CCSP by Alukos

Overview

Purpose and Credit

These notes were created by the following members of the Alukos team:

- Michael Ferullo (Mikef#1337 on Discord)

These notes have been created and shared for the sole purpose of aggregating accurate information from reliable sources in an effort to facilitate studying for the ISACA CISA examination. Alukos works directly with the Certification Station to ensure these notes are continuously maintained using a collaborative effort and are shared publicly with the community.

Additional Resources

For additional resources or to learn more about Alukos, please check out our website.

Disclaimer

1. Most of the information contained within these notes are copyrighted and the sources have been documented accordingly.
2. Any page with an asterisk (*) at the end of the name is either incomplete, potentially inaccurate, irrelevant for the exam, or a combination of these, and may require review.

Index

References

Books

- Gordon, A. (2016). *The Official (ISC)2 Guide to the CCSP CBK.* Sybex.
- Malisow, B. (2017). *CCSP (ISC)2 Certified Cloud Security Professional Official Study Guide.* Sybex.
- Carter, D. (2017). *CCSP Certified Cloud Security Professional All-in-One Exam Guide.* McGraw-Hill Education.

Multimedia

- Handerhan, K. (2019). *Certified Cloud Security Professional (CCSP)* [Video series]. Cybrary. https://www.cybrary.it
- [VMware End-User Computing] (2019, June 7). *Identity and Access Management: Technical Overview* [Video]. YouTube. https://www.youtube.com/watch?v=Tcvsefz5DmA
- [VMware End-User Computing] (2019, December 12). *SAML 2.0: Technical Overview* [Video]. YouTube. https://www.youtube.com/watch?v=SvppXbpv-5k
- [OktaDev] (2019, November 5). *An Illustrated Guide to OAuth and OpenID Connect* [Video]. YouTube. https://www.youtube.com/watch?v=t18YB3xDfXI
- [Programming with Mosh] (2018, January 19). *What is a REST API?* [Video]. YouTube. https://www.youtube.com/watch?v=SLwpqD8n3d0

Online Resources

- Cloud Security Alliance. (2017). *Security Guidance v4.0* [PDF file]. https://cloudsecurityalliance.org/research/guidance/

Terminology*

> This page contains accurate information but may contain redundant definitions or may be missing definitions that should be included.

A-C

B

Bastion Host

A bastion host is a method for remote access to a secure environment. The bastion host is an extremely hardened device that is typically focused on providing access to one application or for one particular usage. Having the device set up in this focused manner makes hardening it more effective. Bastion hosts are made publicly available on the Internet.

> The difference between a jump server and a bastion host is that a jump server is intended to breach the gap between two security zones and have a gateway to obtain access to something inside of the other security zone. A bastion host is outside of your security zone and will require additional security considerations.

C

Capability Maturity Model (CMM)

A way of determining a target's maturity in terms of **process** documentation and repeatability. Contains five levels. Is typically not associated with security, however.

Confidentiality

Protecting information from unauthorized access/dissemination.

Controls

Act as mechanisms designed to restrict a list of possible actions to allowed or permitted actions. If a control is breached, the next step of mitigating is a countermeasure. Proactive.

Cost Benefit Analysis

A cost-benefit analysis (CBA) is the process used to measure the benefits of a decision or taking action minus the costs associated with taking that action. It determines whether certain activities (such as BC/DR) are worth implementing. The CBA will compare the costs of a disaster and the impact of downtime against the cost of implementing the BCDR solution. Another example would be whether the movement to a cloud model would be lower than the cost of not moving to the cloud.

Countermeasure

Countermeasures are what is deployed once a control has been breached. Reactive.

Cross-Cutting Aspects

Cross-cutting aspects are behaviors or capabilities which need to be coordinated across roles and implemented consistently in a cloud computing system.

An example of a cross-cutting aspect is security.

D-F

D

Data Remanence

Any data left over after sanitization and disposal methods have been attempted.

Demand Management

Can we meet our demand requirements (can we scale up and down)? Elasticity in the cloud solves this.

Digital Signatures

Uses the senders private key *and* a hash to guarantee the integrity and the origin (authenticity/non-repudiation) of a message. This method requires a PKI.

E

Enterprise Application

Applications or software that a business would use to assist the organization in solving enterprise problems.

European Economic Area (EEA)

The European Economic Area, abbreviated as EEA, consists of the Member States of the European Union (EU) and three countries of the European Free Trade Association (EFTA) (Iceland, Liechtenstein and Norway; excluding Switzerland). The Agreement on the EEA entered into force on 1 January 1994.

F

Fault Tolerance

Fault tolerance involves the use of specialized hardware that can detect faults and automatically switch to redundant components or systems.

> ⓘ Should be used when the goal is to **eliminate** system downtime as a threat to system availability altogether.

Financial Management

Do we have the funds and can we allocate them appropriately? Are we receiving a good return on investment (ROI)? Are we good stewards of the money entrusted to us? Are we making a profit?

Fraggle Attack

A variation to the Smurf attack is the Fraggle attack. The attack is essentially the same as the Smurf attack but instead of sending an ICMP echo request to the direct broadcast address, it sends UDP packets.

G-I

G

Generally Accepted Accounting Practices (GAAP)

GAAP is a combination of authoritative standards (set by policy boards) and the commonly accepted ways of recording and reporting accounting information. GAAP aims to improve the clarity, consistency, and comparability of the communication of financial information.

These are the industry standards, also recognized by the courts and regulators, that accountants and auditors must adhere to in professional practice. Many of these deal with elements of the CIA Triad, but they also include aspects such as conflict of interest, client privacy, and so forth.

A standard framework of guidelines for financial accounting.

Geofence

A geofence is a virtual perimeter for a real-world geographic area.

H

Hashing

Able to detect corruption.

High Availability

High availability makes use of shared and pooled resources to maintain a high level of availability and minimize downtime. It typically includes the following capabilities:

- Live recovery
- Automatic migration

> ⓘ Should be used when the goal is to **minimize** the impact of system downtime.

I

Inference

An attack technique that derives sensitive material from an aggregation of innocuous data.

Integrity

The process of ensuring that data is real, accurate, and protected from unauthorized modification.

IT Service Management (ITSM)

The activities that are performed by an organization to design, plan, deliver, operate and control IT services offered to customers.

ITSM makes it possible to:

- Ensure portfolio management, demand management, and financial management are all working together for efficient service delivery to customers and effective charging for services if appropriate
- Involve all the people and systems necessary to create alignment and ultimately success

J-L

J

Jump Server

A jump server, jump host or jump box is a system on a network used to access and manage devices in a separate security zone. A jump server is a hardened and monitored device that spans two dissimilar security zones and provides a controlled means of access between them. The most common example is managing a host in a DMZ from trusted networks or computers.

> ⓘ The difference between a jump server and a bastion host is that a jump server is intended to breach the gap between two security zones and have a gateway to obtain access to something inside of the other security zone. A bastion host is outside of your security zone and will require additional security considerations.

M-O

M

Microsoft Best Practice Analyzers (BPA)

Microsoft Deployment Toolkit (MDT)

Microsoft Deployment Toolkit is a computer program that permits network deployment of Microsoft Windows and Microsoft Office.

Middleware

A term used to describe software that works between an operating system and another application or database of some sort. Typically operates above the transport layer and below the application layer.

N

NERC CIP

The NERC (North American Electric Reliability Corporation) CIP (Critical Infrastructure Plan) is a set of requirements designed to secure the assets required for operating North America's bulk electric system.

Nonrepudiation

The assurance that a specific author actually did create and send a specific item to a specific recipient and that it was successfully received. The sender of the message cannot later credibly deny having sent the message, nor can the recipient credibly claim not to have received it.

O

Outage Duration

Outage duration is the length of time of a documented outage and is expressed as an amount of time (minutes, hours, days).

P-R

P

Portfolio Management

A portfolio consists of all endeavors undertaken by an organization. We are looking to show value for each individual endeavor but also the IT program as a whole. For example, are we receiving the value from cloud services that we're looking for?

Q

Quantum Computing

A theoretical technology which allows superposition of physical states to increase both computing capacity and encryption keyspace.

R

Record

A data structure or collection of information that is retained by an organization for legal, regulator, or business reasons.

Return on Investment

A term used to describe a profitability ratio.

Generally calculated by dividing net profit by net assets.

S-U

S

Sandboxing

- Testing untested or untrusted code
- To better understand if an application is working the way it was intended to work

Separation of Duties

Dictates that one person/entity cannot complete an entire transaction alone.

In the case of encryption, a single entity should not be able to administer the issuing of keys, encrypt the data, and store the keys, because this could lead to a situation where that entity has the ability to access or take encrypted data.

Service Oriented Architecture (SOA)

Views software as a combination of interoperable services, the components of which can be substituted at will.

Severity Assessment

Performed by the customer organization to determine the importance of a particular patch or update.

Shadow IT

Money spent on technology to acquire services without the IT department's dollars or knowledge.

For example, on-demand self service promotes and allows the ability to provision computing resources without human interaction, The consumer can provision resources regardless of location and time. This can be a challenge to purchasing departments as the typical purchasing processes can be avoided.

Silos

The configuration when an enterprise deploys applications in dedicated infrastructure.

Having silos in an enterprise deployment could be a precursor to migrating the environments to the cloud.

Smurf Attack

The Smurf attack is a distributed denial-of-service attack in which large numbers of Internet Control Message Protocol (ICMP) packets with the intended victim’s spoofed source IP are broadcast to a computer network using an IP broadcast address. Most devices on a network will, by default, respond to this by sending a reply to the source IP address. If the number of machines on the network that receive and respond to these packets is very large, the victim’s computer will be flooded with traffic.

T

Trust Zones/Security Zones

A trust zone is a network segment within which data flows relatively freely, whereas data flowing in and out of the trust zone is subject to stronger restrictions. These could be physical, logical, or virtual boundaries around network resources, such as:

- DMZ (semi-trusted)
- Department-specific zones/site-specific zones
- Application-defined zones (such as web application tiers)

Before a cloud provider can implement trust zones, they must undergo threat and vulnerability assessments to determine where their weaknesses are within the environment. This will help to determine where trust zones would be most helpful.

When controls implemented by the virtualization components are deemed to be not strong enough, trust zones can be used to segregate the physical infrastructure.

To protect trust zones:

- Implement granular role-based controls on traffic, users, and assets
- Manage inter-zone communications including between sub-zones
- Enforce policy and regulations
- Protect, detect, and contain

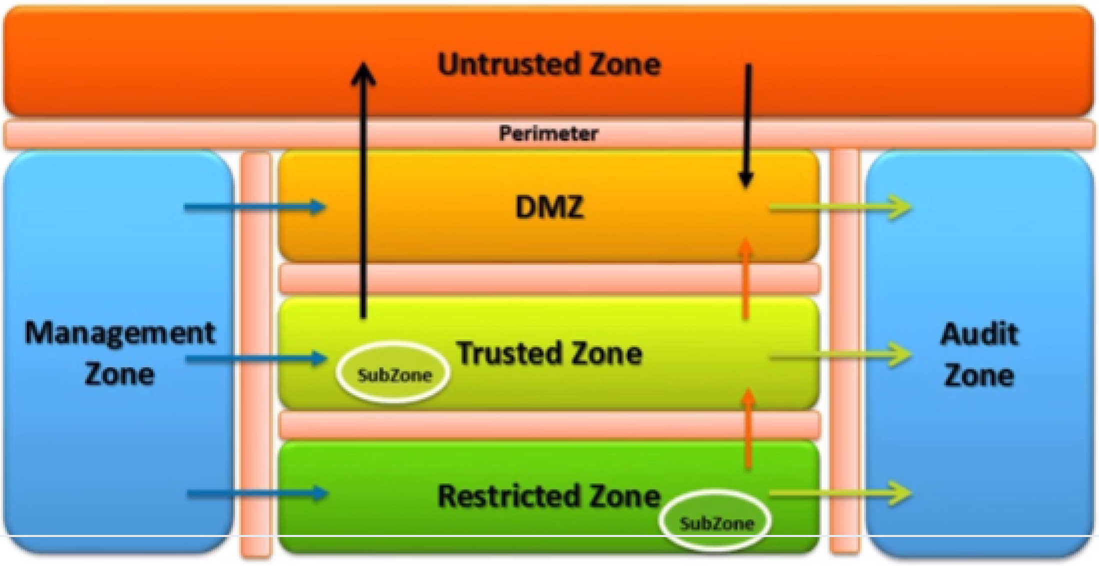

V-Z

W

Wassenaar Arrangement

The **Wassenaar** Arrangement, formally established in July 1996, is a voluntary export control regime whose 42 members exchange information on transfers of conventional weapons and dual-use goods and technologies. This includes import restrictions and data sharing.

Write Once, Read Many (WORM)

A type of long-term storage, meaning it is written to initially and only used for read purposes thereafter.

Standards*

This page contains accurate information but may be missing links to specific standards or formatting may be incorrect.

ISO/IEC

Standard	Category	Description
ISO/IEC 15408:2009	Common Criteria	Information technology - Security techniques - Evaluation criteria for IT security
ISO/IEC 17788:2014	Cloud Computing	Information technology - Cloud computing - Overview and vocabulary
ISO/IEC 17789:2014	Cloud Computing	Information technology - Cloud computing - Reference architecture
ISO/IEC 19086:2016	SLAs	Information technology - Cloud computing - Service level agreement (SLA) framework
ISO/IEC 19941:2017	Cloud Computing	Information technology - Cloud computing - Interoperability and portability
ISO/IEC 19944:2017	Cloud Computing	Information technology - Cloud computing - Cloud services and devices: Data flow, data categories and data use
ISO/IEC 20000-1:2019		Information technology - Service management - Part 1: Service management system requirements
ISO/IEC 27001:2013	ISMS	Information technology - Security techniques - Information security management systems - Requirements

ISO/IEC 27002:2013	Best Practices	Information technology - Security techniques - Code of practice for information security controls
ISO/IEC 27005:2018	Risk	Information technology - Security techniques - Information security risk management
ISO/IEC 27017:2015	Cloud Computing	Information technology - Security techniques - Code of practice for information security controls based on ISO/IEC 27002 for cloud services
ISO/IEC 27018:2014	PII in the Cloud	Information technology - Security techniques - Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors
ISO/IEC 27034:2011	ONF/ANF	Information technology - Security techniques - Application security
ISO/IEC 27036:2014	Supplier Relations	Information technology - Security techniques - Information security for supplier relationships
ISO/IEC 27037:2012	Forensics	Information technology - Security techniques - Guidelines for identification, collection, acquisition and preservation of digital evidence
ISO/IEC 27040:2015	Storage Security	Information technology - Security techniques - Storage security
ISO/IEC 27041:2015	Forensics	Information technology - Security techniques - Guidance on assuring suitability and adequacy of incident investigative method
ISO/IEC 27042:2015	Forensics	Information technology - Security techniques - Guidelines for the analysis and interpretation of digital evidence
ISO/IEC 27043:2015	Forensics	Information technology - Security techniques - Incident investigation principles and processes
ISO/IEC 27050:2019	eDiscovery	Information technology - Electronic discovery

ISO/IEC 28000:2007	Supply Chain	Specification for security management systems for the supply chain
ISO 31000:2009	Risk	Risk management

NIST

Publication	Category	Description
NIST SP 500-292	Cloud Computing	Cloud Computing Reference Architecture
NIST SP 500-293	Cloud Computing	Cloud Computing Technology Roadmap
NIST SP 800-12 Rev. 1		An Introduction to Information Security
NIST SP 800-30	Risk	Guide for Conducting Risk Assessments
NIST SP 800-37	Risk	Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy
NIST SP 800-40 Rev. 3		Guide to Enterprise Patch Management Technologies
NIST SP 800-53		Security and Privacy Controls for Federal Information Systems and Organizations
NIST SP 800-63	Forensics	Digital Identity Guidelines

NIST SP 800-92		Guide to Computer Security Log Management
NIST SP 800-122	Breach Reporting	Guide to Protecting the Confidentiality of Personally Identifiable Information (PII)
NIST SP 800-123		Guide to General Server Security
NIST SP 800-145	Cloud Computing	The NIST Definition of Cloud Computing
NIST SP 800-146		Cloud Computing Synopsis and Recommendations
NIST SP 800-161	Supply Chain	Supply Chain Risk Management Practices for Federal Information Systems and Organizations

(ISC)2

Code of Ethics

Canons

- Protect society, the common good, necessary public trust and confidence, and the infrastructure.
- Act honorably, honestly, justly, responsibly, and legally.
- Provide diligent and competent service to principals.
- Advance and protect the profession.

Exam Outline

The following exam objectives were translated directly from the (ISC)2 Certification Exam Outline. The effective date of these objectives is **August 1, 2019.**

CCSP Examination Weights

Domains	Weight
1 - Cloud Concepts, Architecture and Design	17%
2 - Cloud Data Security	19%
3 - Cloud Platform and Infrastructure Security	17%
4 - Cloud Application Security	17%
5 - Cloud Security Operations	17%
6 - Legal, Risk and Compliance	13%

Domain 1

1.1 Understand Cloud Computing Concepts

- Cloud Computing Definitions
- Cloud Computing Roles
    - Cloud Service Customer
    - Cloud Service Provider
    - Cloud Service Partner
    - Cloud Service Broker
- Key Cloud Computing Characteristics
    - On-Demand Self-Service
    - Broad Network Access
    - Multitenancy
    - Rapid Elasticity
    - Scalability
    - Resource Pooling
    - Measured Service
- Building Block Technologies
    - Virtualization
    - Storage
    - Networking
    - Databases
    - Orchestration/Automation

1.2 Describe Cloud Reference Architecture

- Cloud Computing Activities
- Cloud Service Capabilities
    - Application Capability Types
    - Platform Capability Types
    - Infrastructure Capability Types
- Cloud Service Categories
    - Software as a Service (SaaS)

  - Platform as a Service (PaaS)
  - Infrastructure as a Service (IaaS)
- Cloud Deployment Models
  - Public
  - Private
  - Hybrid
  - Community
- Cloud Shared Considerations
  - Interoperability
  - Portability
  - Reversibility
  - Availability
  - Security
  - Privacy
  - Resiliency
  - Performance
  - Governance
  - Maintenance and Versioning
  - Service Levels and Service Level Agreements
  - Auditability
  - Regulatory
- Impact of Related Technologies
  - Machine Learning
  - Artificial Intelligence
  - Blockchain
  - Internet of Things (IoT)
  - Containers
  - Quantum Computing

1.3 Understand Security Concepts Relevant to Cloud Computing

- Cryptography and Key Management
- Access Control

- Data and Media Sanitization
  - Overwriting
  - Cryptographic Erasure
- Network Security
  - Network Security Groups
- Virtualization Security
  - Hypervisor Security
  - Container Security
- Common Threats

1.4 Understand Design Principles of Secure Cloud Computing

- Cloud Secure Data Lifecycle
- Cloud-based Disaster Recovery (DR) and Business Continuity (BC) planning
- Cost Benefit Analysis
- Functional Security Requirements
  - Portability
  - Interoperability
  - Vendor Lock-in
- Security Considerations for Different Cloud Categories

1.5 Evaluate Cloud Service Providers

- Verification Against Criteria
  - ISO/IEC 27017
  - PCI DSS
- System/subsystem Product Certifications
  - CC
  - FIPS 140-2

Domain 2

2.1 Describe Cloud Data Concepts

- Cloud Data Life Cycle Phases
- Data Dispersion

2.2 Design and Implement Cloud Data Storage Architectures

- Storage Types
    - Long Term
    - Ephemeral
    - Raw-disk
- Threats to Storage Types

2.3 Design and Apply Data Security Technologies and Strategies

- Encryption and Key Management
- Hashing
- Masking
- Tokenization
- Data Loss Prevention (DLP)
- Data Obfuscation
- Data De-identification
    - Anonymization

2.4 Implement Data Discovery

- Structured Data
- Unstructured Data

2.5 Implement Data Classification

- Mapping
- Labeling
- Sensitive Data
    - Protected Health Information (PHI)
    - Personally Identifiable Information (PII)
    - Cardholder Data

2.6 Design and Implement Information Rights Management (IRM)

- Objectives
    - Data Rights
    - Provisioning
    - Access Models
- Appropriate Tools
    - Issuing and Revocation of Certificates

2.7 Plan and Implement Data Retention, Deletion and Archiving Policies

- Data Retention Policies
- Data Deletion Procedures and Mechanisms
- Data Archiving Procedures and Mechanisms
- Legal Hold

2.8 Design and Implement Auditability, Traceability and Accountability of Data Events

- Definition of Event Sources and Requirement of Identity Attribution
- Logging, Storage and Analysis of Data Events
- Chain of Custody and Non-repudiation

Domain 3

3.1 Comprehend Cloud Infrastructure Components

- Physical Environment
- Network and Communications
- Compute
- Virtualization
- Storage
- Management Plane

3.2 Design a Secure Data Center

- Logical Design
  - Tenant Partitioning
  - Access Control
- Physical Design
  - Location
  - Buy or Build
- Environmental Design
  - Heating, Ventilation and Air Conditioning (HVAC)
  - Multi-Vendor Pathway Connectivity

3.3 Analyze Risks Associated with Cloud Infrastructure

- Risk Assessment and Analysis
- Cloud Vulnerabilities, Threats and Attacks
- Virtualization Risks
- Counter-measure Strategies

3.4 Design and Plan Security Controls

- Physical and Environmental Protection
    - On-premise
- System and Communication Protection
- Virtualization Systems Protection
- Identification, Authentication and Authorization in Cloud Infrastructure
- Audit Mechanisms
    - Log Collection
    - Packet Capture

3.5 Plan Disaster Recovery (DR) and Business Continuity (BC)

- Risks Related to the Cloud Environment
- Business Requirements
    - Recovery Time Objective (RTO)
    - Recovery Point Objective (RPO)
    - Recovery Service Level (RSL)
    - Business Continuity/Disaster Recovery Strategy
    - Creation, Implementation and Testing of Plan

Domain 4

4.1 Advocate Training and Awareness for Application Security

- Cloud Development Basics
- Common Pitfalls
- Common Cloud Vulnerabilities

4.2 Describe the Secure Software Development Life Cycle (SDLC) Process

- Business Requirements
- Phases and Methodologies

4.3 Apply the Secure Software Development Life Cycle (SDLC)

- Avoid Common Vulnerabilities During Development
- Cloud-specific Risks
- Quality Assurance
- Threat Modeling
- Software Configuration Management and Versioning

4.4 Apply Cloud Software Assurance and Validation

- Functional Testing
- Security Testing Methodologies

4.5 Use Verified Secure Software

- Approved Application Programming Interfaces (API)
- Supply-chain Management
- Third Party Software Management
- Validated Open Source Software

4.6 Comprehend the Specifics of Cloud Application Architecture

- Supplemental Security Components
  - Web Application Firewall (WAF)
  - Database Activity Monitoring (DAM)
  - Extensible Markup Language (XML) Firewalls
  - Application Programming Interface (API) Gateway
- Cryptography
- Sandboxing
- Application Virtualization and Orchestration

4.7 Design Appropriate Identity and Access Management (IAM) Solutions

- Federated Identity
- Identity Providers
- Single Sign-On (SSO)
- Multifactor Authentication
- Cloud Access Security Broker (CASB)

Domain 5

5.1 Implement and Build Physical and Logical Infrastructure for Cloud Environment

- Hardware Specific Security Configuration Requirements
    - Basic Input Output System (BIOS)
    - Settings for Virtualization and Trusted Platform Module (TPM)
    - Storage Controllers
    - Network Controllers
- Installation and Configuration of Virtualization Management Tools
- Virtual Hardware Specific Security Configuration Requirements
    - Network
    - Storage
    - Memory
    - Central Processing Unit (CPU)
- Installation of Guest Operating System (OS) Virtualization Toolsets

---

5.2 Operate Physical and Logical Infrastructure for Cloud Environment

- Configure Access Control for Local and Remote Access
    - Secure Keyboard Video Mouse (KVM)
    - Console-based Access Mechanisms
    - Remote Desktop Protocol (RDP)
- Secure Network Configuration
    - Virtual Local Area Networks (VLAN)
    - Transport Layer Security (TLS)
    - Dynamic Host Configuration Protocol (DHCP)
    - Domain Name System (DNS)
    - Virtual Private Network (VPN)
- Operating System (OS) Hardening Through the Application of Baselines
    - Windows

  - Linux
  - VMware
- Availability of Standalone Hosts
- Availability of Clustered Hosts
  - Distributed Resource Scheduling (DRS)
  - Dynamic Optimization (DO)
  - Storage Clusters
  - Maintenance Mode
  - High Availability
- Availability of Guest Operating System (OS)

5.3 Manage Physical and Logical Infrastructure for Cloud Environment

- Access Controls for Remote Access
  - Remote Desktop Protocol (RDP)
  - Secure Terminal Access
  - Secure Shell (SSH)
- Operating System (OS) Baseline Compliance Monitoring and Remediation
- Patch Management
- Performance and Capacity Monitoring
  - Network
  - Compute
  - Storage
  - Response Time
- Hardware Monitoring
  - Disk
  - Central Processing Unit (CPU)
  - Fan Speed
  - Temperature
- Configuration of Host and Guest Operating System (OS) Backup and Restore Functions
- Network Security Controls
  - Firewalls
  - Intrusion Detection Systems (IDS)

    - Intrusion Prevention Systems (IPS)
    - Honeypots
    - Vulnerability Assessments
    - Network Security Groups
- Management Plane
    - Scheduling
    - Orchestration
    - Maintenance

5.4 Implement Operational Controls and Standards (e.g., ITIL, ISO/IEC 20000-1)

- Change Management
- Continuity Management
- Information Security Management
- Continual Service Improvement Management
- Incident Management
- Problem Management
- Release Management
- Deployment Management
- Configuration Management
- Service Level Management
- Availability Management
- Capacity Management

5.5 Support Digital Forensics

- Forensic Data Collection Methodologies
- Evidence Management
- Collect, Acquire and Preserve Digital Evidence

5.6 Manage Communication with Relevant Parties

- Vendors
- Customers
- Partners
- Regulators
- Other Stakeholders

5.7 Manage Security Operations

- Security Operations Center (SOC)
- Monitoring of Security Controls
    - Firewalls
    - Intrusion Detection Systems (IDS)
    - Intrusion Prevention Systems (IPS)
    - Honeypots
    - Vulnerability Assessments
    - Network Security Groups
- Log Capture and Analysis
    - Security Information and Event Management (SIEM)
    - Log Management
- Incident Management

Domain 6

6.1 Articulate Legal Requirements and Unique Risks within the Cloud Environment

- Conflicting International Legislation
- Evaluation of Legal Risks Specific to Cloud Computing
- Legal Framework and Guidelines
- eDiscovery
  - ISO/IEC 27050
  - CSA
- Forensics Requirements

6.2 Understand Privacy Issues

- Difference Between Contractual and Regulated Private Data
  - PHI
  - PII
- Country-Specific Legislation Related to Private Data
  - PHI
  - PII
- Jurisdictional Differences in Data Privacy
- Standard Privacy Requirements
  - ISO/IEC 27018
  - GAPP
  - GDPR

6.3 Understand Audit Process, Methodologies, and Required Adaptations for a Cloud Environment

- Internal and External Audit Controls

- Impact of Audit Requirements
- Identify Assurance Challenges of Virtualization and Cloud
- Types of Audit Reports
  - SSAE SOC
  - ISAE
- Restrictions of Audit Scope Statements
  - SSAE
  - ISAE
- Gap Analysis
- Audit Planning
- Internal Information Security Management System (ISMS)
- Internal Information Security Controls System
- Policies
  - Organization
  - Functional
  - Cloud Computing
- Identification and Involvement of Relevant Stakeholders
- Specialized Compliance Requirements for Highly-Regulated Industries
  - NERC/CIP
  - HIPAA
  - PCI
- Impact of Distributed Information Technology (IT) Model
  - Diverse Geographical Locations and Crossing Over Legal Jurisdictions

6.4 Understand Implications of Cloud to Enterprise Risk Management

- Assess Providers Risk Management Programs
  - Controls
  - Methodologies
  - Policies
- Difference Between Data Owner/Controller vs. Data Custodian/Processor
  - Risk Profile
  - Risk Appetite

    - Responsibility
- Regulatory Transparency Requirements
    - Breach Notification
    - SOX
    - GDPR
- Risk Treatment
    - Avoid
    - Modify
    - Share
    - Retain
- Different Risk Frameworks
- Metrics for Risk Management
- Assessment of Risk Environment
    - Service
    - Vendor
    - Infrastructure

6.5 Understand Outsourcing and Cloud Contract Design

- Business Requirements
    - Service Level Agreement (SLA)
    - Master Service Agreement (MSA)
    - Statement of Work (SOW)
- Vendor Management
- Contract Management
    - Right to Audit
    - Metrics
    - Definitions
    - Termination
    - Litigation
    - Assurance
    - Compliance
    - Access to Cloud/Data
    - Cyber Risk Insurance
- Supply-Chain Management

- ISO/IEC 27036

Concepts/Topics

Auditing

Terminology

Acronyms

Acronym	Definition
AUP	Agreed-Upon Procedures

Definitions

AUP

An agreed-upon procedure is a standard a company or client outlines when it hires an external party to perform an audit on a specific test or business process. The procedures, which are called audit standards, are designed and agreed upon by the entity conducting the audit, as well as any appropriate third parties.

The auditor does not provide an opinion; rather, the entities or third parties form their own conclusions based on the report.

Auditability

Auditability is collecting and making available necessary evidence related to the operation and use of the cloud.

Gap Analysis

To create an accurate frame of reference, a gap analysis is conducted. This is like a lightweight audit in that there are generally findings of weaknesses or vulnerabilities, but the purpose is to identify those weaknesses so they can be remediated prior to any actual audit work. It also provides a starting point for those organizations in the early stages of an information system program development, providing them with a clear starting point.

Gap analysis benchmarks and identifies relevant gaps against specified frameworks or standards. This includes reviewing the organization's current position/performance as revealed by an audit against a given standard.

The value of such an assessment is often determined based on what *you did not know* or for an independent resource to communicate to relevant management or senior personnel such risks, as opposed to internal resources saying what you need or should be doing.

Typically, resources or personnel who are not engaged or functioning within the area of scope perform gap analysis. The use of independent or impartial resources is best served to ensure there are no conflicts or favoritism. Perspectives gained from people outside the audit target are invaluable because they may see possibilities and opportunities revealed by the audit, whereas the personnel in the target department may be constrained by habit and tradition.

Purpose

Auditing forms an integral part of effective governance and risk management. It provides both an independent and an objective review of overall adherence or effectiveness of processes and controls. Audits verify compliance by determining whether an organization is following policy. This is not to be confused with verifying whether policy is actually effective. *Testing* is the term used to ensure policy is effective.

Audit Planning

1. Define Objectives

These high-level objectives should interpret the goals and outputs from the audit:

- Document and define the audit objectives.
- Define the audit outputs and format.
- Define the frequency and the audit focus.
- Define the required number of auditors and subject matter experts.
- Ensure alignment with *internal* audit and risk management processes.

---

2. Define Scope

The *organization* is the entity involved in defining the audit scope. The phase includes the following core steps:

- Document the core focus and boundaries of the audit.
- Define the key components of services.
- Define the cloud services to be audited.
- Define the geographic locations that are permitted and required and those that are actually being audited.
- Define the key stages to audit.
- Document the CSP contracts.
- Define the assessment criteria and metrics.
- Document final reporting dates.

---

3. Conduct Audit

When conducting an audit, keep the following issues in mind:

- Adequate staff

- Adequate tools
- Schedule
- Supervision of audit
- Reassessment

---

4. Refine/Lessons Learned

Ensure that previous reviews are adequately analyzed and taken into account, with the view to streamline and obtain maximum value for future audits. To ensure that cloud services auditing is both effective and efficient, several steps should be undertaken either as a standalone activity or as part of a structured framework.

- Ensure that the approach and scope are still relevant.
- Factor in any provider changes that have occurred.
- Ensure that reporting details are sufficient to enable clear, concise, and appropriate business decisions to be made.
- Determine opportunities for reporting improvement and enhancement.
- Ensure that duplication of efforts is minimal (crossover or duplication with other audit and risk efforts).
- Make sure that audit criteria and scope are still accurate, factoring in business changes.
- Have a clear understanding of what levels of information and details can be collected using automated methods and mechanisms.
- Ensure that the right skillsets are available and utilized to provide accurate results and reporting.
- Ensure that the PDCA is also applied to the CSP auditing planning and processing.

These phases may coincide with other audit-related activities and be dependent on organizational structure. They may be structured (often influenced by compliance and regulatory requirements) or reside with a single individual (not recommended).

Audit Responsibilities

Internal Audit

Organizations need ongoing assurances from providers that controls are put in place or are in the process of being identified. Internal audit acts as a **third line** of defense *after* the business or IT functions and risk management functions.

- Audit can provide independent verification of the cloud program's effectiveness giving assurance to the board with regard to the cloud risk exposure.
- Internal audit can also play the role of trusted advisor and proactively work with IT and the business in identifying and addressing the risk associate with third-party providers. This allows a risk-based approach to moving systems to the cloud.

External Audit

An external audit is typically focused on the internal controls over financial reporting. Therefore, the scope of services is usually limited to the IT and business environments that support the financial health of an organization and in most cases doesn't provide specific assurance on cloud risks other than vendor risk considerations on the financial health of the CSP.

BC/DR*

Terminology

Acronyms

Acronym	Definition
BC	Business Continuity
BCM	Business Continuity Management
BCP	Business Continuity Plan
BIA	Business Impact Analysis
DR	Disaster Recovery
DRP	Disaster Recovery Plan
MAD	Maximum Allowable Downtime
MTD	Maximum Tolerable Downtime
RPO	Recovery Point Objective
RSL	Recovery Service Level
RTO	Recovery Time Objective
WRT	Work Recovery Time

Definitions

BC

Business continuity efforts are concerned with maintaining (or "continuing") critical operations during any interruption in service.

Business continuity is defined as the capability of the organization to "continue" delivery of products or services at acceptable predefined levels following a disruptive incident. It focuses primarily on the continuity of *business processes* (as opposed to technical processes).

BCM

Business continuity management is the process by which risks and threats are actively reviewed and managed at set intervals as part of the overall risk management process.

BCM is defined as a holistic management process that identifies potential threats to an organization and the impacts to business operations those threats, if realized, might cause.

It provides a framework for building organizational resilience with the capability of an effective response that safeguards the interests of its key stakeholders, reputation, brand, and value-creating activities.

BCP

The business continuity plan allows a business to plan what it needs to do to ensure that its key products and services **continue to be delivered** in case of a disaster.

Business continuity plans typically outline how to maintain or "continue" *business* operations back to the point of permanent operations. It allows an enterprise to plan what is necessary to ensure that its key products and services will "continue" to be available in the event of a disaster, and that disruption to the business is minimized as much as possible.

> The BCP is not *critical to the continuation of services* in the event of a business interruption. BC, however, is. The BCP is drafted to support BC.

Data Replication

The process of copying data from one location to another. The system works to keep up-to-date copies of its data in the event of a disaster.

DR

Disaster recovery efforts are focused on the resumption of operations after an interruption due to disaster.

Disaster recovery is a subset of business continuity. It is the process of saving data with the sole purpose of being able to recover it in the event of a disaster. Disaster recovery includes backing up systems and IT contingency plans for critical functions and applications.

Disaster recovery focuses on *technology and data policies* (as opposed to business processes).

DRP

The disaster recovery plan allows a business to plan what needs to be done immediately after a disaster to **recover from the event.**

Disaster recovery planning is the process by which suitable plans and measures are taken to ensure that, in the event of a disaster, the business can respond appropriately with the view to recovering critical and essential operations to a state of partial or full level of service in as little time as possible.

DRP is usually part of the BCP and typically tends to be more technical in nature. Addresses what needs to be accomplished during a disaster to restore business processes in order to recover from the event.

Enterprise Cloud Backup

Adds essential features such as archiving and disaster recovery to cloud backup solutions.

Functionality Replication

Used to duplicate processing capability at a secondary location. The secondary location could be with the same CSP or it could be a different CSP. It occurs anytime a needed *function,* including DNS, database, or other functionality is replicated to a CSP's other facilities.

MAD/MTD

A measure of how long it would take for an interruption in service to kill an organization. For example, if a company would fail because it had to halt operations for a week, then it's MAD is one week.

> MAD is measured in **time**.

RTO

The RTO indicates the amount of system downtime defining the total time of the disaster until the business can resume operations.

This is the goal for recovery of operational capability after an interruption in service (i.e., the amount of time it takes to recover). For example, a company might have an MAD of one week, while the company's BCDR plan includes and supports an RTO of six days.

> RTO is measured in **time**. The RTO must be lower than the MAD.

RPO

The RPO indicates the amount of acceptable *data loss* measured in terms of how much data can be lost before the business is too adversely affected.

The point in time at which you would like to restore to. For instance, if an organization performs daily full backups and the BCDR plan includes a goal of resuming critical operations using the last full backup, the RPO would be 24 hours.

Data replication strategies will most affect this metric, as the choice of strategy will determine how much recent data is available for recovery purposes.

> RPO is measured in **time**.

RSL

The recovery service level is a **percentage measurement (0-100%)** of how much computing power is necessary based on the percentage of the production system needed during a disaster.

For example, an RSL of 50% would specify that the DR system would need to operate at a minimum of 50% the performance level of the normal production system.

Server Replication

Concerned more about the processing system rather than the data being replicated.

Storage Replication

Works with a local service to store or archive data to secondary storage using a SAN. This would typically be in the same location.

WRT

The time necessary to *very* restoration of systems once they have been returned to operation.

> ⓘ In DR terms, RTO + WRT < MTD.

Overview

BC/DR protects against the risk of data not being available and the risk that the business processes that it supports are not functional, leading to adverse consequences for the organization. The analysis of this risk leads to the business requirements for BC/DR.

BC/DR starts at risk management since all security decisions are based on risk/risk management. We look at the assets and what they're worth, threats/vulnerabilities, potential for loss versus the cost of the countermeasures. This also helps us identify our critical assets to protect and prioritize. The BIA helps us define our critical assets.

Any BC/DR plan should include the following:

- Required capability and capacity of backup systems
- Trigger events to implement the plan
- Clearly defined roles and responsibilities by name and title
- Clearly defined continuity and recovery procedures
- Notification requirements

Considerations

Notification

Forms

- Telephone call tree rosters
- Website postings
- SMS blasts

Inclusions

- Personnel
- Public
- Regulatory and response agencies

Processes

- Getting the People Out
- Getting the People Out Safely
- Designing for Protection

Continuity

We have to determine what the organization's critical operations are. In a cloud datacenter, that will usually be dictated by the customer contracts and SLAs. The BIA is extremely useful in this portion of the effort, since it informs us which assets would cause the greatest adverse impact if lost or interrupted.

Plan

The authors are big fans of checklists.

- A list of the Items from the Asset Inventory Deemed Critical
- The Circumstances Under Which an Event or Disaster is Declared
- Who is Authorized to Make the Declaration
- Essential Points of Contact
- Detailed Actions, Tasks, and Activities

The plan should be reviewed at least once per year, or as risk dictates.

Kit

There should be a container that holds all the necessary documentation and tools to conduct a proper BC/DR response action.

- A current copy of the plan
- Emergency and backup communication equipment
- Copies of all appropriate network and infrastructure diagrams and architecture
- Copies of all software for creating a clean build and media containing appropriate patches for current versioning
- Emergency contact information
- Documentation tools and equipment
- Emergency essentials (flashlight, water, rations, batteries)

Relocation

- HR and finance should be involved since travel arrangements and payments will be required
- Families should be considered

- Distance needs to be out of impact zone but close enough to not make expenses too high
- Joint operating agreements in the instance that the disaster only affects your organization's campus

Power

- UPS (near-term)
- Generators (short-term)
  - Minimum 12 hours of fuel
  - Should anticipate at least 72 hours

Strategy

- Location
- Data Replication
- Functionality Replication
- Planning, Preparing, and Provisioning
- Failover Capability
- Returning to Normal
- Testing and Acceptance to Production

Risks

- Changes in location
- Maintaining redundancy
- Having proper failover mechanisms
- Having the ability to bring services online quickly
- Having functionality with external services

> ⓘ Budget is *not* a risk since it should be something that is already factored in and accounted for.

Cloud BC/DR

Cloud vs. Traditional

Cloud backup provides many advantages over tape-based backup:

- *Convenience.* As long as you have an Internet connection, data can be backed up as it is saved to disk. Data can be synced across multiple computers so that the data is not only backed up, but it is also instantly shared with other users.
- *Safety.* Local disasters such as fire or flood are no longer concerns.
- *Ease of Recovery.* Online backup systems can be configured to maintain multiple versions of a file. While this may be available with local backup, the ease with which different versions of a file can be restored are superior in the cloud.
- *Ease of Access.* Data can be accessed from anywhere there is an Internet connection.
- *Affordability.* Capital expenditure is reduced as tape drives, libraries, servers, or other hardware is no longer necessary to perform the backup.

Advantages to using a cloud BC/DR include:

- Rapid elasticity
- Broad network connectivity
- On-demand self-service
- Experienced and capable staff
- Measured service

Backup Methodologies

Private Architecture > CSP as BC/DR

The organization maintains its own on-premise IT infrastructure and uses a CSP for BC/DR purposes.

Cloud Operations > Primary CSP as BC/DR

The organization's infrastructure is already hosted in the cloud and they choose to use that same CSP for BC/DR purposes.

In some cases, cloud providers may offer a backup solution as a feature of their service and would ideally be located at another datacenter owned by the provider in case of a local disaster-level event.

Cloud Operations > Third-Party CSP as BC/DR

Regular operations are hosted by the cloud provider, but contingency operations require failover to another cloud provider.

Shared Responsibilities

Declaration

The cloud customer and provider must decide, prior to the contingency, who specifically will be authorized to make decisions for disaster declaration and the explicit process for communicating when it has been made.

Testing

BC/DR testing will have to be coordinated with the cloud provider. This should be planned well in advance of the scheduled testing.

Similarities to Traditional BC/DR

Traditional Hot Site

This would equate to an *active-passive* cloud model.

- In an active-passive deployment, resources are held in a secondary datacenter in standby mode. This would be similar to a hot site in the traditional DR methodology.

