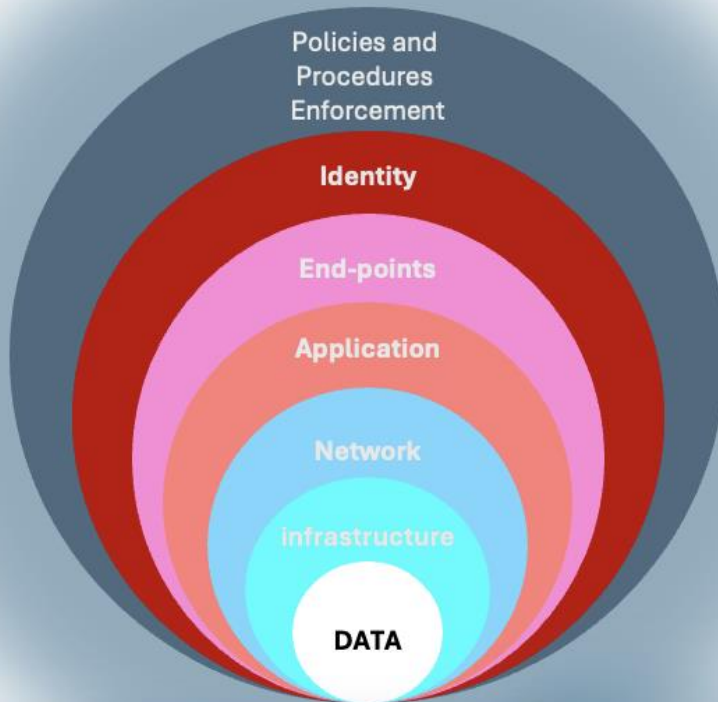


Enterprise Data Security for US Europe and Asia



Prabhat K. Andleigh

[Document title]

Enterprise Data Security for US Europe and Asia

Prabhat K. Andleigh

Enterprise Data Security for US Europe and Asia

Enterprise Data Security for US Europe and Asia

Copyright © 2024 by Prabhat K. Andleigh

All rights reserved. No part of this book may be reproduced in any form or by any electronic or mechanical means including information storage and retrieval systems without permission in writing from its author Prabhat K. Andleigh or its publisher, Amazon Kindle books.

Table of Contents

Chapter 1 – Driving Need for Data Security	11
<i>Cybersecurity vs. Data Security</i>	11
<i>Security Policies</i>	12
<i>Worldwide Need for Data Security</i>	12
<i>Data is an Asset</i>	13
<i>Data Loss Prevention</i>	13
Enterprise DLP	13
Integrated DLP	14
Cloud-Native DLP	14
<i>The Regulatory Landscape along with the emergency threat landscape</i>	14
The impact of Growing to Co-Location and Public Clouds	14
Threat Management	15
Next Steps	15
Chapter 2 – Regulatory Framework	16
Data Protection Laws	16
<i>Legal responsibility arises from a variety of reasons:</i>	17
Cost Liability and Legal Aspects for Data Breaches	17
<i>International Laws and Regulations for Data Security</i>	19
Europe	19
United States	19
South America	20
Asia	20
Australia	20
International	21
<i>Data Protection Laws by Countries and Regions</i>	25
EU General Data Protection Regulation (GDPR)	25
India Privacy Law	27
Laws in the United States of America	28

Enterprise Data Security for US Europe and Asia

OECD – Organization for Economic Cooperation and Development	31
<i>Standards and Compliance</i>	32
Global Standards	32
US Standards	33
International Standards Driving Data Security	37
ENISA – European Union Agency for Network and Information Security	39
Industry Standards	40
<i>Audits and Risk</i>	41
Audit Plan	42
<i>Risk Management</i>	42
<i>Certification Frameworks</i>	43
CSA STAR – Security, Trust Assurance, and Risk	44
SSAE 18	46
SSAE SOC Reports	46
<i>COBIT</i>	47
ITIL and ITSM	48
Chapter 3 – Cloud Impact on Data Security	52
<i>Understanding the Cloud Environments</i>	52
<i>Deployment & Service Models</i>	54
Service Models	54
Deployment Models	57
<i>Cloud Security</i>	59
Logging and Monitoring	59
Cloud Segmentation	60
<i>Cloud Selections Best Practices</i>	60
<i>Major Considerations for Moving to Public Clouds</i>	62
Cloud Migration Challenges	63
<i>Cloud Architecture</i>	64
Build and Configuration Management Automation	64
Cloud Storage	64
Containers and Kubernetes (k8s)	65
Virtual machine (VM) Related Architecture	65
Security Governance Plans	66
General Cloud Risks and Challenges	66

Enterprise Data Security for US Europe and Asia

<i>Enterprise Risk Management</i>	67
<i>Service Level Agreement (SLA)</i>	69
<i>Cloud Audits</i>	70
Chapter 4 – Data Classification Types and Standards	72
Data Life Cycle	72
<i>Data Classification</i>	75
Data Classification by Type	75
<i>Understanding Regulatory Compliance Requirements</i>	77
<i>Steps to Classify Data</i>	77
<i>Data Security Posture Management</i>	78
Chapter 5 – When Data Needs to be Protected	80
<i>The Type of Data That Requires Protection</i>	80
What Is Personal Data	81
PII (Personally Identifiable Information)	81
PI (Personal Information)	82
Sensitive Information	82
Protecting PI, PII and Sensitive Data	83
<i>Protecting Data at Rest, Data in Motion, Data in Use</i>	85
<i>Data Risk Assessment</i>	87
<i>Key Elements of Data Protection</i>	88
<i>Data Protection Framework</i>	89
<i>Data Protection Best Practices</i>	89
Chapter 6 – Who Is Responsible for Data Protection	91
Everyone Has Responsibility for Data Security	91
<i>Responsibility By Deployment Organization Type</i>	92
<i>Responsibility By Role Played by the Employee</i>	94
Can you be a processor of some data and a controller of other data at the same time?	99
<i>Access Controls Based on Responsibility</i>	99

Chapter 7 - Data Storage and Security Architecture	101
<i>Storage Security Management</i>	102
Ensuring data Confidentiality	102
Types of Storage	103
Best Practices for Successful Data Management	103
<i>Systems of Record and Reference</i>	104
<i>Data Storage Types</i>	107
Definitions of a Database	107
<i>Data Architectures</i>	108
Types of Data Architecture	109
Domain and Cross-platform Context	109
Enterprise Data Model	110
Data Integration Architecture	112
<i>Data Lake and Data Lakehouse</i>	113
Data Management Strategy for Hybrid Cloud	114
Recommendations for Hybrid Cloud Implementation	114
Chapter 8 - Using Encryption Technologies	118
<i>Modern Encryption Management</i>	119
Public Keys vs. Private Keys	120
<i>Managing Encryption Keys</i>	121
Hardware Security Module (HSM) vs Trusted Platform Module (TPM)	123
<i>Encryption Types</i>	124
TLS/SSL – Layer 4 – Transport Control Layer	124
<i>Encryption Algorithms</i>	125
Symmetric Key Encryption	125
Asymmetric Encryption	125
Secure Hashing Algorithms	126
Application Level Encryption (ALE)	126
Homomorphic and Polymorphic Encryption	127
Format-preserving Encryption	129
<i>Self-decrypting Encryption</i>	129
Transparent Encryption	130
<i>Data Alteration Treatments for Data Protection</i>	130
Secure User Access vis VPN (Virtual Private Network)	134
Crypto-shredding Before Data or Storage Device Destruction	134

Chapter 9 – Identity and Access Management	136
<i>Scope of IAM Systems</i>	136
Modern IAM Systems	137
IAM Systems as the Common Link	138
<i>IAM Roles, Groups, and Entitlements</i>	139
<i>Key Terms Used for IAM</i>	139
<i>IAM Threats and Risk Management</i>	141
Governance, Risk, and Compliance (GRC)	141
<i>IAM Policy Management</i>	142
IAM policy guidelines	143
Policy Sections for Cloud Environments	145
<i>Departmental View of IAM</i>	145
<i>Selecting IT and Data Security Solutions</i>	147
Neutral and Independent IAM Solution	147
Customization	148
<i>IAM Across the Enterprise</i>	150
<i>Federating IAM</i>	154
CFS (Cloud Federation Service):	155
Federating Access	155
<i>Identity Life Cycle</i>	155
<i>Identity Governance and Administration</i>	158
IAM vs IGA	158
Governance Model	159
Operational Efficiency and Excellence	160
Compliance Management	161
Chapter 10 – Application Security Architecture and Data Exploits	162
<i>Threat Models and Attacks</i>	163
Application Attack Points	163
<i>Comprehensive Application Security Framework</i>	165
Data Exploits and Security Vulnerability	168
Vulnerability Management	169
<i>10 Most Common Web Security Vulnerabilities</i>	171
OWASP 10	171

Enterprise Data Security for US Europe and Asia

Cloud Security Vulnerabilities	177
<i>Exploits in the Cloud</i>	180
Hypervisor Attacks	180
Guarding Against Virtualization Security Vulnerabilities	184
Virtual Security: Developing a Plan and Procedures	185
Best Practices for Improving VM Security	185
Chapter 11 – Defense in Depth and Security Architecture	187
<i>Defense-in-Depth</i>	187
Elements of defense in depth	188
Defense-in-Depth vs. Zero Trust	189
Best Practice for Protecting Workstations	190
<i>Security Architecture</i>	192
Open Systems Interconnection (OSI) Architecture Framework	193
<i>Security Architecture Frameworks</i>	195
TOGAF Framework	195
SABSA Framework	195
OSA Framework	196
Enterprise information Systems Architecture (EISA)	196
<i>Cloud Security Architecture</i>	196
Benefits of Security Architecture	198
Chapter 12 - Data Security Architecture	199
<i>Significant Aspects of Data Security Architecture</i>	200
Protecting the Data Assets	200
Four Elements of Data Security	200
The Three Areas of Data Security	201
<i>Data Security Reference Architecture</i>	202
Security Architecture Framework	203
Common Data Security Architecture	204
Basic Categories of Security Services	205
Selecting CDSA Components	207
<i>Building a Data Security Architecture</i>	208
Use of Layered Technologies for Data Security Architecture	209
Securing User Accounts	209
Securing User Devices	209
Securing Data Centers	209

Enterprise Data Security for US Europe and Asia

Logging and Monitoring of Events	210
The States of Data - Securing Content	210
Data Security Architecture Best Practices	213
<i>Architecting and Designing Data Security Architecture</i>	213
Data Security Architecture to Facilitate Operations	215
Drawing the Security Architecture Diagram?	216
Chapter 13 - Business Continuity & Disaster Recovery	218
<i>Business Continuity</i>	218
Recovery Point Objective) & Recovery Time Objective)	219
<i>BCDR Plan</i>	221
<i>Improving Business Resilience</i>	223
Best Practices for Improving Business Resilience	223
Keep on Top of Business Continuity	226
Chapter 14 - Data Security Controls	228
Standards That Define Security Controls	229
Chapter 15 - Data Security Governance	240
<i>Data Governance Framework</i>	240
Policies and Standards Framework	241
Data Security Governance Framework	241
Data Classification Governance	241
Architecture and Design framework	246
Deployment Framework -- On-Prem and Cloud	246
<i>Benefits of Data Security Governance</i>	246

Chapter 1 – Driving Need for Data Security

Data is what drives the business and profits of every organization. As your organization's data footprint expands across on-prem, colocation service providers and various cloud providers, along with SaaS applications, your risk of a data breach and exposure also increases. Hackers and cybercriminals are seeking to exploit security vulnerabilities to access sensitive data that is spread across multiple cloud data centers and data stores. The hackers and cybercriminals may be state actors with political motives, individuals extracting revenge, or looking for financial gain through ransom activities. Ransoms to release locks on your data can run into huge payouts.

So that is why it is so critical that your organization invest in data security. Especially if you are a financial services company such as a retail or investment bank or even a large corporation which stores personal sensitive data such as date of birth and Social Security number in the US, and Aadhar or PAN card in India, you may be obligated to ensure data security and meet government standards for ensuring data security.

Fraud and cybercrime have been rising year-by-year for many years resulting in not just significant risk of but actual financial loss and loss of reputation for many organizations. Exposure of your confidential financial and sales data is itself sufficient reason to implement robust enterprise data security. But more importantly, potential exposure of your customer and workforce personal data such as IDs, Social Security and Tax IDs and other non-public data can have even more severe detrimental consequences.

The scope of data security includes defining data security policies and frameworks for both building and managing data assets within a company. Your data security officer should enforce compliance to data security policies and standards, and robust security frameworks within the company on an ongoing basis.

Cybersecurity vs. Data Security

Cybersecurity refers to every aspect of protecting your organization and its employees and assets against cyber threats. As cyberattacks become more common and sophisticated, and corporate networks grow more complex with the network

Enterprise Data Security for US Europe and Asia

perimeter extending into the public clouds, a variety of solutions are required to mitigate corporate cyber risk.

Data security is an important subset of cybersecurity. Cyberattacks may have originated due to ransom demands or by state actors looking for political or national disruptions. But one of the primary reasons for cyberattacks is exfiltration of corporate data: employee IDs and password, sensitive customer personal information data and financial records of an organization. These can be sold by bad actors on the dark web to ready buyers who can monetize this data

So, while it is important to protect the corporate extended network infrastructure, it is equally important to secure the corporate confidential and highly confidential data at rest, in transit (or in motion) and in use. Overall cybersecurity is a pre-requisite for data security.

Security Policies

Security policies, and especially data security policies, for enterprise data should be aligned to the overall Information Technology (IT) security policies of the enterprise. The policies should be defined for each architectural component within the enterprise data architecture. The policies should be based on various security related compliances, specifically related to personal identifiable information (PII), confidential data, private data etc. The policies should highlight requirements related to authorization, encryption, masking, audits, and reporting.

Worldwide Need for Data Security

Data security refers to the process of protecting data from unauthorized access and data corruption throughout its lifecycle. Data security is the act of implementing protective digital privacy measures that are applied to prevent unauthorized access to computers, databases, and websites. Data security also protects data from corruption.

Data security includes data encryption, tokenization, and key management practices that protect data across all applications and platforms.

Data security is an essential aspect of IT for organizations of every size and type.

Data is an Asset

The data held by your organization including financial and records, customer information and employee and customer PII is an important asset for your organization. The revenues earned your organization depend on ensuring the protection of these assets from data breaches which can result in damage to the reputation of the firm and very frequently the demand for ransom where the organization's network is held hostage in a cyberattack.

The single most important reason for your organization is to implement data protection strategies is fear of financial loss. Data is recognized as an important corporate asset that needs to be safeguarded. Loss of information can lead to direct financial losses, such as lost sales, fines, or monetary judgments.

Data Loss Prevention

Data loss prevention (DLP) includes all tools, technologies, and initiatives to prevention loss of data due to unauthorized access and exposure of confidential and highly confidential data (collectively sensitive data) that should be accessible only to authorized users with proper permissions, privileges, and justification to access the data. DLP technologies use data classification labels and tags, content inspection techniques, and contextual analysis to identify sensitive data and determine protection actions. They also provide centralized policy management that drive these actions. DLP solutions can be described as Enterprise DLP, Integrated DLP and Cloud-Native DLP.

Enterprise DLP

Financial institutions, and corporate and government organizations that hold PI are required by international regulations and standards to ensure protection of their data assets. Enterprise DLP solutions satisfy the needs for regulatory compliance, internal policy and standards compliance, and overall proprietary and sensitive data protection. They provide centralized policy management and reporting, and implement controls and alerts across endpoints, networks, email, and cloud. They implement advanced content inspection and remediation.

Integrated DLP

These solutions include natively integrated services such as secure email or web gateway and endpoint. They also provide some level of policy management and reporting.

Cloud-Native DLP

These solutions address the cloud-native use cases such as, SaaS applications security and public cloud data security. These solutions are customized for cloud environments.

The Regulatory Landscape along with the emergency threat landscape

Financial institutions, especially banks, brokerage firms, and quite often large publicly traded insurance firms are held to a higher standard wherein regulatory agencies such as departments within the central and state governments, conduct periodic reviews of the organization's cybersecurity and data security posture, and tools used and the level of enforcement and governance.

The impact of Growing to Co-Location and Public Clouds

For smaller firms, the IT infrastructure is quite often limited to the in-house data centers. With growth, firms have extended their data centers to co-location data centers which support multiple client infrastructures. This allows a wide range of actors, not totally under the control of the organization, to access the network infrastructure of the colocation service. Most col-location services offer hosting of servers and provide power and network bandwidth leaving management of the infrastructure to the client.

The next level of growth is in the public clouds such as AWS (Amazon Web Services), GCP (Google Cloud Platform) and Microsoft Azure to take advantage of their utility model (pay for infrastructure used) and the ability to increase or decrease the infrastructure foot-print quickly. While most of the public cloud service providers manage their own infrastructure and provide a use-based utility model, ultimately the legal responsibility for data security rests with the customer organization.

Threat Management

Cyberattacks and threats to the security of data are an ongoing challenge for every organization. These cyberattacks are typically focused on data corruption and data exfiltration. Information fragmentation across the extended infrastructure, enterprise, co-location and cloud, leads to blind spots in security operations that can be exploited by bad actors.

Threat Management is the process your cybersecurity professionals need to use to detect and prevent cyberattacks and respond rapidly to security incidents. Threat and risk assessment tools identify exposure risks by determining potential security weaknesses and taking the appropriate actions to reduce the impact of threatening events and manage the risks.

Next Steps

Now that we understand the driving need for enterprise data security, in the following chapters we will dive deep into the different facets of data security and its underlying foundation, cybersecurity.

Chapter 2 – Regulatory Framework

If you are a financial services company such as a retail or investment bank, brokerage or an insurance company or even a large corporation which stores personal sensitive data such as date of birth and Social Security number in the US and Aadhar or PAN card in India, you may be obligated to ensure data security and meet government standards for ensuring data security.

Compliance with government regulations and standards such as NIST, GDPR, CCPA, SOX, DPDP, PCI, HIPAA and others, all have the same purpose: preventing unauthorized access to customers' and users' sensitive information.

To ensure compliance and prevent any data breaches or exposures, you will need to develop a comprehensive data security strategy and operating governance that spans on-premises, in clouds or in a hybrid cloud mode. The strategy should include how to investigate and remediate cyberthreats. These solutions can enforce security policies and access controls in near real time and help you meet regulatory compliance requirements, improving your data security posture.

Data Loss Prevention (DLP) is required to ensure that sensitive data is not being exposed unnecessarily. Your DLP strategy needs to comply with the laws and regulatory compliance requirements that your organization is subject to because of your range of geographic locations for doing business and the nature of your business and services.

Data Protection Laws

Due to rapid increase in hacking and data breaches over the past decade, many countries and states have passed a number of laws to protect data and prevent data breaches. It is really important to understand the laws and the jurisdictions in which they apply to design a data security strategy and operational governance for data

security, The following section discusses your responsibility to comply with the laws and associated regulatory standards provides a listing of these laws.

Legal responsibility arises from a variety of reasons:

Data that is collected by the organization can identify a person. When the organization controls sensitive or highly confidential data, it is required to protect that data from public disclosure. All national security laws make secure handling of sensitive data mandatory.

In a connected world, it is important to understand the jurisdictions that affect legal responsibility for data, and the operational laws that determine who is responsible and to what extent for inadvertent and malicious disclosure of protected data. The laws determine the penalties (criminal, civil and financial) that maybe brought to bear against the data owners and the senior management executives in the business entity; these can include private, public, and government owned entities.

For example, even though the sensitive data has been created within India and is subject to all legal jurisdictions of India, the data could become subject to other national jurisdictions by way of international commerce or by simply storage of the data in locations governed by laws in those jurisdictions. This is specifically applicable to data about living persons, especially if that living person is a citizen of that nation state, in this case India.

Not all nations in the world have strict laws about data pertaining to living persons.

Cost Liability and Legal Aspects for Data Breaches

Let us start by reviewing some key legal terms and potential actions associated with potential litigations for data loss.

The Preponderance of Evidence Principle: Civil courts are held to a “preponderance of Evidence” principle. For example, in the case of an employee uploading sensitive info to a Cloud Service Provider (CSP such as AWS or GCP), and if a CSP admin exposes it causing loss to the firm, the preponderance principle says that whoever has 51% of responsibility pays 100% on the cost. But this is also affected by the unstated principle that loss of sensitive data makes the owner of that data (the organization) specifically liable for several actions and costs.

Enterprise Data Security for US Europe and Asia

The *Doctrine of Silver Platter* allows law enforcement to use anything presented to them without warrant or court order. The *Doctrine of Plain View* allows them to use any evidence within their presence and is visible to them without any restrictions.

Litigation Hold Notice is issued by law enforcement and courts to prevent possible destruction of evidence required for a case. *Spoilation* is the term for destruction of potential evidence intentionally or otherwise and can be a crime or used in another lawsuit.

When firm's personnel present forensic evidence in courts, they should be trained and certified in the tools they have used. Some jurisdictions (e.g. Texas and Michigan in the US) require them to be licensed. Data must be in original form and not tampered with; data modified for evidentiary purposes becomes inadmissible.

To prevent tampering of data, your security staff should be trained to use *write-blockers* when accessing electronic storage file for forensic purposes. If you are required to provide evidence in a legal proceeding, you need to consider the following best practices.

Snapshotting of virtual machines (VMs) aids in evidence collection for use in a court. Hypervisors provide only log data and are not sufficient.

- File hashes can serve as integrity checks for audit purposes to determine which systems are not configured to the baseline, and also to determine if files have been changed.
- Avoid using Type 2 hypervisors because they have a greater attack surface due to lack of operating system standards as compared to Type 1 hypervisors.
- Getting the CSP to ship hardware hosts (physical servers) for forensic analysis is not good due to potential exposure within the other firm.
- Using cloud storage is considered processing (manipulation, use, movement, masking, hard copy storage, or alteration of data) under most privacy frameworks and should be used carefully.
- Note that warrants, subpoenas, and similar court order are enforceable government requests, but an affidavit is not enforceable by itself.

The following section provides a representative list of nations that have defined and are enforcing laws pertaining to security of data that can identify living persons.

International Laws and Regulations for Data Security

This section provides an overview of the current laws by continent. Note that the laws and regulations are constantly being updated and you need to check the current status more thoroughly.

Europe

The General Data Protection Regulation (GDPR), which came into effect on May 25, 2018, is the successor to the European Union's Data Protection Directive [of] 1995 (Directive 95/46/EC).

Its legal form is a regulation (an order that must be executed) as opposed to a directive (a result to achieve, though the means to achieve aren't dictated).

- when the European Union (EU) enacts a regulation, it becomes national legislation in each EU member state, with member states having no opportunity to change it via national legislation.
- each member state requires national legislation to accompany the GDPR to confirm to its own legal framework and to select and implement permitted exemptions to GDPR.

The GDPR is the toughest privacy and security law in the world. Though it was drafted and passed by the European Union (EU), it imposes obligations onto organizations anywhere, so long as they target or collect data related to people in the EU. The regulation was put into effect on May 25, 2018. The GDPR will levy harsh fines against those who violate its privacy and security standards, with penalties reaching into the tens of millions of euros

Switzerland national privacy law is in line with GDPR and defines conditions for moving data.

United States

US does not have a nationwide privacy law. The FTC (Federal Trade Commission) an independent agency of the United States Government under the Department of Commerce was created to protect American consumers. It enforces the antitrust laws.

Enterprise Data Security for US Europe and Asia

The US Department of Commerce also sets policy for EU-US Privacy Shield program for voluntary compliance with EU Privacy Laws, and manages the EU-US Privacy Shield programs. The program is known as the EU-US Data Privacy Framework (EU-US DPF). The EU-US DPF was accepted by the European Union in March 2022 and the completion of its implementation was announced in July 2023. It replaced the Safe Harbor program.

To join the Data Privacy Framework, a company must self-certify to the Department of Commerce that it complies with the Data Privacy Framework Principles. A participating company's failure to comply with the Principles may violate Section 5 of the FTC Act's prohibition on unfair and deceptive acts. FTC has jurisdiction over most commercial entities and can issue and enforce regulations.

South America

Brazil does not have a national privacy law that complies with EU Directive / Privacy Laws.

Asia

Russia and China have a data localization law for PII data. Russia requires prior consent. China 2017 has a series of obligations including inspections.

Japan APPI – medical coverage for professionals. Requires private consent for data transfers. Japan mandates *Litigation Hold* (hold the data in case litigation is in progress)

Korea and Singapore also have litigation hold.

Some countries require permissions from local data protection commission for moving.

Australia

Australia Privacy Principles (privacy law) has 13 principles (2017) and applies to Australian residents. Australia and New Zealand make it very difficult to move data to CSPs.

International

ISO 27018 was the first international standard for privacy requirements for Cloud Service Providers (CSPs).

Geographical Location	Operational Laws	Description of Responsibility	Scope of
APEC – Asia Pacific Economic Cooperation	Asia-Pacific Economic Cooperation (APEC) Privacy Framework	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	The APEC region countries
Argentina	Personal Data Protection Act	With explicit intent of ensuring adherence and compliance with the EU Data Directive. Because of this, the EU treats entities in Argentina as if they were in the EU	Israeli citizens and residents
Australia and New Zealand	Australian Privacy Act 1988	Laws conform to EU policies. Regulates handling of personal information. It includes details regarding the collection, use, storage, disclosure, access to, and correction of personal information. Consist of fundamental privacy principles	Applicable to all citizens and residents of these countries.
Brazil	The Brazilian General Data Protection Act (in	Similar to EU Privacy Laws, applies if one of the following conditions is met:	LPGD applies to any person or

Enterprise Data Security for US Europe and Asia

	Portuguese, LGPD, Lei Geral de Proteção de Dados)	<ul style="list-style-type: none"> • The data processing takes place in Brazil. • The data processing is related to the offering or provision of goods or services in Brazil, or the processing of data of individuals in Brazil. • The personal data being processed was collected in Brazil 	organization, regardless of where they are based or where the data is located
EFTA - European Free Trade Association (Switzerland, Norway, Iceland, Lichtenstein)	Individual Laws in each country	Similar to EU Privacy Laws Swiss law provides stringent privacy protections, particularly for banking information.	Applicable to all citizens and residents of these countries.
Canada	The Personal Information Protection and Electronic Documents Act (PIPEDA)	Like EU rather than US. Includes such specifics on filing complaints as how and with whom they are filed, and the remedies and enforcements available to the government for redress of such grievances. The EU acknowledges PIPEDA as satisfactorily addressing the principles of the Data Directive and the Privacy Regulation.	Applicable to all citizens and residents of these countries
India	India Digital Personal Data Protection Act 2023 (DPDPA)	Based on Supreme Court landmark decision that Privacy is a fundamental right under the constitution. Similar to GDPR but does not	All residents of India and all Indian citizens worldwide

Enterprise Data Security for US Europe and Asia

		force companies to delete “right to be forgotten” data.	
Israel	Human Dignity and Liberty, 5752 -1992; the Protection of Privacy Law, 5741-1981	Like EU Privacy Laws it protects data collection, transfer, security, on-line privacy, breach notification, electronic marketing and enforcement.	Israeli citizens and residents
Japan	The Act on the Protection of Personal Information ("APPI")	Like EU Privacy Laws. The Personal Information Protection Commission ("PPC"), a central agency acts as a supervisory governmental organization on issues of privacy protection.	Japanese citizens and residents
South Korea	The Personal Information Protection Act (PIPA)	Like EU Privacy Laws, it governs the collection, use, and processing of personal data	South Korean citizens and residents
United States (USA)	HIPAA - Health Insurance Portability and Accountability Act	Protect patient records and data, known as electronic protected health information (ePHI). .	HITECH act in 2010, DHHS, has caused increasing number of audits
	GLBA – Gramm-Leach-Bliley Act (Financial Services Modernization Act of 1999)	Allowed banks and financial institutions to merge. Specifies the kinds of protections and controls that financial institutions are required to use for securing customers’ account information. Requires institutions to have a written information security plan (ISP). Requires that the bank ask the person, in writing, in hard copy, if the person wants to opt out of any data-sharing	

Enterprise Data Security for US Europe and Asia

	activity, at least once a year for every year that bank account stays open.
SOX – Sarbanes Oxley Act 2002	Executives can no longer hide behind an excuse of ignorance. They have become culpable for any financial wrongdoings including insider malicious actions perpetrated within their companies.
CCPA - California Consumer Privacy Act	Provides California residents with the right to : know what personal information is being collected about them, whether their personal information is sold or disclosed, and to whom, say no to the sale of personal information, access their personal information, and enjoy the same prices irrespective of their exercise of the law.
DORA	The Department of Regulatory Agencies (DORA) is Colorado’s umbrella regulatory agency, charged with managing licensing and registration for multiple professions and businesses, implementing balanced regulation for Colorado industries, and protecting consumers.

Table 2-1: International Regulatory Laws

Data Protection Laws by Countries and Regions

To fully understand the impact of the laws in each country, it is essential to review in detail the content of the laws and the implications of how they are applied. The following sections describes the specific data protection laws specific to each country or region. For this book, the coverage is on EU, India and the US.

EU General Data Protection Regulation (GDPR)

Once again, GDPR is the toughest privacy and security law in the world and it imposes obligations onto organizations anywhere, so long as they target or collect data related to people in the EU.

If you process the personal data of EU citizens or residents, or you offer goods or services to such people, then **the** GDPR applies to you even if you're not in the EU. The fines for violating the GDPR are very high. There are two tiers of penalties, which max out at €20 million or 4% of global revenue (whichever is higher), plus data subjects have the right to seek compensation for damages.

The GDPR defines an array of legal terms at length listed here briefly:

- **Personal Data** – information that relates to an individual who can be directly or indirectly identified. Names and email addresses are obviously personal data. Location information, ethnicity, gender, biometric data, religious beliefs, web cookies, and political opinions can also be personal data.
- **Data Processing** – Any action performed on data, whether automated or manual.
- **Data Subject** – The person whose data is processed. These are your customers or site visitors.
- **Data Controller** – the entity who decides why and how personal data will be processed.
- **Data Processor** -- A third party that processes personal data on behalf of a data controller.

The GDPR law is a comprehensive document with 99 articles that lays down the rights and responsibilities of the data subjects, the data controllers and the data processor. Thee 99 articles also define the define the key principles.

- **Lawfulness, fairness and transparency** — Processing must be lawful, fair, and transparent to the data subject.
- **Purpose limitation** — You must process data for the legitimate purposes specified explicitly to the data subject when you collected it.

Enterprise Data Security for US Europe and Asia

- **Data minimization** — You should collect and process only as much data as absolutely necessary for the purposes specified.
- **Accuracy** — You must keep personal data accurate and up to date.
- **Storage limitation** — You may only store personally identifying data for as long as necessary for the specified purpose.
- **Integrity and confidentiality** — Processing must be done in such a way as to ensure appropriate security, integrity, and confidentiality (e.g. by using encryption).
- **Accountability** — The data controller is responsible for being able to demonstrate GDPR compliance with all of these principles.

The following describes the categories for the 99 articles:

- General Provisions – Articles 1 -- 4
- Principles – Articles 5 -- 11
- Rights of the Data Subject – Articles 12 – 23
- Controller and Processor – Articles 24 – 43
- Transfers of Personal Data to Third Countries or International Organizations - - Articles 44 -- 50
- Independent Supervisory Authorities – Articles 51 – 59
- Cooperation and Consistency – Articles 60 – 76
- Remedies, Liability and Penalties – Articles 77 – 84
- Provision Relating to Specific Processing Situations – Articles 85 – 91
- Delegate Acts and Implementing Acts – Articles 92 – 93
- Final Provisions – Articles 94 -- 99

The rights of individuals are embodied in articles 12 through 23 as listed below:

Article #	Description
Article 12	Transparent information, communication and modalities for the exercise of the rights of the data subject.
Article 13	Controller shall provide to data subject: DPO contact and if intent to transfer to data to third country.
Article 14	Information provided to data subject where personal data has not been obtained from the data subject.

Enterprise Data Security for US Europe and Asia

Article 15	Right to complaint and to know processing purpose, storage period, third party disclosure, erasure, obtain copy.
Article 16	Right to obtain from the controller without undue delay the rectification/completion of inaccurate personal data.
Article 17	Right to be forgotten – erasure from all storage and attempt to erase from all third parties where shared.
Article 18	Right to obtain from the controller restriction of processing if data incorrect, processing unlawful, not needed.
Article 19	The controller shall communicate in writing any rectification or erasure of personal data or restriction.
Article 20	Portability -- right to receive personal data, which data subject has provided to a controller, in a structured, commonly used and machine-readable format and have the right to transmit those data to another controller.
Article 21	Right to object – to processing for direct marketing, or scientific/historical research or statistical purposes.
Article 22	Right not to be subject to decision based solely on automated processing which produces legal effects, profiling.
Article 23	Restrictions protecting rights based on national security, defence, public safety, criminal investigations.

Table 2-2: GDPR Articles for Rights of Individuals

So, if you conduct your business that includes EU jurisdictions or includes EU citizens, it is critical to study in detail the 99 articles of the GDPR law and ensure that your applications and processes are in compliance.

India Privacy Law

The Digital Personal Data Protection (DPDP) Act, 2023 was passed by both houses of the Indian Parliament and has received Presidential Assent. As specified in the

Enterprise Data Security for US Europe and Asia

Gazette of India “ An Act to provide for the processing of digital personal data in a manner that recognises both the right of individuals to protect their personal data and the need to process such personal data for lawful purposes and for matters connected therewith or incidental thereto.”

Citizen Rights: The Act grants "data principals" (the equivalent of "data subjects" under the GDPR) four key rights against "data fiduciaries," including companies and government entities, including: (1) the right to confirm whether and how their data is being used; (2) the ability to correct misleading or false data; (3) data portability rights; and (4) a "right to be forgotten," or the authority to restrict companies from using data they previously shared. Unlike the GDPR, the DPDP *does not* require companies to delete such data altogether.

Establishment of Regulatory Agency: The Act also calls for the establishment of a Data Protection Authority ("DPA"), now adopted by the Ministry of Information Technology, which would have the power to investigate, enjoin, and fine non-compliant entities. All data fiduciaries would be required to disclose data breaches to the DPA. The DPA would also have the authority to label certain entities as "Significant Data Fiduciaries" ("SDFs") based on the volume and sensitivity of data they process, thus subjecting such entities to additional transparency, auditing, and reporting requirements. SDFs would be required to appoint a Data Protection Officer to oversee compliance with the law; comply with annual independent audits of their processing of personal data; and conduct impact assessments for new technologies or large-scale profiling or use of personal data. The DPA also would have the authority to mandate that other data fiduciaries be subject to these requirements, even if they are not categorized as SDFs. Finally, all data fiduciaries would be required to notify the DPA in the event of a data breach, which would also have the authority to mandate individual notification.

