Education

Cloud Storage – Securing CDMI

Eric A. Hibbard, CISSP, CISA, ISSAP, ISSMP, ISSEP, SCSE
Hitachi Data Systems

SNIA Legal Notice

- 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.

Abstract

Cloud Storage – Securing CDMI

With SNIA's publication of the Cloud Data Management Interface (CDMI) specification, cloud storage implementations can now offer a standard set of features and capabilities. Security is part of this feature set and some believe that it is a make-or-break element of cloud storage, and cloud computing in general.

This session will overview the security of the new CDMI standard, which includes protective measures employed in the management and access of data and storage. These measures span transport security, authentication, authorization and access controls, data integrity, sanitization, data retention, protections against malware, data at rest encryption, and security capability queries.

Cloud Computing Overview

NIST Cloud Computing Attributes

- *On-demand self-service.* A consumer can unilaterally provision computing capabilities as needed automatically without requiring human interaction with each service's provider.
- *Broad network access.* Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms.
- *Resource pooling.* The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- *Rapid elasticity.* Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in.
- *Measured Service.* Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service.

NIST Cloud Service Models

- *Cloud Software as a Service (SaaS).* The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure.
- *Cloud Platform as a Service (PaaS).* The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider.
- *Cloud Infrastructure as a Service (IaaS).* The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications.

NIST Cloud Deployment Models

- *Private cloud.* The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.
- *Public cloud.* The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- *Hybrid cloud.* The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

NOTE: NIST also defines a *community cloud* model, which is not included here because it not very common.