BC/DR Planning*

1. Define Scope

2. Gather Requirements

In migrating to a cloud service architecture, your organization will want to review its existing BIA and consider a new BIA, or at least a partial assessment, for cloud-specific concerns and the new risks and opportunities offered by the cloud.

→ BIA /concepts/business/requirements/bia

Potential emergent BIA concerns include, but are not limited to, the following:

- New Dependencies
- Regulatory Failure
- Data Breach/Inadvertent Disclosure
- Vendor Lock-In/Lock-Out

3. Analyze

Will our plan meet the metrics specified in the previous step?

4. Assess

→ Assessing Risk /concepts/risk/risk-management/assessing-risk

5. Design

Should address technical alternatives, procedures, workflow, staff, other business necessities.

---

6. Implement

Implement plan, exercising, assessing, and maintaining the plan.

---

7. Test

Any BCDR plan should be tested at **regular intervals.**

- Tabletop Exercise
- Walk-Through Drill
- Functional Drill
- Full-Interruption

There are two reasons to conduct a test of the organization's recovery from backup in an environment other than the primary production environment:

- You want to approximate contingency conditions, which includes not operating in the primary location. Assuming your facility is not available during contingency operations allows you to better simulate an emergency situation, which adds realism to the test.
- The risk of negative impact to both production and backup is too high. A recovery from backup into the *production* environment carries the risk of failure of both data sets (the production and the backup set).

Tabletop Exercise

Essential participants work together at a scheduled time to describe how they would perform their tasks in a given BCDR scenario.

> ⓘ This has the **least impact** on production of the testing alternatives, but is also the **least thorough.**

Walk-Through Drill

Simulates a disaster scenario but only includes operational and support personnel. It is more complicated than a tabletop exercise. Attendees practice certain functional steps to ensure that they have the knowledge and skills needed to complete them. Acting out the critical steps, recognizing difficulties, and resolving problems is critical for this type of test.

Moves beyond the involvement of a tabletop exercise. Chooses a specific event scenario and applies the BCP to it.

Specific characteristics include:

- Practice and validation of specific functional response capabilities
- Demonstration of knowledge as well as team interaction
- Role playing with simulated response at alternate locations
- Mobilization of the crisis management and response team
- Actual resource mobilization to reinforce the content of the plan

Functional Drill

Involves moving personnel to the recovery site(s) to attempt to establish communications and perform real recovery processing. The drill will help the organization determine whether following the BCP will successfully recover critical systems at an alternate processing site. Because a functional drive fully tests the BCP, all employees are involved. It demonstrates emergency management capabilities and tests procedures for evacuation, medical response, and warnings.

This test is also sometimes considered a "parallel" test. Parallel tests indicate that both the DR site and the production site are processing transactions, which results in heightened risk.

Full-Interruption Test

The entire organization takes part in an unscheduled, unannounced practice scenario, performing their full BCDR activities.

Provides the highest level of simulation, including notification and resource mobilization. A real-life emergency is simulated as closely as possible. It is important to properly plan this type of test to ensure that business operations are not negatively affected. This usually includes processing data and transactions using backup media at the recovery site. All employees must participate in this type of test, and all response teams must be involved.

> ⓘ As this could include system failover and facility evacuation, this test is the most useful for detecting shortcomings in the plan, but it has the greatest impact on productivity.

8. Report

9. Revise

Business

Governance

Terminology

Definitions

Loss of Governance

The risk that the client cedes control to the cloud provider.

Overview

Governance is the system by which the provisioning and usage of cloud services are directed and controlled. Governance defines actions, assigns responsibilities, and verifies performance. Governance includes the policy, process, and internal controls that comprise how an organization is run. Everything from the structures and policies to the leadership and other mechanisms for management.

Governance provides oversight, foundation, direction, and support for the organization as a whole. This includes policies, mission statements, how issues are addressed, and so on. The fact that any issue actually requires addressing is outline by governance. Governance identifies what the organization needs to do to please their stakeholders, prioritize performance vs. security, and so on.

Corporate Governance

Corporate governance is a broad area describing the relationship between the shareholders and other stakeholders in the organization versus the senior management of the corporation.

Third-Party Governance

Like our organization, CSPs also have governance. We do not want to sacrifice our governance in favor of theirs. For example, think of third-party governance for a CSP.

- Who reviews SLAs?
- How is this accomplished?
- Who reviews whether the metrics outlined in the SLA are actually being met?

Governance is making sure we have the right goals and that we have the supporting structure in place to achieve those goals. Do *we* have good third-party governance? Does our CSP have good governance?

Policies

Overview

Policies are one of the foundational elements of a governance and risk management program. They guide the organization, based on standards and guidelines. Policies ensure that the organization is operating within its risk profile. Policies actually define, or are the expression of, the organization.

The designing and implementing of security policies is carried out with input from senior management.

Organizational Policies

Organizational policies take the form of those intended to reduce exposure and minimize risk of financial and data losses, as well as other types of damages such as loss of reputation.

Organizational policies form the basis of functional policies that can reduce the likelihood of the following:

- Financial loss
- Irretrievable loss of data
- Reputational damage
- Regulatory and legal consequences
- Misuse and abuse of systems and resources

Functional Policies

Functional policies are particularly useful for organizations that have a well-engrained and fully operational ISMS:

- Information security policy
- Information technology policy

- Data classification policy
- Acceptable usage policy
- Network security policy
- Internet use policy
- Email use policy
- Password policy
- Virus and spam policy
- Software security policy
- Data backup policy
- Disaster recovery (DR) policy
- Remote access policy
- Segregation of duties policy
- Third-party access policy
- Incident response and management policy
- Human resources security policy
- Employee background checks
- Legal compliance guidelines

Cloud Policies

As part of the review of cloud services, either during the development of the cloud strategy or during vendor reviews and discussions, the details and requirements should be expanded to compare or assess the required criteria *as per existing policies.*

- Password policies
- Remote access
- Encryption
- Third-party access
- Segregation of duties
- Incident management
- Data backup

All of these policies are an expression of management's strategic goals and objective with regard to managing and maintaining the risk profile of the organization.

Contracts

Terminology

Definitions

Memorandum of Agreement/Understanding (MOA/MOU)

A *nonbinding* agreement between two or more parties outlining the terms and details of an understanding, including each parties' requirements and responsibilities.

Letter of Intent (LOI)

A LOI outlines the general plans of an agreement between two or more parties before a legal agreement is finalized.

Overview

The contract will spell out all the terms of the agreement: what each party is responsible for, what form the services will take, how issues will be resolved, and so on. The contract will also state what the penalties are (usually financial) when the cloud provider fails to meet the SLA for a given period.

> The provider will ensure the customer has constant, uninterrupted access to the Customer's data storage resources. Customer's monthly fee will be waived for any period following a calendar month in which any service level has not been attained by the Provider.

The contract defines all the terms of an agreement with a cloud provider, including what each party is responsible for, what form the services will take, how issues will be resolved, and so on.

> ⓘ Contract management is all about third-party governance.

From a contractual, regulated, and PII perspective, the following should be reviewed and fully understood with regard to any hosting contracts (along with other overarching components within an SLA):

- *Scope of processing.* Clear understanding of the permissible types of data processing should be provided. The specifications should also list the purpose for which the data can be processed or utilized.
- *Use of subcontractors.* Understanding where any processing, transmission, storage, or use of information will occur.
- *Deletion of data.* Where the business operations no longer require information to be retained for a specific purpose, the deletion of information should occur in line with the organizations data retention policies and standards.
- *Appropriate or required data security controls.* Where processing, transmission, or storage of data and resources is outsourced, the same level of security controls should be required for any entry's contracting or subcontracting services.
- *Locations of data.* Where information is being stored, processed, and transmitted in the event of daily actions or failover.
- *Return or restitution of data.* For both contractors and subcontractors where a contract is terminated, the timely and orderly return of data has to be required both contractually and within the SLA.
- *Right to audit subcontractors.* Right to audit clauses should allow for the organization owning the data (not possessing) to audit or engage the services of an independent party to ensure that contractual and regulatory requirements are being satisfied by either the contractor or the subcontractor.

At a minimum, you should ensure that your cloud service contract states that you must be notified in the event of a subpoena or other similar actions to ensure that you do not lose access to your data as a result of lawsuits.

Immediate notification from a CSP to a customer should be required for all security related events. This is especially true for security breaches. This should be explicitly addressed in the service contract between the client and the CSP and should be non-negotiable.

Operations Management

Information Security Management

Overview

Information security management **deals with the CIA of data.**

- Security management
- Security policy
- Information security organization
- Asset management
- Human resources security
- Physical and environmental security
- Communications and operations management
- Access control
- Information systems acquisition, development, and maintenance
- Provider and customer responsibilities

> ⓘ This section looks much like what you'd see as the foundation for ISO 27001 as part of the ISMS (Information Security Management System or Program).

Configuration Management

Terminology

Acronyms

Acronym	Definition
CI	Configuration Item

Definitions

CI

Can be applied to anything designated for the application of the elements of configuration management and treated as a single entity in the configuration-management system. Examples of CI types include:

- Hardware/Devices
- Software/Applications
- Communications/Networks
- Location
- Database
- Service
- Documentation
- People (Staff and Contractors)

Overview

Configuration management **deals with the documentation of processes.**

Configuration management aims to maintain information about CIs required to deliver an IT service, including their relationships. Configuration management occurs when the configuration of an item, such as a network device, must be changed.

The process should include policies and procedures for each of the following:

- The development and implementation of new configurations that should apply to the hardware and software configurations of the cloud environment
- Quality evaluation of configuration changes and compliance with established security baselines
- Changing systems, including testing and deployment procedures, that should include adequate oversight of all configuration changes
- The *prevention* of any unauthorized changes in system configurations

Operational Relationships

In relation to **availability management:**

- If an existing configuration were to have negative impacts on system availability, they would have to be identified, monitored, and remediated as per the existing SLAs for the services and systems affected.

In relation to **change management:**

- Change management has to approve modifications to all production systems prior to them taking place. There should never be a change that is allowed to take place to a CI in a production system unless change management has approved the change first.
- Configuration management aims to *prevent unauthorized changes* whereas change management aims to *allow changes through a formal approval process.*

In relation to **release and deployment management:**

- Once the release is officially live in the production environment, the existing configurations for all systems and infrastructure affected by the release have to be updated to accurately reflect their new running configurations and status within the configuration management database (CMDB).

Change Management

Overview

Change management **deals with ensuring you are following the appropriate processes.**

Change management manages all changes to CIs, including any devices. All changes must be tested and formally approved prior to deployment in the live environment. It also deals with modifications to the *network*, such as the acquisition and deployment of new systems and components and the disposal of those taken out of service.

Change management is an approach that allows organizations to manage and control the impact of change through a structured process. The primary goal of change management is to create and implement a series of processes that allow changes to the scope of a project to be formally ***introduced*** and approved.

Baselines

Baselining

Baselining creates a general-purpose map of the network and systems, based on the required functionality as well as security (such as required security controls)

The baseline should be a reflection of the risk appetite of the organization and provide the optimum balance of security and operational functionality.

Deviations and Exceptions

Deviations and exceptions to the baseline should be documented.

Roles and Processes

The change management process should be formalized in the organization's governance:

- Composition of the change management board
- The process

- Documentation requirements
- Instructions for requesting exceptions
- Assignment of change management tasks (validation scanning, analysis, deviation notification)
- Procedures for addressing deviations upon detection
- Enforcement measures and responsibility

The process has two forms:

- One that will occur once
- One that will occur repetitiously

The initial process should look something like this:

1. *Full asset inventory.* May be aided by information pulled from other sources such as the BIA.
2. *Codification of the baseline.* Can use risk management framework, enterprise and security architecture, and so on.
3. *Secure baseline build.*
4. *Deployment of new assets.*

In the normal operations mode, the change management process is slightly different:

1. *Change management board meeting.*
2. *Change management testing.*
3. *Deployment.*
4. *Documentation.*

Objectives

- Respond to a customer's changing business requirements while maximizing value and reducing incidents, disruption, and rework.
- Respond to business and IT requests for change that aligns services with business needs.
- Ensure that changes are recorded and evaluated.
- Ensure that authorized changes are prioritized, planned, tested, implemented, documented, and reviewed in a controlled manner.

- Ensure that all changes to CIs are recorded in the configuration management system.
- Optimize overall business risk. It is often correct to minimize business risk, but sometimes it is appropriate to knowingly accept a risk because of the potential benefit.

Process

A change-management process focused on the cloud should include policies and procedures for each of the following:

- The development and acquisition of new infrastructure and software
- Quality evaluation of new software and compliance with established security baselines
- Changing systems, including testing and deployment procedures; they should include adequate oversight of all changes
- Preventing the unauthorized installation of software and hardware

---

Other Functions

Patch Management

Patch management is a function of change management. It is the process of identifying, acquiring, installing, and verifying patches for products and systems. Patches correct security and functionality problems in software and firmware.

Applicability Assessment

Performed by the customer organization to determine whether a particular patch or update applies to the customer's deployment.

Operational Relationships

In relation to **configuration management:**

- Change management has to approve modifications to all production systems prior to them taking place. There should never be a change that is allowed to take place to a CI in a production system unless change management has approved the change first.
- Configuration management aims to ***prevent*** *unauthorized changes* whereas change management aims to ***allow*** *changes through a formal approval process.*

In relation to **problem management:**

- Problem management works with change management to ensure fixes are properly tested and approved. Once the fix is approved, change management will deploy the fix, and the fix will be marked as completed in both the change management and problem management processes.

In relation to **release and deployment management:**

- Change management must approve any activities that release and deployment management will be engaging in prior to the release.
    - Change management must approve requests to carry out the release, and then deployment management can schedule and execute the release.

In relation to **service level management:**

- Change management must approve changes to all SLAs as well as ensure that the legal function has a chance to review them and offer guidance and direction on the nature and language of the proposed changes prior to them taking place. There should never be a change that is allowed to take place to an SLA that governs a production system unless change management has approved the change first.

Incident Management*

Terminology

Acronyms

Acronym	Definition
IMT	Incident Management Team
IRT	Incident Response Team

Definitions

Event

According to the ITIL framework, an event is defined as a **change of state** that has significance for the management of an IT service or other CI. This could be any unscheduled adverse impact to the operating environment.

Events are anything that occur in the IT framework. As a result, the term can also be used to mean an alert or notification created by an IT service, CI, or monitoring tool.

Events often require IT operations staff to take actions and lead to incidents being logged.

> ✓ **Fact.** Not all events are incidents, but all incidents are events.

Incident

According to the ITIL framework, an incident is defined as an unplanned interruption to an IT service or a reduction in the quality of an IT service.

Essentially, incidents are unscheduled events.

> ⓘ An event is distinguished form a disaster by the **duration** of impact. We consider events impact to last 3 days or less.

> ✓ **Fact.** Not all events are incidents, but all incidents are events.

Overview

Incident management **deals with minimizing the impact to the business.**

Incident management describes the activities of an organization to identify, analyze, and correct hazards to prevent a future reoccurrence. Incident management is typically involved in an initial attack and resolution of the attack. Identifying the *root cause* of the attack and deploying a fix to a known error is part of problem management.

Within a structured organization, an IRT or IMT typically addresses these types of incidents.

> ⓘ Incident management should be focused on the identification, classification, investigation, and resolution of an incident, with the ultimate goal of **returning the effected systems to normal as soon as possible.**

Purpose

- Restore normal service operation as quickly as possible
- Minimize the adverse impact on business operations
- Ensure service quality and availability are maintained

Objectives

- Ensure standardized methods and procedures are used for efficient and prompt response, analysis, documentation of ongoing management, and reporting of incidents
- Increase visibility and communication of incidents to business and IT support staff
- Enhance business perception of IT by using a professional approach in quickly resolving and communicating incidents when they occur
- Align incident management activities with those of the business
- Maintain user satisfaction

Plan

Should include:

- Definitions of an incident by service type or offering
- Customer and provider roles and responsibilities for an incident
- Incident management process from detection to resolution
- Response requirements
- Media coordination
- Legal and regulatory requirements such as data breach notification

Incident Prioritization

Incident prioritization is made up of the following items (displayed in a matrix of 1-5 where 1 is highest and 5 is lowest).

Impact

Effect on the business

Urgency

Extent to which the resolution can be delayed

Priority

```
Urgency * Impact
```

Process

1. Incident Occurs
2. Incident is Reported
3. Incident is Classified
4. Investigate and Collect Data
5. Resolution
6. Approvals
7. Implement Changes
8. Review
9. Reports

Problem Management

Terminology

Definitions

Problem

A problem is the unknown root cause of one or more *incidents,* often identified as a result of multiple similar incidents.

Known Error

A known error is an identified *root cause* of a *problem*.

Workaround

A workaround is a *temporary* way of overcoming technical difficulties (such as incidents or problems).

Overview

The objective of problem management is to *minimize the impact of problems* on the organization by identifying the **root cause** of the problem at hand. It aims to prevent reoccurrence of a problem.

Operational Relationships

In relation to **change management:**

- Problem management works with change management to ensure fixes are properly tested and approved. Once the fix is approved, change management will deploy the fix, and the fix will be marked as completed in both the change management and problem management processes.

In relation to **release and deployment management:**

- If anything were to go wrong with the release, incident and problem management would need to be involved to fix whatever went wrong (if it happened once, this would be more of an incident; if it happened several times, this would be more of a problem).

Release and Deployment Management

Overview

In one sentence, this section deals with procedures for testing, certification, accreditation, and moving to operations (including monitoring and upgrades).

Release and deployment management aims to plan, schedule, and control the movement of releases to test and live environment. The primary goal of release and deployment management is to ensure that the integrity of the live environment is protected and that the correct components are released.

Deployment management plans, schedules, and controls the movement of releases to test and live environments. Deployment management is used when a new software version needs to be released or a new system is being deployed.

Objectives

- Define and agree upon deployment plans
- Create and test release packages
- Ensure the integrity of release packages
- Record and track all release packages in the Definitive Media Library (DML)
- Manage stakeholders
- Check delivery of utility and warranty (utility + warranty = value in the mind of the customer)
    - Utility is the functionality offered by a product or service to meet a specific need; it's what the service does.
    - Warranty is the assurance that a product or service will meet agreed-upon requirements (SLA); it's how the service is delivered.
- Manage risks
- Ensure knowledge transfer

JML

Joiners

Are users and software licenses up to date?

Movers

Are software licenses being automatically re-harvested?

Leavers

How much are you overpaying for what you no longer need?

Operational Relationships

In relation to **change management:**

- Change management must approve any activities that release and deployment management will be engaging in prior to the release.
    - Change management must approve requests to carry out the release, and then deployment management can schedule and execute the release.

In relation to **problem management:**

- If anything were to go wrong with the release, incident and problem management would need to be involved to fix whatever went wrong (if it happened once, this would be more of an incident; if it happened several times, this would be more of a problem).

In relation to **configuration management:**

- Once the release is officially live in the production environment, the existing configurations for all systems and infrastructure affected by the release have to be updated to accurately reflect their new running configurations and status within the configuration management database (CMDB).

In relation to **availability management:**

- If the release were not to go as planned, any negative impacts on system availability would have to be identified, monitored, and remediated as per the existing SLAs for the services and systems affected.

Service Level Management

Terminology

Acronyms

Acronym	Definition
UC	Underpinning Contract

Definitions

Efficient Service Management

Are we doing the right things properly and are we getting the benefits? Are we going about it in the most efficient manner to achieve the most benefits?

Overview

Service level management **deals with third-party governance.**

Service-level management aims to negotiate agreements with various parties and to design services in accordance with the agreed-upon service-level targets. Typical negotiated agreements include:

- UCs are external contracts negotiated **between the organization and vendors or suppliers.**

Operational Relationships

In relation to **change management:**

- Change management must approve changes to all SLAs as well as ensure that the legal function has a chance to review them and offer guidance and direction on the nature and language of the proposed changes prior to them taking place. There should never be a change that is allowed to take place to an SLA that governs a production system unless change management has approved the change first.

SLAs

Terminology

Acronyms

Acronym	Definition
SLA	Service-Level Agreement

Definitions

Service Levels

Service levels indicate the minimum expected performance.

Overview

The SLA will set specific, quantified goals for these services and their provision over a certain timeframe.

One use of the SLA is to determine whether a customer is actually receiving the services outlined in the SLA. An independent third-party can objectively affirm whether the requirements outlined in the SLA are being met.

The SLA should contain elements of the contract that can be subject to discrete, objective, repeatable, numeric metrics. For example:

- There will be no interruption of connectivity to data storage longer than three (3) seconds per calendar month.

SLAs are negotiated with **customers.**

- Assessment of risk environment
- Risk profile
- Risk appetite
- Responsibilities
- Regulatory requirements
- Risk mitigation
- Risk frameworks

An SLA typically includes items that can be measured quantitatively:

- *Availability*: What is the uptime?
- *Performance*: What is the response time?
- *Privacy of the data*: Does the data need to be encrypted?
- *Logging and reporting*: What needs to be logged?
- *DR expectations*: What is the RTO, RPO, and MTD?
- *Location of the data*: What regulations may dictate where the data is stored?
- *Data format and structure*: How will data be stored, saved, and retained?
- *Portability of the data*: Will data be maintained such that it can be moved among CSPs?
- *Identification and problem resolution*: Who is called when there are problems?
- *Change-management process*: What is the process for changes to be submitted and verified?
- *Dispute-mediation process*: Exit strategy with expectations on the provider to ensure a smooth transition.

Having an SLA separate from the contract allows the company to revise the SLA without revising the contract. The contract would then refer to the SLA. The term for the contract could be multiple 2-year periods while the SLA would be reviewed quarterly, for example. This reduces the need to engage legal for contract review.

> I like to consider the SLA as the "performance contract" since it is more concerned with the performance-related aspects than the contract itself.

PLAs

Terminology

Acronyms

Acronym	Definition
PLA	Privacy-Level Agreement

Overview

PLAs are an agreement whereby the cloud provider states the level and types of personal data protection(s) in place. A PLA would require the cloud provider to document expectations for the cloud customer's data security, which could be an explicit admission of liability for CSPs, which is why this typically isn't common.

The CSA allows for four items in a PLA:

1. A PLA requires the CSP to clearly describe the level of privacy and data protection that is undertaken to maintain with respect to the relevant data processing.
2. The CSP must adopt a common structure or outline of the PLAs in which it can promote a powerful global industry standard.
3. A PLA should offer a clear and effective way to communicate the level of data protection offered by a CSP.
4. The eventual goal for a PLA is to provide customers with a tool to baseline personal data protection legal requirements across an environment and to evaluate the data protection offered by different CSPs.

Characteristics

- Provides a clear and effective way to communicate the level of personal data protection offered by a service provider

- Works as a tool to assess the level of a service provider's compliance with data protection legislative requirements and leading practices
- Provides a way to offer contractual protection against possible financial damages due to lack of compliance

OLAs

Terminology

Acronyms

Acronym	Definition
OLA	Operational-Level Agreement

Overview

OLAs are SLAs negotiated between internal business units **within the enterprise.** OLAs are a contract that defines how various IT groups within a company plan to deliver a server or set of services.

For example, an IT department may guarantee system uptime to a separate group within the same organization.

Availability Management

Overview

Availability management aims to define, analyze, plan, measure, and improve all aspects of the availability of IT services. It is also responsibility for ensuring that all IT infrastructure, processes, tools, roles, and so on, are appropriate for the agreed-upon availability targets.

> ⓘ Most virtualization platforms allow for the management of system availability and can act in the event of a system outage (such as vMotion).

Operational Relationships

In relation to **release and deployment management:**

- If the release were not to go as planned, any negative impacts on system availability would have to be identified, monitored, and remediated as per the existing SLAs for the services and systems affected.

Capacity Management

Overview

Capacity management **deals with scalability and elasticity.**

Capacity management is focused on ensuring that the business IT infrastructure is adequately provisioned to deliver the agreed service-level targets in a timely and cost-effective manner.

Capacity management considers all resources required to deliver IT services within the scope of the defined business requirements. System capacity must be monitored and thresholds must be set to prevent systems from reaching an over-capacity situation.

Business Continuity Management

Terminology

Acronyms

Acronym	Definition
BCM	Business Continuity Management

Overview

Business continuity management **deals with planning, redundancy, criticality analysis, and sustaining the long-term operations of the business.**

BCM is focused on the planning steps that businesses engage in to ensure that their mission-critical systems are able to be restored to service following a disaster or service interruption event.

To focus BCM activities correctly, a prioritized ranking or listing of systems and services must be created and maintained. This is accomplished through the BIA, which identifies and produces a prioritized listing of systems and services critical to the normal functioning of the business.

Continual Service Improvement Management

Overview

Continual service improvement management **deals with identifying how the organization can perform more efficiently.**

Metrics on all services and processes should be collected and analyzed to find areas of improvement using a formal process. You can use various tools and standards to monitor performance. One example is the ITIL framework.

Business Requirements

Terminology

Definitions

Business Requirement

A business requirement is an operational driver for decision making and input for risk management.

Business Rules

Business rules are lists of statements that tell you whether you may or may not do anything or that give you the criteria and conditions for making a decision.

Scoping

Refers to including only departments or business units impacted by any (cloud) engagement.

Overview

Security activities actually hinder business efficiency (because generally the more secure something is, be it a device or a process, the less efficient it will be). This is why the business needs of the organization drive security decisions, and not the other way around.

1. Gather business requirements
2. Perform business impact analysis

Involvement and Alignment

- Involvement from business units and alignment of IT processes/services with those units.

- The success of our services is based on the success of the business.
- Are we delivering the value we're expected to?

Quantifying Benefits and Opportunity Cost

- Reduction in Capital Expenditure
- Reduction in Personnel Costs
- Reduction in Operational Costs
- Transferring Some Regulatory Costs
- Reduction in Costs for Data Archival/Backup Services

Types of Requirements

Functional Requirements

Those performance aspects of a device, process, or employee that are necessary for the business task to be accomplished. Example: A salesperson in the field must be able to connect to the organization's network remotely.

Nonfunctional Requirements

Those aspects of a device, process, or employee that are not necessary for accomplishing a business task but are desired or expected. For example: The salesperson's remote connection must be secure.

BIA

Overview

The business impact analysis gathers asset valuation information that is beneficial for risk analysis and selection of security controls, and criticality information that helps in BC/DR planning by letting the organization understand which systems, data, and personnel are necessary to continuously maintain. It is an assessment of the priorities given to each asset and process within the organization.

It is based off the assumption that there are certain things the business needs to know in order to decide how they will handle risks within our organization.

The BIA is the basis of almost everything done with the organization. Business processes are based on the criticality of assets identified as a result of the BIA. A proper analysis should consider the effect (impact) any harm or loss of each asset might mean to the organization overall. Assets can be tangible or intangible.

Goals

- *Identify critical business processes and dependencies.* Such as determining RPOs and RTOs.
- *Identify risks and threats.* Such as CSP failure.
- *Identify requirements.* These may come from senior management, regulations, or a combination of both.

Analysis

Inventory of Assets

We must understand what assets exist before we begin to determine their value.

Valuation of Assets

We determine a value for every asset (usually in terms of dollars), what it would cost the organization if we lost that asset (either temporarily or permanently), what it would cost to replace or repair that asset, and any alternate methods for dealing with that loss.

A proper analysis should consider the effect ("impact") any harm or loss of each asset might mean to the organization overall. Special care should be paid to identifying critical paths and single points of failure.

Determination of Criticality

Criticality denotes those aspects of the organization without which the organization could not operate or exist.

Tangible Assets

The organization is a rental car company; cars are critical to its operations-if it has no cars to rent to customers, it can't do business.

Intangible Assets

The organization is a music production firm; music is the intellectual property of the company-if the ownership of the music is compromised (for instance, if the copyright is challenged and the company loses ownership, or the encryption protecting the music files is removed and the music can be copied without protection), the company has nothing of value and will not survive.

Processes

The organization is a fast-food restaurant noted for its speed; the process of taking orders, preparing and delivering food, and taking payment is critical to its operations-if the restaurant cannot complete the process for some reason (for instance, the registers fail so that the restaurant cannot accept payment), the restaurant cannot function.

Data Pathways

The organization is an international shipping line; matching orders to cargo carriers is critical to its operations. If the company cannot complete its logistical coordinationassigning cargo requests to carriers with sufficient capacity-it cannot provide its services, and will not survive.

Personnel

The organization is a surgical provider; the surgeon is critical to the existence of the company-if the surgeon cannot operate, there is no company.

Business Responsibilities*

Responsibilities in the Cloud

Shared Responsibilities

Shared Administration

Cloud Provider Responsibilities

The Physical Plant

- Buy versus lease.
- Secure Hardware Components
- Manage hardware configuration.
- Set hardware to log events and incidents.
- Determine compute component composition by customer need.
- Configure secure remote administrative access.

Secure Logical Framework

- Installation of Virtual OSs
- Secure Configuration of Various Virtualized Elements

Secure Networking

- Firewalls
- IDS/IPS
- Honeypots
- Vulnerability Assessments
- Communication Protection
    - Encryption
    - Virtual Private Networks (VPNs)
    - Strong Authentication

Mapping and Selection of Controls

Data Access

Customer Directly Administers Access

- User contacts the organization's administrator to request an account
- The administrator confirms the account necessity and permissions (perhaps with the user's manager or HR)
- The administrator logs onto the cloud system and makes the necessary modification to the ACL(s)
- The administrator notifies the user, and the user now has access

Provider Administers Access on Behalf of the customer

- User contacts the provider's administrator to request an account
- The administrator confirms the account necessity and permissions with the cloud customer POC
- The customer POC confirms the account necessity and permissions with the user's office/HR
- The customer POC passes verification to the cloud administrator
- The cloud administrator makes the necessary modification to the ACL(s)
- The administrator notifies the user, and the user now has access

Third-Party (CASB) Administers Access on Behalf of the Customer

- User contacts the CASB to request an account
- The CASB confirms the account necessity and permissions with the cloud customer POC
- The customer POC confirms the account necessity and permissions with the user's office/HR
- The customer POC passes verification to the CASB
- The CASB logs onto the cloud system and makes the necessary modification to the ACL(s)
- The CASB notifies the user, and the user now has access

Physical Access

Audits

- SOC

Policy

Monitoring and Testing

Cloud

Terminology

Definitions

Apache CloudStack

Apache CloudStack creates, manages, and deploys clouds. It is an open-source application. It is software utilized to deploy and manage large networks of virtual machines, that need to be highly available.

It is deployed as a highly scalable IaaS computing platform.

Capacity

Capacity is the measurement of the degree to which the cloud can support or provide service.

Cloud App

A cloud application is a software application accessed via the internet and which may include an agent or applet installed locally on the user's device.

Cloud Application Management for Platforms (CAMP)

CAMP is a specification geared towards PaaS. The specification indicates that for consumers this will provider for "portability between clouds." This is accomplished by standardization of the management API, which allows use cases for deploying, stopping, starting, and updating applications.

Cloud Appropriateness

A pitfall in which application development in cloud environments is much different since applications are built on web service frameworks and typically do not support legacy systems and programming languages.

Cloud Bursting

Augmenting internal, private datacenter capabilities with managed services during times of increased demand.

The organization might have datacenter assets it owns, but it can't handle the increased demand during times of elevated need (crisis situations, heavy holiday shopping periods, and so on), so it rents the additional capacity as needed from an external cloud provider.

Cloud Deployment

Deals with *which type* of cloud you will be leveraging: private, public, community, or hybrid.

Cloud Migration

Cloud migration is the process of transitioning all or part of a company's data, applications, and services from onsite premises to the cloud, where the information can be provided over the Internet on an on-demand basis. The steps in a cloud migration include:

- Choosing a provider
- Planning
- Migrating
- Testing and validation
- Maintaining

Concerned with the actual *movement* of the data, application, and services to the cloud.

Cloud Provisioning

A term used to describe the deployment of a company's cloud computing strategy, which typically first involves selecting which applications and services will reside in the public cloud and which will remain on-site behind the firewall or in the private cloud.

Cloud Standards Customer Council (CSCC)

The Cloud Standards Customer Council (CSCC) is an end-user advocacy group. It is dedicated to accelerating cloud's successful adoption, as well as to drilling down into the standards, security, and interoperability issues that surround the transition to the cloud.

Cloud Testing

Cloud testing is load and performance testing conducted on the cloud applications and services, to ensure optimal performance and scalability under a wide variety of conditions.

Cloud Washing

The act of adding the name "cloud" to a non-cloud service and selling it as a cloud solution.

Compliance as a Service (CompaaS/CaaS)

Includes a variety of compliance services such as data encryption, disaster recovery, reporting, and vulnerability scanning.

Data Science as a Service (DSaaS)

Involves an outside company providing advanced analytics applications (gathered using data science) to corporate clients for their business use.

Dynamic Optimization

The process in which cloud environments are constantly monitored and maintained to ensure that the resources are available when needed and that nodes share the load equally so that one node doesn't become overloaded.

Elasticity

The flexibility of allocating resources as needed for immediate usage, instead of purchasing resources according to other variables.

Eucalyptus

Eucalyptus is a paid and open-source computer software building AWS-compatible private and hybrid cloud computing environments.

Multitenancy

Multitenancy refers to the notion of hosting multiple cloud tenants on a single host while sharing resources.

Networking as a Service (NaaS)

Includes network services from third-parties to customers that do not want to build their own networking infrastructure.

Scalability

New computing resources can be assigned and allocated without any significant additional capital investment on the part of the cloud provider, and at an incremental cost to the cloud customer.

Simplicity

Usage and administration of cloud services ought to be transparent to cloud customers and users; from their perspective, a digital data service is paid for and can be used, with very little additional input other than what is necessary to perform their duties.

Sprawl

A phenomenon that occurs when the number of VMs on a network reaches a point where the administrators can no longer manage them effectively.

Sprawl is a virtualization risk that occurs when the amount of content grows to such a degree that management is near impossible.

To prevent sprawl, the administrator should define and enforce a process for the deployment of VMs and create a library of standardized VM image files.

Tenancy Separation

Tenants, while running on the same host, are maintained separately in their virtual environments. This is known as *tenancy separation.*

Vertical Cloud Computing

Refers to the optimization of cloud computing and cloud services for a particular vertical (e.g., a specific industry) or specific-use application.

Cloud Computing

Overview

Cloud computing, as defined by NIST in SP 800-145 is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

There are three pillars of cloud services:

- **Processing** data (CPU)
- **Moving** data (networking)
- **Preserving** data (storage)

To determine the effectiveness of a cloud security program:

1. Business/Information Technology Functions
2. Risk Management Functions
3. Internal Audit

Infrastructure

Building block technologies of cloud services include:

- Servers
- Virtualization
- Storage
- Network
- Management
- Security
- Backup and recovery
- Infrastructure systems (converged infrastructures)
- Databases
- Memory (RAM)
- Processing (CPU)

Add-on services that are not considered building blocks might include:

- Encryption
- SSO

Cloud environments do not have a static definition for the perimeter. The perimeter could be the demarcation point, it could be the borders around the individual customers services, it could be nearly no perimeter at all. The standard definition of what constitutes a network perimeter takes on different definitions and deployment models.

Cloud Roles

Cloud Access Security Broker (CASB)

Typically a third-party entity or company that looks to extend or enhance value to multiple customers of cloud-based services through relationships with multiple cloud service providers. It acts as a liaison between cloud services customers and cloud service providers, selecting the best provider for each customer and monitoring the services.

A third-party entity offering independent identity and access management (IAM) and key management services to CSPs and cloud customers, often as an intermediary. This can take the form of a variety of services, including SSO, certificate management, and cryptographic key escrow.

Cloud Administrator

Cloud Application Architect

Responsible for adapting, porting, or deploying an application to a target cloud environment.

The cloud application architect is responsible for prepping and deploying an application to the cloud environment. This person will work alongside development and implementation resources to ensure that the performance, reliability, and security of an application are sustained over the lifecycle of the application. Knowledge and application of the phases of the SLDC, including assessment, verification, and testing are required.

Cloud Architect

Responsible for an organization's cloud computing strategy.

- Responsibilities include determining when and how a private cloud meets the policies and needs of an organization's strategic goals.
- Architecting the private cloud, designing/deploying hybrid cloud solutions.
- Understand and evaluate technologies, vendors, services, to support private and hybrid clouds.

The cloud architect is responsible for an organization's cloud computing strategy. Part of the responsibilities include determining when and how a private cloud meets the policies and needs of an organization's strategic goals. Further, responsibilities include:

- Architecting the private cloud
- Designing/deploying hybrid cloud solutions

The architect has a key role in understanding and evaluating technologies, vendors, services, to support private and hybrid cloud. The architect will be required to build relationships between customers and team members.

Cloud Carrier

The cloud carrier is the intermediary that provides connectivity and transport of cloud services between the CSPs and the cloud service consumers.

The intermediary that provides connectivity and transport of cloud services between the CSPs and the cloud service consumers.

Cloud Data Architect

Ensures the various storage types and mechanisms utilized within the cloud environment meet and conform to the relevant SLAs and that the storage components are functioning according to their specified requirements.

A cloud architect that focuses on storage. A focal point is to make sure storage types and mechanisms used in the cloud environment meet requirements and the relevant SLAs and that the storage components function according to their specifications.

Cloud Developer

Full stack engineering, automation and development of the infrastructure are key aspects of the role. Interactions with cloud administrators and security practitioners will be required for debugging, code reviews, and security assessment requirements.

Cloud Service Auditor

A cloud service auditor can provide value by determining the effectiveness of the CSP, identify control deficiencies within the consumer organization, and provide an assessment of the quality of service. This would include determining if the SLA is being met.

Third-party organization that verifies attainment of SLAs.

Cloud Service Broker (CSB)

A third party that looks to add value to customers of cloud-based services working with multiple CSPs. Value can come from activities such as customizing the services or monitoring the services.

A third-party entity which acts as a liaison between customers and CSPs ideally selecting the best provider for each customer. The CSB acts as a middleman to broker the best deal and customize services.

Cloud Service Customer

Anyone who is purchasing a cloud service. This could be an individual or an organization.

Cloud Service Consumer

Individual or entity that utilizes or subscribes to cloud-based services or resources.

Cloud Service Integrator

Someone who connects (or integrates) existing systems and services to the cloud for a cloud customer.

Cloud Service Manager

Responsible for policy design, business agreement, pricing models, and SLAs. This role interacts with cloud management and customers. In addition, the role will work with the cloud administrator to implement SLAs and policies.

Cloud Service Operations Manager

A role within a CSP that provides audit data when requested or required, manages inventory and assets, prepares systems for the cloud, and manages and maintains services.

Cloud Service Provider (CSP)

The service provider sets the governance.

A service provider that offers customer storage or software solution available via a public network, usually the Internet. The cloud provider dictates both the technology and operational procedures involved.

The CSP will own the datacenter, employ the staff, own and manage the resources (hardware and software), monitor service provision and security, and provide administrative assistance for the customer and the customer's data and processing needs.

> ⓘ Examples include Amazon Web Services, Rackspace, and Microsoft Azure.

Cloud Service User

The cloud service user has the legal responsibility for data processing that is carried out by the CSP.

The cloud service user is also known as the **data controller.**

Cloud Storage Administrator

Focuses on user groups and the mapping, segregations, bandwidth, and reliability of storage volumes assigned. Additionally, this role may work with network and cloud administrators to ensure SLAs are met.

Cloud User

Someone using cloud services. It could be an employee of a company who is a cloud customer or just a private individual.

> Not all cloud users are staff of cloud customers. Many cloud users are simply individuals who are using publicly available cloud services for their personal purposes, such as a person who has a OneDrive account to sync their data.

Managed Service Provider (MSP)

The difference between using a CSP and an MSP is that when using an MSP, the cloud customer has full control over governance.

Distinguishing characteristics of MSPs include:

- Network Operations Center (NOC)
- Help Desk
- Remote Monitoring
- Proactive Maintenance
- Predictive Billing

Cloud Characteristics

Overview

- Broad Network Access
- On-Demand Services
- Resource Pooling
- Rapid Elasticity
- Measured Service

Characteristics

Broad Network Access

There should never be network bandwidth bottlenecks. This is generally accomplished with the user of such technologies as advanced routing techniques, load balancers, multisite hosting, and other technologies.

On-Demand Services

On-demand services refer to the model that allows customers to scale their compute and/or storage needs with little or no intervention from or prior communications with the provider. The services happen in real time.

Resource Pooling

Resource pooling is the characteristic that allows the cloud provider to meet various demands from customers while remaining financially viable. The cloud provider can make capital investments that greatly exceed what any single customer could provide on their own and can apportion these resources, as needed, so that the resources are not under-utilized (which would mean a wasteful investment) or overtaxed (which would mean a decrease in level of service).

Rapid Elasticity

Allows the user to obtain additional resources, storage, compute power, and so on, as the user's need or workload requires. This is more often transparent to the user, with more resources added as necessary seamlessly.

Threats include:

- *Abuse or nefarious use of cloud services.* Even when using the cloud for legitimate purposes, from a management perspective, users in a cloud customer organization often do not pay directly for cloud services (and are often not even aware of the cost of use) and can create dozens or even hundreds of new virtual systems in a cloud environment for whatever purposes they need or desire.

Measured Service

Measured or metered service simply means that the customer is charged for only what they use and nothing more.

> ⓘ While it is true that multitenancy is quite often an aspect of most cloud service offerings, it is not exactly a defining element of the field. There are cloud services that do not include multitenancy, as customers can purchase, rent/lease, and stand-alone resources.

Cloud Drivers

There are many drivers that may move a company to consider cloud computing. These may include the costs associated with the ownership of their current IT infrastructure solutions as well as projected costs to continue to maintain these solutions year in and year out.

Reduction in Capital Expenditure (CapEx)

Capital Expenditure

- Buildings
- Computer Equipment

Operational Expenditure

- Utility Costs
- Maintenance

Reduction in IT Complexity

- *Risk Reduction.* Users can use the cloud to test ideas and concepts before making major investments in technology.
- *Scalability.* Users have access to a large number of resources that scale based on user demand.
- *Elasticity.* The environment transparently manages a user's resource utilization based on dynamically changing needs.

Consumption-Based Pricing

- *Virtualization.* Each user has a single view of the available resources, independent of their arrangement in terms of physical devices.
- *Cost.* The pay-per-usage model allows an organization to pay only for the resources it needs with basically no investment in the physical resources available in the cloud. There are no infrastructure maintenance or upgrade costs.

