

Minimizing Risk and Lowering Cost with Cloud Data Migrations

EMC Corporation and DataTrust Solutions Vapor

ABSTRACT

This white paper explains how EMC® Elastic Cloud Storage (ECS™) Appliance and DataTrust Solutions Vapor cloud backup/migration software enable companies and organizations to back up cloud-based data or move from one public cloud to another or from a public cloud to a private cloud.

August, 2015

REDEFINE
EMC WHITE

EMC²

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller, visit www.emc.com, or explore and compare products in the [EMC Store](#)

Copyright © 2015 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided "as is." EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

Part Number H14284.1

TABLE OF CONTENTS

EXECUTIVE SUMMARY	4
AUDIENCE.....	4
CUSTOMER CHALLENGE	4
Reliability	4
Cost	4
Security and Compliance.....	4
Bringing the Data Back Home	4
EMC SOFTWARE-DEFINED STORAGE	5
EMC ECS Appliance	5
INCREASE SPEED TO MARKET	5
DATATRUST SOLUTIONS VAPOR	5
Simple Web Administration	6
Scheduling	6
Robust Reporting System with Customer Reports.....	6
Cloud to ECS Appliance Backup.....	7
Wide Range of Cloud Backup and Recovery Options	7
Flexible Data Migration	7
Migration Chain of Custody.....	7
Utilities for Updating the Source.....	7
CONCLUSION	7

EXECUTIVE SUMMARY

Public clouds have offered a number of compelling benefits including low initial cost, simple APIs, fast deployment, and ability to bypass bureaucratic IT processes.

One potentially negative outcome of the proliferation of public cloud use is the movement of sensitive data outside the corporate firewall. The potential for security issues has made many organizations uncomfortable.

As a consequence, many IT organizations are implementing private clouds behind the corporate firewall. The private cloud model provides cloud simplicity and flexibility, but the data is securely on premises inside the corporate datacenter. This enables IT to regain administrative control over sensitive data, service levels, and storage infrastructure.

To make this a reality requires moving cloud object data from one cloud to another, whether public or private. This is not a trivial task, particularly when some of the data is under compliance or governance requirements.

EMC ECS Appliance supports private cloud implementations and works with software from DataTrust Solutions to move the cloud data.

This white paper will provide an overview of the cloud data migration solution from DataTrust Solutions and EMC Corporation.

AUDIENCE

This whitepaper is intended for systems architects, corporate IT architects, storage engineers, and CIOs. It focuses on backup/migration software that allows organizations to back up cloud-based data and move content from one public cloud to another or from a public cloud to a private cloud.

CUSTOMER CHALLENGE

Cloud storage is becoming a mature commodity. There are a substantial number of funded cloud vendors in the market and consolidation has already started to occur. And as it does, companies and organizations that were early adopters of public cloud storage are reevaluating their usage based on factors including reliability, cost, and security and compliance.

Reliability

Public cloud storage is subject to outages. The inability to access data for periods of time has called into question the notion that public cloud is really enterprise class. Also user loss of access to their data during an outage results in loss of productivity, which can be expensive.

Cost

The instant-on and then pay-as-you consume model is appealing at first. However, many companies and organizations are realizing that cloud storage at some volume becomes more expensive than physical implementations or private cloud. It is not expensive to put data in the public cloud but it can be to change it or get it out. In some cases, less than a year of cloud fees could pay for a whole new hardware system.

Security and Compliance

At some companies, the user community has adopted cloud storage behind the back of IT. Initially, someone starts to use it because it is faster and easier to get storage from a cloud provider than it is to fill out forms and wait for the IT department to provision it.

Ironically, it can be people inside IT involved in this shadow activity. A developer, for instance, might want to do some tests and has a choice: provision storage via a public cloud provider in five minutes or wait two or three weeks for internal storage to be provisioned. Once someone knowingly gets away with it, then others will view this as a low-risk option and follow suit. This is a grave danger in industries where data is subject to compliance, retention, and governance.