SPI Service Model

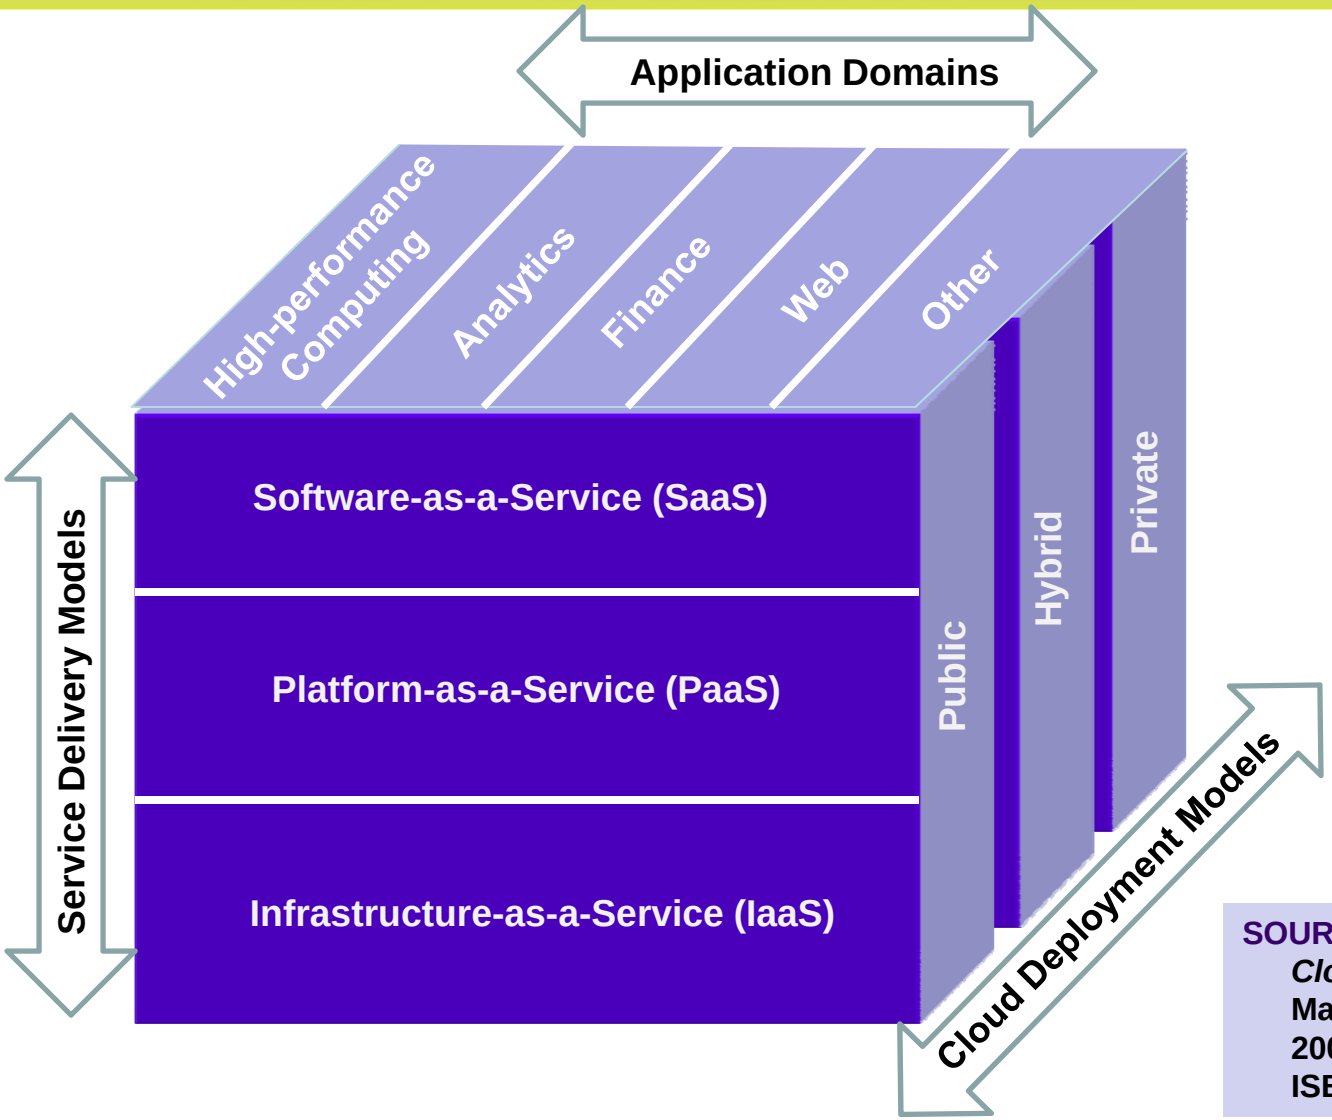

SOURCE:
Cloud Security and Privacy, **Mather, Kumaraswamy, Latif, 2009, O'Reilly, ISBN: 978-0-596-80276-9.**

Top Threats to Cloud Computing

Education SNIA

#1: Abuse and Nefarious Use of Cloud Computing (IaaS & PaaS)

#2: Insecure Interfaces and APIs (IaaS, PaaS, SaaS)

#3: Malicious Insiders (IaaS, PaaS, SaaS)

#4: Shared Technology Issues (IaaS)

#5: Data Loss or Leakage (IaaS, PaaS, SaaS)

#6: Account or Service Hijacking (IaaS, PaaS, SaaS)

#7: Unknown Risk Profile (IaaS, PaaS, SaaS)

SOURCE: Cloud Security Alliance, *Top Threats to Cloud Computing*, Version 1.0, 2010, http://www.cloudsecurityalliance.org/topthreats.

Cloud Computing Security Guidance

Governance	Operations
Governance and Enterprise Risk Management	Traditional Security, Business Continuity and Disaster Recovery
Legal and Electronic Discovery	Data Center Operations
Compliance and Audit	Incident Response, Notification and Remediation
Information Lifecycle Management	Application Security
Portability and Interoperability	Encryption and Key Management
	Identity and Access Management
	Virtualization

SOURCE: Cloud Security Alliance, *Security Guidance for Critical Areas of Focus in Cloud Computing*, Version 2.1, 2009, http://www.cloudsecurityalliance.org/guidance.

Overview of CDMI

Check out SNIA Tutorial:

The Cloud Data Management Interface (CDMI) - The Cloud Storage Standard

Cloud Data Management Interface

- Applicable to three types of Cloud Storage:
  - Cloud Storage for Cloud Computing
    - Whitepaper at snia.org/cloud – the management interface for the lifecycle of storage in a compute cloud
  - Public Storage Cloud
    - Both a Data Path for the Cloud and a Management Path for the Cloud Data
  - Private Cloud Storage
    - As well as hybrid clouds
    - An API for Storage Vendors selling into Cloud based solutions
- Semantics
  - Simple Containers and Data Objects with tagged Metadata
  - Data System Metadata expresses the data requirements
- Protocol
  - RESTful HTTP as "core" interface style
  - JSON (JavaScript Object Notation)– format of the representations are extensible