Business Agility

- *Mobility.* Users can access data and applications from around the globe.
- *Collaboration and Innovation.* Users are starting to see the cloud as a way to work simultaneously on common data and information.

Cloud Security

Terminology

Acronyms

Acronym	Definition
API	Application Programming Interface
DAM	Database Activity Monitoring
FAM	File Activity Monitoring
WAF	Web Application Filter

Definitions

API Gateway

An API gateway translates requests from clients into multiple requests to many microservices and delivers the content as a whole via an API it assigns to that client/session.

API gateways can provide access control, rate limiting, logging, metrics, and security filtering services.

DAM

Captures and records, at a minimum, all SQL activity in real time or near real time, including database administrator activity, across multiple database platforms; and can generate alerts on policy violations.

A DAM operates at layer 7 of the OSI model.

FAM

FAM monitors and records all activity within designated file repositories at the user level, and generate alerts on policy violations.

Honeynet

Grouping multiple honeypot systems to form a network that is used in the same manner as the honeypot, but with more scalability and functionality.

Honeypot

Honeypots are computer systems setup to look like production systems using the modern concept of deception. They contain an operating system and can mimic many common systems such as Apache or IIS web servers, Windows file shares, or Cisco routers. A honeypot could be deployed with a known vulnerability that an attacker would be enticed to exploit. While it appears vulnerable to attack, it is in fact protection the real systems from attack while gathering defensive information such as the attacker's identity, access, and compromise methods.

Used to detect, deflect, or in some manner counteract attempts at unauthorized use of information systems.

Enticement vs. entrapment. The real term to be used should be "distract".

WAF

A WAF is a type of firewall that filters HTTP traffic and can help prevent DoS attacks.

A WAF operates at layer 7 of the OSI model.

XML Gateway

XML gateways transform how services and sensitive data are exposed as APIs to developers, mobile users, and the cloud. They can be either hardware or software based and they can implement security controls such as DLP, antivirus, and antimalware. XML gateways can also act as a reverse proxy and perform content inspection on many traffic protocols, including SFTP.

Overview

Enterprise security architecture provides the conceptual design of network security infrastructure and related security mechanisms, policies, and procedures. It links components of the security infrastructure as a cohesive unit with the goal of protecting corporate information.

The following principles should be adhered to at all times:

- Define protections that enable trust in the cloud.
- Develop cross-platform capabilities and patterns for proprietary and open source providers.
- Facilitate trusted and efficient access, administration, and resiliency to the customer or consumer.
- Provide direction to secure information that is protected by regulations.
- Facilitate proper and efficient identification, authentication, authorization, administration, and auditability.
- Centralize security policy, maintenance operation, and oversight functions.
- Make access to information both secure and easy to obtain.
- Delegate or federate access control where appropriate.
- Ensure ease of adoption and consumption, supporting the design of security patterns.
- Make the architecture elastic, flexible, and resilient, supporting multitenant, multilandlord platforms.
- Ensure the architecture addresses and supports multiple levels of protection, including network, OS, and application security needs.

Standards

In the absence of cloud-specific security standards that are universally accepted by providers and customers alike, you'll deal with a patchwork of security standards, frameworks, and controls that are being applied to cloud environments.

Standard	Description

ISO/IEC 27001	Information security management systems
ISO/IEC 27002	Code of practice for information security controls
ISO/IEC 27017	Code of practice for information security controls based on ISO/IEC 27002 for cloud services
NIST SP 800-53	Security and Privacy Controls for Federal Information Systems and Organizations

Supplemental Security Devices

- Generic firewall
- Web application firewall (layer 7)
- Database activity monitoring (detect and stop malicious commands)
  - A DAM can determine if malicious commands are being executed on your organization's SQL server and can prevent them from being executed.
- Deception technology
- API gateways (access control, rate limiting, logging, metrics, and security filtering)
- XML gateways (DLP, AV, antimalware)
- IPS/IDS (host-based, network-based; signature, anomaly, stateful)

DNSSEC

Terminology

Acronyms

Acronym	Definition
KSK	Key Signing Key
ZSK	Zone Signing Key

Overview

DNSSEC is a suite of extensions that adds security to the DNS protocol by enabling DNS response validation using a process called zone signing.

Specifically, DNSSEC provides:

- Origin authority
- Data integrity
- Authenticated denial of existence

With DNSSEC, the DNS protocol is much less susceptible to certain types of attacks:

- Cache poisoning/spoofing
- Pharming
- MITM

DNSSEC uses digital signatures embedded in the data. The DNSSEC digital signature verifies you are communicating with the site you intended to visit (authenticity). DNSSEC puts additional records in DNS alongside existing records. The new record types, such as RRSIG and DNSKEY, can be accessed the same way as common records such as A, CNAME, and MX. A public and private key will exist for each zone. When a request is made, the request include information signed with a private key; the recipient then unlocks it with the public key.

It's important to note that DNSSEC does *not* provide encryption.

DNSSEC does *not* protect against:

- Confidentiality
- DDoS

Components

ZSK

Used to sign and validate the individual record sets within the zone.

KSK

Used to sign the DNSKEY records in the zone.

Cloud Shared Considerations

Overview

Components

Interoperability

Interoperability defines how easy it is to **move and reuse application components** regardless of the provider, platform, OS, infrastructure, location, storage, format of data or APIs, how well applications work together, and how well new applications work with other solutions present in the business, organization or provider's existing architecture.

Standards-based products, processes, and services are essential for entities to ensure the following:

- Investments do not become prematurely technologically obsolete.
- Organizations are able to easily change CSPs to flexibly and cost effectively support their mission.
- Organizations can economically acquire commercial and develop private clouds using standards-based products, processes, and services.

Interoperability mandates that those components should be replaceable by new or different components from different providers and continue to work, as should the exchange of data between systems.

Portability

Portability is the ability to **move applications and associated data between one cloud provider and another** or between legacy and cloud environments/public and private cloud environments.

Portability can help both prevent vendor lock-in and deliver business benefits by allowing identical cloud deployments to occur in different CSP solutions, either for the purposes of DR or for the global deployment of a distributed single solution.

Portability is the measure of how difficult it might be to move the organization’s systems/data from a given cloud host to another cloud host.

Reversibility

The process for customers to **retrieve their data and application artifacts and for the provider to delete data after an agreed period**, including contractually specified cloud service-derived data. This is important when moving from one CSP to another.

The ability of a cloud customer to quickly remove all data, applications, and anything else that may reside in the cloud provider's environment, and move to a different cloud provider with minimal impact to operations.

Involves aspects such as technical, operational, long-term support for the workload.

> ⓘ I sometimes refer to this as "removability" since it seems to be focused more on the aspect of ensuring your data is deleted from a cloud customer than it does on actually migrating data successfully like that of portability.

Availability

Systems and resource availability defines the success or failure of a cloud-based service. As a SPOF for cloud-based services, where the service or cloud deployment loses availability, the customer is unable to access target assets or resources, resulting in downtime.

Security

For many customers and potential cloud users, security remains the biggest concern, with security continuing to act as a barrier preventing them from engaging with cloud services.

Privacy

In the world of cloud computing, privacy presents a major challenge for both customers and providers alike. The reason for this is simple: no uniform or international privacy directives, laws, regulations, or controls exist, leading to a separate, disparate, and segmented mesh of laws and regulations being applicable depending on the geographic location where the information may reside (data at rest) or be transmitted (data in transit).

Resiliency

Cloud resiliency represents the ability of a cloud services data center and its associated components, including servers, storage, and so on, to continue operating in the event of a disruption, which may be equipment failure, power outage, or a natural disaster. It represents how adequately an environment can withstand duress.

Performance

Governance

The term *governance* relating to processes and decisions looks to define actions, assign responsibilities, and verify performance. The same can be said and adopted for cloud services and environments, where the goal is to secure applications and data when in transit and at rest. In many cases, cloud governance is an extension of the existing organizational or traditional business process governance, with a slightly altered risk and controls landscape.

Although governance is required from the commencement of a cloud strategy or cloud migration roadmap, it is seen as a recurring activity and should be performed on an ongoing basis.

SLAs

Auditability

Auditability allows for users and the organization to access, report, and obtain evidence of actions, controls, and processes that were performed or run by a specified user.

Regulatory Compliance

Cloud Service Models

Overview

Cloud service models affect the **level of control** an organization has over their resources.

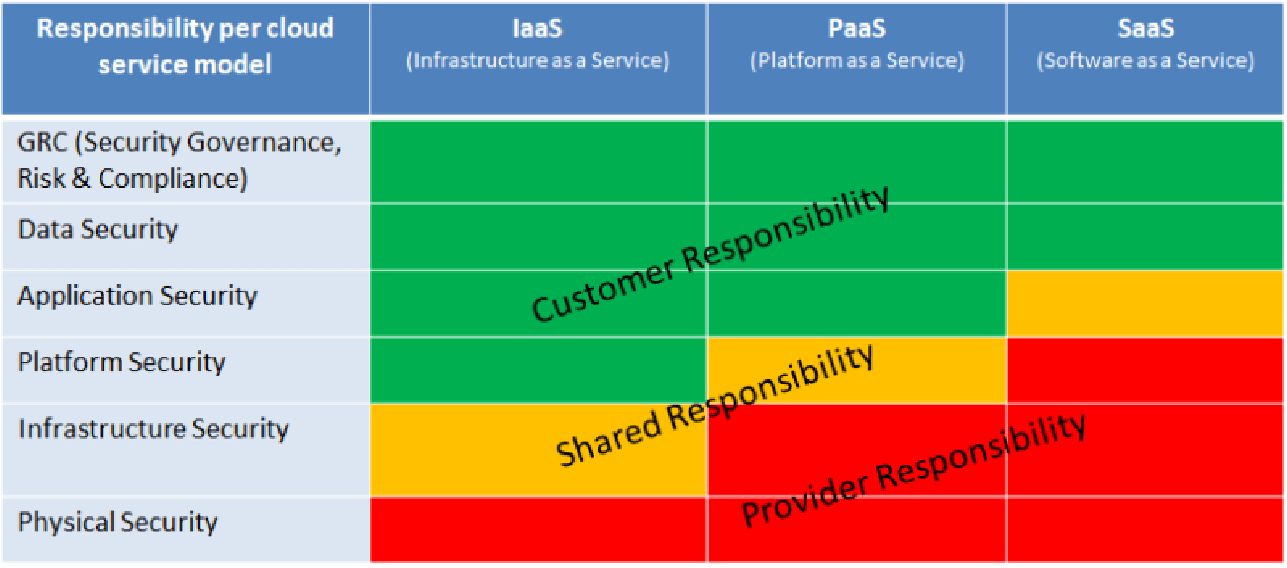

Responsibility per cloud service model	IaaS (Infrastructure as a Service)	PaaS (Platform as a Service)	SaaS (Software as a Service)
GRC (Security Governance, Risk & Compliance)	🟩	🟩	🟩
Data Security	🟩	🟩	🟩
Application Security	🟩	🟩	🟨
Platform Security	🟩	🟨	🟥
Infrastructure Security	🟨	🟥	🟥
Physical Security	🟥	🟥	🟥

Customer Responsibility

Shared Responsibility

Provider Responsibility

There are three major cloud service models:

→ IaaS /concepts/cloud/cloud-service-models/untitled

→ PaaS /concepts/cloud/cloud-service-models/paas

→ SaaS /concepts/cloud/cloud-service-models/saas

> **Fact.** In all three models, the customer is giving up an essential form of control: physical access to the devices on which the data resides.

IaaS

Overview

IaaS offers only hardware and administration, leaving the customer responsible for the OS and other software.

IaaS consists of a facility, hardware, an abstraction layer, an orchestration (core connectivity and delivery) layer to tie together the abstracted resources, and APIs to remotely manage the resources and deliver them to consumers

IaaS allows the customer to install all software, including operating systems (OSs) on hardware housed and connected by the cloud vendor. In this model, the cloud provider has a datacenter with racks and machines and cables and utilities, and administers all these things. However, all logical resources such as software are the responsibility of the customer.

> ⓘ Some examples of IaaS would include datacenters that allow clients to load whatever operating system and applications they choose. The cloud provider simply supplies the compute, storage, and networking functions.

Boundaries

Provider

The provider is responsible for the buildings and land that compose the datacenter; must provide connectivity and power; and creates and administers the hardware assets the customer's programs and data will ride on.

Customer

The customer, however, is in charge of everything from the operating system and up; all software will be installed and administered by the customer, and the customer will supply and manage all the data.

Characteristics

- Scale
- Converged network and IT capacity pool
- Self-service and on-demand capacity
- High reliability and resilience

Risks

- Personnel Threats
- External Threats (malware, hacking, DoS/DDoS, MITM, etc.)
- Lack of Specific Skillsets

Benefits

- Usage metered
- Dynamic scalability
- Reduced cost of ownership
- Reduced energy and cooling costs

PaaS

Overview

PaaS allows a way for customers to rent hardware, operating systems, storage, and network capacity over the Internet from a cloud service provider.

PaaS contains everything included in IaaS, with the addition of OSs. The cloud vendor usually offers a selection of OSs, so that the customer can use any or all of the available choices. The vendor will be responsible for patching, administering, and updating the OS as necessary, and the customer can install any software they deem useful.

> Some examples of PaaS include hosting providers that offer not only infrastructure but systems already loaded with a hardened operating system such as Windows Server or a Linux distribution.

Boundaries

Provider

The provider is responsible for installing, maintaining, and administering the OS(s).

Customer

The responsibilities for updating and maintaining the software will remain the customer's.

Characteristics

- Support multiple languages and frameworks
- Multiple hosting environments

- Flexibility
- Allow choice and reduce lock-in
- Ability to auto-scale

Risks

- Interoperability Issues (OS and OS updates may not function with customer applications)
- Persistent Backdoors
- Virtualization
- Resource Sharing

> ⓘ Risks impacting IaaS also affect PaaS:
>
> - Personnel Threats
> - External Threats (malware, hacking, DoS/DDoS, MITM, etc.)
> - Lack of Specific Skillsets

Benefits

- Operating systems can be upgraded frequently
- Distributed teams can work together
- Services are not bound by national border
- Optimization of expenditures by leveraging a single vendor

> ⓘ PaaS and SaaS often include data replication in their services.

SaaS

Overview

SaaS is a software delivery method that provides access to software and its functionality remotely as a web-based service. Allows organizations to access business functionality at a cost typically less than paying for licensed applications because SaaS pricing is based on a monthly fee.

SaaS includes everything listed in IaaS and PaaS, with the addition of software programs. The cloud vendor becomes responsible for administering, patching, and updating this software as well. The cloud customer is basically only involved in uploading and processing data on a full production environment hosted by the provider.

Within SaaS, two delivery models are currently used:

Hosted Application Management

The provider hosts commercially available software for customers and delivers it over the web

Software on Demand

The CSP gives customers network-based access to a single copy of an application created specifically for SaaS distribution

> ⓘ Some examples of SaaS would include things like customer relationship manager (CRM) software or accounting software hosted in the cloud. The provider takes care of all the infrastructure, compute, and storage needs as well as providing the underlying operating systems and the application itself. All of this is completely transparent to the end user who only sees the application they have purchased.

Boundaries

Customer

The customer only supplies and processes data to and in the system.

Characteristics

- Overall reduction of costs
- Application and software licensing
- Reduced support costs

Risks

- Proprietary Formats
- Virtualization
- Web Application Security

> ⓘ Risks impacting IaaS also affect SaaS:
>
> - Personnel Threats
> - External Threats (malware, hacking, DoS/DDoS, MITM, etc.)
> - Lack of Specific Skillsets
>
> Additionally, the risks impacting PaaS also affect SaaS:
>
> - Interoperability Issues
> - Persistent Backdoors
> - Virtualization
> - Resource Sharing

Benefits

- Limited administration
- Always running latest version of software
- Standardized software distribution
- Global accessibility

> ⓘ PaaS and SaaS often include data replication in their services.

Cloud Deployment Models

Overview

Cloud deployment models affect the **extent of resource sharing** the organization will be subjected to.

There are four major cloud deployment models:

→ **Private Cloud** /concepts/cloud/cloud-deployment-models/private

→ **Public Cloud** /concepts/cloud/cloud-deployment-models/public

→ **Community Cloud** /concepts/cloud/cloud-deployment-models/community

→ **Hybrid Cloud** /concepts/cloud/cloud-deployment-models/hybrid

Private Cloud

Overview

> **NIST SP 800-145**
>
> The cloud infrastructure is provisioned for **exclusive use by a single organization** comprising multiple consumers (e.g., business units). It may be owned, managed, and operated **by the organization, a third party, or some combination of them**, and it may exist **on or off premises.**

Examples of private clouds include such things as what used to be called intranets. These often host shared internal applications, storage, and compute resources. One example is an internally hosted SharePoint site.

Risks

- Personnel Threats
- Natural Disasters
- External Attacks
- Regulatory Noncompliance
- Malware

Threats

- Malware
- Internal Threats
- External Attackers
- Man-in-the-Middle Attacks
- Social Engineering
- Theft/Loss of Devices
- Regulatory Violations

- Natural Disasters

> ⓘ The terms "public" and "private" can be confusing, because we might think of them in the context of who is offering them instead of who is using them. Remember: A public cloud is owned by a specific company and is offered to anyone who contracts its provided services, whereas a private cloud is owned by a specific organization but is only available to users authorized by that organization.

Public Cloud

Overview

> **NIST SP 800-145**
>
> The cloud infrastructure is provisioned for **open use** by the general public. It may be owned, managed, and operated **by a business, academic, or government organization, or some combination of them**. It exists **on the premises of the cloud provider.**

The public cloud is what we typically think of when discussing cloud providers. The resources (hardware, software, facilities, and staff) are owned and operated by a vendor and sold, leased, or rented to anyone (offered to the public-hence the name).

Examples of public cloud vendors include Rackspace, Microsoft Azure, and Amazon Web Services (AWS).

Risks

Multitenant Environments

- Conflict of Interest
- Escalation of Privilege
- Information Bleed
- Legal Activity

Vendor Lock-In

Vendor lock-in occurs when the organization creates a dependency on the provider.

- Limitations to moving (bandwidth from old provider or otherwise)
- Unfavorable contracts
- Using proprietary data formats
- Regulatory constraints (where other providers may not be able to meet needs)

> ⓘ While there are several ways to avoid vendor lock-in, the best way is through favorable contract language. To avoid lock-in, the organization has to think in terms of *portability.*

Vendor Lock-Out

When the cloud provider goes out of business, is acquired, or ceases operation for any reason.

To avoid lock-out, the organization should consider the provider's:

- Longevity
- Core Competency
- Jurisdictional Suitability
- Supply Chain Dependencies
- Legislative Environment

Threats

- Loss of Policy Control
- Loss of Physical Control
- Lack of Audit Access
- Rogue Administrator
- Escalation of Privilege
- Contractual Failure

> ⓘ Threats impacting the private model also affect the public model:
>
> - Malware
> - Internal Threats
> - External Attackers
> - Man-in-the-Middle Attacks

- Social Engineering
- Theft/Loss of Devices
- Regulatory Violations
- Natural Disasters

ⓘ The terms "public" and "private" can be confusing, because we might think of them in the context of who is offering them instead of who is using them. Remember: A public cloud is owned by a specific company and is offered to anyone who contracts its provided services, whereas a private cloud is owned by a specific organization but is only available to users authorized by that organization.

Community Cloud

Overview

> **NIST SP 800-145**
>
> The cloud infrastructure is provisioned for **exclusive use by a specific community of consumers from organizations that have shared concerns** (e.g., mission, security requirements, policy, and compliance considerations). It **may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them**, and it **may exist on or off premises.**

A community cloud features infrastructure and processing owned and operated by an affinity group; disparate pieces might be owned or controlled by individuals or distinct organizations, but they come together in some fashion to perform joint tasks and functions.

Gaming communities might be considered community clouds. For instance, the PlayStation network involves many different entities coming together to engage in online gaming: Sony hosts the identity and access management (IAM) tasks for the network, a particular game company might host a set of servers that run digital rights management (DRM) functions and processing for a specific game, and individual users conduct some of their own processing and storage locally on their own PlayStations.

Risks

- Resiliency Through Shared Ownership
  - Several points of entry
  - Configuration management unity
  - Baseline enforcement
  - Shared decision making
- Shared Costs
  - Shared access and control
- No Need for Centralized Administration for Performance and Monitoring

  - Removes reliability of centralized standards for performance and security monitoring

---

Threats

- Malware
- Internal Threats
- External Attackers
- Man-in-the-Middle Attacks
- Social Engineering
- Theft/Loss of Devices
- Regulatory Violations
- Natural Disasters
- Loss of Policy Control
- Loss of Physical Control
- Lack of Audit Access

Hybrid Cloud

Overview

> **NIST SP 800-145**
>
> The cloud infrastructure is a composition of two or more distinct cloud infrastructures (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).

A hybrid cloud contains elements of the other models. For instance, an organization might want to retain some private cloud resources (say, their legacy production environment, which is accessed remotely by their users), but also lease some public cloud space as well (maybe a PaaS function for DevOps testing, away from the production environment so that there is much less risk of crashing systems in operation).

An example of a hybrid cloud environment might include a hosted internal cloud such as a SharePoint site with a portion carved out for external partners who need to access a shared service. To them it would appear as an external cloud; therefore, it would be operating as a hybrid.

Cloud Infrastructure Components

Compute

Terminology

Definitions

Affinity

Grouping of resources.

Anti-affinity

Separation of resources.

Compute Parameters

A cloud server’s compute parameters depend on the number of **CPUs** and the amount of **RAM** used. The ability to allocate these resources is a vital compute concern.

Limits

A limit creates a **maximum** ceiling for a resource allocation.

Reservations

A reservation creates a guaranteed **minimum** resource allocation that the host must meet.

Shares

The concept of shares is used to arbitrate the issues associated with compute resource **contention** situations. Share values are used to prioritize compute resource access for all guests assigned a certain number of shares. Shares allow the cluster's *reservations* to be allocated and then addresses any remaining resources that may be available for use by members of the cluster through a prioritized percentage-based allocation mechanism.

Network

SDN

Terminology

Acronyms

Acronym	Definition
NBI	Northbound Interface
SBI	Southbound Interface
SDN	Software-Defined Networking

Overview

SDN is the idea of *separating* the network control plane from the actual network forwarding plane. This allows for greater control over networking capabilities and for the integration of such things as APIs.

Characteristics

- *Programmatically configured.* Allows network managers to configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs, which they can write themselves because the programs do not depend on proprietary software.
- *Open standards-based and vendor-neutral.* When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.
- *Directly programmable.* Network control is directly programmable because it is decoupled from forwarding functions.
- *Agile.* Abstracting control from forwarding lets administrators dynamically adjust network-wide traffic flow to meet changing needs.

- *Centrally managed.* Network intelligence is (logically) centralized in software based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch.

Elements

- *Controller.* Enables centralized management and control, automation, and policy enforcement across physical and virtual network environments.
- *SBI.* Relays information between the controller (control layer) and the **individual network devices** (such as *physical* switches, access points, routers, and firewalls).
- *NBI.* Relays information between the controller (control layer) and the **applications** and policy engines, to which an SDN looks like a single logical network device

> ⓘ The SBI and NBI are considered the SDN architecture APIs, as they define the communication between the applications, controllers, and networking systems.

Benefits

The following are primary benefits are observed by using SDN:

- Hardware agnostic
- Management plane is separated from the data plane

Layers

In the SDN architecture, the splitting of the control and data forwarding functions is referred to as "disaggregation," because these pieces can be sourced separately, rather than deployed as one integrated system. This architecture gives the applications more information about the state of the entire network from the controller, as opposed to traditional networks where the network is only application-aware.

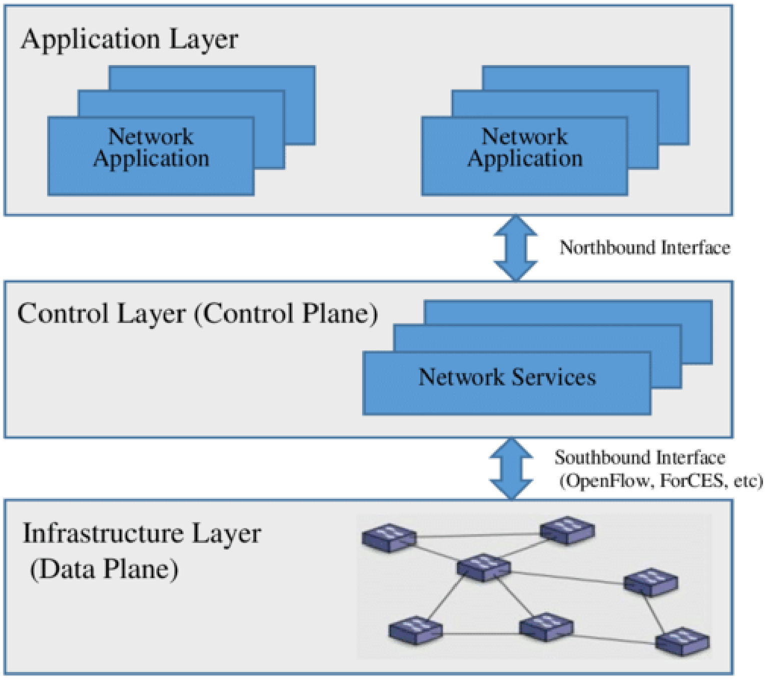
Application Layer
Network Application
Network Application
Northbound Interface
Control Layer (Control Plane)
Network Services
Southbound Interface (OpenFlow, ForCES, etc)
Infrastructure Layer (Data Plane)

SDN Architecture

Application Layer

This layer provides:

- Business applications

> APIs act as a bridge between the application and control layers using the NBI.

Control Layer (Control Plane)

The control plane refers to the processes that control the work done by the network device but do not directly impact the forwarding of individual frames or packets.

This layer provides:

- Network services
- SDN control software

Infrastructure Layer (Data Plane)

The data plane refers to the actions that devices take to forward data.

This layer provides:

- Physical network devices (which may be geographically disparate)

Network Models

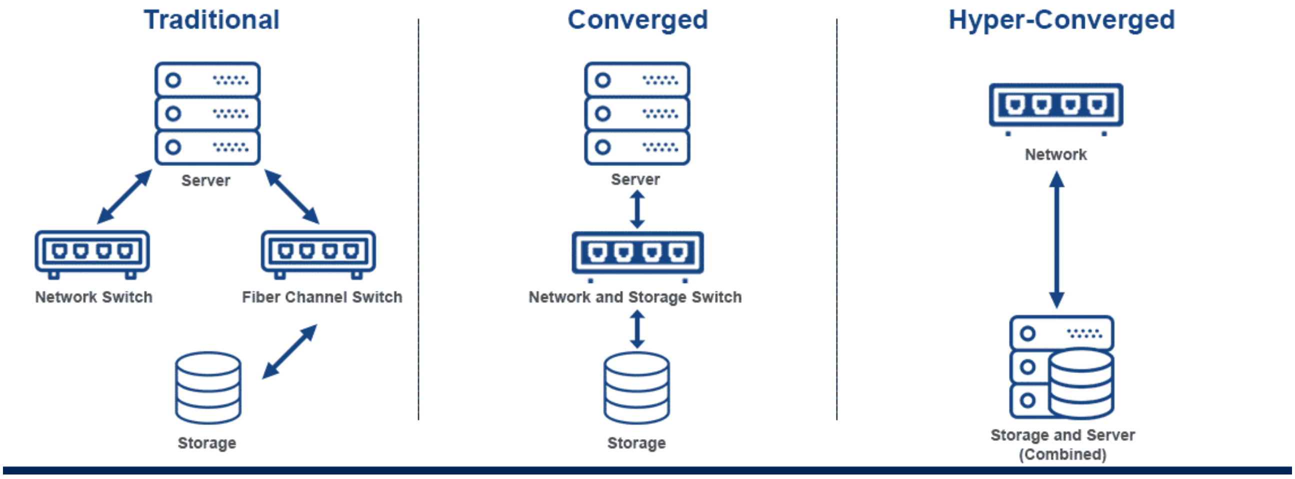

Network Model Differences

Traditional Networking Model

The traditional model is a **layered** approach with physical switches at the top layer and logical separation at the hypervisor level.

> This model allows for the use of traditional network security tools. There may be some limitations on the visibility of network segments between VMs.

Converged Networking Model

The converged model is optimized for cloud deployments and utilizes **standard** perimeter protection measures. The underlying storage and IP networks are converged (across the same infrastructure, as opposed to having a separate LAN and SAN) to maximize the benefits for a cloud workload.

You can think of a converged network model as being a super network, one that is capable of carrying a combination of data, voice, and video traffic across a single network that is optimized for performance.

> ⓘ This method facilitates the use of virtualized security appliances for network protection.

Storage

Storage Architectures

Overview

IaaS	PaaS	SaaS
Volume Storage	Structured Storage	Information Storage and Management (long-term)
Object Storage	Unstructured Storage	Content and File Storage (long-term)
		Raw Storage
		Ephemeral Storage
		CDN

> ⓘ If storage is accessible via API, then it's considered PaaS.

Storage Architectures

Volume Storage

With volume storage, the customer is allocated a storage space within the cloud; this storage space is represented as an attached drive to the user's virtual machine. From the customer's perspective, the virtual drives performs very much in the same manner as would a physical drive attached to a tangible device.

Volume storage contains two subset storage types:

- File Storage
- Block Storage

File Storage

With file storage, the data is stored and displayed just as with a file structure in the legacy environment (as files and folders), with all the same hierarchical naming functions.

Features of file storage include:

- File sharing
- Local archiving
- Data protection

File storage is commonly implemented in:

- Big data analytical tools and processes
- NAS

Block Storage

Block storage is a blank volume that the customer or user can put anything into.

Features of block storage include:

- Flexibility
- Performance

Challenges of block storage include:

- Requires greater administration
- May require OS or application to store, sort, and retrieve data

Use cases for block storage include:

- Data of multiple types and kinds, such as enterprise backup services

Common implementations of block storage include:

- iSCSI
- SAN

- RAID
- VMFS
- Email servers (such as Microsoft Exchange)

Object Storage

This type of cloud storage arrangement involves the use of associating metadata with the saved data. Data objects (files) are saved in the storage space along with relevant metadata such as content type and creation date. Data is stored as objects, not as files or blocks. Objects include the content, metadata describing the content and object, and a unique address identifier for locating that object.

An object storage system typically comes with minimal features. It gives the ability to store, copy, retrieve, and delete files and also gives authority to control which user can perform these actions. If you want to be able to search or have an object metadata central repository for other apps to draw on, you have to do it by yourself. Many storage systems such as Amazon S3 provide REST APIs and web interfaces to allow programmers to work with objects and containers.

Challenges of object storage include:

- Write-once, read many (WORM) makes object storage unsuitable for databases.
- Replication is not complete until all versions have been synchronized, which takes time. This makes object storage unsuitable for data that constantly changes, but a good solution for stagnant data.

Use cases for object storage include:

- When significant levels of description are required, including marking, labels, and classification and categorization.
- When data requires indexing capabilities.
- When data requires data policy enforcement (such as IRM/DRM or DLP).
- Centralization of some data management functions.
- Unstructured data such as music, images, and videos.
- Backup and log files.
- Large sets of historical data.
- Archived files.

- Big data endeavors.

Common implementations of object storage include:

- Amazon S3
- CDNs

Security for object storage can be provided by using:

- *IRM/DRM (file-level or file-based encryption).* Protects against *hardware* theft. Any process or user that has access to the OS still has access to the data.

Structured Storage

Structured storage includes information with a high degree of organization, such that inclusion in a relational database is seamless and readily searchable by simple, straightforward search engine algorithms or other search operations.

Databases

Databases are considered structured data. Data will be arranged according to characteristics and elements in the data itself, including a specific trait required to file the data known as the primary key. In the cloud, the database is usually backend storage in the datacenter, accessed by users utilizing online apps or APIs through a browser.

Databases may be installed on object (undesirable) or volume storage.

> ⓘ Databases are most often configured to work with PaaS and SaaS.

Unstructured Storage

Unstructured storage includes information that does not reside in a traditional row-column database.

Examples

- Email messages
- Word processing documents
- Videos
- Photos
- Audio files
- Presentations
- Web pages

Although these sorts of files may have an internal structure, they are still considered unstructured because the data they contain does not fit neatly in a database.

Information Storage and Management

Data is entered into the system through a web interface and stored within the SaaS application (usually a back-end database). This data storage utilizes database, which in turn are installed on object or volume storage.

Content and File Storage

File-based content is stored within the application and made accessible via the web-based user interface.

Ephemeral Storage

This type of storage exists only as long as the instance is online. It is typically used for swap files and other temporary storage needs and is terminated with its instance.

CDN

With a CDN, content is stored in *object* storage, which is then distributed to multiple geographically distributed nodes to improve Internet consumption speed.

A CDN is a form of data caching, usually near geophysical locations of high use demand, for copies of data commonly requested by users.

Use cases for CDNs include:

- Online multimedia streaming services; rather than dragging data from a datacenter to users at variable distances across a continent, the streaming service provider can place copies of the most requested media near metropolitan areas where those requests are likely to be made, thus improving bandwidth and delivery quality.

Raw Storage

Raw storage or raw device mapping (RDM) is an option in the VMware server virtualization environment that enables a storage LUN to be directly connected to a VM from the SAN.

Storage Types

Primary Storage

- RAM

Secondary Storage

- Media (HDD, CD/DVD, tape)

Storage Networks

Terminology

Definitions

Challenge Handshake Authentication Protocol (CHAP)

CHAP is used to periodically verify the identity of the client using a three-way handshake. This is done upon initial link establishment and may be repeated any time after the link has been established.

Initiators

The *consumer* of storage, typically a server with an adapter card in it called a HBA. The initiator commences a connection over the fabric to one or more ports on your storage system, which are called target ports.

Kerberos

Kerberos is a network authentication protocol that uses secret-key (symmetric) cryptography.

Network-Attached Storage (NAS)

A NAS is a network file server with a drive or group of drives, portions of which are assigned to users on that network. The user will see a NAS as a file server and can share files to it. NAS commonly uses TCP/IP.

Secure Remote Password (SRP)

SRP is a secure password-based authentication and key-exchange protocol. SRP uses a strong secret to enable parties to security communicate.

Storage Area Network (SAN)

A SAN is a group of devices connected to the network that provide storage space to users. Typically, the storage apportioned to the user is mounted to that user's machine, like an empty drive. The user can then format and implement a filesystem in that space according to their own preference. SANs usually use iSCSI or Fibre Channel protocols.

Simple Public-Key Mechanism (SPKM)

SPKM provides authentication, key establishment, data integrity, and data confidentiality in an online distributed application requirement.

SPKM is good for any application that uses GSSAPI calls.

The use of SPKM requires a PKI, which generates digital signatures for ensuring nonrepudiation.

Targets

The ports on your storage system that deliver storage volumes (called target devices or LUNs) to the initiators.

Protocols

iSCSI

iSCSI is an IP-based (layer 3) storage networking standard for linking data storage facilities. It provides *block-level* access to storage devices by carrying SCSI commands over a TCP/IP network. Because iSCSI is a layer 3 solution, it will permit routing of the traffic.

Best Practices

- Storage network traffic should not be shared with other network traffic (management, fault tolerance, or vMotion). A dedicated, local LAN or private VLAN should be provisioned to segregate iSCSI traffic.
- iSCSI does not handle overprovisioning of resources in a graceful manner. This practice should be avoided.

- iSCSI is unencrypted. Encryption must be added separately through IPsec (tunneling) and IKE (security).
- iSCSI supports authentication from the following protocols:
  - Kerberos
  - SRP
  - SPKM
  - CHAP

Storage Operations*

Tightly Coupled

All the storage devices are directly connected to a shared physical backplane, thus connecting all of them directly. Each component of the cluster is aware of the others and subscribes to the same policies and rulesets.

- Proprietary
- Shared physical backplane and network fixes the maximum size of the cluster
- Delivers a high-performance interconnect between servers
  - Allows for load-balanced performance
  - Allows for maximum scalability as the cluster grows (array controllers, I/O ports, and capacity can be added into the cluster as required to service the load)
- Fast, but loses flexibility*

> ⓘ A tightly coupled cluster is usually confined to more restrictive design parameters, often because the devices might need to be from the same vendor (proprietary) in order to function properly. Although this may be a limiting factor, a tightly coupled architecture will also enhance performance as it scales.

Loosely Coupled

A loosely coupled cluster will allow for greater flexibility. Each node of the cluster is independent of the others, and new nodes can be added for any purpose or use as needed. They are only logically connected and don't share the same proximate physical framework, so they are only distantly physically connected through communication media.

A loose cluster offers performance, I/O, and storage capacity **within the same node**. As a result, performance scales with capacity and vice versa.

- Cost-effective

- Can start small and grow as demand requires
- Performance, I/O, and storage capacity are all contained within the same node
    - This allows performance to scale with capacity
- Scalability is limited by the performance of the interconnect

In a loosely coupled storage cluster, each node acts as an independent data store that can be added or removed from the cluster without affecting other nodes. This, however, means that the overall cluster's **performance/capacity depends on each node's own maximum performance/capacity**. Because each node in a loosely coupled architecture has its own limitations, the number of nodes will not affect overall performance.

- Distributed.
- Built for servers in multiple locations.

Virtualization*

Virtualization can help with the following business problems:

- *Auditing.* The ability to create baselines for virtual machines aids in verifying the appropriate security controls are in place.

Data transforming from raw objects to virtualized instances to snapshotted images back into virtualized instances and then back out to users in the form of raw data may affect the organization's current classification methodology; classification techniques and tools that were suitable for the traditional IT environment might not withstand the standard cloud environment. This should be a factor of how the organization considers and perceives the risk of cloud migration.

Hypervisors

A hypervisor can be software-based, hardware-based, or firmware-based.

Type 1

Type 1 hypervisors reside directly on the host machine, often as bootable software.

> ⓘ Type 1 hypervisors are often called bare-metal or hardware hypervisors.

Type 2

Type 2 hypervisors are software hypervisors that run on top of the operating system already running on a host device.

Secure Configuration

If you are using VMware's distributed power management (DPM) technology, ensure you disable any power management settings in the host BIOS to avoid conflicting with the proper operation of DPM.

Recommendations

- *Secure build.* Every OS vendor has developed a list of best practices to securely deploy their OS.
- *Secure initial configuration.* Standard baselines are available to determine what a secure initial configuration will look like for an organization. The standard baselines should be scoped and tailored to the specific risk appetite of the organization to develop a secure initial configuration.

Best Practices

- *Host hardening.* Removing unnecessary services, changing default passwords, and renaming default accounts.
- *Host patching.* Contact the vendor to determine all patches that are available, review the patches to ensure all are appropriate, deploy the patches. Patches typically include firmware updates (least deployed), driver updates, and OS updates.
- *Host lockdown.*
- *Secure ongoing configuration maintenance.*

Best practices to secure the tools used to manage virtual hosts

- *Defense in depth.* Security tools should always be deployed in layers so that access can be controlled even if one layer is compromised.
- *Access control.* Control who has access to the tools.
- *Auditing and monitoring.* Log who is using the tools and what actions are being taken.
- *Maintenance.* Update and patch the tools as vulnerabilities are found or as vendors release critical patches.

Components

VLANs

Benefits provided by VLANs:

- Performance
    - VLANs can break up broadcast domains limiting superfluous traffic from being propagated across the network.
- Establishment of virtual workgroups
    - Workstations can be moved from one VLAN with simple changes. People working together on a particular project can easily be put into a single VLAN.
- Flexibility
    - As users move around within a campus, the switchport can be updated with their VLAN allowing them to maintain the same IP address.
- Ease of partitioning resources
    - VLANs can be setup in software versus different physical networks.

Virtual Switches

Secure Configuration

- Physical NIC redundancy to redundant physical switches
- Port channeling
- Network isolation (management plane vs. virtual switches vs. VM traffic)
- Internal and external network isolation
- Use security applications that are virtual network aware (IPS, etc.)
- vMotion is sent in clear

Application Virtualization

Allows the ability to test applications while protecting the OS and other application on a particular system. Common implementations include:

- Linux WINE
- Microsoft App-V
- XenApp

Features

Distributed Resource Scheduling (DRS)

A method for providing HA, workload distribution, and balancing of jobs in a cluster.

Live Migration

Live migration is the term used to describe the movement of functional virtual instances from one physical host to another and how VMs are moved prior to maintenance on a physical device.

- VMs are moved as "live instances" when they are transitioned from one active host to another.
- VMs are migrated in an *unencrypted* form.

VMs are moved as image snapshots when they are transitioned from production to *storage.* This is *not* live migration.

Virtual Machine Introspection (VMI)

Allows for **agentless** retrieval of the guest OS stage, such as the list of running processes, active network connections, and opening files.

An agentless means of ensuring a VM's security baseline does not change over time. It examines such things as physical location, network settings, and installed OS to ensure that the baseline has not been inadvertently or maliciously altered.

Used for malware analysis, memory forensics, and process monitoring and for externally monitoring the runtime state of a virtual machine. The introspection can be initiated in a separate virtual machine, within the hypervisor, or within another part of the virtualization architecture. The runtime state can include processor registers, memory, disk, network, and other hardware-level events.

VMI is typically used:

- With external monitoring using an IPS
- To conduct malware analysis
- To perform memory forensics

Threats

- Attacks on the Hypervisor
- Guest Escape
- Information Bleed
- Data Seizure

Hyperjacking

The installation of a rogue hypervisor that can take complete control of a host through the use of a VM-based rootkit that attacks the original hypervisor, inserting a modified rogue hypervisor in its place.

Guest Escape

Information Bleed

Data Seizure

Instant-On Gaps

Vulnerabilities that exist when, after VM has been powered off from an extended period of time and may be missing security patches, it is powered back on.

Remedies for this would include ensuring these systems are isolated until fully patched.

Operations

Personnel Isolation

- Brewer Nash might come in here

Hypervisor Hardening

Instance Isolation

Host Isolation

Management Plane

Overview

The management plane allows for cloud providers to manage all resources from a **centralized location** instead of needing to log into each individual server when needing to perform tasks. The management plane is typically hosted on its own dedicated server.

