



Education

CLOUD STORAGE SECURITY

INTRODUCTION

Gordon Arnold, IBM

- The material contained in this tutorial is copyrighted by the SNIA.
 - Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
 - This presentation is a project of the SNIA Education Committee.
 - Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
 - The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

➤ Cloud Storage Security Introduction

- ◆ Introduction of computing and data services in a virtualized and service provider context exposes the customer's information to a new set of threats and vulnerabilities. This session provides an introduction to those threats and what techniques are available to mitigate the threats.

At a NIST key management workshop
in the last year:

NIST Cloud security presentation got the reaction from
Whitfield Diffie –

“It took us 70 years to solve the problems created by
introducing radio – how long will it take us to solve
problems of security in the “cloud”?

Does that lead us to conclude that cloud security is
impossible?

- Having put our data on the airwaves – have we learned the right lessons and do we have the right methods for securing the data?

- Will the cloud be used for data that is low value or public - and therefore has no risk of exposure?
 - ◆ All data may have some value
 - ◆ Will clouds be a security forcing function – enterprises will have to classify data

- Will we use cloud/service providers appropriately and build in security from the ground up?

- You need to decide:
 - ✦ What, When, and How to Move to the Cloud
- To roughly assess your risk tolerance and importance of information being considered for moving to a cloud here are some rough steps.
- - Identify the asset for the cloud deployment - What kind of data is it?
- - Evaluate the asset - How important is it?
- - Map the asset to potential deployment models - Public, Private, Hybrid?
- - Evaluate potential cloud service models and providers
- - Understand the potential data flow - and where the data will be

- <http://cloudsecurityalliance.org/csaguide.pdf>

- Compliance and Audit - Investigative support
- Information Lifecycle Management - most relevant to storage
 - Data Security - Privileged user access
 - Location of the data - Data Location
 - Co-location of data - Data isolation and segregation
 - Discovery of ESI
 - Data Breach implications
 - Data Backup and Recovery - Recovery point and time objectives
- Legal and Electronic Discovery - Point to tutorial on this subject.
- Encryption and Key Management - DAR and in-flight
- Virtualization

Topics for Discussion

- Privileged user access
- Regulatory compliance
- Data location
- Data isolation and segregation
- Recovery point and time objectives
- Investigative support

Privileged User Access

- When your data is “in the cloud” instead of under your employee’s direct control
 - ◆ Do you trust service providers to keep that data private and protected?
 - ◆ Once the data is in the cloud – can you even identify what administrators could potentially access or what data they could potentially alter?
 - › For instance are they sub-contracting services
- You would like for the cloud to provide the services effectively without being able to read the sensitive or personally identifiable information in your data
 - ◆ Do your service level agreements cover who has access to your data?

Privileged User Access

- Do administrators of the cloud infrastructure have access to the customer data?
 - ◆ What storage administration tasks would expose customer data? – there should be a way to manage the data and recovery without being able to read the data
- One technical approach is to encrypt the data in transit and at rest
 - ◆ Is the data encrypted in the cloud and do the customers maintain control over the encryption keys?
 - › This can be part of the solution
 - ◆ Can the encryption keys be safely and securely provided to encryption processes without an opportunity for compromise or the cloud having to retain those keys?
- Can the cloud perform all the operations related to data availability and integrity without being able to read the data – except in encrypted form?
 - ◆ For instance can a copy of media be made without having to decrypt data?

- You and your customers will be subject to various regulations
 - ◆ Will the service provider be able to supply the necessary documentation for compliance?
 - Review the documentation and audit information which the service provider is able to provide
 - ◆ Is the documentation related to treatment of customer data sufficient to meet requirements for compliance?
 - For instance can a proof of encryption document be produced for any case of loss of physical storage media?
 - Ask if the documentation which would be provided is sufficient to invoke the safe harbor provisions of the privacy disclosure laws you may be subject to

Regulatory Compliance

- What is the documentation of compliance by the cloud provider and what are the compensating controls required to meet the compliance requirements?
 - ◆ How much trust can you put in representations – is the service agreement language sufficient to protect your interests?
- If the service provider is not able to meet your compliance requirements can you segregate the data and applications which can not be hosted in the cloud versus those you have to separate services?
 - ◆ You may be doing data classification to decide what data is appropriate for hosting in the cloud