CDMI Data Objects

- Stored data can be accessed using native protocols:
  - HTTP, CIFS, NFS, iSCSI, SQL, etc.
- Stored data can also be accessed using CDMI as a Data Path in a standardized manner. This facilitates:
  - Cloud-to-cloud migration
  - Cloud federation
  - Cloud backup
  - Cloud virus scanning
  - Cloud search
  - And more.
- Desired cloud storage characteristics can be associated with stored data:
  - Replication, Compression, Placement, Retention, QoS, etc.

The Complete CDMI Picture

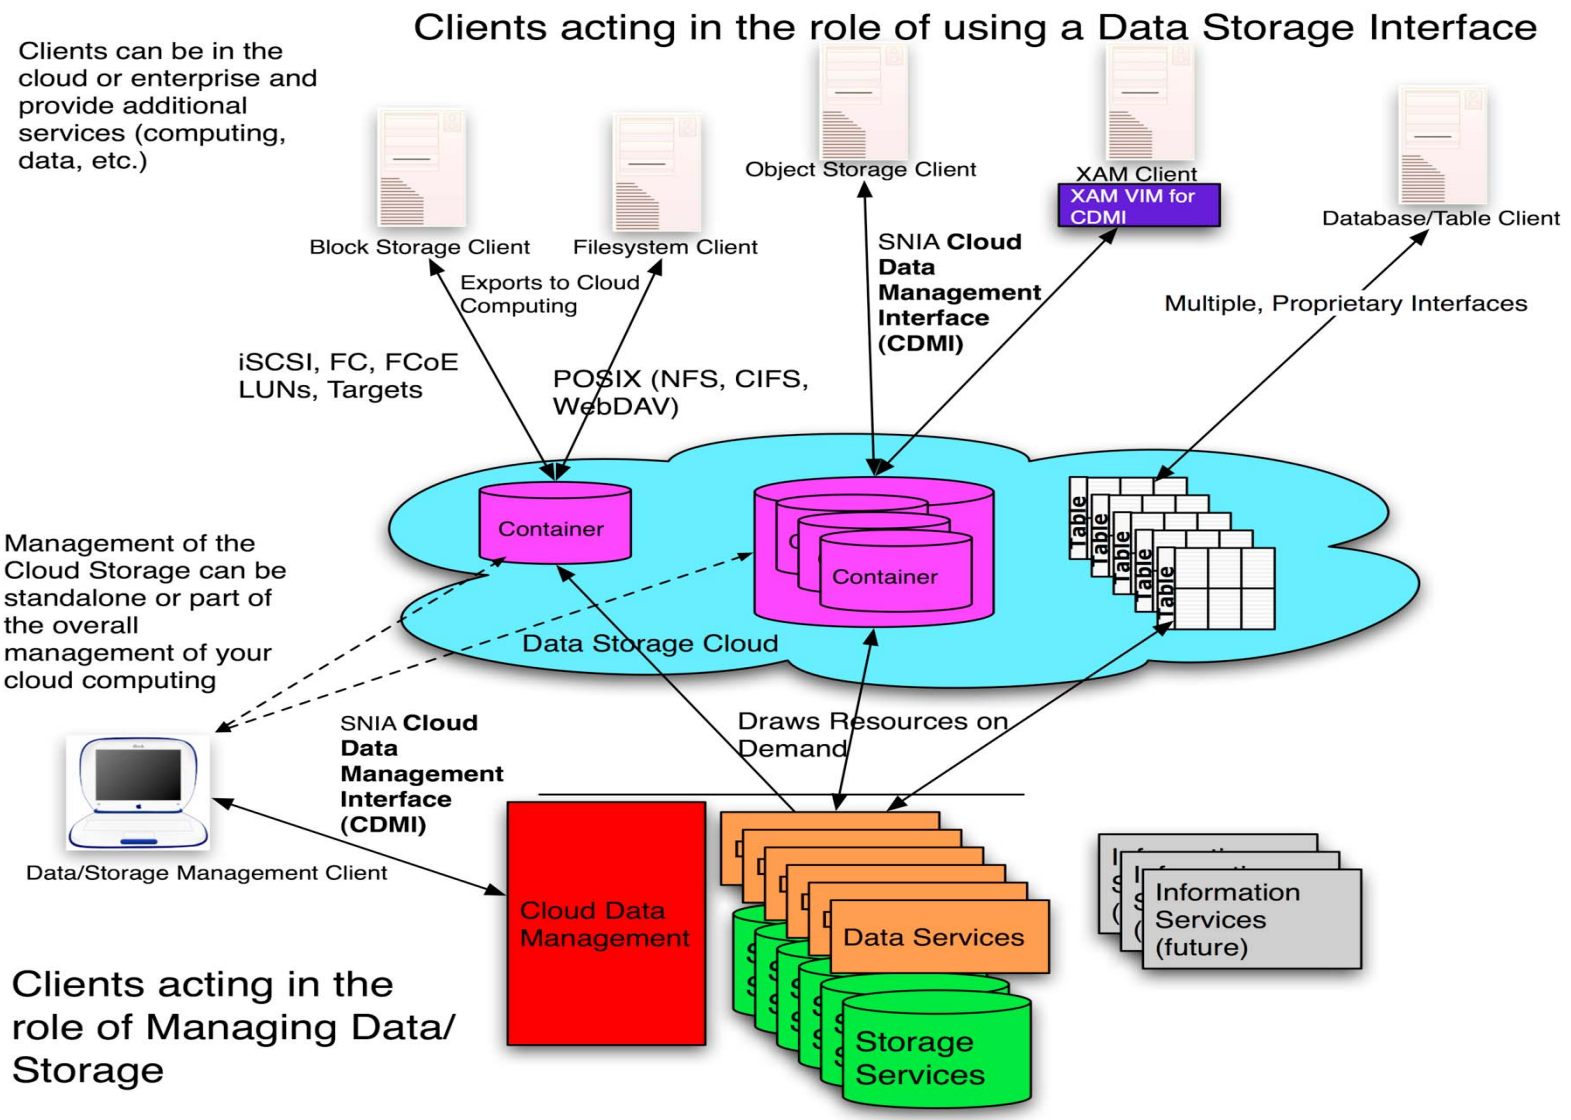

CDMI Security

Check out SNIA Tutorial:

Cloud Storage Security

CDMI Security Summary

- Security refers to the protective measures employed in managing and accessing data and storage.
- Security measures:
  - Include transport security, user and entity authentication, authorization and access controls, data integrity, data and media sanitization, data retention, protections against malware, data at rest encryption, and security capability queries.
  - Take the form of mandatory, optional, and vendor extensions
- The transport security and security capability queries are mandatory for all implementation; all other security mechanisms are optional to implement.
- Client use of security is always optional, but encouraged.

CDMI Security Objectives

- Provide a mechanism that assures that the communications between a CDMI client and server cannot be read or modified by a third party
- Provide a mechanism that allows CDMI clients and servers to provide an assurance of their identity
- Provide a mechanism that allows control of the actions a CDMI client is permitted to perform on a CDMI server
- Provide a mechanism for records to be generated for actions performed by a CDMI client on a CDMI server
- Provide mechanisms to protect data at rest
- Provide a mechanism to eliminate data in a controlled manner
- Provide mechanisms to discover the security capabilities of a particular implementation

Transport Security

- HTTP is the mandatory transport mechanism for CDMI
- HTTP over TLS (HTTPS), on TCP port 443, is the mechanism used to secure the communications between CDMI entities
  - TLS 1.1 must be implemented by CDMI entities and TLS 1.2 is strongly encouraged
  - Mandatory cipher suites include: TLS_RSA_WITH_AES_128_CBC_SHA, TLS_RSA_WITH_AES_256_CBC_SHA256, and TLS_RSA_WITH_NULL_SHA
  - The DER encoded X.509, Base64 encoded X.509, and PKCS#12 certificate formats must be supported.
  - Certificate Revocation Lists must be supported in the DER encoded X.509 and Base64 encoded X.509 formats.

