Use of Storage Security in the Cloud

David Dodgson

Unisys

Purpose

- Address the concern about security of storage in a cloud
- Define what secure storage means to different users of different clouds

Objectives

- Describe common use cases for storage in the cloud
- Compare security needs for different use cases
- Compile a list of requirements for storage in clouds, both public and private
- Prioritize the requirements

Use Cases

Storage Use Cases

- Infrastructure as a Service (IaaS)
- Platform as a Service (PaaS)
- Storage as a Service
- Hybrid Storage

Infrastructure as a Service

- Uses virtual machines (VMs) that are similar to physical machines
- Accesses storage using traditional methods (files and directories)

Tenant and Management

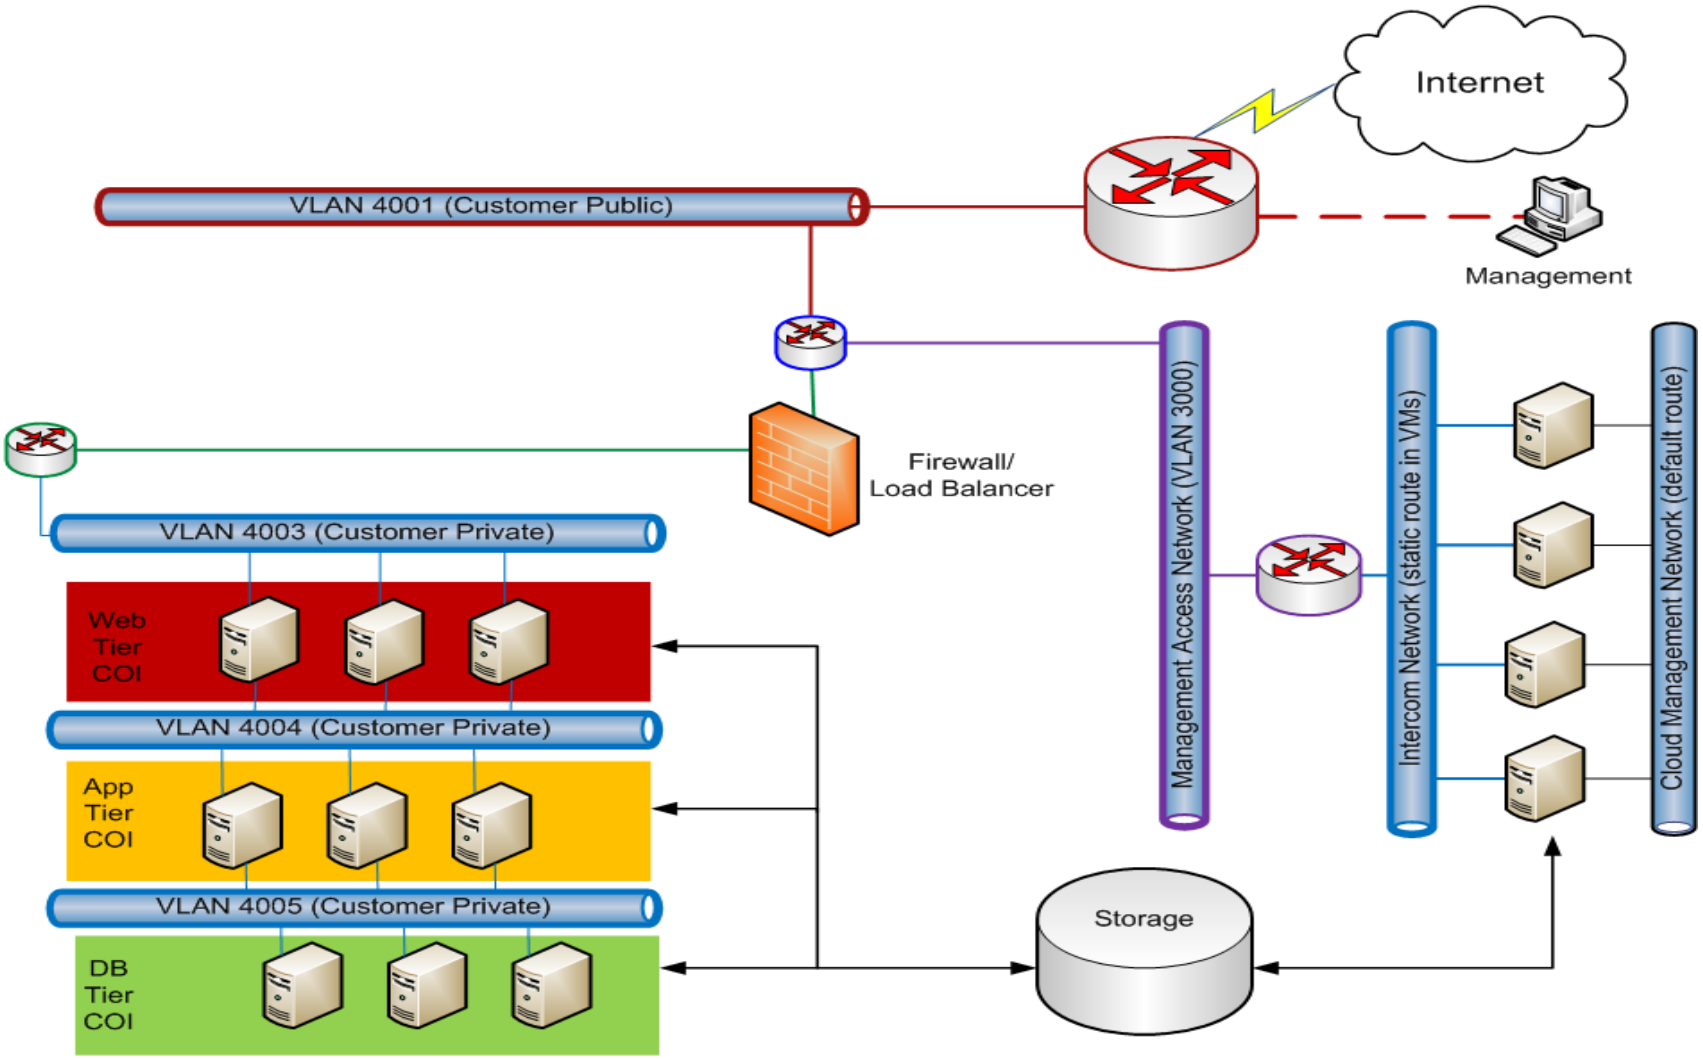

Communities of Interest

- The VMs are grouped into separate communities-of-interest (COIs).
- They share storage with other VMs in the same COI but will be separate and secure from other COIs.
- A VM may be a member of multiple COIs, but data from each COI will be presented distinctly (e.g. as separate disks).
- Cloud management data is a separate COI.

Secure COIs

- Storage requirements for a COI are:
  - Encryption
  - Key management
  - Network configuration
  - Data verification
  - Storage virtualization

Platform as a Service

- Pre-defined virtual machines
- Application support (I/O through cloud services)
- Secure COIs rely more on identity management than network configuration
- Auditing may be affected by an application moving from one platform to another

Storage as a Service

- Storage objects that reside in the cloud
- Each object has its own unique identity
- Accessed through network interface (usually a RESTful web service)
- Each object has its own metadata
- Requirements usually depend on the service level agreement (SLA)