The cloud management plane will allow monitoring and administration of the cloud network. Functions of the cloud management plane include:

- Configuration management and services lifecycle management
- Services registry and discovery
- Monitoring, logging and auditing
- SLA management
- Security services and infrastructure management

Cloud Infrastructure Design

Overview

Location is the major and primary concern when building a data center. It's important to understand the jurisdiction where the datacenter will be located. This means understanding the local laws and regulations under that jurisdiction. In addition, the physical location of the datacenter will also drive requirements for protecting data during threats such as natural disasters.

The industry standard for uptime in cloud service provision is "five nines," which means 99.999% uptime. This equates to less than six minutes per year.

> A customer's ability to connect to a datacenter may be limited by a failure within the own customer's ISP. This would be a lack of availability from the customer's perspective but not a lack of uptime on the part of the provider.

Terminology

Chicken Coop

A design methodology in which a datacenter arranges racks within long rectangles with a long side facing the wind to provide natural cooling.

Redundancy

Deploying duplicate devices that can take over active operation if the primary device fails.

Resiliency

The ability to restore normal operations after a disruptive event. Redundancy is the foundation of resiliency.

Design Principles

Hardening

Encryption

Layered Defenses

Also referred to as "defense-in-depth", this is the practice of having multiple overlapping means of securing the environment with a variety of methods. These should include a blend of administrative, logical, technical, and physical controls.

Logical Design

Overview

The following is true about the logical design for a network:

- It lacks specific details such as technologies and standards while focusing on the needs at a general level.
- It communicates with abstract concepts, such as a network, router, or workstation, without specifying concrete details.

Abstractions for complex systems, such as network designs, are important because they simplify the problem space so humans can manage it (such as a WAN diagram).

Logical designs are often described using terms from the customer's *business* vocabulary.

> ⓘ An important aspect of a logical network design is that it is part of the requirements set for a solution to a customer problem.

Virtualization will leverage a hypervisor to assist with logical separation. A number of items should be considered in the logical design. These include:

- *Communications access.* What is allowed and what is not allowed?
- *Secure communications across the management plane.*
- *Secure storage.*
- *Disaster recovery.*

Physical Design

Terminology

Definitions

Cable Mining

The process of reviewing, identifying, and removing cables that are no longer being used.

Mean Time Before Failure (MTBF)

A measure of component reliability. It provides the average time between system or component failures.

Mean Time to Repair (MTTR)

Represents the average time required to repair a device that has failed or requires repair.

Mean Time to Switchover (MTTS)

The average time to switch over from a service failure to a replicated failover instance (backup).

Plenum

In building construction, a plenum is a separate space provided for air circulation for heating, ventilation, and air-conditioning (sometimes referred to as HVAC) and typically provided in the space between the structural ceiling and a drop-down ceiling.

> ⓘ Cold air is usually circulated through plenums.

Overview

The basic idea of physical design is that it communicates decisions about the hardware used to deliver a system. The following is true about a physical network design:

- It is created from a logical design.
- It often expands elements found in a logical design.

For instance, in terms of networking, a WAN connection on a logical design diagram can be shown as a line between two buildings. When transformed into a physical design, that single line can expand into the connection, routers, and other equipment at each end of the connection. The actual connection media might be shown on a physical design, along with manufacturers and other qualities of the network implementation.

Four items that would be considered in datacenter design:

- MTBF
- MTTR
- **Automating** service enablement
- Consolidating monitoring capabilities

Facilities and Redundancy

The following factors should be taken into consideration when designing a datacenter:

- Regulatory issues
- Geographic location
- Redundancy issues

Geography

Rural Design

- Availability of emergency services is a concern.

Urban Design

- Municipal codes can restrict building design.

Redundancy

Power Redundancy

- Power Provider Redundancy
- Power Line Redundancy
- Power Conditioning and Distribution Redundancy

Communications Redundancy

Personnel Redundancy

- Cross-Training
- Water
- Egress
- Lighting

Security Redundancy

- Perimeter defenses should use layered approach
- Vehicular approach/access, to include driveways that wind and curve and/or include speed bumps as well as bollards
- Guest/visitor access through a controlled entry point
- Proper placement of hazardous or vital resources
- Interior physical access controls
- Specific physical protections for highly sensitive assets
- Fire detection and suppression systems
- Sufficient power for all these functions

Holistic Redundancy

→ Uptime Institute /standards/data-center-design/uptime-institute

External Redundancy

- Power feeds/lines
- Power substations
- Generators
- Generator fuel tanks
- Network circuits
- Building access points
- Cooling infrastructure

Internal Redundancy

- Power distribution units (PDUs)
- Power feeds to rack
- Cooling units
- Networking
- Storage units
- Physical access points

Efficiency

A minimum effective (clear) height of 24 inches should be provided for raised floor installations. Additional clearance can help achieve a more uniform pressure distribution in some cases, but is not necessary.

The power requirements for cooling a datacenter depend on the amount of heat being removed (not generated) and the temperature delta between the data center and the outside air.

- Temperature: 64.4 to 80.6 F (18 to 27 C)
- Humidity: 40%-60%
- Dew point: 41.9 to 59 F (5.5 to 15 C)

Safety

Fire Suppression

FM-200

FM-200 is used as a replacement for older Halon systems specifically because it (unlike Halon) does not deplete the ozone layer. It is discharged into the room within 10 seconds and suppresses fire immediately.

- Odorless
- Colorless
- Liquefied compressed gas (at room temperature)
- It is stored as a liquid and dispensed in to the hazard as a colorless, electrically non-conductive vapor that is clear and does not obscure vision.
- Classified as a "clean agent" which means that it is safe to use within occupied spaces. It is nontoxic at levels used for fire suppression.
- Does not leave a residue after discharge.
- Does not deplete the ozone and has a minimal impact on the environment relative to the impact a catastrophic fire may have.

Halon

It is considered a good practice to avoid all unnecessary exposure to Halon.

- Effective gaseous fire suppression agent.
- It depletes the ozone and is harmful to the environment.

Fire/Smoke Detection

Codes require detectors under the floor or above the ceiling where HVAC piping, electrical feeders, or IT cables are placed within these plenum spaces.

Spot-type Detectors

- *Ionization-based*. Uses a small amount of radioactive material.
- *Photoelectric*. Uses a light source and a photosensitive sensor.

Aspirating/Air Sampling Smoke Detectors (ASSD)

Isolation

Overview

- Restricted physical access to devices
- Secure KVMs
- Restricted logical access to devices

Secure KVMs

Secure kernel-based VMs: allows you to turn Linux into a hypervisor that allows a host machine to run multiple, isolated environments called guests or VMs.

- Secure KVMs differ from their counterparts in that they are designed to deter and detect tampering.
  - Isolated data channels
  - Tamper-warning labels
  - Housing intrusion detection
  - Fixed firmware
  - Tamper-proof circuit board
  - Safe buffer design
  - Selective USB access
  - Push-button controls

> ⓘ It's important to note that KVM could mean secure Linux-based kernel virtual machine or Keyboard, Video, Mouse. Ensure you understand the difference between each of these.

Data

Data Policies

Data Archiving Policy

Needs to include the ability to perform eDiscovery and granular retrieval. The capability to retrieve data by date, subject, and author is very useful. A good archiving policy should include eDiscovery capability. Data monitoring should also be included in a data archiving policy. Cloud storage allows for data to be moved and replicated frequently. This provides for high availability and high resiliency, while requiring good data governance.

- *Data encryption*: the encryption procedure needs to consider the media used, restoration options, and how to eliminate issues with key management. Loss of encryption keys could directly lead to the loss of data. The following also need to be included in a data archiving policy:
- *Data monitoring*: cloud storage allows for data to be moved and replicated frequently. While this provides for HA and resiliency, it also creates a challenge for data governance.
- *Data restoration*: having a process to backup data is critical for data protection. Having data in a backup that is unable to be restored is useless. The restoration process needs to be tested and verified working.
- *eDiscovery process*: data stores continue to grow. Finding data in the cloud could be considered finding a needle in a haystack without a good eDiscovery process.
- *Data backup*: backing up data could be considered the foundation of a data archiving policy.
- *Data format*: numerous tape formats have been developed over the years. File formats and media types can change over time. Consideration must be given to all file formats to ensure data is not left orphaned.

Data Audit Policy

The organization should have a policy for conducting audits of its data. The policy should include detailed descriptions of:

- Audit periods

- Audit scope
- Audit responsibilities (internal and/or external)
- Audit processes and procedures
- Applicable regulations
- Monitoring, maintenance, and enforcement

The data audit policy should include the following:

- The process for data disposal
- Applicable regulations
- Clear direction of when data should be destroyed

> ⓘ The data audit policy addresses activities that take place in *all* phases of the data lifecycle.

Data Retention Policy

Organizations must demonstrate compliance with a well-defined data retention policy. The policy should ensure that only data that is not subject to regulatory or business requirements is deleted. It should also include a repeatable and predictable process.

The policy needs to consider:

- Retention periods
- Data formats
- Data security
- Data retrieval procedures

Data Classification

Overview

Data is classified based on its **value or sensitivity level.** This is performed in the **create phase** of the data lifecycle.

Data classification can be defined as a tool for categorization of data to help an organization effectively answer the following questions:

- What data types are available?
- Where is certain data located?
- What access levels are implemented?
- What protection level is implemented, and does it adhere to compliance regulations?

Virtualization has the potential to affect data classification processes and implementations in the cloud. Data transforming from raw objects to virtualized instances to snapshotted images back into virtualized instances and then back out to the users in the form of raw data may affect the organization's current classification methodology. Techniques and tools that were suitable for the traditional IT environment might not withstand the standard cloud environment.

> The purpose of classification is to dictate *how to protect data.* You protect top-secret data differently than you protect secret data.

Datasets	Input Entities
Primary set	P&DP law Scope and purpose of the processing Categories of the personal data to be processed Categories of the processing to be performed

Secondary set	Data location allowed Categories of users allowed Data retention constraints Security measures to be ensured Data breach constraints Status

Classification Process

The data classification section describes how and when data should be classified, and gives security procedures and controls for handling the data classifications.

Ask yourself, "how much damage could it cause if this data got inadvertently exposed?" (This is harm.) Data's value includes harm, time to create data, liability/compromise, etc. Similarly, what would the impact be in the following scenarios:

- If the information was widely distributed (such as SSNs or government information).
- If an employee of the CSP accessed the data.
- If the data was manipulated by an outsider or was unexpectedly changed.
- If the information was unavailable for a period of time.

The following items help determine classification:

- Sensitivity
- Jurisdiction
- Criticality

The following process should be followed:

- Execute data discovery
- Define data classification policies
- Execute data classification process
- Implement enforcement technologies to protect classified data

> Data is classified by a certain trait. For example "to encrypt" or "not to encrypt" when using encryption; or "internal use" and "limited sharing" when using DLP. It can be manual (a task assigned to the user creating the data) or automatic based on policy rules (according to location, creator, content, and so on).

> The relationship between data classification and data labeling is important. Data labeling is usually referred to as tagging the data with additional information (department, location, and creator). One of the labeling options can be classification according to certain criteria: top secret, secret, classified. **Classification is usually considered part of data labeling.**

Data Labeling

Labeling is a technology which can be used to group data elements together.

When the data owner creates, categorizes, and classifies the data, it also needs to be labeled. It is the data owner's job to label data, *not* the CSP.

Labels might include the following types of information:

- Data owner
- Date of creation
- Date of scheduled destruction/disposal
- Confidentiality level
- Handling directions
- Dissemination/distribution instructions
- Access limitations
- Source
- Jurisdiction
- Applicable regulation

> Data *classification* and labeling are most likely to affect the **Create, Store, Use, and Share** phases. Data *disposal* is most likely to affect the **Destroy** phase.

Data Protection/Control

- Data retention
- Data deletion
- Data archiving

Data Retention

The retention periods section details how long the different data classifications should be retained.

The data retention policy should include the following:

- *Retention periods.*
- *Applicable regulations.*
- *Retention/data formats.* The retention formats section details the medium on which the different data classifications should be stored. It also contains any handling procedures that should be followed.
- *Data security.*
- *Data classification.*
- *Archiving and retrieval procedures.*
- *Monitoring, maintenance, and enforcement.*

> ⓘ The data *retention* policy addresses the activities that take place in the **Archive** phase of the data lifecycle.

Data Disposal

Disposal options in the legacy environment:

- Physical destruction of media and hardware

- Degaussing
- Overwriting
- Cryptoshredding

> The data disposal policy addresses activities that take place in the Destroy phase of the data lifecycle.

Data Archival

The archiving and retrieval procedures section of the data retention policy will contain information on how data should be sent into storage to support later recovery if needed.

- Data encryption procedures
- Data monitoring procedures
- Ability to perform e-discovery and granular retrieval
- Backup and DR options
- Data format and media type
- Data restoration procedures

Data Categorization

Overview

Data is categorized based on its **use and organization.** This is performed in the **create phase** of the data lifecycle.

The data owner will be in the best position to understand how the data is going to be used by the organization. This allows the data owner to appropriately categorize the data.

- Regulatory Compliance
- Business Function
- Functional Unit
- Project

Data Roles

Overview

Clarifications

- Data **ownership** is based on **possession**.
- Data **type** is based on data **format**.
- Data **context** refers to the environment in which the data **resides** and is manipulated.
- Data **jurisdiction** refers to the standardizing bodies and regulatory entities that control the data.

Roles

Cloud Customer	Cloud Service Provider
Data Owner	Data Processor
Data Controller	Data Custodian

Data Owner

The entity that holds the legal rights and control over a set of data. Data owners define distribution and associated policies.

In most cases, this is the organization that has collected or created the data. This is also the individual with rights and responsibilities for that data; this is usually the department head or business unit manager for the office that has created or collected a certain dataset.

> ⓘ In the cloud context, the data owner is usually the **cloud customer.** From an international perspective, the data owner is also known as the **data controller.**

Data Controller

The person who either alone or jointly with other persons determines the purposes for which and the manner in which any personal data is processed; this entity determines the "why" and "how" personal data is processed.

> In the cloud context, the data controller is usually the **cloud customer.** From an international perspective, the data controller is also known as the **data owner.**

Data Custodian

Data custodians are responsible for the safe custody, transport, data storage, and implementation of business rules. This is any organization or person who manipulates, stores, or moves the data on behalf of the data owner.

The custodian is usually a specific entity in charge of maintaining and securing the privacy-related data on a *daily basis*, as an element of the data's use; for example, this could be a database administrator hired by the CSP.

The data custodian must adhere to any policies set forth by the data owner in regard to the use of the data.

> In the cloud context, the data custodian is usually the **cloud service provider.** From an international perspective, the data custodian is also known as the **data processor.**

Data Processor

Any person other than the data owner who processes the data on behalf of the data owner/controller.

> ⓘ In the cloud context, the data processor is usually the **cloud service provider.** From an international perspective, the data processor is also known as the **data custodian.**

Data Steward

The person responsible for data content, context, and associated business rules.

While the data owner maintains sole responsibility for the data and the controls surrounding that data, there is sometimes the additional role of data steward, who will oversee data access requests and the utilization of the data.

Data Subject

The individual who is the focus of personal data.

Data Lifecycle

Terminology

Acronyms

Acronym	Definition
DLP	Data Loss Prevention
DRM	Digital Rights Management
IRM	Information Rights Management

Overview

Being able to destroy data, or render it inaccessible, in the cloud is critical to ensuring confidentiality and managing a secure lifecycle for data.

1. Map the different lifecycle phases.
2. Integrate the different data locations and access types.
3. Map these into functions, actors, and controls.

Functions, Actors, and Controls

Functions

- *Access/Read.* View/read the data, including creating, copying, file transfers, dissemination, and other exchanges of information.
- *Process.* Perform a transaction on the data; update it; use it in a business processing transaction, etc. This would *not* include viewing, since that is a component of

accessing/reading.

- *Store*. Hold the data (in a file, database, etc.).

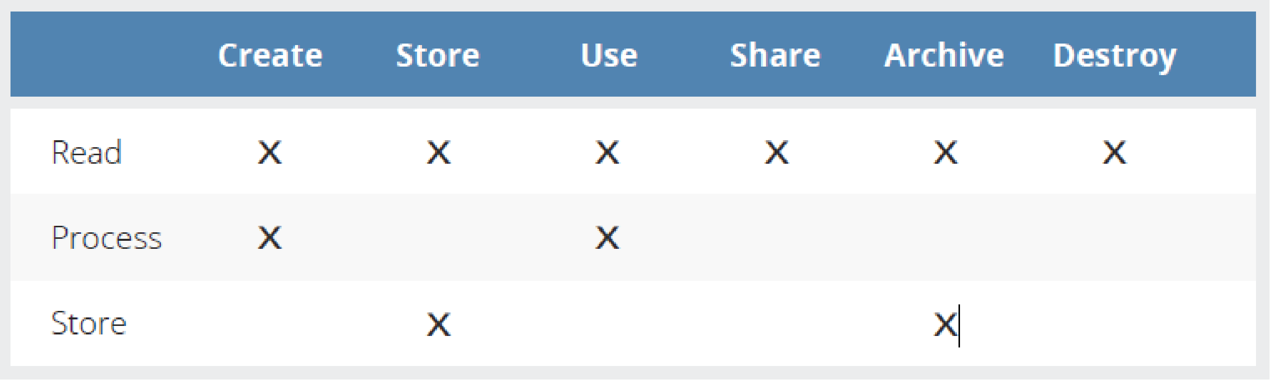

	Create	Store	Use	Share	Archive	Destroy
Read	X	X	X	X	X	X
Process	X		X			
Store		X			X	

Information Lifecycle Phases

Controls

Controls act as a mechanism to restrict a list of possible actions to allowed or permitted actions. These controls can be of a preventative, detective (monitoring), or corrective nature.

To determine the necessary controls to be deployed, you must first understand the:

- Functions of the data
- Locations of the data
- Actors upon the data

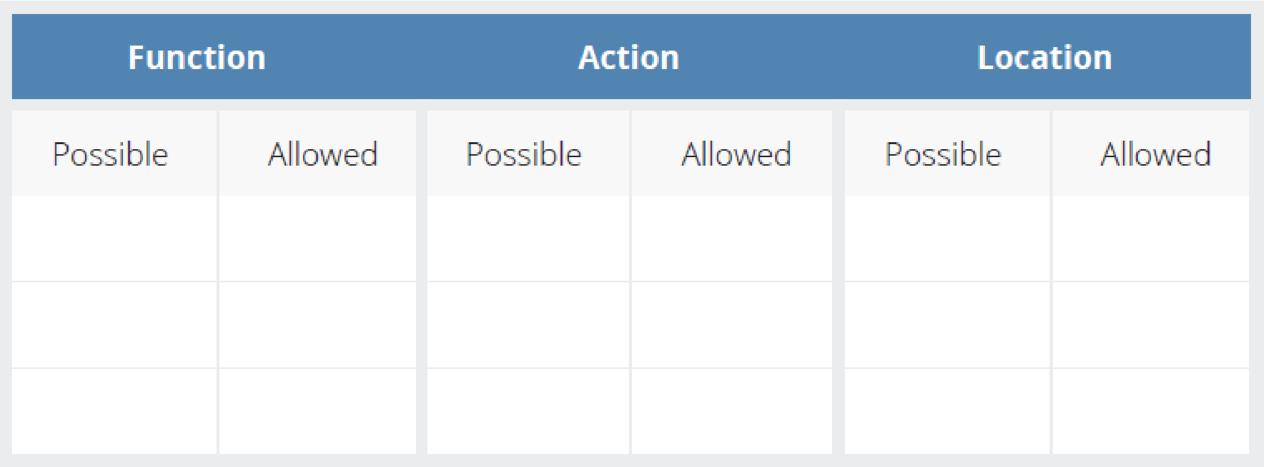

Function		Action		Location	
Possible	Allowed	Possible	Allowed	Possible	Allowed

Mapping the Lifecycle to Functions and Controls

Phases

Create/Update

The create phase is the initial phase of the data lifecycle. Data is created any time it is considered new. This encompasses data which is newly created, data that is being imported from elsewhere, and also data that already exists but has been modified into a new form. This phase could be considered "create/update".

- The **data owner** is defined.
- Data is **categorized**.
- Data is **classified**.
- Data is **labeled, tagged, and marked**.

The create phase is an ideal time to implement technologies such as SSL/TLS with the data that is inputted or imported. It should be done in the create phase so that the data is protected initially before any further phases.

Data Created Remotely

- Data should be encrypted.
- Connections should be secured (VPN).
- Secure key management practices should be in place.

Data Created within the Cloud

- Data should be encrypted.
- Secure key management practices should be in place.

Store

Usually meant to refer to near-term storage (as opposed to long-term storage). Occurs almost concurrently with the Create phase.

As soon as data enters the store phase, it's important to immediately employ:

- The use of backup methods on top of security controls to prevent data loss.

- Additional encryption for data at rest.
- DLP and IRM technologies are used to ensure that data security is enforced during the Use and Share phases of the cloud data lifecycle. They may be implemented during the Store phase, but do not enforce data security because data is not accessed during this phase.

> ⓘ While security controls are implemented in the create phase in the form of SSL/TLS, **this only protects data in transit and not data at rest**. The store phase is the first phase in which security controls are implemented to protect *data at rest*.

Use

Data is vulnerable in this state since it must be unencrypted.

- Technologies such as **DLP** and **IRM/DRM** could be leveraged to assist with monitoring access.

User Side

- Connections should be secured (VPN).
- The platforms with which users connect to the cloud should be secured.
- Permissions for modifying and processing should be implemented.
- Logging and auditing should be implemented.

Provider Side

- Strong protections in the implementation of virtualization.
- Personnel and administrative controls should be implemented.

> ⓘ Due to the nature of data being actively used, viewed, and processed in the use phase, it is more likely to be leaked in this phase than in others.

Share

- *IRM/DRM.* Can control who can share and what they can share.
- *DLP.* Can identify and prevent unauthorized sharing.
- *VPNs/encryption.* For confidentiality.
- *Restrictions based on jurisdiction.* Export or import controls, such as ITAR, EAR, or Wassenaar.

Archive

- Data should be encrypted.
- Key management is of utmost importance.
- Physical security.
    - Location (environmental, jurisdictional, geographical)
    - Format (medium, portability, weaknesses, age)
    - Staff Procedure (recovery procedures, backups)
- Retention policies
    - Retention period
    - Applicable regulations
    - Retention formats
    - Data classification
    - Archiving and retrieval procedures
    - Monitoring, maintenance, and enforcement

Many cloud providers will offer archiving services as a feature of the basic cloud service; realistically, most providers are already performing this function to avoid inadvertent loss of customer data. Because the customer is ultimately responsible for the data, the customer may elect to use another, or an additional, archive method. The **contract** will stipulate specific terms, such as archive size, duration, and so on and will determine who is responsible for performing archiving activities in a managed cloud environment.

Destroy

- Cryptoshredding (cryptographic erasure)

> It's very important to note that cryptoshredding requires **two cryptosystems:** one to encrypt the *target data* and one to encrypt the resulting *encryption keys.*

Data Privacy

Terminology

Definitions

Personal Data

Any information relating to an identified or identifiable natural personal data subject; an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his or her physical, physiological, mental, economic, cultural, or social identity.

Privacy

The protection of PII.

Processing

Processing is any *manipulation* of the data, to include security or destroying it, in electronic or hard-copy form. Viewing data is not considered processing.

Security

The owner's right to determine to whom information is disclosed. Security protects privacy.

Direct and Indirect Identifiers

Direct identifiers and indirect identifiers form the two primary components for identification of individuals, users, or personal information.

Direct Identifiers

Legally defined PII elements are sometimes referred to as *direct identifiers.* Direct identifiers are those data elements that **immediately reveal** a specific individual (the person's name, Social Security or credit card number, and so on).

Indirect Identifiers

Indirect identifiers are the characteristics and traits of an individual that **when aggregated could reveal** the identity of that person (the person's birthday, library ID card number, and so on).

User Data Types

Sensitive Data

Sexual orientation and religious affiliation fit within the sensitive data category. Other information include health information and political beliefs.

Personal Data

Personal data includes address, phone number, date of birth, and gender. Personal data can usually be discovered with a minimal amount of investigation.

Internet Data

Internet data includes browsing habits, cookies, and other information regarding an individual's internet usage.

Biometric Data

Biometric data includes fingerprints, finger scans, retina scans, and other biometric data that would need to be captured using a biometric scanner or software.

Contractual and Regulated PII

PII relates to information or data components that can be utilized by themselves or along with other information to identify, contact, or locate a living individual.

NIST SP 800-122 defines PII as any information about an individual...

> ...that can be used to distinguish or trace an individual's identity, such as name, Social Security Number, date and place of birth, mother's maiden name, or biometric records; and any other information that is linked or linkable to an individual, such as medical, education, financial, and employment information.

Contractual PII

Where an organization or entity processes, transmits, or stores PII as part of its business services, this information is required to be adequately protected in line with relevant laws.

Failure to meet or satisfy contractual requirements may lead to penalties through to termination of contract at the discretion of the organization to which services are provided.

Regulated PII

The key focus and distinct criteria to which the regulated PII must adhere is required under law and statutory requirements, as opposed to the contractual criteria that may be based on best practices or organizational security policies.

Key differentiators from a regulated perspective involve satisfying regulatory requirements (such as HIPAA and GLBA).

Failure to supply these can result in sizable and significant financial penalties and restrictions around processes, storing, and providing of services.

Mandatory Breach Reporting

Another key component and differentiator related to regulated PII is mandatory breach reporting requirements.

Mandatory breach reporting requires both private and government entities to notify and inform individuals of any security breaches involving PII.

Security breaches should be reported immediately to customers; however, 72 hours is defined in GDPR for informing the authorities.

Data Security

Terminology

Definitions

Snarfing

The action of grabbing data and using it without the owner's consent.

Overview

Individuals who gain unauthorized access to data are the most common and well understood threat to storage. The unauthorized access can be from an outside attacker, a malicious insider, or a user who may not be malicious but still has access to something he or she shouldn't.

Strategies

Data Security Strategies

- Understand data type
- Understand data structure and format
- Understand the cloud service models
- Understand the cloud storage options
- Understand CSP data residency offering
- Plan data discovery and classification
- Define data ownership
- Plan protection of data controls
- Plan for ongoing monitoring

Data Separation Strategies

Data separation should be implemented in a layered approach. The five layers addressed are:

- Compute nodes
- Management plane
- Storage nodes
- Control plane
- Network

Data Dispersion

Terminology

Acronyms

Acronym	Definition
AONT-RS	All-or-Nothing-Transform with Reed-Solomon
SMSS	Secret Sharing Made Short

Overview

Data dispersion is a technique that is commonly used to improve data security, but without the use of encryption mechanisms.

- Bit splitting
- Erasure coding

If the data is static (doesn't change), creating and distributing the data is a one-time cost. If the data is dynamic (continually changing), the erasure codes have to be re-created and the resulting data blocks redistributed.

Data dispersion is much like traditional RAID technologies; spreading the data across different storage areas and potentially different cloud providers spread across geographic boundaries. However, this comes with inherent risk. If data is spread across multiple cloud providers, there is a possibility that an outage at one provider will make the dataset unavailable to users, regardless of location. This would be a threat to availability, depending on the implementation.

Components

Bit Splitting

Bit splitting is like adding encryption to RAID. The data is first encrypted, then separated into pieces, and the pieces are distributed across several storage areas.

Bit splitting carries the following **disadvantages:**

- Processing is CPU intensive
- Availability concerns; all parts of the data need to be available to decrypt and use the information
- Storage requirements and costs are higher than other storage systems

Bit splitting can use different encryption methods, a large percentage of which are based on two secret sharing cryptographic algorithms:

- SMSS
- AONT-RS

SMSS

SMSS uses a three-phase process:

- Encryption of information
- Use of information dispersal algorithm
- Splitting the encryption key using the secret sharing algorithm

The different fragments of data and encryption keys are then signed and distributed to different cloud storage services. This makes it impossible to decrypt without both arbitrarily chosen data and encryption key fragments.

AONT-RS

AONT-RS integrates the AONT and erasure coding. This method first encrypts and transforms the information and the encryption key into blocks in a way that the information cannot be recovered without using all the blocks, and then it uses the IDA to split the blocks into shares that are distributed to different cloud storage services (similar to SMSS).

An AONT is an encryption mode which allows the data to be understood only if all of it is known.

> Depending on how the system is implemented, some or all of the data set is required to be available to unencrypt and read the data.

Erasure Coding

Erasure coding is like using parity bits for RAID striping; in RAID, parity bits help you recover missing data if one striped drive gets lost. Erasure coding helps you recover missing data if a cloud data are is unavailable/lost while your data is dispersed.

Erasure coding is a method of data protection in which data is broken into fragments that are expanded and encoded with a configurable number of redundant pieces of data and stored across different locations such as disks, storage nodes, or geographical locations. This allows for the failure of 2 or more elements of a storage array, thus offering more protection than RAID. This is commonly referred to as *chunking.* When encryption is used, this is referred to as *sharding.*

> Erasure coding is good for latency tolerant, large capacity stores and is generally found in the context of *object storage* with very large volume-cloud operators.

Data Deidentification

Overview

In certain scenarios, we may find it necessary to obscure actual data and instead use a representation of that data. Data de-identification is a method of creating a structurally similar but inauthentic version of an organization's data, typically used for purposes like software testing and user training or:

- In **test environments** where new software is tested, actual production data should *never* be used. However, in order to determine the actual functionality and performance of the system, it will be necessary to use data that approximates closely the same traits and characteristics of the production data.
- When **enforcing least privilege** we do not always need to show user's all elements of a dataset. For instance, a customer service representative might need to access a customer's account information, and be shown a screen with that information, but that data might be an abridged version of the customer's total account specifics.
- For **secure remote access** when a customer logs onto a web service, the customer's account may have some data abridged to avoid risks such as hijacked sessions, stolen credentials, or shoulder-surfing.

> **Deidentification** is needed when you want to give a dataset to someone for statistical analysis or for testing new software but parts of the actual data (usually data that could be used to identify a person) must be hidden from them (often required by regulatory legislation).
>
> **Encryption** is needed when you want to ensure that people who can access the files containing the data or backups can't read any of the data unless they are able to log in to the database and have database user permissions to see what they want to see. You may want to encrypt deidentified data for testing and performance measurement purposes, so that the testers can't see the identities but do find any problems arising out of the encryption as well as other problems - that's the case where both are needed.

Techniques

Masking

Substitution/Randomization

Substitution/randomization is the replacement of the data (or part of the data) with random characters. Usually, other traits are left intact: length of the string, character sets, special characters, case sensitivity, etc.

There are two primary methods of substitution:

- *Random substitution.* Does not allow for reconstruction of the original data stream.
- *Algorithmic substitution.* Allows the real data to be regenerated.

> ✓ **Fact.** This is the most effective method of applying data masking *and* still being able to preserve the authentic look and feel of the data records.

Hashing

Using a one-way cryptographic function to create a digest of the original data. Ensures data is unrecoverable and can be used as an integrity check.

Because hashing converts variable-length messages into fixed-length digests, you lose many of the properties of the original data. Additionally, you can no longer recover the original message.

Shuffling

Using different entries from within the same data set to represent the data. This has the obvious drawback of using actual production data.

Masking

Hiding the data with useless characters; for example, showing only the last four digits of a Social Security number.

Data masking is a technology that keeps the format of a data string but alters the content. For example, if you are storing development data for a system that is meant to parse SSNs, it is important that the format remain intact. Data masking ensures that the data retains its original format without being actionable by anyone who manages to intercept the data.

- *Static data masking (SDM)/static obscuring.* A new dataset is created as a copy from the original data. Only the obscured copy is used. This would be a good method to create a sample data set for testing purposes. Typically efficient when creating clean nonproduction environments.
    - Primarily used to provide high quality (i.e., realistic) data for development and testing of applications without disclosing sensitive information.
    - Includes protecting data for use in analytics and training as well as facilitating compliance with standards and regulations that require limits on the use of data that identifies individuals.
    - Facilitates cloud adoption because DevOps workloads are among the first that organizations migrate to the cloud. Masking data on-premise prior to uploading it to the cloud reduces risks for organizations concerned with cloud-based data disclosure.
- *Dynamic data masking (DDM)/dynamic obscuring.* Data is obscured as it is called (such as when the customer service agent or customer is granted authorized access, but the data is obscured as it is fed to them). Efficient when protecting production environments. It can hide the full credit card number from customer service representatives, but the data remains available for processing.
    - Primarily used to apply role-based (object-level) security for databases/applications.
    - In practice, the complexities involve in preventing masked data from being written back to the database essentially means DDM should only be applied in read-only contexts such as reporting or customer service inquiry functions.
    - Not well suited for use in a dynamic (read/write) environment such as an enterprise application because masked data could be written back to the database, corrupting the data.

> ✓ **Fact.** Masking is not very effective for test systems but is very useful for billing scenarios. This is a common dynamic (DDM) method of masking data.

Deletion/Nulls

Deleting the raw data from the display before it is represented, or displaying null sets.

> ⓘ Other possible masking techniques include:
>
> - *Encryption.* Most complex and most significant performance decrease.
> - *Numeric variance.* For financial and data drive information fields.

Anonymization

Unlike masking or obfuscation, in which the data is replaced, hidden, or removed entirely, anonymization is the process of removing any **identifiable characteristics** from data. It is often used *in conjunction* with another method such as masking.

Obscuring more information about an individual to prevent aggregation or inference.

Anonymization strips out identifying information from a record.

The process of anonymization is similar to masking and includes identifying the relevant information to anonymize and choosing a relevant method for obscuring the data.

Tokenization

Tokenization is the practice of utilizing a random or opaque value to replace what would otherwise be sensitive data.

The practice of tokenization involves having two distinct databases: one with the live, actual sensitive data, and one with nonrepresentational tokens mapped to each piece of data. In this method, the user or program calling the data is authenticated by the token server, which pulls the appropriate token from the token database, and then calls the actual data that maps to that token from the real database of production data, and finally presents it to the user or program.

1. An application collects or generates a piece of sensitive data.
2. Data is sent to the tokenization server; it is not stored locally.
3. The tokenization server generates the token. The sensitive data *and* the token are stored in the token database.
4. The tokenization server returns the token to the application.
5. The application stores the token rather than the original data.
6. When the sensitive data is needed, an authorized application or user can request it.

Tokenization generates a token that is used to substitute sensitive data, which itself is stored in a secured location such as a database. When accessed by a nonauthorized entity, only the token string is shown, not the actual data.

> ✓ **Fact.** Tokenization is often used when the *format* of the data is important (e.g., replacing credit card numbers in an existing system that requires the same format text string). Tokenization is often implemented to satisfy the PCI DSS requirements for firms that process credit cards.

DLP

Terminology

Definitions

Content Discovery

Includes the tools and processes to identify sensitive information in storage.

Overview

DLP solutions are designed to protect your documents when they are *inside* your organization (unlike IRM/DRM which can protect an object anywhere it travels). It is a set of controls that is used to ensure that data is only accessible to those who should have access to it.

DLP describes the controls put in place by an organization to ensure that certain types of data (structured and unstructured) remain under organizational controls, in line with policies, standards, and procedures.

> ⓘ According to CSA, DLP is typically a way to monitor and protect data that your employees access via monitoring local systems, web, email, and other traffic. It is not typically used within datacenters, and thus is more applicable to SaaS than PaaS or IaaS, where it is typically not deployed.

Components

DLP consists of three components:

1. Discovery and Classification
2. Monitoring
3. Enforcement

1. Discovery and Classification

The discovery process usually maps data in cloud storage services and databases and enables classification based on data categories (regulated data, credit card data, public data, and more).

2. Monitoring

Monitors for violations.

3. Enforcement

Many DLP tools provide the capability to interrogate data and compare its location, use, or transmission destination against a set of policies to prevent data loss. If a policy violation is detected, specified relevant enforcement actions can automatically be performed.

DLP policies are enforced and violations which were observed during the monitoring phase are addressed.

DLP provides the following benefits:

- Additional Security
- Policy Enforcement
- Enhanced Monitoring
- Regulatory Compliance

Types of DLP

DAR

In order to protect DAR, DLP solutions must be deployed on each of the systems (typically as an agent) that house data, including any servers, workstations, and mobile devices.

Mostly associated with **storage.**

- Storage based data
- Installed on the actual storage subsystems, file servers, or application servers

DIM/DIT

Mostly associated with **network traffic.**

- Deployed near the gateway
    - Proxy, network tapping, SMTP relays
    - Requires SSL interception to inspect HTTPS traffic

DIU

Mostly associated with **endpoints.**

- Offers insights into how users access/process data files
- More complex due to the large number of devices, users, and potential for multiple locations

Encryption

Terminology

Acronyms

Acronym	Definition
DAR	Data at Rest
DIM	Data in Motion
DIT	Data in Transit
DIU	Data in Use
EFS	Encrypting File System
FDE	Full Disk Encryption
FPE	Format-Preserving Encryption
WDE	Whole Disk Encryption

Definitions

Diffie-Hellman

The Diffie-Hellman **key exchange process** is used for **asymmetric encryption** and is designed to allow two parties to create a shared secret (symmetric key) over an untrusted medium.

> ⓘ Diffie-Hellman is not a symmetric algorithm; it is an asymmetric algorithm used to establish a shared secret for a symmetric key algorithm.

Hardware Security Module (HSM)

A device that can safely *store* and *manage* encryption keys used in servers, data transmission, log files, and so forth.

> ⓘ The key difference between HSM and TPM is that an HSM manages keys for several devices, whereas a TPM is specific to a single device.

Homomorphic Encryption

Intended to allow for processing of encrypted material without decrypting it first. Since the data is never decrypted, the provider and anyone trying to intercept communication between the user and the cloud would never have the opportunity to view the data in plaintext form.

IPsec

IPsec is a framework for providing secure transmission of sensitive information over unsecured networks such as the Internet. IPSec works at the network layer, protecting and authenticating IP packets between participating IPsec endpoints.

IPsec can provide the following network security services:

- Confidentiality by encrypting packets before sending them across a network.
- Integrity by authenticating packets sent to ensure that the data has not been altered during transmission.
- Authentication of the source of the IPsec packets sent.
- Antireplay by detected and rejected replayed packets.

Downsides to IPsec:

- IPsec adds overhead to network traffic
- Not NAT friendly

IPsec does not always tunnel traffic. GRE is a technology leveraged by IPsec that performs tunneling. IKE is another aspect of IPsec that provides encryption.

Trusted Platform Module (TPM)

A physical chip on a host device which stores RSA encryption keys specific to that host for hardware authentication. The purpose is to provide WDE in the event that a hard drive is removed from its host.

> ⓘ The key difference between HSM and TPM is that an HSM manages keys for several devices, whereas a TPM is specific to a single device.

Overview

Encryption will be used to protect data at rest, in transit, and in use. Encryption will be used on the remote user end to create the secure communication connection (VPN), on the cloud customer's enterprise to protect their own data, and within the datacenter by the cloud provider to ensure various cloud customers don't accidentally access each other's data.

Without encryption, it would be impossible to use the cloud in any secure fashion.

> ⓘ **Deidentification** is needed when you want to give a dataset to someone for statistical analysis or for testing new software but parts of the actual data (usually data that could be used to identify a person) must be hidden from them (often required by regulatory legislation).
>
> **Encryption** is needed when you want to ensure that people who can access the files containing the data or backups can't read any of the data unless they are able to log in to the database and have database user permissions to see what they want to see. You may want to encrypt deidentified data for testing and performance measurement purposes, so that the testers can't see the identities but do find any problems arising out of the encryption as well as other problems - that's the case where both are needed.

Using encryption should always be directly related to business considerations, regulatory requirements, and any additional constraints that the organization may have to address.

There are three components of an encryption system:

- *The data.* This is the data object or objects that need to be encrypted.
- *The encryption engine.* Performs the mathematical process of encryption.
- *Key management.* All encryption is based on keys. Safe-guarding the keys is a crucial activity, necessary for ensuring the ongoing integrity of the encryption implementation and its algorithms.

Encryption by Data Architecture

DAR

Mostly associated with **storage.**

DAR is data that is stored on some type of media such as a hard disk drive or solid state drive. DAR deals with storage-based data.

Availability and integrity of DAR is controlled through the use of redundant storage.

Common Protections

- **FDE/WDE** (BitLocker) to encrypt the entire drive
- **Database encryption** (transparent encryption)
- **Volume encryption** to encrypt a specific volume
- **File or directory encryption** (EFS) to encrypt individual files, folders, or directories
- **FPE**
- **DLP**

> ⓘ DAR is *best* protected using encryption methods such as **AES** (which is used by most **FDE/WDE** encryptions, including **BitLocker**).

DIM/DIT

Mostly associated with **network traffic.**

DIM deals with moving data across network or communication links. Transferring data from one computer to another such as via SFTP, or by using an interactive user session such as submitting data via a browser.

- DIM is typically associated with network traffic and is deployed near the gateway. For this reason, SSL interception would be necessary to inspect encrypted (HTTPS) traffic. The topology is a mixture of proxy, network tapping, or SMTP relays.

Common Protections

- **VPN** (IPsec) to provide confidentiality
- **TLS/SSL/HTTPS** to prevent eavesdropping or tampering
- **VLANs** help provide confidentiality and integrity
- **DLP** provides control between demarcation network zones by using activity monitoring and egress filtering

> ⓘ DIM/DIT is *best* protected by using encrypted transport methods such as **TLS**.

DIU

Mostly associated with **endpoints.**

DIU is data that is **actively being manipulated** by a specific user endpoint, such as if a file is opened and being worked on. This is drastically different than DIM; DIU is data that is being processed by the CPU in real-time.

DIU is mostly concerned with the actual endpoint. It requires leading edge technology and is not commonly protected. It offers insights into how users access/process data files. Client-based technology is typically more complex due to the large number of devices, users, and potential for multiple locations.

Common Protections

- **Digital signatures and encryption** to protect APIs.
- **IRM** can be used as a means for **data classification and control.**
- **DRM** is an extension of normal data protection which is encapsulated within the concept of IRM. In DRM, advanced security controls such as extra ACLs and permission requirements are placed onto the data.

→ IRM/DRM /concepts/data/data-security/irm-drm

> ⓘ DIU is *best* protected through secure API calls and web services, which make use of technologies such as **digital signatures**.