Data Fiduciary Obligations: All data fiduciaries would be required to implement "appropriate security safeguards," including de-identification, encryption, and tools to prevent misuse, unauthorized access, modification, disclosure, or destruction of personal data. The bill establishes temporal limitations on the processing and retention of personal data, prohibiting data fiduciaries from retaining personal data longer than "reasonably necessary" to satisfy its intended purpose or comply with legal obligation, and requires fiduciaries to undertake periodic review to ensure they are not unnecessarily retaining personal information.

Laws in the United States of America

These include Administrative Law, Criminal Law, Civil Law and Case Law.

Enterprise Data Security for US Europe and Asia

Criminal Law: Under Criminal Law, the state (federal or a state) is the plaintiff and it can result in fines and incarceration. Generally, the issues of jurisdiction and subsequent prosecution are worked out in advance between law enforcement and court jurisdictional bodies.

Civil Law: Civil law is the body of law, statutes, and so on that deals with personal and community-based law such as marriage and divorce as opposed to criminal or military law. 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. Typical examples of civil law cases are disputes (in the United States) over property boundaries, mineral rights, and marital divorce. Civil cases are called lawsuits or litigation. Punitive measures for civil cases can include restitution of monetary damages or requirement to perform actions (usually in response to a breach-of-contract case, which we'll address later in this section), but not imprisonment.

Case Law: Collection of “precedents” resulting from judgments in Federal or Civil Courts that are considered as prior interpretations of the Criminal or Civil Laws and fully applicable until overturned.

data breach protection applies to a set of personal data that is narrower than that protected in the more general privacy protections.)

Perhaps the primary issue that businesses are contending with is that the law's requirements that could threaten established business models throughout the digital sector. For instance, companies that generate revenue from targeted advertising over internet platforms — such as Facebook, Twitter, and Google — must, as the law is currently written, allow California residents to delete their data or bring it with them to alternative service providers. This restriction could extend to internet service providers such as AT&T and Verizon, which collect broadband activity data (web browsing data) and could attempt to use it to generate behavioral profiles for marketing purposes.

The US does not have a single national law comparable to Europe's GDPR but has a combination of federal and state laws.

Federal and National Laws and Industry Standards:

The following lists the applicable laws and standards that apply nationwide.

- GLBA (Gramm-Leach-Bliley Act) requires financial institutions and companies that offer consumer financial products and advice such as banks, brokerages and insurance companies to explain their information sharing practices to their customers and to safeguard sensitive data.

Enterprise Data Security for US Europe and Asia

- HIPAA (Health Insurance Portability and Accountability Act) is a federal law that requires creation of national standards to protect sensitive patient health information from being disclosed without the patients' consent or knowledge. It specifies the formats and content structuring for storage and transmission of patient identification and clinical data.
- COPPA (Children's on-line privacy protection). Administered by the FTC, 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.
- PCI-DSS (Payment Card Industry Data Security Standard) is an industry consortium compliance agreement, not a law. It is designed to reduce payment card fraud by ensuring security controls around cardholder data. For example, organizations collecting and using credit information about consumers are not allowed to store the three-to-four-digit CSC (card security code) printed on the back (usually in the signature area) in their databases. It is used for validating that the cardholder is physically in possession of the credit card.

Common Law Privacy Act – based on consent orders/decrees by FTC is used by the FTC and the state AGs (Attorney Generals), is used to prevent unfair and deceptive trade practices. It provides that consumers have the right to know about the personal information a business collects about them and how it is used and shared. It also embodies the right to delete personal information collected from them (with some exceptions). You will commonly find it used in the form of disclosures by banks, brokerages and other institutions that collect your information on how they store and how they share your personal data.

Contract Law regulates the obligations established by agreement, whether express or implied, between private parties in the United States. The law of contracts varies from state to state, although there is a nationwide federal contract law applied to certain types of contracts, such as contracts entered into pursuant to Federal Reclamation Law.

State Laws

Several states have enacted their own laws, many of them derived from the basic principles of GDPR. The following states that have already enacted laws or are in the process of enacting them. At time of the publishing of this book, the California CCPA is the most stringent and adheres very closely to the basic tenets of GDPR.

- Massachusetts: Standards for protection of Personal Information of residents.

Enterprise Data Security for US Europe and Asia

- California: CCPA (California Consumer Privacy Act) – does not have strict 72 hour breach disclosure of GDPR.
- Colorado – DORA – Department of Regulatory Agencies also covers Colorado Privacy Act for consumer data protection.
- Washington: Less stringent but it is working through the legislatures.

CCPA - The California Consumer Privacy Act

The law notably establishes a broad definition of personal information including a consumer's personal identifiers, geolocation, biometric data, internet browsing history, psychometric data, and inferences a company might make about the consumer. The protections over this data are to be enforced by the state's attorney general, though consumers will maintain a private right of action should companies fail to maintain reasonable security practices, resulting in unauthorized access to the personal data. The following lists the rights embodied in CCPA:

- The right to know about the personal information a business collects about them and how it is used and shared;
- The right to delete personal information collected from them (with some exceptions);
- The right to opt-out of the sale or sharing of their personal information; and
- The right to non-discrimination for exercising their CCPA rights.
- The right to correct inaccurate personal information that a business has about them; and
- The right to limit the use and disclosure of sensitive personal information collected about them.

OECD – Organization for Economic Cooperation and Development

OECD applies to 36 member countries including US. It publishes model tax convention. Highlight include:

- Primary tax is country where investment originates.
- OECD privacy guidelines: Collection limitations, data quality (data is valid and accurate and allows data subject to correct it).
- Individual participation (limitations on collecting data with knowledge and permission of individual).

Enterprise Data Security for US Europe and Asia

- Purpose specification (Clearly state explicit purpose for collecting data) and it follows the use collection limitation principle (restrict data to info necessary for an allowed transaction),
- Security safeguards (protect data against unauthorized access and modification).
- Openness (allow the data subject to access the information)
- Accountability

OECD influences the European Union in their law-making,

Standards and Compliance

Data Security Standards are guidelines or criteria that organizations must follow to protect sensitive and confidential information by government and corporate entities who collect and use personal data of employees and customers. These standards are designed to prevent unauthorized access, use disclosure, disruption, modification, or destruction of data.

A variety of standards are in active use driven by national governments and international consortiums. The following sections describe the major standards currently in use.

Note that Custom IAM solutions can become weak and are not an industry best practice even though you may follow industry standards. If you design your own IAM solution for your organization, you need to know these standards to ensure compliance according to your geographic location and business activities.

Global Standards

Global standards include:

- CSA CCM
- COBIT5
- NIST RMF
- ISO/IEC 27017
- HIPAA
- EU GDPR
- And other country specific standards.

In addition, industries such as personal credit cards define their own standard – PCI-DSS.

Enterprise Data Security for US Europe and Asia

The following sections provide a description of the major international standards that your organization needs to be fully aware of as you conduct your business..

US Standards

The US Federal government established the security practices for federal computer systems and requires federal agencies to conduct periodic assessments of the risk and magnitude of harm that could result from the unauthorized access, use, disclosure, disruption, modification, or destruction of information and information systems that support the operations and assets of the agency and address information security throughout the life cycle of each agency.

The standards used in the US include US standards as well as European and International Standards. These include the following:

- FIPS 140-2
- FedRAMP
- NIST
- ISO/IEC
- ISO
- ENISA
- PCI-DSS.

We will describe each of these in the following sections.

FISMA – Federal Information Security Modernization Act

FISMA was originally passed as the Federal Information Security Management Act in 2002 as part of the E-Government Act. Today, FISMA refers to the Federal Information Security Modernization Act of 2014, passed in response to the increasing amount of cyber attacks on the federal government.

The Federal Information Security Modernization Act 2014 codifies the Department of Homeland Security's role in administering the implementation of information security policies for federal Executive Branch civilian agencies, overseeing agencies' compliance with those policies, and assisting OMB in developing those policies. It defines the framework of guidelines and security standards to protect government information and operations. FISMA requires all federal agencies to develop, document, and implement agency-wide information security programs.

FIPS 140-2 -- Approve Cryptographic H/W and S/W Modules

This Federal Information Processing Standard (140-2) specifies the security requirements that will be satisfied by a cryptographic module, providing four increasing, qualitative levels intended to cover a wide range of potential applications and environments. It consists of the following four levels:

- Level 1 is the lowest – Basic security requirements with production grade equipment and coatings and at least one tested encryption algorithm.
- Level 2 - Raises the level to tamper evidence feature - features that show tampering- seals, and should run on an Operating System that has been approved by Common Criteria at EAL2.
- Level 3 – Requires strong enclosures and tamper resistant circuitry which includes attack detection such as tamper/detection response circuitry, identity based authentication.
- Level 4 – Highest level - physical envelope of protection around modules, high level of detection and response, useful for operation in physically unprotected areas. Used for cryptographic modules. Includes erasure of device in case of an attack.

Controlled Unclassified Information (CUI) level of information may be processed by cryptographic modules certified for FIPS 140-2 Level 1 but not at the Top Secret, Secret, or SCI (Sensitive classified information) level. Note that Federal classifications are Unclassified, Public Trust, Confidential, Secret and Top Secret at increasing levels of security.

Areas covered by FIPS 140-2 include:

- Cryptographic module specification and parts
- Roles
- Services
- Authentication

It provides for a finite state model, physical security, operational environment, cryptographic key management, prevention of electromagnetic interference, and attack mitigation.

Cryptographic modules for compliance with FIPS 140-2 can be provided by private sector or public sector for use by the US federal government. The module is provided along with a compliance certificate. The module vendors pay for cryptographic module certification. FIPS 140-2 Certification is not provided by the US government. FIPS compliance is also recognized around the world as one of the

Enterprise Data Security for US Europe and Asia

best ways to ensure cryptographic modules are secure. FIPS 140-2 is not a regulatory framework for sensitive data.

FedRAMP (Federal Risk and Authorization Management Program)

Derived from NIST 800-53, it is adopted for the cloud. It is a US federal government-wide program that provides a standardized approach for security assessment, authorization and continuous monitoring for cloud products and services.

FedRAMP was created by a Memorandum by OMB (Office of Management and Budget) GSA (General Services Administration) established the FedRAMP Program Management Office (PMO). FedRAMP resides within the GSA.

FedRAMP consists of two primary entities that authorize services:

- Joint Authorization Board (JAB) which includes the CIOs (Chief Information Officers) from the Department of Defense, Department of Homeland Security, and General Service Administration. JAB serves as the primary governance and decision -making body for FedRAMP.
- Program Management Office (PMO) assists and guides agencies through the FedRAMP Authorization process.

3PAOs (Third Party Assessment Organizations) play a critical role as independent assessors checking the cloud provider's security implementation and risk posture.

CIA, US Post Office, Department of Defense, and Department of Homeland Security, and other federal agencies are required to use FedRAMP. They are required to ensure that federal data is maintained within US boundaries

National Institute of Standards and Technology (NIST)

FISMA established the NIST Special Publications 800-53, 800-59, and 800-60. Additional security guidance documents are being developed in support of the project including NIST Special Publications 800-37, 800-39, 800-171, 800-53A and NIST Interagency Report 8011.

NIST maintains a number of standards that define the security requirements for on premises (data centers) and cloud environments. These standards are also required to be followed by all entities doing business with the federal government as well as all financial institutions that collect and manage employee and customer personal sensitive data. Please refer to <https://www.nist.gov/> for a detailed description of all NIST standards.

The following is a list of the NIST information technology related standards.

Enterprise Data Security for US Europe and Asia

- *NIST 500 -292* – NIST Cloud Computing Reference Architecture (RA) and Taxonomy.
- *NIST 500-294* - Cloud computing security reference architecture – includes conceptual models, reference architecture, and controls framework.
- *NIST 800-30* – Risk management guide for IT systems and data centers – it defines vulnerabilities/threats.
- *NIST 800-37* - Is a guide for applying a Risk Management Framework (RMF). It defines automation of controls wherever possible, with a focus on continual improvement and near real time risk management, and limiting risk level at the process level and management level are part of the framework. Cost metrics and perceived threats are not a part of the framework.
- *NIST 800-39* -- Offers a risk management framework that is general enough to be applicable to both the public and private sectors
- *NIST 800-53* – Security and Privacy Controls for Federal Information Systems and Organizations. Firms choose this because standards are in public domain and available for free and have been widely accepted within US for use by all financial institutions and large corporations that have a business interest in supplying software to the US government or are regulated by the government consumer protection laws. There is no official international acceptance of NIST. However, NIST compliance is viewed as a very positive basis for security compliance.
- *NIST 800-57* – Provides cryptographic key management guidance. It consists of three parts – part 2 is the most important.
- *NIST 800-61* – Computer security Incident handling/response lifecycle
- *NIST 800-63* – Suite provides technical requirements for federal agencies implementing digital identity services.
- *NIST 800-64* - guideline is to assist agencies in building security into their IT development processes.
- *NIST 800-92* -- Guide to Computer Security Log Management," establishes guidelines and recommendations for securing and managing sensitive log data.
- *NIST 800-122* – Protecting the Confidentiality of Personally Identifiable Information
- *NIST 800-145* - Working definition for cloud computing.
- *NIST 800-160* Vol 1 – Systems Security Engineering: Considerations for multi-disciplinary approach in the Engineering of Trustworthy Secure Systems.
- *NIST 800-171* -- safeguarding sensitive information on federal contractors' IT systems and networks

NIST standards have very extensive coverage and need to be studied in depth to ensure full compliance.

International Standards Driving Data Security

ISO 270001 and ISO 270002 establish the requirements and procedures for creating an information security management system (ISMS). This serves both, as an operational controls regimen as well as an important audit and compliance guideline. The vocabulary and an overview of the ISO 270001 and ISO 270002 are provided in the ISO27000 standard.

A related standard, the ISO27701, is the international standard for privacy information management and its objective is to protect private information assets and to demonstrate compliance with privacy and data protection regulations.

ISO/IEC and ISO

ISO/IEC standards that apply to information technology include the following:

- *ISO/IEC 17788* — Documents provide an overview of cloud computing along with a set of terms and definitions.
- *ISO/IEC 17789* — Specifies the cloud computing reference architecture including cloud roles, activities, functional components, diagrams, and their relationships.
- *ISO/IEC 20000* -- the international Service Management standard.
- *ISO/IEC 27001* — Describes requirements for certification of ISMS (Information Security Management System), the organization's entire security program.
- *ISO/IEC 27008* – Guidance for auditors on ISMS.
- *ISO/IEC 27014* – IT – Guidance on principles for Security techniques for governance of information security by which organizations can evaluate, direct, monitor, and communicate the information security related activities within the org.
- *ISO/IEC 27017* – Information Technology – Code of Practice for Security techniques based on ISO/IEC 27002 for cloud services.
- *ISO/IEC 27034* – Offers guidance on information security to those specifying, designing and programming or procuring, implementing, and using application systems. Defines concepts frameworks and processes. Defines Application Security Management Process (ASMP) that assess security risks and lays out ANFs (Application Normative Frameworks). It also Lays out the ONF (Organization Normative Framework) for all components/applications

Enterprise Data Security for US Europe and Asia

defining business, regulatory, technical contexts, specifications, roles, processes, and Application Security Control (ASC) library. Multiple ANFs can be used with one ONF for each application.

- *ISO/IEC 27035* Preparation- Defines the steps: 1) Detection and Analysis, 2) Containment, Eradication, Recovery, 3) Post Mortem, and 4) Preparation which handles communication, hardware, documentation (lists, baselines, network diagrams). Detection Analysis covers alerts (end-point protection, network, host monitoring, SIEM). Eradication, Containment and Recovery takes system off-line & cleans up compromised devices; Post Mortem is a review to determine how to do better. For Cloud Preparation, it requires defining SLAs (Service Level Agreements), Governance with the CSP (Cloud Service Provider), how to get data specific to SaaS/PaaS, cloud jump kits for remote cloud access, collecting logs from clouds, and data server access. Note that for Detection, approaches used include logs from APIs (based on use of immutable servers), application stacks maps, threat modeling, table-top exercise, monitoring cloud Management Plane. For Detection Analysis approaches use of logs include auto-scaling and cloud management activities, flow records, packet capture, instrument applications for logs, CSP Chain of Custody, automated forensics, does CSP support processing machine, check network isolation from network flows, raw data logs, management plane logs. For Containment, Eradication, Recovery ensure that the management plane meta structure is free of contamination. In IaaS create new infrastructure to allow time to understand attacker in the compromised one after containment.
- *ISO/IEC 27040* — provides security guidance for storage systems and ecosystems and information risks associated with the confidentiality, integrity and availability of information on various data storage technologies.
- *ISO/IEC 27050* – Electronic Discovery – Part 1 – Overview and Concepts
- *ISO/IEC 29100* – For PI (personal Information) Security -- Information technology – security techniques – privacy framework. Is applicable to natural persons and orgs involved in specifying, procuring, architecting, designing, developing, testing, maintaining, administering, and operating information and communications systems for PI.
- *ISO/IEC 38500* – IT - Governance of IT for the org – Provides a framework for effective governance of IT to assist those at the highest level of orgs to understand and fulfill their legal, regulatory and ethical obligations.

IEC 642443 is an international series of standards that address cybersecurity for operational technology in automation and control systems. It is mentioned here for awareness although it does not directly apply to data systems security for banking and financial institutions or large corporations.

ISO 270001 and ISO 270002 Standards

- *ISO 27001* – Dictates creation of an organizational information security management system (ISMS). Details what orgs can be *certified* for ISMS. This is preferred due to international acceptance but it is not easy to implement. Furthermore, it does not favor cloud based, PC or open source technologies. It has 35 control objectives, 114 controls across 14 domains including: InfoSec policies, organization of information security, HR security, asset management, access control, cryptographic, physical environment security, operations security, system acquisition, supplier relationship, information security incident management, information security business continuity management, and compliance.
- *ISO 27002* - Implementation guidance for ISO 27001 - electing, implementing, and managing information security controls mapped to an ISMS. An *audit* against it shows if the org has adequate controls to meet ISO 27001. Note that SAS 70 (outdated) and SSAE 16 are audit standards for services providers and include some review of audit controls but not a cohesive program.
- *ISO 27018* – Privacy requirements for CSPs
- *ISO 27037:2012* – Evidence preservation

ISO 31000 And Related Standards

- ISO 31000:2018 – Risk Management - Principles and Guidelines, a framework for managing risk. PDCA (Plan, Do, Check, Act).
- ISO 31010:2019 — Risk management – Risk assessment techniques – provides guidance on the selection and application of techniques for assessing risk in a wide range of situations.

Note that there are equivalent ISO/IEC standards to the ISO standards. ISO/IEC 27001 and ISO 27001 are similar targeted at an ISMS. ISO and ISO/IEC are Truly international standards but NIST 800-37 (Risk) and SOX (Financial) are US, and GDPR is European

ENISA – European Union Agency for Network and Information Security

ENISA does not spell out the types of assurance due to the large variety of providers and regulators. AS per ENISA, cloud risk assessment should provide means for customers to assess risks due to cloud migration, compare offerings from

CSPs, reduce assurance burdens on CSPs but NOT: reduce the risk of regulatory compliance.

ENISA includes programmatic management as a defining trait of cloud computing even specifying through “WS API”. This is not included in the ISC(2) or NIST definition which includes metered service, shared resources and scalability.

Industry Standards

Industry standards are voluntary agreements that establish the security standards for specific industries and activities.

Payment Card Industry Data Security Standard (PCI-DSS)

The PCI-DSS is an information security standard designed to reduce payment card fraud. Following the standard requires increasing security controls around cardholder data. The standard covers processing digital transactions and payments using credit or debit cards, storage of credit card data, transmission of cardholder information to another entity and controlling who has contracted protected cardholder data.

The following lists the key highlights of the PCI-DSS standard:

- For PCI-DSS an organization can never store Card Verification Code (CVC) number for any duration. You can store only consumer personal data, credit card number, address, email, and phone number, date of birth can be stored.
- PCI-DSS does not cover companies that offer credit card debt repayment counseling.
- PCI-DSS covers banks issuing credit cards, retailers accepting credit card payments, businesses processing credit card payment on behalf of themselves of a retailer.
- A PCI-DSS violation can result in fines, suspension of credit card processing privileges, increased audit frequency. Jail time can be awarded only on a related crime. Fines can range for \$5,000 to \$10,000 per month.
- PCI-DSS merchant levels 1 through 4 are based on the number of transactions over the course of a year, and not on dollar value of transactions, location, or the combination of dollar value and number of actions:
 - 1. — > 6 million transactions/year
 - 2. — 1-6 M transactions/year
 - 3. — 20K – 1M transactions/year
 - 4. — <20K transactions/year

Enterprise Data Security for US Europe and Asia

- The standard has over 200 controls. It has a mandatory breach reporting requirement for regulated sensitive data (PI).
- Tokenization or encryption are required for storing any credit card holder information.

All merchants are required maintain a firewall to protect card data, avoid vendor supplied defaults, protect card data, update antivirus software, maintain secure systems, restrict access on the need to know basis, assign unique ID to each person accessing cardholder data, restrict physical access to storage units, track and monitor all accesses to cardholder data, actively use test systems to monitor security posture, and maintain policies and governance of these policies.

Audits and Risk

Audits are defined as the on-site verification of the security management of systems storing and using confidential and highly confidential sensitive personal data. Verification activities can include inspection or examination of processes, policies, IT infrastructure in use on premises and in clouds, access controls for systems, and protection of data storage. The audit can apply to the entire organization or to a subset of the organization.

The following describes the highlights of a typical audit in the US, such as those conducted by OCC (Office of the Comptroller of the Currency) which is the primary regulator for banks, or FFIEC (Federal Financial Institution Examination Council) responsible for inspecting federally supervised financial institutions.

- The audit starts with a scope statement described in an enforcement request statement known as “Matter Requiring Attention” (MRA.). It describes the audit scope and the evidence production timeline.
- Audit scope statement for a project may include deliverables, statement of purpose, classification, time spent in production, business areas to be reviewed but it does not include costs. Audits are allowed in the actual production environment but there are limitations on destructive techniques, and prohibition on direct access to production.
- Note that auditors can interview personnel and have an ethical, not legal, responsibility to report on illicit activities discovered during the audit.
- Lack of physical access to ephemeral hosts in clouds makes understanding security controls more difficult and so cloud data center audits are typically less trustworthy. So, the auditor has to rely on whatever the cloud service

Enterprise Data Security for US Europe and Asia

provider to disclose and business being audited chooses to provide. They are also considered less trustworthy because they frequently rely on third parties,

- A specified baseline configuration built to defined standards can be used to demonstrate that all VMs include essential controls – very useful for audits. The baseline can be checked to confirm that necessary security controls are in place.
- Pass through Audits, also called inherited audits, can occur when the CSP is certified for PCI-DSS, SOC1, SOC2, HIPAA, and best practices CCM and GDPR. This is called compliance inheritance.
- CSP's assets are not within the scope of custom audits but everything they build is within the audit assessment scope.
- Assessments are Point-in-time compliance. Internal audits should be continuous, not point-in time.

Audit Plan

Audit Plans include lessons learned, defining objectives and scope; but does not include remediation, fixing issues that caused the audit.

Audit Scoping Statement

- An Audit Scoping Statement should include constraints on limitation of destructive techniques, mandate of a particular time zone review, prohibition of access to production environment; but it should not prohibit all personnel interviews.
- The statement should include areas to be audited, personnel who can be interviewed, governance practice, risk management controls effectiveness, financial and management reporting, and operations effectiveness.
- Scope limitation is a restriction on the applicability of an auditor's report that may arise from inability to obtain sufficient appropriate evidence, e.g in financial statements, or the ephemeral assets used in the CSP.

Risk Management

For a business, risk management is a very important consideration to ensure proper security controls to prevent data exposure. It is even more important when breach of sensitive data or exposure of employee or customer personal information can result in significant penalties and loss of reputation. Risk assessment is the responsibility of every organization.

Enterprise Data Security for US Europe and Asia

Risk appetite is the extent of risk in financial and loss of reputation terms that an organization is willing to accept in the pursuit of its business, financial and strategic perspectives. Risk appetite reflects the risk management posture of the Board of Directors that drives the security controls and the security culture of the organization.

Risk management is qualitative and quantitative. Use qualitative assessments with new, untrained, or insufficient data. The documentation for risk assessments should include: the information security program details including policies, standards, procedures, guidelines and baselines; technical information including network diagrams, application documents, BCDR plan, incident response, etc.

The following describes the key considerations for risk management:

- Risk should always be considered from a business perspective often balanced by opportunity and not be compromised for profit, performance or costs. In the case of cloud-based environments, risk is a shared responsibility with the CSP, but the organization cannot abdicate its own responsibility of ensuring proper identity and access management controls.
- Extending the network perimeter to the public clouds enhances opportunities for collaboration mostly by giving external partners and customers limited access to the data owner's data in the cloud. While this is risky, the comparable risk on legacy on-premises data is even higher in sending data out to partners or inviting partners into the protected corporate data centers.
- Some on-premises risks cannot be transferred to the cloud such as liability for data exposure. Your organization is ultimately responsible for liability from any data exposure. Some risks can be mitigated even in the cloud, but these must be clearly understood and negotiated when the CSP contract is established.
- While on premises data centers may experience resource exhaustion, it is conceivable that in the cloud also, a business continuity and disaster recovery threat can emerge and risk assessment must be performed for a contingency such as, all clients of the CSP demand extra resources at the same time or if the data centers experience a catastrophic failure.

Certification Frameworks

Certification frameworks are essential for an organization to not only self-assess its policies and procedures but also to convince its customers that their data is safe within the organization, and required data loss prevention policies and procedures

Enterprise Data Security for US Europe and Asia

are being enforced. The following sections describe some of the certifications that should be considered for compliance for data protection.

CSA STAR – Security, Trust Assurance, and Risk

The Cloud Security Alliance (CSA) STAR (Security, trust, Assurance, and Risk) Registry program covers a full range of aspects for cloud security. Is an assurance program based on CSA Cloud Capability Matrix (CCM), Consensus Assessments and Initiative Questionnaire (CAIQ) and CSA's Code of Conduct for GDPR Compliance for CSPs and their cloud customers.

The program is assessed at the following three levels.

- Level 1 -- Self - Assessment can choose one or both of Self-assessment (security) and GDPR CoC (privacy) and also continuous self-assessment (no privacy). CSA STAR provides listings to solution providers who have integrated CAIQ and CCM and other GRC stack components; it helps GRC monitoring and reporting.
- Level 2 – Third Party Certification – You can choose 3rd party certification (security) and GDPR CoC (Privacy) and also continuous self-assessment (no privacy). STAR attestation is based on Type 1 or Type 2 SOC attestations supplemented by CCM. It serves as the standard for SOC 2 and SOC 3 reporting and provides recognition with the AICPA logo. Allows comparison of orgs and shows how mature their processes are and what areas they need to consider improving.
- Level 3 (highest transparency and monitoring) – Continuous Auditing and monitoring enables automation of the current security practices of CSPs who publish their security practices according to CSA formatting and specifications. Cloud customers and tool vendors can retrieve this information in many contexts. This has no GDPR CoC for privacy. It is an improvement over the point-in-time certification for trust and transparency.

Each level also has a continuous auditing option that allows increasing the transparency. CSPs can submit CAIQ (295 questions) also CAIQ lite with 135 questions) and or CCM reports, CSA encourages both.

CCM (Cloud Capability Matrix)

The CCM is a meta framework of cloud specific security controls, mapped to leading standards, best practices and regulations. It provides structure, detail and clarity for a comprehensive assessment.

Enterprise Data Security for US Europe and Asia

It is a Security Control Framework that provides mapping and cross-relationships with the main industry accepted security standards, regulations and control frameworks including ISO 27001/27002, ISACA's COBIT and PCI-DSS. The 136 controls specified by it are arranged in the following domains:

1. Application and Interface Security
2. Audit Assurance and compliance
3. Business Continuity Management and Operational Resilience
4. Change Control and Configuration Management
5. Data Security and Information Lifecycle Management
6. Data Center Security
7. Encryption and Key Management
8. Governance and Risk Management
9. Human Resources
10. Identity and Access Management
11. Infrastructure and Virtualization Security
12. Interoperability and Portability
13. Mobile Security
14. Security Incident Management, e-Discovery and Cloud
15. Supply Chain Management, Transparency and Accountability
16. Threat and Vulnerability Management

These are controls are applied across (the other axis in the matrix) :

- Fulfillments towards data subjects (notice, consent, rights).
- Fulfillments towards data protection authority (notification for breach, DPA checking for specific privacy risks, authorization for specific processing).
- Organizational contractual measures (controller/processor privacy agreement, data transfer agreement, training for data processing personnel).
- Technical procedural measures (security measures, data breach identification/management, data retention for specific processing).

The CCM addresses security architecture constructs, application security. and physical security but it does not address business drivers.

CSA CCM lists security controls from the laws: DMCS (Digital Millenium Copyrights Act), HIPAA (Health Information Portability and Accountability Act), and FERPA (Family Education Rights and Privacy Act).

SSAE 18

Statement on Standards for Attestation Engagements # 18 (SSAE 18) is a series of enhancements aimed to increase the usefulness and quality of Security Operations Center (SOC) reports. It supersedes SSAE 16.

SSAE SOC Reports

SSAE is designed to audit public entities or privately held companies. The reports are of the following types:

- SOC1 reports are used to assess reliability of financial reporting mechanisms and can be presented to an investor or major client and also auditors but not to a potential client.
- SOC2 Type 1 report describes the security controls at a point-in-time but not how they function and, yes, whether the design is suitable to meet relevant trust principles. These audits are considered for restricted use vs. being for a more broad audience (including internal, external, regulatory) such as SOC2 and SOC 3,
- SOC2 Type 2 Reports describes operational effectiveness and is used to assess the implementation and effectiveness of security controls within the org. It must always report on security over and above confidentiality, processing integrity and privacy. Most commonly seen for public CSPs

The Security principles of SOC2 reports consist of 5 categories (SAPCoP):

Security – refers to protection of systems against unauthorized access (access Controls)

Availability — refers to the accessibility of the system, products or services as stipulated by the contract or SLA. The minimum acceptable level is set by both parties.

Processing Integrity — addresses whether or not a system achieves its purpose. Does not necessarily imply data integrity.

Confidentiality – data is considered confidential if its access and disclosure is restricted to a specified set of persons or orgs. Encryption is an important control for protecting data at rest and in motion.

Privacy – system's collection, use, retention, disclosure. And disposal of sensitive PII data is in conformity with the org's privacy notice.

SOC3 Report is not a report — it **is an attestation** that security controls are adequate. Most useful for a quick analysis to short-list acceptable CSPs

COBIT

COBIT is the acronym for Control Objectives for Information and Related Technologies. The COBIT framework was created by ISACA (Information Systems Audit and Control Association) to bridge the crucial gap between technical issues, business risks and control requirements. COBIT is a **comprehensive I&T (Information & Technology) Governance and Management framework**. COBIT is not a full description of the whole IT environment of the firm and **it is not a framework for organized business processes**. It does **not make or proscribe IT related decisions**.

COBIT is designed to provide important information about the state of the IT environment to **internal (Boards of Directors, executives, business managers, IT managers, assurance providers and risk managers), and external (regulators, business partners, IT vendors) entities**.

COBIT 19 updated COBIT 5 and included the following enhancements:

- Flexibility and Openness.
- Currency and relevance.
- Prescriptive application, Performance management of IT integrated into conceptual model.

COBIT describes six *Governance* system principles which ensure that stakeholder needs conditions and options are evaluated to determine balanced, agreed on direction and objectives set through prioritization and decision making (used for measuring compliance). COBIT defines the components to build and sustain a governance system : processes, organization structures, culture/behavior, skills, and infrastructure. The six principles of the Governance System include:

1. Provide Stakeholder Value,
2. Holistic Approach,
3. Dynamic Governance system,
4. Governance Distinct from Management,
5. Tailored to Enterprise need,
6. End-to-end Governance system.

The three Governance Framework Principles include:

Enterprise Data Security for US Europe and Asia

1. Based on Conceptual Model,
2. Open and Flexible,
3. Aligned to Major Standards.

Governance Objectives are: EDM – Evaluate, Direct, Monitor

Management Plan

The Management plans, builds, runs and monitors activities in alignment with directions set by the governance body. The Management objectives include:

1. APO – Align, Plan and Organize (overall org strategy),
2. BAI – Build, Acquire, and Implement,
3. DSS – Deliver, Service, Support,
4. MEA – Monitor, Evaluate, Assess

Risk Optimization

Risk optimization addresses business risk associated with the use, ownership operations of I&T within an enterprise.

Resource Optimization

Resource Optimization – ensures appropriate capabilities are in place to execute the strategic plan.

ITIL and ITSM

Information Technology Infrastructure Library (ITIL) is a framework of best practices for delivering IT services. ITIL's systematic approach to IT Service Management (ITSM) can help businesses manage risk, strengthen customer relations, and build an IT environment geared for growth, scale, and change.

ITIL is focused on customer needs for services rather than IT systems, and stresses continual improvement. It describes processes, tasks, and checklists which can be applied to establish a baseline from which the organization can plan, implement and measure its IT infrastructure. It is used to demonstrate compliance and measure improvement.

There is no formal third-party compliance assessment or certification.

ITIL 4, released in February 2019 is based on ITIL 2011. It underpins ISO/IEC 20000 the international Service Management standard. As of 2021, this is the latest version of ITIL in use. It is easier for companies to align ITIL version 4 with Agile,

Enterprise Data Security for US Europe and Asia

DevOps and Lean work methods because it is more inclusive for modern digital environment.

ITIL consists of 5 books representing stages:

1. Service Strategy,
2. Service Design,
3. Service Transition,
4. Service Operation,
5. Continual Service Improvement

ITIL describes its Service Value Chain as: Plan, Improve, Engage, Design/Transition, Obtain/Build, Deliver and Support

Its guiding principles are: Focus on value, start where you are, progress iteratively, Collaborate and promote visibility, think and work holistically, keep it a simple/practical, optimize and automate

Change Management is described as the process designed to understand and minimize risks while making IT changes. Change request is via RFC (Request for Change) ticket.