Typical Use of Storage as a Service

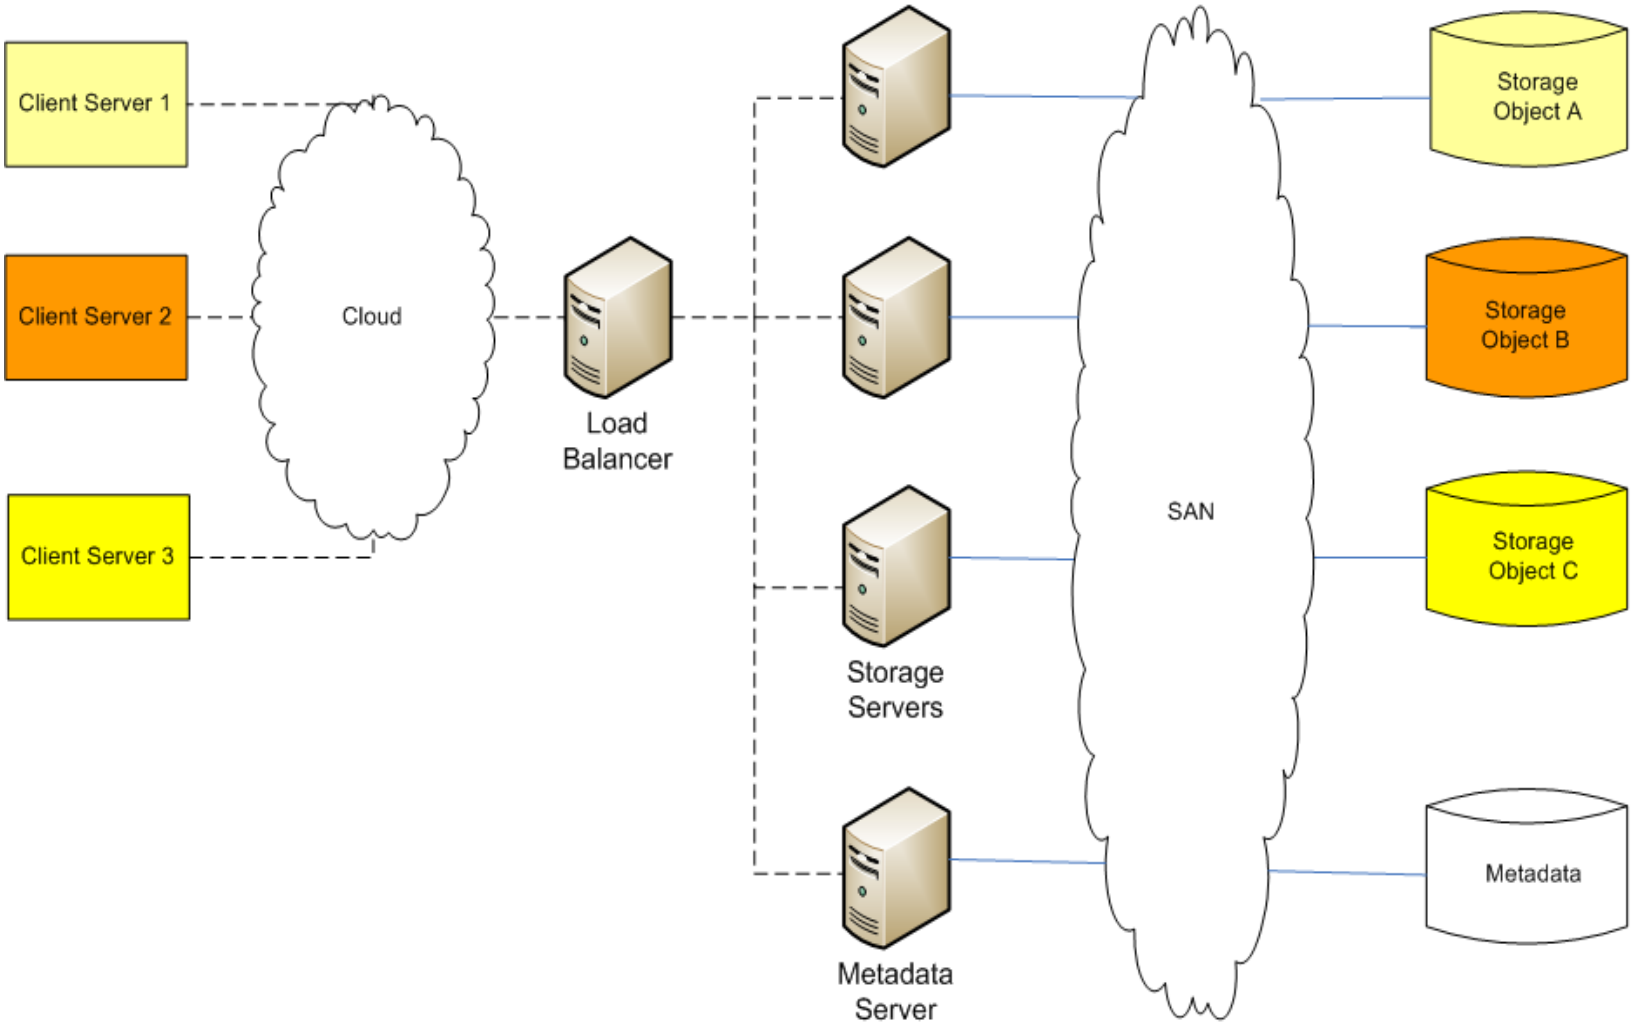

Hybrid Storage

- Combines IaaS with Storage as a Service
- Legacy applications continue to use directories and files while storage resources are provided from the cloud
- Local storage cache used to improve performance and reduce charges