Encryption by Service Model

Overview

IaaS

- Volume encryption
  - Instance-managed encryption
  - Externally-managed encryption
- Object and file storage
  - Client-side encryption
  - Server-side encryption
  - Proxy encryption

PaaS

- Application layer encryption
- Database encryption

- Other

SaaS

- Provider-managed encryption
- Proxy encryption
- [DLP]
- [Dedicated appliance/server]
- [Virtual appliance]
- [Endpoint agent]
- [Hypervisor agent]
- [DAM/FAM]

IaaS

Volume Storage Encryption

- *Instance-managed encryption.* The encryption engine runs within the instance, and the key is stored in the volume but protected by a passphrase or keypair.
- *Externally managed encryption.* The encryption engine runs in the instance, but the keys are managed externally and issued to the instance on request.

Protects from the following risks:

- Snapshot cloning/exposure
- Exploration by the cloud provider (and private cloud admins)
- Exposure by physical loss of drives

Object and File Storage Encryption

- *Client-side encryption.* When object storage is used as the back-end for an application, encrypt the data using an encryption engine embedded in the application or client.
- *Server-side encryption.* Data is encrypted on the server (cloud) side after being transferred in. The cloud provider has access to the key and runs the encryption engine.
- *Proxy encryption.* In this model, you connect the volume to a special instance or appliance/software, and then connect your instance to the encryption instance. The proxy handles all crypto operations and may keep keys either onboard or externally.

> ⓘ Object storage encryption is best able to protect against hardware theft. Since encryption occurs at the OS level, anyone with access to the OS already has access to that data.

PaaS

- *Application layer encryption.* Data is encrypted in the PaaS application or the client accessing the platform.
- *Database encryption.* Data is encrypted in the database using encryption that's built in and is supported by a database platform like Transparent Database Encryption (TDE) or at the field level.
- *Other.* These are provider-managed layers in the application, such as the messaging queue. There are also IaaS options when that is used for underlying storage.

SaaS

SaaS providers may use any of the options previously discussed. It is recommended to use per-customer keys when possible, in order to better enforce multitenancy isolation.

- *Provider-managed encryption.* Data is encrypted in the SaaS application and generally managed by the provider.
- *Proxy encryption.* Data passes through an encryption proxy before being sent to the SaaS application.

TLS

Terminology

Acronyms

Acronym	Definition
TLS	Transport Layer Security

Definitions

TLS Authentication

When the server proves its identity to the client.

Overview

In TLS, the parties will establish a shared secret, or **symmetric key**, for the duration of the session. The session key is uniquely generated each time a new connection is made. Asymmetric encryption is used in establishing a secure TLS connection.

TLS relies on PKI certificates authenticated and issued by a trusted *third-party.*

Characteristics

- NAT friendly
- More performance intensive than IPsec
- More devices support TLS than IPsec

> Remember, TLS is *communication* based (DIM/DIT), not storage based (DAR).

Components

Handshake Protocol

Allows the client and the server to authenticate each other and to **negotiate an encryption algorithm** and **cryptographic key exchange** *before* data is sent or received.

The TLS handshake protocol takes care of the **authentication and key exchange** necessary to establish or resume secure sessions. To establish a secure session, the Handshake Protocol handles the following:

- Cipher suite negotiation
- Authentication of the server and client
- Session key information exchange

This protocol is what negotiates and establishes the actual TLS connection.

Record Protocol

The TLS Record Protocol *uses the keys* setup during the TLS handshake to **secure application data**.

Provides **connection security** and ensures that the connection is private and reliable. Used to encapsulate higher-level protocols, among them the TLS Handshake Protocol. The TLS Record Protocol is the actual **secure communication method for transferring the data**.

Provides connection security, reliability, and ensures the connection will be private. It is also leveraged to verify integrity and origin of the application data.

- Fragmentation and reassembly of messages.
- Compression and decompression of outgoing/incoming blocks.
- Applies a Message Authentication Code (MAC) to outgoing messages, and verifies incoming messages using the MAC.
- Encrypts and decrypts messages.

Encryption Types

Storage Encryption

Basic Storage-Level Encryption

- The encryption engine is located on the storage management level, with the keys usually held by the CSP.
- The engine encrypts data written to the storage and decrypts it when exiting the storage.
- Only protects from hardware theft or loss; does not protect from CSP administrator access or any unauthorized access from the layers above the storage.

Volume Storage Encryption

- Typically accomplished through an encrypted container, which is mapped as a folder or volume.
- Only allows access through the volume OS and provides protection against physical loss or theft, external administrator access, snapshot and storage-level exfiltration.
- Does not protect against access made through the instance or an attack within the application running on the instance.
- Key is managed by customer in IaaS.

Volume storage can use the following types of encryption:

Instance-Based Encryption

The encryption engine is located on the instance. Keys can be guarded locally but should be managed external to the instance.

Proxy-Based Encryption

The encryption engine is running on a proxy instance or appliance. The proxy instance is a secure machine that handles all cryptographic actions, including key management and storage. Keys can be stored on the proxy or via the external key storage, with the proxy providing the key exchanges and required safeguarding of keys in memory.

Object Storage Encryption

Object storage can use the following types of encryption:

- *File-level encryption.* This is typically in the form of IRM/DRM. The encryption engine is commonly implemented at the client side (in the form of an agent) and preserves the format of the original file.
- *Application-level encryption.* The encryption engine resides in the application that is using the object storage. This type of encryption can be used with:
    - Database encryption
    - Object storage encryption
    - Can leverage proxy encryption (although this is more commonly seen in volume storage)

> ⓘ Application encryption is associated with object storage.

Database Encryption

File-Level Encryption

Database servers typically reside on volume storage. For this deployment, you are encrypting the volume or folder of the database, with the encryption engine and keys residing on instances attached to the volume.

Provides **protection** from:

- Media theft
- Lost backups
- External attacks

Does not provide protection from:

- Attacks with direct access to the application, OS, or database

Transparent Encryption

The encryption engine resides within the database, and it is transparent to the application. Keys usually reside within the instance, although processing and managing them may be offloaded to an external KMS.

Provides **protection** from:

- Media theft
- Backup system intrusions
- Certain database and application-level attacks

> ⓘ Oracle and Microsoft refer to this as Transparent Database Encryption (TDE). Sybase references this as Application Security Encryption (ASE).

Application-Level Encryption

The encryption engine resides at the application that is utilizing the database.

Used with:

- Database encryption
- Object storage encryption
- Proxy encryption

Provides **protection** from:

- Application-level attacks
- Hardware loss
- Compromised database and administrative accounts

> ✓ **Fact.** Application-level encryption involves encryption the data *before* it enters the fields of the database; it is much more difficult to search and review data that has been encrypted, so this reduces the functionality of the database.

Key Management

Overview

The main considerations for key management are performance, accessibility, latency, and security. Can you get the right key to the right place at the right time while also meeting your security and compliance requirements?

Challenges

- *Key storage.* Keys need to be stored in a secure manner. The optimal storage method is via a hardware solution such as an HSM.
- *Key replication.* Keys stored in the cloud via software will likely be backed up and replicated multiple times. This will require a good management system.
- *Key access.* Access to keys should be monitored and only permitted to authorized users.

Considerations

- *Level of protection.* Encryption keys should *always* be secured at the same level or higher as the data they protect. The sensitivity of the data dictates this level of protection, according to the organization's **data security policies**. The strength of the cryptosystem is only valid if private keys are never disclosed.
- *Key recovery.* This usually entails a procedure that involves multiple people (separation of duties/two-person integrity), each with access to only a portion of the key.
- *Key distribution.* Keys should *never* be distributed in the clear. Often, passing keys out of band is a preferable, yet cumbersome and expensive, solution.
- *Key revocation.* A process must exist for revoking keys.
- *Key escrow.* Keys are held by a trusted third party in a secure environment. This can aid in many management efforts.
- *Outsourcing.* It is preferable to have keys stored somewhere other than with the cloud provider or within the cloud provider's datacenter. Keys should *never* be stored with the data they're protecting.

Key Management Methods

There are four potential options for handling key management:

- *HSM/appliance.* Use a traditional HSM or appliance-based key manager, which will typically need to be on-premises, and deliver the keys to the cloud over a dedicated connection.
- *Virtual appliance/software.* Deploy a virtual appliance or software-based key manager in the cloud.
- *Cloud provider service.* This is a key management service offered by the cloud provider. Before selecting this option, make sure you understand the security model and SLAs to understand if your key could be exposed.
- *Hybrid.* You can also use a combination, such as using a HSM as the root of trust for keys but then delivering application-specific keys to a virtual appliance that's located in the cloud and only manages keys for its particular context.

Internally Managed

Keys are stored on the component that is also acting as the encryption engine.

- Storage-level encryption
- Internal database encryption
- Backup application encryption

Helpful for mitigating against the risks associated with lost media. However, this method alone is not ideal.

Externally Managed

Keys are maintained separate from the encryption engine and data. They *can* be on the same cloud platform, internally within the organization, or on a different cloud.

> ⓘ Externally managed keys could also be referred to as "customer-managed" keys.

Managed by a Third Party

Outsourcing key management to a third-party, such as a CASB.

Hosting the keys within the organization is expensive and complicated and attenuates some of the benefit we get from offloading our enterprise to the cloud. Another option is using a CASB. The cost of using a CASB should be lower than trying to maintain keys within the organization, and the CASB will have core competencies most cloud customer's wont.

Key Management Systems

The following are ways to manage keys that allows the customer to generate, hold, and retain the keys. The foundational aspect is the customer must control the keys to retain maximum control of the data.

- Remote Key Management System
- Client-Side Key Management

IRM/DRM

Terminology

Acronyms

Acronym	Definition
DRM	Digital Rights Management
IRM	Information Rights Management

Overview

DRM was originally associated with digital media content/multimedia whereas IRM was associated with business. The terms are now used interchangeably.

- Provides content retrieval
- Provides user authentication
- Provides key storage
- Does *not* provide certificate generation; certificates are usually generated by a CA in the PKI

> ⓘ DRM can be used as a means for data classification and control since it is required to understand the type of data it needs to protect.

DRM encrypts content and then applies a series of *rights.* Rights can be as simple as preventing copying, or as complex as specifying group or user-based restrictions on activities like cutting and pasting, emailing, changing the content, etc. Any application or system that works with DRM protected data must be able to interpret and implement the rights, which typically also means integrating with the key management system.

DRM solutions are used to protect **intellectual property (usually copyrights)**, in order to comply with the relevant protections, and to maintain ownership rights.

Permissions are bound to the actual object, not to a particular share. Permissions are embedded into the actual object, allowing granularity.

- DRM can protect documents, emails, web pages, and database columns.
- DRM ACLs determine who can open, edit, copy, save, and even print the document.
- DRM baseline policies should be used to ensure that the appropriate policies are applied to all documents created.

> ⓘ Usually, material protected by DRM need some form of labeling or metadata associated with the material in order for the DRM tool to function properly.

Implementation

- *Rudimentary reference checks*. Require the input of some information that can only be acquired if you purchased a licensed copy of the application. This input is usually a word or phrase, and is usually entered when the application is launched and in use.
- *Online reference checks*. Implemented when the application requires a product key at installation and checks that key against the vendor's license database to verify validity.
- *Local agent checks*. Implemented when an agent must be downloaded to install the application. The agent checks the application's license.
- *Presence of licensed media (CD/DVD)*.
- *Support-based licensing*. Unlicensed versions of applications could be installed, but would be unable to obtain any kind of software updates, patches, or hot fixes.

Functions

- *Persistent Protection.* Ensures that data is safeguarded or protected wherever it resides, including in copies.
- *Dynamic Policy Control.* Allows data owners to modify the permissions for their protected data.

- *Automatic Expiration.* Allows administrators to set expiration dates for access that has been granted.
- *Continuous Auditing.*
- *Replication Restrictions.* Ensures that illegal or unauthorized copying of protected data is prohibited.
- *Remote Rights Revocation.*

Challenges

- Replication restrictions
- Jurisdictional conflicts
- Agent/enterprise conflicts
- Mapping IAM & DRM
- API conflicts

Restrictions

- Printing
- Copying
- Saving
- Editing
- Screenshots

Categories

There are two broad categories of DRM:

- Consumer DRM
- Enterprise DRM

Consumer DRM

Used to protect broadly distributed content like audio, video, and electronic books destined for a mass audience. There are a variety of different technologies and standards, and the emphasis is on one-way distribution.

Consumer DRM offers good protection for distributing content to customers, but does have a sordid history with most technologies being cracked at some point.

Enterprise DRM

Enterprise DRM is used to protect the content of an organization internally and with business partners. The emphasis is on more complex rights policies and integration within business environments.

Enterprise DRM can well-secure content stored in the cloud, but requires deep infrastructure integration. It's most useful for document based content management and distribution.

Data Discovery

Terminology

Definitions

Mapping

Enables an organization to know all of the locations where data is present within an application and within other storage. Allows for the ability to implement security controls and policies by understanding what type of data is present in the system.

Overview

Discovery is a process for identifying and providing visibility into the location, volume, and context of structured and unstructured data stored in a variety of data repositories.

Data discovery is a term that can be used to refer to several kinds of tasks:

- The organization is attempting to create that initial inventory of data it owns.
- The organization is involved in electronic discovery ("eDiscovery" is the legal term for how electronic evidence is collected as part of an investigation or lawsuit).
- The modern use of datamining tools to discover trends and relations in the data already in the organization's inventory.

The goal of data discovery is to work with and enable people to use their intuition to find meaningful and important information in data.

> ⓘ The difference between *data discovery* and *eDiscovery* is that data discovery is typically used for big data or analytics whereas eDiscovery is used for evidence.

Process

- *Implement data discovery.* Helps determine where data is stored.
- *Classify discovered data.* Helps determine the types of data that must be protected.
- *Map and define controls.* Helps ensure the controls that are implemented comply with data privacy acts and provide coverage of their tenets
- *Apply controls.*

Issues

- *Poor data quality.*
- *Dashboards.*
- *Hidden costs.*

Challenges

- *Identifying where your data is.* Not knowing where the data is can make finding your data a challenge.
- *Accessing the data.* Who has access to what data? Once you find the data to be analyzed, does the analysis process have access to it in a useable way?
- *Performing preservation and maintenance.* How will the data be preserved and by whom? Preservation needs to be spelled out in an SLA.

Data Discovery Types

Label-Based Discovery

Labels are used to group data elements together and provide information about those elements.

With accurate and sufficient labels, the organization can readily determine what data it controls, and what amounts of each kind. Based on examining labels created by the data owners during the Create phase (during content creation). Labels can be used with **databases** as well as **flat files**. The *indexed sequential access method* is used with labels.

> Similar to metadata, but less formal. This form of discovery is similar to a Google search, with the greater the number of similar labels, the greater likelihood of a match. This typically is more used with flat files in ISAM or quasi-relational data storage.

Metadata-Based Discovery

Metadata is a listing of traits and characteristics about specific data elements or sets (colloquially referred to as "data about data"). It is often automatically created at the same time as the data, often by the hardware or software used to create the parent data. Data can be retrieved on a specific set of metadata.

> You could examine column attributes to determine whether the name of the column or the size and data type resembles a credit card number (for example, if the column is a 16-digit number or the name resembles "CC"). This remains the most common analysis technique.

Content-Based Discovery

Discovery tools can be used to locate and identify specific kinds of data by delving into the content of datasets. This technique can be as basic as term searches or can use sophisticated pattern-matching technology. Can often take longer than the other two methods of discovery.

> In the credit card example, a common method is to perform a Luhn check on the number. This is a numeric checksum used by credit card companies to verify if a number is valid (resembles a credit card number).
>
> This is a growing trend also used successfully in DLP and web content analysis products.

Content analysis utilizes pattern matching, hashing, statistics, lexical, or other forms of probability analysis. DLP uses content analysis.

Content-based discovery uses characteristics such as:

- Keywords
- Pattern matching
- Frequency

Data Analytics

Datamining

Real-Time Analytics/Real-User Monitoring (RUM)

Allows for reactive and predictive operations based on customers' current and past shopping behavior.

- Most closely measures actual activity.
- Harvests information from actual user activity, making it the most realistic depiction of user behavior.
- Can sometimes reveal personal information.
- Allows for reactive and predictive operations (such as recommending other, related products) based on customers' current and past shopping behavior.

Agile Business Intelligence

A data discovery approach that offers insight to **trends of trends**, using both historical and predictive approaches.

The Agile approach to data analysis offers greater insight and capabilities than previous generations of analytical technologies.

Synthetic Performance Monitoring

Synthetic agents can simulate user activity in a much faster manner than real-user monitoring and perform these actions without rest. Synthetic performance monitoring approximates user activity and thus, is not as accurate as RUM.

Forensics

Terminology

Definitions

Cloud Computing Forensic Science

The application of scientific principles, technological practices, and derived and proven methods to reconstruct past cloud computing events through identification, collection, preservation, examination, interpretation, and reporting of digital evidence.

Digital Forensics

Digital forensics is generally considered the application of science to the identification, collection, examination, and analysis of data while preserving the integrity of the information and maintaining a strict chain of custody for the data.

Forensic Science

Forensic science is generally defined as the application of science to the law.

Network Forensics

Network forensics is defined as the capture, storage, and analysis of network events. The idea is to capture every packet of network traffic and make it available in a single searchable database so that the traffic can be examined and analyzed in detail.

Standards

The goal of the following standards is to promote best practices for the acquisition and investigation of digital evidence.

Standard	Description

ISO/IEC 27037	Guide for collecting, identifying, and preserving electronic evidence
ISO/IEC 27041	Guide for incident investigations
ISO/IEC 27042	Guide for digital evidence analysis
ISO/IEC 27043	Incident investigation principles and processes
ISO/IEC 27050	Overview and principles for eDiscovery

Overview

Challenges

- Control over data (trustworthiness)
- Multitenancy
- Data volatility
- Evidence acquisition

Data Access

Access to data will be decided by the following:

- The service model
- The legal system in the country where data is legally stored

> ⓘ There are certain jurisdictions where forensic data/IT analysis requires licensure (Texas, Colorado, and Michigan, for example).

Investigation Process

1. Identify

Identify and preserve evidence and begin chain of custody documentation.

---

2. Collect

Label, record, acquire evidence, and ensure that modification does not occur.

1. Develop a plan to acquire the data; important factors for prioritization include:
    1. Likely value
    2. Volatility
    3. Amount of effort required
2. Acquire the data
3. Verify the integrity of the data

Network forensics has various use cases for data acquisition and collection:

- Uncovering proof of an attack
- Troubleshooting performance issues
- Monitoring activity for compliance with policies
- Sourcing data leaks
- Creating audit trails for business transactions

Prioritization of Volatile Data

1. Network connections
2. Login sessions
3. Contents of memory
4. Running processes
5. Open files
6. Network configuration
7. Operating system time

Alternative Prioritization

1. CPU cache, registers, RAM
2. Virtual memory
3. Disk drives
4. Backups and printouts

---

3. Examine

After data has been collected, the next phase is to examine the data, which involves assessing and extracting the relevant pieces of information from the collected data.

Yields data. Just the facts. For example:

- File opened at 10:23 AM
- DNS stopped at 7:02 AM
- etc.

---

4. Analyze

The analysis should include identifying people, places, items, and events and determining how these elements are related so that a conclusion can be reached.

Often, this effort includes correlating data among multiple sources. For instance, a NIDS log may link an event to a host, the host audit logs may link the event to a specific user account, and the host IDS log may indicate what actions that user performed.

How to identify who completed an event:

- Source address
- User identity (if authenticated or otherwise known)
- Geolocation
- Service name and protocol

- Window, form, or page (such as URL address)
- Application address
- Application identifier

Information. Taking data and putting it into context. For example:

- DNS stopped at 7:02 AM *but* nobody should have had access to DNS at 7:02 AM...

5. Report

The final phase is reporting, which is the process of preparing and presenting the information resulting from the analysis phase. Many factors affect reporting, including the following:

- Alternative explanations
- Audience consideration
- Actionable information

> ⓘ The ultimate recipient of all forensic evidentiary collection and analysis-the entity getting the reports-will be **the court**, in order to make a final determination of its merits and insights.

6. Lessons Learned

Document what was learned. Something is always learned.

eDiscovery

Terminology

Acronyms

Acronym	Definition
ESI	Electronically Stored Information

Overview

eDiscovery refers to any process in which electronic data is sought, located, secured, and searched with the intent of using it as evidence in a civil or criminal legal case.

Typically, eDiscovery refers to the process of identifying and obtaining electronic evidence for either **prosecutorial or litigation purposes**.

You do not want to use unique testing techniques because those may not be repeatable or accepted by other experts (or the court).

The Federal Rules of Civil Procedure require a party to litigation to preserve and produce electronically stored information in its possession, custody, or control. This can be very difficult when data is dispersed across many geophysical locations.

> ⓘ The difference between *data discovery* and *eDiscovery* is that data discovery is typically used for big data or analytics whereas eDiscovery is used for evidence.

Types of eDiscovery

SaaS-based eDiscovery

An eDiscovery software vendor hosts their app on their own network and delivers it to customers via the Internet. Customers use the app for various tasks such as analysis or review (collection, preservation, review).

Hosted eDiscovery

The *customer collects* relevant data in response to an eDiscovery matter, processes it, and sends it via the Internet to their hosting provider. The *provider stores* customer data on their site or in a colocation facility, *and runs various levels of eDiscovery on the data.*

Third-party eDiscovery

The customer may hire a third-party with expertise in performing eDiscovery in the cloud.

Evidence Management

Evidence Collection

Evidence is only admissible if it has no probative value (that is, if it has no bearing on the case). Modified data *is still admissible,* as long as the modification process was documented and presented along with the evidence.

The **evidence custodian** is the person designated to maintain the chain of custody for the duration of an investigation.

Chain of Evidence

The chain of *evidence* is a series of events that, when viewed in sequence, account for the **actions of a person** during a particular period of time **or the location of a piece of evidence** during a specified time period. The chain of evidence can be thought of as the details that are left behind to tell the story of what happened.

Chain of Custody

The chain of *custody* is the practice and methods of documenting control of evidence from the time it was collected until it is presented to the court.

All evidence needs to be tracked and monitored from the time it is recognized as evidence and acquired for that purpose. Clear documentation must record which people had access to the evidence, where the evidence was stored, what access controls were placed on the evidence, and what modifications or analysis was performed on the evidence from the moment it was collected until the time it reaches the court. The chain of custody should be maintained for digital evidence, including the physical medium as well as the data contained on it (bits).

Everything should be recorded with detail:

- When an item is *gathered*
- When an item is *stored*
- When an item is *removed*

- When an item is *transported*
- Whenever any action, process, test, or other handling of an item is *to be* performed
- Whenever any action, process, test, or other handling of an item *is* performed

The reasons for this include:

- The documentation of evidence ensures that the evidence can be properly traced back to its origin.
- The analysis of evidence ensures that all the data contained in the evidence is identified.
- The preservation of evidence ensures that the evidence is stored properly and able to be retrieved when needed.
- The collection of evidence ensures that all the evidence needed is properly obtained.

Goals

- Be able to prove that evidence was secure and under the control of some particular party at all times
- Take steps to ensure that evidence is not damaged in transit or storage

Logging

Terminology

Acronyms

Acronym	Definition
SIEM	Security Information and Event Management

Definitions

Aggregation

The ability of a SIEM to collect logs from several data sources.

Correlation

The ability of a SIEM to analyze and search logs.

Overview

Four key subsystems should be monitored in cloud environments:

- *Network.* Dropped packets.
- *Disk.* Slow reads/writes.
- *Memory.* Excessive usage.
- *CPU.* High utilization.

To maintain reasonable investigation capabilities, auditability, and traceability of data, it is recommended that you specify data access requirements in the cloud SLA or contract with the CSP.

Logging should suffice for the purpose of reconstructing the pertinent information (who, what, where, when, etc.) necessary to form a narrative of what transpired. This will be different for every organization and environment, and thus is *not* usually driven by standards or frameworks.

SIEM

A SIEM *aggregates* data from many sources and is able to make *correlations* based on that data.

Characteristics

- Data aggregation (centralized collection of log data)
- Data correlation (enhanced analysis capabilities)
- Alerting (automated response)
- Dashboards (for management)
- Compliance (regulatory compliance requirements such as HIPAA or SOX)
- Retention (log management and retention)
- Forensic analysis (continuous monitoring and incident response)

Benefits

- Regulatory compliance requirements (HIPAA or SOX)
- Gaining or maintaining certifications (ISO 27001)
- Log management and retention
- Continuous monitoring and incident response
- Policy enforcement validation and policy violation detection
- Can detect repeated performance issues

Logging by Service Model

SaaS

- Webserver logs
- Application server logs
- Database logs
- Guest OS logs
- Host access logs
- Virtualization platform logs and SaaS portal logs
- Network captures
- Billing records

PaaS

Application data that can be extracted and monitored is typically defined by the developers when building their PaaS application. At a minimum, however, OWASP recommends the following logs be available:

- Session management failures
- Application errors
- System events
- Application and related systems startups and shutdowns
- Logging initialization (starting, stopping, or pausing)
- Use of higher-risk functionalities, such as network connections and the addition or deletion of users
- Legal and other opt-ins, such as permissions for mobile phone capabilities and terms of use

IaaS

- Cloud or network provider perimeter network logs
- Logs from DNS servers
- Virtual machine manager logs
- Host OS and hypervisor logs
- API access logs
- Management portal logs
- Packet captures
- Billing records

IAM

Terminology

Acronyms

Acronym	Definitions
ABAC	Attribute-Based Access Control
RBAC	Role-Based Access Control
SCIM	System for Cross-domain Identity Management
SPML	Service Provisioning Markup Language

Definitions

Access Control

Access control is restricting access to a resource.

Access Management

Access management is the process of managing access to the resources.

Attributes

Facets of an identity. Attributes can be relatively static (like an organizational unit) or highly dynamic (IP address, device being used, if the user authenticated with MFA, location, etc.).

Authentication

The process of confirming an identity. When you log in to a system you present a username (the identifier) and password (an attribute we refer to as an authentication factor).

> ⓘ Authentication is often referred to as AuthN.

Authoritative Source

The “root” source of an identity, such as the directory server that manages employee identities.

Authorization

Allowing an identity access to something (e.g. data or a function).

> ⓘ Authorization is often referred to as AuthZ.

Entitlement

Mapping an identity (including roles, personas, and attributes) to an authorization. The entitlement is what they are allowed to do, and for documentation purposes we keep these in an entitlement matrix.

Entity

The person or “thing” that will have an identity. It could be an individual, a system, a device, or application code.

Federation

An association of organizations that come together to exchange information as appropriate about their users and resources to enable collaborations and transactions.

Identifier

The means by which an identity can be asserted. For digital identities this is often a cryptological token. In the real world it might be your passport.

Identity

The unique expression of an entity within a given namespace. An entity can have multiple digital identities, such as a single individual having a work identity (or even multiple identities, depending on the systems), a social media identity, and a personal identity. For example, if you are a single entry in a single directory server then that is your identity.

Persona

The expression of an identity with attributes that indicates context. For example, a developer who logs into work and then connects to a cloud environment as a developer on a particular project. The identity is still the individual, and the persona is the individual in the context of that project.

Policy Decision Point

Policy Enforcement Point

Access decisions can be enforced at various points with various technologies.

Policy Management

Establishes the security and access policies based on business needs and the degree of acceptable risk.

Role

Identities can have multiple roles which indicate context. “Role” is a confusing and abused term used in many different ways. For our purposes we will think of it as similar to a persona, or as a subset of a persona. For example, a given developer on a given project may have different roles, such as “super-admin” and “dev”, which are then used to make access decisions.

SCIM

SCIM is a standard for exchanging identity information between domains. It can be used for provisioning and deprovisioning accounts in external systems and for exchanging attribute information.

- Standardized
- Open standard for automating the exchange of user identity information between identity domains or IT systems
- Newer than SPML

SPML

A method for automating account creation.

- Standardized
- Seldom implemented due to inflexibility and lack of vendor support
- Older and uses XML, which is slow

Overview

IAM is about the people, processes, and procedures used to create, manage, and destroy identities of all kinds. It is the security discipline that enables the right individuals to access the right resources at the right times for the right reasons.

Usually, **regulatory** compliance drives IAM efforts, not business requirements or needs.

> ✓ **Fact.** IAM is always the *responsibility* of the cloud customer.

An IAM system includes the following components:

- Identity Management
- Access Management
- Identity Repository and Directory Services

An IAM system should carry out two functions:

1. Ensure the identity of an entity

2. Once authenticated, the entity should be given the correct level of access to the systems they are trying to access

IAM sets out to define what a subject (active entity) is allowed to do to an object (passive entity)?

> *Accountability* is the end purpose of all IAM efforts. Identification, authentication, and authorization are the elements of IAM that support this effort.

Key Phases

- Provisioning
- Centralized Directory
- Privileged User Management
- Authentication
- Access

Components

Identity Management

Identity management is the process whereby individuals are given access to system resources by associating user rights with a given identity. This includes registering, provisioning, and deprovisioning identities for all relevant entities and their attributes, while making that information available to the proper audit.

Access Management

Access management is the process that deals with controlling access to resources once they have been granted. Access management tries to identify who a user is and what they are allowed to access.

This stage is where the real risk decisions are made. It is more important to control access rights than it is to control the number of identities.

Identity Repositories and Directory Services

Identity Repository

An identity repository is the store of information or attributes of identities.

Directory Service

A directory services is how identities and attributes are managed. There are several directory services available today:

- X.500 and LDAP
- Microsoft Active Directory
- Novell eDirectory
- Metadata replication and synchronization

Provisioning and Deprovisioning

Provisioning Process

Individual components of provisioning and deprovisioning process could be considered parts of identity and access management. In some cases, the process is referred to as "I triple A", starting at the identification phase and ending with accountability.

- *Proofing/registering.* Typically performed by HR.
- *Provisioning.* Traditionally performed by IT but strides are being made to automate this.
- *Identification.* The process of *claiming* an identity (asserting who you are). This should be unique for accountability (such as a user ID, account number, RFID, or IP/MAC).
- *Authentication.* The the process that establishes with adequate certainty the identity of an entity. Supports the identification claim by *proving* by some means (such as a

username and password) that this is the user they say they are.; "Who are you?" and "How do I know I can trust you?"

- *Authorization.* The process of granting access to resources. Enforced at the policy enforcement point.; "What do you have access to?"
- *Accountability/auditing.* Based on identification information; the ability to match a subject's action to their identity.
- *Deprovisioning.* The process whereby a user account is disabled. This process should be standardized using SPLM or SCIM.

Provisioning Methods

Account provisioning or identity provisioning technology creates, modifies, disables, and deletes user accounts and their profiles across IT infrastructure and business applications.

- Discretionary account provisioning
- Self-service account provisioning
- Workflow-based account provisioning
- Automated account provisioning

Discretionary Account Provisioning

Administrators determine which applications and data a user should have access to.

Self-Service Account Provisioning

Users participate in some aspects of the provisioning process, thus reducing administrative overhead. Often, users are allowed to request an account and manage their passwords.

Workflow-Based Account Provisioning

Gathers the required approvals from the designated approves before granting a user access to an application or data. For example, the business rules in a finance application might require that every new account request be approved by the company's CFO.

Automated Account Provisioning

Requires every account to be added the same way through an interface in a centralized management application. This streamlines the process of adding and managing user credentials and provides administrators with the most accurate way to track who has access to specific applications and data sources.

Automated account provisioning is usually accomplished using one of the following two methods:

- SPML
- SCIM

> ⓘ Cloud platforms tend to have greater support for the ABAC model for IAM, which offers greater flexibility and security than the RBAC model. RBAC is the traditional model for enforcing authorizations and relies on what is often a single attribute (a defined role). ABAC allows more granular and context aware decisions by incorporating multiple attributes, such as role, location, authentication method, and more. When available, ABAC is the preferred model for cloud-based access management.

FIM

Terminology

Acronyms

Acronym	Definition
FIM	Federated Identity Management
IAM	Identity and Access Management
IdP	Identity Provider
OIDC	OpenID Connect
SAML	Security Assertion Markup Language
SCIM	System for Cross-domain Identity Management
SPML	Service Provisioning Markup Language
SSO	Single Sign-On

Overview

FIM is the process of asserting an identity *across different systems or organizations.* This is the key enabler of SSO and also core to managing IAM in cloud computing.

FIM is an arrangement that can be made among multiple enterprises that allows subscribers to use the same identification data to obtain access to the networks of all enterprises in the group.

> ⓘ FIM could essentially be considered SSO but for multiple organizations.

Federation Responsibilities

Cross-Certification Federation / Web-of-Trust

In a cross-certification federation, each member of the federation has to review and approve every other member for inclusion in the federation. This does not scale well, and once the number of organizations gets fairly substantial, it becomes unwieldy.

Third-Party Certification Federation (Proxy Federation)

In a third-party certification federation, member organizations outsource their responsibilities to review and approve each other to some external party who will take on this responsibility on behalf of all the members.

- In federations where the participating entities are sharing data and resources, *all of those entities are usually the service providers* because they're providing the services.
- In a third-party certification model, the *third-party is the identity provider; this is often a CASB* because they are maintaining the identity repository.
- The cloud provider is neither a federated identity provider nor a federated service provider, unless the cloud provider is specifically chosen as the third-party providing this function.

This is popular in the cloud environment, where the identifier role can often be combined with other functions and outsourced to a CASB.

Federation Standards

SAML

SAML is an OASIS standard for federated identity management that supports both **authentication** and **authorization**. It uses XML to make assertions between an IdP and a relying party. Assertions can contain authentication statements, attribute statements, and authorization decision statements. SAML is very widely supported by both enterprise tools and cloud providers but can be complex to initially configure. It is a means for users from outside organizations to be verified and validated as authorized users inside or with another organization without the user having to create identities in both locations.

SAML can include *attributes* like specific features of an app based on job role.

1. SCIM or SPML syncs accounts to IdP
2. User identity using `alukos.com` accesses an app (such as O365); this is the **service provider**
3. Service provider recognizes `alukos.com` and sends request to IdP
4. IdP verifies whether user can access the application
5. IdP sends back SAML token that asserts the user's legitimacy

In some cases, instead of the service provider sending the request to the IdP, it can send a redirect instead. This means the service provider is informing the client to request their token from the IdP:

1. User tries to access the web application
2. The application redirects the user to the IdP
3. IdP issues a claims token and redirects the user back to the application
4. The application validates the token and authorizes the user by asserting claims, allowing the user to access the authorized protected resources
5. The token is then stored in the session cookie of the user's browser, ensuring the process doesn't have to be repeated for every request.

> **Fact.** SAML 2.0 is the most widely used standard today.

OAuth

OAuth is an IETF standard for **authorization** that is very widely used for web services (including consumer services). OAuth is designed to work over HTTP. It is most often used for delegating access control/authorizations between services.

OAuth provides the next step after authentication, specifically authorization, and allows for delegation of permissions. It enables a third-party application to obtain limited access to an HTTP service on behalf of a resource owner by managing an approval interaction between the resource owner and the HTTP service, or by allowing the third-party application to obtain access on its behalf. For example, it may allow Spotify to update Facebook that you're listening to a particular song.

- OAuth is *not* designed for SSO
- OAuth provides delegation of rights to applications

OIDC

OpenID is a standard for **federated authentication** that is very widely supported for web services. It is based on HTTP with URLs used to identify the identity provider and the user/identity (e.g., identity.alukos.com). OIDC uses REST and JSON.

OIDC allows developers to authenticate their users across websites and applications without having to manage usernames and passwords; it allows information from an IdP to be used instead.

> ⓘ OIDC is more of an Open Provider (OP) rather than an IdP. For the apps, they're called relying parties.

WS-Federation

WS-Federation defines mechanisms to allow different security *realms* to federate, such that authorized access to resources managed in one realm can be provided to security principals whose identities reside in other realms.

- WS-Federation is reliant on SOAP

> ⓘ WS-Federation was used mostly by Microsoft and IBM.

Federation Roles

When a system or user who is part of a federation needs access to an application and the local system has accepted the credentials, the system or user making the request must obtain local tokens through their own authentication process.

IdP

An IdP is responsible for providing identifiers for users looking to interact with a system, asserting to such a system that an identifier presented by a user is known to the provider, and possibly providing other information about the user that is known to the provider. This can be achieved via an authentication module, which verifies a security token that can be accepted as an alternative to repeatedly explicitly authenticating a user within a security realm.

The IdP is usually referred to as the *source* of the identity in federation. The identity provider isn’t always the authoritative source, but can sometimes rely on the authoritative source, especially if it is a broker for the process.

The IdP *holds all the identities* and generates a token for known users. The IdP is usually the **customer.**

Relying Party

The relying party is the entity that takes the authentication tokens from an identity provider and *grants access* to resources in federation. The relying party is usually the **service provider** and consumes these tokens.

Said another way, the relying party is the system that *relies* on an identity assertion from an IdP.

> ⓘ In a cloud environment, it is desirable that the organization itself continues to maintain all identities and act as the IdP.

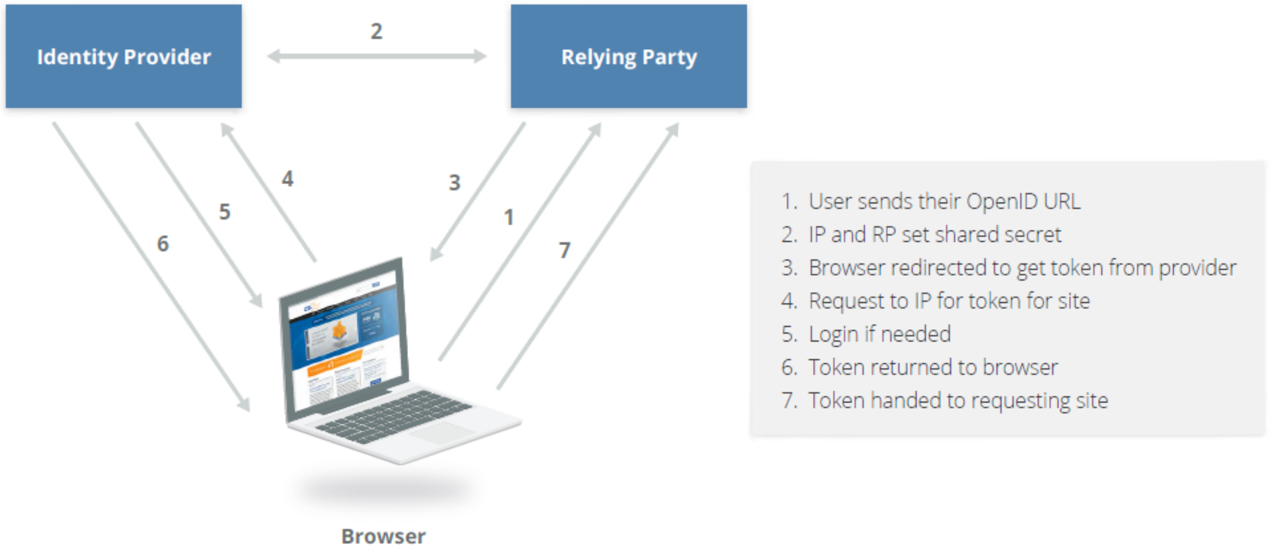

1. User sends their OpenID URL
2. IP and RP set shared secret
3. Browser redirected to get token from provider
4. Request to IP for token for site
5. Login if needed
6. Token returned to browser
7. Token handed to requesting site

How Federated Identity Management Works

SSO

Terminology

Acronyms

Acronym	Definition
FIM	Federated Identity Management
RSO	Reduced Sign-On
SCIM	System for Cross-domain Identity Management
SPML	Service Provisioning Markup Language
SSO	Single Sign-On

Definitions

Federated SSO

Federated SSO is typically used for facilitating **interorganizational** and **intersecurity** domain access to resources leveraging federated identity management.

Reduced Sign-On (RSO)

Not to be confused with SSO or federated SSO, RSO refers to not having to sign into each piece of data or store once authorization has been granted. It generally operates through some form of credential synchronization. RSO introduces security issues not experienced by SSO because the nature of SSO eliminates usernames and other sensitive data from traversing the network.

> The foundation of federation relies on the existence of an identity provider; therefore, RSO has no place in a federated identity system.

Overview

SSO refers to a situation where the user signs in once, usually to an authentication server; then when the user wants to access the organization's resources, each resource will query the authentication server to determine if the user is logged in and properly authenticated; the authentication server then approves the request and the resource server grants the user access.

SSO ensures that a single user authentication process grants access to multiple systems or even organizations.

SSO is a subset of FIM, as it relates only to *authentication* and technical interoperability. It leverages federated trusts to implement a streamlined process of SSO in the cloud account provisioning process (using SPML or SCIM).

- **Authentication** occurs a single time using SAML, OIDC, or WS-Federation.
- **Authorization** occurs using SAML or OAuth.

MFA

Terminology

Acronyms

Acronym	Definition
MFA	Multifactor Authentication
TOTP	Time-Based One Time Passwords

Definitions

Step-Up Authentication

Step-up authentication is an additional factor or procedure that validates a user's identity by using the following means:

- Challenge questions
- Out-of-band authentication (phone call or SMS)
- Dynamic knowledge-based authentication (questions unique to the end user)

Overview

Authentication requires a match against a template or a known quantity. Therefore, variables (such as variable keystrokes) aren't useful for authentication.

Multifactor authentication requires at least two of the following:

- *Type 1:* Something you know
- *Type 2:* Something you have
- *Type 3:* Something you are

In some cases, a fourth factor is mentioned:

- *Type 4:* Something you do (linked to something you are)
    - Physiological and behavioral, such as how you write or type.

This field is constantly evolving, and additional considerations include:

- TOTP
- Somewhere you are (location based)

Legal

Terminology

Definitions

Applicable Law

This determines the legal standing of a case or issue.

Common Law

The existing set of rulings and decisions made by courts, informed by cultural mores and legislation. These create *precedents,* which each party will cite in court as a means to sway the court to their own side of a case.

Culpable Negligence

Often used to prove liability.

Due Care

A company practices due care by **developing** (taking action) security policies, procedures, and standards. Due care shows that a company has taken responsibility for the activities that take place within the corporation and has taken the necessary steps to help protect the company, its resources, and employees from possible risks.

Due care is the duty owed by one entity to another, in terms of a reasonable expectation.

Due care is the minimal level of effort necessary to perform your duty to others; in cloud security, that is often the care that the cloud customer is required to demonstrate in order to protect the data it owns.

The lack of due care is often considered negligence.

> Due diligence is **understanding (researching)** the current threats and due care is **implementing countermeasures (taking action)** to provide protection from those threats.

Due Diligence

Due diligence is the act of **investigating** and understanding the risks the company faces.