Bringing the Data Back Home

Senior IT managers are shocked at how shadow IT has taken on a life of its own. They want to get control of their data again—or at least get a copy of it. Perhaps they have decided to go to an on-premises private cloud or another public cloud—but this time on their terms. This decision is important. EMC and DataTrust Solutions can help companies and organizations make it the right one.

EMC Software-Defined Storage and DataTrust Solutions Vapor backup/migration software enable you to back up cloud-based data or move from one public cloud to another or from a public cloud to a private cloud. You can rein in shadow IT and make smart moves to lower cost, improve service levels, and strengthen regulatory compliance.

EMC SOFTWARE-DEFINED STORAGE

EMC Software-Defined Storage helps organizations lower traditional storage costs by abstracting hardware and software resources to allow compute and storage resources to scale independently. This paves the way for rapid deployment of modern scale-out storage architectures.

EMC ECS Appliance

A turnkey, software-defined object storage cloud platform, ECS Appliance integrates powerful software on EMC-certified commodity hardware. It combines the cost advantages of commodity infrastructure with the reliability, availability, and serviceability of traditional arrays. The ECS Appliance uses multiple data types and access protocols to support the broadest range of applications on a single platform.

ECS Appliance also provides a single management view of a globally distributed infrastructure. What's more, it supports multi-tenancy and detailed metering so enterprises and service providers can easily deploy storage-as-a-service.

Other features include:

- Complete, turnkey commodity hardware platform provides up to 65% lower TCO than public cloud storage services
- Assures the broadest application support—object and HDFS—from a single platform
- Patent-pending object storage engine delivers an unmatched combination of storage efficiency and global access to data
- Terabyte-to-exabyte architecture delivers flexibility and control to scale-out as dictated by business need
- Commodity hardware deployment with enterprise-level reliability, availability, and serviceability
- Cloud storage in-a-box enables seamless management, multi-tenancy, metering, and self-service access

INCREASE SPEED TO MARKET

ECS Appliance provides a variety of REST-based API protocols to create and manipulate object data storage. DataTrust Solutions connects to ECS Appliance using the S3 API, a commonly understood industry-standard protocol that allows developers to create, update, and delete objects. By using the S3 API, DataTrust Solutions replaces the complexity of proprietary protocols.

As EMC's next generation Object Storage Platform, ECS Appliance supports multiple object storage protocols including Swift, S3, Atmos REST, and Centera CAS allowing customers to quickly adopt the new storage platform whether for legacy applications or for the most modern new object applications. DataTrust Solutions complements the flexibility of the ECS Appliance platform by supporting multiple types of migrations to ECS Appliance including NAS, CAS, Atmos, and S3 speeding the time to deployment for any application that is selected for migration.

Now let's turn our focus on what DataTrust Solutions Vapor software can do on top of the capabilities of ECS Appliance.

DATATRUST SOLUTIONS VAPOR

DataTrust Solutions Vapor software uses the ECS Appliance platform to help companies and organizations move cloud object data within public clouds, to private clouds, to object systems, and to NAS while meeting compliance requirements. Vapor with ECS can function as a cloud backup solution or as a cloud storage replacement solution. Table 1 provides an overview of the software characteristics.

Table 1

DataTrust Solutions Vapor Software Overview	
Public clouds supported	Amazon S3 OpenStack More coming soon
Private clouds supported	EMC ECS Appliance EMC Atmos Hitachi Cloud
Object storage systems supported	EMC Centera

	EMC Atmos Hitachi HCP
Protocols	CIFS NFS API REST
Systems	EMC Isilon (Smart Lock) EMC VNX (FLR) EMC Data Domain (FLR) EMC Atmos EMC ECS Appliance
Cloud data protection	Transparent to applications and users Out-of-band copy Many target options
Cloud backup and restore options	Cloud to NAS Cloud to Object Cloud to Cloud

Simple Web Administration

DataTrust Solutions Vapor application is accessed and administered via a standard web browser. The browser uses a REST interface to connect to the core application, which was written in JAVA and designed to manage billions of objects. Vapor runs in a 64-bit Linux operating environment in a VM or on a physical server, depending upon the desired performance.