Typical Hybrid Storage

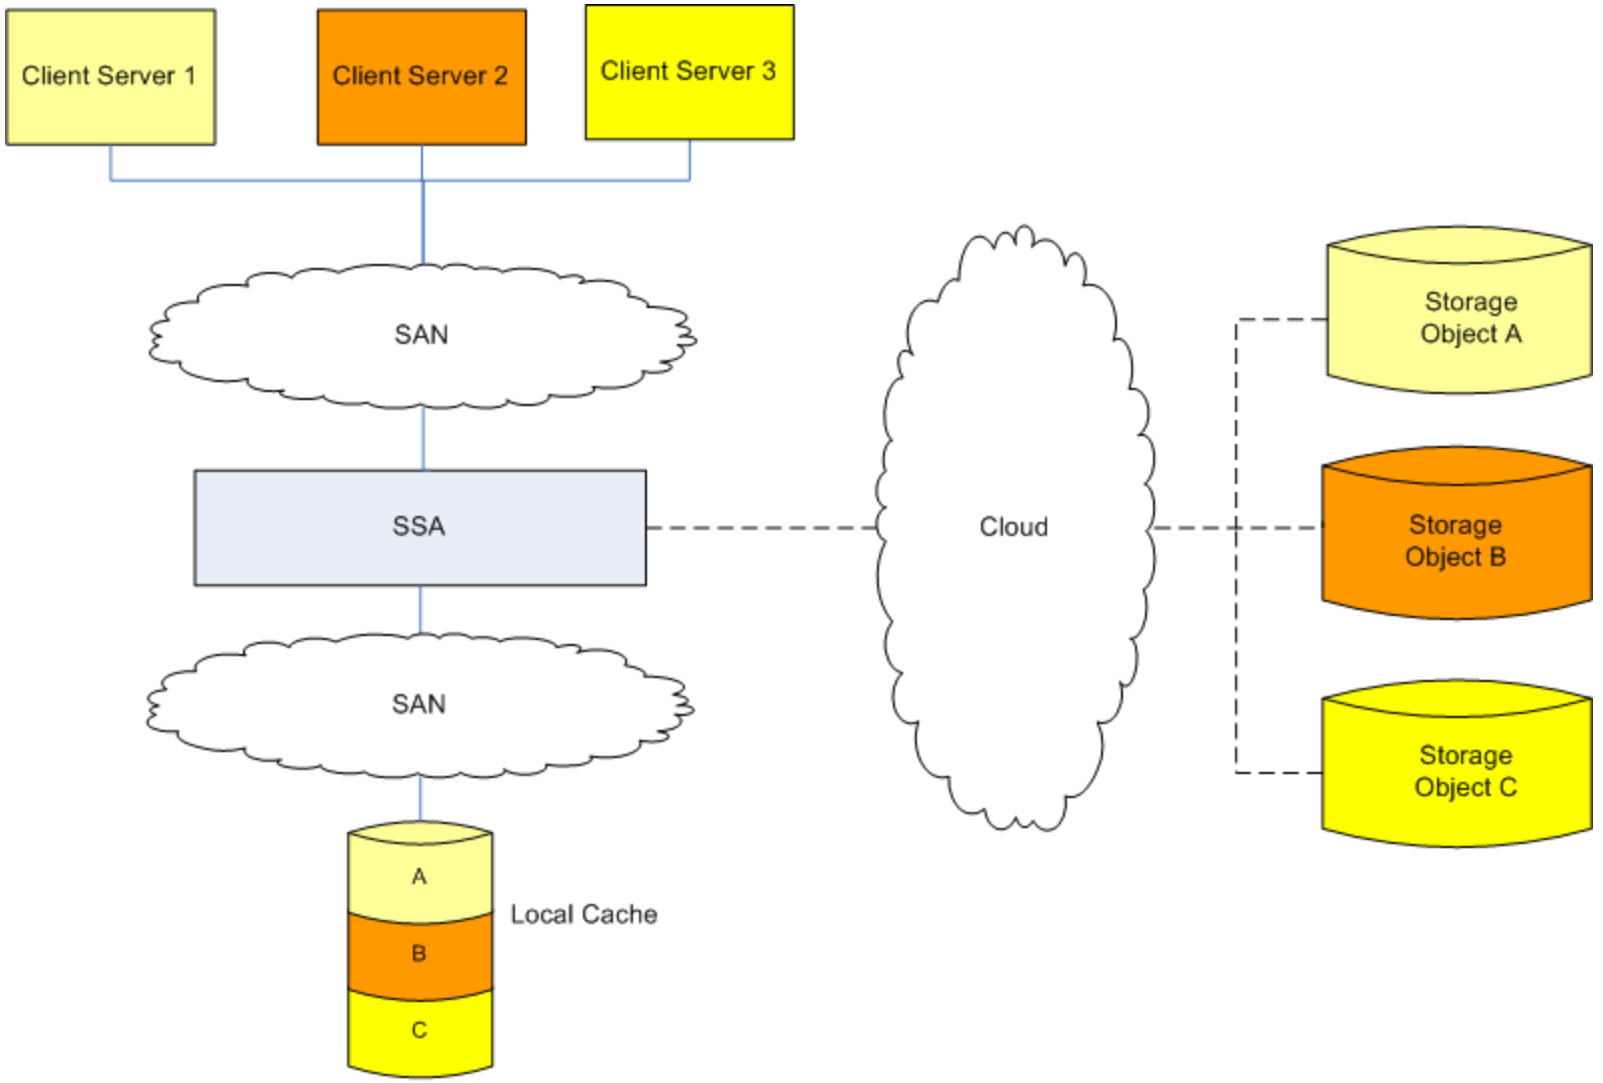

Storage Requirements

Storage Requirements Chart

Requirement	*Priority*			
	IaaS	*PaaS*	*Storage*	*Hybrid*
Encryption	1	1	3	1
Key management	2	2	4	2
Network configuration	3	8	10	10
Data verification	6	7	5	6
Storage virtualization	4	5	7	5
Identity management	7	3	2	4
Auditing	8	9	6	8
Regulatory compliance	9	10	11	11
Storage objects	12	4	1	7
Object cache	13	13	9	3
Storage management	5	6	8	9
Physical security	10	12	12	13
Software management	11	11	13	12

Encryption

- Prevents one tenant from reading another's data.
- Prevents against an insider copying data.
- Data is encrypted using a standard encryption method, e.g. AES-256. All sensitive information, both data and metadata is encrypted.
- The encryption mechanism must satisfy compliance (PCI, FIPS 140-2).
- It must be possible to rekey data at need, preferably without denying access to the data.