Due diligence is the legal concept that describes the actions and processes a cloud customer uses to ensure that a reasonable level of protection is applied to the data in their control.

Due diligence requires that an organization continually scrutinize their own practices to ensure they are always meeting or exceeding requirements for protection of assets and stakeholders.

Due diligence is any activity taken in support or furtherance of due care.

> Due diligence is **understanding (researching)** the current threats and due care is **implementing countermeasures (taking action)** to provide protection from those threats.

Jurisdiction

This usually determines the ability of a national court to decide a case or enforce a judgment or order.

Liability

The measure of responsibility an entity has for providing due care. An organization can share risk, but it cannot share liability.

Prudent Person Rule

Based on a judge's discretion, can we demonstrate we've acted responsibly as a prudent person would? What type of action would a prudent person exercise in a particular situation?

Spoliation

The term used to describe the destruction of potential evidence (intentionally or otherwise); in various jurisdictions, it can be a crime, or the grounds for another lawsuit.

Legal Foundations

Under current laws, no cloud customer can transfer risk or liability associated with the inadvertent or malicious disclosure of personally identifiable information (PII). Your organization is ultimately responsible for any breaches or releases of data, even if you are using a cloud service and the breach/release results from negligence or attack on the part of the cloud provider. Legally and financially, in the eyes of the court, your organization is always responsible for any unplanned release of PII.

Privacy Law

Privacy can be defined as the right of an individual to determine when, how, and to what extent they will release personal information. Privacy law also typically includes language indicating that personal information must be destroyed when its retention is no longer required.

Criminal Law

Criminal law involves all legal matters where the government is in conflict with any person, group, or organization that violates *statutes.*

Statutes are rules that define conduct prohibited by the government and are designed to provide for the safety and well-being of the public. Statutes are legislated by lawmakers.

Enforcement of criminal law is called *prosecution.* Only the government can conduct law enforcement activity and prosecutions.

The burden of proof for criminal law is *beyond a reasonable doubt.*

> Examples of criminal laws include traffic laws, robbery, theft, and murder. Each have their own related set of consequences.

State Laws

State law typically refers to the law of each U.S. state, with their own state constitutions, state governments, and state courts.

Typically, federal laws supersede state laws; however, the general rule is that the most stringent of the laws apply to any situation unless there are other compelling reasons.

> Examples of state laws include speed limits, tax laws, the criminal code, etc.

Federal Laws

> Examples of federal laws include those against kidnapping and bank robbery.

Civil Law

Civil law is the body of laws, statutes, and so on that deals with personal and community-based law such as marriage and divorce. It is the set of rules that govern private citizens and their disputes.

As opposed to criminal law, the parties involved in civil law matters are strictly private entities, including individuals, groups, and organizations.

Burden of proof is typically a *preponderance of evidence.* A preponderance of evidence means that an entity that has a simple majority of fault (51 percent or more) is responsible for the *full weight* of a breach.

Contracts

A contract is an agreement between parties to engage in some specified activity, usually for mutual benefit.

Disputes that arise from failure to perform according to activity specified in the contact is known as a *breach.* In the event of a breach, a party to the contract can *sue*.

Example of contractual items that contract law applies includes:

- Service-level agreements (SLAs)
- Privacy-level agreements (PLAs)
- Operational-level agreements (OLAs)
- Payment Card Industry Data Security Standards (PCI DSS) contracts

Tort Law

Tort law refers to the body of rights, obligations, and remedies that set out reliefs for persons who have been harmed as a result of wrongful acts by others. Tort actions are not dependent on an agreement between the parties to a lawsuit.

Tort law and case precedence is what guides the courts in the handling of these civil cases whereby relief of some sort is sought.

Tort law serves four objectives:

- It seeks to compensate victims for injuries suffered by the culpable action or inaction of others.
- It seeks to shift the cost of such injuries to the person or persons who are legally responsible for inflicting them.
- It seeks to discourage injurious, careless, and risky behavior in the future.
- It seeks to vindicate legal rights and interests that have been compromised, diminished, or emasculated.

> ⓘ Negligence is the most common type of tort lawsuit.

Administrative Law

Administrative law are laws not created by legislatures, but by executive decision and function. Many federal agencies can create, monitor, and enforce their own administrative law.

> Examples of administrative law includes the federal tax law, which is administered by the IRS, who creates those laws; investigates and enforces them with IRS agents; and decides outcomes in cases tried by lawyers and adjudicated by judges both in the employ of the IRS.

International Law

International law deals with the rules that govern interactions between countries. International law is based on the premise that all nations are sovereign and equal. The value and authority of these laws is dependent upon the participation of nations in the design, observance, and enforcement.

International laws determine how to settle disputes and manage relationships between countries. These include the following:

- Conventions establishing rules expressly recognized by member countries
- Customs are they are practices in a country and accepted as law
- General principles of law recognized by civilized nations
- Judicial decisions or precedent as it has developed over time in a particular instance
- Trade regulations, including import agreements, tariff structures, and so forth
- Treaties, which can be created to solve a dispute or to create alliances

Legal Concepts

Silver Platter Doctrine

Allows law enforcement entities to use material presented voluntarily by the owner as evidence in the prosecution of crimes, without a warrant or a court order.

Doctrine of Plain View

Allows law enforcement to act on probable cause when evidence of a crime is within their presence.

Intellectual Property

Intellectual property describes creations of the mind. Intellectual property rights give the individual who created an idea an exclusive right to that idea for a defined period of time.

Copyrights

The legal protection for expressions of ideas. In the United States, copyright is granted to anyone who first creates an expression of an idea. Usually, this involves literary works, films, music, software, and artistic works.

> ⓘ Examples of copyrights include artistic works, such as books, plays, movies, songs, video games, and so on. Software is also protected by copyright.

In the United States, copyrights last for either 70 years after the author's death, or 120 years after the first publication of a work for hire.

There are a family of exceptions to copyright exclusivity. Limitations to copyrights include:

- *First sale:* selling a purchased book at a yard sale.
- *Fair use:* copies of songs can be made within reason (not well defined).

It is not mandatory to register works in order to own them. The U.S. Copyright Office allows copyright holders to register their works as means of *securing proof*.

Trademarks

Trademark protection is for intellectual property used to immediately identify a brand. It is intended to be applied to specific words and graphics.

> A trademark can be the name of an organization, or a logo, a phrase associated with an organization, even a specific color or sound, or some combination of these.

In order to have a trademark protected by law, it must be registered within a jurisdiction. Commonly, that is the U.S. Patent and Trademark Office (USPTO), the federal entity for registering trademarks. Trademarks registered with the USPTO can use the (R) symbol to signify registration. States also offer trademark registration, and trademarks registered with state offices often use the TM symbol.

Trademarks last for as long as the property they protect is still being used commercially.

Patents

Patents protect formulas, processes, materials, decorations, patterns, inventions, and plants. This includes *cryptographic algorithms*.

> Examples of patentable items include formulas for drugs, smelting processes for alloys, fabrics or textile patterns, genetically modified spices, etc.

Patents last for 20 years from the date the application was submitted, with a few exceptions.

Patent Cooperation Treaty (PCT) has been adopted by over 130 countries to provide the international protection of patents.

Trade Secrets

Trade secrets are intellectual property that involve processes, formulas, commercial methods, and so forth. Trade secrets are acknowledged as the ownership of private business material.

> Examples of trade secrets include the formula to Coca-Cola.

Protections exist for trade secrets upon creation, without any additional requirement for registration.

Intellectual property retains trade secret protection for as long as the business continues efforts to use it in commercial enterprise and maintains efforts to prevent its disclosure.

> For all intellectual property disputes, it is often up to the owners to enforce these rights. In the United States, the USPTO handles patents. Globally, the World Intellectual Property Organization (WIPO) handles patents.

Doctrine of the Proper Law

The *Doctrine of the Proper Law* is a term used to describe the processes associated with determine **what legal jurisdiction** will hear a dispute when one occurs.

Restatement (Second) Conflict of Law

The *Restatement (Second) Conflict of Law* refers to a collation of developments in common law (that is, judge made law, not legislation) that help the courts stay up with changes. Many states have conflicting laws, and judges use these *restatements* to assist them in determining **which laws should apply** when conflicts occur. The conflicting legal rules may come from federal law, state law, or laws from other countries. The factors relevant to the choice of the applicable rule of law are used. Whichever state's laws fit the situation the best or are the most restrictive are what ultimately influence the decisions.

- A *restatement* is a collation of developments in the common law that informs the judicial system of updates.
- *Conflict of laws* relates to a difference/variance between the laws.

Differentiators

Laws

Laws are legal rules that are created by *government entities* such as a congress or parliament.

> Failure to properly follow laws can result in punitive procedures that can include fines and imprisonment.

Regulations

Regulations are rules that are created by either other departments of government or external entities empowered by government.

Regulators are entities that ensure organizations are in compliance with the regulatory framework for which they are responsible. These can be government agencies, certification bodies, or parties to a contract (FTC, SEC, etc.).

The burden of proof for regulations are *more likely than not.*

> Failure to properly follow regulations can result in punitive procedures that can include fines and imprisonment.

Standards

Standards are created by other, nongovernmental organizations that provide frameworks and guidelines for business to follow. These are generally embraced by industries to provide a recognized, respectable standard for responsible, professional behavior.

> On occasion, some standards are strong enough and respected enough that laws or regulation are passed that establish them as the de facto standard of legal expectations as well.

Contracts and Frameworks

Contracts do not derive authority from the government. For instance, **Payment Card Industry (PCI)** compliance is wholly voluntary (a contractual standard), but is also a regulated requirement for those who choose to participate in credit card processing. Those participants agree to submit to PCI regulation, including audits and controls. It's not a law, but it is a regulatory framework, complete with regulators.

Clarification

United States

In the United States, there is no single federal law governing data protection. The FTC and other associated U.S. regulators hold that the applicable U.S. laws and regulations apply to data after it leaves its jurisdiction, and U.S. regulated entities remain liable for the following:

- Data exported out of the United States
- Processing of data overseas by subcontractors
- Subcontractors using the same protections for the regulated data when it leaves the country

EEA

Regulations

Regulations have binding legal force throughout *every* Member State and enter into force on a set date in all the Member States. It mandates that all countries comply with the regulation.

Directives

Directives lay down certain results that must be achieved but each Member State is free to decide how to transpose directives into national laws.

A directive allows member states to create their own laws; it allows every member country to create its own law this is compliant with the directive.

Decisions

Decisions are laws relating to specific cases and directed to individual or several Member States, companies or private individuals. They are binding upon those to whom they are directed.

Risk

Terminology

Acronyms

Acronym	Definition
ERM	Enterprise Risk Management
KRI	Key Risk Indicator

Definitions

Asset

Anything of value to a company.

ERM

ERM includes managing overall risk for the organization, aligned to the organization’s governance and risk tolerance. Enterprise risk management includes all areas of risk, not merely those concerned with technology. It is the overall management of risk for an organization.

Exploit

An instance of compromise.

Natural Disasters

Natural disasters are a category of risks to cloud deployments that is based *solely* on geography and location. Regulatory violations are not always based on location. They can also be based on the type of industry.

Reputational Risk

Reputational risk is the loss of value of a brand or the ability of an organization to persuade.

Residual Risk

The risk that exists *after* controls have been implemented.

Risk Appetite

Risk appetite is the total risk that the organization can bear in a given risk profile, usually expressed in aggregate. Risk appetite is set by senior management and is the level, amount, or type of risk that the organization finds acceptable.

- Organizations accept a level of risk that allows operations to continue in a successful manner.
- It is legal and defensible to accept risks higher than the norm, or greater than your competitors, except risks to health and human safety; these risks *must* be addressed to the industry standard or whatever regulator motif to which your organization adheres.

As risk appetite or tolerance increases, so does the willingness to take greater and greater risks.

Risk Owners and Players

These are the individuals in the organization who together determine the organization's overall risk profile. For example, while one department may be willing to take moderately high risks in engaging cloud activities, another may have a lower risk tolerance. It is the aggregate of these individual tolerances that determines the organization's overall risk appetite.

Risk Profiles

The risk profile of the organization is a comprehensive analysis of the possible risks the organization is exposed to. It lists the identified risks and their potential effects.

The risk profile is determined by an organization's willingness to take risks as well as the threats to which it is exposed. The risk profile should identify the level of risk to be accepted, the way risks are taken, and the way risk-based decision making is performed.

Additionally, the risk profile should take into account potential costs and disruptions should one or more risks be exploited.

Risk Tolerance

Risk tolerance is the level of risk that an organization can accept per *individual* risk.

Secondary Risk

When one risk response triggers another risk event. For example, a fire suppression system that displaces oxygen is a means to mitigate the original risk (fire) but adds a new risk (suffocating people).

Threat

Something that could cause loss to all or part of an asset.

Threat Agent

Something or someone that carries out the attack.

Total Risk

The risk that exists *before* any controls are implemented.

Types of Risk

Policy and Organization Risks

- Provider lock-in
- Loss of governance
- Compliance risks
- Provider exit (vendor lock-out)

General Risks

- Impact of SPOF
- Increased need for technical skills
- Provider assumes more control over technical risks (loss of governance)

Virtualization Risks

- Guest breakout
- Snapshot and image security
- Sprawl

Cloud-Specific Risks

- Management plane breach
- Resource exhaustion
- Isolation control failure
- Insecure or incomplete data deletion
- Control conflict risk
- Software-related risks

Legal Risks

- Data protection
- Jurisdiction
- Law enforcement
- Licensing

Non-Cloud Specific Risks

- Natural disasters
- Unauthorized access
- Social engineering
- Default passwords
- Network attacks

Risk Identification

Terminology

Definitions

Risk Modeling

Risk modeling is based on an asset or a threat.

Risk Register

A risk register is a table to consolidate information about risks.

Overview

Risk = Asset * Threat * Vulnerability (in some cases, the asset is excluded but shouldn't be)

What the above calculation displays is that if there is no threat, no there is no vulnerability. Likewise, if there is no vulnerability, there is no risk. Finally, if the asset has no value, there is also no risk.

Risk Management

Overview

Every decision we make should be based on risk. Risk management in the cloud is based on the **shared responsibilities** model.

There are four steps to adequately managing risk:

→ Framing Risk /concepts/risk/risk-management/framing-risk

→ Assessing Risk /concepts/risk/risk-management/assessing-risk

→ Responding to Risk /concepts/risk/risk-management/responding-to-risk

→ Monitoring Risk /concepts/risk/risk-management/monitoring-risk

Framing Risk

Overview

Addresses how organizations describe the environment in which risk-based decisions are made. Refers to determine what risk and levels are to be evaluated.

Risk framing is designed to produce a risk-management strategy intended to address how organizations assess, respond to, and monitor risk. This allows the organization to clearly articulate the risks that it needs to manage, and it establishes the boundaries for risk-based decisions within organizations.

Assessing Risk

Terminology

Acronyms

Acronym	Definition
ALE	Annualized Loss Expectancy
ARO	Annual Rate of Occurrence
AV	Asset Value
EF	Exposure Factor
SLE	Single Loss Expectancy

Definitions

ALE

The annualized loss expectancy is a product of the yearly estimate for the exploit (ARO) and the loss in value of an asset after an SLE.

```
ALE = SLE * ARO
```

ARO

The annual rate of occurrence is an estimate of how often a threat will be successful in exploiting a vulnerability over the period of a year.

While previous activity is not a great predictor of future outcomes, historical data can sometimes help you predict the annualized rate of occurrence for a specific loss.

AV

The asset value is represented as a dollar value.

Delphi Technique

The Delphi technique is a qualitative method of group decision-making that involves successively collating the judgment of the group.

EF

The exposure factor is the estimated loss that will result if the risk occurs. The threat vector is the multiplier involved in determining exposure factor.

The exposure factor is represented as a percentage.

Impact

Impact includes loss of life, loss of dollars, loss of prestige, loss of market share, and other facets. Unlike risk, impact uses definitions rather than an ordinal scale for measuring.

Risk

The probability of a threat materializing; the likelihood an impact will be realized.

- Risk can be reduced but never eliminated.
- Risks arise from the loss of CIA of information or information systems.
- Risks reflect the potential adverse impacts to organizational operations (mission, functions, image, or reputation), organizational assets, individuals, or other organizations.

SLE

Single-loss expectancy must be calculated to provide an estimate of loss. SLE is defined as the difference between the original value and the remaining value of an asset after a single exploit.

The monetary value of the asset is the most objective, discrete metric possible and the most accurate for the purposes of SLE determination. In the formula below, SLE is a dollar value.

```
SLE$ = AV$ * EF%
```

Threat

Any **circumstance or event with the potential to adversely impact** organizational operations and assets, individuals, other organizations, or the Nation through an information system via unauthorized access, destructions, disclosure, or modification of information, and/or denial-of-service.

Any potential danger that is associated with the exploitation of a vulnerability.

Threat Source

Either intent and method targeted at the intentional **exploitation of a vulnerability** or a situation and method that may **accidentally trigger a vulnerability.**

Threat Vector

A threat vector is the multiplier involved in determining exposure factor.

Vulnerability

An inherent **weakness** in an information system, security procedures, internal controls, or implementation that could be exploited by a threat source.

A lack of a countermeasure or a weakness in a countermeasure that is in place.

Overview

Risk assessment is the process used to identify, estimate, and prioritize information security risks.

Risks must be communicated in a way that is clear and easy to understand. It may also be important to communicate risk information outside the organization. To be successful in this, the organization must agree to a set of risk-management metrics.

Risk is determined as the by-product of likelihood and impact. For example, if an exploit has a likelihood of 1 (high) and an impact of 100 (high), the risk would be 100. As a result, 100 would be the highest exploit ranking available.

> ⓘ No countermeasure should be greater in cost than the risk it mitigates, transfers, or avoids.

The purpose of engaging in risk assessment is to identify:

- Threats to organizations (i.e., operations, assets, or individuals) or threats directed through organizations against other organizations.
- Vulnerabilities internal and external to organizations.
- The harm (i.e., adverse impact) that may occur given the potential for threats exploiting vulnerabilities.
- The likelihood that harm will occur.

There are additional cloud-specific risk concerns that should also be considered:

- Risk of service failure and associated impact
- Insider threat risk impact
- Risk of compromised customer to other tenants in the cloud environment
- Risk of DoS attacks
- Supply chain risk to the CSP

Identifying these factors helps to determine risk, which includes the likelihood of harm occurring and the potential degree of harm. Using a risk scorecard is recommended. The impact (consequence) and probability (likelihood) of each risk are assessed separately, and then the results are combined.

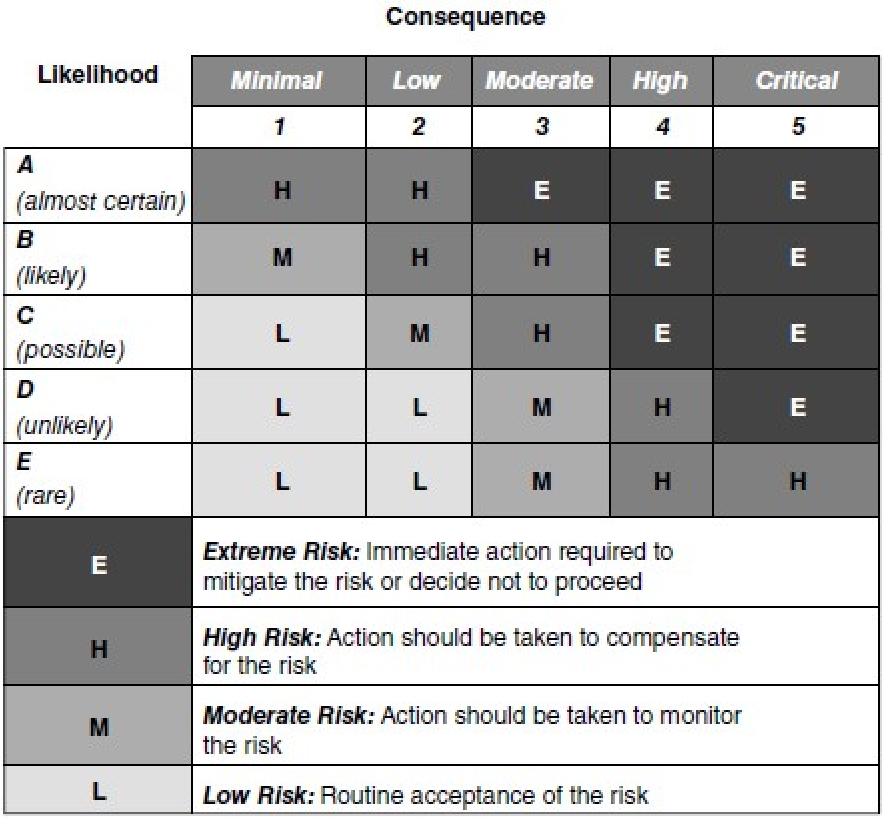

Consequence

Likelihood	*Minimal*	*Low*	*Moderate*	*High*	*Critical*
	1	*2*	*3*	*4*	*5*
A *(almost certain)*	H	H	E	E	E
B *(likely)*	M	H	H	E	E
C *(possible)*	L	M	H	E	E
D *(unlikely)*	L	L	M	H	E
E *(rare)*	L	L	M	H	H

E	***Extreme Risk:*** Immediate action required to mitigate the risk or decide not to proceed
H	***High Risk:*** Action should be taken to compensate for the risk
M	***Moderate Risk:*** Action should be taken to monitor the risk
L	***Low Risk:*** Routine acceptance of the risk

Example of a Risk Scorecard

Threat Categories

- Human
- Natural
- Technical
- Physical
- Environmental
- Operational

Risk Assessments

> Risk assessments should be performed *periodically*.

Qualitative Risk Assessments

Qualitative risk assessments typically employ a set of methods, principles, or rules for assessing risk based on **non-numerical** categories or levels (such as high, medium, or low). Qualitative risk assessments use subjective analysis to help prioritize probability and impact of risk events.

The opinion of an expert in the field is a prime source for qualitative analysis. Interviews and risk workshops are two ways in which you can work with experts on managing qualitative risk.

Characteristics

- Assessors have limited expertise in quantitative assessments (assessors do not require as much experience when performing a qualitative assessment)
- The timeframe to complete the assessment is short
- Implementation is easier
- The organization doesn't have enough data for a quantitative assessment
- The assessors are long-term employees who have experience with the *business* and *critical systems*
- Results that are descriptive versus measurable

Process

- Management approval is obtained and management is kept informed
- A risk-assessment team can be formed. Members may include staff from senior management, IS, legal or compliance, internal audit, HR, facilities and safety coordination, IT, and business owners, as appropriate.
- The assessment team requests documentation
- The team sets up interviews with organizational members to identify vulnerabilities, threats, and countermeasures. All levels of staff should be represented.
- The analysis of the data gathered can be completed, which typically includes matching the threat to a vulnerability, matching threats to assets, determining how likely the threat is to exploit the vulnerability, and determining the impact to the organization in the event an exploit is successful. Analysis also includes matching of current and planned countermeasures to the threat-vulnerability pair.

- When matching is completed, risk can be calculated. In qualitative analysis, the product of likelihood and impact produces the level of risk.
- Once risk is determined, additional countermeasures can be recommended to minimize, transfer, or avoid the risk.
- When this is completed, the risk that is left over-after countermeasures have been applied to protect against the risk-is also calculated. This is the **residual risk.**

Quantitative Risk Assessments

Quantitative risk assessments typically employ a set of methods, principles, or rules for assessing risk based on the use of **numbers.**

The hallmark of a quantitative assessment is the numeric nature of the analysis. Frequency, probability, impact, countermeasure effectiveness, and other aspects of the risk assessment have a discrete mathematical value in a pure quantitative analysis.

Characteristics

- Allows the assessor to determine whether the cost of the risk outweighs the cost of the countermeasure
- Requires a lot of time
- Must be performed by assessors with a significant amount of experience and particular skillset
- Subjectivity is introduced because the metrics may need to be applied to qualitative measures
- This type of assessment most effectively supports **cost-benefit analyses**

> ⓘ Most organizations are not in a position to authorize the level of work required for a quantitative risk assessment.

Process

Three steps are undertaken in a quantitative risk assessment:

1. Management approval

2. Define risk assessment team
3. Review information available within the organization

> ⓘ Quantitative risk assessments often use values such as AV, EF, SLE, ARO, and ALE to quantify risk. Quantitative assessments lead to the proper mitigation strategy by specifying a **dollar value** for risks.

> ⓘ Often, the risk assessment an organization conducts is a combination of qualitative and quantitative methods. Fully quantitative risk assessment may not be possible because there is always some subjective input present, such as the value of information. Value of information is often one of the most difficult factors to calculate.

Vulnerability Assessments

Unlike a risk assessment, vulnerability assessments tend to focus on the technology aspects of an organization, such as the network or applications. Data gathering for vulnerability assessments typically includes the use of software tools.

Responding to Risk

Overview

Risk response provides a consistent, organization-wide response to risk in accordance with the organizational risk frame by taking these steps:

- Developing alternative courses of action for responding to risk
- Evaluating the alternative courses of action
- Determining appropriate courses of action consistent with organizational risk tolerance
- Implementing risk responses based on selected courses of action

Monitoring Risk

Terminology

Acronyms

Acronym	Definition
KGI	Key Goal Indicator
KPI	Key Performance Indicator
KRI	Key Risk Indicator

Definitions

KGIs

KGIs are examined *after* the fact. For example:

- Did we meet our goal of 99.99% uptime?

KPIs

KPIs are examined *before* you meet your goals. For example:

- Our SLA states that we will have 99.99% uptime.
- Did we meet this SLA for the quarter?

If you did not meet your SLA for the quarter, you will likely not reach your performance goals for the year.

KRIs

KRIs are those items that will be the first things that let you know something is amiss. You need to identify and closely monitor the things that will most quickly alert you to a change in the risk environment. In cloud computing, this might be the announcement of the discovery of a new vulnerability that could impact your cloud provider.

KRIs examine what might cause you to not meet your performance. For example:

- You may receive alerts for processor utilization of 60% sustained for 5 minutes or longer on a particular server to detect a potential DoS attack.

The idea is that if you can receive an indication that a risk is about to materialize, you can proactively move to counter that risk so you can meet your goals.

Overview

Risk monitoring is the process of keeping track of identified risks. It should be treated as an ongoing process and implemented throughout the system life cycle.

The most important elements of a risk monitoring system include the ability to:

- Clearly identify a risk
- Classify or categorize the risk
- Track the risk over time

There are three purposes of the risk-monitoring components:

- Determine the ongoing effectiveness of risk responses
- Identify risk-impacting changes
- Verify that planned risk responses are implemented and inline with organizational mission

Risk Response

Organizations have four primary ways to address risk, all of which should take into effect the cost-benefit analysis (the potential cost should not outweigh the benefit). For example, we do not put a $10 lock on a $5 bicycle.

Avoidance

Leaving a business opportunity because the risk is simply too high and cannot be compensated for with adequate control mechanisms-a risk that exceeds the organization's appetite.

> This is the only surefire method for eliminating a specific risk.

Acceptance

The opposite of avoidance; the risk falls within the organization's risk appetite, so the organization continues operations without any additional efforts regarding the risk.

The only reason organizations accept any level of risk is because of the potential benefit also afforded by a risky activity. Risk is often balanced by corresponding **opportunity**.

> If the risk of the activity is estimated to be within the organization's risk appetite, then the organization might choose risk acceptance.

Transference

The organization pays someone else to accept the risk, at a lower cost than the potential impact that would result from the risk being realized; this is usually in the form of insurance. This type of risk is often associated with things that have a low probability of occurring but a high impact should they occur.

This typically takes the form of insurance, but contracts and SLAs are another method of transferring risk.

Mitigation

The organization takes steps to decrease the likelihood or the impact of the risk (and often both); this can take the form of controls/countermeasures, and is usually where security practitioners are involved.

When we choose to mitigate risk by applying countermeasures and controls, the remaining, leftover risk is called *residual risk.* The task of the security program is to reduce residual risk until it falls within the acceptable level of risk according to the organization's risk appetite.

Rejection

Risk rejection isn't a legitimate form of risk response. Rejection involves ignoring risks and continuing with operations. This is essentially the same thing as accepting a risk without mitigating it to a tolerable level. In most cases, this would indicate failure of due diligence and may put the organization in a position of liability.

Risk Controls

Administrative Controls

Administrative controls are those processes and activities that provide some aspect of security.

> ⓘ Examples include personnel background checks, scheduled routine log reviews, mandatory vacations, robust and comprehensive security policies and procedures, and deciding business processes so that there are no single points of failure and so that proper separation of duties exists.

Technical Controls

Technical (logical) controls, are those controls that enhance some facets of the CIA triad, usually operating within a system, often in electronic fashion.

> ⓘ Technical controls include encryption mechanisms, access control lists to limit user permissions, and audit trails and logs of system activity.

Physical Controls

Physical controls are controls that limit physical access to assets or that operate in a manner that reduces the impact of a physical event.

> ⓘ Examples include locks on doors, fire suppression equipment in datacenters, fences, and guards.

Compensating Controls

Controls to catch the failure of a first control.

- Intent and rigor of the original equipment
- Provide a similar level of defense as the original requirement
- Be above and beyond other requirements
- Be commensurate with the additional risk imposed by not adhering to the requirement

Software

Development Types

Agile

Characteristics

- Often involves daily meetings called Scrums
- Favors customer collaboration and prototyping instead of an elaborate contract mechanism
- Works in short, iterative work periods (between a week and month in duration)
- Prototyping is favored over testing
- Relies on cooperative development instead of expertise
- Does not depend on planning

Cloud Applications

Terminology

Definitions

Forklifting

The idea of moving an existing legacy enterprise application to the cloud with little or no code changes.

Overview

Applications can be broken down into the following subcomponents:

- Data
- Functions
- Processes

Deployment Pitfalls

The rapid adoption of cloud services has caused a disconnect between providers and developers on how to best meet development requirements. Pitfalls of moving applications to the cloud:

- *On-premise does not always transfer.* On-premises transfer is a pitfall that occurs when a company transfers its applications and configurations to the cloud. APIs that were developed to use on-premises resources may not function properly in the cloud or provide the appropriate security.
- *Not all apps are cloud-ready.*
- *Lack of training and awareness.*

- *Lack of documentation and guidelines.* For example, the rapid adoption of ever-evolving cloud services may result in outdated documentation or no documentation whatsoever.
- *Complexities of integration.* Results from when developers and administrators are used to having full access to on-premise components, servers, and network equipment to make integration of services and systems seamless moving to the cloud where access to these types of systems and services are limited. When new applications need to interface with old applications, developers often do not have unrestricted access to the supporting services.
- *Overarching challenges.* Tenancy separation and the use of secure, validated APIs.

SDLC*

Terminology

Acronyms

Acronym	Definition
IaC	Infrastructure as Code
SDLC	Software Development Lifecycle

Definitions

Application Accreditation

The application has been approved by management.

Application Certification

The application meets technical requirements.

IaC

A type of IT setup wherein developers or operations teams automatically manage and provision the technology stack for an application through software, rather than using a manual process to configure discrete hardware devices and operating systems.

Overview

The purpose of the SDLC to ensure that applications are properly secured prior to an organization using them. Several software development lifecycles have been published with most of them containing similar phases.

For the purposes of the exam I have consolidated a few popular models into a single model using the following phases:

1. Analysis (planning and requirements analysis)
2. Define
3. Design
4. Develop (implement)
5. Test (verify)
6. Deploy/Release/Maintain
    1. Secure Operations (could also be considered part of deploy/release)
    2. Secure Disposal

Lifecycle Process

Verification and validation should occur at **every stage** of development and during the software development lifecycle, and in line with change management components.

1. Analysis

In this phase, business and security requirements and standards are being determined. Calls for all business requirements are to be defined even before initial design begins:

- Functional requirements are determined
- Nonfunctional requirements are determined
- Planning for QA requirements and identification of risks
- Feasibility study (especially if bound by regulations such as HIPAA)
- **Security** requirements

The requirements are then analyzed for their validity and the possibility of incorporating them into the system to be developed.

The analysis phase of the SDLC is when requirements of the project are put into a project plan. This plan will outline the specifications for the features and functionality of the software or application to be created. ~~At the end of the analysis phase, there will be formal~~

~~requirements and specifications ready for the development team to turn into actual software.~~

2. Define

Identify the business needs of the application. Refrain from choosing any specific tools or technology at this phase, as it's too early to make these decisions.

We're trying to determine the purpose of the software, in terms of meeting the user's needs; therefore, we may solicit input from the user community in order to determine what they want. Develop *user* stories. The following questions should be answered:

- What will the *user* want to accomplish and how will you approach it?
- What will the user interface look like?
- Will it require the use or development of any APIs?

The defining phase is meant to clearly define and document the **product requirements** to place them in front of the customers and get them approved. This is done through a requirement specification document, which consists of all the product requirements to be designed and developed during the project lifecycle. For example, if we determined during the analysis phase that a regulation such as HIPAA is required, here is where we specify that we need 256-bit encryption.

> User involvement is most crucial in this stage. While some development models allow for user involvement in the entirety of the process, user input is most necessary in this phase, where developers can understand the user requirements-what the system/software is actually supposed to produce, in terms of function and performance.

3. Design

Business requirements are most likely to be mapped to software construction in this phase.

System design helps in specifying hardware and system requirements and helps in defining overall system architecture. Formal requirements for risk mitigation/minimization are integrated with the programming designs. The system design specifications serve as input.

> ⓘ While the requirements for risk mitigation and minimization may be *determined* during the *analysis* phase of the SDLC, they are not integrated with the programming designs until the design phase of the SDLC.

This is the phase where we would want to identify what **programming language** and architecture we will use as well as specific hardware and system requirements. Threat modeling and secure design elements should also be undertaken and discussed here. For example, since we need to use 256-bit encryption, we should choose AES. Additionally, MFA should be implemented using passwords and retina scans.

Logical design is the part of the design phase of the software development lifecycle in which all *functional* features of the system chosen for development in the analysis phase are described independently of any computer platform.

4. Develop

Upon receiving the system design documents, work is divided into modules or units and **actual coding starts**. This is typically the **longest** phase of the software development lifecycle.

Activities include:

- Code review
- Unit testing
- Static analysis

As each portion of code is created and completed, *functional testing* is done on it by the *development* team. This testing is done to ensure that it compiles correctly and operates as intended.

5. Test

After the code is developed, it is tested against the requirements to make sure that the product is actually solving the needs gathered during the requirements phase.

In the testing phase, we will use techniques and tools such as:

- Unit testing
- Integration testing
- System testing
- Acceptance testing (users)
- Certification and accreditation (management)

This includes DAST and SAST testing, vulnerability assessments, and penetration tests. Functional *and* nonfunctional (security) testing are performed. If the application does not work **securely,** it does not work at all.

- *Functional testing.* Does the application do what we designed it for?
- *Nonfunctional testing.* Does the application do what we designed it for *in a secure manner, with cool graphics, and bells and whistles?*

→ Application Security Testing /concepts/software/security-testing

6. Maintain

Most software development lifecycle models include a maintenance phase as their endpoint. Overall, this phase includes continuous monitoring and updates as needed, as well as disposal.

The maintenance phase will go on through the entire lifetime of the software or application. The maintenance phase includes pushing out continual updates, bug fixes, security patches, and anything else needed to keep the software running securely and operating as it should.

6.1 Secure Operations

We enter this phase after thorough testing has been successfully completed and the application and its environment are deemed secure.

Proper software configuration management and versioning are essential to application security. The following applications can be useful:

- Puppet
- Chef

This phase calls for the following activities to take place:

- Dynamic analysis
- Vulnerability assessments and penetration testing
- Activity monitoring
- Layer 7 firewalls (such as WAFs)

6.2 Secure Disposal

Once the software has completed its job or has been replaced by a newer or different application, it must then be securely disposed of.

APIs

Terminology

Acronyms

Acronym	Definition
API	Application Programming Interface
REST	Representational State Transfer
SOAP	Simple Object Access Protocol

Overview

APIs allow secure communication from one web service to another. This is typically an interface that allows for proper communication to occur. They are the coding components that allow applications to speak to one another, generally through a web interface of some kind.

> ✓ **Fact.** APIs consume tokens rather than traditional usernames and passwords.

> ⓘ Regardless of what type of API you use to offer web services, you are granting another application access to the primary application and any data it may have access to.

Types of APIs

REST

REST is a software architecture style consisting of guidelines and best practices for creating scalable web services. It allows web applications to access other applications, databases, and so on in order to extend their functionality. REST is *not a protocol.* It is a class/category of APIs.

The server does not need to store any temporary information about clients, so **sessions are not required**. Credentials are used to allow authentication between clients and servers. Generally, a REST interaction involves the client asking the server for data, sometimes as the result of processing; the server processes the request and returns the result. In REST, an enduring session, where the server has to store some temporary data about the client, is not necessary.

REST performs web service requests using uniform resource identifiers (URIs).

Request verbs describe what you will do with a resource. The most common HTTP verbs are POST, GET, PUT, PATCH, and DELETE. Some would say these are properly called "methods." REST HTTP methods correspond to CRUD methods:

- [C]reate (POST)
- [R]ead (GET)
- [U]pdate (PUT)
- [D]elete (DELETE)

REST is based on 5 principles:

- RESTful client-server
- Stateless (sessionless)
- Cache
- Layered System
- Uniform Contract

Characteristics

- Lightweight (no envelopes required)
- Caching (performant)
- Scalable
- Uses simple URLs
- Not reliant on XML
- Efficient (smaller messages than XML)
- Reliant on HTTP/S (HTTP/S-only)
- Supports multiple data formats (CSV, JSON, YAML, XML, etc.)
- Widely used

Uses

- When bandwidth is limited
- When *stateless* operations are used
- When caching is needed

Advantages

- No expensive tools required to interact with the web service
- Smaller learning curve
- Efficient (SOAP uses XML for *all* messages, REST can use smaller message formats); REST uses what is called "postcards"
- Fast (no extensive processing required)
- Closer to other web technologies in design philosophy

> ⓘ REST is seemingly the most widely used since it uses HTTP, which is everywhere.

SOAP

SOAP is a protocol specification for exchanging structured information in the implementation of web services in computer networks.

SOAP allows programs to operate independently of the client operating system.

- SOAP envelope and then HTTP, FTP, or SMTP
  - Since everything must be "put in an envelope and addressed properly" it adds overhead
- Provides WS-* features
- Should only be used when REST is not available

SOAP uses message-level encryption.

Characteristics

- Standards-based
- Reliant on XML (XML-only)
- Reliant on SAML (SAML-only)
- Slower (no caching)
- Highly intolerant of errors
- Built-in error handling

Uses

- Asynchronous processing
- Format contracts
- *Stateful* operations

Advantages

- Language, platform, and transport independent
- Works well in distributed enterprise environments
- Standardized
- Provides significant pre-build extensibility in the form of the WS* standards
- Built-in error handling
- Automation when used with certain language products

> ⓘ SOAP should only be used when REST is not possible. Banks typically use SOAP because they can hide the business logic using the envelope and the envelope can be encrypted, which adds a layer of security as it is able to hide some of the logic in the envelope.

API Security

- In order to detect possible erroneous or malicious modification of the organization's data by unauthorized or security-deficient APIs, it's important to take representative samples of the production data on a continual basis and perform integrity checks.
- Because untrusted APIs may not be secured sufficiently, increased vigilance for the possibility of introducing malware (such as by using antimalware detection capabilities) into the production environment is essential.

Application Security Testing

Terminology

Acronyms

Acronym	Definition
DAST	Dynamic Application Security Testing
RASP	Runtime Application Self-Protection
SAST	Static Application Security Testing

Definitions

Black-Box Testing

Testing the program as it functions, in runtime.

Event-Driven Security

Automates detection and remediation of security issues.

Functional Testing

Functional testing is performed to confirm the functional aspect of a product. It verifies that a product is performing as expected based on the initial requirements laid out.

OWASP Dependency-Check

Dependency-Check is a Software Composition Analysis (SCA) tool that attempts to detect publicly disclosed vulnerabilities contained within a project's dependencies. It does this by determining if there is a Common Platform Enumeration (CPE) identifier for a given dependency. If found, it will generate a report linking to the associated CVE entries.

A utility that identifies project dependencies and checks whether there are any known, publicly disclosed, vulnerabilities.

Regression Testing

Re-running functional and non-functional tests to ensure that previously developed and tested software still performs after a change. If not, that would be called a regression.

Source Code Analysis

Performing an analysis of the source code, byte code, and binaries.

Software Assurance

Encompasses the development and implementation of methods and processes for ensuring that software functions as intended while mitigating the risks of vulnerabilities, malicious code, or defects that could bring harm to the end user.

Software-Defined Security

Automates security controls.

White-Box Testing

Reviewing the source code.

Overview

Security of applications must be viewed as a holistic approach in a broad context that includes not just software development considerations but also the business and regulatory context and other external factors that can affect the overall security posture of the applications being consumed by an organization.

OWASP makes several recommendations for testing:

- Identity management testing

- Authentication testing
- Authorization testing
- Session management testing
- Input validation testing
- Testing for error handling
- Testing for weak cryptography
- Business logic testing
- Client-side testing

Types of Testing

SAST

Overview

Source code, byte code, and binaries are all tested without executing the application. This type of testing is often used in the early stages of application development as the full application is not testable in any other way at that time. SAST is a **white-box** test, meaning that the tester has knowledge of and access to the source code.

- Performs an analysis of the source code, byte code, and binaries; SCA.
- Performed in an offline manner; SAST does *not* execute the application.

Uses

- Used to determine coding errors and omissions that are indicative of security vulnerabilities (XSS, SQL injection, buffer overflows, unhandled error conditions, and potential backdoors).
- SAST is often used as a test method while the tool is *under development* (early in the development lifecycle).
- SAST is excellent for DevOps or CI/CD (continuous integration/continuous development) environments. It is becoming the most popular option due to the growing needs of the new development era.
- SAST can help identify the following flaws:
  - Null pointer reference

    - Threading issues
    - Code quality issues
    - Issues in dead code
    - Insecure crypto functions
    - Issues in back-end application code
    - Complex injection issues
    - Issues in non-web app code

Advantages

- Because SAST is a white-box test, it usually delivers more results and more accuracy than DAST.

Goals

- Attempt to catch flaws prior to going into production.