System-wide Security Capabilities

- Capabilities defined for the cloud storage system (found in the capabilities objects for domains, data objects, containers, and queues):
  - **cdmi_domains** – If present and "true", the cloud storage/computing system supports domains.
  - **cdmi_queues** – If present and "true", the cloud storage/computing system supports queue objects.
  - **cdmi_security_audit** – If present and "true", the cloud storage system supports audit logging.
  - **cdmi_security_data_integrity** – If present and "true", the cloud storage system supports data integrity/authenticity.
  - **cdmi_security_encryption** – If present and "true", the cloud storage system supports data at rest encryption.
  - **cdmi_security_https_transport** – If present and "true", the cloud storage system supports HTTPS communications.
  - **cdmi_security_immutability** – If present and "true", the cloud storage system supports data immutability/retentions.
  - **cdmi_security_sanitization** – If present and "true", the cloud storage system supports data/media sanitization.

Domains (Administrative Ownership)

- A CDMI implementation may optionally include a hierarchy (parent-child) of administrative ownership of stored data within a CDMI storage system
- Each domain corresponds to logical groupings of objects that are to be managed together
- Domain measurement information about objects that are associated with each domain flow up to parent domains, facilitating billing and management operations
- The Domain membership capability provides information about and allows the specification of end users and groups of users that are allowed to access the domain via CDMI and other access protocols
- Domains provide a single unified place to map identities and credentials to principals used by ACLs within the context of a domain

Access Controls

- A CDMI implementation may optionally implement access control lists (ACLs)
- CDMI specifies three types of privileged users: “administrator”, “backup_operator”, and “cross_domain”
- ACLs are lists of permissions-granting or permissions-denying entries called *access control entries* (ACEs).
  - An ALLOW ACE grants some form of access to a *principal*. Principals are either users or groups, and are represented by *identifiers*.
  - A DENY ACE denies access of some kind to a principal.
- ACEs are composed of five fields: “type”, “who”, “flags”, “access_mask”, and “timestamp”
- When evaluating whether access to a particular object *O* by a principal *P* shall be granted, the server traverses the object's logical ACL (its ACL after processing inheritance from ancestor containers) in list order

CDMI Logging

- Logging is divided into three functional areas:
  - CDMI object functions
  - Security events
  - Data management events
- CDMI clients can access log data by creating a logging queue that defines the scope of log messages to be received
- Queues are a special class of container (i.e., FIFO); logging queues are persistent
- Access controls can be applied to queues

Cryptographic Support

- A CDMI implementation may optionally implement cryptographic capabilities
- Data Integrity
  - A CDMI client can determine the hashing options available by checking the ***cdmi_value_hash*** data systems metadata
  - A CDMI client can request a particular hashing option, using the ***cdmi_value_hash*** data systems metadata
  - A CDMI client can determine the actual hashing option used by the CDMI implementation by checking the ***cdmi_value_hash_billed***
  - A CDMI client can determine the actual hash value by checking the ***cdmi_hash*** metadata
- Data At Rest Encryption
  - A CDMI client can determine the encryption options available by checking the ***cdmi_encryption*** data systems metadata
  - A CDMI client can request a particular encryption option, using the ***cdmi_encryption*** data systems metadata
  - A CDMI client can determine the actual hashing option used by the CDMI implementation by checking the ***cdmi_encryption_billed***

Immutability

- A CDMI implementation may optionally implement retention management disciplines
  - Retention – uses retention time criteria to determine the time period deletions are prohibited; only one per object and extensions of the object metadata are allowed
  - Hold – enforces read-only data object access and prohibition of object deletion; multiple holds are allowed
  - Deletion – manual and/or automatic
- Enforcements associated with value changes to the retention duration are not a CDMI responsibility
- Releases from holds are performed out-of-band or by vendor extension

CDMI Security Guidance

Exploiting the Mandatory and Optional Features

- Always check the security capabilities of your cloud service provider's CDMI implementation
  - Ensure it has adequate protective measures
  - Make a "risk" based decision to use a particular implementation
- Use TLS (preferably TLS 1.2) to
  - Authenticate CDMI entities (certificates for servers; HTTP authentication for clients)
  - Encrypt sensitive information communicated between CDMI entities.

Exploiting the Mandatory and Optional Features (cont.)

- Use Domains to provide a place for authentication mappings to external authentication providers
- Audit logging within the context of CDMI
  - Establish logging queues and restrict access
  - Capture messages for all security and data management events
  - Make sure the CDMI client retrieves the messages on a regular basis

Exploiting the Mandatory and Optional Features (cont.)