Scheduling

You can copy or restore a single file, a list of files, a date range, or everything. The scheduler can be used to migrate, back up, or restore data during specified hours. It can also be used to selectively throttle the bandwidth during peak operation times. Some clouds charge less if you recall data during off-peak hours. If you are going to copy a lot of data out of cloud, then do it during the off-peak hours, as it is the least costly time. DataTrust Vapor allows you to control the migration or backup schedule.

When using Vapor to back up up your existing cloud, it has labor saving low-touch set-and-forget policies to automate repetitive tasks.

Schedule Periodic or near continuous backups or restores:

Immediate

- Schedule start time
- Schedule a repetitive start time
- Schedule an end time
- Run continuous

Backup Options

- Full
- Incremental

Robust Reporting System with Customer Reports

DataTrust Vapor has a customizable report engine and supports a range of standard and custom reports:

- Point-in-time status
- Source Cloud Inventory
- Migrated or backed-up queue
- Completed migrated and backup file list
- Restored files
- Log files
- Chain-of-custody report with file level verification

All actions in Vapor are logged and tracked in the integrated database. Creating custom reports is as simple as writing a SQL statement. All the information gathered such as file metadata and time stamps are available and can be used to create custom reports at the desired level of detail.

Cloud to ECS Appliance Backup

DataTrust Vapor will inventory the various buckets in the public cloud service. The inventory can be scheduled just as backup software is. In this case, the source is the cloud and the target is ECS Appliance. You can do periodic or continuous backups with support for file systems, object systems, and other cloud/REST options.

Wide Range of Cloud Backup and Recovery Options

- Single file/object
- Date or range of dates
- List of files/objects
- Full
- Recovery to different system than source
- Single object, date range, list, and full

Flexible Data Migration

If you are moving the cloud content to a different infrastructure, DataTrust Vapor has the ability to cross-map the original content handles to the new ones. If the original content is an object, Vapor can continue to use that object or redo it based on the originating application's requirements. If the original content was a file system, Vapor can continue to use that path information. What's more, if it was written as an object it can be written out as a file system. This conversion can be done on the fly during the migration process. This provides a lot of flexibility for moving data since it doesn't have to stay in original format.

Migration Chain of Custody

The objective of DataTrust Solutions is to meet the needs of companies in regulated industries when they have to transition content from a legacy system to a new system.

Data, metadata, streams, extended attributes, permissions, and ACLs not only have to be moved, but preserved and verified during the migration or data conversion process. In addition, compliance adds requirements including chain of custody, data retention, and data disposition.

DataTrust Solutions provides individual file/object verification by using cryptographic hashes on both the source and the target content. Vapor creates a chain-of-custody report detailing where the data was originally stored, any file attributes and time stamps and a content-based hash. Then it shows where the content is now and provides a hash of the content and the file attributes.

From this, DataTrust Solutions is able to provide the necessary documentation to prove that the content did not change during the course of the transition from the old to the new storage infrastructure.

Utilities for Updating the Source

DataTrust Solutions Vapor software also provides utilities for updating the source application when data is moved to a new type of storage system, whether it is private cloud or NAS. If the reference ID or GUID or path/file name changed, Vapor maintains a cross-reference mapping of the new ID/file to the old. This is used to update the application pointers with the new object ID or UNC path.

CONCLUSION

If your data stored on public clouds needs to be protected then you should look to EMC and DataTrust Solutions for cloud data protection.

Whether migrating data from legacy hardware or a public cloud, companies and organizations should look to EMC and DataTrust Solutions for comprehensive data migration solutions.

As this white paper has shown, EMC ECS Appliance and DataTrust Solutions Vapor software provide a high degree of flexibility in moving and managing cloud object data. This capability provides companies and organizations with the opportunity to lower costs, increase service levels, and enhance regulatory compliance.