Thoroughness

SAST uses **code coverage** as a measure of how thorough testing was. For example, "SAST covered 90% of the source code."

DAST

Overview

DAST is considered a black-box test since the code is not revealed and the test must look for problems and vulnerabilities while the application is running. The tester is not given any special information about the systems they are testing. DAST is performed on live systems.

- Must discover individual execution paths in the application.
- Used to analyze code in it's running state.

Uses

- DAST is most effective when testing exposed interfaces of web applications:
    - HTTP, HTML
- DAST is well-suited for waterfall environments.

- DAST can help identify the following flaws:
    - Environment configuration issues
    - Patch level issues
    - Runtime privileges issues
    - Authentication issues
    - Protocol parser issues
    - Session management issues
    - Issues in 3rd party web components
    - Malware analysis

Thoroughness

DAST uses **path coverage** as a measure of how thorough testing was. The objective is to test a significant sample of the possible logical paths from data input to output.

> ⓘ Web vulnerability testing and fuzzing are considered DAST tests.

RASP

Overview

RASP is generally considered to focus on applications that possess self-protection capabilities built into their runtime environments, which have full insight into application logic, configuration, and data and event flows.

RASP prevents attacks by self-protecting or reconfiguring automatically without human intervention in response to certain conditions (threats, faults, and so on).

Unlike firewalls which rely solely on network data, RASP leverages the applications intrinsic knowledge of itself to accurately differentiate attacks from legitimate traffic.

Vulnerability Assessments

Vulnerability scanning is a test that is run on systems to ensure that systems are properly hardened and there are not any *known* vulnerabilities on the system.

Most often, vulnerability assessments are performed as **white-box** tests, where the assessor knows the application and the environment the application runs in.

> Any vulnerabilities not currently *known* and included in the scanning tool will remain in place and can later become zero-day exploits.

Penetration Testing

Penetration testing is a type of test in which the tester attempts to break into systems using the same tools that an attacker would to discover vulnerabilities. These tests usually begin with a vulnerability scan to identify system weaknesses and then move on to the penetration phase.

This is a process used to collect information related to system vulnerabilities and exposures, with the view to *actively exploit* the vulnerabilities in the system.

Penetration is often a **black-box** test, in which the tester carries out the test as an attacker, has no knowledge of the application, and must discover any security issues within the application or system being tested.

> To assist with targeting and focusing the scope of testing, independent parties also often perform **gray-box** testing, with only **some level of information provided.**

Secure Code Reviews

Open Source Review

Can detect flaws that a structured testing method might not.

Fuzzing

Fuzzing is an automated software testing technique that involves providing invalid, unexpected, or random data as inputs to a computer program. The program is then monitored for exceptions such as crashes, or failing built-in code assertions, or for finding potential memory leaks.

Fuzzing is a **black-box** test.

Threat Modeling

Threat modeling is the practice of viewing the application from the perspective of a potential attacker. In this way, threat modeling can be seen as an **application-specific** form of penetration. The idea is to identify specific points of vulnerability and then implement controls or countermeasures to protect or thwart those points from successful exploitation.

There are many ways to perform threat modeling. For instance, a vulnerability scan could be perceived as a low-level sort of threat model since attackers will be looking for the same (known) vulnerabilities that a vulnerability scan looks for.

Threat modeling also makes use of use/misuse cases. There are typically charts that display actions or functions that a user will take as well as an attacker. You then define the controls required to mitigate them.

→ Threat Models /models-and-guidance/threat-models

Technology

AI

Terminology

Definitions

Analytical Artificial Intelligence

Solely cognitive-based, focusing on a system-s ability to analyze past data and make future decisions.

Artificial Intelligence

Artificial intelligence is the ability of devices to perform human-like analysis. Artificial intelligence operates by consuming a large amount of data and recognizing patterns and trends in the data.

Contains humanized and analytical types.

Humanized Artificial Intelligence

Incorporates emotional intelligence, cognitive learning and responses, and also expands to include social intelligence.

Blockchain

Models

Private

A private blockchain is an invitation-only network governed by a single entity. Entrants to the network require permission to read, write or audit the blockchain. There can be different levels of access and information can be encrypted to protect commercial confidentiality

Public

A public blockchain has an open network. The information is available in a public domain. Due to its permissionless nature, any party can view, read, and write data on the blockchain and the data is accessible to all. No particular participant has control over the data in a public blockchain.

Consortium

A consortium blockchain operates in a semi-private manner. It requires permission to join, but it can be shared and utilized by numerous different organizations that are working together.

Hybrid

A hybrid blockchain combines benefits of public and private blockchains: an application or service can be hosted on an independent permissioned blockchain while leveraging a public blockchain for security and settlement.

Containers

Overview

In containerization, the underlying hardware is **not emulated**; the container(s) run on the same underlying kernel, sharing the majority of the base OS.

Includes:

- OS replication
- A single kernel
- The possibility for multiple containers

Training

Types of Training

Training

The formal presentation of material, often delivered by internal subject matter experts. It addresses and explains matters of the organization's policies, content mandated by regulation, and industry best practices for the organization's field.

Education

The formal presentation of material in an academic setting, often for credit toward a degree.

Awareness

The additional, informal, often voluntary presentation of material for the purpose of reminding and raising attention among staff.

Training Program Categories

Initial Training

Initial training is delivered to personnel when they first enter the employ of the organization. Often thorough and comprehensive, this should be mandatory for all personnel, regardless of their position or role. The content should be broad enough to address the security policies and procedures all staff will be expected to understand and comply with, but it should have sufficient specificity so that everyone knows hot to perform basic security functions.

Topics that might be covered could include the following:

- Password policy
- Physical security
- The use of any security credentials or tokens
- How to report security concerns
- The acceptable use policy (AUP)

Recurring Training

Recurring training is for continual updating of security knowledge that builds on the fundamentals taught in the initial training session. This should be done on a regular basis, on a schedule according to the needs of the organization, regulatory environment, and industry fluctuations. At the very least, each employee should receive recurring training annually.

- Any updates and modifications to security practices and procedures
- Changes to regulations and policies
- Introduction of any new elements in the infrastructure

Refresher Training

Refresher training sessions are offered to those personnel who have demonstrated a need for additional lessons. This might include those personnel who have had an extended absence from the workplace or who have missed a recurring training session.

Additional Training Insights

- Live Training
- Online Courseware

Laws

Argentina

PDPA 25.326

Terminology

Acronyms

Acronym	Definition
PDPA	Personal Data Protection Act

Overview

In 2000, Argentina passed the Data Protection Act with the explicit intent of ensuring adherence and compliance with the EU Data Directive. PDPA, consistent with EU rules, prohibits transferring personal data to countries that do not have adequate protections, such as the United States. Argentina has also enacted a number of laws to supplement the 2000 act.

Australia

Privacy Act 1988

Terminology

Acronyms

Acronym	Definition
APP	Australian Privacy Principles
NPP	National Privacy Principles

Overview

The Australian Privacy Act regulates the handling of personal information. It includes details regarding the collection, use, storage, disclosure, access to, and correction of personal information. It consists of fundamental National Privacy Principles (NPP) covering such issues as:

- Transparency in the handling of personal information
- The rules on collecting information from solicitation
- Correctness and integrity of collected data

Within the privacy principles, the following components are addressed for personal information:

- Collection
- Use
- Disclosure
- Access
- Correction
- Identification

Since 2014, the revised Privacy Amendment Act has introduced a new set of principles focusing on the handling of personal information, now called the Australian Privacy Principles (APP). The Privacy Amendment Act requires organizations to put in place SLAs, with an emphasis on security. These SLAs must list the right to audit, reporting requirements, data locations permitted and not permitted, who can access the information, and cross-border disclosure of personal information.

Components

APP8 (cross-border disclosure of personal information)

Focuses on regulating the disclosure or transfer of personal information to a separate entity offshore or overseas.

APP11.1 (security of personal information)

Requires that an organization take reasonable steps to protect the personal information it holds from misuse, interference, and loss from unauthorized access, modification, or disclosure.

Canada

PIPEDA

Terminology

Acronyms

Acronym	Definition
PIPEDA	Personal Information Protection and Electronics Document Act

Overview

PIPEDA is a **Canadian** law relating to data privacy. It governs how **private sector organizations** collect, use and disclose personal information in the course of commercial business. In addition, the Act contains various provisions to facilitate the use of electronic documents. The act was also intended to reassure the European Union that the Canadian privacy law was adequate to protect the personal information of European citizens.

EU/EEA

Directive 95/46/EC

Terminology

Acronyms

Acronym	Definition
EU	European Union
OECD	Organization for Economic Cooperation and Development

Overview

The EU Data Protection Directive of 1995 was the first major EU data private law. This overarching regulation describes the appropriate handling of personal and private information of all EU citizens. Any entity gathering the PII of any citizen of the EU is subject to the Data Directive. This includes by either automated or paper means.

It does **not** apply to the processing of data in these instances:

- By a natural person in the course of purely personal or household activities
- In the course of an activity that falls outside the scope of community law, such as operations concerning public safety, defense, or state security

Companies must meet special conditions to ensure that they provide an adequate level of data protection if they plan to transfer data to a jurisdiction outside the EEA. Some countries have been approved by the EU commission by virtue of the country's domestic law or of the international commitment it has entered into. These countries include:

- Switzerland
- Australia
- New Zealand
- Argentina

- Israel

U.S. companies that have subscribed to the Privacy Shield principles are also approved for this purpose.

Principles

The Data Directive addresses individual personal privacy by codifying these seven principles:

- Notice
- Choice
- Purpose
- Access
- Integrity
- Security
- Enforcement

Notice

The individual must be informed that personal information about them is being gathered or created.

Choice

Every individual can choose whether to disclose their personal information. No entity can gather or create personal information about an individual without that individual's explicit agreement.

Purpose

The individual must be told the specific use the information will be put to. This includes whether the data will be shared with any other entity.

Access

The individual is allowed to get copies of any of their own information held by any entity.

Integrity

The individual must be allowed to correct any of their own information if it is inaccurate.

Security

Any entity holding an individual's personal information is responsible for protecting that information and is ultimately liable for any unauthorized disclosure of that data.

Enforcement

All entities that have any personal data of any EU citizen understand that they are subject to enforcement actions by EU authorities.

> ⓘ This list largely conforms to a set of principles created by the OECD.

Directive 2002/58/EC

Overview

The ePrivacy Directive is concerned with the processing of personal data and the protection of privacy in the electronic communications sector.

GDPR

Terminology

Acronyms

Acronym	Definition
EEA	European Economic Area
EU	European Union
GDPR	General Data Protection Regulation
SA	Supervisory Authority

Overview

The EU adopted the GDPR in 2016, which is binding on all EU member states, as well as members of the European Economic Area (EEA). GDPR is also directly binding on any corporation that processes the data of EU citizens.

Principles

- The concept of consent
- Transfers abroad
- The right to be forgotten (also known as *right of erasure*)
- Establishment of the role of the data protection officer
- Access requests
- Home state regulation
- Increased sanctions

To be able to demonstrate compliance with the GDPR, the data controller should implement measures, which meet the principles of data protection by design and data protection by default.

Privacy by design and by default require data protection measures to be designed into the development of business processes for products and services. Such measures include pseudonymizing personal data, by the controller, as soon as possible.

Each member state will establish an independent SA to hear and investigate complaints, sanction administrative offenses, etc.

Breaches of Security

The GDPR requires companies to report that they have suffered a breach of security. The reporting requirements are risk-based, and there are different requirements for reporting the breach to the Supervisory Authority and to the affected data subjects. Breaches must be reported within 72 hours of the company becoming aware of the incident.

Sanctions

Violations of the GDPR expose a company to significant sanctions. These sanctions may reach up to the greater of 4% of their global turnover or gross income, or up to EUR 20 million.

Compliance

Nation	GDPR Compliance
EU	Yes
United States	**No**
Australia and New Zealand	Yes
EFTA	Yes

Israel	Yes
Japan	Yes
Canada	Yes

ENISA NIS

Terminology

Acronyms

Acronym	Definition
EEA	European Economic Area
ENISA	European Union Agency for Cybersecurity
EU	European Union
NIS	Network and Information Security

Overview

From a security standpoint, the NIS Directive is paving the way to more stringent security requirements. The NIS Directive requires EU/EEA member states to implement new information security laws for the protection of critical infrastructure and essential services by May 2018.

International

APEC

Terminology

Acronyms

Acronym	Definition
APEC	Asia-Pacific Economic Cooperation

Overview

The APEC is a regional organization meant to work toward economic growth and cooperation of its member nations. APEC agreements are not legally binding and are followed only with voluntary compliance by those entities (usually private companies) that choose to participate. It aims to address privacy as it relates to the following:

- Privacy as an international issue
- Electronic trading environments
- Effects of cross-border data flows

The APEC intent is to enhance the function of free markets through common adherence to PII protection principles. APEC members understand that consumers will not trust markets if their PII is not protected during participation in those markets. Therefore, APEC principles offer reassurance to consumers as an effort to increase faith in trading practices and thereby ensure mutual benefit for all involved.

- Individuals know when their data is used, transmitted, or stored.
- Limitations on usage are based on what is known to the individuals.
- The entity collecting or creating PII has responsibilities toward maintaining data accuracy and integrity.

ⓘ APEC promotes a consistent approach to information privacy to ensure the free flow of information.

Components

Framework

The APEC privacy framework is a principles-based privacy framework that is made up of four parts, as noted here:

- Part 1: Preamble
- Part 2: Scope
- Part 3: Information Privacy Principles
- Part 4: Implementation

Principles

The nine principles that make up the framework are as follows:

- Preventing Harm
- Notice
- Collection Limitations
- Use of Personal Information
- Choice
- Integrity of Personal Information
- Security Safeguards
- Access and Correction
- Accountability

EFTA

Terminology

Acronyms

Acronym	Definition
EFTA	European Free Trade Association
EU	European Union

Overview

Switzerland is not technically a member of the EU. Instead, it is a member of a four-nation smaller body known as the EFTA along with Iceland, Lichtenstein, and Norway.

OECD

Terminology

Acronyms

Acronym	Definition
GDPR	General Data Protection Regulation
OECD	Organization for Economic Cooperation and Development

Overview

The OECD is a standards organization made up of representatives from many countries, and it publishes policy suggestions. Its standards are *not* legally binding and do not have the effect of a treaty or other law (such as GDPR).

The OECD published the **first set of internationally accepted privacy principles** and recently published a set of revised guidelines governing the protection of privacy and trans-border flows of personal data**.** These revised guidelines focus on the need to globally enhance privacy protection through improved interoperability and the need to protect privacy using a practical, risk-management-based approach.

According to the OECD, several new concepts have been introduced in the revised guidelines, including the following:

- *National privacy strategies.* While effective laws are essential, the strategic importance of privacy today also requires a multifaceted national strategy coordinated at the highest levels of government.
- *Privacy management programs.* These serve as the core operational mechanism through which organizations implement privacy protection.
- *Data security breach notification.* This provision covers both notice to an authority and notice to an individual affected by a security breach affecting personal data.

Principles

The OECD identified the following **Fair Information Practices:**

- *Collection Limitation Principle.* There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.
- *Data Quality Principle.* Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.
- *Purpose Specification Principle.* The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.
- *Use Limitation Principle.* Personal data should not be disclosed, made available or otherwise used for purposes other than with the consent of the data subject or by the authority of the law.
- *Security Safeguard Principle.* Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorized access, destruction, use, modification or disclosure of data.
- *Openness Principle.* There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.
- *Individual Participation Principle.* An individual should have several rights surrounding their data, such as obtaining the data without excessive financial charges and challenging data relating to the subject, and if the challenge is successful to have the data erased, rectified, completed or amended.
- *Accountability Principle.* A data controller should be accountable for complying with measures which give effect to the principles stated above.

> ⓘ The OECD principles do not include "the right to be forgotten" as the EU does.

Russia

Data Localization Law

Overview

Under the Data Localization Law, businesses collecting data of Russian citizens, including on the Internet, are obliged to record, systematize, accumulate, store, update, change, and retrieve the personal data of Russian citizens in databases located within the territory of the Russian Federation.

Switzerland

DPA

Terminology

Acronyms

Acronym	Definition
DPA	Data Protection Act

Overview

In accordance with Swiss data protection law, the basic principles of which are in line with EU law, three issues are important:

- The conditions under which the transfer of personal data processing to third parties is permissible
- The conditions under which personal data may be sent abroad
- Data security

Principles

Data Processing by Third Parties

In systems of law with extended data protection, as is the case for EU and Switzerland, it is permissible to enlist the support of third parties for data processing.

Transferring Personal Data Abroad

Exporting data abroad is permissible if legislation that ensures adequate data protection in accordance with Swiss standards exists in the country in which the recipient of the data is located.

Data Security

CIA of data must be ensured by means of appropriate organizational and technical measures.

United States

COPPA

Terminology

Acronyms

Acronym	Definition
COPPA	Children's Online Privacy Protection Act

Overview

COPPA imposes certain requirements on operators of websites or online services directed to children under 13 years of age, and on operators of other websites or online services that have actual knowledge that they are collecting personal information online from a child under 13 years of age.

> ⓘ COPPA forbids the collection of personal information or cookies.

DMCA

Terminology

Acronyms

Acronym	Definition
DMCA	Digital Millennium Copyright Act

EAR

Terminology

Acronyms

Acronym	Definition
CCL	Communications Control List
EAR	Export Administration Regulations
USML	United States Munitions List

Overview

In general, the EAR govern whether a person may export a thing from the U.S., reexport the thing from a foreign country, or transfer a thing from one person to another in a foreign country. The EAR apply to physical things (sometimes referred to as "commodities") as well as technology and software.

> ⓘ ITAR covers the *control* of all defense articles and services, while EAR covers the *restriction* of commercial and dual-use items and technologies. You can find ITAR-covered items on the USML, while EAR items are listed on CCL.

ECPA

Terminology

Acronyms

Acronym	Definition
ECPA	Electronic Communication Privacy Act

FERPA

Terminology

Acronyms

Acronym	Definition
FERPA	Family Educational Rights and Privacy Act

Overview

FERPA is a federal law that protects the privacy of **student education records**. FERPA gives parents certain rights with respect to their children’s education records. These rights transfer to the student when he or she reaches the age of 18 or attends a school beyond the high school level. Students to whom the rights have transferred are “eligible students.”

FISMA

Terminology

Acronyms

Acronym	Definition
FISMA	Federal Information Security Management Act
NIST	National Institute of Standards and Technology
OMG	Office of Management and Budget

Overview

FISMA is a United States federal law that made it a requirement for **federal agencies** to develop, document, and implement an information security and protection program. FISMA defines a comprehensive framework designed to protect U.S. government information, operations, and assets against natural or fabricated threats. FISMA requires agencies to comply with NIST guidance.

GLBA

Terminology

Acronyms

Acronym	Definition
GLBA	Gramm-Leach-Bliley Act
ISO	Information Security Officer
ISP	Information Security Plan

Overview

GLBA was created to allow **banks** and **financial institutions** (such as **insurance companies**) to merge. GLBA includes several provisions specifying the kinds of protections and controls that financial institutions are required to use for security customers' account information. The act also requires financial institutions to give customers written privacy notices that explain their information-sharing practices.

Individuals must *opt out* for data sharing.

> ⓘ GLBA is also known as the "Financial Services Modernization Act of 1999."

Components

Financial Privacy Rule

Regulates the collection and disclosure of private financial information.

Safeguards Rule

Stipulates that financial institutions must implement security programs to protect such information.

Pretexting Provisions

Prohibits the practice of pretexting (accessing private information using false pretenses).

> Some of the provisions include requiring all financial institutions to have a written ISP, and later revisions of FDIC guidance require that an ISO be named and given adequate resources in order to implement the ISP.

HIPAA

Terminology

Acronyms

Acronym	Definition
DHHS	Department of Health and Human Services
HIPAA	Health Insurance Portability and Accountability Act
PHI	Personal Health Information
OCR	Office for Civil Rights

Overview

HIPAA is a set of federal laws governing the handling of PHI.

Individuals must *opt in* for data sharing.

Enforcement

- The DHHS is in charge of **managing** HIPAA.
- The OCR is the federal **enforcement** arm of the DHHS.

Components

Privacy Rule

The first primary regulation promulgated by the DHHS was the Privacy Rule. It contained language specific to maintaining the privacy of patient information as it was traditionally stored and used, on **paper**.

The patient controls the distribution of their information.

Security Rule

With the explosion of networking, digital storage, and the Internet came the Security Rule, followed by the Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009, which provided financial incentives for medical practices and hospitals to convert paper record-keeping systems to **digital**.

The healthcare provider has to protect information based on the privacy the patient has specified.

HITECH

Terminology

Acronyms

Acronym	Definition
EHR	Electronic Health Records
HITECH	Health Information Technology for Economic and Clinical Health

Overview

The HITECH Act is a legislation that was created to stimulate the adoption of EHR and the supporting technology in the United States.

ITAR

Terminology

Acronyms

Acronym	Definition
CCL	Commercial Control List
ITAR	International Traffic in Arms Regulations
USML	United States Munitions List

Overview

International Traffic in Arms Regulations (ITAR) is a United States regulatory regime to restrict and **control** the **export** of defense and military related technologies to safeguard U.S. national security and further U.S. foreign policy objectives.

> ⓘ ITAR covers the *control* of all defense articles and services, while EAR covers the *restriction* of commercial and dual-use items and technologies. You can find ITAR-covered items on the USML, while EAR items are listed on CCL.

> ✓ **Fact.** The State Department is involved with controlling exports.

Privacy Shield

Terminology

Acronyms

Acronym	Definition
DoC	Department of Commerce
DoT	Department of Transportation
FTC	Federal Trade Commission

Overview

The EU-U.S. Privacy Shield is a framework for regulating transatlantic exchanges of personal data for commercial purposes between the EU and the U.S., with the goal of protecting EU citizens' personal information.

Enforcement

- The DoC **manages** the Privacy Shield program in the United States.
- The program is **administered** by the DoT and the FTC.
- The FTC is the U.S. **enforcement** arm for most Privacy Shield activity.

Principles

The following principles are *requirements* for participation:

- Notice

- Choice
- Accountability for Onward Transfer
- Security
- Data Integrity and Purpose Limitation
- Access
- Recourse, Enforcement and Liability

If American companies don't want to subscribe to Privacy Shield, they can create internal policies called "binding corporate rules" and "standard contractual clauses" that explicitly state full compliance with the Data Directive and Privacy Regulation.

Safe Harbor

Terminology

Acronyms

Acronym	Definition
DoC	Department of Commerce
DoT	Department of Transportation
FTC	Federal Trade Commission

Overview

To allow some U.S.-based companies to operate legitimately inside the EU, the EU Data Directive included a set of *safe harbor* privacy rules designed to outline what American companies must do in order to comply with EU laws. These rules outlined the proper handling of storage and transmission of private information belonging to EU citizens.

- Must voluntarily agree to comply with the Data Directive.
- Must sign up with a federal enforcement entity in the United States that would administer the program.
- Must agree to allow auditing and enforcement by the program administrators.

Any U.S. organization subject to the FTC's jurisdiction and some transportation organizations subject to the jurisdiction of the DoT can participate in the Safe Harbor program. Certain industries, such as telecommunication carriers, banks, and insurance companies, may not be eligible for this program.

Enforcement

- For most companies, the Safe Harbor program was **administered** by the DoC.

- For the specific industries of airlines and shipping companies, the program was **administered** by the DoT.
- The **enforcement** arm of the DoC is the FTC.

Principles

Under the Safe Harbor program, U.S. companies have been able to voluntarily adhere to a set of seven principles:

- Notice
- Choice
- Transfers to third parties
- Access
- Security
- Date integrity
- Enforcement

Organizations must also be subject to enforcement and dispute resolution proceedings.

> ⓘ In May 2018, an update to the Data Directive took effect known as the "Privacy Regulation". The Privacy Regulation supersedes the Data Directive and ends the Safe Harbor program, replacing it with a new program known as Privacy Shield.

SCA

Terminology

Acronyms

Acronym	Definition
SCA	Stored Communication Act

Overview

Enacted as part of Title II of the ECPA, the SCA addresses both voluntary and compelled disclosure of *stored wire and electronic communications and transactional records* held by third parties. It further provides for privacy protection regarding certain electronic communications and computing services from unauthorized access or interception by government entities.

SOX

Terminology

Acronyms

Acronym	Definition
SEC	Securities and Exchange Commission
SOX	Sarbanes-Oxley Act

Overview

In 2002 the Sarbanes-Oxley Act (SOX) was enacted as an attempt to prevent fraudulent accounting practices, poor audit practices, inadequate financial controls, and poor oversight by governing boards of directors. It applies to all publicly traded corporations. SOX protects individuals from **accounting errors** and **fraudulent practices** in **publicly traded companies**. It provides for corporate accountability.

SOX is not a set of business practices and does not specify how a business should store records; rather, it defines which records (such as financial records) are to be stored and for how long.

> ⓘ SOX is also known as the "**public company** accounting reform and investor protection act".

Enforcement

- The SEC is the organization responsible for establishing standards and guidelines and conducting audits and imposing subsequent fines should any be required.

Standards

Auditing and Assurance

AICPA SOC

Terminology

Acronyms

Acronym	Definition
AICPA	American Institute of Certified Public Accountants
IAASB	International Auditing and Assurance Standards Board
SAS	Statement on Auditing Standards
SOC	System and Organization Controls
SSAE	Statement on Standards for Attestation Engagements

Overview

SOC is part of the SSAE reporting format created by the AICPA. SSAE 18 is the current audit standard. These are uniformly recognized as being acceptable for regulatory purposes in many industries, although they were specifically designed as mechanisms for ensuring compliances with the Sarbanes-Oxley Act, which governs publicly traded corporations.

> ⓘ The SAS 70 was replaced by SOC Type 1 and Type 2 reports in 2011 following changes and a more comprehensive approach to auditing being demanded by customers and clients. For years, SAS 70 was seen as the de facto standard for datacenter customers to obtain independent assurance that their datacenter service provider had effective internal controls in place for managing the design, implementation, and execution of customer information.

Components

SOC 1

SOC 1 reports are strictly for auditing the financial reporting instruments of a corporation. There are two subclasses of SOC 1 reports: Type 1 and Type 2.

> ⓘ The international equivalent to the AICPA SOC 1 is the IAASB issued and approved ISAE 3402.

SOC 2

SOC 2 reporting was specifically designed for IT-managed service providers and cloud computing. The report specifically addresses any number of the so-called **Trust Services** principles, which follow:

- Security
- Availability
- Processing Integrity
- Confidentiality
- Privacy

The SOC 2 is an examination of the design and operating effectiveness of controls that meet the criteria for principles set forth in the AICPA's Trust Services principles.

Typically, SOC 2 reports are only provided to *customers* and require an NDA be signed.

> ⓘ A cloud provider intending to prove its trustworthiness would look to an SOC 2 report as the artifact that demonstrated it.

Type 1

SOC 2 Type 1 reports only reviews the *design* of controls, not how they are implemented and maintained, or their function.

A Type 1 report presents an auditor's opinion at a *specific date.*

> The SOC 2 Type 1 is not extremely useful for determining the security and trust of an organization.

Type 2

SOC 2 Type 2 reports are a thorough review of the target's controls, including how they have been implemented and their efficacy. It gives the customer a realistic view of the provider's security posture and overall program.

A type 2 report presents an auditor's opinion over a declared period, generally between 6 months and 1 year.

> The SOC 2 Type 2 *is* extremely useful for determining the security and trust of an organization. It provides a true assessment of an organization's security posture.

> Cloud vendors will probably never share an SOC 2 Type 2 report with any customer or even release it outside the provider's organization. The SOC 2 Type 2 report is extremely detailed and provides exactly the kind of description and configuration that the cloud provider is trying to restrict from wide dissemination.

SOC 3

The SOC 3 is the "seal of approval". It contains no actual data about the security controls of the audit target and is instead just an assertion that the audit was conducted and that the target organization passed.

SOC 3 reports are usually publicly available. Thus, if you are not a customer, you would likely receive a SOC 3 report.

> ⓘ This is currently the practice accepted in the industry.

BITS Shared Assessments

Overview

An organization that provides firms with a way to obtain a detailed report about a service provider's controls (people, processes, and procedures) and a procedure for verifying that the information in the report is accurate. They offer the tools to assess third-party risk, including cloud-based risks.

Shared Assessments is a third party risk membership program that provides organizations with a way to obtain a detailed report about a service provider's controls (people, process and procedures) and a procedure for verifying that the information in the report is accurate.

CSA STAR

Terminology

Acronyms

Acronym	Definition
CAI	Consensus Assessments Initiative
CAIQ	Consensus Assessments Initiative Questionnaire
CCM	Cloud Controls Matrix
CSA	Cloud Security Alliance
CTP	CloudTrust Protocol
STAR	Security, Trust, and Assurance Registry

Definitions

CAIQ

A self-assessment performed by cloud providers, detailing their evaluation of the practice areas and control groups they use in providing their services.

→ CSA CCM /standards/security-management-and-controls/csa-ccm

Overview

The CSA STAR program appeared as demand for a single consistent framework for evaluating cloud providers developed.

The CSA STAR program was designed to provide an independent level of program assurance for **cloud** consumers. Provides a mechanism for users to assess the security of **cloud security providers.** It allows customers to perform a large component of due diligence and allow a single framework of controls and requirements to be utilized in assessing CSP suitability and the ability to fulfill CSP requirements.

Components

The CSA STAR program consists of three levels based on the Open Certification Framework.

STAR Level 1

> **Cloud Security Alliance**
>
> At level one organizations can submit one or both of the security and privacy self-assessments. For the security assessment, organizations use the Cloud Controls Matrix to evaluate and document their security controls. The privacy assessment submissions are based on the GDPR Code of Conduct.

Requires the release and publication of due diligence assessments against the CSA's *Consensus Assessment Initiative Questionnaire and/or Cloud Controls Matrix (CCM).*

A free designation that allows cloud computing providers to document their security controls by performing a self-assessment against CSA best practices. The results are made publicly available to customers.

Self-Assessment

GDPR Code of Conduct (CoC) Self-Assessment

STAR Level 1 Continuous

Continuous Self-Assessment

> **Cloud Security Alliance**
>
> A CSP that uses a CAIQ to achieve Self-Assessment, a point-in-time assessment, can use a Continuous Self-Assessment to demonstrate effectiveness of controls over a period of time, to achieve STAR Continuous Level 1.

STAR Level 2

Requires the release and publication of available results of an assessment carried out by an independent third party based on CSA CCM and ISO 27001:2013 or an AICPA SOC 2

> **Cloud Security Alliance**
>
> Organizations looking for a third-party audit can choose from one or more of the security and privacy audits and certifications. An organization’s location, along with the regulations and standards it is subject to will have the greatest factor in determining which ones are appropriate to pursue.

Attestation, Certification, C-Star

- Attestation includes a third-party individual assessment against SOC 2 standards.
    - This level is a collaboration between CSA and the AICPA to provide guidelines for CPAs to conduct SOC 2 engagements using criteria from the AICPA (Trust Service Principles, AT 101) and the CSA CCM.
- Certification includes a third-party individual assessment against ISO/IEC 27001
    - This level is a technology-neutral certification that is based on a rigorous, independent third-party assessment of a cloud service provider's security. The certification leverages the requirements of the ISO/IEC 27001 management system standard together with the CSA CCM.
- C-Star Assessment complies with Chinese national standards.
    - This level is an assessment that is based on an independent third-party assessment of a cloud service provider's security for the Greater China market. The assessment harmonizes CSA best practices with Chinese national standards.

GDPR Code of Conduct (CoC) Self-Assessment

STAR Level 2 Continuous

Level 2 + Continuous Self-Assessment

> **Cloud Security Alliance**
>
> A CSP, who holds a third-party audit, can achieve STAR Level 2 Continuous by adding a Continuous Self-Assessment, which allows them to quickly inform customers of changes to their security programs, instead of communicating those until the next audit period in normal STAR Level 2.

STAR Level 3 Continuous Auditing

Continuous Certification

Requires the release and publication of results related to the security properties of monitoring based on the CloudTrust Protocol.

The CloudTrust protocol is intended to establish a digital trust between a cloud computing customer and provider and create transparency about the providers configurations, vulnerabilities, access, authorization, policy, accountability, anchoring and operating status conditions.

Providers will publish their security practices according to CSA formatting and specifications, including validation of CCM, CTP, and CloudAudit (A6) standards. Customers and tool vendors will retrieve the information in a variety of contexts.

> **Cloud Security Alliance**
>
> Automate the current security practices of cloud providers. Providers publish their security practices according and customers and tool vendors can retrieve and present this information in avariety of contexts.
>
> A CSP is the most transparent through a continuous, automated process that ensures that security controls are monitored and validated at all times.

EuroCloud StarAudit

Overview

The StarAudit scheme evaluates cloud services according to a well-defined and transparent catalogue of criteria. The result of this audit process shows the respective maturity and compliance levels of a service. The certification procedure is based on best practices and provides answers to the fundamental questions managers are likely to ask when looking for a suitable cloud service provider. Unlike pure security or data protection audits, it covers the entire range of cloud service functions and validates compliance against the requirements in clearly understandable terms.

ISO/IEC 15408:2009

Terminology

Acronyms

Acronym	Definition
CC	Common Criteria
EAL	Evaluation Assurance Level
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
ST	Security Target
TCSEC	Trusted Computer System Evaluation Criteria
UTM	Unified Threat Management

Overview

ISO/IEC 15408, better known as the **Common Criteria (CC)**, is an international set of guidelines and specifications developed for evaluating information security products, with the view to ensuring they meet an agreed-upon security standard for government entities and agencies.

CC looks at certifying a product only and does **not include administrative or business processes**.

> ⓘ The certification of the product only certifies its capabilities. If misconfigured or mismanaged, it is no more secure than anything else the customer might use.

Components

Protection Profiles

Protection profiles define a standard set of security requirements for a specific type of product, such as a firewall, IDS, or UTM.

Evaluation Assurance Levels

EALs define how thoroughly the product is **tested**. EALs are rated using a sliding scale from 1-7, with 1 being the lowest-level evaluation and 7 being the highest.

> The higher the level of evaluation, the more QA tests the product would have undergone; however, undergoing more tests does not necessarily mean the product is more secure.

Level	Description
EAL1	Functionally tested
EAL2	Structurally tested
EAL3	Methodically tested and checked
EAL4	Methodically designed, tested, and reviewed
EAL5	Semiformally designed and tested
EAL6	Semiformally verified design and tested
EAL7	Formally verified design and tested

Process

There are three steps to successfully submit a product for evaluation according to the Common Criteria:

1. The vendor must detail the security features of a product using what is called a security target.
2. The product, along with the Security Target, goes to a certified laboratory for testing according to evaluate how well it meets the specifications defined in the protection profile.
3. A successful evaluation leads to an official certification of the product.

NIST FIPS 140-3

Terminology

Acronyms

Acronym	Definition
CST	Cryptographic and Security Testing
CMVP	Cryptographic Module Validation Program
FIPS	Federal Information Processing Standard
NIST	National Institute of Standards and Technology
NVLAP	National Voluntary Laboratory Accreditation Program
SBU	Sensitive But Unclassified

Definitions

SBU

FIPS is only for SBU data.

Overview

A federal standard for accrediting and distinguishing secure and well-architected cryptographic modules produced by private sector vendors who seek to or are in the process of having their solutions and services certified for use in U.S. government departments and regulated industries (this includes financial services and healthcare) that collect, store, transfer, or share data that is deemed to be "sensitive" but not classified (i.e., Secret/Top Secret).

FIPS 140 Publication Series was issued by NIST to coordinate the requirements and standards for cryptography modules covering both hardware and software components for cloud and traditional computing environments.

The primary goal for the FIPS 140-3 standard is to accredit and distinguish secure and well-architected cryptographic modules produced by private sector vendors who seek to or are in the process of having their solutions and services certified for us in U.S. government departments and regulated industries that collect, store, transfer, or share data that is deemed to be sensitive but not classified.

> ⓘ FIPS 140-3 is essentially a list of approved cryptosystems.

Components

The FIPS 140-3 standard provides **four distinct levels of qualitative security** intended to cover a range of potential applications and environments with emphasis on secure design and implementation of a cryptographic module.

Includes **11 sections:**

- Cryptographic module specification
- Cryptographic module interfaces
- Roles, services, and authentication
- Software/firmware security
- Operating environment
- Physical security
- Non-invasive security
- Sensitive security parameter management
- Self-tests
- Life-cycle assurance
- Mitigation of other attacks

The CMVP validates cryptographic modules to FIPS 140-3 and other cryptography-based standards. The goal of the CMVP is to promote the use of validated cryptographic modules and provide federal agencies with a security metric to use in procuring equipment containing validated cryptographic modules. In the CMVP, vendors of cryptographic modules use independent, accredited CST laboratories to have their modules tested. NVLAP accredited laboratories perform cryptographic module compliance/conformance testing.

The amount of physical protection provided by the product, in terms of tamper resistance, is what distinguishes the security levels for cryptographic modules.

> ⓘ Some important sections that were included in FIPS 140-2 but are no longer explicitly referenced in FIPS 140-3 include:
>
> - Finite state model
> - Cryptographic key management
> - Electromagnetic interference/electromagnetic compatibility (EMI/EMC)

Security Level 1

Requires production-grade equipment and externally tested algorithms.

> ⓘ This is the lowest level of security.

Security Level 2

Adds requirements for physical tamper-evidence and role-based authentication. Software implementations must run on an Operating System approved to Common Criteria at EAL2.

- Detection and tampering.

Security Level 3

Adds requirements for physical tamper-resistance and identity-based authentication. There must also be physical or logical separation between the interfaces by which “critical security parameters” enter and leave the module. Private keys can only enter or leave in encrypted form.

- Prevention and Detection

Security Level 4

This level makes the physical security requirements more stringent, requiring the ability to be tamper-active, erasing the contents of the device if it detects various forms of environmental attack.

> This is the highest level of security.

Cloud Computing

NIST SP 800-145

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology
SP	Special Publication

Overview

Outlines both the cloud computing deployment and service models and their definitions.

NIST SP 800-146

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology
SP	Special Publication

Overview

Reprises the NIST-established definition of cloud computing, describes cloud computing benefits and open issues, presents an overview of major classes of cloud technology, and provides guidelines and recommendations on how organizations should consider the relative opportunities and risks of cloud computing.

Focused on risk components and the appropriate analysis of such risks.

Cloud Computing Reference Architecture

CSA TCI

Terminology

Acronyms

Acronym	Definition
CSA	Cloud Security Alliance
TCI	Trusted Cloud Initiative

Overview

The TCI Reference Architecture is both a methodology and a set of tools that enable security architects, enterprise architects and risk management professionals to leverage a common set of solutions that fulfill their common needs to be able to assess where their internal IT and their cloud providers are in terms of security capabilities and to plan a roadmap to meet the security needs of their business.

TCI helps cloud providers develop industry-recommended, secure and interoperable IAM configurations, and practices. The CSA TCI helps CSPs develop identity, access, and compliance management guidelines.

Mission Statement

To promote research, development, and education of best practices and methodologies around a reference architecture for a secure and trusted cloud.

Principles

- Define protections that enable trust in the cloud.
- Develop cross-platform capabilities and patterns for proprietary and open-source providers.
- Facilitate trusted and efficient access, administration and resiliency to the customer/consumer.
- Provide direction to secure information that is protected by regulations.
- Facilitate proper and efficient governance, identification, authentication, authorization, administration and auditability.
- Centralize security policy, maintenance operation and oversight functions.
- Access to information must be secure yet still easy to obtain.
- Delegate or federate access control where appropriate.
- Must be easy to adopt and consume, supporting the design of security patterns.
- Must be elastic, flexible and resilient supporting multitenant, multi-landlord platforms.
- Must address and support multiple levels of protection, including network, operating system, and application security needs.

ISO/IEC 17789:2014

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

ISO/IEC 17789:2014 defines the Cloud Computing Reference Architecture (CCRA) which includes the cloud computing roles, cloud computing activities, and the cloud computing functional components and their relationships.

An activity is a set of tasks that accomplish a goal.

According to ISO/IEC 17789, the three main groups of cloud service activities are:

- Activities that use services
- Activities that provide services
- Activities that support services

ISO/IEC 17789 defines three major roles that perform groups of activities in cloud computing. The activity roles are:

- Cloud service customer (CSC)
- Cloud service provider (CSP)
- Cloud service partner (CSN)

The CSC role includes the cloud service user, cloud service administrator, cloud service business manager, and cloud service integrator sub-roles. These sub-roles govern several activities.

The CSP role includes the cloud service manager, cloud service operations manager, cloud service deployment manager, cloud service business manager, cloud service security and risk manager, customer support and care representative, inter-cloud provider, and the network provider sub-roles. These sub-roles govern several activities.

The CSN role includes cloud service developer, cloud auditor, and cloud service broker sub-roles. These sub-roles govern several activities.

Data Center Design

ANSI/BICSI 002-2014

Terminology

Acronyms

Acronym	Definition
ANSI	American National Standards Institute
BICSI	Building Industry Consulting Services International

Overview

The ANSI/BICSI 002-2014 standard covers cabling design and installation.

Created the Data Center Design and Implementation Best Practices standard, which includes specifications for items such as hot/cold aisle setups, power specifications, and energy efficiency.

IDCA Infinity Paradigm

Terminology

Acronyms

Acronym	Definition
IDCA	International Data Center Authority

Overview

A framework intended to be used for operations and datacenter design. The Infinity Paradigm covers data center location, facility, structure, and infrastructure and applications.

NFPA 70

Terminology

Acronyms

Acronym	Definition
NFPA	National Fire Protection Association

Overview

The NFPA publishes standards regarding fire protection.

NFPA standard 70 requires the implementation of an emergency power-off button to protect first responders in the data center in case of emergency.

NFPA 75 and 76

Terminology

Acronyms

Acronym	Definition
NFPA	National Fire Protection Association

Overview

The NFPA publishes standards regarding fire protection.

NFPA 75 and 76 standards specify how hot or cold aisle containment is to be carried out.

Components

Hot Aisle Containment

The process of isolating exhaust heat into a single aisle for two separate rows of server cages.

In this configuration, racks are configured such that the backs of the devices face each other.

Cold Aisle Containment

The process of isolating intake air into a single aisle for two separate rows of server cages.

In this configuration, racks are configured such that the fronts of the devices face each other.