From a data security perspective, it is important to follow the practices recommended by ITIL for any financial business or a large corporation storing and managing sensitive employee and/or customer personal information. The following lists the best practices;

General Management Practices:

1. Strategy Management
2. Portfolio Management
3. Architecture Management
4. **Service Financial Management**
5. Workforce and Talent Management
6. Continual Improvement – operational day-to-day part, for service value chain.
7. Measurement Reporting
8. Risk Management (linked to ISO/IEC 31000:2018)
9. Information Security Management
10. Knowledge Management
11. Organizational Change Management
12. **Project Management**
13. Relationship Management
14. Supplier Management

Service Management Practices

Enterprise Data Security for US Europe and Asia

- Business Analysis
- Service Design
- Service Level Management
- Availability Management
- Capacity and Performance Management
- Service Continuity Management
- Monitoring and Event Management
- Service Desk
- Incident Management
- Service Request Management (includes Fulfillment and Access Management)
- Problem Management
- Release Management (waterfall and DevOps)
- Change Control (formerly Change Management)
- Service Validation and Testing
- Service Configuration Management
- IT Asset Management

The three technical management practices are:

- Deployment Management (aligned with Release Management and change control)
- Infrastructure and Platform Management
- Software Development and Management

Service Transitions: List of ITIL Processes:

- Transition Planning and Support
- Change Management
- Service Asset and Configuration Management
- Release and Deployment management
- Service Validation and Testing
- Change Evaluation
- Knowledge management

Other best practices that need to be managed to ensure security of data include the following.

Change Management – aims to ensure standardized methods and procedures for efficient handling of changes. Goal: minimal disruption of service, reduction in back-out activities, economic use of resources. Types of changes: 1) Standard, 2) Normal 3) Urgent/Emergency

Enterprise Data Security for US Europe and Asia

- RFC (Request for Change) –form to record details of requested change
- FSC (Forward Schedule of Change – contains details of all forthcoming changes
- CS (Change *Schedule*) : Details of forthcoming changes

Configuration management is an ongoing iterative process of tracking all deployed configured resources in data center and cloud; List:

- Identification
- Planning
- Change control
- Change management
- Release management
- Maintenance

Event Management: - Event is something not functioning correctly leading to an incident being logged. Event management detects issues, monitoring checks status even if no problem. An event can lead to an incident, problem or change.

Incident Management – An unplanned interruption to IT Service or reduction in quality. Aim is to restore service ASAP and minimize adverse impact on business operations.

Problem Management - for an incident with unknown underlying element, identified as a result of multiple incidents. Different from Incident Management (which has the goal of returning to service at the earliest. Aimed at reducing the number and severity of incidents. Process used include: Trend analysis, targeting support action, and providing info to org. Error Control process iteratively diagnoses known errors. Problem Control includes problem identification and recording, classification, investigation, and diagnosis.

Chapter 3 – Cloud Impact on Data Security

The cloud or the hybrid cloud model facilitates building up and enhancing infrastructure much faster with associate cost savings provided by the pay-as-you-use utility model of the clouds. This has encouraged rapid migration of applications from on premises to public clouds such as AWS, GCP and Azure.

This extension of the network perimeter to include the cloud environment creates an added challenge due to the increase in the attack surface, and access to data stored in the cloud which is not completely under the control of the organization. The application and data owner needs to ensure additional security controls for data moved to, stored in and used in the cloud environments.

If your organization migrates applications to the cloud, you need a secure way not only for you but also your business partners and customers who you allow access to your cloud environment to view their data. Cloud security ensures your data and applications are protected from unauthorized access and use but are readily available to authorized users. With properly designed security you will have a reliable method to access your cloud applications and information, helping you quickly act on any potential security issues.

There are several steps involved in selecting and securely building safe applications in the cloud. These start with selecting the cloud service provider (CSP) and adhering to the best practices for cloud application development and deployment.

Understanding the Cloud Environments

Most organizations have the four common types of environments:

1. DEV – for software development and building the application for deployment.

Enterprise Data Security for US Europe and Asia

2. QA – for testing the application.
3. NON-PROD – testing the application in almost production like environment.
4. PROD – production environment running the commercial applications.

Data security on-premises and in the cloud follows the same general guidelines. Only the production environments (PROD) should be allowed to have PII data. Consequently, proper security controls must be implemented for PROD to ensure that the least privilege principle is followed and access to PROD is granted on need basis. Furthermore, access to PROD should require privilege escalation assuring the request to access PROD is granted by a person who oversees the work of the requester.

Sandbox

The cloud environments typically have another common environment called SANDBOX used for all development activity in the cloud. The Sandbox should never have any production data such as employee or customer personal information. A hosted cloud environment isolated as a sandbox is a good test bed for an application developed on-premises.

Dynamic Optimization and Elasticity

The cloud environments can be optimized dynamically using technology for resources and cloud services using telemetry, algorithms, service and resource analytics, and policies. This analysis can drive actions to dynamically adjust resources in the environments to reduce waste, cost and risk exposure, while simultaneously improving service levels.

Elasticity is defined as the degree to which a system is able to adapt to workload changes by provisioning and de-provisioning resources in an automatic manner, such that each point-in-time the resources match the current demand as closely as possible. Elasticity in the cloud, the ability to increase or decrease the resources a cloud-based application uses, allows you to scale computer processing, memory, and storage capacity to meet changing demands.

Multi-tenancy

Multitenancy refers to a single cloud instance and infrastructure built to enable multiple cloud customers (tenants) to efficiently share scalable computing resources in a public or private cloud. Appropriate data classifications of data elements become very important and presents a challenge for a move to the cloud. The potential for bad actors, who as legitimate tenants in the public cloud environments

sharing the same hosts, can increase the risks for data exposure. Data loss protection is of utmost importance in the cloud for this reason.

Disaster Recovery

Cloud deployment model for disaster recovery for PROD is typically HYBRID for on-premises and cloud. Disaster recovery should never be to a community cloud or to a private cloud in the on-premises data centers.

Location of Cloud Data Centers

Location jurisdiction is probably a lower importance issue for CSP SLA compared to Bandwidth, Availability and storage. However, multiple jurisdictions pose a serious complication for an organization for compliance from an international perspective. Other considerations such as different certifications, capabilities or operational procedures are not as compelling from a decision making perspective.

Deployment & Service Models

These are two important terms for considering migration to a public cloud.

1. Service Models – (SPI Tiers) — SaaS, PaaS, IaaS
2. Deployment Models - Private (within firewall), Public, Community, Hybrid (includes on-prem)

Service Models

Different service models are used for different purposes. In summary they are:

Infrastructure as a Service (IaaS) is a self-service model for managing remote data center infrastructures. IaaS provides virtualized computing resources over the Internet hosted by a third party such as AWS, GCP or Azure. This provides you with greater control of managing your cloud computing resources and building your own security architectures within the IaaS.

Platform as a Service (PaaS) allows organizations to build, run and manage applications without investing in their own IT infrastructure. This makes it easier and faster to develop, test and deploy applications. Developers can focus on writing code and create applications without worrying about time-consuming IT infrastructure activities such as provisioning servers, storage and backup. PaaS

Enterprise Data Security for US Europe and Asia

reduce your management overhead and lower your costs. PaaS also makes it easier for you to innovate and scale your services on demand.

Software as a service (SaaS) replaces the traditional on-device software with software that is licensed on a subscription basis. It is centrally hosted in the cloud by the SaaS application provider. A good example is Salesforce.com. Most SaaS applications can be accessed directly from a web browser without any downloads or installations required. However, some SaaS applications require plugins.

The following provides a more detailed description of how these models are used.

IaaS

NIST defines IaaS as the “ Capability provided to customer to consume provision processing, storage, network, and other fundamental computing resources, where the consumer is able to deploy and run arbitrary SW which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over the operating systems, storage and deployed applications, and possible limited control of select networking components (e.g. host firewalls).”

In the case of the IaaS model, the customer is fully responsible for how they define and implement network security for their environments such as IP address groups, and also for provisioning processing, storage, and networks.

Key benefit to customer is metering and price on the basis of units consumed. But it does not represent transferring cost of ownership, ability to scale infrastructure, or increased cooling and energy efficiencies. CSPs never allow clients direct access to servers — only to VMs. Hypervisors used by the CSPs have multi-tenancy even for this service model.

Private IaaS provides the customer the most control of their cloud assets. The customer can configure cloud memory at “VOLUME” level to install programs and partition storage for their teams and at “OBJECT” level to allow sharing of data in a structured manner. For data-at-rest, the default encryption methods include file level encryption. It also provides DRM (Digital Rights Management) at the application level. The cloud customer accepts responsibility for securing their cloud based applications.

Typical IaaS attacks include VM attacks, virtual network (virtual switch), Hypervisor, VM rootkits (malicious hypervisor on the fly), DoS (denial-of-service), and other colocation and multi-tenancy related exposures. The customer has to really design the full extent of security of sensitive data used by their applications. The CSP can provide volume-based storage encryption for which the CSP

Enterprise Data Security for US Europe and Asia

maintains the key. The encryption engine is located on the instance itself for instance-based encryption, and for proxy-based encryption, proxy has to run on a separate secure machine.

PaaS

NIST defines PaaS as “ the capability provided to the consumer is to deploy into the cloud infrastructure consumer-created or acquired apps created using programming languages, libraries, services, and tools supported by the CSP. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems and storage, but has control over the deployed apps and possibly configuration settings for the app-hosting environment.”

For PaaS, the CSP provides the operating system and a set of services. You, the customer, are responsible for security for everything you implement through your on-premises IAM extended to your cloud environment. You can, but do not have to, build your PaaS on your IaaS. A fully negotiated contract is greater likelihood than for SaaS.

PaaS is the most preferred method for migrating on-premises legacy applications to cloud environments.

Knowledge of the CSP systems is the biggest challenge for the data custodian to properly adhere to policies over and above that for IaaS. This is not a contractual requirement. The data custodian also needs to have knowledge of data classification rules, or access controls to systems.

A PaaS allows you, as the customer, to deploy into cloud applications you acquired or created using programming tools supported by the CSP. You do not manage or control the underlying infrastructure but have control over the deployed applications and, potentially, configuration settings for your application hosting environment. You will most likely be responsible for administration and data security of your applications in production in your PaaS cloud environment.

Each of your PaaS instances should have its own user-level permissions. When instances share common policies and controls the cloud security policies should be carefully implemented to reduce authorization creep and inheritance.

Database storage allows you as the customer to install and run applications in the cloud and store data. For data-at-rest, encryption methods include file level encryption. DRM is available at the application level.

Typical PaaS attacks have a higher likelihood of suffering backdoor vulnerabilities installed during software product development. Backdoors are left by developers inadvertently, or on purpose by malicious developers.

SaaS

NIST defines SaaS as “ capability provided to the consumer to use the CSPs apps running on a cloud infrastructure. The apps are accessible through either a thin-client interface, such as a web browser (or a web-based email), or a program interface. The consumer does not manage or control the underlying infrastructure including networks, operating systems, storage or even individual App capabilities with the possible exception of limited user specific app config settings.”

SaaS is a well-known for hosted application management and software on demand and has been in extensive use for some time. The major benefits of this service model are cost, ease of administration, and automated software updates managed by the vendor of the SaaS application and the CSP. SaaS software can be built on PaaS and IaaS. SaaS is the most cost-effective cloud model.

If you deploy a SaaS application in the cloud, you as the customer are responsible for IAM for your team and your customers, training your team, and managing user access.

Common attacks in the SaaS model are security risks due to misconfigurations which is a common source for attacks for on-premises applications, Shadow IT within your own organization, storage exposures, and access management based attacks.

Deployment Models

Cloud deployment is the installation of hardware and software accessible over the Internet on a specialized platform. Each deployment model is defined according to where the infrastructure for the environment is located. The four deployment models are:

1. Private -- within the data centers inside the firewall.
2. Public – on public clouds such as AWS, GCP, Azure etc.
3. Community – where a group of organizations create a shared cloud
4. Hybrid -- includes on-premises, public and shared

Private Cloud

NIST defines private cloud as “Cloud infrastructure provided for exclusive use by a single organization comprising of multiple consumers (business units). It may be owned, managed and operated by a third party or some combination of them and it may exist off-premises.”

Public Cloud

NIST defines public cloud as “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.”

Hybrid Cloud

NIST defines hybrid cloud as “ cloud infrastructure is a component 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).”

For a hybrid cloud, you will need to connect your private cloud in the data center to the public cloud via WAN or VPN. To minimize conflicts the two clouds should not use the same IP address ranges. Furthermore, they should be at equivalent security levels.

To manage access to the internal or private network from an external network you will need to use a “transit” virtual network such as a “Bastion” host. This is a single dedicated virtual network for the hybrid connection. You can then peer the other network via this designated Bastion. It is important to deploy firewalls rulesets, ACLs, and security tools via the Bastion host.

Community Cloud

NIST defines community cloud as “ the cloud infrastructure is provided 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 organizations within the community, a third party, or some combination of them, and it may exist on or off premises.”

Cloud Security

Cloud specific risks for developers to address include multi-tenancy and third-party administrators. Attacks such as denial-of-service (DoS) or distributed-denial-of-service (DDoS) are the responsibility of the CSP. The organization's security team should not resort to default configurations and access controls. DDoS in the cloud prevents users from connecting to their cloud assets while DDoS on-prem allows users to still connect on personal devices.

While excessive use of security controls can lead to user dissatisfaction, for Bring-your-Own-Device (BYOD) such as user personal laptop or mobile device, a cloud security professional must always be cognizant of concern due to potential for broad unmanaged access. BYOD users may utilize unapproved applications and APIs to enhance productivity not realizing risks and simply assuming no malicious intent of the app vendor but that can pose significant data security risks. Because cloud access is all remote, controls between users and the environment should include all of the following:

- Logging and audits to identify unauthorized users.
- Encryption to secure data.
- Secure login with complex passwords.

It should not allow “once in all in”. Every new access must be authenticated.

Contractual controls help extend your Governance ability to CSP. Your Security controls help to implement data governance. CSPs do not list their security models publicly to avoid attackers. They are available on request as reports and you should obtain them to develop your own Cloud Data Security Plan.

While migrating to cloud for a SaaS solution hosting a new concern is data disclosure through insufficiently isolated resources typically by just selecting default configurations. Inadvertent disclosure by internal company resources and malicious intrusion by external entities should also be addressed. It is also important to address any regulatory security compliance for financial or credit card processing transactions.

Logging and Monitoring

Logging in the cloud may be a challenge depending on the CSP. IP addresses may not reflect a particular workflow because multiple VMs share IP addresses. Logs can be lost if the controller shuts down the instance. Sending all logs to SIEM may

be cost prohibitive. So your organization needs to invest in a central logging processing approaches.

Security operations procedures for the cloud for audit logging should include adding new audit/logging rules and amending rules to filter out false positives, with clear definitions of security events and incidents. You can use either the CSP's audit logging tool or implement one you have selected for its features in providing services such as searching and collating logs for reporting. Logging should also be a part of your DRP (data recovery plan). NIST SP 800-92 describes policies and procedures for log management, and prioritizing log management throughout the organization. Log management tasks include creating and maintaining the log management infrastructure and tools, providing support staff, and establishing standard log management procedures.

Cloud Segmentation

Logical segmentation in the cloud is achieved by VLANs, NAT (Network Address Translator), and Bridging. Lack or ambiguity of physical end-points as individual network components in the cloud make it difficult to uniformly define, manage, and protect IT assets

SANS Institute's CIS (CIS/SANS) Security Controls for effective Cyber Defense are a recommended set of actions for cyber defense that provide specific and actionable ways to stop today's most pervasive and dangerous attacks. A principal benefit of the controls is that they prioritize and focus a smaller number of actions with high pay-off results. Control2 includes app whitelisting, managing authorized software by system, alerts on finding unexpected software, software inventory tools usage, dangerous file types, virtualized operating systems use. You should allow usage only of software that allows signed software ID tags.

Cloud Selections Best Practices

Even a small organization without extensive funds can, with a limited staff, build and manage cost-effective services by moving to the cloud. But all organizations, big or small, must invest in data security. They need to follow these best practices guidelines to achieve adequate data security in their cloud environments.

NIST's definition of a cloud carrier is "an intermediary that provides connectivity and transport of cloud services from Cloud Providers (CSPs) to Cloud Customers."

Enterprise Data Security for US Europe and Asia

An Internet Service Provider (ISP) such as, Verizon or Airtel is a carrier which provides access through dial-up, DSL, cable, and fiber-optic links so it is not a cloud carrier.

A Managed Service Provider (MSP) with a Network Operations Center (NOC) manages the customer's IT infrastructure but the customer dictates the technology and operating procedures for the computer servers within their infrastructure.

Key Players Within the Cloud Customer Organization

A number of players within the public cloud customer organization play key roles. These include the following:

- Cloud Administrator,
- Cloud Application Architect,
- Cloud Infrastructure Architect,
- Cloud Data Architect,
- Cloud Application Developer,
- Cloud Data Steward
- Cloud Storage Administrator

Why Should you Migrate Applications to the Cloud

The leading factor in a decision to move to the cloud is overall cost reduction due to outsourcing administration. Other factors that play a role include network scalability, global accessibility and offsite backup capability. Reduced cost of administering the operating systems and servers in the cloud is also a reason to go to cloud but the more important reasons are: reduced cost of ownership, reduced energy costs and metered usage (the utility model) for use of computing resources and moving data.

Key factors to consider while migrating to the cloud include:

- Interoperability - this is most crucial in choosing a cloud provider (consider this more compelling than resiliency, portability and governance).
- Portability – is to prevent the CSP from holding client hostage. You can test portability conducting service trials in another cloud.
- Auditability – CSP will allow customer organizations to determine with assurance that all contract terms are met.
- Proximity of the CSP to the customer's employees should be the least important factor in your CSP selection process compared to depreciation of assets, utility consumption costs, and shift from IT focus to business focus. Your employees can operate all cloud functions remotely.

- Cost of software licensing is reduced by moving to the cloud for the services provided by the CSP. Consider using cloud-native services to the extent feasible.
- Utility costs for the on-premises data centers are reduced by moving to the cloud but not necessarily the cost for data protection and security. The organization is responsible for the security of the data owned by the organization even after it is moved to the cloud.
- Loss due to depreciation of assets is reduced by moving to the cloud.

Major Considerations for Moving to Public Clouds

Migrating to the public clouds starts with a CSP contract. A CSP/Customer contract is the only guarantee of service and commitment. Contracts help extend governance to the cloud.

Service level Agreement (SLA) satisfaction surveys about the CSP under consideration are an important decision factor for your cloud CSP selection decision, especially as it relates to data security evaluations.

The following lists some key considerations when you migrate your products to the public cloud services providers (such as AWS and GCP).

Interoperability – how easy is it to operate with applications spread across clouds. This is important for business continuity (BC) and disaster recovery (DR).

Portability – prevent Vendor Lock-in – identical app components can exist in multiple clouds. This ensures that you can transfer your cloud assets to another CSP should the need arise. This also facilitates BC/DR.

Availability – defines success/failure of cloud services, CSPs provide 99.9%. There should be no single-point-of-failure (SPOF).

Security – CSPs require NDAs to share SoC reports. SoC reports are not listed for public to avoid hackers.

Privacy – No uniform international privacy laws or directives, or any geographical challenges are typically addressed by CSPs. You need to develop your own security models.

Resiliency – CSP should be able to continue operations in the event of a disruption or a natural disaster in one region.

Enterprise Data Security for US Europe and Asia

Performance – CSPs must maintain performance to stated SLA levels.

Governance – processes and decisions defining actions, assigning responsibilities, and verifying performance, with documented procedures for addressing shortfalls.

SLAs – created by CSPs are heavily weighted in favor of CSP but include downtime, upgrades, patching, vulnerability. Most CSPs resist any modifications to their SLAs.

Auditability – allows users and organizations to access, report, and obtain evidence of actions, controls, etc.

Regulatory Compliance – Confirm that the CSPs are compliant with regulations and laws, such as, HIPAA, GDPR, SOX, PCI-DSS. You may need to fill the gaps.

Cloud Migration Challenges

Managing file storage systems in cloud environments is challenging because virtual machines (VMs) are stored as snapshot files when not in use. The following describes the typical storage types in cloud environments. Ensuring data security and access controls for sensitive data within all of these storage structures is the responsibility of the customer organization. The CSP provides only perimeter level security.

Favorable contract language is probably the best tool to avoid vendor lock-in and establish how data can be moved to another CSP vendor easily. Best way to avoid vendor lock-out is to use another CSP for backups right from the start. Lack of standards among CSPs makes migrations across them complex which leads to unforeseen vendor lock-in problems.

Prohibitions by the CSP on port scanning and penetration testing for testing data security within the cloud environment can hamper your ability as the customer organization to protect your cloud data. Cloud administration violates the Brewer Nash (Chinese Wall) principle because the same cloud CSP administrative staff can operate on competing customers. In large cloud providers it may happen unknowingly.

Cloud Architecture

Architecting a cloud-safe application to ensure security of your sensitive data involves a number of key considerations. These are explored in the following sections.

Build and Configuration Management Automation

Infrastructure-as-Code (IaC) has become a popular way to provision infrastructure predictably and consistently. It also gives teams the ability to guarantee application security at the earliest stages in the development lifecycle, a process known as shift cloud security left. Automation is a means to achieve consistent repeatable configuration management. It is achieved through orchestration of configurations, coordination, and management of cloud services.

The cloud management plane abstracts and centralizes administrative management. It is responsible for managing the assets of the resource pool. The CSP is responsible for the security of the management plane and exposing its features. The management plane controls the metastructure (protocols and mechanisms that provide the interaction between the infrastructure and other layers). It is accessed via APIs and web consoles. The management plane runs on its own servers with dedicated connectivity to hosts being managed. The management plane should be accessible to and used by the designated most privileged users of your organization who use an API interface to manage your cloud environment.

Cloud computing has two main layers of infrastructure:

1. Fundamental resources pooled together to create a cloud (raw, physical and logical compute processes, etc. and
2. Virtual abstracted infrastructure managed by you as the cloud customer organization (computer, network and storage assets they use from resource pools.

Cloud Storage

Cloud volume and object storage use RAID magnetic hard disks and high speed SSDs. The typical storage types provided by CSPs are:

Object storage which stores objects in a hierarchy such as a file tree with minimal features. The object stores in cloud use an opaque value or descriptor to categorize and organize data. It is not structured, unstructured or volume data.

Volume storage and block storage do not have file structures. You, as the user, are assigned a logical storage area in which raw data and files can be stored. Objects are stored with related metadata such as creation date and content type.

Typically, SSDs are more expensive than spinning platters, but they are faster. SSDs are not vulnerable to degaussing (they are not magnetic but solid state). The most potential problem point for object storage with high reading/writing dependencies is replication of objects and integrity. Data can be replicated at block level, file level, and DB level. The operating systems on hosts within the cloud environment handle formatting, allocation and security controls for volume storage

Database Management Systems are offered by CSPs as their own versions of structured data management applications and they also provide some industry standard RDBMSs as well as CDNs (Content Delivery Networks) often used for streaming data to users. Data security within these CDNs is also your responsibility as the customer organization.

Containers and Kubernetes (k8s)

Containers are used for serverless computing and are code execution environments that run within an operating system and share and leverage resources. VMs on the other hand are a full abstraction of the operating system. Containers are a constrained place to run segregated processes while utilizing the kernel and other resources. Multiple containers can run within a VM, and can be launched rapidly. K8s has become the de-facto technology for public cloud containers. Data security within the containers requires special design considerations due to potential sharing by multiple applications within your own organization.

Container security includes:

1. Security of compute, network and storage infrastructure,
2. Security of management plane,
3. securing the image repository,
4. Building security in the tools and code inside the containers.

Containers do not provide full isolation security but provide task isolation.

Virtual machine (VM) Related Architecture

VMs snapshotted when not in use do not get updates so it is really important to have vulnerability scans when they are re-instantiated so that new security updates are applied to them. It is also important to have automatic registration with the

configuration management system which would ensure that each VM has the correct updated version. Event logging must also be turned on for incident management.

Immutable workloads enable security. These cannot be patched or change code while running within a VM. Security should be checked while creating images using automation to ensure that required security constraints have been met.

Some constructs such as shared databases are platform-based workloads running in their own virtual private clouds. They are neither VMs or containers. Access Controls in those are quite often based on using the individual user login access controls for the database system.

Security Governance Plans

The CSP crafts, creates and promulgates Governance that will determine which controls are selected for the environment and how they are deployed because the CSP runs and operates the cloud data centers. The customer and regulators may require additional controls that may be crafted on the customer side or as special contractual requirements with the CSP.

Defense in Depth (layered defense), a fundamental aspect of security principles should be implemented rigorously on-premises as well as in cloud environments.

Organizations moving to a cloud needs to know the Service Model (Private, Public, Hybrid, Community) and understand and plan for the risks and security constraints for the model they use for resource sharing. They also need to fully understand the security parameters of the Deployment Model (SaaS, PaaS, and IaaS) for instituting proper controls.

Negotiating to conduct a trial run in the cloud before a permanent migration is the best method for reducing the risk of the application not delivery functionality and performance.

Software developers creating cloud apps need to be aware of encryption of data at rest, in motion, and in use, and data masking and obfuscation options to ensure that they meet the required standards. Only applications using databases may need hashing of fields.

General Cloud Risks and Challenges

The cloud management plane breach is probably the most significant breach because this breach allows the hacker full access to the cloud environment.

Enterprise Data Security for US Europe and Asia

On moving to cloud, the organization should weigh the risk of allowing external entities to access the cloud data for collaborative purposes against sending the data outside the legacy environment for collaboration.

Multi-tenancy is a new added risk in the cloud because bad actors can have legitimate access to the cloud and can in-fact have access to the same hosts running their VMs alongside your VMs. The bad actor may be able to overcome the protections within the host and be able to access your VM. Data Security and cyber security in general are very important to prevent data exfiltration in such a case. VM Hopping is a type of attack that a flawed hypervisor can allow when inter-VM isolation or trust levels are misconfigured.

A new management risk in the cloud is virtual sprawl due to uncontrolled growth because of the ease of use in adding hosts and new services. Your users and developers do not see the costs and may add new hosts but not release them when no longer needed. It is a resource management problem (technical and logical issues are contributory to the resource management issue). Additional risk of sprawl is inadvertent activity. Setting policies on who can instantiate new hosts and using automation to ensure that the correct and safe VM images are used is essential to mitigate security risks. There should also be automation to warn users of unused assets that can be removed after a certain period of unuse.

The benefit of a private cloud is retaining control of governance, but it could have the same cost issue as on-premises data centers.

Enterprise Risk Management

Cloud risk should be linked to corporate governance. You should apply Enterprise Risk Management (ERM). ERM is an organization-wide strategy to identify risks and shape the organization's risk posture by mandating who can perform certain business activities and the policies they must follow.

Risk Categories

Risk categories include the following:

- *Policy* – The purpose of organization policies is to ensure that all employees and departments and business units follow the corporate policies for secure operations ensuring the safety of corporate data assets.
- *Organization* – Migrating applications and data to a CSP poses risks in the form of provider lock-in (inability to move to another CSP due to contractual terms or use of proprietary services, governance loss, (because the CSP would

Enterprise Data Security for US Europe and Asia

not permit invasive security policies and procedures), compliance (in case the organization is required to ensure compliance with government or industry standards not subscribed to by the CSP) and provider exit (the CSP goes out of business or is acquired).

- *General* – Migrating to a CSP (technical control is shifted to the CSP),
- *Virtualization risks* -- (Guest breakout attacks, uncontrolled sprawl, lack of VM snapshot and image security),
- *Cloud-specific risks* - (management plane breach, resource exhaustion within the CSP, DoS, Isolation control failure across hosts, insecure or incomplete data deletion, control conflict across tenants, untested or insecure software in CSP),
- *Non-cloud specific (same as on-prem), Legal* (data protection, jurisdiction, law enforcement exposure, licensing on porting). New attack vectors in cloud.

Countermeasures for Managing Risk

Your risk mitigation plan should always have a number of countermeasures built-in to avoid or minimize risk. These countermeasures should include:

Defense-in-depth – create multiple layered security measures to protect your corporate data assets such that if one line of defense is compromised, additional layers exist as a backup to ensure that threats are stopped along the way.

Preventive Controls – put in place data security controls to prevent or mitigate unauthorized access to sensitive data.

Compensating controls – use of compensating controls is a mechanism that is put in place to invoke an alternative security control if the primary control fails such as a secondary signature to authorize critical or sensitive transactions or the creation of exception reports to catch unauthorized access.

Automation – automating the VM build process using an open source cloud agnostic IaC automation software tool such as Terraform minimizes the risk of unauthorized creation and configuration of VMs.

Access Controls -- Streamline end-to-end governance, risk, and compliance (GRC) by automating access risk analysis and identifying user risks before access is granted.

Note that cloud computing runs on real hardware and software services: hypervisor, storage controllers, volume mgmt., IP address management, security group management, VM image service, identity service, message queue, management databases, and guest operating systems. You need to clearly evaluate what is under

the purview of the CSP that you have no control over other than the contractual terms and what is your responsibility for data protection.

Access Management

Two important constructs for access management are trust zone management and access management.

Trust zone management – is for creating secure trust zones which are separated by airgaps from non-secure zones which are prevented from accessing sensitive data. Use a *jump server* (called a jump host) such as a Bastion host in AWS to compartmentalize and control access to trust zones. This is the only host exposed for access outside the VPC.

Access Management -- is the set of practices that enables only those users permitted the ability to perform an action on a particular resource. The three most common Access Management services are: Policy Administration, Authentication, and Authorization. The real risk decisions are managed at the PDP (Policy Decision Point) and PEP (Policy Enforcement Point).

Risk Management Guidelines in NIST Cloud Technology Roadmap

The NIST Cloud Computing Standards Roadmap is a comprehensive document that describes all definitions related to cloud computing and operations. If you operate any part of your software in the US, a thorough study of this document is very important. Even otherwise it provides an excellent understanding of cloud enablement of applications.

Service Level Agreement (SLA)

An SLA is the agreement that a CSP has with its client customers on the level and type of service it provides. This is a very important document that determines the level of control you as the customer have on creating and managing data security controls for your products deployed in the cloud. You need to keep the following in perspective before signing an SLA:

- The incentive for the CSP to perform is financials penalty for not meeting service levels. There is no penalty or liability for the CSP for risk exposure nor liability for violating data breach notification to your customers. You as the client is responsible for that.

Enterprise Data Security for US Europe and Asia

- The SLA is best to ensure that you as the client or the cloud customer receive dependable, consistent performance. The SLA typically does not cover audits performed by the CSP, training of your resources, or responses to regulators.
- Main SLA components typically include undocumented single points of failure (SPOFs), migration to alternate providers and components they need to support, if data backups included, and any other actions that might impact the SLA.
- SLA issues you need to watch out for: CSP provided uptime guarantees, CSP liability, data protection requirements if any (typically your liability), disaster recovery, suspension of service (due to delayed or unacknowledged payments for service), security recommendations, SLA penalty exclusions (when downtime starts and scheduled downtime). Disaster Recovery is a tricky part of SLA.
- QoS (Quality of Service) issues: Availability, outage duration, MTBF, capacity metric, performance, reliability, storage device capacity, server capacity, instance startup time, response time, completion time, mean-time to switchover, mean-time system recovery, storage and server scalability. Ask for specific reports to cover these aspects.

Cloud Audits

Cloud Audits are typically mandated by government entities and sometimes by industry consortiums. For an audit associated by a legal order of compliance such as an MRA (Matters Requiring Attention), data discovery about IT assets, their configurations and access controls and evidence collection of this data are big challenges due to the ephemeral nature of the assets, multi-tenancy which limits CSP permitted access to the assets. The regulators will determine if the cloud migration is satisfactory, and not the CSP or management and they can find satisfactory evidence. To ensure that evidence will be available the continuous operations tracking principles must be employed. These include audit logging (new event detection, new rules, configuration changes, changes in hardware and software assets allocated), contract and authority maintenance, secure disposal of assets, and incident response legal preparation.

Chain of Custody refers to a process used to track the movement and control of an asset through its lifecycle by documenting each person and organization who handles an asset, the date and time it was collected or transferred, and the purpose of the transfer. This is a big challenge for compliance with audits across multiple data centers in many jurisdictions.

Enterprise Data Security for US Europe and Asia

The key challenges related to CSP operations are in collecting evidence: privacy issues with seizure of servers from many customers, trustworthiness of evidence from CSP, unqualified technicians collecting evidence, contamination of data, lack of physical access to host OS and Guest OS (VM), collecting metadata exposure (data bleed), etc. Also, location of storage, recording in evidence log, undocumented removal from storage, undocumented transport of evidence, lack of procedures for running tests, etc.

Chapter 4 – Data Classification Types and Standards

For protecting data it is important to know what the data represent. Not all data can or should be protected. Data classification is used to identify data that has sensitive, confidential or highly confidential information that needs to be protected. The protection of data starts from the very moment it is created until it is removed or destroyed. This protection must be carried through all stages of data storage (at rest) or movement of data (in motion) or in use.

Data Life Cycle

Any data created for an information system can undergo six phases during the course of its existence. These can best be remembered by the acronym CSU-SAD where the alphabets of the acronym are:

C S U – S A D – Create Store Use -- Share Archive Destroy

Data security applies to the full cycle of data through the six phases listed below. It must be protected during every phase of this cycle.

Phase	On-Prem Management	On Public Cloud Management
Create	Data is created by a user when that person executes an application and enters information. The data must be classified at the time of creation and sensitive data (such as PI information) should be encrypted, masked or tokenized immediately before storing that data. The user may be a data subject or some other	Data will most often be created by users accessing the cloud remotely. Depending on the use case, the data might be created locally, by users at their remote workstation, and then uploaded to the cloud, or in the cloud datacenter via remote manipulation of the data residing there. When data is created remotely, the data created by the user should be encrypted before uploading to the cloud unless the transport mechanism is based on TLS 1.3 protocol at the transport

	<p>user acting as agent on behalf of user agent (such as a banker or stock broker).</p>	<p>layer. We want to protect against obvious vulnerabilities, including man-in-the-middle attacks and insider threat at the cloud data center.</p>
<p>Store</p>	<p>Store is usually meant to refer to near-term storage. The activity in the Store phase occurs almost concurrently with the Create phase— that is, Store will happen as data is created. Sensitive data should be encrypted when stored (data at Rest) for mitigating exposure to threats within the datacenter</p>	<p>Storing data in the cloud requires special consideration, especially, if it includes sensitive data. Data is stored in logical pools spread across physical storage that spans multiple servers (sometimes in multiple locations), and the physical environment is typically owned and managed by a hosting cloud service provider company. Advantages of cloud storage include usability, sharing data and accessibility, disaster recovery and cost-efficient scalability. Data Security is an additional concern for data stored in the cloud, especially sensitive data.</p>
<p>Use</p>	<p>Data in use is transferred to an application and converted back to ciphertext if transmitted in encrypted form. It is more vulnerable than data at rest because, by definition, it must be accessible to those who need it. The more people and devices that have access to the data, the greater the risk that it will end up in the wrong hands at some point. Securing data in use requires user authentication and controlling access as tightly as possible</p>	<p>Operations in the cloud environment will necessitate remote access, so those connections will all have to be secured, usually with an encrypted tunnel. Best practices in include:</p> <ul style="list-style-type: none"> Not storing sensitive information Encrypt data, explore using Cloud Services to encrypt data Use strong passwords for cloud access, especially if sensitive data is stored. Separate encryption key storage from data storage.