- Do the customer's requirements limit where their data can be physically located?
 - ◆ Primary and any copies for recovery and availability
 - ◆ You may know about the primary data center location (or not) but there may be copies of the data outside of the country for recovery or archive functions
- National and regional laws may limit the ability to move data outside of national boundaries
 - ◆ EU privacy directive has specific restrictions on personally identifiable information being moved outside of the EU
 - ◆ State specific laws like those in Massachusetts and Nevada have specific restrictions on treatment of their citizens' data

Data Isolation and Segregation

- In the cloud – can different customers' data be isolated despite the shared infrastructure?
- Is the security implemented such that shared infrastructure or a multi-tenancy is supported without introducing compromises of the customer's data security?
- Do you pay extra for dedicated infrastructure or share infrastructure with the security able to partition the infrastructure?
 - ◆ virtual private cloud?

- Consider the workflow in the case where:
 - ◆ You want to rent computational capacity for business intelligence analytics on a monthly basis
 - ◆ You have to transfer the data to cloud for the processing
 - ◆ After the processing is complete how is the data cleansed from the cloud?
 - > Data over writes
 - > Cryptographic erasure

- Are there defined service level objectives for RPO and RTO?
- Is the availability of the service defined?
- Clouds can be complex environments with complex multi-layered applications and infrastructure environments
 - ◆ Do the service level agreements cover end to end responsiveness and availability?

- When processing and storing the customer's data
 - ◆ Are there sufficient audit trails around creation, access, modification, destruction of data?
 - ◆ Can the accuracy of the data be attested?
 - ◆ Is there a chain of custody defined for retrieval of data?
 - ◆ Is there a facility for a trusted third party for dispute resolution over data?
 - ◆ Is there a provision for snapshots of all customer data at a point in time?

- Are there sufficient audit trails around creation, access, modification, destruction of data?
 - ◆ Whenever anyone touches the data
 - > Is there an audit record written?
 - > Are those records kept in a tamper proof way?
 - Could a privileged user cover their tracks?
 - > Is there an attestation available that data destruction is complete?

- Can the accuracy of the data be attested?
 - ◆ Are digital signatures or hashes available to verify if data has been modified without the owner's knowledge?
 - › Originator's application may support digital signatures,
 - › Enterprises may provide an edge-of-enterprise gateway that adds digital signatures, or
 - › The service provider may provide a hash of the data as it is ingested
 - ◆ Encryption may not be sufficient
 - › Some encryption mode of operations (cypher blocking modes) may not be able to detect substitution of blocks if the encryption key was compromised

Chain of Custody

- Is there a chain of custody defined for retrieval of data?
 - ◆ Provenance
 - › Can you prove who created the data and that the data retrieved was exactly the data supplied the originator?
 - › When data is modified – can you prove who modified the data and that the data that is subsequently retrieved is unaltered?
 - For instance some encryption modes of operation detect changes or substitutions of blocks of data and some do not
 - Some storage like tape permits additional data to be added for each block to make sure that data has not been altered
 - ◆ If there was a legal discovery or law enforcement request can you provide the data in a way that will satisfy a court?
 - › You can't just pull the disk drives out of a common service...

Dispute resolution

- Is there a facility for a trusted third party for dispute resolution over data?
 - ◆ Is there the legal framework in place for a trusted third party to escrow and make data available in case there is a dispute?
 - > Who sent what when?
 - > Was it actually delivered – is the delivery receipt mechanism or proof of retrieval?
 - > Is there an ability to restrict access to the matter in question?
 - > Is there an ability to filter out PII or other data which may not be appropriate to disclose?
- This may come with cloud services where before there was not the possibility of the 3rd party being involved in a business to business transaction

Investigative Support

- Is there a provision for snapshots of all customer data at a point in time?
 - ◆ Fraud or other types of investigations may require a snapshot of the state of a company's representations, for instance
 - › The state of messages which have flowed back and forth
 - › The statements on a portal or web site
 - › Files which may have been shared
 - › Users' interactions with SaaS applications
 - ◆ For example – a dispute over an order when price, quantity, delivery dates, etc. are in question

Refer to Other Tutorials

- Please use this icon to refer to other SNIA Tutorials where appropriate.



**Check out SNIA Tutorial:
Enter Tutorial Title Here**

- Please send any questions or comments on this presentation to SNIA: tracksecurity@snia.org

**Many thanks to the following
for their contributions to this tutorial.**

- SNIA Education Committee

**Gordon Arnold, IBM
Eric Hibbard, HDS
Larry Hoffer, Emulex**

SNIA Security TWG

SNIA Cloud TWG