Key Management

- Provide keys to authorized users securely
- Manage life-cycle of keys
  - Creation
  - Storing (perhaps for years)
  - Re-encryption
  - Deletion
- Key management must allow for users to be added or removed from COIs easily.

Key Management

- Separate management role for the crypto-administrator.
- Tenant should be a crypto-administrator
  - Cloud component determines COI membership
  - Key management system in tenant's data center
  - Cloud stores keys in memory keystore
- Transmission of keys between the key management system and the cryptographic system must be secure.

Network Configuration

- The storage area network (SAN) should be configured to limit access to the data
  - Fibre-channel uses zones and masking
  - IP uses vLANs and firewalls

Network Configuration

- Relevant SAN protocols must be supported depending on the cloud requirements (FC, iSCSI, NAS, FCoE).
- It must be possible for the SAN to be configured so that access is restricted to the appropriate users, preferably from a single pane-of-glass.
- It should be possible to control access levels based upon the configuration (e.g. hypervisor vs. VM).

Typical SAN

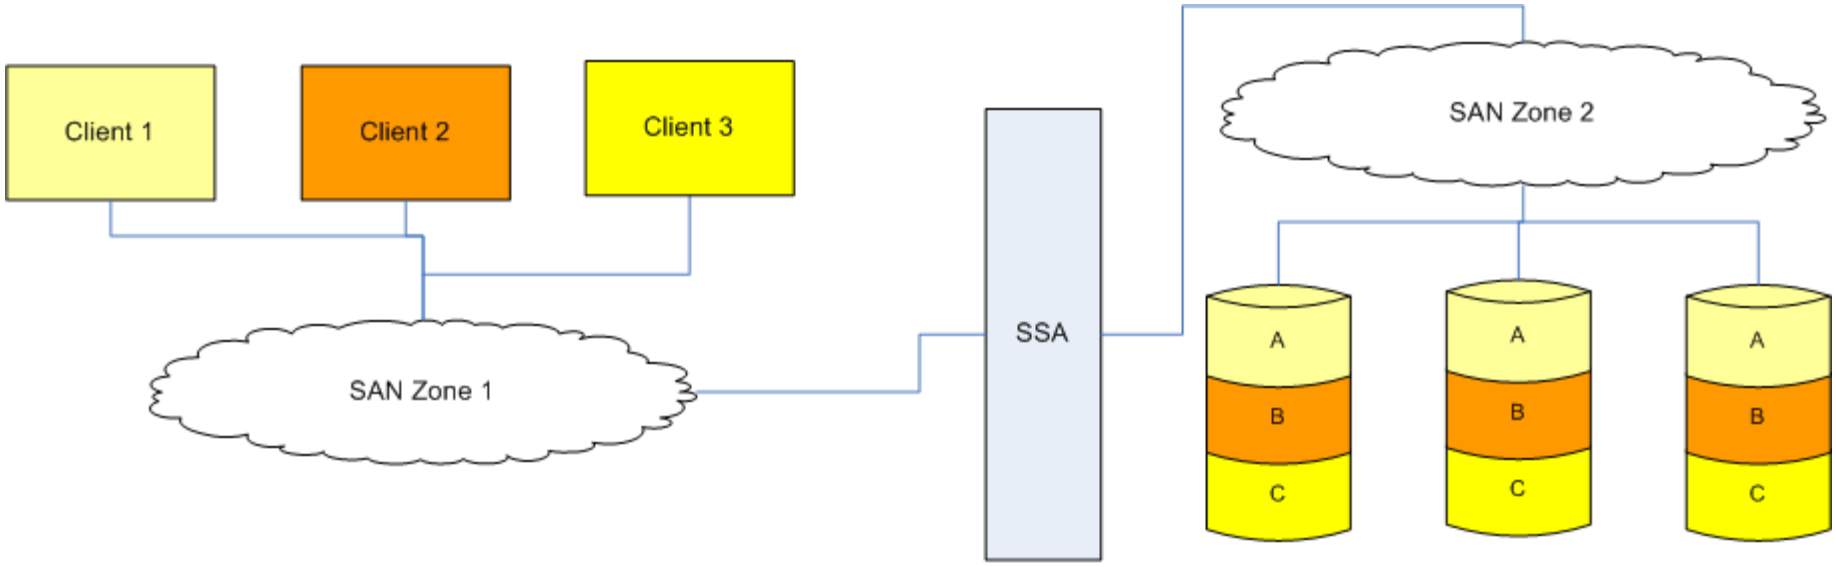

Data Verification

- Verify that data that is read is the same as that which was stored.
- Protects against data tampering, either accidental or intentional.
- Possible techniques include SHA-2, checksums, homomorphic tags.
- Digitally signing at least part of the data (e.g. metadata) allows verification that it was used to perform the initial encryption.

Storage Virtualization

- Allocate available storage among the COIs.
- Needed to handle small storage COIs efficiently.
- Allocate only the storage needed. Cloud users don't want to pay for excess storage.
- Present the storage to the VM using a method understood by its OS (e.g. FC would use a virtual disk).

Identity Management

- Limit the visibility of storage in a COI to authenticated actors (user, application, or resource).
- An identity is a verifiable set of credentials.
- An actor is authenticated and authorized.
- Access control limits what an individual actor can do with a resource through the use of roles.
- Certificates may be used by an actor to protect sensitive information (keys).

Auditing

- Auditing is the recording of significant events, usually in a log.
- Events from multiple sources may be combined in a single log. This may make time synchronization a necessity.