<p>Share</p>	<p>Sharing data with downstream applications is a common practice, which leaves a lot of questions to be answered depending on if the data is shared among applications on-prem, to the cloud or to third parties. Key issues include data lineage, data owner rights and ability to define usage and ensuring that entities receiving the data will honor security requirements.</p>	<p>Data moved into the cloud is typically used for sharing with other applications and needs special protections because the data owner has no control on the receiver of the data.</p> <p>Key approaches to protecting such data such as encryption and Digital Rights Management (DRM) become even more important for the cloud.</p>
<p>Archive</p>	<p>Data archiving is the process of moving data that is no longer actively used, but may be important to be retained long-term for future reference or regulatory reasons, to a separate storage device for long-term retention.</p>	<p>Cloud archiving is often completely in a public cloud. Some key points to consider:</p> <ul style="list-style-type: none"> Does the archive include sensitive data and is it safe in the archive. Is there a backup for the archive in case the cloud instance has a disaster. Are all cloud data security rules applied to the archive.
<p>Destroy</p>	<p>Sensitive data should be kept only as long as it is needed. Any other data should also be destroyed when it is not needed to minimize the risk of disclosure and to reclaim storage space. Typical methods include:</p> <p>Wipe and remote wipe (this requires multiple passes at replacing real data with fake data and</p>	<p>The public cloud does not offer the option of degaussing or destroying the storage media by any other means. So for data that is no longer needed and needs to be destroyed, an effective approach is crypto-shredding; the process of encrypting the data once and then encrypting the encrypted data a second time with a different key, and then deleting both keys.</p>

	reformatting multiple times)
	Degaussing to remove all magnetic field
	Physical shredding of the hard disk or optical device

Table 4-1 - the CSU-SAD Data Life cycle

Data Classification

The Data Classification Standards within the organization define the data that needs to be protected.

The following are the key reasons why the Data Classification standards need to be adhered to for protecting valuable data assets collected and used by the organization. It assists in identifying those assets that are most critical or valuable to the organization.

- It lends credence to the selection of protection mechanisms.
- It is often required for regulatory compliance or legal restrictions.
- It helps to define access levels, types of authorized uses, and parameters for declassification and/ or destruction of resources that are no longer valuable.

Data Classification by Type

Data classification is commonly seen as four levels: highly confidential (or restricted), confidential, internal only, or public. Note that PI (personal information) is a broad term referring to any data related to an individual while PII (personally identifiable information) is a subset of that which can uniquely identify a person.

A individual can has an identity (who they are, such as name, government ID, email address, etc. which can uniquely identify that person) and identification (such as an account or credentials assigned to that individual such as those by an employer or a bank.).

Enterprise Data Security for US Europe and Asia

For better understanding the table below breaks down further the restricted data into Highly Confidential, Authentication Credentials and Authentication Identifiers and the internal only data into proprietary and private classes:

Classification	Description	Reason for Protection
Highly Confidential	Customer identity, Non-public information personally identifiable information (PII) government identification, driver license, passport, biometrics, or other government issued ID	Regulatory, Legal
Authentication Credentials	Username, password, credit cards, and mobile phone number.	Regulatory, Legal, reputation
Authentication Identifiers	DOB, mother's maiden name, phone # on account, parties to an account, etc,	Regulatory, legal, customer data
Confidential	Customer information, sales records, financial records, email, addresses, bank balances, etc.	Regulatory, Legal
Proprietary	Company's org charts, design documents, telephone directories, department reports, financial reports, emails, communications, etc.	Company's crown jewels
Private	Company's private data such as customer lists, customer account details other than PII.	Competitive critical info
Public	Published data visible to all – marketing literature, web-sites, etc.	No protection required.

Table 4-2: Data Classification Types

Note: Authentication credentials and identifiers are also considered Highly Confidential.

Sensitive Data: Sensitive data is any data that if breached would cause damage to the mission of the organization. All of the classes of data listed in the table above

except Private data are considered sensitive data that need to be protected. Note that the Highly Confidential and Authentication data needs special care and protections for the cases of data at rest, data in transit and data in use. A data breach is any event in which an unauthorized entity is able to view or access sensitive data.

Understanding Regulatory Compliance Requirements

It is important to understand what compliance requirements your organization is subjected to. If you are a large financial institution or a bank, you may be subject to FFIEC or OCC audits in the US. If in Europe, you may be required to fully understand the compliance requirements for GDPR. Similarly, India and some other countries have national requirements for banking and credit card usage. You may also be subject to state compliance requirements. In the US, PCI-DSS (Payment Card industry – Data Security Standard) compliance is required for all institutions collecting and storing and credit card data. . If you manage patient medical records, you need to ensure compliance with HIPAA standards.

Steps to Classify Data

Several steps are involved in properly classifying data. These include the following:

1. Define Data Classification objectives.
2. Identify Data.
3. Data Governance – the regulatory requirements.
4. Define Categories and categorize data.
5. Define Usage outcomes for the data categories.
6. Monitor and maintain.

These steps are described in the table below.

#	Step Title	Description
1.	Define Data Classification Objective	Define the goals and objectives for which data classification is being performed; i.e. is it for voluntary protection or is required by a regulator.

Enterprise Data Security for US Europe and Asia

2.	Identify data	Discover where your data resides, how valuable is it, how many copies exist, and who has accessed the data.
3.	Data governance – regulatory requirements	Determine if there are regulatory or standards compliance requirements for your data and if any audits will be performed by an external agency.
4.	Define categories and categorize data	Define the data categories as per the classification discussed earlier and assign values and associated risk in case of a loss of the data.
5.	Define usage outcomes for data categories	Determine how the data should be organized and how it will be used by various stakeholders and applications that need access to your data.
6.	Monitor and maintain	Install tools and procedures to monitor all activities associated with the classes of data that need to be protected.

Table 4-3: Steps to Classify Data

Data Security Posture Management

Continuously assess your data and security posture, automatically remediate issues with data policies, and maintain compliance with industry regulations ensuring that data is known, protected, and governed depending on jurisdiction and data type.

- Use tools including AI based tools to discover data across your infrastructure, including cloud, SaaS, and on-premises environments.
- Monitor changes and detect data drift from what you expect with continuous scanning and data mapping,
- Discover unknown data and identify whether it is sensitive or exposed to privacy or data security risks.

With data scattered across multiple clouds and data stores, security teams need continuous visibility into where their sensitive data resides, who has access, and which data is at risk. Data security tools need to provide automated data discovery and integrating data intelligence with on-premises data centers and other cloud

Enterprise Data Security for US Europe and Asia

risks. They should deliver DSPM telemetry without requiring additional tools and integrations. The tools should cover the areas such as on-premises or cloud misconfigurations, vulnerabilities and other potential threats.

Chapter 5 – When Data Needs to be Protected

Data protection is important for your organization, because it prevents the critical confidential information of your organization from fraudulent activities, hacking, phishing, and identity theft. To work effectively and ensure that your business interests and the reputation of your organization is maintained at the highest standards, you need to ensure the safety of your information by implementing a detailed data protection strategy and operational plan that is rigorously enforced throughout your organization.

As the amount of data stored and created increases, so does the importance of data protection. Data breaches and cyberattacks can cause devastating damages not only to your business interest and reputation but also to your most valued customers. So as a data hosting organization you need to proactively protect your data, especially workforce and customer personally identifiable data, data and regularly test and update your protective measures.

Ultimately, the key principle and importance of data protection is safeguarding and protecting data from different threats and under different circumstances. As the data hosting organization, you are totally responsible for it, even if your data is stored at a third-party location such as a colocation data center or a cloud.

The Type of Data That Requires Protection

Vital information of the workforce (employees) and customers, such as names, addresses, emails, phone numbers, health information, or bank details, are all data that should be carefully stored and protected. Data protection gains special importance when the information concerns customers. If such information gets in the wrong hands, it can compromise people's safety in many forms, including personal integrity, physical safety, and financial security. Stolen information can also be used by bad actors to create fake profiles and commit fraud. There can be significant penalties levied against your organization by the regulatory agencies in case of exposure of your data, even if it is due to misconfiguration by a co-location data center partner or a cloud service provider.

The following three terms are very important for you to understand from a data protection perspective.

- Personally identifiable information (also known as PII)

- Personal information (also known as PI)
- Sensitive information

What Is Personal Data

Personal data refers to digital or analog information that can be used to uniquely identify a specific person. Personal data can include a person's name, address, email address, IP address, phone number, driver license, Social Security number, government issued ID, banking information, health records, and depending on the comprehensiveness of the regulatory jurisdiction it can include more data items.

Some data in isolation may not represent personal data but in the context of other related information, it can be construed as PI. For example, John Smith, by itself without any other information attached to the name, would not be considered personal data in most jurisdictions for any legal reason. There can be thousands of people with the John Smith name, so the name by itself does not identify a specific person.

However, attaching a Social Security Number (SSN) in the US or an Aadhaar card in India uniquely identifies a person. Aggregated other information such as street address and phone number or email address can also uniquely identify a person. Most jurisdictions would call such aggregations personal data. As such, if you are in possession of such data about a person, you may be required by regulatory jurisdictions or national laws to protect such data and comply with data privacy regulations.

To ensure your organization is compliant with all relevant data privacy laws and regulations, understanding the definition of *personal data* across jurisdictions is imperative. As discussed in Chapter 2, this definition can change with regulatory and national jurisdictions. So you must take into consideration that based on the jurisdictions you operate in, how will personal data be interpreted.

PII (Personally Identifiable Information)

Technically, all personally identifiable information (PII) is considered personal data, but not all personal data is considered PII.

They're not mutually exclusive. So it is important to understand the differences and implications of these differences in how you address your data protection strategy and operational plan.

PII consists of any information about a person — including data that can trace or distinguish their identity — and any information that can be linked to them (like medical, financial, or employment data). But *PI* (*personal data*) on its own doesn't always consist of all those identifiers.

Distinguishing a person's identity means identifying one individual over another using specific data.

Tracing that individual means you're processing enough data to understand aspects of that person's status or activities. As such, personal information like name, email, phone number, Social Security number, etc. are considered *PII*.

To put it in perspective, the key difference between *PI* (personal data) and *PII* (personally identifiable information) is that *PII* is often used to differentiate one person from another, while personal data includes any information related to a living individual, whether it distinguishes them from another individual or not. The key differentiator is identity.

PI (Personal Information)

The term *personal information*, or *PI*, was first used in one particular data protection law: the California Consumer Privacy Act (CCPA) but is now commonly being used interchangeably with *PII* in the US.

CCPA defines *PI* as information that identifies, relates to, describes, is capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household.

However, this does not include information that has been made publicly available by the local, state, or federal government. Still, identifiers that can be linked to a California resident include those similar to *PII* identifiers: a person's name, address, email, Social Security number, etc.

A key difference is that unlike *PII*, *PI* under CCPA, also constitutes data like IP addresses; biometric, location, or audio information; and personal device identifiers.

Sensitive Information

The definition of *sensitive information* — also known as *sensitive data* — differs from one data privacy law to another.

Enterprise Data Security for US Europe and Asia

It is an overarching data definition such that sensitive information is personal data that most jurisdictions believe should be treated with a higher standard of care. To protect it, your organization may need to apply greater security measures. And, depending on the law, it's possible you'll need different kinds of consent to collect it. So, it is important to evaluate the requirements for sensitive data protection across jurisdictions in which you conduct your business.

If your organization allows unauthorized access to a data subject's sensitive information, you face a greater risk of being penalized by data protection regulatory agencies. Permitting access to sensitive data leaves a data subject open to various forms of harm and/or discrimination based on, for example, their sexual orientation, religious beliefs, private health matters, and the like.

Additionally, depending on location, sensitive information may comprise data collected from children. The latest GDPR regulations allow children 16 years and older to provide consent on their own for their data to be processed. Parental consent is required still for children 13 to 15; children under 13 cannot, under any circumstances, provide consent themselves.

Like PI, sensitive data does not include any data that the government makes available to the public. It is totally secure data in your organization's possession.

Like the other terms previously listed, the way each data privacy jurisdiction and national law interprets sensitive data varies. For that reason, it is important to check the privacy laws that apply to your regulatory and national jurisdictions before your organization collects personal information.

Protecting PI, PII and Sensitive Data

Corporate trade secrets, national security information, personal medical records, Social Security and credit card numbers are all stored, used, and transmitted online and through connected devices.

This proliferation of valuable data presents criminals with an increasingly wide range of opportunities to monetize stolen information and intellectual property. In addition, foreign governments, corporations, and organized crime rings have embraced hacking as one of the most potent tools at their disposal for pursuing objectives – be they geo-political, military, commercial or criminal.

Data needs to be protected for the following considerations:

- Workforce or employee data
- Customer data

Enterprise Data Security for US Europe and Asia

- Corporate financial data

Workforce (Employee) Data Protection: Workforce data protection is the act of ensuring that applicable privacy and data protection laws are followed regarding employee information. This means, of course, ensuring the data is stored securely, but also informing employees when the organization shares all or any of their data with a third party. Note that employee data is covered by GDPR.

Generally, employee data protection includes:

- Safeguarding the information: ensuring the information is protected and shielded from cyberattacks.
- Keeping employees informed about who is given access to their data, and why.
- Ensuring employees are aware and informed about their rights.
- Ensuring minimal access: access to personal data should only be accessible on a need-to-know basis.
- Have retention policies place for all employees including those who have left the organization and applicants who were not hired.

Customer Data Protection: Your customers are a valuable resource for your organization. You need to protect their data not only to prevent your competitors getting a list of your customers but also because the customer data is protected by regulatory standards and industry standards that govern your business. The requirements are very onerous for financial institutions that save customer identification and financial data such as for GDPR, CCPA, India Digital Personal Data Protection and others as described in detail in Chapter 2. The protections for customer data are very similar to those for employees, including:

- Safeguarding the information: ensuring the information is protected and shielded from cyberattacks.
- Keeping customers informed about who is given access to their data, and why.
- Ensuring customers are aware and informed about their rights.
- Ensuring minimal access: access to personal data should only be accessible on a need-to-know basis.
- Other specific rights provided by the various regulatory laws on industry standards.

Location of Data: Location of data is an important consideration because the location determines the types of protection available. The data can be at rest in on-premises data centers or in cloud environments, or it can be in motion when it is

being transferred from one location to another, or it can be use by an application. Let us review in detail the protections for data in these three states.

Protecting Data at Rest, Data in Motion, Data in Use

Effective data protection requires knowing where your PI, PII or sensitive data is at all times. You need to keep track of the data collected and used by your organization at all times. *Data Discovery Tools* can make the job of locating and tracking movement of data easier.

Data needs to be protected in three states: at rest, in use, and in motion. Each state presents unique security challenges.

Data State	On-Prem Protections	In Cloud Protections	Comments
Data at Rest	When data is stored on a hard drive (or SSD) it needs to be protected. On-premises information is primarily protected by conventional perimeter-based defenses such as firewalls, passwords and anti-virus programs. However, these barriers are not impenetrable. additional layers of defense to protect sensitive data from intruders in the event that the network is compromised.	Encryption is the only real way to protect data at rest in the cloud because the cloud customer has no way to directly access or protect the hardware media.	Encrypting Hard Drives or SSDs is the best way to ensure the security of data at rest. Other approaches include storing individual data elements in separate locations to decrease the likelihood of attackers gaining enough information to commit fraud or other crimes.

Data in Motion	<p>Data is at its most vulnerable when it is in motion, and protecting information in this state requires specialized capabilities. Our expectation of immediacy dictates that a growing volume of sensitive data be transmitted digitally— forcing many organizations to replace couriers, faxes, and conventional mail service with faster options such as email.</p> <p>An email typically takes many internet ‘hops’ across the labyrinth of public servers and broadband pipes that make up the virtual fabric of the public internet. Anyone with the right tools can intercept an email as it moves along this path.</p> <p>Data in motion from on-premises to the cloud must use secure transit paths and all data must have TLS 1.3 encryption at the transport layer.</p>	<p>Data transmission within the public clouds is encrypted by default by the cloud service provider.</p> <p>Data transmission from one cloud to another if managed by your organization needs to follow the same general rules as data transmitted from on-premises to the public cloud.</p> <p>The same rule applies to data being transmitted from the public cloud to your or your partner’s on-premises locations.</p>	<p>Email can be made more secure by ensuring that messages and attachments remain confidential by sending them through an encryption platform that integrates with the organization’s existing systems and workflows. Optimally, users should be able to send and receive encrypted messages directly from their standard email service. Policy filters that automatically detect sensitive information in email and attached files can automate the encryption process to further ensure security (and regulatory compliance).</p> <p>It is also important for the encryption service your organization uses to cover mobile email applications.</p>
-----------------------	---	--	---

<p>Data in Use</p>	<p>Data in use is more vulnerable than data at rest because, by definition, it must be accessible to those who need it. Of course, the more people and devices that have access to the data, the greater the risk that it will end up in the wrong hands at some point. The keys to securing data in use are to control access as tightly as possible and to incorporate some type of authentication to ensure that users aren't hiding behind stolen identities.</p>	<p>Data in use in public clouds needs to follow the same rules as the on-premises data in use in terms of least privilege and role-based permissions for accessing and using the data.</p>	<p>Organizations also need to be able to track and report relevant information so they can detect suspicious activity, diagnose potential threats, and proactively improve security. For example, an account being disabled due to a certain number of failed login attempts could be a warning sign that a system is under attack.</p>
---------------------------	---	--	---

Table 5-1: States of Data

Data Risk Assessment

Your organization needs to always assess risk from a workflow perspective.

1. Do employees access corporate systems from their personal devices or use company-issued devices to work from home?
2. What happens when employees take their devices on business trips? How is data transferred between devices or communicated to other stakeholders?
3. What your customers or business partners do with any sensitive files you send them?

Information will end up spreading across multiple devices and networks with varying degrees of security and risk. The following need to be addressed:

- What types of sensitive data does the organization store, use, or transmit?

- Who has access to this data and where, when, and why are they using it?
- How is data stored when it is not in use?
- How is access to databases controlled?
- What mechanisms are used to transport data?
- What are the pertinent laws, regulations, and standards?

Key Elements of Data Protection

One very important data protection model is the CIA triad, where the three letters of the name represent the three elements of data protection: confidentiality, integrity, and availability. This model was developed to help individuals and organizations develop a holistic approach to data protection. The three elements are defined as follows:

- **Confidentiality:** The data is retrieved only by authorized operators with appropriate credentials.
- **Integrity:** All the data stored within an organization is reliable, precise, and not subject to any unjustified changes.
- **Availability:** The data stored is safely and readily available whenever needed.

Your organization needs to build a framework of best practices that meet these three requirements of the data protection CIA triad. These best practices need to be built around:

- Access controls
- Accountability
- Data accuracy and integrity
- Purpose limitation
- Data minimization
- Encryptions based on location and use
- Secure storage management
- Data lifecycle management
- Regulatory compliance
- Auditing and configuration management
- Incident response and remediation
- Data risk management
- Data breach prevention
- Data backup and recovery

Data Protection Framework

As the number of organizations that process the personally identifiable information (known as PI) increases, so does the need for such organizations to ensure the safety and privacy of data. Training their workforce and their data management staff, and certification of the staff they have adequate training presents best practices related to the protection of the PII.

It is essential for your organizations to implement a data protection framework that provides guidance on the protection of PII. The framework will help your organization to ensure that all data stored in your servers in your data centers and at co-location partners or cloud service providers is protected and used appropriately. It will also give the organization guidance and structure on any changes needed and the specific use of such changes.

Additionally, using a well-known data protection framework may decrease the risk of data breach and exposure incidents, and regulators may require greater effort to protect the data in such cases. A data protection framework may also adapt to meet the evolving data protection requirements, while data protection laws may be subject to changes which are then captured in the framework. Data protection standards may help you and your organization to better manage your customer's data.

Data Protection Best Practices

There are different data protection management practices. Some of the most commonly used include:

- *Data loss prevention (DLP)*: this consists of a set of tools and processes used to secure data from theft, loss, misuse, deletion, or other illegal or inappropriate forms of contact, or modification, or exfiltration.
- *Firewalls*: Typically, software tools used for monitoring and filtering the network traffic to ensure data is transferred or accessed only by authorized users. Firewalls can block access by unauthorized entities.
- *Encryption*: Altering the content of data based on an algorithm that can be reversed only with the right encryption password or key. See Chapter 9 for details on algorithms. Encryption protects data even if it gets stolen, since the data would be unreadable without the key required to decrypt the data.

Enterprise Data Security for US Europe and Asia

- *Data erasure*: Deleting data that is no longer needed or relevant. Although no longer needed the data may contain PI, PII or sensitive data which, if exposed, can harm the organization. This is also a requirement of the GDPR.
- *Data resiliency*: Building resiliency systems within the software and hardware of an organization's system to ensure the security in case of natural disasters or power outages. Disaster Recovery in case of a catastrophic data center failure or a ransomware attack is of paramount importance for your organization.
- *Data backups*: A plan to securely back up data on a daily and weekly basis is essential to address loss of data in case of failure or breach of the data centers. Such backup plans may include a separate physical disk or cloud.

Chapter 6 – Who Is Responsible for Data Protection

As a general rule of thumb, you as the organization who is the owner of the data is responsible for ensuring that appropriate steps are taken to protect the data. You need to install in your organization the data protection strategy and an operational plan that is rigorously administered and governed to ensure compliance. You need to protect your confidential data by applying the appropriate security at all levels and through all departments within your organization.

Within your organization, typically, your organization's Chief Information Security Officer (CISO) is the corporate executive leader and the face of data security in your organization. The person in this role is responsible for creating the policies and strategies to secure data from threats and vulnerabilities, as well as devising the response plan if the worst happens resulting in data exposure or a ransomware demand.

Many organizations, in pursuit of GDPR compliance, also appoint a Data Protection Officer (DPO). The primary role of the DPO is to ensure that the organization processes the personal data of its staff, customers, and, if needed, providers of services or any other individuals (collectively referred to as data subjects) in compliance with the applicable data protection rules.

If your application is deployed in a third party hosting service provider (HSP) or a Cloud Service Provider (CSP), a good rule of thumb you need to follow is that the HSP or the CSP is responsible for providing a secure hosting environment (the computing platforms, the services, and network access) but you as the customer organization is responsible for all data stored within the platforms of the hosting organization. No matter whose fault causes a breach or data exposure, you as the data owner has the ultimate liability for the data breach and exposure; and consequently, will also bear the cost of the loss, damage to reputation, and any potential financial loss to your customer whose data has been exposed.

Everyone Has Responsibility for Data Security

Data is only as secure as its weakest link, therefore data security is the responsibility of both, your organization overall and, in particular, the individuals within your organization who have access to the data. A consistent and active effort is required by both parties to prevent any loss of data. █

Consequently, every employee in the organization has the responsibility for data protection. And, more specifically, the responsibility can be assigned by departments within the organization and the role of the employee.

Responsibility By Deployment Organization Type

As businesses move into a hybrid deployment model consisting of On-Premises data centers, Private Cloud, Community Cloud, and Public Cloud, the perceptions may not be consistent with the reality. In all cases, whether a disclosure (exposure or data breach) occurred on-premises or in the cloud, you the business (as the cloud consumer) have full responsibility, even if the data breach was due to the negligence of the Cloud Service Provider (CSP) for a public cloud deployment or the Hosting Service Provider (HSP) in case of a private cloud deployment.

The responsibility matrix in the table below illustrates the legal responsibility with respect to data exposure, notifications to your organization’s customers affected by the breach, costs associated with remediation, and penalties if imposed by regulatory agencies. While this matrix describes some aspects of legal responsibility, you as the data owner (and more particularly, executives such as eh CEO, COO and CISO of the organization that created and or acquired the data) always have responsibility and legal liability in the case of a data breach or exposure even though the exposure may have been the result of negligence by a private cloud hosting provider or a public cloud service provider or your own workforce member.

The table below delineates the roles and responsibilities:

Data Processor	Role	Legal Responsibility
The business on-premises IT	The IT Staff managing the entire physical network and software infrastructure within the organization. The organization must have defined policies and standards based on ISO, NIST, GDPR and other regulatory jurisdictions tailored according to the organization’s risk profile.	The organization has responsibility for the security of ALL data assets. In addition, the organization has the responsibility for notifying the data subjects within the time period days of the occurrence of the disclosure based on the legal requirements of

<p>Hosting Service Provider (HSP)</p>	<p>This is a shared responsibility structure because the hosting service provider has the network, storage and access control responsibilities.</p>	<p>the governing jurisdictions.</p> <p>However, The legal responsibility for the PII data rests always with your organization and not the HSP. The terms of the contract drive the responsibility for the HSP to ensure that they maintain a secure environment for the PII data.</p>
<p>Cloud Service Provider (CSP)</p>	<p>The deployment model (IaaS, PaaS, SaaS) determines the shared responsibility structure. The organization has more responsibility for data security under the IaaS model than it does under the PaaS model and the least responsibility under the SaaS model.</p>	<p>However, the ultimate legal responsibility for the PI or PII data rests always with the organization and not the CSP. The CSP typically has a menu for security operations. The terms of the contract drive the responsibility for the CSP to ensure that they maintain a secure environment for the PII data. However, your organization may not have much leverage on the CSP.</p>
<p>Business Partners</p>	<p>Business partners who may use the PII data of your organization must not retain it or ensure that the data is maintained at the same level of security as your organization. Your organization must contractually bind the business partner to ensure data privacy and security.</p>	<p>Despite the contract, the legal and notification responsibility for the PII data rests with your organization always, and not the business partner. The terms of the contract drive the responsibility for the business partner to ensure that they maintain a secure environment for the PII data and the</p>

	damages sharing they are responsible for in case of a breach.
--	---

Table 6-1: Data Management Roles and Responsibilities

Responsibility By Role Played by the Employee

Roles play a very important part in the determination of responsibilities of the various entities involved in the creation, storage, and use of data within and outside of the business.

Ownership, custody, rights, responsibilities, and liability are all relative to the dataset in question, and therefore are only specific to that data in that circumstance. For instance, a cloud provider is usually the data processor for a cloud customer’s data, but the cloud customer, that is your organization, is the data owner for the information that you store in the provider’s platforms while the service provider is responsible for data that the provider collects and creates, such as the provider’s own customer list, asset inventory, and billing information.

The table below describes the various organization roles associated with creating, organizing and administering access to the data created within the organization or acquired by the organization from external sources.

Role Name	Role Description	Comments	When Responsibility ends
Data Owner	An organization or individual who has collected or created the data asset and is accountable for it. This is typically an executive role that goes to the department, team or business unit that owns a data asset. The following are examples of	Legally responsible for all data they own. This is true even if data is compromised by a data custodian several times removed from the data owner.	Responsibility ends when the data has been destroyed from all locations where it was deployed.

Enterprise Data Security for US Europe and Asia

Data Controller	responsibilities associated with the data owner role	Identified in the Create phase.	
	A data controller is a person, company, or other body that determines the purpose and means of processing personal data (this can be determined alone, or jointly with another person/company/body).	The data owner determines the classification of the data that data controller manages.	Responsibility ends when the data is no longer required and the data under its control has been destroyed from all locations where it was deployed.
Data Steward	Responsible for data content, context, and associated business rules. Responsible for utilizing the organization's data governance processes to ensure fitness of data elements - both the content and metadata.	The data steward is responsible for ensuring the quality of data throughout its existence.	The role ends when the data has been destroyed from all locations where it was deployed.
Data Custodian	Data custodian is usually the organization or person responsible for, or the person with administrative control over, granting access to an organization's documents or electronic files while protecting the data as defined by the organization's security policy or its standard IT practice. On-prem it is the DBA. In the cloud, it is the CSP.	Do not necessarily all have direct relationships with data owners; custodians can be third parties, or even further removed down the supply chain.	The role ends when the data has been destroyed from all locations where it was deployed.
Data Processor	A data processor is the entity who	Processes personal data on behalf of a	Responsibility ends when role

Enterprise Data Security for US Europe and Asia

	<p>processes data on behalf of a data controller. The data controller decides the purpose and manner to be followed to process the data, while data processors holds and process data.</p>	<p>data controller but does not have any responsibility or control over that data.</p>	<p>is no longer required and data under its control has been destroyed from all locations where it was deployed by the data processor.</p>
<p>Data User</p>	<p>Data users can be any person or entity with a personal or business purpose to view and use the data including researchers, academics working in research institutions and employees in State/Territory government agencies.</p>	<p>Data users have a responsibility to manage the data with appropriate data security controls depending on the nature of the data and the terms of use.</p>	<p>Responsibility ends when the use purpose is no longer applicable and the data user has destroyed all copies of the data under their control.</p>
<p>Data Aggregator</p>	<p>Data aggregators are data mining systems that spread business information online. They collect and share business data with a multitude of sources, including search engines like Google.</p>	<p>Should maintain security of the data in their possession, especially data obtained under the agreements with the data providers or owners.</p>	<p>When responsibility ends is unclear because many data aggregators hold on to data for long unspecified durations.</p>
<p>Data Subject</p>	<p>“identified or identifiable natural person[s].” In other words, data subjects are just people—human beings from whom or about whom you collect information in connection with your business and its operations</p>	<p>Depending on geographical jurisdictions, data subjects may have rights to privacy, know how the data is being used and the right for the data to be erased under the regulatory laws</p>	<p>Data subjects have a perpetual right to their own data and their own identity.</p>

	and standards of that jurisdiction.
--	-------------------------------------

Figure 6-2: Data Management Responsibility by employee Role

Data Owner: The data owner is the person who has ultimate organizational responsibility for data. The owner is typically the CEO, president, or a department head (DH). Data owners identify the classification of data and ensure that it is labeled properly. They also ensure it has adequate security controls based on the classification and the organization's security policy requirements. Owners may be liable for negligence if they fail to perform due diligence in establishing and enforcing security policies to protect and sustain sensitive data.

Every data field in every database in the organization should be *owned* by a *data owner*, who is in the authority to ultimately decide on the access to, and usage of, the data. The data owner could be the original producer of the data, one of its consumers, or a third party. The data owner should be able to fill in or update its value which implies that the data owner has knowledge about the meaning of the field and has access to the current correct value (e.g. by contacting a customer, by looking into a file, etc.). Data owners can be requested by data stewards to check or complete the value of a field, as such correcting a data quality issue.

A Data Owner has administrative control and has been officially designated as accountable for a specific information asset dataset. This is usually the senior most officer in a division. Some examples of Data Owners include the Registrar for student data; the Treasurer for financial data; the VP of Human Resources for employee data. But in all cases, the responsibility for data security rests with the CEO.

System Owners: The system owner is the person who owns the system that processes sensitive data and develops a system security plan in coordination with information owners, the system administrator, and functional end users. The System Owner (SO) develops and maintains the system security plan and ensures that the system is deployed and operated according to the agreed-upon security requirements. The SO ensures that system users and support personnel receive appropriate security training, and assists in the identification, implementation, and assessment of the common security controls. The SO is typically the same entity as the data owner, but it can sometimes be someone different, such as a different department head (DH).

Business/ Mission Owners: The business or mission owner is a role that is viewed differently in different organizations. The responsibilities of the business or mission

Enterprise Data Security for US Europe and Asia

owner can overlap with the responsibilities of the system owner or be the same role. Business owners might own processes that use systems managed by other entities. As an example, the sales department could be the business owner but the IT department and the software development department could be the system owners for systems used in sales processes.

Data Steward: The data stewards are the Data Quality (DQ) experts in charge of ensuring the quality of both the actual business data and the corresponding metadata. They assess DQ by performing extensive and regular data quality checks. These checks involve, amongst other evaluation steps, the application or calculation of data quality indicators and metrics for the most relevant DQ dimensions. Clearly, they are also in charge of taking initiative and to further act upon the results of these assessments. A first type of action to be taken is the application of *corrective measures*. However, data stewards are not in charge of correcting data themselves, as this is typically the responsibility of the data owner. The second type of action to be taken upon the results of the data quality assessment involves a deeper investigation into the *root causes* of the data quality issues that were detected. Understanding these causes may allow designing *preventive measures* that aim at eradicating data quality problems. Preventive measures may include modifications to the operational information systems where the data originate from (e.g., making fields mandatory, providing drop-down lists of possible values, rationalizing the interface, etc.). In addition, values entered into the system may immediately be checked for validity against predefined integrity rules and the user may be requested to correct the data if these rules are violated. For instance, a corporate tax portal may require employees to be identified based upon their social security number, which can be checked in real-time by contacting the social security number database. Implementing such preventive measures obviously requires the close involvement of the IT department in charge of the application. Overall, preventing erroneous data from entering the system is often more cost-efficient than correcting errors afterwards. However, care should be taken not to slow down critical processes because of non-essential data quality issues in the input data.

Data Processors: A data processor is any system used to process data. However, in the context of the EU Data Protection law, data processor has a more specific meaning. The EU Data Protection law defines a data processor as “a natural or legal person which processes personal data solely on behalf of the data controller.” In this context, the data controller is the person or entity that controls processing of the data.

Custodians: Data owners often delegate day-to-day tasks to a custodian. A custodian helps protect the integrity and security of data by ensuring it is properly

stored and protected. For example, custodians would ensure the data is backed up in accordance with a backup policy. If administrators have configured auditing on the data, custodians would also maintain these logs. In practice, personnel within an IT department or system security administrators would typically be the custodians.

Can you be a processor of some data and a controller of other data at the same time?

Yes. Many companies that are data processors of some personal data are also data controllers of other personal data. The concept of whether you are a controller or processor is based on your processing actions as to a particular type of personal data, not to your company as a whole. For example, your business could be a processor of your customers' data, but a data controller when it comes to your own employees' data.

Generally, businesses are going to be data controllers of their own employees' personal data, used for human resources operations, as well as their own customer relationship data that they use for customer relationship management and support functions. It is harder to generalize about when businesses function as a data processor. Some organizations that process personal data may only be controllers and never act as data processors.

Access Controls Based on Responsibility

Four types of Access Controls are used in organizations:

1. Access Control List (ACL) is a table listing the permissions attached to computing resources. It tells the operating system which users can access an object, and which actions they can carry out. There is an entry for each user, which is linked to the security attributes of each object. ACL is commonly used for traditional DAC systems.
2. Attribute-Based Access Control (ABAC): ABAC evaluates a set of rules and policies to manage access rights according to specific attributes, such as environmental, system, object, or user information. It applies Boolean logic to grant or deny access to users based on a complex evaluation of atomic or set-valued attributes and the relationship between them.
3. Role-based access control (RBAC), also known as role-based security, is a mechanism that restricts system access. It involves setting permissions and

Enterprise Data Security for US Europe and Asia

privileges to enable access to authorized users based on their roles in the organization.

