keep in mind that the price of the disk is not included.

# **Backup to Virtual Machine Recovery**

Whatever transferring tool you choose, it will take some time to download all data. It is a good idea to recover your infrastructure to EC2 virtual machines. The price for the demo office for 7 days is as follows:

- Fileserver on EC2 t2.medium instance powered by Linux with 2 CPUs and 4GB RAM costs \$8.88.
- MS Exchange server on EC2 m4.large instance with 2 CPUs and 8GB RAM costs \$42.01.
- Server with Windows and Web SQL Database aboard on EC2 m4.large instance costs \$40.99.

#### Note: software licensing price is included





Setting up Elastic IP and attaching Amazon EBS (Elastic Block Storage) volumes will require an additional fee, which depends strongly on your system networking demands. Find out more on Amazon pricing page.

To weigh all pros and cons, we have created a table with major figures in it:

	Cost	Term		
SOFTWARE				
CloudBerry Lab Licensing	\$939.77	Forever		
CloudBerry Lab Maintenance	\$219.95	1 year		
STORAGE ON AWS				
7 TB of "hot" S3 Standard data	\$2,541.84	1 year		
3 TB of "cool" S3 Standard IA data	\$583.68	1 year		
3 TB of "cold" Glacier data	\$258.12	1 year		
<b>RECOVERY TO A VIRTUAL MACHINE</b>				
3 servers (File Server, SQL, Exchange)	\$91.88*	7 days		
BACKUP UPLOAD COST				
Standard Internet Upload	Free			
S3 Transfer Acceleration	\$400 (\$800 outside the USA)			
AWS Snowball	\$200**			
AWS Import/Export Disk	\$94.49***			
RECOVERY DATA DOWNLOAD COST				
Standard Internet Upload	\$952.24			
S3 Transfer Acceleration	\$1321.52			

AWS Import/Export Disk

\*EBS and networking cost is not included.

\*\*Delivery cost is not included

AWS Snowball

\*\*\*Delivery and device cost is not included.





\$500\*\*

\$94.49\*\*\*

# Conclusion

Ideal backup strategy hinges on planning. Files should be maintained more often than full system images, and configurations should be maintained even more frequently to be ready for cloud recovery because every minute of downtime means money lost. Moreover, data transfer possibilities ought to be explored. Network transfer features must be tested so disk drives are prepared for data import and export.

Amazon Web Services provides a universal platform for backup storage, virtual machine deployment and data transfer. High levels of customization and a large array of tools optimize your IT operations on the cloud. CloudBerry Lab's solutions empower you to master your infrastructure.

### About AWS

For 10 years, Amazon Web Services has been the world's most comprehensive and broadly adopted cloud platform. AWS offers over 90 fully featured services for compute, storage, databases, analytics, mobile, Internet of Things (IoT) and enterprise applications from 42 Availability Zones (AZs) across 16 geographic regions in the U.S., Australia, Brazil, Canada, China, Germany, India, Ireland, Japan, Korea, Singapore, and the UK. AWS services are trusted by millions of active customers around the world monthly -- including the fastest growing startups, largest enterprises, and leading government agencies -- to power their infrastructure, make them more agile, and lower costs.

To learn more about AWS, visit aws.amazon.com

### About CloudBerry Lab

Established in 2008 by a group of experienced IT professionals, CloudBerry Lab<sup>™</sup> provides cloud-based backup and file management services to small and mid-sized businesses (SMBs). CloudBerry's offerings include powerful, easy-to-use backup management capabilities and military-grade encryption using customer-controlled keys. Customers can choose to store their backup data with more than 20 online storage providers, including Amazon S3 and Amazon Glacier. CloudBerry also collaborates with thousands of VARs and MSPs to provide them with turnkey, white-la bel data protection services. It has been an Amazon Web Services Advanced Technology Partner since 2012. CloudBerry Lab has also achieved Storage Competency Partner status in the AWS Partner Network.

For more information, visit www.cloudberrylab.com. Follow us on Twitter at @cloudberrylab.

© 2017 CloudBerry Lab, Inc. All rights reserved. CloudBerry Lab and the "Cloud" logo, are trademarks or registered trademarks of CloudBerry Lab, Inc. All other third party brands, products, service names, trademarks, or registered service marks are the property of and used to identify the products or services of their respective owners. All specifications are subject to change without notice. © 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.