- Multiple views of the log may be required. The audit administrator may see a different view than a tenant.
- Operations may stop when the log is full.
- Alerts can be used for events that require human intervention.

Regulatory Compliance

- Compliance with regulations is mandatory in certain circumstances (PCI, HIPAA).
- Legal enforcement across different legislative areas can be a problem. Exposure may be limited by:
  - Tenant control of keys
  - Location specification
- Cloud providers may provide building blocks to knowledgeable third-party vendors.

Storage Objects

- Storage objects contain data that is addressed in a manner independent of the platform.
- Configured at creation time based on metadata. The set of supported metadata traits may expand over time or by SLA levels.
- Standard (e.g. CDMI) vs. proprietary (e.g. Amazon, Google, …) interfaces.

Object Cache

- Cache storage objects locally, rather than accessing the cloud for every I/O.
- Improves performance and reduces costs.

Storage Management

- Quality-of-service of the storage is determined by resiliency, performance, and administration.
- Resiliency is determined by RAID, backup and snapshots, and disaster recovery.
- Performance is determined by throughput and latency.
- Administration is determined by how storage is provisioned and managed, preferably using a GUI from a single pane of glass.
- It must be possible to move data between storage of different QoS characteristics.

Physical Security

- Security of the physical location
- Reliability and cost of power
- Environmental concerns (e.g. cooling)
- Ruggedness of the hardware itself (especially in military situations)

Software Management

- Software errors should not stop the system.
- Neither should software updates.
- Error reporting can be a problem in secure situations. Defense and commercial applications may have different requirements.

Summary

- Secure storage COIs can be used to provide security for storage in the cloud.
- The large number of possible requirements for storage in the cloud makes prioritization mandatory.
- The best way of getting agreement between customers and vendors is by examining use cases.

Links

- Unisys Corporation
  http://www.unisys.com
- SNIA Cloud Storage Initiative
  http://www.snia.org/forums/csi
- Representational State Transfer (REST)
  http://en.wikipedia.org/wiki/Representational_State_Transfer
- Homomorphic codes, see "Provable Data Possession at Untrusted Stores"
  http://portal.acm.org/citation.cfm?id=1315318

Questions

- ??