4. Discretionary Access Control (DAC): The owner of a protected system or resource sets policies defining who can access it. DAC can involve physical or digital measures, and is less restrictive than other access control systems, as it offers individuals complete control over the resources they own.
5. Mandatory Access Control (MAC): A central authority regulates access rights based on multiple levels of security. MAC involves assigning classifications to system resources and the security kernel or operating system. Only users or devices with the required information security clearance can access protected resources.

Most large organizations use role-based access controls to provide their employees with varying levels of access based on their job roles and responsibilities.

Employees are only allowed to access the information necessary to effectively perform their job duties. This protects sensitive data and ensures employees can only access information and perform actions they need to do their jobs. RBAC is driven by business need so understanding your business needs is important for implementing RBAC. Define the roles based on your analysis of the business needs accounting for the job and responsibility for individuals and groups of individuals. Start with coarse-grained access controls and then expand this with fine grain controls defining the read, write and delete access controls. Access should be allowed only to those who need it to perform their jobs and should not be provided to all employees irrespective of roles or job functions.

Access controls must be re-assessed and modified when an employees' job role changes. Similarly, their membership in access control groups needs to be changed accordingly.

Access Controls have to be defined for each data asset and its associated hosting IT asset that you need to control access to because the data asset has sensitive data such as PI or PII.

Chapter 7 - Data Storage and Security Architecture

Modern datacenter architectures are designed to provide the agility, availability, and security that workloads require to support an organization's digital objectives amid the rise in cybersecurity attacks. Flexible software-defined networking (SDN) infrastructure is an approach being used increasingly to meet emerging edge computing, distributed data, and mobile/hybrid work strategies.

At the same time, modern application architectures and increased data volumes require better control and visibility at the network edge, where the network and compute merge, to optimize resource usage efficiently. This includes managing CPU, memory, network fabric, inline firewalls, load balancers (ADCs), racks, and power.

The speed and volume of traffic in virtualized and containerized application environments require new security requirements, especially in multitenant scenarios where zero trust security is imperative. Automated and policy-based network segmentation and micro-segmentation are necessary, and centralized security appliances are inefficient for expanding traffic flows.

To address these challenges, a simpler, scalable, and consistent approach to network automation is necessary. Such an approach should be aligned with cloud agility and streamlined operations, allowing organizations to unify, automate, and secure overlay and underlay networking across their distributed application landscape. This unified approach can significantly reduce operational costs and increase business agility by ideally providing a common and unified network fabric across switches and hosts that support heterogeneous infrastructure comprising multiple hypervisors, Kubernetes, and bare metal workloads. In addition, this fabric should encompass distributed security services, eliminating the inefficient and costly clutter of appliance sprawl.

From a data security perspective, managing the storage locations, and more importantly, the security of the storage locations where sensitive data is stored, is of utmost importance. Zero trust strategies based on segmentation and micro-segmentation play a crucial role in such secure storage management. This chapter focuses on the data storage requirements and storage architectures required to support the data storage requirements.

This chapter also addresses the strategies for secure data storage management strategies across hybrid on-premises and cloud environments.

Storage Security Management

Storage security management is the process of ensuring an organization's storage systems and its data are fully protected in accordance with the organization's security requirements. This includes data that resides within the storage systems, as well as data in transit to and from those systems.

Ensuring data Confidentiality

To protect important data from loss or inappropriate disclosure, your organization needs to follow these approaches and best practices for effective ways to ensure the confidentiality of sensitive data in your organization:

- Restrict access to data strictly on a need to know and need-to-use basis.
- Implement a confidentiality policy that is actively socialized and rigorously applied and managed.
- Implement a data retention policy to remove unused or unneeded confidential data to avoid inadvertent exposure. Delete sensitive data you no longer need. Storage of unused or unneeded sensitive data increases chances of inadvertent exposure.
- Take physical security measures to protect all data in data centers and access controls for all confidential data on-premises and in the cloud environments.
- Develop and implement a cybersecurity program and assign a team for the governance of the program.
- Protect all confidential and sensitive data through binding non-disclosure agreements.
- Protect data at rest. Encrypt your sensitive data at rest. Enable full disk encryption on all devices.
- Restrict confidential data to the office. Employees should not be allowed to carry confidential or sensitive data on their mobile devices.
- Protect data in transit. Don't transfer unencrypted data over the Internet. Man-in-the-middle attacks can compromise data in transit. This implies that any transfers should enforce TLS 1.3 encryption at the transport layer.
- Encrypt backups. All confidential or sensitive data at rest, including backups should be encrypted.

- Store more than one copy in a remote location such that ransomware lockout attacks can be thwarted and data protected.

Types of Storage

Data can be stored in a variety of different data storage mechanisms such as block storage, NoSQL databases and relational data bases.

Databases are typically a durable and reliable type of data store. Anything that you need to store permanently can go in a database.

For structured data relational databases continue to be an industry standard tool for the reliable storage of important data.

Generally, Cloud storage is safe from hackers when properly secured. Reputable cloud service providers invest heavily in security measures like encryption, access controls, and monitoring. However, no system is completely immune to threats. Confidential and sensitive data stored with the cloud environments must be rigorously protected through access controls and encryption.

Best Practices for Successful Data Management

There are several aspects for managing data. Note the following best practices that you need to adopt for your organization:

- Build strong file naming and cataloging conventions so that the confidential and sensitive data can be easily identified and protected.
- Carefully consider metadata for data sets that describes the purpose and security requirements for the data.
- Data Storage should be carefully considered to ensure that confidential and sensitive data is stored in secure storage as needed.
- Commitment to a data security culture across your entire organization is a necessary step for protecting your confidential and sensitive data.
- Ensure that data quality can be trusted at all times through security and privacy that is comprehensive and ensures integrity of data.
- Invest in quality data-management software and best practices for tracking systems of record and systems of reference for all data.
- Implement the concepts of Systems of Record and Systems of Reference for data to ensure integrity of the data.

Let us explore further what we mean by Systems of Record and Systems of Reference.

Systems of Record and Reference

Your organization's data drives your business and is constantly being transferred to downstream applications and systems for use and is stored back after modifications in those downstream applications. To ensure durable integrity of the data at all times, it is important to define a system of record which is kept secure and fully up to date with all authorized changes.

Any downstream applications should be able to obtain a copy of the data from the System of Record and maintain a System of Reference for its own use and the use of its downstream applications and third-party client applications. All data modified in a System of Reference should be written back to the System of Record with a timestamp for the change. A third type of system is the decentralized System of Engagement used for certain types of social media and peer to peer applications. The table below explores these different systems and the role they play in greater detail.

System Type	Used As	Data at Rest Protections
System of Record	An information storage and retrieval system that provides a centralized, authoritative source of data elements in an IT environment containing multiple points of data generation.	As the centralized authoritative data source, it is imperative that the integrity of this data is totally protected from unauthorized alteration when the data is modified in a system of reference and written back. Only authorized transactions should be permitted to update the system of record.
System of Reference	A system of reference for an entity is the authoritative system that is expected to contain correct, complete and current data about an entity which can be used as the source of truth for transactional, operational or analytical purposes.	The system of reference used typically by downstream applications captures the protected form, and stores it. When a user searches for the records containing the data, if the user is authorized to see the plaintext representation of the data and has a business need to, the downstream system of reference unprotects the data by calling an

		API exposed by the system of record.
System of Engagement	Systems of engagement are decentralized IT components that incorporate technologies such as social media and the cloud to encourage and enable peer interaction.	Data at rest in the systems of engagement should always reflect the exact replication of the systems of record except to the extent it is modified in use and temporarily stored back.
Managed third party data transfers	An application or utility exports the protected form of the data from the system of record or the system of reference to a downstream third-party application	The third-party application may use or store the data or pass it on to other applications or users. The data owner has to specify the protections and security capabilities that must be maintained by the third-party applications and may insist on a contract for it. This may apply to business partners who may require PII data in plaintext form.

Table 7-1: Roles of Systems in Data Management

Managing System of Record (SoRec) vs. Reference (SoRef)

As we said a System of Record is the authoritative master data store for your organization where data is created, captured and stored. For example, for conducting your business, you have a system being used by the sales team which uses customer data to track sales lead tracking and order activities, you have an order management system to track orders, an order fulfillment system to track order delivery, a billing system to rack customer billing and payments. All of these systems use customer data that maybe changed in various ways by the different systems. The System of Record is the common authoritative source of data for all aspects of the customer including confidential and sensitive data in addition to the order management and accounting data. A common protected database is treated as the System of Record. The System of Record is, typically, the source system of record and also, typically, the system of origin or initial entry.

A system can be designated as the System of Reference as the system that is the authoritative complete and current data based on certain considerations or efforts

that confirm the reliability of the data. For the customer data example above, you may define a finance-team system as the system of reference because it contains the most accurate and current data about customers; or we may create a new system which contains the best current version of the customer data from the finance and sales teams' systems. The System of Reference is supposed to be the source of truth and typically, it is the single source of truth.

The data from the System of Reference is written back to the System of Record to ensure integrity so that all downstream applications are synchronized with the changes in the System of Reference.

Data Source System (SysSrc)

A data source is the initial location where data is born or where physical information is first digitized. However, even captured and highly processed data may serve as a data source, as long as another process utilizes it. A data source may be a database, a flat file, live measurements from physical devices, scraped web data, or any of the myriad static streaming data captured from the internet.

Databases remain the most common data sources, as the primary stores for data in ubiquitous relational database management systems (RDBMS). Data sources can be from machines or from files depending on where the data is being created.

Operational Data Store (ODS)

An operational data store (ODS) is a central database that provides a snapshot of the latest data captured and consolidated from multiple transactional systems for operational reporting. It enables your organizations to combine data in its original format from various sources into a single destination to make it available for business reporting. It is often, created as a data warehouse.

Storage Management

Now that we understand the importance of Systems of Record, Systems of Reference and Systems of Engagement, let us look at the data storage types and how they impact the level of security you can achieve with each data storage type and the best practices for using them.

Data Storage Types

Data storage is the retention of information using technology specifically developed to keep that data and have it as accessible as necessary. Data storage refers to the use of recording media to retain data using computers or other devices.

The most prevalent forms of data storage are file storage, block storage, and object storage, with each being ideal for different purposes. In file storage inexpensive and simply constructed data is stored in files and folders. This is commonly found on hard drives and means that the files look exactly the same to the hard drive as they do to the user. In object storage, Data is stored as objects with metadata and unique identifiers. Although it is generally less expensive to store data this way, object storage is only ideal for data that doesn't need to be edited. In block storage, Data is stored in evenly-sized blocks. Although more expensive and complex and less scalable, block storage is ideal for data that must be frequently accessed and edited.

Data can be stored and organized in a variety of different data storage mechanisms according to the data structures such as block storage, NoSQL databases and relational data bases.

Definitions of a Database

A database is an organized collection of structured information, or data, typically stored electronically in a computer system. Databases are used to store and manage large amounts of structured and unstructured data, and they can be used to support a wide range of activities, including data storage, data analysis, and data management.

Embedded Database: If data access is exclusive to that application and in effect hidden from end-users or other applications, it can be called embedded. This represents a single database tier regardless of the number of databases utilized.

Shared Data Tier: In a shared data tier, a single data tier supplies data to several applications. It may consist of multiple databases each dedicated to an application or databases that supply data to multiple applications. In this case, each shared database entity such as table, view, index, stored procedure, trigger, etc. will need to be identified to be accessible to multiple applications.

Database Management Systems (DBMS)

In addition to the data within the database, a DBMS has management software that manage data stores or files and centrally maintain several databases spread across

multiple computer systems (or a cluster). The DBMS allocates storage space, access control, and data replication. A DBMS also provides functionality for organizing the data in tables and managing view and index structures as well as stored procedures and triggers.

Database as a Service (DBaaS)

Databases maintained on a single hardware/software environment shared by many applications without their having to manage them are participating in a DBaaS. The databases are maintained and administered by the database (DB) Administration staff on behalf of the applications teams sharing the DBaaS. The DB Administrators have to know and document the association of the data objects with the applications. In a DBaaS the central DB Administrators are responsible for instantiating, populating, backing up and managing access controls for the DBMSs within their control in the DBaaS.

Data Architectures

Data architecture describes how data is managed--from collection through to transformation, distribution, and consumption. It sets the blueprint for data and the way it flows through data storage systems. It is foundational to data processing operations. It includes defining the flow of data across various business applications and the infrastructure used for these business applications.

With the evolution of migrating on-premises applications to public clouds and the development of cloud-native applications, migration of data from on-premises to the clouds and use within the clouds is also an important architectural consideration.

The following are important considerations for any data architecture effort that your organization embarks on:

6. Data is a shared asset. A modern data architecture needs to eliminate departmental data silos and give stakeholders a view of the data they need to access.
7. Provide adequate interfaces for users to access their data. The data architectures need to provide interfaces that make it easy for users to consume the data they need using tools appropriate for their jobs.
8. Security of the data is essential. Modern data architectures must be designed for data security, and they must support data policies and access controls directly on the raw data.

9. Common vocabularies ensure common understanding. Shared data assets require a common vocabulary to allow all applications to access and use the data.
10. Data flows should be optimized for speed and integrity. Your organization should architect your data migration strategy to reduce the number of times data must be transformed, and to ensure that the data is current. The data architecture approach should optimize enterprise data management.

Types of Data Architecture

In the modern IT architecture, it is important to understand how your business collects, uses, and disposes of data. From a storage perspective, three types of data architecture stand out:

1. Application Data – Data stored within a data store managed by the application; this is typically used only by that application.
2. Data Warehouses – A compilation of data from several related applications such as customer data, order entry data, billing and payment data, and so on. This compiled data allows correlation and advanced reporting using artificial intelligence (AI) techniques. It acts as the single version of truth for all applications and can be stored for the duration of the required data retention period.
3. Data Lakes – These can be thought of as a high-volume storage of all current and historical data, usually captured from the data warehouse(s) and retained for long periods of time in lower cost storage.

Your organization-wide data architecture strategy needs to consider the strategies and tools used to migrate the data from one type of data to the next and when such migrations happen along with any data transformation for the purposes of the new type of data storage.

Domain and Cross-platform Context

Data domains are high-level categories of data for the purpose of assigning accountability and responsibility for the data. Some examples of data domains include:

- Customer
- Product or Service
- Location or address
- Vendor or supplier

Enterprise Data Security for US Europe and Asia

- Transaction or order
- Billing and accounting
- Legal

Different business areas can have a different set of domains for data relevant to their business. The domains can have sub-domains. For example, the Customer domain can have sub-domains such as: individual, business, government, and so on.

Data domains determine the strategies and policies that must be applied in the data architecture such as access controls, data security, data retention, and so on.

Enterprise Data Model

An enterprise data model is a visual representation, or graph, of your organization's business data. Enterprise data modeling is a process for conceptualizing the relationships between different types of information in an organization. Enterprise data models help users across disciplines store and interact with data more effectively for a variety of use cases.

The enterprise data model can be organized into three layers of abstraction: the subject area model, the enterprise conceptual data model, and the enterprise logical data model. The subject area model is simply a list or hierarchy of the subject areas within the enterprise data model. The conceptual data model defines the highest levels of relationships between the entities used in your business. The enterprise data model addresses the unique requirements of your organization's business. This model provides a detailed understanding of how your business operates and the data that drives the operations. This model is a high-level vision of your organization driving your data architecture.

Database models

Database Models define how your data is organized based on its use. The common data models used in businesses include:

- *Hierarchical* model organizes data into a tree-like structure, where each record has a single parent or root. Sibling records are sorted in a particular order. That order is used as the physical order for storing the database.
- *Relational* model as in relational database management systems which sorts data into tables, also known as relations, each of which consists of columns and rows. Each column lists an attribute of the entity in question.
- *Network* model builds on the hierarchical model by allowing many-to-many relationships between linked records, implying multiple parent records. Based

on mathematical set theory, the model is constructed with sets of related records

- *Object-oriented* model combines the simplicity of the relational model with some of the advanced functionality of the object-oriented database model. In essence, it allows designers to incorporate objects into the familiar table structure.
- *Entity-relationship* model captures the relationships between real-world entities much like the network model, but it isn't as directly tied to the physical structure of the database. Instead, it's often used for designing a database conceptually.
- *NoSQL* database models include graph database and unstructured document models.

Please refer to a database management book to under how to use these models. The key point here is that each of these models has its own data security requirement.

Metadata

Metadata (or metainformation) is data that provides information about other data, but not the content of the data itself, such as the text of a message or the image itself. Metadata describes how, when, and by whom a particular set of data was collected, and how the data is formatted. Metadata is created when files are created and when edited. This information can contain revisions, comments, template information, file properties and summary information.

This can be useful for cyber forensics in data security breach cases because it provides information that may not be immediately clear from the file itself, such as when and how it was modified or accessed by an attacker.

It is important to keep metadata secure because unprotected metadata can be revealed to the wrong people (e.g., hackers, cybercriminals, or malicious competitors). These adversaries may then take advantage of the metadata to steal further data and also aggregate information from the metadata about your business that can be used in a manner detrimental to your organization.

Master Data

Master data is the consistent and uniform set of identifiers and extended attributes that describe the data about the business entities that provide context for business transactions. The most commonly found categories of master data are parties (customers, vendors, individuals), products, and financial structures and locational (address) concepts. Master data acts as a foundational reference for various business processes, enabling companies to make informed decisions, and especially

in our context, data security, and maintain consistency across different systems and departments.

Business Intelligence

Business intelligence greatly enhances how a company approaches its decision-making by using data to answer questions of the company's past and present. It can be used by teams across an organization to track key metrics and organize on goals. Business intelligence refers to the processes used to generate data reports. In contrast, business analytics refers to the processes used to apply data-driven insights to make decisions or take required actions.

Data Integration Architecture

Your organization has multiple data sources, so your data architecture needs to be designed to aggregate data from various sources like websites, applications, social media, and business databases to make data easily accessible. Data integration architecture is based on technologies that can transform the data and transport datasets to their target locations. All these ingestion and transformation processes involve data of various sizes, structures, and types, thereby dealing with transformational complexity. A data integration architecture aims to solve the heterogeneity feature from various data sources, locations, and interfaces.

Data integration architecture changes with the advances in cross-platform utility and other development paradigms for new kinds of digital operations. Integration software breaks down data application silos and enables various software applications to communicate with each other.

In order to connect the different applications with each other, application programming interfaces (APIs) are used, which are specially designed to enable this kind of integration. Integration Architecture enables you to process these applications and data objects within your information technology architecture while providing and using interfaces.

The diagram below shows a highly simplified snapshot of data integration architecture.

Enterprise Data Security for US Europe and Asia

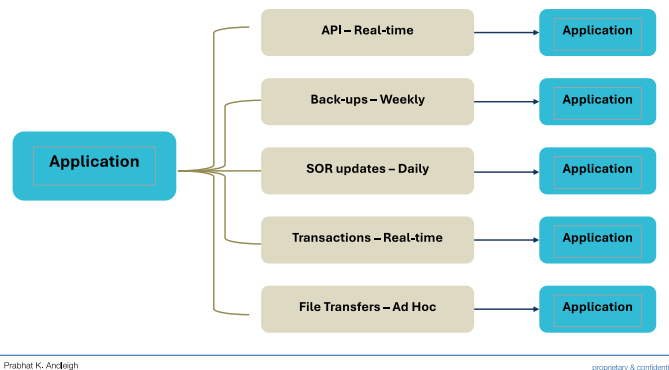


Figure 7-1: Data Integration Architecture

As part of the design process it is important to Ensure checks for data quality. Observability features should be essential to your integration architecture. Data from disparate sources usually contain anomalies like null values or duplicate references.

It is equally important to maintain consistency during integration. Consistency helps prevent confusion and creates a single source of truth for data usage, which makes collaboration between teams easier. For example, maintaining the formats of customer information as it flows between groups will help prevent confusing scenarios later.

And finally, documenting the integration process helps standardize your processes and makes identifying the cause of errors easier. Additionally, if proper documentation follows every cycle, it becomes easier to maintain consistency and spot useless data.

Data Lake and Data Lakehouse

A data lake is a centralized repository designed to store, process, and secure large amounts of structured, semi-structured, and unstructured data. It can store data in its native format and process any variety of it, ignoring size limits. A data lake is a central location that holds a large amount of data in its native, raw format. Compared to a hierarchical data warehouse, which stores data in files or folders. A data lake is typically used for analytics processing. A data lake uses a flat architecture and object storage to store the data.

The Data Lakehouse concept has captured the hopes of modern enterprises that seek to combine the best of the data warehouse with the best of the data lake. Like a data warehouse, it transforms and queries data at high speed. Like a data lake, it consolidates multi-structured data in flexible object stores. Together these elements can support both business intelligence (BI) and data science workloads.

While a data lake holds data of all structure types, including raw and unprocessed data, a data warehouse stores data that has been treated and transformed with a specific purpose in mind, which can then be used to source analytic or operational reporting.

Many organizations have implemented data lakehouses to streamline their architectures, reduce cost, and assist the governance of self-service analytics. Common use cases include data mesh support, a unified access layer for analytics, data warehouse consolidation, data modernization for the hybrid cloud, and departmental lakehouses.

Data Management Strategy for Hybrid Cloud

Architectural considerations for dealing with a hybrid database management system (DBMS) cloud environment are neither inherently obvious nor consistent, which has financial and performance implications for data and analytics leaders.

Data and analytics leaders are challenged to understand how data flows, both in volume and direction, and how data location impacts performance, application latency, SLAs, high availability and disaster recovery (HA/DR) strategies, and financial models in hybrid DBMS cloud scenarios.

Recommendations for Hybrid Cloud Implementation

Cloud provider lock-in, makes migration of data to another provider difficult and costly. most data and analytics use cases will require connecting to distributed data sources, leading enterprises to double their investments in metadata management. It is important to have a well thought-out strategy for implementing hybrid cloud environments. The following key considerations are critical for this decision:

Use Case Alignment. Select hybrid cloud deployment architectures that align with specific use-case requirements such as compatibility with on-premises options (if available), network latency and throughput requirements.

Data Flow based colocation. Guide your decision making on colocation of application components for hybrid cloud architectures by monitoring data flow, in

both volume and direction, between these components. Volume will impact the expected latency and performance, while direction will impact the cost of a hybrid deployment.

Compelling reason. Adopt multi-cloud architectures when there is a compelling reason to do so. These may include vendor lock-in risk mitigation, application availability or specific services that are not available in your primary cloud provider's environment. Make sure the payoff is worth the overhead!

Use Case Alignment

With respect to DBMS deployments, the following four core scenarios of hybrid cloud deployments can be defined as the driving considerations:

- Architecture spanning
- Use-case specific
- Multi-cloud
- Intercloud

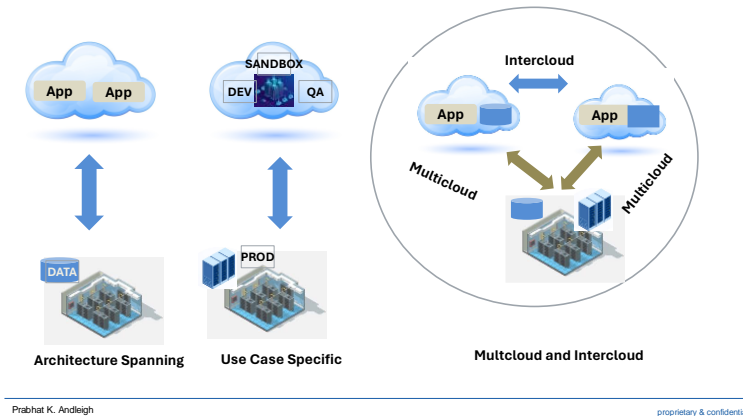


Figure 7-2: Hybrid Cloud Deployment Architecture

Architecture Spanning

Different components of an application architecture may reside on-premises and/or in the cloud. The DBMS might reside on-premises and the applications that connect to it may reside in the cloud — for example, a business intelligence (BI) dashboard application. This also includes architectures with data residing both in the cloud and

Enterprise Data Security for US Europe and Asia

on-premises, such as the ability of the DBMS to have some replicas, partitions or shards residing on-premises and some in the cloud for the same database.

Data can be partitioned by age, frequency of access, or geography. Demand patterns and data request surges also play an important part.

It is critical to understand the characteristics around the flow of data (for example, whether data is flowing into the cloud or out of the cloud) and the expected volumes of data. There may be issues with latency — that is, the time it takes to move the data between on-premises and cloud. SLAs should be defined and tested. This may lead to a requirement of a special communications link between the on-premises and cloud components, leading to greater financial cost implications.

Use-case Specific

Components are segmented by their development life cycle function. For example, any of the development, test, quality assurance (QA), disaster recovery (DR), or production instances of a DBMS may reside on-premises or in the cloud. Although financial and latency considerations remain important, in this scenario, compatibility is the primary concern.

There is a difference in API compatibility, which will impact how the application functions, and deployment environment compatibility, which will always differ from cloud PaaS to on-premises.

Multicloud

In the case of multicloud each cloud runs an application independently and the application is not spread across the clouds. Services from multiple cloud providers are used. A DBMS offering may be deployed on-premises and on one or more clouds. As such, all of the considerations of use-case-specific hybrid cloud apply with the added considerations of deploying software in multiple cloud environments. It is very common where an organization uses a public cloud for customer application deployment as well as a private cloud for internal application deployment.

The multicloud scenario generally appeals to those end users who are concerned about cloud vendor lock-in and want to be able to move their applications easily to a different cloud provider, or even repatriate them to on-premises deployments. It can also be brought into use due to an acquisition of another organization that is using a different cloud environment.

Note that multi-cloud environments can play an important role in your business continuity and disaster recovery planning.

Intercloud

In the case of a intercloud, an application runs across multiple clouds which significantly adds to complexity. Not only are services from multiple cloud providers used, but there is also intercloud communication where data is moving almost in real-time between two different clouds. A new complexity added in intercloud use is access control and identity management which may be more difficult to implement in a unified manner, and may require pushing this functionality down to the DBMS or application level rather than relying on cloud provider services. Data flow management is another potential complication.

Co-locating Application Components by Monitoring Data Flow

In any enterprise moving to the cloud, more than one or all of these scenarios may exist simultaneously. As the cloud footprint expands to encompass the different hybrid DBMS cloud scenarios, information and application architects will need to closely monitor and understand the implications of each, and how they affect application performance and latency, specific use cases like HA/DR, and even the impact of cloud brokering services for financial arbitrage and their impact on data flow. Data needs to have a natural home — either in the cloud or on-premises. Referred to as “data gravity”, this will impact the design of the hybrid cloud architecture. All of the concerns associated with the architecture spanning, use-case-specific, multicloud and intercloud scenarios apply in the design of dbPaaS (database platform as a service) and data architecture.

Considerations for Multicloud/Intercloud Architectures Usage

When adopting a multicloud or intercloud hybrid cloud architecture, compatibility friction may occur for auxiliary services as different cloud providers have a different range of services available, with differing degrees of functionality. There will be different capabilities for monitoring, provisioning, payment and financial governance. Different cloud infrastructure may present different performance.

It is important to run a proof of concept (POC) not only to prove the compatibility of the capabilities and tooling, but also to show that the desired SLAs are met.

Chapter 8 – Using Encryption Technologies

Data encryption has been used in some form or another since ancient times to communicate with armed forces and for espionage. Ciphers were developed for that purpose. A well-known example in the mid-20th century is the use of the Enigma machine by the German military to transmit and receive coded messages. The Enigma's encryption key changed every day, making the messages hard to crack.

Claude E. Shannon, who worked for several years at Bell Labs during the 1940s is considered to be a major contributor to mathematical cryptography.

In the early 1970s: IBM formed a 'crypto group,' which designed a block cypher to protect its customers' data. In 1973, the US NIST (National Bureau of Standards) adopted it as a national standard - the Data Encryption Standard, or DES; its draft was published in the U.S. Federal Register on 17 March 1975. After advice and modification by the National Security Agency (NSA), acting behind the scenes, it was adopted and published as a Federal Information Processing Standard (FIPS) and published in 1977. FIPS 46-3 is a recent version of it. DES remained in use until it was cracked in 1997.

DES was officially replaced by the Advanced Encryption Standard (AES) in 2001 when NIST announced FIPS 197. After an open competition, NIST selected Rijndael, submitted by two Belgian cryptographers, to be the AES. DES, and more secure variants of it are still in current use due to its incorporation in many national and organizational standards where it provides an acceptable level of security.

Data encryption algorithms scramble plaintext so that only the holder of the decryption key can read it. This process provides data security for personal information that users receive, send, and store. Data encryption works by securing digital data on computer systems in data centers, in cloud storage and in transmission across the Internet.

We have seen in a previous chapter that cybersecurity refers to the implementation of techniques and procedures which helps to protect the data from unauthorized entities. Cryptography, on the other hand, refers to the encryption and decryption of coded language so that only the sender and receiver can decipher that.

Encryption works with cybersecurity to defend against brute-force and cyber-attacks, including malware and ransomware. Data encryption works by securing stored data on-premises or in cloud environments as well as during storage and during data transmission among end-user devices, on-premises data centers and

cloud environments. So, cryptography helps protect private information, sensitive data, and can enhance the security of communication.

Most data transmitted over a network even now is sent in clear text making it easy for unwanted persons to capture and read sensitive information. Encryption algorithms protect data from intruders and make sure that only the intended recipient can decrypt and view and use the data.

Encryption is commonly used to protect data in transit and data at rest. Every time someone uses an ATM or buys something online with a smartphone, encryption is used to protect the information being exchanged between the user and the bank's systems. Financial transactions and private messaging communications should and most often use encryption to increase security.

Modern cryptography and the encryptions are achieved by using algorithms that have a key to encrypt and decrypt information. These keys convert the messages through encryption to make them unreadable and then return them to the original readable form through decryption. In general, the longer the key is, the more difficult it is to crack the code. This holds true because deciphering an encrypted message by brute force would require the attacker to try every possible key. To put this in context, each binary unit of information, or bit, has a value of 0 or 1. An 8-bit key would then have 256 or 2^8 possible keys. A 56-bit key would have 2^{56} , or 72 quadrillion, possible keys to try and decipher the message. With modern technology, cyphers using keys with these lengths are becoming easier to decipher. DES has an effective key length of 56 bits.

Modern Encryption Management

As technology advances, so does the quality of encryption and the current AES uses a symmetric-key block cipher with the same key for encryption and decryption. It operates on fixed-size data blocks of 128 bits and supports key sizes of 128, 192, or 256 bits. The introduction of the asymmetric key cyphers (sometimes termed public-key cyphers based on algorithms which use two mathematically related keys for encryption of the same message has had a major impact and the best example is its use in your access to sites such as Amazon. Some of these algorithms permit publication of one of the keys, because it is extremely difficult to determine one key simply from knowledge of the other.

Key management is an important topic for data security, and needs us to dive deeper into it. To understand modern encryption and keys management, let's start by understanding public keys vs. private keys and their usage.

Public Keys vs. Private Keys

Public keys and private keys are the working parts of Public Key Infrastructure (PKI) cryptography. Together, they encrypt and decrypt data that resides or moves in a network. A public key is a large numerical value that is used to encrypt data. The key can be generated by a software program, but more often, it is provided by a trusted, designated authority and made available to everyone through a publicly accessible repository or directory.

In a Private Keys system, private keys are used for both encryption and decryption of data and are shared between sender and receiver. They offer faster performance and are part of symmetric cryptography. A private key is a secret key that is shared between two parties in symmetric cryptography and both parties keep it and use it to exchange data,

Only one party keeps a private key in asymmetric cryptography. The owner of the private key publishes a public key that can be used by anyone to encrypt data intended for the private key owner, who uses the private key to decrypt the information encrypted with their corresponding public key. This is also used to create the digital signature of a file or certificate. Public key is truly public and can be shared widely while the private key is known only to the owner.

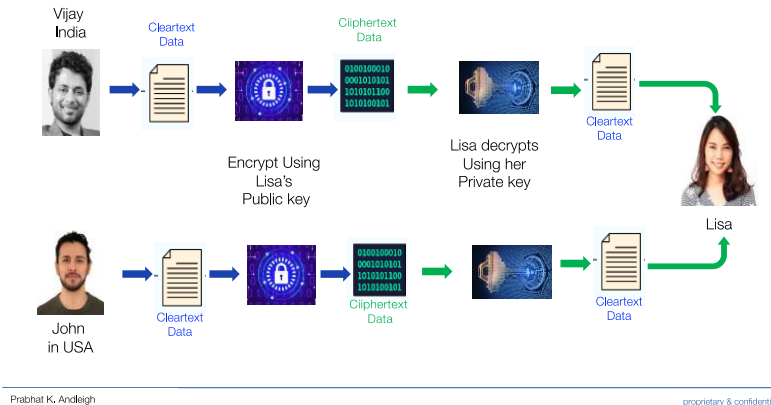


Figure 8-1: Example of Public and Private Keys for Messages

Let's use a simple example to understand the concept, Vijay (who could be in India) wants to send Lisa (who could be in the US) an encrypted message (for example, an email). If Lisa wishes to receive encrypted messages; she publishes

one of the keys, her public key, and anyone, say Vijay in this example, can use it to encrypt a message and send it to her. To do this, Vijay takes Lisa's public key and encrypts his message sent to her. Then, when Lisa receives the message, she takes the private key that is known only to her to decrypt the message from Vijay. John (also in the US) can use the same public key owned by Lisa to encrypt a message sent to her. So, any number of users can use the same public key to encrypt messages that only the owner of the private key can decrypt. As we can see, public keys are shared with everyone.

Public key cryptography is extremely useful for establishing secure communications over the Internet (via HTTPS). A website's SSL/TLS certificate, which is shared publicly, contains the public key, and the private key is installed on the origin server — it's "owned" by the website. For example, Amazon, who in this case owns the private key can use this method for secure communications with their customers who have access to Amazon's public key.

Managing Encryption Keys

Encryption keys (also called cryptographic keys) play a very important role in data security and therefore must be stored in a very safe manner. Encryption key management is the administration of policies and procedures for protecting, storing, organizing, and distributing encryption keys.

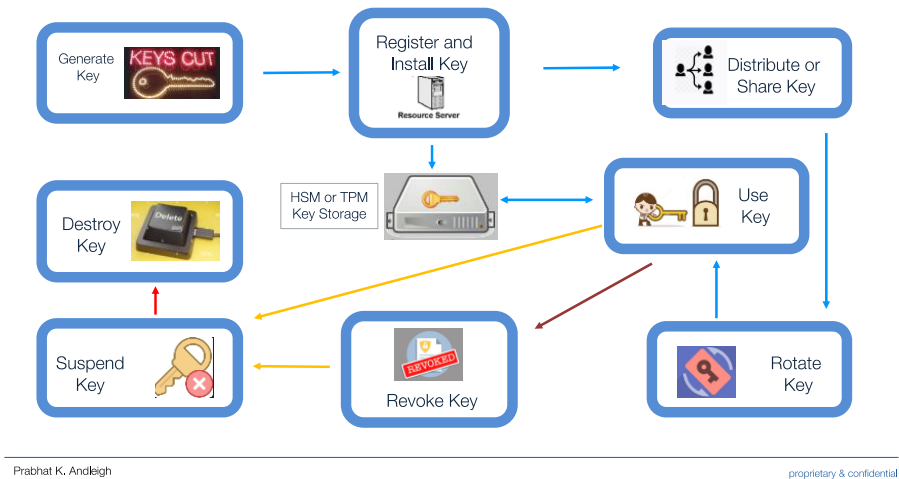
It is important to define automatic key rotation so that new versions are created on a pre-defined schedule. This reduces the chance for hackers to find the keys and use them for extended periods for exfiltrating data. Keys may be automatically deleted once they are expired, kept forever for archiving purposes, or set to delete after they have been expired for a certain amount of time.