Uptime Institute

Overview

The Uptime Institute is an advisory organization for matters related to IT service. UI publishes a standard for datacenter design, and it also certifies datacenters for compliance with this standard.

Components

The UI standard is split into four tiers, in ascending durability of the datacenter.

Tier 1

Tier 1 is a simplistic datacenter, with little or no redundancy and is labeled `Basic Site Infrastructure`.

Minimum Requirements

- Dedicated space for IT systems
- An uninterruptable power supply (UPS) system for line conditioning and backup purposes
- Sufficient cooling systems to serve all critical equipment
- A power generator for extended electrical outages, with at least 12 hours of fuel to run the generator at sufficient load to power the IT systems

Features

- Scheduled maintenance will require systems (including critical systems) to be taken offline.
- Both planned and unplanned maintenance and response activity may take systems (including critical systems) offline.
- Untoward personnel activity (both inadvertent and malicious) will result in downtime.

- Annual maintenance is necessary to safely operate the datacenter and requires full shutdown (including critical systems). Without this maintenance, the datacenter is likely to suffer increased outages and disruptions.

Tier 2

Tier 2 is slightly more robust than Tier 1 and is named `Redundant Site Infrastructure Capacity Components`. It features all the attributes of the Tier 1 design, with additional elements.

Additional Features

- Critical operations do not have to be interrupted for scheduled replacement and maintenance of any of the redundant components; however, there may be downtime for any disconnection of power distribution systems and lines.
- Contrary to Tier 1, where untoward personnel activity *will* cause downtime, in Tier 2 it *may* cause downtime.
- Unplanned failures of components or systems might result in downtime.

Tier 3

The Tier 3 design is known as a `Concurrently Maintainable Site Infrastructure`. The facility features both the redundant capacity components of a Tier 2 build and the added benefit of multiple distribution paths.

Additional Features

- There are dual power supplies for all IT systems.
- Critical operations can continue even if any single component or power element is out of service for scheduled maintenance ore replacement.
- Unplanned loss of a *component may* cause downtime; the loss of a single *system,* on the other hand, *will* cause downtime.
- Planned maintenance will not necessarily result in downtime; however, the risk of downtime may be increased during this activity. This temporary elevated risk does not make the datacenter lose its Tier 3 rating for the duration.

Tier 4

The Tier 4 design is known as a `Fault-Tolerant Site Infrastructure`. Each and every element and system of the facility has integral redundancy such that critical operations can survive both planned and unplanned downtime at the loss of any component or system.

Additional Features

- There is redundancy of both IT and electrical components, where the various multiple components are independent and physically separate from each other.
- Even after the loss of any facility infrastructure element, there will be sufficient power and cooling for critical operations.
- The loss of any single system, component, or distribution element *will not* affect critical operations.
- The facility will feature automatic response capabilities for infrastructure control systems such that the the critical operations will not be affected by infrastructure failures.
- Any single loss, even, or personnel activity will not cause downtime of critical operations.
- Scheduled maintenance can be performed without affecting critical operations. However, while one set of assets is in the maintenance state, the datacenter may be at increased risk of failure due to an event affecting the alternate assets. During this temporary maintenance state, the facility does not lose its Tier 4 rating.

Forensics

ISO/IEC 27050:2019

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

The ISO/IEC 27050 standard provides guidelines for eDiscovery processes and best practices. ISO/IEC 27050 covers all steps of **eDiscovery** processes including:

- Identification
- Preservation
- Collection
- Processing
- Review
- Analysis
- Final Production of the Requested Data Archive

Privacy

AICPA/CICA GAPP

Terminology

Acronyms

Acronym	Definition
AICPA	American Institute of Certified Public Accountants
CICA	Canadian Institute of Chartered Accountants
GAPP	Generally Accepted Privacy Principles

Overview

The GAPP is an AICPA standard that consists of **10 key privacy principles** and 74 privacy objectives and associated methods for measuring and evaluating criteria. It is focused on managing and preventing threats to privacy.

> ⓘ GAPP is a component of SOC 2. This is the standard in the accounting industry.

Components

According to GAPP, following are the 10 main privacy principles:

- Management
- Notice
- Choice and consent
- Collection

- Use, retention, and disposal
- Access
- Disclosure to third parties
- Security for privacy
- Quality
- Monitoring and enforcement

ISO/IEC 27018:2019

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

ISO/IEC 27018 addresses the privacy aspects of **cloud computing** for consumers. It is the **first international set of privacy controls** *in the cloud.*

Components

ISO/IEC 27018 focuses on five key principles:

- *Consent.* CSPs must not use the personal data they receive for advertising and marketing unless expressly instructed to do so by the customers. In addition, a customer should be able to employ the service without having to consent to the use of their personal data for advertising or marketing.
- *Control.* Customers have explicit control over how CSPs are to use their information.
- *Transparency.* CSPs must inform customers about items such as where their data resides. CSPs also need to disclose to customers the use of any subcontractors who will be used to process PII.
- *Communication.* CSPs should keep clear records about any incident and their response to it, and they should notify customers.

- *Independent and yearly audit.* To remain compliant, the CSP must subject itself to yearly third-party reviews. This allows the customer to rely upon the findings to support their own regulatory obligations.

> ⓘ Trust is key for consumers leveraging the cloud; therefore, vendors of cloud services are working toward adopting the stringent privacy principles outlined in ISO 27018.

Risk Management

ENISA Cloud Computing: Benefits, Risks, and Recommendations for Information Security

Terminology

Acronyms

Acronym	Definition
ENISA	European Union Agency for Network and Information Security

Overview

ENISA is a standard and model developed in Europe and is responsible for producing *Cloud Computing: Benefits, Risks, and Recommendations for Information Security.* It identifies **35 types of risks** organizations should consider but goes further by **identifying the top eight security risks** based on likelihood and impact:

- Loss of governance
- Lock-in
- Isolation failure
- Compliance risk
- Management interface failure
- Data protection
- Malicious insider
- Insecure or incomplete data deletion

ISACA COBIT

Terminology

Acronyms

Acronym	Definition
COBIT	Control Objectives for Information and Related Technologies
ISACA	Information Systems Audit and Control Association

Overview

Designed for all types of business, regardless of their purpose. COBIT is a framework for managing IT controls, largely from a process and governance perspective.

It is a framework created by the ISACA for IT governance and management. It was designed to be a supportive tool for managers—and allows bridging the crucial gap between technical issues, business risks, and control requirements.

The framework helps companies follow law, be more agile and earn more.

Components

- *Framework*: Organizes IT governance objectives and good practices by IT domains and processes and links them to business requirements.
- *Process descriptions*: A reference process model and common language for everyone in an organization. The processes map to responsibility areas of plan, build, run, and monitor.
- *Control objectives*: Provides a complete set of high-level requirements to be considered by management for effective control of each IT process.

- *Management guidelines*: Helps assign responsibility, agree on objectives, measure performance, and illustrate interrelationship with other processes.
- *Maturity models*: Assesses maturity and capability per process and helps to address gaps.

ISO 31000:2018

Terminology

Acronyms

Acronym	Definition
ISO	International Organization for Standardization

Overview

Key components of ISO 31000 are designing, implementing, and reviewing **risk management**. The key requirement of ISO 31000 is management endorsement, support, and commitment. A key concept in ISO 31000 involves risk management being an embedded component as opposed to a separate activity.

ISO 31000 is an international standard that focuses on designing, implementing, and reviewing risk management processes and practices. It is *not* intended for certification purposes; implementing it does not address specific or legal requirements related to risk assessments, risk reviews, and overall risk management. The standard explains that proper implementation of a risk management process can be used to:

- Create and protect value
- Integrate organizational procedures
- Be part of the decision-making process
- Explicitly address uncertainty
- Be a systematic, structured, and timely risk management program
- Ensure the risk management program is based on the best available information
- Be tailored to the organization's business requirements and actual risks
- Take human and cultural factors into account
- Ensure the risk management program is transparent and inclusive
- Create a risk management program that is dynamic, iterative, and responsive to change

- Facilitate continual improvement and enhancement of the organization

NIST's Framework for Improving Critical Infrastructure Cybersecurity

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology

Overview

The framework provides a common taxonomy and mechanism for organizations to:

- Describe their current cybersecurity posture
- Describe their target state for cybersecurity
- Identify and prioritize opportunities for improvement within the context of a continuous and repeatable process
- Assess progress toward the target state
- Communicate among internal and external stakeholders about cybersecurity risk

> This framework is ideal to use as a supplement to best coding practices and code reviews and testing, but also functions as an approach to risk mitigation.

Components

Framework Core

Cybersecurity activities and outcomes divided into five functions:

- Identify
- Protect
- Detect
- Respond
- Recover

Identify (ID)

Contains two categories:

- Asset Management (ID.AM)
  - *ID.AM-2.* Software platforms and applications within the organization are inventoried.
  - *ID.AM-3.* Organizational communication and data flows are mapped.
  - *ID.AM-5.* Resources (such as hardware, devices, data, and software) are prioritized based on their classification, criticality, and business value.
- Risk Assessment (ID.RA)
  - *ID.RA-1.* Asset vulnerabilities are identified and documented.
  - *ID.RA-5.* Threats, vulnerabilities, likelihoods, and impacts are used to determine risk.

Framework Profile

Used to assist the organization in aligning activities with business requirements, risk tolerance, and resources.

Framework Implementation Tiers

Used to identify where the organization is with regard to their particular approach.

NIST SP 800-37

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology
SP	Special Publication

Definition

NIST SP 800-37 is the Guide for Implementing the **Risk Management Framework (RMF)**. This particular risk management framework is a methodology for handling all organizational risk in a holistic, comprehensive, and continual manner. This RMF supersedes the old "Certification and Accreditation" model of cyclical inspections that have a specific duration.

This RMF relies heavily on the use of automated solutions, risk analysis and assessment, and implementing controls based on those assessments, with continuous monitoring and improvement.

Components

- **Categorize** information systems
- **Select** security controls
- **Implement** security controls
- **Assess** security controls
- **Authorize** information systems

- **Monitor** security controls

Secure Architecture and Design

ISACA ITIL

Terminology

Acronyms

Acronym	Definition
ISACA	Information Systems Audit and Control Association
ITIL	Information Technology Infrastructure Library (formerly)

Overview

ITIL is a group of documents that are used in implementing a framework for IT service management. ITIL forms a customizable framework that defines how service management is applied throughout an organization. It focuses on aligning IT services with the needs of business.

ITIL was specifically designed to address service delivery entities (in particular, British telecommunications providers), and how they provide service to their customers.

Components

ITIL is organized into a series of five volumes:

- Service Strategy
- Service Design
- Service Transition
- Service Operation
- Continual Service Improvement

> According to the CSA Enterprise Architecture, ITIL is part of Information Technology Operation and Support.

Jericho

> The Jericho forum is now part of The Open Group Security Forum.

SABSA

Terminology

Acronyms

Acronym	Definition
SABSA	Sherwood Applied Business Security Architecture
BOSS	Business Operation Support Services

Overview

SABSA is a means of looking at **security capabilities** from a **business perspective**.

SABSA is a proven framework and methodology used successfully around the globe to meet a wide variety of enterprise needs including risk management, information assurance, governance, and continuity management. Although copyright protected, SABSA is an open-use methodology, not a commercial product.

> **Fact.** SABSA is a means of looking at security capabilities from a business perspective.

Components

SABSA includes the following **components**, which can be used separately or together:

- Business Requirements Engineering Framework

- Risk and Opportunity Management Framework
- Policy Architecture Framework
- Security Services-Oriented Architecture Framework
- Governance Framework
- Security Domain Framework
- Through-Life Security Service Management and Performance Management Framework

> ⓘ According to the CSA Enterprise Architecture, SABSA is part of the BOSS.

TOGAF

Terminology

Acronyms

Acronym	Definition
TOGAF	The Open Group Architecture Framework

Overview

TOGAF is a means to incorporate **security architecture** with the overall **business architecture**.

TOGAF is one of the many frameworks available to the cloud security professional for designing, planning, implementing, and governing an enterprise Information Technology architecture. TOGAF provides a standardized approach that can be used to address business needs by providing a common lexicon for business communication.

TOGAF is based on **open methods** and approaches to enterprise architecture, allowing the business to avoid a lock-in scenario from the use of proprietary approaches. TOGAF also provides for the ability to quantifiably measure ROI so that the business can use resources more efficiently. Thus, TOGAF is a means to incorporate security architecture with the overall business architecture.

According to the CSA Enterprise Architecture, TOGAF provides four high-level services:

- Presentation Services (Business)
- Application Services (Application)
- Information Services (Data)
- Infrastructure Services (Technology)

> ⓘ TOGAF relies heavily on modularization, standardization, and already existing, proven technologies and products.

Secure Application Development

ISO/IEC 27034:2011

Terminology

Acronyms

Acronym	Definition
ANF	Application Normative Framework
ASMP	Application Security Management Process
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization
ONF	Organizational Normative Framework

Overview

ISO/IEC 27034 provides one of the most widely accepted set of standards and guidelines for **secure application development**. ISO/IEC 27034 is a comprehensive set of standards that cover many aspects of application development. It defines concepts, frameworks, and processes to help organizations integrate security within their software development lifecycle.

A few of the key elements include the organizational normative framework (ONF), the application normative framework (ANF), and the application security management process (ASMP).

ISO/IEC 27034 supports the concepts defined in ISO/IEC 27001 and provides a framework to implement the controls found in ISO/IEC 27002.

Components

The ONF and ANF are used to build the ASMP.

ONF

The ONF defines the organizational security best practices for *all* application development, and include several sections.

The ONF includes the application lifecycle reference model as well as roles and responsibilities (as shown below). It also contains an application specifications repository, which details the functional requirements for all applications.

Part of ISO/IEC 27034 lays out the ONF for all of the components of best practices with regard to application security. It is the container for all subcomponents of application security best practices catalogued and leveraged by an organization. The standard is composed of the following categories:

- Business Context
- Regulatory Context
- Technical Context
- Specifications
- Roles, Responsibilities, and Qualifications
- Processes
- Application Security Control (ASC) Library

ISO/IEC 27034 defines an ONF management process. This bidirectional process is meant to create a continuous improvement loop. Innovations that result from security a single application are returned to the ONF to strengthen all organization application security in the future.

Business Context

Includes all application security policies, standards, and best practices adopted by the organization.

Regulatory Context

Includes all standards, laws, and regulations that affect application security.

Technical Context

Includes required and available technologies that are applicable to application security.

Specifications

Documents the organization's IT functional requirements and the solutions that are appropriate to address these requirements.

Roles, Responsibilities, and Qualifications

Documents the actors within an organization who are related to IT applications.

Processes

Relates to application security.

Application Security Control Library

Contains the approved controls that are required to protect an application based on the identified threats, the context, and the targeted **level of trust**.

- An ASC is a control that mitigates a security weakness within an application.
  - Each ASC is paired to an application based on its contexts. These can be technical, regulatory, and business contexts.
  - Each ASC must include a verification measurement.

The Application Level of Trust is designed since not every application has the same need for security controls. Each ASC can fit within one or more levels of trust, commonly referred to as an "application level of trust."

ASC > ANF > ONF

ANF

The ANF uses the applicable portions of the ONF on a specific application to achieve the needed security requirements or the target trust level.

The ANF is used together with the ONF in that it is created for a **specific application.** The ANF shares the applicable parts of the ONF needed to achieve an application's required level of security and the level of trust desired.

> ⓘ The ANF-to-ONF relationship is a one-to-one relationship; every application has an ANF that maps back to the ONF. However, the ONF-to-ANF relationship is one-to-many. The ONF has many ANFs, but the ANF only has one ONF.

ASMP

ISO/IEC 27034 defines an ASMP to manage and maintain each ANF. The ASMP is created in five steps:

1. Specifying the application requirements and environment
2. Assessing application security risks
3. Creating and maintaining the ANF
4. Provisioning and operating the application
5. Auditing the security of the application

Security Management and Controls

CSA CCM

Terminology

Acronyms

Acronym	Definition
CCM	Cloud Controls Matrix
CSA	Cloud Security Alliance

Overview

The CCM is designed to provide **guidance** for cloud vendors and to assist cloud customers with assessing the overall security risk of a CSP. Can be used to perform security control audits.

The CSA CCM is an essential and up-to-date security controls framework that is addressed to the cloud community and stakeholders. A fundamental richness of the CCM is its ability to provide **mapping** and **cross relationships** with the main industry-accepted security standards, regulations, and controls frameworks (such as ISO 27001/27002, ISACA COBIT, and PCI DSS).

The CSA CCM framework gives organizations the necessary structure relating to information security tailored to the **cloud** industry.

The CCM allows you to note where specific controls (some of which you might already have in place) will address requirements listed in multiple regulatory and contractual standards, laws, and guides.

> ⓘ The CSA CCM can be particularly useful for assisting with **supply chain** reviews.

Components

Domains

The CCM can be seen as an inventory of cloud service security controls, arranged in the following 16 *separate* security domains:

- Application and Interface Security
- Audit Assurance and Compliance
- Business Continuity Management and Operational Resilience
- Change Control and Configuration Management
- Data Security and Information Lifecycle Management
- Data Center Security
- Encryption and Key Management
- Governance and Risk Management
- Human Resources
- Identity and Access Management
- Infrastructure and Virtualization Security
- Interoperability and Portability
- Mobile Security
- Security Incident Management, E-Discovery, and Cloud
- Supply Chain Management, Transparency, and Accountability
- Threat and Vulnerability Management

Inclusions

- AICPA 2009 TSC Map
- AICPA Trust Service Criteria (SOC 2SM Report)
- AICPA 2014 TSC
- BITS Shared Assessments AUP v5.0
- BITS Shared Assessments SIG v6.0
- BSI Germany
- Canada PIPEDA
- CCM V1.X
- CIS-AWS-Foundation v1.1
- COBIT 4.1

- COBIT 5.0
- COPPA
- CSA Enterprise Architecture (formerly Trusted Cloud Initiative)
- CSA Guidance V3.0
- ENISA IAF
- 95/46/EC - European Union Data Protection Directive
- FedRAMP Security Controls (Final Release, Jan 2012) –LOW IMPACT LEVEL–
- FedRAMP Security Controls (Final Release, Jan 2012) –MODERATE IMPACT LEVEL–
- FERPA
- GAPP (Aug 2009)
- HIPAA / HITECH Act
- HITRUST CSF v8.1
- ISO/IEC 27001:2013
- ISO/IEC 27002:2013
- ISO/IEC 27017:2015
- ISO/IEC 27018:2015
- ITAR
- Jericho Forum
- Mexico - Federal Law on Protection of Personal Data Held by Private Parties
- NERC CIP
- NIST SP 800-53 R3
- NIST SP 800-53 R4 App J
- NZISM
- NZISM v2.5
- ODCA UM: PA R2.0
- PCI DSS v2.0
- PCI DSS v3.0
- PCI DSS v3.2
- Shared Assessments 2017 AUP
- IEC 62443-3-3:2013
- C5
- NIST SP 800-R4 Moderate
- AICPA TSC 2017
- FedRAMP R4 Moderate

FedRAMP

Overview

The Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. FedRAMP enables Agencies to rapidly adapt from old, insecure legacy IT to mission-enabling, secure, and cost effective cloud-based IT. FedRAMP created and manages a core set of processes to ensure effective, repeatable cloud security for the government.

> ✓ **Fact.** The FedRAMP standard dictates that American federal agencies must retain their data within the boundaries of the United States, including data within cloud datacenters.

ISO/IEC 27001:2013

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

ISO/IEC 27001 was originally developed and created by the British Standards Institute, under the name of BS 7799. ISO 27001 is the standard to which organization's certify, as opposed to ISO 27002, which is the best practice framework to which many others align.

The standard provides "established guidelines and general principles for initiating, implementing, maintaining, and improving information security management with an organization." It looks for the information security management system (ISMS) to address the relevant risks and components in a manner that is appropriate and adequate based on the risks.

Components

ISMS

An ISMS typically ensures that a structured, measured, and ongoing view of security is taken across an organization, allowing security impacts and risk-based decisions to be taken. ISO/IEC 27001 is a standard framework for implementing and managing an ISMS based on the PDCA model.

Plan

Establish all the necessary objectives and processes to deliver results in accordance with the expected output.

Do

Implement the new processes.

Check

Measure the results of the new processes and hold them against the expected results in order to determine the differences.

Act

Analyze the differences generated in the check stage to determine their cause and decide where to apply changes.

ISO/IEC 27001 consists of 35 control objectives and 114 controls spread over 14 domains. The controls are mapped to address requirements identified through a formal risk assessment. The following domains make up ISO 27001:

Annex	Control
A.5	Information security policies
A.6	Organization of information security
A.7	Human resource security
A.8	Asset management
A.9	Access control
A.10	Cryptography

A.11	Physical and environmental security
A.12	Operations security
A.13	Communication
A.14	System acquisition, development, and maintenance
A.15	Supplier relationships
A.16	Information security incident management
A.17	Information security aspects of business continuity management
A.18	Compliance

ISO/IEC 27002:2013

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

ISO/IEC 27002 provides **guidelines** for organizational information security standards including the selection, implementation, and management of controls taking into consideration the organization's information security risk environments.

It is designed to be used by organizations that intend to select controls within the process of implementing an ISMS based on ISO/IEC 27001.

> ✓ **Fact.** ISO/IEC 27002 is how ISO/IEC 27001 is accomplished.

ISO/IEC 27017:2015

Terminology

Acronyms

Acronym	Definition
IEC	International Electrotechnical Commission
ISO	International Organization for Standardization

Overview

ISO/IEC 27017 offers guidelines for information security controls applicable to the provision and use of **cloud services** by providing additional implementation guidance for relevant controls specified in ISO/IEC 27002 and additional controls with implementation guidance that specifically relate to **cloud services.**

> ⓘ ISO 27017 provides controls and implementation guidance for both CSPs and cloud service customers.

NIST SP 800-53

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology
SP	Special Publication

Overview

The primary goal and objective of the NIST SP 800-53 standard is to ensure that appropriate **security requirements and security controls** are applied to all U.S. **federal** government information and information management systems.

Components

Although the NIST Risk Management Framework provides the pieces and parts for an effective security program, it is aimed at **government agencies** focusing on the following key components:

- 2.1 Multitiered Risk Management
- 2.2 Security Control Structure
- 2.3 Security Control Baselines
- 2.4 Security Control Designations
- 2.5 External Service Partners
- 2.6 Assurance and Trustworthiness
- 2.7 Revisions and Extensions

PCI DSS

Terminology

Acronyms

Acronym	Definition
AoC	Attestation of Compliance
ASV	Approved Scanning Number
DSS	Data Security Standards
ISA	Internal Security Assessor
PAN	Primary Account Number
PCI	Payment Card Industry
PFI	PCI Forensic Investigator
QSA	Qualified Security Assessor
RoC	Report of Compliance
SAD	Sensitive Authentication Data
SAQ	Self-Assessment Questionnaire

Definitions

ISA

An ISA is an individual who has earned a certificate from the PCI Security Standards Company for their sponsoring organization. This certified person has the ability to perform PCI self-assessments for their organization.

QSA

QSAs are the independent groups/entities which have been certified by PCI SSC for compliance confirmation in organization procedures.

RoC

A RoC is a form that has to be filled by *all level 1 merchants* undergoing a PCI DSS audit. The ROC form is used to verify that the merchant being audited is compliant with the PCI DSS standard.

SAD

Examples include CVV or PIN numbers.

SAQ

The PCI DSS SAQs are validation tools intended to assist merchants and service providers report the results of their PCI DSS self-assessment.

Overview

VISA, MasterCard, and American Express established PCI DSS as a security **standard** to which all organizations or merchants that accept, transmit, or store cardholder data, regardless of size or number of transactions, **must comply.**

> ⓘ PCI DSS is a standard, *not* a regulation.

Components

PCI DSS is a comprehensive and intensive security standard that lists both technical and nontechnical requirements based on the number of credit card transactions for the applicable entities.

Merchant Levels

Level	Description
1	Any merchant-regardless of acceptance channel-processing more than 6 million transactions per year. Any merchant that the credit card issuer, at its sole discretion, determines should meet the Level 1 requirements to minimize risk to the credit card issuer's system.
2	Any merchant-regardless of acceptance channel-processing 1-6 million credit card transactions per year.
3	Any merchant processing 20,000 to 1 million credit card e-commerce transactions per year.
4	Any merchant processing fewer than 20,000 credit card e-commerce transactions per year and all other merchants-regardless of acceptance channel-processing up to 1 million credit card transactions per year.

> ⓘ Merchants at different tiers are required to have more or fewer audits in the same time frame as merchants in other tiers, depending on the tier. Businesses at level 1 are required to undergo a **yearly PCI audit** conducted by a QSA.

Merchant Requirements

All merchants, regardless of level and relevant service providers, are required to comply with the following **12 domains/requirements:**

- Install and maintain a firewall configuration to protect cardholder data
- Avoid using vendor-supplied defaults for system passwords and other security parameters.
- Protect stored cardholder data.
- Encrypt transmission of cardholder data across open, public networks.
- User and regularly update antivirus software.
- Develop and maintain secure systems and applications.

- Restrict access to cardholder data by business need-to-know.
- Assign a unique ID to each person with computer access.
- Restrict physical access to cardholder data.
- Track and monitor all access to network resources and cardholder data.
- Regularly test security systems and processes.
- Maintain a policy that addresses information security.

The 12 requirements list more than 200 controls that specify required and minimum security requirements for the merchants and service providers to meet their compliance obligations.

Supply Chain

ISO 28000:2007

Terminology

Acronyms

Acronym	Definition
ISO	International Organization for Standardization
PDCA	Plan-Do-Check-Act

Overview

Focus is to assist organizations in the identification and implementation of controls to safeguard people, products, and assets.

In line with other ISO security-related management systems, ISO 28000 focuses on the use of PDCA as a lifecycle of continual improvement and enhancement.

ISO 28000 defines a set of security management requirements and provides for certification against certain relevant elements that relate to risk:

- Security management policy
- Organizational objectives
- Risk management practices
- Documented practices and records
- **Supplier relationships**
- Roles, responsibilities, and authorities
- Use of PDCA
- Organizational procedures and processes

> ⓘ Because ISO 28000 defines a set of security management requirements, the onus is on the organization to establish a security management system that meets the standard's requirements.

Models and Guidance

Application Risk Management

OWASP Top Ten

Terminology

Acronyms

Acronym	Definition
OWASP	Open Web Application Security Project

Overview

To address these vulnerabilities, organizations must have an application risk-management program in place. Implementation of an application risk-management program addresses not only vulnerabilities but also all risks associated with applications.

Risks from 2020

1. Injection
2. Broken Authentication
3. Sensitive Data Exposure
4. XML External Entities (XXE)
5. Broken Access Control
6. Security Misconfiguration
7. Cross-Site Scripting (XSS)
8. Insecure Deserialization
9. Using Components with Known Vulnerabilities
10. Insufficient Logging and Monitoring

1. Injection

An injection attack occurs when an attacker sends malicious statements to an application via data input fields. Another way to say this is that untrusted data is sent to an interpreter as part of a command or query. These could be SQL queries, LDAP queries, or other forms of injection.

Prevention

- Whitelisting input validation/bounds checking (preventing what types of data can be input)
- Using prepared statements
- Escaping all user supplied input

2. Broken Authentication

Occurs when authentication and session management application functions are not implemented correctly.

Impact

- Allows attackers to compromise passwords, keys, or session tokens
- Exploitation of other implementation flaws to assume other identities

Prevention

- Do not use custom authentication schemes.
- Rotate session IDs after a successful login.
- Do not allow simple passwords to be used.
- Ensure the connection is encrypted so credentials aren't exposed.

3. Sensitive Data Exposure

Commonly allowed when web applications do not properly protect sensitive data, such as credit cards, tax IDs, and authentication credentials. Attackers may steal or modify poorly protected data to initiate credit card fraud, identity theft, or other felonies. Even with proper encryption methods put in place, sensitive data is still at risk if the client's browser is insecure.

Impact

- Attackers may steal or modify weakly protected data to conduct fraud, theft, or other crimes.

Prevention

- Encryption (at rest or in transit)
- Perform checks against client browsers to ensure they meet security standards. If the browser doesn't meet the security standards, it can be prevented access to the web application.

4. XML External Entities (XXE)

XML external entities refer to references, such as the application directory structure or the configuration of the hosting system, that should be removed from the code, but are left in by accident. These items can provide information to an attacker that may allow them to circumvent authentication measures to gain access.

Attackers can exploit vulnerable XML processors if they can upload XML or include hostile content in an XML document, exploiting vulnerable code, dependencies, or integrations.

5. Broken Access Control

> This risk is also known as *missing function-level access control.*

Web applications typically verify function level access rights before allowing that functionality from the UI. Best practice requires applications to perform access control checks on the server when each function is accessed. Broken access control occurs when requests are not verified.

Impact

- Attackers can forge requests allowing access to functionality without proper authorization.

Prevention

- Set the default to deny all access to functions, and require authentication/authorization for each access request. This ensures that no particular function may be run without explicitly ensuring that it was called by an authorized user.
- Run a process as both a user and privileged user, compare results, and determine similarity.

6. Security Misconfiguration

Prevention

- Secure settings should be defined, implemented, and maintained, as defaults are well known to attackers.
- Software should be patched regularly to keep it up to date.
- Software settings should be kept up to date.

7. Cross-Site Scripting (XSS)

Occurs when an application receives untrusted data and then sends it to a web browser without proper validation.

Impact

- Allows attackers to execute scripts in a victim's browser which can hijack user sessions, deface websites, or redirect the user to malicious websites.

Prevention

- Use an auto-escaping template system
- Put untrusted data in only allowed slots of HTML documents
- HTML escape when including untrusted data in any HTML elements
- Use the attribute `escape` when including untrusted data in attribute elements
- Sanitize HTML markup with a library designed for the purpose

> ⓘ This is not a threat to the back-end database, but a threat to the client.

8. Insecure Deserialization

Serialization is the process of turning some object into a data format that can be restored later. People often serialize objects in order to save them to storage, or to send as part of communications.

Deserialization is the reverse of serialization; taking data structured from some format and rebuilding it into an object. The features of these native deserialization mechanisms can be repurposed for malicious effect when operating on untrusted data. Attacks against deserializers have been found to allow DoS, access control, and remote code execution attacks.

Today the most popular data format for serializing data is JSON. Before that, it was XML.

9. Using Components with Known Vulnerabilities

Components, such as libraries, frameworks, and other software modules, almost always run with full privileges.

Impact

- This could allow an attack to undermine application defenses and launch unpredictable attacks.
- Data loss or server takeover.

Prevention

- Applications using components with known vulnerabilities should be quarantined or, at an absolute minimum, have special monitoring to prevent application attacks.

> ⓘ It's important to remember that these are *known* vulnerabilities, not unknown. Developers are willingly using these components *knowing* that they have vulnerabilities. This could be for a variety of reasons, including the fact that they

> may not actually be leveraging a "vulnerable" aspect of a particular component in their application.

10. Insufficient Logging and Monitoring

Risks from 2013

1. Injection
2. Broken Authentication and Session Management
3. Cross-Site Scripting (XSS)
4. **Insecure Direct Object References**
5. Security Misconfiguration
6. Sensitive Data Exposure
7. Missing Function-Level Access Control
8. **Cross-Site Request Forgery (CSRF)**
9. Using Components with Known Vulnerabilities
10. **Invalidated Redirects and Forwards**

4. Insecure Direct Object References

When a developer allows a reference to an internal implementation object to be exposed. The exposure could be a file, directory, or database key. An example could be the following URL:

```
www.sybex.com/authoraccounts/benmalisow
```

The URL reveals location of specific data as well as the format for potential other data (such as other authors' pages/accounts).

Impact

- Attackers may be able to manipulate these references to access unauthorized data.

Prevention

- Refrain from including direct access information in URLs.
- Check access each time a direct object reference is called by an untrusted source.
- Run a process as both user and privileged user, compare results, and determine similarity; this will help you determine if there are functions that regular users should not have access to and thereby demonstrate that you are missing necessary controls.

8. Cross-Site Request Forgery (CSRF)

Occurs when a logged-on user's browser sends a forged HTTP request along with cookies and other authentication information, forcing the victim's browser to generate a request that the application thinks is a legitimate request from the user.

Impact

- The attacker could have the user log into one of the user's online accounts.
- The attacker could collect the user's online account login credentials to be used by the attacker later.
- The attacker could have the user perform an action in one of the user's online accounts.

Prevention

- Ensure that all HTTP resource requests include a unique, unpredictable token.
- Include a CAPTCHA code as part of the user resource request process.

10. Invalidated Redirects and Forwards

Redirection to unauthorized pages, often in conjunction with a social engineering/phishing aspect.

Prevention

- Don't use redirects/forwards in your applications.
- Train users to recognize invalidated links.

OWASP WSTG

Testing Methods

1. Information Gathering

2. Configuration and Deployment Management Testing

3. Identity Management Testing

4. Authentication Testing

5. Authorization Testing

6. Session Management Testing

7. Input Validation Testing

8. Testing for Error Handling

9. Testing for Weak Cryptography

10. Business Logic Testing

11. Client Side Testing

Cloud Computing

NIST SP 500-293

Terminology

Acronyms

Acronym	Definition
NIST	National Institute of Standards and Technology
SP	Special Publication

Overview

The NIST Cloud Technology Roadmap helps CSPs develop industry-recommended, secure, and interoperable identity, access, and compliance management configurations and practices. It offers guidance and recommendations for enabling security architects, enterprise architects, and risk-management professionals to leverage a common set of solutions that fulfill their common needs to be able to assess where their internal IT and CSPs are in terms of security capabilities and to plan a roadmap to meet the security needs of their business.

Cloud Computing Certification

ENISA CCSL

Terminology

Acronyms

Acronym	Definition
CCSL	Cloud Certification Schemes List
CCSM	Cloud Certification Schemes Metaframework
ENISA	European Union Agency for Cybersecurity

Overview

Includes the CCSL and CCSM which, together, aid in selecting a certification scheme for cloud computing customers.

Components

CCSL

The Cloud Certification Schemes List (CCSL) provides an overview of different existing **certification** schemes. It describes the main characteristics relevant to cloud computing and cloud computing customers.

CCSL answers questions like these:

1. Part 1 - General information about the certification scheme
2. Part 2 - Which are the underlying security standards or best practices?

3. Part 3 - Assessments and certification of compliance (such as which organizations are accredited to issue certificates)
4. Part 4 - Current adoption and usage (such as how many providers have obtained certification)
5. Part 5 - Security objectives

The schemes that make up the CCSL include:

- Certified Cloud Service - TUV Rhineland
- CSA Attestation - OCF level 2
- CSA Certification - OCF level 2
- CSA Self-Assessment - OCF level 1
- EuroCloud Self-Assessment (ECSA Self-Assessment)
- EuroCloud Star Audit Certification (ECSA Audit)
- ISO/IEC 27001 Certification
- PCI DSS v3.1
- LEET Security Rating Guide
- AICPA SOC 1
- AICPA SOC 2
- AICPA SOC 3

CCSM

The Cloud Certification Schemes Metaframework (CCSM) is an extension of the CCSL designed to provide a high-level mapping of security requirements of the customer to security objectives in existing cloud security schemes. There are 27 CCSM security objectives that the customer can select to cross-reference against.

> ⓘ The CCSM is basically a comparison matrix/table of the CCSL schemes and the CCSM security objectives which displays whether certain framework contains a certain objective (such as risk management). The matrix allows for easy visibility into whether a particular framework contains what a cloud customer may require.

Cloud Computing Risk Management

CSA Treacherous Twelve

Terminology

Acronyms

Acronym	Definition
CSA	Cloud Security Alliance

Overview

The CSA Treacherous Twelve is a report of threats related specifically to cloud computing.

Components

1. *Data Breaches:* If a multitenant cloud service database is not properly designed, a flaw in one client's application can allow an attacker access not only to that client's data but to every other client's data as well.
2. *Insufficient Identity, Credential and Access Management*
3. *Insecure Interfaces and APIs:* Cloud computing providers expose a set of software interfaces or APIs that customers use to manage and interact with cloud services. Provisioning, management, orchestration, and monitoring are all performed using these interfaces. The security and availability of general cloud services is dependent on the security of these basic APIs. From authentication and access control to encryption and activity monitoring, these interfaces must be designed to protect against both accidental and malicious attempts to circumvent policy.
4. *System Vulnerabilities*
5. *Account Hijacking:* If attackers gain access to your credentials, they can eavesdrop on your activities and transactions, manipulate data, return falsified information, and

redirect your clients to illegitimate sites. Your account or service instances may become a new base for the attacker.

6. *Malicious Insiders:* European Organization for Nuclear Research (CERN) defines an insider threat as "A current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems."
7. *Advanced Persistent Threats*
8. *Data Loss:* Any accidental deletion by the CSP, or worse, a physical catastrophe such as a fire or earthquake, can lead to the permanent loss of customers' data unless the provider takes adequate measures to back it up. Furthermore, the burden of avoiding data loss does not fall solely on the provider's shoulders. If a customer encrypts their data before uploading it to the cloud but loses the encryption key, the data is still lost.
9. *Insufficient Due Diligence:* Too many enterprise jump into the cloud without understanding the full scope of the undertaking. Without a complete understanding of the CSP environment, applications, or services being pushed to the cloud, and operational responsibilities such as incident response, encryption, and security monitoring, organizations are taking on unknown levels of risk in ways they may not even comprehend but that are a far departure from their current risks.
10. *Abuse and Nefarious Use of Cloud Services:* It might take an attacker years to crack an encryption key using his own limited hardware, but using an array of cloud servers, he might be able to crack it in minutes. Alternatively, he might use that array of cloud servers to stage a DDoS attack, serve malware, or distribute pirated software.
11. *Denial of Service:* By forcing the victim cloud service to consume inordinate amounts of finite system resources such as process power, memory, disk space, and network bandwidth, the attacker causes an intolerable system slowdown.
12. *Shared Technology Issues:* Whether it's the underlying components that make up this infrastructure that were not designed to offer strong isolation properties for a multitenant architecture (IaaS), redeployable platforms (PaaS), or multicustomer applications (SaaS), the threat of shared vulnerabilities exists in all delivery models. A defense-in-depth strategy is recommended and should include compute, storage, network, application and user security enforcement, and monitoring, whether the service model is IaaS, PaaS, or SaaS. The key is that a single vulnerability or misconfiguration can lead to compromise across an entire provider's cloud.

Security Management and Controls

CIS CSC

Terminology

Acronyms

Acronym	Definition
CIS	Center for Internet Security
CSC	Critical Security Controls

Overview

A recommended set of actions for cyber-defense that provide specific and actionable ways to stop today's most pervasive attacks. The guidelines consist of 20 key actions, called critical security controls (CSC), that organization should implement to block or mitigate known attacks.

Controls

- CSC 1: Inventory of Authorized and Unauthorized Devices
- CSC 2: Inventory of Authorized and Unauthorized Software
- CSC 3: Continuous Vulnerability Assessment and Remediation
- CSC 4: Controlled Use of Administrative Privileges
- CSC 5: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations, and Servers
- CSC 6: Maintenance, Monitoring, and Analysis of Audit Logs
- CSC 7: Email and Web Browser Protections
- CSC 8: Malware Defenses
- CSC 9: Limitation and Control of Network Ports, Protocols, and Services

- CSC 10: Data Recovery Capability
- CSC 11: Secure Configurations for Network Devices such as Firewalls, Routers, and Switches
- CSC 12: Boundary Defense
- CSC 13: Data Protection
- CSC 14: Controlled Access Based on the Need to Know
- CSC 15: Wireless Access Control
- CSC 16: Account Monitoring and Control
- CSC 17: Security Skills Assessment and Appropriate Training to Fill Gaps
- CSC 18: Application Software Security
- CSC 19: Incident Response and Management
- CSC 20: Penetration Tests and Red Team Exercises

An underlying theme for the controls is support for large-scale, standards-based security automation for the management of cyber defenses.

CIS CSC 2

Effectiveness Metrics

When testing the effectiveness of the automated implementation of this control, organizations should determine the following:

- The amount of time it takes to detect new software installed on the organization's systems
- The amount of time it takes the scanning functions to alert the organization's administrators when an unauthorized application has been discovered on a system
- The amount of time it takes for an alert to be generated when a new application has been discovered on a system
- Whether the scanning function identifies the department, location, and other critical details about the unauthorized software that has been detected

Automation Metrics

Organizations should gather the following information to automate the collection of relevant data from these systems:

- The total number of unauthorized applications located on the organization's business systems
- The average amount of time it takes to remove unauthorized applications from the organization's business systems
- The total number of the organization's business systems that are not running whitelisting software
- The total number of applications that have been recently blocked

Threat Models

DREAD

Overview

DREAD is part of a system for risk-assessing computer security threats previously used at Microsoft and although currently used by OpenStack and other corporations it was abandoned by its creators. It provides a mnemonic for risk rating security threats using five categories.

The categories are:

- **D**amage
- **R**eproducibility
- **E**xploitability
- **A**ffected users
- **D**iscoverability

When a given threat is assessed using DREAD, each category is given a rating from 1 to 10. The sum of all ratings for a given issue can be used to prioritize among different issues.

RISK_DREAD = (Damage + Reproducibility + Exploitability + Affected Users + Discoverability) / 10

Categories

[D]amage

[R]eproducibility

Reproducibility is the measure of how easy an exploit is to reproduce.

[E]xploitability

[A]ffected Users

[D]iscoverability

STRIDE

Overview

Created by Microsoft, the STRIDE threat model provides a standardized way of describing threats by their attributes. It is used to parse the various types of attacks and malicious techniques that might be used against software.

> ⓘ STRIDE is particularly useful as part of the software development lifecycle in attempting to identify vulnerabilities throughout the build process. These six concepts help in identifying and classifying threats or vulnerabilities and help form a common language used to describe them.

The attributes include:

- **S**poofing
- **T**ampering
- **R**epudiation
- **I**nformation Disclosure
- **D**enial of Service
- **E**levation of Privilege

Attributes

[S]poofing

Attacker assumes identity of subject.

[T]ampering

Data or messages altered by an attacker.

[R]epudiation

Illegitimate denial of an event.

[I]nformation Disclosure

Information obtained without authorization.

[D]enial of Service

Attacker overloads system to deny legitimate access.

[E]levation of Privilege

Attacker gains a privilege level above what is permitted.