- Align the automatic deletion capability with the organization's data retention policy
- Prior to using Holds, understand the process and mechanism for lifting the Holds
- For cryptographic functionality, it is always important to verify that the implementation has complied with the requested algorithm; something other than what was requested may be used

Addressing CSA Top Threats with CDMI

#1: Abuse and Nefarious Use of Cloud Computing

- Monitor public blacklists for one's own networks and URIs.

#2: Insecure Interfaces and APIs

- Implement and use TLS for encrypted communications
- Implement and use authentication (with CDMI Domains)
- Implement and use access control lists (CDMI Access Control)
- Implement and use security logging (CDMI Logging Queues)
- Only use exported protocols with appropriate security mechanisms.

#3: Malicious Insiders

- All elements defined for #2.
- Encrypt data before storing in the CDMI implementation.

Addressing CSA Top Threats with CDMI (cont.)

Education SNIA

#4: Shared Technology Issues

- Enforce service level agreements for patching and vulnerability remediation.
- Conduct vulnerability scanning and configuration audits.

#5: Data Loss or Leakage

- All elements defined for #3
- Contractually demand providers sanitize persistent media before releasing it for reuse.
- Contractually specify provider backup and retention strategies.

#6: Account or Service Hijacking

- All elements defined for #2.
- Leverage strong two-factor authentication techniques where possible.

#7: Unknown Risk Profile

- Implement and use TLS v1.2 when possible.
- Due diligence on the providers' infrastructure and approach.

Final Thoughts

Final Thoughts

- Security and legal issues will persist as challenges for organizations that choose to use cloud computing, but there are promising signs that some of these issues will be addressed.
- It is, however, extremely important to understand the risks and to enter the cloud with your eyes wide open (i.e., select a cloud service provider that offers an appropriate set of contractual terms and conditions as well as demonstrable risk mitigations).

Q&A / Feedback

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

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

- SNIA Education Committee

Eric A. Hibbard, CISSP, CISA
Larry Hofer, CISSP, PE
Mark Carlson
Gianna DaGiau

SNIA Cloud Storage TWG

For More Information

Cloud Security Resources

- SNIA Cloud Storage Initiative, http://www.snia.org/cloud
- Cloud Security Alliance (CSA), *Security Guidance for Critical Areas of Focus in Cloud Computing*, *Top Threats to Cloud Computing*, http://www.cloudsecurityalliance.org
- European Network and information Security Agency (ENISA), *Cloud Computing – Benefits, risks and recommendations for information security*, http://www.enisa.europa.eu/
- Information Systems Audit and Control Association (ISACA), *Cloud Computing: Business Benefits With Security, Governance and Assurance Perspectives*, http://www.isaca.org
- *Cloud Security and Privacy*, Mather, Kumaraswamy, Latif, 2009, O'Reilly Publishing, ISBN: 978-0-596-80276-9

Relevant Standards Activities

- *Open Grid Forum (OGF)* is developing on an Open Cloud computing Interface (OCCI)
- *Storage Networking Industry Association (SNIA)* is developing the Cloud Data Management Interface (CDMI) specification
- *Cloud Computing Interoperability Forum (CCIF)* is developing Unified Cloud Interfaces and APIs
- *Distributed Management Task Force (DMTF)* has established the "Open Cloud Standards Incubator" to develop a set of informational specifications for Cloud resource management
- *Open Cloud Consortium (OCC)* is researching the creation of inter-Cloud interfaces with the aim of developing compatibility standards
- *Cloud Security Alliance (CSA)* to promote the use of best practices for providing security assurance within Cloud computing
- *Object Management Group (OMG)* to establish a uniform vocabulary for Cloud Computing, as well as to synchronize standards development
- ISO/IEC JTC 1 Subcommittee 38 (SC38) on Distributed Application Platforms and Services (DAPS) has a focus on Web services, SOA, and cloud computing

SNIA Security Organizations

- SNIA Security Technical Work Group (TWG)
  - Focus: Requirements, architectures, interfaces, practices, technology, educational materials, and terminology for storage networking.
  - http://www.snia.org/tech_activities/workgroups/security/
- Storage Security Industry Forum (SSIF)
  - Focus: Educational materials, customer needs, whitepapers, and best practices for storage security.
  - http://www.snia.org/ssif