Encryption Key Management Systems (KMS)

The generation, exchange, storage, use, crypto-shredding (destruction) and replacement of keys is known as encryption key management. You can use key management system utilities available in on-site installed as well as SaaS versions. Automating key management ensures not only the privacy of the keys but also helps you rotate the key on schedule and know when the keys are no longer needed and should be deleted.

Master keys, which are stored in secure hardware, are used to encrypt all other keys on the system. All other keys that are encrypted under these master keys are stored

outside the protected area of the secure hardware or software cryptographic key storage.



Prabhat K. Andleigh

proprietary & confidential

Figure 8-2: Lifecycle of a Cryptographic Key

The typical encryption key lifecycle described in the diagram above entails key generation, key registration along with installation and protection in an HSM or TPM, sharing of the key with users, key use for application access, key rotation, revocation and suspension, and finally, retirement and destruction of these keys. The retirement of these keys must be handled with the utmost care, especially when they are used for protecting sensitive or valuable information such as, financial transactions, credit card data, etc. A Key Management System (KMS) makes it possible to proactively manage the keys throughout their lifecycle. A KMS is capable of safely handling both requests for inbound and outbound key distribution. Furthermore, for security and compliance purposes, these systems can keep track of audit logs for these keys. To properly generate and protect the keys, the KMS needs to be supported by its dedicated HSM described below for the key storage management.

Hardware Security Module (HSM) vs Trusted Platform Module (TPM)

A hardware security module (HSM) is a tamper-proof physical device that protects asymmetric and symmetric key cryptography's secret digital keys. They are used to achieve a high level of data protection and trust when implementing PKI or SSH. By keeping the decryption keys apart from the encrypted data, HSMs have an additional layer of security. In this manner, encrypted data is kept private even in the event of a hack. All master keys should be stored in secure HSMs.

Most Hardware security modules (HSMs) are plug-in devices that can be connected directly to a computer or a network server. HSMs are typically hardened, tamper-resistant hardware devices that secure cryptographic processes by generating, protecting, and managing keys used for encrypting and decrypting data and creating digital signatures and certificates. An HSM is specifically purposed to handle issuance distribution and storage of cryptographic keys.

The keys that HSM tools are managing are frequently safely backed up outside the HSM. HSMs are frequently used by Certificate Authorities (CAs) to generate, store, and manage asymmetric key pairs.

HSMs are different from trusted platform modules (TPMs) even though both are physical devices and involve data encryption. An HSM is a removable unit that runs on its own, while a TPM is a chip on your motherboard that can encrypt an entire laptop or desktop disk.

A TPM, or a trusted platform module, is a physical or embedded security technology but unlike an HSM it has a limited scope, and it is typically a microcontroller that resides on a computer's motherboard or in its processor. TPMs use cryptography to help securely store essential and critical information on computers to enable platform authentication.

As we have seen, a KMS is employed to provide efficient management of the entire lifecycle of cryptographic keys in accordance with specific compliance standards, whereas an HSM serves as the core component for the secure generation, protection, and usage of the keys.

Encryption Types

Encrypting adds overhead; encrypting everything in the cloud causes additional overhead and time delay and the protection cost may be disproportionate to the value of the assets.

In the following sections we will dive deeper into the vast array of encryption technologies in current use and what role they typically play in data security.

TLS/SSL – Layer 4 – Transport Control Layer

Transport Layer Security (TLS) is designed to provide security at the transport layer. TLS was derived from a security protocol called Secure Socket Layer (SSL). The SSL is a secure socket layer, whereas the TLS is a transportation layer protection. The SSL and TLS cryptographic protocols authenticate server-to-device data transfers. For example, a cryptographic protocol encrypts data exchanged among the Web server and a user.

SSL aims to provide a safe and secure way to transmit sensitive data, including personal information, credit card details, and login credentials. The SSL protocol can only be used by websites with an SSL certificate, a digital document that validates a site's identity. SSL is being deprecated in favor of TLS which is much more secure.

The TLS (and SSL) protocols are located between the application protocol layer and the TCP/IP layer, where they can secure and send application data through the transport layer. Because the protocols work between the application layer and the TCP/IP layer, TLS and SSL can support multiple application layer protocols.

TLS uses both symmetric encryption and public key encryption for securely sending private data, and adds additional security features, such as authentication and message tampering detection.

Transport Layer Security, or TLS, is a widely adopted security protocol designed to facilitate privacy and data security for communications over the Internet. TLS ensures that no third party may eavesdrop or tamper with any message. TLS uses both asymmetric encryption and symmetric encryption. During a TLS handshake, the client and server agree upon new keys to use for symmetric encryption, called "session keys." Each new communication session will start with a new TLS handshake and use new session keys.

Transport Layer Security (TLS) encrypts data sent over the Internet to ensure that eavesdroppers and hackers are unable to see what you transmit. This is particularly

useful for private and sensitive information such as passwords, credit card numbers, and personal correspondence.

As said earlier, an important use case of TLS is encrypting the communication between web applications and servers, such as web browsers loading a website. TLS is also used in applications such as email, file transfers, video, and audio conferencing. TLS is compatible with a significant number of protocols including HTTP, SMTP, FTP, XMPP, and others. Note that TLS isn't designed to secure data on end systems, only data transferred over the internet.

TLS 1.3 is the current TLS level and is designed to prevent man-in-the middle attacks. It is replacing TLS 1.2 and earlier versions of TLS and SSL.

Encryption Algorithms

An encryption algorithm is the method used to transform cleartext data into ciphertext. An algorithm will use the encryption key to alter the data in a predictable way, so that even though the encrypted data will appear random, it can be converted back into plaintext by using the decryption key. Different algorithms are used for applications based on the requirements of the application and the required security characteristics.

Let's review the common encryption methods currently used and what are they used for.

Symmetric Key Encryption

Symmetric Key Encryption involves a shared secret (single key encrypts and decrypts). It does not require signed certificates (certificates require public/private keys). Parties typically know each other and already have shared secret keys or use another mechanism to dynamically share a secret key. That is, it involves passing keys out of band (that is keys are passed using a different mechanism than the one used for the sharing the encrypted data. This also does not involve hashing (one-way mathematical algorithm for checking messages authenticity) or key pair or key exchange used for asymmetric encryption.

Asymmetric Encryption

Asymmetric Encryption uses mathematically related key pairs (used for asymmetric encryption). Diffie Hellman (DH) key exchange process is designed to allow two

parties unknown to each other to create a shared secret by sending cryptographic keys over a public channel using public and when mixed with a private key arrive at the shared secret. Although the result is symmetric the method is asymmetric. Typically, the public key is registered with a certificate authority and is available to the other party. So any message encrypted by the public key can be decrypted only by the related private key. This method is typically used for digital signatures and for exchanging secure symmetric session keys.

Secure Hashing Algorithms

Secure Hashing Algorithm SHA-256 and the emerging *SHA-512* are hash functions developed by the United States National Security Agency (NSA) are commonly used for protecting passwords. They replace SHA-1 and SHA-2. SHA 256 refers to the bit-size of the hash value generated (256 bits) making it computationally unfeasible to reverse engineer the original input data. It is, therefore, a one-way transformation. The SHA function generated will always be the same for the same message, but even if one character in the message is changed, the SHA function will be different. The SHA value is generated through a series of mathematical transformations. Message authenticity across transit can be established by comparing the hash function created before the message is sent with the hash function generated by the receiving party. Identical hash values, also called digests, indicate that the message has been unaltered during transit.

For digital signature, the hash value generated from the content being signed is encrypted with the signer's private key and attached to the message. The recipient can decrypt the signature using the sender's public key and confirm that the attached hash value from the sender is the same as the calculated hash value create at the receiver end.

Application Level Encryption (ALE)

In ALE (*Application Level Encryption*) as the name implies, the encryption engine resides within the application. Note that encryption is easy in this case but key management can be complex due to need for a key management infrastructure and additional supporting processes, which should be aligned with your systems architecture.

ALE can be implemented in various ways to address different security requirements such as, end-to-end encryption, zero trust architectures, partial field-level database encryption and so on. The encryption subsystem works better when integrated with

others to form defense-in-depth with access control, logging, intrusion detection, request authentication, and data leakage prevention.

ALE does a better job of protecting from risks than transport and at-rest encryption, but at the cost of tradeoffs. With the data encrypted, searching encrypted data becomes difficult. This better data protection is driving the use of application-level encryption, in the fields of finance, healthcare where it is used application-level end-to-end encryption for sensitive patient data, and others. The locations where this encryption is employed include:

- ***Client-side encryption which is implemented on the client*** .
- ***End-to-end encryption*** when implemented on clients in a way that no secrets or keys are available to servers.
- ***Field-level encryption where certain fields in a database record are encrypted based on the context. Note that this makes searching for data difficult.***
- ***Zero Trust Encryption where the end-to-end encryption can operate under full zero trust assumptions, making the application compliant to the zero-trust architecture principles.***

In short, application-level encryption only points to an architectural choice of where encryption happens. But if we look closer, that means many things for your distributed application.

When and why use ALE: every security requirement should be driven by a risk model and a threat model that justifies the choice of security control, the scope of its application, and details. ALE addresses that for the specific use cases in domains such as, finance and healthcare.

Homomorphic and Polymorphic Encryption

Homomorphic Encryption is an extension of public-key cryptography with additional capacity for computing encrypted data without access to the secret key; that is, it does not need to decrypt data to process it. It can perform computations directly on ciphertext (the encrypted data). Homomorphic encryption converts data into ciphertext that can be analyzed and worked with as if it were still in its original form, so it enables complex mathematical operations to be performed on encrypted data without compromising the encryption.

While Homomorphic encryption creates only one single set of encrypted data and provides the user with only one key for decryption, Polymorphic encryption allows

data to be encrypted in multiple forms, providing those who access it with multiple keys for each function of the encrypted data sets.

In Homomorphic encryption applied to Data-in-Use, a code snippet that encrypts the data is included in the transmitted data. This code also decrypts, so both ends need that code snippet. With this approach, the ciphertext does not need to be decrypted for use because the decryption happens on the fly while the ciphertext is being used.

A common use case for Homomorphic encryption is for organizations to securely store, process, and exchange sensitive data, while also ensuring compliance with data governance regulations, such as in the case of customer data or personal data. It allows the data to be transmitted securely to the destination without the complexity associated with decryption key management.

Types of Homomorphic Encryption

Homomorphic encryption uses the concept of mathematical gates such as addition and multiplication for computation. The following lists the common types of Homomorphic encryptions.

- Partially Homomorphic encryption that evaluates circuits consisting of only one type of mathematical gate.
- Somewhat Homomorphic Encryption can evaluate two types of gates, but only for a subset of circuits. The limit on somewhat homomorphic encryption comes when a ciphertext generates too much noise in the data.
- Fully Homomorphic Encryption enables computations directly on encrypted data, or ciphertext, to keep data protected at all times. The benefits of Fully Homomorphic encryption are significant, from enabling the use of untrusted networks to enhancing data privacy, securing data stored in the cloud, enabling data analytics in Regulated Industries, and improving election security and transparency.

Use Cases for Homomorphic Encryption

Homomorphic encryption can help organizations maintain a high level of data security without reducing productivity or violating protocols. It can greatly increase data privacy and security in a variety of applications.

Electronic cash systems might also use homomorphic encryption. A business transferring funds may not want the system, or a user within the system with access to the data, to know the extent of the funds being transferred.

Government systems also use Homomorphic encryption for secret data being shared with government agencies.

Homomorphic Encryption Advantage

With homomorphic encryption, organizations can establish a higher standard of data security without breaking business processes or application functionality. These organizations can ensure data privacy, while still deriving intelligence from their sensitive data.

Homomorphic Disadvantage

One of the biggest drawbacks of homomorphic encryption is that it is computationally intensive and slow. Encrypting, decrypting, and performing operations on ciphertexts requires more resources and time than on plaintexts.

Format-preserving Encryption

Format preserving encryption is a method of encryption where the resulting ciphertext has the same form as the input cleartext. The form of the text can vary according to use and the application, for example a 16-digit credit card number. It preserves the format of the information while it is being encrypted. It is weaker than standard Advanced Encryption Standard (AES), but it can preserve the length of the data as well as its format. This is very useful in cases where the size and format of the field are important such as, in the case of storage of data in a database field.

Self-decrypting Encryption

Self-decrypting storage disks use software programs that can encrypt the data on your computer's storage drive. These programs encrypt the data while it is being written to the storage and decrypt it while it is being read from your storage drive.

As an organization, any sensitive data stored within your business systems on-premises as well as in the cloud environments is your responsibility and your organization can be held liable in case of a breach. Customer credit card information, personal identification numbers, email lists, internal policies, product roadmaps, and intellectual property are stored on your storage systems, so these systems are vulnerable to accidental loss, hackers, and data thieves. Self-decrypting storage reduces the risk of exposure due to a deliberate attack.

A performance disadvantage is because the processor is working to encrypt and decrypt the data on the fly and can slow down storage and retrieval. There are also ways that hackers may be able to overcome the encryption and recover data that has been encrypted.

Transparent Encryption

Transparent Encryption is a form of self-encryption and refers to a method of encrypting data at rest, where the encryption and decryption process is transparent to the user and the application. This means that the user or application does not need to take any explicit action to encrypt or decrypt the data. The encryption and decryption are performed automatically by the underlying processes and software within the storage or database management system.

Transparent Encryption is often used in databases, file systems, and storage devices to protect sensitive data without modifying the existing application or infrastructure. This allows for data to be encrypted without any changes to how the application interacts with the data,

Some examples of where Transparent Data Encryption solutions are used include:

- Full-Disk Encryption (FDE) for encrypting entire storage devices. This is totally transparent to the applications accessing the data.
- File-level Encryption for encrypting files and directories. This is used by a wide variety of software vendors.
- Database Encryption for encrypting database data, like column-level encryption or TDE (Transparent Data Encryption).

Transparent Encryption can be used for specific database tables, but not entire database files, applications, or objects.

Data Alteration Treatments for Data Protection

At a very basic level, data alteration refers to altering the raw data just before it is processed by a computer and then changing it back after the processing is completed. Data alteration is also recognized as a process involved in the process of creating, recording, encoding, examining, checking, converting, or transporting computer data.

From a data security perspective, an unauthorized alteration is the unauthorized modification of a secure document and can occur for a variety of reasons, primarily

by hackers for ransom demands or sometimes even by mistake by workforce members.

Data diddling is a type of cybercrime in which data is altered as it is entered into a computer system, most often by a computer virus or by a workforce member trying to cover up a financial crime such as diverting funds into their private accounts. Computerized processing of the altered data results in a fraudulent benefit.

Our focus here is to discuss, how you can use data alteration to prevent data being compromised or in the case of exfiltration, the data becoming useful to the hackers.

Data Alteration Treatments for Data Protection

The legitimate uses of data alteration include several common approaches that make the data unreadable to an unauthorized user. The following section describes the common approaches for data alteration.

Data Masking

Data Masking. The common types of data masking include Random substitution, Algorithmic substitution, Shuffle Masking (shifting data in columns), anonymization based deidentification (used for credit card masking by replacing real credit card numbers by virtual equivalents), Deletion.

Data making can be Static or Dynamic. In static data masking the sensitive data is permanently replaced by altering data at rest. This is a common methodology for generating masked test data that looks just like real data. In dynamic data masking, the decision to mask the data is determined at the time of the data access request and it is based on attribute values in the data being requested.

Data masking is a method where the masked data created is structurally similar but different in content from the version of an organization's original data that can be used for purposes such as software testing and user training. The purpose is to protect the actual data while having a functional substitute for occasions when the real data is not required or not used for security and privacy reasons. The masked data is essentially a fake, but a realistic version of your organizational data. The goal is to protect sensitive data, while providing a functional alternative.

Data masking replaces sensitive data with “decoy” data that looks authentic. Only authorized users have access to the real data – so even when a breach occurs, the private data is still protected.

Benefits of Data Masking

There are several benefits from data masking. These include:

Enterprise Data Security for US Europe and Asia

- Enhanced security. Data masking helps mitigate the risk of data breaches, and risks from compromising by different types of malwares and cyberattacks.
- Compliance with regulations when the regulatory agencies specifically require data to be masked when not protected and in test environments.
- Data privacy during transit to other environments.
- Secure third-party data sharing for data stored in third party storage.
- Datalakes used for customer behavior analytics.

The difference between hiding and masking is that hidden variables don't appear at all, while masked values appear but are replaced with asterisks in logging tracing and debug sessions.

Some drawbacks from data masking include the following:

- Data masking may not always provide complete protection for sensitive data, and it can be defeated through brute force techniques.
- Data masking may not always preserve the usefulness or value of the original data, especially when you need to run analytics in production on masked data and the masked value is an important sorting criterion.
- Data masking may not always be reversible if the original data cannot be found or is corrupted.
- Data masking may not always be compatible or interoperable with other data sets or systems limiting its usefulness in some cases.

Data Anonymization

Data Anonymization – direct anonymization is often used with (PII) and indirect anonymization is often used with identifiers (demographic and other which can be aggregated to identify). In that case, anonymization removes the indirect identifiers. Data anonymization is the process of protecting private or sensitive information by erasing or encrypting identifiers that connect an individual to stored data.

One example of anonymized data is a dataset that has been stripped of any original personally identifiable information such as names, addresses, and phone numbers. This type of data can be used to analyze trends and patterns without the risk of exposing any individual's personal information.

Financial services companies, such as banks, brokerages, and insurance companies, who may be under regulatory controls, employ data anonymization to protect sensitive information such as financial histories, PII, and transaction information.

It is important to note that Data anonymization removes classified, personal, or sensitive information from datasets, while data masking obscures confidential data

with altered values. So you need to decide which is more appropriate for your use case.

Tokenization

Tokenization is used to secure many different types of sensitive data, including:

- Payment card data.
- Government identification numbers such as U.S. Social Security numbers or Indian Aadhaar card numbers and similar national identification numbers.
- Telephone numbers.
- Passport numbers.
- Driver's license numbers.
- Email addresses.
- Bank account numbers.
- Names, addresses, birth dates.

Tokenization is the replacement of actual data by token. For example, replace credit card or bank account number with an alternate code called the “token”, which shall be unique for a credit card or bank account number; the token requestor (i.e. the entity which accepts requests from the customer) tokenizes the credit card or bank account number and passes it on to the downstream application.

For payment processing, tokenization requires replacing a credit or debit card or one's account details with a token to protect their identify and to prevent identity theft. Tokens, in themselves without the availability and association with the data they represent, have no value because they cannot be associated with any account or person. For example, the customer's 16-digit main Indian national account number (PAN) is replaced with a randomly generated, bespoke alphanumeric ID. While the token is used for display on-screen, the real value behind the token is used for transaction processing.

As we have seen, tokenization replaces valuable/sensitive data with tokens and saves both; so it needs at least 2 databases one for the token and another one referencing the token with the real data it represents.

SSMS – Secret Sharing Made Short

SSMS (Secret Sharing Made Short) encrypts the data set, then splits the already encrypted data into pieces, and splits the encryption keys into pieces. It then uses digital signatures to sign the encrypted data pieces and the key pieces and distributes them to various cloud storage locations. This makes exfiltration difficult

because the hacker must figure out the exact sequences of the broken pieces of the data and the encryption keys to reassemble it to its original form.

All-or-Nothing Transform with Reed-Solomon

All-or-Nothing Transform with Reed-Solomon (AONT-RS): It integrates AONT with Erasure Coding. In simple words Erasure coding (EC) is one of the methods of data protection through which the data is broken into sectors. Then they are expanded to the sector size with redundant data and encoded by including the redundant data pieces in ciphertext and stores them across different storage media. Erasure coding adds the redundancy to the system that tolerates failures. In this methodology, it first transforms the data and the encryption key into blocks and then uses the information dispersal algorithm to split the blocks into multiple shares distributed to different cloud storage services (similar to SSMS).

Bit Splitting

Bit splitting is the process of encrypting the data in AES-256, then splitting the encrypted data into smaller chunks distributed to several locations and encrypted again using SHA-256 hash to ensure integrity of the encrypted ciphertexts).

A negative aspect associated with it is that it may require trust in additional third parties beyond the primary CSP because the chunks may be in different clouds. Note that end-users have no involvement in this, and vendors have no challenges. Furthermore, senior management typically has no policy concerns in using this methodology.

Secure User Access vis VPN (Virtual Private Network)

VPN (e.g. IPSec Gateway) – encrypts data packets. VPNs use encryption to create a secure connection over unsecured Internet infrastructure. VPNs protect data as users interact with apps and web properties over the Internet, and they can keep certain resources hidden.

Crypto-shredding Before Data or Storage Device Destruction

Crypto-shredding is used for data destruction. This is achieved by encrypting the data and also encrypting the first set of resulting keys with a different encryption engine. Then by destroying the original set and the second keys, no keys are available to decrypt the data. The media can be safely destroyed after this action.

Digital Signatures

A digital signature is an electronic, encrypted, stamp of authentication on digital information such as email messages, macros, or electronic documents. A signature confirms that the information originated from the signer and has not been altered in storage or in transit. It is a cryptographic output used to verify the authenticity of data.

A digital signature algorithm allows for two distinct operations: a signing operation, which uses a signing key to produce a signature over raw data. And a data retrieval and verification operation

When a signer digitally signs a document, a cryptographic hash is generated for the document. That cryptographic hash is then encrypted using the sender's private key, which is stored in a secure HSM box. It is then appended to the document and sent to the recipients along with the sender's public key.

Verifying a signature will tell you if the signed data has changed or not. When a digital signature is verified, the signature is decrypted using the public key to produce the original hash value. The data that was signed is hashed. If the two hash values match, then the signature has been verified.

Chapter 9 – Identity and Access Management

Identity and Access Management (IAM) is a framework of policies, processes, and technologies that enable organizations to manage digital identities and control user access to critical information of the organization. This is very critical not only for cybersecurity in general, but it also plays a very important role for data security in particular. As we have seen, the security of your organization's data assets is the express responsibility of your organization. Loss of important data and unauthorized exposure of sensitive data risks loss of reputation along with loss of business. Your organization's executives can even face criminal liability in case of a breach and exposure of sensitive data.

IAM ensures that the right people get the right access to the right resources at the right times for the right reasons, enabling the right business outcomes.

While the sentence appears to be a mouthful, there is an important truth and even more architecture, design, and implementation implications to meet every concept described in this sentence.

The concepts and use of IAM have been around for years and have continued to rise in prominence due to increasing threats to data security due to hacking for ransom, revenge or even state sponsored hacking and cyber terrorism.

Well thought out IAM implementation and governance improves security and mitigates risks from external hackers as well as insider threats, such as that from disgruntled employees. IAM controls provide a variety of security benefits, including data security benefits, to organizations.

Scope of IAM Systems

IAM systems govern all users of an organization's IT infrastructure under a standard policy for identity management and access controls to ensure proper authentication of the users and authorization to access certain types or classes of data.

IAM platforms provide workforce identity management or customer identity and access management (CIAM), or both. Well defined access controls and policies

govern both classes of users to ensure that they have the level of access to data to perform their tasks or business functions.

Modern IAM Systems

The business world has become exponentially more digital as a whole, which leads to an organization's data being accessible from the Internet. The traditional network perimeter that covered on-premises is still very relevant but has undergone a substantial change as a model with increasing use of SaaS (Software-as-a-Service) applications and application migrations to the public and private clouds. It has become a hugely expanded security model. Identity is the new perimeter, and therefore, not surprisingly, there has been a dramatic increase in Identity-based threats and attacks.

With sophisticated security breaches continuing to rise, organizations need more systems, protocols, and standards to safeguard their data. And while regulations for privacy, security and compliance have intensified, the time, resources and money to adhere to these standards have increased. Organizations are increasingly adopting a Zero Trust security model in which identity verification plays a major role. As digital roadmaps accelerated, risks and associated costs have risen as well. It is important to implement enterprise wide security in a comprehensive manner the first time.

Modern security requires IAM

Businesses have automated IT systems for decades but the challenge of ensuring cybersecurity and more importantly, data security, are more recent. Our need for speed and agility, our requirement for services to be available on-demand, our desire for flexibility to scale up and down as needed on-premises and in the cloud, and ultimately our strong desire to constantly innovate, makes moving IAM systems and authorizations to the cloud as an extension of the on-premises systems an imminent obligation.

Modern IAM systems and control programs bring together identity management, identity governance, access management and privileged access management. It is critical to establish security as it fosters trust, a critical component of building brand loyalty. However, more security imposes more demands on users to authenticate themselves, and quite often, frequently when single-sign-on (SSO) is not designed well. So, organizations walk a tight line in balancing security with seamless, convenient user experience.

With the wide ranging on-premises and cloud deployments, it has become important to have a governance administration into one unified, on-premises and cloud-based solution that is built for the required level of security and almost unlimited scalability in place to manage an increasing workforce and customer base.

These modern IAM architectures provide the frameworks for protecting data, especially sensitive data, and the organization's other resources. The network perimeter used to be the organization's data center but with increasing deployment of cloud-based applications, the perimeter is expanding beyond the internal networks, which establish security boundaries in on-premises systems, to cloud environments. This has increased the challenge in protecting the expanded perimeter networks. Just firewalls are no longer sufficient for managing access to applications and data in all applications deployment environments, especially, the cloud.

IAM Systems as the Common Link

A modern IT infrastructure supports a wide variety of applications, including those developed in-house, developed by third party vendors running within your organization's infrastructure (which also includes cloud deployments), and SaaS applications that your users and customer's log into to perform their functions.

So, a modern IT infrastructure brings together a wide variety of cloud infrastructure providers, SaaS applications, developer and management tools, analytics services, and security platforms that comprise their solutions. They all have one common link that allows users to seamlessly access and use all of these disparate solutions. That link is user identity. Modern IAM solutions enable organizations to empower all of their users to easily and safely access all of their applications, devices, and technologies. Legacy IAM solutions that were built for our on-premises world, and typically siloed and difficult to maintain, need significant retrofits and in many cases total overhaul and rebuild of these Identity platforms to provide the required coverage and meet expectations.

The impact of IAM architectures and implementations are pervasive and far reaching across your organization and customer relationships.

IAM Roles, Groups, and Entitlements

All workforce members in your organization have a role to play depending on their title and job description. The type of access to data, especially sensitive data, and the types of functions they can perform in manipulating that data is determined by the needs of their role. For each role, the workforce members need permissions to access data and use applications to manipulate the data.

Assigning permissions to each workforce member individually can be quite daunting in a large organization. It is easier to define roles by job types and assign the access permissions by job types. These job types are created as groups and each workforce member can be assigned to one or more groups depending on their access requirements. The groups they are members of can change from time to time as changes take place in their roles such as job transfers, promotions or temporary project assignments.

Roles and Groups management is a major IAM activity for your organization on an ongoing basis, and administration of these roles and groups has to be managed carefully and performed on a regular basis.

In a large organization Identity and Access Management (IAM) roles are entities you create and assign specific permissions to trusted identities such as the workforce identities and applications that allow data access. When your trusted identities assume IAM roles, they are granted only the permissions scoped by those IAM roles.

It is important for your organization to periodically review the roles and group memberships of each of your workforce members to ensure that they remain consistent with any changes in their job responsibilities. Such audits/reviews must be conducted quarterly and in the worst case annually. Attestation by the manager of the workforce member should be required for such reviews.

Key Terms Used for IAM

The following lists the key terms associated with IAM that you should become very familiar with:

- *Workforce identity*. A typical business makes use of a wide variety of applications. Some of these applications touch and use sensitive data. The access to these applications should be governed by the roles and group

Enterprise Data Security for US Europe and Asia

assignments of the workforce members based on their identity and the permissions associated with that identity.

- *Customer Identity and Access Management (CIAM)*. Assigning an identity to each customer who needs access to applications within your extended network ensures that they can securely perform functions and access only the data they are permitted to access, download, and use.
- *Federated Identity*. Your workforce members may use applications across many systems and applications with their own identity management systems. Federated identity is a method of linking a user's identity across multiple separate identity management systems. It allows users to quickly move between systems while maintaining security.
- *B2B identity*. Using Business-to-Business identity management ensures that your organization can deliver secure and seamless experiences for your business partners, vendors, and clients by providing fast and frictionless access to your SaaS applications while protecting business critical systems against malicious attacks and data breaches.
- *Single Sign-On (SSO)*. Single sign-on is an authentication scheme that allows a user to log in with a single ID to login and access any of several related, yet independent, software systems. This helps generate good user experience. True single sign-on allows the user to log in once and access services without re-entering authentication factors.
- *Multi-factor authentication (MFA)*. Multi-factor authentication is an electronic authentication method in which a user is granted access to a website or application only after successfully presenting two or more pieces of evidence to an authentication mechanism. These two pieces of evidence can include any two of: a) something the user has such as a security token, a bank card, a key, etc., b) : something the user knows uniquely such as a password, passphrase, PIN, etc., and c) something the user is that uniquely defines the user such as biometrics (fingerprint, eye iris, voice), behavioral patterns (even such as typing speed), etc.
- *Anomaly detection*. Anomaly detection is the process of examining specific data points and detecting rare occurrences that seem suspicious because they're different from the established pattern of behaviors. Anomaly detection works by taking a baseline of the normal traffic and activity taking place on the network and detecting unusual activities such as spikes, dips, deviations from cyclic patterns, and trend changes.

IAM Threats and Risk Management

Cybersecurity threats against an organization's sensitive data has been increasing exponentially and affects all classes of businesses.

A threat is a potential for a threat agent to exploit a vulnerability in your expanded network perimeter due to misconfigurations that allow a bad actor to exploit it. A risk is the potential for loss when the threat happens and a threat agent is able to exploit the vulnerability and compromise or exfiltrate data, especially sensitive data.

Some of these threats from bad actors include malware, ransomware, phishing attacks and social engineering. Malware is a type of software that is designed to harm or damage a computer system making it potentially inoperable to perform the tasks it is designed for; these can include viruses and data deletion or alteration. Phishing and social engineering are used to gain the credentials of a workforce member to use those credentials to gain access to the data within your on-premises and cloud environments that allow the bad actors to compromise the security and exfiltrate data.

Cybersecurity risks include loss of confidentiality, data integrity, availability of information for use, or corruption of information (or control) systems and reflect the potential adverse impacts to organizational operations (i.e., mission, functions, image, or reputation) and loss of assets, or exposure of individuals' personal data, and so on.

Data risk management includes all practices for identifying risks, assessing risks, and mitigating risks to an acceptable level, called the organization's risk appetite. While it is not practical to reduce data risk completely, the goal is to reduce it to the level that the organization deems as acceptable as a risk versus the cost eliminating risk tradeoff.

Governance, Risk, and Compliance (GRC)

GRC is a comprehensive approach to managing cybersecurity that incorporates three key components: governance, risk management, and compliance. Governance refers to the policies, processes, and procedures that an organization has in place to manage cybersecurity risk. These must be documented, and the workforce must be trained to understand the implications of these in their day-to-day activities.

Governance, risk, and compliance (GRC) is not a part of IAM. GRC and IAM are two separate disciplines with different reporting structures and distinct goals. GRC

Enterprise Data Security for US Europe and Asia

is a structured way to align IT with business goals while managing risks and meeting all industry and government regulations. It refers to an organization's strategy for handling the interdependencies among corporate governance policies, enterprise risk management programs, and regulatory and standards compliance.

Insufficient data access controls, improper identity lifecycle management, identity misconfigurations and privilege abuse can lead to compliance violations, customer privacy breach, and total business disruption. GRC is intended to ensure that policies and risk management address how to mitigate these.

Some of the common and effective ways to lower the risk of data exposure are:

- Know where your sensitive data is and prioritize resources to protect it.
- Verify what users and groups have access to which data and if they are authorized to access it.
- Ensure implementation of appropriate security policies to specify what workflow is needed to access different types of data.

Your organization needs to assess the risks you are facing and from what situations and have a firmly documented and rigorously followed set of guidelines and procedures that your workforce members need to adhere to. Also important is the documented process for reporting attacks and breach attempts.

IAM Policy Management

Identity and access management policies are documents that explain how users should connect to resources within the organization's network. IAM policies cover critical security areas including authenticating user accounts, assigning access privileges, and password usage. IAM policies define permissions for an action performed by a user regardless of the method that the user used to perform the operation. So, they actually assist the user by explaining how to safely access the organization's resources within the expanded network perimeter as well as how they should access external resources. These are called the rules of access management. For the administrators, this ensures that all users follow the documented rules of access management ensuring required cybersecurity and data security controls.

IAM policies lay the foundations to address a wide range of access management challenges. Security concerns addressed by policies should include:

- Credential thefts due to weak authentication practices.
- The use of weak passwords that can cause exposure of user credentials.

Enterprise Data Security for US Europe and Asia

- Internal attacks by disgruntled workforce members who may be able to escalate in an unauthorized manner their privileges to administrator account privileges.
- Security risks posed by improper and unauthorized use of shared accounts.
- Outsider attacks via third-party users and vendors who may be able to obtain login credentials that allow them privileges and permissions that they should not have.
- Orphaned accounts that have not been deactivated and deleted and can be used by hackers for cyber-attacks.
- Insecure remote access via home devices or public wi-fi. Your organization should set policies on how workforce members use personal devices and unsecure public WiFi connections.
- Compliance issues due to poor auditing of access requests that did not adhere to policies.
- Poor perimeter protection due to chaotic access controls. Lack of a unified enterprise-wide access management policy.

An IAM policy aims to create certainty and reduce confusion. It provides easy-to-follow rules that apply across the whole organization. Access policies make cybersecurity more robust by hardening the network perimeter. Robust multi-factor authentication (MFA) and privileges management systems deliver enhanced protection for confidential information such as sensitive data. This lowers the risk of data loss and can help your organization meet your compliance goals. Active IAM policy management may, in fact, be a requirement by the regulatory agencies or entities that require to meet standards (such as PCI-DSS)

IAM policy guidelines

IAM policies depend on the needs of the organization and responsibilities imposed on them by regulatory agencies and market forces for maintaining confidential and sensitive data. Financial services firms, for example, need to have policies that ensure that they meet the data security requirements imposed by regulatory agencies such as FFIEC, OCC, GDPR and other country and state laws.

Good IAM policy documentation and rigorous governance boosts security and provides greater control of user access to your system. This helps your organization mitigate data breaches, and identity theft and illegal access to sensitive corporate information.

IAM policy is typically a subset of a larger corporate IT policy document and has some shared characteristics that may include most, if not all, of the following:

Enterprise Data Security for US Europe and Asia

- *Version history.* It should provide the details of previous versions along with date and time information and the identity of the authors, and also describe why the policy was changed.
- *Purpose/Scope.* What the policy document aims to achieve, and why it is important from a cybersecurity and data security perspective,
- *Audience.* Who the policy applies to, and who is liable for penalties if policies are not followed.
- *Definitions.* A glossary of terms used in the document that describes exactly what they mean in the context of the policy. This will assist readers in understanding authentication requirements.
- *Exceptions.* Users may need access in situations that breach specific access rules. Explain how to manage such exceptions. This should be short because the policy itself should cover most practical situations. Exceptions should be rare, so a short referral is all that's needed.
- *References.* Details of any regulatory frameworks, countrywide federal and state laws, industry standards and internal documents referenced by the policy. This enables readers to access further sources of information.
- *Enforcement.* How the policy will be enforced and details on penalties for breaching the Identity and Access Management policy. This includes both internal sanctions and the possibility of civil or criminal penalties.

In addition to these general common policy statements, the specific sections for the IAM policy should describe the following:

- *Access control.* Rules relating to log-in procedures and account creation.
- *Account management.* Rules for system administrators. Includes account data, shared accounts, and logging user activity. It should also describe how and when accounts are de-provisioned.
- *Administrator/special access.* Administrators and special account holders typical have wide unhindered access to corporate resources including potentially, confidential data. The policy should describe the guardrails for using administration or special accounts.
- *Access Rights.* Policies should describe the access rights of various classes of users and verification methods employed to manage such access. This should include topics such as password and passphrase management, MFA, and logging policies at the user verification stage.
- *Rules for Managing Access Privileges.* These rules should describe how administrators should manage access privileges. These rules should focus on the principle of least privilege.

Enterprise Data Security for US Europe and Asia

- *Remote Access.* Policies relating to remote connections and securing remote work. Focuses on device security and authentication practices.
- *Vendor access.* Policies should describe the rights and limitations of third-party vendor access.
- *Data Collection Rules.* Rules governing data collection are derived from the requirements of regulatory agencies industry standardization bodies,

The final document should reflect the needs of each department whin your organization. It should deliver clarity while covering every relevant identity and access management need.

Policy Sections for Cloud Environments

The IAM policy sections focused on hybrid or cloud native applications are extensions of the general IAM policy definitions.

Deployment processes for cloud environments and on-premises infrastructures may look similar, but there are differences that must be accounted for. Extending on-premises capabilities to cloud can also result in some security blind spots due to misconfigurations such as IAM configuration drift across a multi-cloud environment that are not very obvious. To avoid risk, it's critical to understand the most common misconfigurations and ensure that the policy documentation specifically addresses this issue.

Departmental View of IAM

The departments within your organization have different business goals and therefore different processes for collecting and managing data. Consequently, their security goals and how they achieve them are also very specific to the nature of their business. The table below explores the differences across the major business departments that collect and process workforce or customer data.

Department	Security Goals	How to Achieve
Corporate Security e.g. CSO, CISO	<ul style="list-style-type: none">• Reduce the risk of a security incident.• Accelerate detection and response time.	Fully integrate with all apps, domains and devices and bring them all into one place to manage, view and administer security.

Enterprise Data Security for US Europe and Asia

	<ul style="list-style-type: none"> • Mitigate the impact of a security incident. 	
IT & Business Technology e.g. CIO	<ul style="list-style-type: none"> • Drive operational security efficiency. • Improve workforce data access controls, productivity, and satisfaction. • Enable business innovation and digital transformation 	Utilize automation in your IAM solution for administrative actions such as user account creation, granting permissions, and speeding up account management services to workforce members. Increase software updates to include security releases.
Infrastructure and Operations (I&O)	<ul style="list-style-type: none"> • Migration from on-premises to cloud architecture to support cost-effective infrastructure growth. • Ensuring all systems are highly available, resilient, and performant (no planned downtime) 	Have a secure cloud-first architecture that easily integrates with legacy on-premises components to provide a highly resilient and available service with at least a 99.99% SLA.
Product Engineering	<ul style="list-style-type: none"> • Drive operational efficiency to focus on secure development efforts on core business • Accelerate time to build and release for their customers 	Provide an advanced, developer-first, identification and authorization platform for developers, and training developers to practice secure development.
Marketing & Sales e.g. CDO, VP or Marketing & Sales	<ul style="list-style-type: none"> • Enhance customer confidence in security practices to accelerate revenue generation. • Provide secure collaboration with business partners. • Enhance customer login and SSO experience 	Your IAM solution should have flexible single-sign-on login solutions while removing friction from customer journeys with advanced Identity capabilities such as progressive profiling and biometric authentication

Table 9-1: Corporate Department Roles for IAM and Security Management

Selecting IT and Data Security Solutions

Selecting the most appropriate and efficient IT and data security solutions is a very important function of the IT and Cybersecurity Services teams. While providing an advanced, developer-first, identification and authorization platform for developers enhances their productivity and the ability perform secure development, it is also important to enhance the productivity and frictionless user experience for the rest of the workforce.

The following lists a set of attributes that should drive your technology choices and implementations to achieve data security goals while maintaining positive user experience.

Neutral and Independent IAM Solution

Your IAM solution must be technology agnostic to maintain business agility, allowing you to choose the best software solutions based on your business needs. Most departments choose the best-of-breed solutions, but if not properly coordinated they can become unmanageable with a real lack of integration. Cloud migrations also result in new tech stacks that may vary by cloud service provider (CSP). The challenges for identity management can be immense. To align all departments, it helps to maintain neutrality and select IAM solutions that are independent and supported by all departments without constraining yourself. The table below describes the key attributes of your IAM solution.

Your Solution Should Provide	Target Result
<ul style="list-style-type: none">• Broad and deep set of pre-built integrations that support the latest and most commonly used open standards.• Integration of user directories and protocols that enable the automation of end-to-end Identity lifecycle.• Support for all open authentication standards.• Secure hybrid IT access for mission-critical business applications and multi-cloud environments across a single Identity platform.	<p>A strong security posture based on freedom to select best of breed solutions now and in the future that allow you to maximize your investments; along with automation for faster and efficient rollouts and increased operational and developer efficiency.</p>

Table 9-2: Characteristics of an IAM Solution

You need to consider how to get departments to align with your technology and tech stack choices and adjust to your rollout frequency.

Customization

Organizations undergo significant changes over time and so should your IAM solution as the needs of your organization change. Access management for your workforce should not only enhance workforce efficiency with fast access to the data they need but also ensure data security. Your IAM solution should consider every user (workforce, customers, partners, and vendors) across their entire Identity and use cases.

Your Solution Should Provide	Target Result
<ul style="list-style-type: none">• An API-first architecture with a comprehensive extensibility framework that ensures flexibility.• Dynamic reconfigurable workflows that allow automation in user lifecycle and access policies.• Visibility into all device identities to create contextual security policies across all environments.• Policy engine to allow dynamic access policies tailored to user, use case, device and more	Adapt your IAM solution continually enhance your customer experience and security posture while adjusting to rapid changes in your organization without significant programming investment.

Table 9-3: Customizability of Your IAM Solution

Evaluate every IAM solution to determine how much efficiency is improved in programming effort, especially when adapting rapidly to regulatory compliance and audit requirements by regulators such as, FedRAMP, SOX, CCPA, GDPR, etc. Also consider adopting a Zero Trust security model as you enhance your IAM solutions.

Ease of use

Your IAM solution should be user-friendly, so you create a positive experience for your workforce as well as your customers and vendors. Ensure that the solutions

Enterprise Data Security for US Europe and Asia

you choose are easily customizable and maintainable so that they easily integrate with all of your applications, users, and devices across their identity lifecycle across all environments including on-premises and clouds. Ensure that they provide a good end-user experience and provide self-help guidance to reduce friction in its use.

Your Solution Should Provide	Target Result
<ul style="list-style-type: none">• Users should have access to self-service tools to help them bootstrap identity into any development project quickly and efficiently and the ability to automate Identity centric processes.• It should provide integration wizards and a centralized console for administrators to manage all users, apps, policies, and identity and access management controls and passwordless authentication.• Administrators and IAM solution developers should have Widgets, APIs & SDKs which span user authentication, resource configuration and access controls and should have access to a comprehensive set of resources without code• Insights and visibility to stay on top of security with user, app, and device level activity and reporting	Allow enhancing the security posture through better IAM Solution development tools and operational controls for administrators and IAM Solution developers to expedite and automate enhancements, and additionally include user self-help to manage identity to drive increased adoption

Table 9-4: Usability Features of Your IAM Solution

Note that ease of use is extremely important to your organization's IAM solution to drive user adoption. Furthermore, customization and change over time is a fact of life and you should plan on it in selecting the tools and their implementation.

Reliable and secure

Your IAM solution must be trustworthy: reliable and secure because IAM is a business-critical function that is built on a foundation of trust. Your strategic IAM solution safely connects people to technology and should flawlessly manage your organization's most valuable data, their user identities and credentials. Trust can be lost instantly due to security risks that result in a data exposure.

Your IAM Solution Should Provide	Target Result
<ul style="list-style-type: none"> • Your goal should be 99.99% SLA for uptime so you should have zero planned downtime achieved through self-healing nodes. • Your shared security responsibility model should include your workforce, vendors, cloud service providers and customers. Each should know their role. • Your security analytics tools should proactively recommend how to improve security posture and define security perimeters around which access can be limited or restricted 	<p>It is important to build trust with both the workforce as well as your customers by ensuring a high uptime with an enhanced security posture that builds confidence across the key stakeholders,</p>

Table 9-5: Security and Reliability Characteristics of Your IAM Solution

You need to determine the key drivers for enhancing user experience and customer confidence in your IAM Solution, and the effect of frequent and/or prolonged system outages causing lost productivity, revenue, or customers.

IAM Across the Enterprise

IAM is responsible for managing identities and controlling access to an organization’s systems, applications, and data while Active Directory (AD) is a centralized directory service that stores and manages information about users and other assets in and/or on a network, such as their role and associated network privileges.

Lightweight directory access protocol (LDAP) is a protocol that helps users find data about organizations, persons, and more. LDAP has two main goals: to store data in the LDAP directory and authenticate users to access the directory.

Microsoft has long used Active Directory (AD) not just for single sign-on but for managing policies. SharePoint permissions are a good example of such policy driven adoption but you should use AD more directly: permitting access to email servers, VPN servers, and so on. As Microsoft moves AD to the cloud via Azure Active Directory (Azure AD) and managed AD on other public clouds, the opportunity to take advantage of policies is more apparent. For example, with IT concerned over access to and distribution of content and apps in mobile cloud settings, a cloud-based identity manager like Azure AD for more than mere sign-on validation becomes appropriate.

Enterprise Data Security for US Europe and Asia

Identity based security has existed for a long time. It has increased in adoption in recent times starting with the move to more distributed and complex systems and in the early 2020s due to the increasing partial or full migration to applications from on-premises to clouds. Adding value to the concept of identity-based security is the notion of centralized identity management, or what is called centralized trust. In short, this is the ability to provide credential validation services delivered from a central source.

The concept of centralized trust based on a central source is that each “actor” in a system — device, person, database, server, or queue — goes to the central database of all identity servers to validate its credentials and to be allowed access. This approach allows the ability to have common identity validation for systems both inside and outside the enterprise, such as those hosted on public clouds. Solving the identity problems on a central basis such as identifying and neutralizing security problems offers the ability to spend less on enterprise security by relying on the centralized trust model to deal with identity management across external and internal systems.

A secure, robust, and effective identity and access management system is one of the most important investments your organization can make. A common question often asked is, “Should these systems be deployed on-premises, or in the cloud?”

Depending on your specific needs, an IAM solution can be deployed either on-premises, or in the cloud, or both. Each of these scenarios comes with certain advantages or tradeoffs that need to be considered. On-premises systems are those hosted within your organization’s own infrastructure, whereas a cloud deployment means that the resources (i.e., software, servers, data, etc.) are accessed as services delivered over the internet. Cloud IAM would be provided by third-party vendors, typically in a subscription model, with offerings often referred to as Identity-as-a-service (IDaaS).

On-Premises IAM

For your on-premises identity and access management system, your organization will own and be responsible for all aspects of the solution, from design to deployment, maintenance, and improvements, hosted and usually accessed within your network, and all in your own datacenters or even in a third-party hosted data center. It involves deploying hardware, software, and dedicated servers on-site to manage user identities, access privileges, and authentication. You will be managing user identities and access privileges for your organization within your physical infrastructure, which gives your organization complete control over your data, and

you can customize your IAM solutions to meet your specific cybersecurity, data security and regulatory compliance needs.

GCP IAM

For GCP IAM , you need to use Google *Cloud Identity*, Google Cloud's built-in managed identity to create or sync user accounts across applications and projects. It provides the capability to provision and manage users and groups, set up single sign-on, and configure two-factor authentication (2FA) directly from the Google Admin Console. You also get access to the Google Cloud Organization, which enables you to centrally manage projects via the *Resource Manager*. The Resource Manager is designed to hierarchically manage resources by project, folder, and organization.

Google Cloud Identity is A unified identity, access, application, and endpoint management (IAM/EMM) platform. Cloud Identity integrates with hundreds of cloud applications out of the box. It is Google's approach to a secure, and flexible approach to identity and device management. Out of the box it provides:

- *Single sign-on (SSO)*. It improves employee experience and increases productivity with one-click access to thousands of pre-integrated apps, both in the cloud and on-premises.
- *Multi-factor authentication (MFA)*. It protects your user accounts and company data with a wide variety of MFA verification methods such as phishing-resistant security keys, mobile push notifications, and one-time passwords (OTPs).
- *Mobile device management (MDM)*. It keeps employees productive and data more secure with mobile management for Android and iOS devices.

You can set up user and group provisioning between Microsoft AD and your Cloud Identity or Google workspace account by using Google Cloud Directory Sync (GCDS).

Google's *Workforce Identity Federation* lets you use an external identity provider (IdP) to authenticate and authorize a workforce—a group of users, such as workforce members, vendors, partners, and contractors—using IAM, so that the users can access Google Cloud services. Workforce Identity Federation uses an identity federation approach instead of directory synchronization, eliminating the need to maintain separate identities across multiple platforms. It supports attributes defined in external identity provider and uses the attribute information to determine the scope of user access to Google Cloud resources.

AWS IAM

AWS IAM is the AWS provided IAM solution. When you first create an AWS account, you begin with a single sign-in identity that has complete access. You use IAM to control who is authenticated (signed in) and authorized (has permissions)

Access to AWS resources and use of AWS is managed through AWS IAM authentication. It manages access to all AWS services and resources in the account. This identity is called the AWS account root user and is accessed by signing in with the email address and password that you used to create the account.

AWS provides Secure access to AWS resources for applications that run on Amazon EC2. You can use IAM features to securely provide credentials for applications that run on EC2 instances. These credentials provide permissions for your application to access other AWS resources. AWS IAM provides the following features:

- *Shared access* to your AWS account. You can grant other people permission to administer and use resources in your AWS account without having to share your password or access key.
- *Granular permissions*. You can grant different permissions to different people for different resources. For example, you might allow some users complete access to Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), Amazon DynamoDB, Amazon Redshift, and other AWS services. For other users, you can allow read-only access to just some S3 buckets, or permission to administer just some EC2 instances, or to access your billing information but nothing else.
- *Multi-factor authentication (MFA)*. You can add two-factor authentication to your account and to individual users for extra security. With MFA you or your users must provide not only a password or access key to work with your account, but also a code from a specially configured device.
- *Identity federation*. You can allow users who already have passwords elsewhere—for example, in your corporate network or with an internet identity provider—to get temporary access to your AWS account.

Additional features of AWS IAM include the following:

- If you use AWS CloudTrail, you receive log records that include information about those who made requests for resources in your account. That information is based on IAM identities.
- PCI DSS Compliance. AWS IAM supports the processing, storage, and transmission of credit card data by a merchant or service provider and has

Enterprise Data Security for US Europe and Asia

been validated as being compliant with Payment Card Industry (PCI) Data Security.

- Standard (DSS). For more information about PCI DSS, including how to request a copy of the AWSPCI Compliance Package.

AWS IAM is integrated with many AWS services.

AWS IAM is designed to be *Eventually Consistent*. Like many other AWS services, AWS IAM is eventually consistent such that the update will migrate to all other regions that you have subscribed to. IAM achieves high availability by replicating data across multiple servers within Amazon's data centers around the world. If a request to change some data is successful, the change is committed and safely stored. However, the change must be replicated across IAM in all active regions for your accounts, which can take some time. Such changes include creating or updating users, groups, roles, or policies. We recommend that you do not include such IAM changes in the critical, high-availability code paths of your application. Instead, make IAM changes in a separate initialization or setup routine that you run less frequently. Also, be sure to verify that the changes have been propagated before production workflows depend on them.

AWS Identity and Access Management (IAM) and AWS Security Token Service (AWS STS, an AWS web service that enables you to request temporary limited-privilege credentials for users) which are features of your AWS account offered at no additional charge. You are charged only when you access other AWS services using your IAM users or AWS STS temporary security credentials.

Federating IAM

Federated identity allows authorized users to access multiple applications and domains using a single set of credentials and links the user's identity across multiple identity management systems.

FID (Federated Identity and Directory Service) refers to the use of a directory service such as AD for federated identity. An FID provides the following capabilities:

- Creates a common logical and up to date view and storage of all identity sources, typically in directory system capable of the High Availability Directory Access Protocol (HDAP).
- Acts as the identity hub to authenticate a user and retrieve profile information (attributes and/or groups) for authorization.

Enterprise Data Security for US Europe and Asia

- Is the global reference image for Identity Governance and Administration (IGA) to provision and administrate on-premises and cloud applications.
- Saves time by simplifying the management of users and groups (automating configuration across multiple identity systems instead of scripting and endless customization)
- Supports stronger authentication and security methods—MFA (specific to gov/defense sector) PIV/CAC cards, biometrics, FIPS framework, etc.

CFS (Cloud Federation Service):

Provides a secure Federated Access Management (supporting SAML and OpenID Connect) for inside-out access (IdP-initiated) or outside in access (SP-initiated). Creates one access and audit point to connect all your internal identity and authentication sources to the cloud.

Federating Access

Through proprietary or standard based protocols, access can be federated but you need to install federated identity to be able to use it effectively.

Identity Life Cycle

Identity Governance Administration (IGA) tracks and governs the user identity through its lifecycle. The following describes the key stages of the identity lifecycle including:

1. Identity Creation.
2. Entitlements management.
3. Policy and role management.
4. Access request certification.
5. Fulfillment and provisioning.
6. Audit reporting.
7. Identity analytics and machine learning.

Identity Creation

Every entity that connects to your organization's network to perform its tasks and functions must have an Identity that is created, provided data access and role

entitlements, governed through its life, and destroyed when no longer needed. Identities should be created for the following entities:

- Workforce.
- Contractors.
- Customers.
- Nonhumans (RPA bots, service accounts).

As part of Identity creation, it is also important to manage identities for new joiners such as new workforce members customers, vendors, etc. Additionally as people move (due to job transfers) or leave (such as due to resignations), their Identity needs to be updated.

Entitlements Management

Once an Identity has been created, entitlements such as access permissions need to be added to the entity. For example, for a workforce member, their job defines a basic set of permissions to common corporate websites and capabilities. In addition to the basic set, their specific job functions within the job classification or the project they are working on may entitle them to additional privileges and permissions.

IGA is concerned with discovering the privileges and entitlements of all entities in the organization that have an Identity. Fine-grained entitlements are more specific privileges and permissions for entities that play a special role or are working on a special project.

Policy and Role Management

We have seen the definition of Roles and Groups earlier in this chapter. It should be noted that there are two dimensions to roles that differentiate them. These are business roles (based on the business function hat of the workforce user) and resource roles (also known as buckets) used by an application for performing a task on behalf of a user.

Assigning Policies to roles and groups management ensures consistency across the organization and facilitates consistent management of entitlements for the roles and groups in terms of provisioning resources. Policies help in simplifying entitlements management, and mitigating security risk by minimizing the potential of incorrect allocation of privileges and entitlements. While also improving Identity and entitlement efficiency, a major role Policies play is in demonstrating compliance, especially if required by regulators.

Assigning roles and entitlements follow a well-defined, and preferably a documented workflow that steers requests for approvals by managers and tracks status. The workflow should include any delegation or escalations due to improperly constructed requests.

Access Request Certification

Access Requests are often managed as annual campaigns to update all workforce members as well as on an as needed basis. The managers certify role compositions and entitlements through an automated system (or via emails in some not so automated systems).

Fulfillment and Provisioning

Once the identity has been created, the next step is to start provisioning. This can be direct via email, via service desk, a SaaS access management tool, or other means.

Auditing and Reporting

Auditing is a very important function to ensure security by regularly checking to see if the Identity is correctly provisioned based on the role of the entity and updated as per the changes in the person's job role. Auditing may be a specific requirement by regulatory agencies. Policy based controls and monitoring are used to check if there are any segregation of duties violations such as by developers accessing production servers without specific documented privilege escalation via a jump server. It also checks integrity of the attributes associated with the Identity, and presence of compensating controls in case there are no direct controls in place.

Identity Analytics and Machine Learning

The final set is Identity Analytics to track who used the Identity, when it was used, the purposes for which it was used such as what data was accessed. Specific topics for analytics include the following:

- Risk modeling.
- Clean-up analytics which include login analysis and dormant and orphan account discovery and peer analysis and manage an access outliers report.
- Governance on-the-fly with advanced role mining and recommendations, policy analysis and modeling, and policy effectiveness and violation analysis

Identity Governance and Administration

Identity Governance and Administration (IGA) consists of two parts, Identity Governance and Identity Administration. Broadly speaking, Identity Governance is about visibility, segregation of duties, role management, attestation of access permissions, analytics and reporting, while Identity Administration refers to account administration, credentials administration, user and device provisioning and managing entitlements.

Compromised identities caused by weak, stolen or default user credentials are a growing threat to organizations. Centralized visibility creates a single authoritative view of “who has access to what,” allowing authorized administrators to promptly detect inappropriate access, policy violations or weak controls that put organizations at risk.

With IGA solutions, security personnel can track and control user access for both on-premises and cloud-based systems as part of the cloud governance efforts. They can secure users by ensuring that the right user accounts have the right access to the right systems and detect and prevent inappropriate access. By implementing the right controls with IGA, your organization can minimize risk and maintain regulatory compliance.

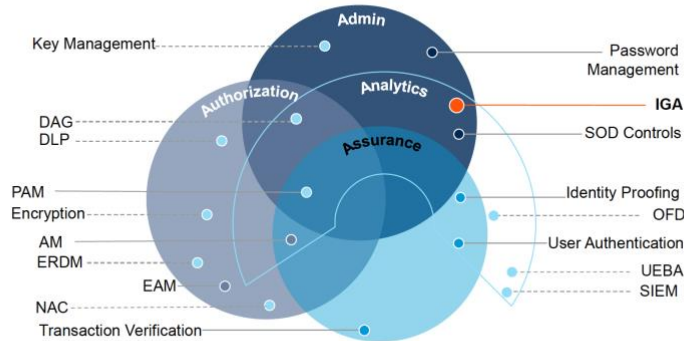
IGA tools are vital if your organization, due to the nature of its business, is required to meet stringent regulatory requirements. They provide a centralized view of access rights across your organization, simplifying the monitoring, reporting, and adjustment of permissions.

IAM vs IGA

While IAM primarily focuses on granting access rights, IGA expands the purview by encompassing crucial identity governance aspects such as role-based access control, segregation of duties, and policy enforcement. Please see the following Gartner diagram which shows the differences and the overlap between IAM and IGA.

IGA — In relation to IAM

IGA and the Bigger Scope of IAM



Gartner.

Diagram source: Gartner

Figure 9-1: IGA and IAM Relationship and Overlap

Governance Model

A governance model is a framework that outlines an organization's general leadership accountabilities and describes how leaders and members interact with other parties.

It is important to adopt an IGA Governance Model that meets the requirements of your organization. Identity governance solutions enable business and IT users to identify risky workforce populations, policy violations and inappropriate access privileges and to remediate these risk factors. The model should allow security administrators to efficiently manage user identities and access across the enterprise. It should improve their visibility into identities and access privileges and help them implement the necessary controls to prevent inappropriate or risky access.

The Governance Model should be designed for your organization's security leadership to verify that the right controls are in place to meet the security and privacy requirements of regulations like SOX, HIPAA and GDPR.

They provide consistent business processes for managing passwords and passphrases, as well as reviewing, requesting and approving access, all underpinned by a common policy, role and risk model. With role-based access control,

companies significantly reduce the cost of compliance, while managing risk and establishing repeatable practices for a more consistent, auditable and easier-to-manage access certification efforts.




Figure: 9-2: IGA Governance Model

The diagram above shows the role of Identity Analytics in the governance and management of entity accounts as part of IGA and compliance management.

Operational Efficiency and Excellence

The goal of operational efficiency and excellence is to reduce manual administration and streamline business processes by centralizing access controls driven through centralized identity repositories.

Operational efficiency also entails business enablement to empower the departmental business functions to use self-service tools to enhance their own user experiences and have easy access to the resources they are entitled to at the right time.

Compliance Management

A very important role played by IGA is overall compliance management, especially if your organization is subjected to regulatory compliance. The compliance goals are to ensure that segregation of duties, if required, is maintained and all user access beyond the least privilege is certified for the appropriateness and documented so it can be tracked through forensic tools. The IGA tools provide oversight that ensures accountability of user access controls.

Chapter 10 – Application Security Architecture and Data Exploits

Application security, commonly known in short as *AppSec*, secures software at the application level. It is the practice of using security software, hardware, techniques, best practices and procedures to protect computer applications from external security threats and aim to prevent data or code within the app from being stolen or hijacked.

Secure development practices include security measures and improving security practices in the software security lifecycle. The goal of AppSec is to minimize the likelihood that malicious actors can gain unauthorized access to systems, applications or data. The ultimate goal of application security is to prevent attackers from accessing, modifying or deleting sensitive or proprietary data.

Because most vulnerabilities are introduced during the development and publishing stages, application security includes many types of cybersecurity solutions to help identify flaws during the design and development phases that could be exploited and alert teams so they can be fixed.

Any action taken to ensure application security is called a countermeasure or security control. The National Institute of Standards and Technology (NIST) defines a security control as: "A safeguard or countermeasure prescribed for an information system or an organization designed to protect the confidentiality, integrity, and availability of its information and to meet a set of defined security requirements."

Application security controls are techniques to enhance the security of an application at the coding level, making it less vulnerable to threats. Many of these controls deal with how the application responds to unexpected inputs that a cybercriminal might use to exploit a weakness.

There are various kinds of application security programs, services, and devices an organization can use. Firewalls, antivirus systems, and data encryption are just a few examples to prevent unauthorized users from entering a system. Application security is the process of developing, adding, and testing security features within applications to prevent security vulnerabilities against threats such as unauthorized access and modification.

Modern application security uses controls and code scans to ensure that users can only access files and data based on defined authorization rules and cannot elevate their privileges using various exploits.

Vulnerability Management tracks risks across the enterprise (i.e. patch levels, exposed applications, vulnerabilities, etc.). This chapter describes threats and vulnerabilities, and the approaches and tools used to manage vulnerabilities.

Threat Models and Attacks

An important step in understanding how to achieve application security is to identify the threats faced by the application data. Threat modeling is a method of optimizing network security by locating vulnerabilities, identifying objectives, and developing countermeasures to either prevent or mitigate the effects of cyber-attacks against the system.

Threat modeling is a structured process with these objectives: identify security requirements, pinpoint security threats and potential vulnerabilities, quantify threat and vulnerability criticality, and prioritize remediation methods.

Threat modeling methods create an abstraction of the system and the profiles of potential attackers. Based on that the model catalogs potential threats that these attackers pose.

Application Attack Points

The impact from an attack risk being realized is measured in financial and reputation terms, and not in the amount of data or property lost, or number of people affected. Ultimately the financial loss and loss of reputation are the key drivers for your business.

In modern application development, the virtual machines (VMs) play an important role. A Virtual Machine (VM) is a compute resource that uses software abstraction instead of a physical computer to run programs and deploy applications. It provides memory and storage, and it can be moved from one host to another host on-premises or in the cloud by the cloud service provider. One or more virtual “guest” machines run on a physical “host” machine. Each VM runs its own operating system and functions separately from the other VMs, even when they are all running on the same host. This means that, for example, a virtual MacOS virtual machine can run on a physical PC. VMs are used on on-premises hosts as well as in cloud environments.

A hypervisor, also known as a virtual machine monitor or VMM, is software that creates and runs virtual machines (VMs). A hypervisor enables computer virtualization by creating, running and managing resources on virtual machines (VMs). A hypervisor allows one host computer to support multiple guest VMs by virtually sharing its resources, such as memory and processing.

Because hypervisors are software, they run the risk of malicious attacks that can compromise data security.

Types of Hypervisors

Hypervisors can be Type 1 which run directly on the host computer and are often called bare metal hypervisors or Type 2 where the hypervisor is hosted on another operating system such as VMware. You may want to study the hypervisor architecture in-depth on reference material outside the scope of this book.

The attack surface presented to a hacker depends on the type of hypervisor. Type 2 has greater attack surface because of the underlying operating system.

Mitigating Attack Incidents

Events are anything that can occur in an IT environment; incidents are unscheduled events. The concept of two-person integrity (commonly used in banks, for example, for vault management) forces collusion between two parties for unauthorized access. The chances of two bad actors colluding are significantly lower than a single bad actor. An example of this is escalation of privilege that requires a manager to approve access to production environments.

The following describes examples of common attacks and how they can be mitigated from becoming incidents from a data security perspective.

- SQL Injection threat can be reduced by DAMs (Database Activity Monitors which can be either agent based or network monitors).
- An XML Gateway can perform content inspection on SFTP (Secure File Transfer Protocol) and other protocol traffic and can act as a reverse proxy. WAF (Web Access Firewall) and DAM do not handle SFTP traffic.
- Denial of Service can be addressed by WAF, a Layer 7 firewall that understands HTTP traffic and is effective on DoS.
- An *XML Accelerator* is often used to assist legacy applications that may not have the programmatic capability to process assertions from modern web services; but are not effective on XML firewall or WAF.
- An API Gateway can filter API traffic and can implement access control rate limiting, logging, metrics, and security filtering.

- Firewalls in cloud environments are typically software firewalls and can be configured across SaaS, PaaS, and IaaS. Note that loss of availability due to Distributed Denial of Service (DDoS) is not an issue that you need to address for your cloud environments because your cloud service provider addresses that.
- Threat monitoring tools (SIEM, SEM, SIM) aid in identifying and mitigating threats and also optimizing performance; but they do not help in decreasing the size of log files or reducing work for the production personnel.
- RUM (Real User Monitoring) tools capture real-time performance data about how users interact with your web applications. They keep track of important metrics to help you understand how well your websites perform for users in different parts of the world. Customers may raise privacy issues with RUM due to surveillance so synthetic monitoring that simulates RUM may be preferable to RUM because it is faster and more comprehensive, but it is less accurate than RUM and more expensive to generate the traffic test data.
- Inference attacks are inferred from resource calls if the CSP does not secure the hypervisor sufficiently and in that one user on a VM can see the resource calls of the user's VM and infer activity from the volume of traffic. This could be used for a social engineering attack.
- STRIDE (S- Spoofing, T-Tampering, R-Repudiation, I-Information Disclosure, D-Denial of Service, E-Elevation of Privilege) is a model for identifying the type of security threat and helps determine the types of mitigation required.
- Whether on-premises or cloud, physical and logical access control methods are necessary. Your organization manages your internal controls and works with your managed service provider and even your customer for external controls.
- DNS Attacks: DNSSEC (DNS Security) adds security to the DNS responses to be validated. DNSSEC provides origin authority, data integrity, and authenticated denial of existence. Protects against DNS spoofing attacks. DNS attacks include: Foot-printing (DNS domain name exfiltration), DoS on DNS, Data modification of IP addresses, redirection to DNS names of attacker's servers, Spoofing (malicious cache poisoning).

Comprehensive Application Security Framework

To protect information, a solid, comprehensive application security framework is needed for analysis and improvement. This application security framework should

Enterprise Data Security for US Europe and Asia

be able to list and cover all aspects of security at a basic level. It should incorporate security controls to address, at a minimum, the following:

- Security elements that need to be preserved: availability, utility, integrity, authenticity, confidentiality, nonrepudiation
- Sources of loss of these elements: abuse, misuse, accidental occurrence, and natural forces.
- Acts that cause loss: use of false data, disclosure, interference with use, copying, misuse or failure to use correct information.
- Safeguard functionality used to protect from acts that cause loss including audits, risk avoidance, detection, prevention, recovery, mitigation, and investigation.
- Methods you use for selecting safeguard functionality such as due diligence on compliance with regulations and standards and whether they meet the business needs.
- Objectives to be achieved by the application security framework such as avoiding negligence in configurations and security controls setup, protection of data privacy, and minimizing performance impact.

The following sections describe the security elements that need to be preserved. Each of the six elements can be violated independently of the others. The elements are unique and independent and often require different security controls. Maintaining availability of information does not necessarily maintain its utility; information may be available, but useless for its intended purpose.

Availability

Availability for computer systems is the ability to access information or resources in a specified location and in the correct format. When a system is regularly not functioning, information and data availability is compromised and it will affect the users. Besides functionality, another factor that effects availability is time. If a computer system cannot deliver information efficiently, then availability is compromised again. Data availability can be ensured by storage, which can be local or offsite.

Utility

While utility is not considered a pillar in information security in a scenario like where you encrypt the only copy of valuable information and then accidentally delete the encryption key. The information in this scenario is available, but in a form that is not useful. To preserve utility of information, you should require mandatory backup copies of all critical information and should control the use of

protective mechanisms such as cryptography. Test managers should require security walk-through tests during application development to ensure unusable forms of information.

Integrity

Integrity for our context refers to methods of ensuring that the data is real, accurate and guarded from unauthorized user modification. Data integrity is a major information security component because users must be able to trust information. Untrusted data compromises integrity. Stored data must remain unchanged within a computer system, as well as during transport. It is important to implement data integrity verification mechanisms such as checksums and data comparison.

Authenticity

Authenticity of users accessing data, called authentication, refers to a process that ensures and confirms the user's identity. The process begins when the user tries to access data or information. The user must prove access rights and identity. Commonly, usernames and passwords are used for this process. However, this type of authentication can be circumvented by hackers. A better form of authentication is biometrics, because it depends on the user's presence and biological features (retina or fingerprints). The PKI (Public Key Infrastructure) authentication method uses digital certificates to prove a user's identity. Other authentication tools can be key cards or USB tokens. Unsecured emails that seem legitimate pose a great security risk.

Confidentiality

Confidentiality of data ensures that the data remains private or secret and access to sensitive and protected information is controlled. Sensitive information and data should be disclosed to authorized users only. Confidentiality can be enforced by using a classification system. The user must obtain a certain clearance level to access specific protected data or information. Confidentiality can be ensured by using role-based security methods to ensure user or viewer authorization (data access levels may be assigned to a specific department) or access controls that ensure user actions remain within their roles (for example, define user to read but not write data).

Nonrepudiation

Nonrepudiation refers to a method of guaranteeing message transmission between parties using digital signature and/or encryption. Proof of authentic data and data origination can be obtained by using a data hash. Nonrepudiation can be achieved

by using digital signatures to prove the delivery and receipt of messages where the sender cannot deny having sent the message,

Data Exploits and Security Vulnerability

A security exploit is an unintended and unpatched flaw in software code that exposes it to potential exploitation by hackers or malicious software code such as viruses, worms, Trojan horses and other forms of malware.

Security exploits may result from a combination of software bugs, weak passwords or software already infected by a computer virus or worm, and misconfigurations. Mitigating these security exploits require patches, or software fixes, to prevent the potential for unauthorized access or compromised integrity.

A security vulnerability is a weakness an adversary could take advantage of to compromise the confidentiality, availability, or integrity of a resource. In this context a weakness refers to implementation flaws or security implications due to design choices. For instance, being able to overrun a buffer's boundaries while writing data to it introduces a buffer overflow vulnerability.

Zero-Day Vulnerability

A zero-day vulnerability is a newly discovered software security flaw or vulnerability that is known to the software vendor but doesn't have a patch in place to fix the flaw. It has the potential to be exploited by cybercriminals because an official patch or update to fix the issue hasn't been released. The "zero-day" refers to the fact that the developers have "zero days" to fix the problem that has just been exposed — and perhaps already exploited by hackers. Once the vulnerability becomes publicly known, the vendor must work quickly to fix the issue to protect its users. But the software vendor may fail to release a patch before hackers manage to exploit the security hole. That's known as a zero-day attack

There are several public vulnerability repositories available that allow you to have easy access to information regarding known vulnerabilities. The most prominent vulnerability repositories are CVE, NVD and OVAL. CVE has established a referencing system for registering vulnerabilities called the CVE identifier (CVE-ID). CVE-IDs usually include a brief description of the security vulnerability and sometimes advisories, mitigation measures and reports.

Vulnerability Management

Vulnerability management identifies, classifies, evaluates, and mitigates vulnerabilities. IT security professionals perform the vulnerability management process in an organized and timely manner by following the steps described below:

1. *Preparation.* Define the scope of the vulnerability management process.
2. *Vulnerability Scanning.* Vulnerability scanners are automated tools that scan a system for known security vulnerabilities providing a report with all the identified vulnerabilities sorted based on their severity. Known vulnerability scanners include Nexpose, Nessus and OpenVAS.
3. *Identification, Classification and Evaluation of the Vulnerabilities.* The vulnerability scanner provides a report of the identified vulnerabilities.
4. *Remediating Actions.* The asset owner determines which of the vulnerabilities will be mitigated.
5. *Rescan.* Once the remediating actions are completed, a rescan is performed to verify their effectiveness.

Penetration Testing

Penetration testing is the assessment of the security of a system against different types of attacks performed by an authorized security expert. The tester attempts to identify and exploit the system's vulnerabilities. The difference between a penetration test and an actual attack is that the former is done by a tester who has permission to assess the security of the system and expose its security weaknesses. In addition the tester is given certain boundaries to operate and perform this task.

There exists some confusion in the mind of the public over penetration testing and vulnerability scanning. The two approaches actually complement each other, with vulnerability scanning being one of the first steps of a penetration test.

Exploits and Exploit Prevention

An *exploit* is a program, or piece of code, designed to find and take advantage of a security flaw or vulnerability in an application or computer system, typically for malicious purposes such as installing malware. An exploit is not malware itself, but rather it is a method used by cybercriminals to deliver malware.

An **exploit kit** is simply a collection of exploits. Exploit Kits are tools embedded in compromised web pages which automatically scan a visitor's machine for vulnerabilities and attempt to exploit them. If the exploit succeeds the kit injects malware to the user's system. The ease of use and the friendly interface of many Exploit Kits allow non-expert users to deploy them as well.

Exploit Classification

There are several methods of classifying exploits. The most common is by how the exploit communicates to the vulnerable software. A *remote exploit* works over a network and exploits the security vulnerability without any prior access to the vulnerable system. A *local exploit* requires prior access to the vulnerable system and usually increases the privileges of the person running the exploit past those granted by the system administrator. Exploits against client applications also exist, usually consisting of modified servers that send an exploit if accessed with a client application.

Exploits against client applications may also require some interaction with the user and thus may be used in combination with the social engineering method. Another classification is by the action against the vulnerable system; unauthorized data access, arbitrary code execution, and denial of service are examples.

Many exploits are designed to provide superuser-level access to a computer system. However, it is also possible to use several exploits, first to gain low-level access, then to escalate privileges repeatedly until one reaches the highest administrative level (often called "root").

Mitigating Exploits

After an exploit is discovered, the vulnerability often needs to be fixed through a patch to make the exploit unusable. That is the reason why some black hat hackers as well as military or intelligence agencies hackers do not publish their exploits but keep them private.

Pivoting

Pivoting refers to a method used by penetration testers to simulate an attack that uses a compromised system to attack other systems on the same network to avoid restrictions such as firewall configurations, which may prohibit direct access to all machines. For example, if an attacker compromises a web server on a corporate network, the attacker can then use the compromised web server to attack other systems on the network. These types of attacks are often called multi-layered attacks. Pivoting is also known as *island hopping*.

Pivoting can further be distinguished into *proxy* pivoting and *VPN* pivoting. *Proxy pivoting* generally describes the practice of channeling traffic through a compromised target using a proxy payload on the machine and launching attacks from the computer. This type of pivoting is restricted to certain TCP and UDP ports that are supported by the proxy.

VPN pivoting enables the attacker to create an encrypted layer to tunnel into the compromised machine to route any network traffic through that target machine; for example, to run a vulnerability scan on the internal network through the compromised machine, effectively giving the attacker full network access as if they were behind the firewall.

10 Most Common Web Security Vulnerabilities

For your organization, it is important that web security best practices become a priority before a breach happens. An effective approach to web security threats must, by definition, be proactive and defensive. The following describes the 10 common and significant web security pitfalls to be aware of, including recommendations on how they can be mitigated. The focus is on the *Top 10 Vulnerabilities* identified by the *Open Web Application Security Project (OWASP)*, an international, non-profit organization whose goal is to improve software security across the globe.

OWASP 10

The following lists the ten common web security mistakes listed in the 2021 update of OWASP for the Common Web Security Mistakes that can compromise your data security and expose your data to unauthorized exfiltration.

A01-2021: Broken Access Control

This was determined to be the vast majority of failures and has been bumped up from previous versions of OWASP. The following is a sampling of the issues that were discovered as part of broken access control. New paragraph

- Violation of the principle of least privilege or deny by default.
- Bypassing access control checks by modifying the URL, internal application state, or the HTML
- Permitting viewing or editing someone else's account by providing its unique identifier.
- Accessing API with missing access controls.
- Elevation of privilege without properly logging in.
- Metadata manipulation such as replaying or tampering with a JSON web token.

- CORS (Cross Origin Resource Sharing) misconfiguration allows the API access from unauthorized or untrusted origins.
- Force browsing to authenticated pages as an unauthenticated user or to privileged spaces as a standard user.

This is a classic case of trusting user input and paying the price in a resulting security vulnerability. A direct object reference means that an internal object such as a file or database key is exposed to the user. The problem with this is that the attacker can provide this reference and, if authorization is either not enforced (or is broken), the attacker can access or do things that they should be precluded from.

Another common vulnerability example is a password reset function that relies on user input to determine whose password we're resetting. After clicking the valid URL, an attacker can just modify the `username` field in the URL to say something like "admin".

Accessing API with missing access controls also is an authorization failure. It means that when a function is called on the server, proper authorization was not performed. A lot of times, developers rely on the fact that the server side generated the user interface, and they think that the functionality that is not supplied by the server cannot be accessed by the client. It is not as simple as that, as an attacker can always forge requests to the "hidden" functionality and will not be deterred by the fact that the user interface doesn't make this functionality easily accessible.

Prevention: except for public resources, denied by default, and implement access control mechanisms once and use them throughout the application including minimizing CORS. Perform user authorization properly and consistently and whitelist the choices. More often than not though, the whole problem can be avoided by storing data internally and not relying on it being passed from the client via common gateway interface (CGI) parameters. Session variables in most frameworks are well suited for this purpose.

On the server side, authorization must *always* be performed and stateful identifiers should be invalidated on the server after logout. JWT tokens should be short lived and for long lived tokens follow the OAuth standards to revoke access.

A02-2021: Cryptographic Failures

Sensitive data should be encrypted at all times, including in transit and at rest. Credit card information and user passwords should *never* travel or be stored unencrypted, and passwords should always be hashed. Obviously the crypto/hashing algorithm must not be a weak one – when in doubt, web security standards recommend at minimum AES 256 and RSA 2048 bits. Furthermore,

session IDs and sensitive data should not be traveling in the URLs and sensitive cookies should have the secure flag on.

Prevention:

- *In transit:* Use HTTPS with a proper certificates. Do not accept anything over non-HTTPS connections. Have the secure flag on cookies.
- *In storage:* First and foremost, you need to lower your exposure. If your organization does not need the sensitive data any more, you should shred it. Data you don't have can't be stolen. If you store credit cards, you need to be PCI-DSS compliant. So, if you have sensitive data that you actually do need to keep, store it encrypted and make sure all passwords are hashed. And *do not store the encryption keys next to the protected data.*

A03-2021: Injection

Injection flaws result from a classic failure to filter untrusted input. Some of the more common injections are SQL, NoSQL, OS command, Object Relational Mapping (ORM), LDAP, Expression Language (EL) or Object Graph Navigation Library (OGNL) injection. It can happen when you pass unfiltered data to the SQL server (SQL injection), to the browser (XSS), to the LDAP server (LDAP injection), or anywhere else. The problem here is that the attacker can inject commands to these entities, resulting in loss of data and hijacking clients' browsers.

Anything that your application receives from untrusted sources must be filtered, preferably according to a whitelist. You should almost never use a blacklist, as getting that right is very hard and usually easy to bypass because pattern matching does not work very well.

Prevention: protecting against injection requires filtering your input properly and determining whether an input can be trusted. Note that *all* input needs to be properly filtered, unless it can unquestionably be trusted. For example, in a system with 1,000 inputs, successfully filtering 999 of them is not sufficient, as this still leaves one field that can allow a hacker to bring down your system. And you might think that putting an SQL query result into another query is a good idea, as the database is trusted, but if the perimeter is not, the input comes indirectly from bad actors. This is called *second order SQL injection*. It is best to rely on your framework's filtering functions: they are proven to work and are thoroughly scrutinized.

Another example is Cross Site Scripting (XSS) which is a fairly widespread input sanitization failure. An attacker gives your web application JavaScript tags on input. When this input is returned to the user un-sanitized, the user's browser will

execute it. It can be as simple as crafting a link and persuading a user to click it, or it can be something much more sinister. On page load the script runs and, for example, can be used to post your cookies to the attacker.

Prevention: The simple web security solution is to not return HTML tags to the client. This has the added benefit of defending against HTML injection, a similar attack whereby the attacker injects plain HTML content (such as images or loud invisible flash players). And always use positive server-side input validation. Additionally, for any residual dynamic queries, escape special characters using the specific escape center syntax for that interpreter.

A04-2021: Insecure Design

This category of mistakes focuses on risks related to design and architectural flaws, with a call for more use of threat modeling, secure design patterns, and reference architectures. This is a broad category and representing a variety of weaknesses expressed as missing or ineffective control in the design note that design flaws and implementation defects have different roots carrot root causes and remediations. A secure design can still have implementation defects leading to vulnerabilities that may be exploited. Also note that insecure design cannot be fixed by a perfect implementation.

Prevention: It is important to left shift the design far enough to ensure that secure patterns and threat modeling are used as part of the development. Established a secure development life cycle and integrate security language and controls into all development user stories as well as checks at each tier of the application along with fully defined test use cases for every tier.

A05-2021: Security Misconfiguration

Web servers and applications can be misconfigured in many ways including the following:

Running the application with debug enabled in production.

1. Having directory listing enabled on the server, which leaks valuable information.
2. Running outdated software (think WordPress plugins, old PhpMyAdmin).
3. Having unnecessary services running on the machine.
4. Not changing default keys and passwords.
5. Leaving error handling information unprotected from attackers, such as stack traces.

Prevention: Have a good (preferably automated) “build and deploy” process, which can run tests on deploy. The poor man’s security misconfiguration solution is post-commit hooks, to prevent the code from going out with default passwords and/or development stuff built in.

A06-2021: Vulnerable and Outdated Components

This is really more of a maintenance/deployment problem where before incorporating new code, it has not been checked out thoroughly for known documented issues with it. Using code from an unreliable source can result in the risk of serious web security vulnerability.

The lesson here is that software development does not end when the application is deployed. There has to be documentation, tests, and plans on how to maintain and keep it updated, especially if it contains 3rd party or open source components.

Prevention:

Exercise caution. Beyond obviously using caution when using such components, do not be a copy-paste coder. Carefully inspect the piece of code you are about to put into your software, as it might be broken beyond repair (or in some cases, intentionally malicious—web security attacks are sometimes unwittingly invited in this way).

Stay up-to-date. Make sure you are using the latest versions of everything that you trust, and have a plan to update them regularly. At least subscribe to a newsletter of new security vulnerabilities regarding the product.

A07-2021: Identification and Authentication Failures

This is a collection of multiple problems that might occur during broken authentication, but they don’t all stem from the same root cause. Potential authentication pitfalls include:

- The URL might contain the session ID and leak it in the referrer header to someone else.
- Weak or predictable passwords used or stored in plain text. The passwords might not be encrypted either in storage or transit.
- Not using multifactor authentication.
- The session IDs might be predictable, thus making gaining access trivial.
- Session fixation (hijacking a valid user session) might be possible if the session ID is not managed properly by the web application.
- Session hijacking might be possible also if timeouts not implemented correctly or using HTTP (with no SSL security), etc.

Prevention: The most straightforward way to avoid this web security vulnerability is to use a framework. This should include:

- And force multifactor authentication.
- Do not ship or deploy with any default credentials, especially for admin users.
- Ensure compliance but NIST 800 63-B (Digital Identity Guidelines).
- Use server side secure, built-in session manager that generates a new random session ID after login.

A08-2021: Software and Data Integrity Failures

These failures relate to code and infrastructure that does not protect against integrity violations. An example of this is where an application relies upon plugins, libraries, or modules from untrusted sources, repositories, and content delivery networks (CDNs)

Prevention: some of the prevention mechanisms include using digital signatures or similar mechanisms to verify the source of the software and its authenticity. Also ensure that libraries are being consumed from trusted repositories, and implement a review process for code and configuration changes to minimize inadvertent entry of malicious code. Furthermore, ensure that CI CD pipelines properly segregated with full access controls and ensure that unsigned signed or unencrypted serialized data is not sent to untrusted clients.

A09-2021: Security Logging and Monitoring Failures

The goal of this category is to help detect, escalate, and respond to active breaches. Logging, detection, monitoring an active response are important for all auditable events such as logins, failed logins, high value transactions. Quite frequently warnings and errors generate none or unclear log messages. Log storage and review of alerting thresholds and response escalation processes are not often implemented. And penetrating testing and scans by dynamic application security testing tools are not properly managed.

Prevention: Ensure proper log management and review of logs and alerts. All high-value transactions should have an audit-trails and integrity controls. Monitoring should actively alert and prevent attacks in real-time. Furthermore, establish or adopt an incident response and recovery plan (see NIST 800-61r2 or later)

A10-2021 – Server-Side Request Forgery (SSRF)

SSRF flaws occur whenever a web application is fetching a remote resource without validating the user-supplied URL. It allows an attacker to coerce the

application to send a crafted request to an unexpected destination, even when protected by a firewall, VPN, or another type of network access control list (ACL).

Prevention: follow defense in depth controls and segment remote resource access functionality is separate networks. Also enforce denied by default firewall policies or network access control rules and manage all logs that are accepted and block blocked network flows on firewalls. From the application layer sanitize and validate all client supplied input data and enforce URL, schema, port, and destination with the positive allowed list. Furthermore, do not send raw responses to clients and disable all http://redirections.

Cloud Security Vulnerabilities

The Cloud Security Alliance (CSA) is dedicated to defining and raising awareness and best practices to help secure the cloud computing environment. CSA has aimed to provide organizations with up-to-date, expert-informed understanding of cloud security threats in order to make educated risk-management decisions regarding cloud adoption strategies. The CSA Notorious Nine list of threats specifically defines threats related to the shared, on-demand nature of cloud computing. The following section describes the Treacherous 12 defined in 2016.

CSA Treacherous Twelve

CSA published the “The Treacherous 12 - Cloud Computing Top Threats in 2016” to provide organizations with an up-to-date, expert-informed understanding of cloud security concerns in order to make educated risk-management decisions regarding cloud adoption strategies. The following describes the threats and vulnerabilities covered in CSA Treacherous Twelve:

1. *Data Breach* - org’s sensitive data falls into the hands of competitors or malicious outsiders, Data breaches can be internal or external.
2. *Insufficient Identity, Credential and Access Management (new)* – can cause data breaches when attacks occur due to lack of scalable IAM systems, failure to use MFA, weak password use, lack of ongoing password rotation or cryptographic keys, passwords and certificates. It recommends to not embed keys in public facing repositories such as GitHub and to store and protect keys appropriately. Also de-provision accounts no longer needed because they can be used by malicious actors. You need to periodically test and monitor smartcard, OTP, MFA, phone authentication tools and methodologies in use in your organization.
3. *Insecure Interfaces and APIs* – Covers provisioning, management, orchestration, and monitoring using APIs. Note that APIs can be compromised

Enterprise Data Security for US Europe and Asia

due to lack of access control, lack of encryption, and inadequate monitoring against accidental and malicious threats. As additional risk is because APIs are also extended to third parties which can result in additional risk.

4. *System Vulnerabilities* - are exploitable bugs in system program that an attacker can use. Additional risk in cloud is due to multi-tenancy because it exposes new attack surfaces. Mitigate this risk with regular vulnerability planning, patch updates, and secure design and architecture.
5. *Account Hijacking* - where the attacker has hijacked an account and can rack up major charges and watch the victim organization's traffic. This poses a significant risk of affecting some portion of CIA (Confidentiality, Availability and Integrity). A really bad practice that should be avoided is the practice of sharing account credentials between users and services. Hijacking is a term used for interception siphoning of data-in-transit. Remote access for cloud environments makes service traffic hijacking more susceptible. XSS is an example of hijacking data-in-transit. Cloud environments add a new threat: an attacker gaining credentials can eavesdrop, redirect clients to illegitimate sites, and use it as base for further attacks.
6. *Malicious Insiders* – include a current or former employee, contractor or business partner who has or had authorized access to an organization's network, systems and data. This is often overlooked. It is important to mitigate this by using least privilege access controls, full disk encryption, securing keys, and logging and monitoring.
7. *Advanced Persistent Threats* - *APTs* - (new) – A parasitical form of cyberattack that infiltrates systems to establish a foothold in the computing infrastructure to smuggle data and IP, often adapting to new protections. These attacks can be deliver via USBs, and through partner networks. APTs can move laterally through the data center and blend with normal work. Mitigate APTs by educating users to recognize spear phishing (email spoofing attack). The best defense is instituting regular mandatory awareness programs.
8. *Data Loss* – data can be lost in the cloud due to reasons other than malicious attackers; possibly fire or earthquake, loss of encryption key, and so on. EU requires handling such risks as a data breach. These risks can also include malicious or inadvertent, possibly casual or explicit, web-based or standalone, managed or independent. Cloud consumers should review contracted back-up and data loss provisions. You should also ask about your CSPs geographic redundancy and understand who is responsible for the various types of risks and threats.
9. *Insufficient Due Diligence* – lack of review of the CSPs contract for obligations and liability and pushing applications dependent on internal network security controls without appropriate cloud controls. Lack of knowledge of designers and application developers unfamiliar with cloud increases threat exposure. Impacts of insufficient due diligence include commercial, technical (unknown

architectural/operational issues), legal (data at rest, in motion or in use) in regulatory challenge foreign locations, and compliance.

10. *Abuse and Nefarious Use of Cloud Services* – are caused by poorly secured cloud services, free unsecure cloud service trials to your workforce, and fraudulent account sign-ups. Mitigations: CSP detection of payment instrument fraud and misuse of cloud offerings.
11. *DoS – Denial of Service* and *DDoS (Distributed Denial-of-Service)*– prevent users of cloud services from accessing data or apps by forcing the victim cloud service to use inordinate amounts of disk, CPU, network, or memory. DDoS attacks cause intolerable system slowdown.
12. *Shared Technology Vulnerabilities* – caused due to isolation issues in multi-tenancy in infrastructure (CPU caches, GPUs), compromising hypervisor or VM escapes. Use Defense-in-depth – compute, network, storage and end-user security to mitigate such vulnerabilities.

To counter these threats, your organization must implement a robust cloud security strategy. This includes:

- **Regular Security Audits:** Conduct comprehensive security audits to identify and rectify vulnerabilities.
- **Strong Access Controls:** Implement stringent access controls and authentication mechanisms.
- **Continuous Monitoring:** Employ advanced monitoring solutions to detect and respond to security incidents promptly.
- **Data Encryption:** Ensure data is encrypted both in transit and at rest, with strong and secure key management practices.
- **Workforce Training:** Provide regular security awareness training to all workforce members to mitigate insider threats from malicious insiders.

Securing Data in Cloud Applications & Infrastructure

With cloud-based applications and infrastructure, there is always a data security risk. Most SaaS-based solution are deployed directly on the public Internet. As a result, they need to be built with security in mind from the foundation up. This is especially true for directory services, but applies to virtually all web-based solutions. SaaS providers know that they are exposed to greater risk, and as a result ensure that their solutions are protected at all architectural layers as described below:

- **Data layer** – most cloud solutions encrypt data at rest to ensure that should a database or storage system be compromised so that it is difficult for hackers to expose the data. Even more important than encrypting data is one-way hashing and salting of passwords. With this process, the hashing and salting algorithm

should be strong enough to make it virtually impossible to reverse the password.

- **Application layer** – the application should be coded securely and checked for vulnerabilities. The most common cloud application vulnerabilities tend to be those at the user interface (presentation) layer with cross site scripting errors. These errors can expose the database or create a hole that allows a hacker to gain control over the application or server.
- **Server layer** – for cloud solutions, the servers are often hosted by a provider such as AWS or GCP, so physical controls are not really a concern for your organization. As a result, patching and updating to latest software versions is necessary. Any extraneous ports or services should be disabled where possible.
- **Network layer** – the network layer is protected through a secure connection. There are a number of different technologies to help here, including TLS. No communication between components whether internal or external should be allowed over an unsecure connection when cloud infrastructure is being leveraged.
- **Presentation layer** – access to the application itself should be protected via authentication. If possible, the applications should enforce long passwords and also leverage MFA (multi-factor authentication.)

Exploits in the Cloud

Given that the cloud environments are not totally under your organization's control but your organization is still responsible for any data exposure or exfiltration, it is important to understand the type of attacks and how to mitigate them. The following sections describe the types of attacks that your organization may face, especially, in your cloud applications.

Hypervisor Attacks

A hypervisor attack is an exploit in which an intruder takes advantage of vulnerabilities in the program used to allow multiple operating systems to share a single hardware processor. Most often, the attacker uses hypervisor services such as *create/delete*, *clone* and *migrate* to execute and extend a threat. The potential for success increases with the software stack's size, number of APIs and if there are lower degrees of security assurance in the code. Generally, a larger software stack provides a greater attack surface than a smaller code stack, because a large amount of code will likely have more coding errors. A large number of APIs, including in

third-party applications give an intruder more attack paths and a larger attack surface; typical APIs include hypervisor calls, interrupts generated by the hardware, processor instructions with parameters that the hypervisor can process, etc.

What a hypervisor attack looks like

When a hypervisor is compromised, a hacker can attack each virtual machine (VM) on a virtual host. One possible outcome of a hypervisor attack: the resource usage of a virtual machine can increase, resulting in a denial of service across the host or even a collection of servers. This problem is exacerbated when multiple virtual servers are involved. But commercial monitoring tools can detect and prevent these types of attacks.

Hyperjacking

Hyperjacking is an attack in which a hacker takes malicious control over the hypervisor that creates the virtual environment within a VM host. The goal of the attack is to target the operating system underlying the VMs so that the attacker's program can run and the applications on the VMs above the operating system are completely unaware of its presence.

Hyperjacking involves installing a malicious, fake hypervisor that can manage the entire server system. Regular security measures are ineffective because the operating system will not be aware that the machine has been compromised. In hyperjacking, the hypervisor specifically operates in stealth mode, and it makes it more difficult to detect, and more likely gain access to computer servers where it can affect the operation of the entire organization. If the hacker gains access to the hypervisor, everything that is connected to that server can be manipulated. The hypervisor represents a single point of failure when it comes to the security and protection of sensitive information.

For a hyperjacking attack to succeed, an attacker would have to take control of the hypervisor by the following methods:

- Injecting a rogue hypervisor beneath the original hypervisor
- 6. Directly obtaining control of the original hypervisor
- 7. Running a rogue hypervisor on top of an existing hypervisor

Some basic design features in a virtual environment that can help mitigate the risks of hyperjacking include the following:

- Security management of the hypervisor must be kept separate from regular traffic. This is a more network related measure than hypervisor related.

Enterprise Data Security for US Europe and Asia

- Guest operating systems should never have access to the hypervisor. Management tools should not be installed or used from guest operating systems which could have been compromised.
- Regularly patching of the hypervisor is important to ensure that all security related updates are in place.

Rootkit Attacks

Rootkits are a potential means of hypervisor attack, although that method is less common. A rootkit is a program or, more often, a collection of software tools that gives a threat actor remote access to and control over a computer or other system. While there have been legitimate uses for this type of software, such as to provide remote end-user support, most rootkits open a backdoor on victim systems to introduce malicious software, such as viruses, ransomware, keylogger programs or other types of malware, or to use the system for further network security attacks. Rootkits often attempt to prevent detection of malicious software by endpoint antivirus software.

Rootkits can be installed in many ways, including phishing attacks or social engineering tactics to trick users into giving the rootkit permission to be installed on the victim's system, often giving remote cybercriminals administrator access to the system.

Once installed, a rootkit gives the remote actor access to and control over almost every aspect of the operating system (OS). Older antivirus programs often struggled to detect rootkits, but most anti-malware programs today have the ability to scan for and remove rootkits hiding within a system.

How rootkits work

Since rootkits can't spread by themselves, they depend on clandestine methods to infect computers. Typically, they spread by hiding in software that may appear to be legitimate and could actually provide legitimate functions.

When users give a rootkit installer program permission to be installed on their system, the rootkit surreptitiously installs itself and conceals itself until a hacker activates it. Rootkits can contain malicious tools, including banking credential stealers, password stealers, keyloggers, antivirus disablers and bots for DDoS attacks.

Rootkits are typically installed through the same common vectors as any malicious software, including by email phishing campaigns, executable malicious files, crafted malicious PDF files or Word documents, connecting to shared drives that

have been compromised or downloading software infected with the rootkit from risky websites.

Rootkit detection and removal

Rootkits are designed to be difficult to detect and remove; rootkit developers attempt to hide their malware from users and administrators, as well as from many types of security products. Once a rootkit compromises a system, the potential for malicious activity is very high. Your organization needs to use commercial software products designed to monitor and prevent hypervisor attacks.

Stopping Hypervisor Attacks Before They Start

A hypervisor attack can hand hackers the keys to your virtual kingdom. But, with the proper precautions and tools, you can minimize the risk. It is paramount for your organization to build in security right at the hypervisor level and monitor the system rigorously for hypervisor attacks. It is better to design security measures into a system (i.e., an operating system, hypervisor), rather than adding them after the fact.

Note that security implications are often ignored in migrating applications to cloud environments because the addition of virtualized servers, storage and networking in a data center and with the extended perimeter in the cloud environments has created new security dependencies that were not found in the physical environments of the past.

Intrusions detection and prevention is essential for data security on-premises as well as in cloud environments. Hypervisor vendors provide good intrusion-detection capabilities, because they can place detection code in the places where intruders are most likely to attack. The Hypervisor vendors can provide information on where a hypervisor attack might occur and which types of security tools are available to prevent them. Some hypervisors are integrated into an OS kernel and others run on bare metal. Your organization needs to consider the best security options for both types of hypervisors.

Implementing Hypervisor security after a deployment

If you have already deployed one or more hypervisors, there are other ways to enhance the protection of your virtualized environments such as:

- Use security tools to monitor the virtual environment, including the virtual servers as well as the network traffic between VMs and hosts. These tools must include oversight and visibility into the virtual administration activities.

- Integrate hypervisor monitoring into your overall system management/monitoring infrastructure.
- Continuously validate your virtual environment to ensure the integrity and security of your virtual servers.

With an increasing workload of mission-critical applications that store sensitive data within virtualized infrastructures, the importance of hypervisor security will continue to grow. While these precautions will not guarantee the security of your virtual environment, they can minimize the risks of a hypervisor attack.

Guarding Against Virtualization Security Vulnerabilities

You can limit virtualization security risks and vulnerabilities by understanding attack vectors. This section of our virtualization security guide helps you prevent attacks.

An effective way to guard against virtualization security risks and unauthorized access is to think like an intruder. With this mindset -- as well as knowledge of attack vectors and vulnerabilities -- you can allocate your resources to weak spots in your environment.

The following section outlines various virtualization security risks, concerns and scenarios that can arise in any data center. Getting up to speed on the latest virtualization security vulnerabilities and exploits helps prevent attacks and ensures that you're prepared in the face of a security breach.

Risks of Hackers Stealing a VM and Its Data

Because a virtual machine (VM) is reduced to a single file, it offers administrators ample flexibility. The tradeoff, however, is that VMs are easy to steal.

- Consolidation security issues -- having several VMs share the same network interface card, for example, leaves your virtual environment susceptible to spamming attacks.
- Protecting storage networks from virtual machine security risks - when attached to virtualization hosts, storage networks can be attack vectors for hackers. To prevent this, you should keep your hypervisor and VM storage separated.
- A service console in a demilitarized zone, or DMZ depends solely on SSL (Secure Sockets Layer) protection which may leave the service console vulnerable.

Virtualization shouldn't change core security principles, but it may alter specific procedures. When you turn a hardware switch with 50 cables into a virtual switch that connects multiple virtualization hosts, for example, intrusion protection policies should be modified.

Virtual Security: Developing a Plan and Procedures

Your organization should create a virtual security plan that protects your investment in your virtual infrastructure. While sound practices can prevent sensitive information from landing in the wrong hands, you also want to guard against simple human error, which can subject your virtual infrastructure to data security breaches and other vulnerabilities.

How to build security into a virtualized server environment

Sometimes, the best way to secure a virtualized security environment is to think like an attacker. Using this approach, you should secure your high-value assets and work your way down to lower-priority assets.

Virtualization server security best practices

On internal networks, virtualization has redrawn the topology of assets by moving dispersed servers and applications that were once separated by hardware and network filters onto a single server. While this process alleviates traditional security threats, it also poses new challenges.

Considerations for Securing Virtual Environments

Getting an entire IT department thinking about protection against attacks is crucial for executing a virtual security plan. Confusion surrounding virtual environment semantics, guest OS management and virtualization administration network access can create virtual environment vulnerabilities.

Ensure that your IT staff has the required training and knowledge to manage the security of physical and virtual machines in a distributed network that includes cloud environments and have a detailed understanding of the data security threats faced by your organization.

Best Practices for Improving VM Security

Even if you already have a virtual security plan in place, it's never too late to strengthen it. How do you assess risks? Do you have the proper permissions set up?

How do you manage moving targets? You should routinely ask these kinds of questions, and this tip explains how to incorporate these considerations into your strategy to improve VM security. The following lists some key guidelines for improving VM security.

- Server virtualization security should be a top priority for your organization if you maintain a virtual environment.
- Ensure that your information security teams are not excluded from the planning and architecture stages of virtualization projects because retrofitting security in an existing virtualization infrastructure complicates an already-complex problem.
- Clearly document the differences of how to effectively protect different on-premises and cloud virtualization environments and also that the roles and responsibilities of IT staff are clearly defined in terms of responsibility for what they control and can modify.
- While your organization moved towards a virtual infrastructure for hardware cost-effectiveness, virtualization can also improve security by isolating unstable or compromised applications, creating cheaper intrusion detection tools and offering powerful forensic analysis capabilities.
- Before deploying a virtualization environment, you need to understand how a virtual infrastructure affects data center security. Is isolation and abstraction, worth the risks of malware or a compromised physical host and how should you prevent that.
- Securing a virtual environment doesn't have to differ dramatically from fortifying a traditional, physical infrastructure, but server virtualization security requires a few adjustments that must be carefully evaluated and implemented.
- Consider balancing virtual machine workloads to improve security and performance.

Chapter 11 – Defense in Depth and Security Architecture

Defense in Depth (DiD) is an approach to cybersecurity in which a series of defensive mechanisms are layered in order to protect valuable data and information. If one mechanism fails, another steps up immediately to thwart an attack. This multi-layered approach with intentional redundancies increases the security of a system as a whole and addresses many different attack vectors.

Defense-in-depth layered security architecture has three core parts — administrative controls, physical controls and technical controls. These controls coupled with the concepts of Zero Trust cybersecurity framework based on the idea that no entity – inside or outside of a network – should be implicitly trusted are the bedrock of what should define your security architecture. They are designed to protect your network's administrative, physical and technical aspects.

Defense-in-Depth

Defense in Depth is commonly referred to as the "castle approach" because it mirrors the layered defenses of a medieval castle. Before you can penetrate a castle, you are faced with the moat, ramparts, draw-bridge, towers, battlements and so on. Once you breach the outer walls, the attackers may be faced with one or more inner security walls before they can reach the inner parts of the castle.

The modern digital world while revolutionized by how users access and use the data provided by the information systems is constantly open to attack. The range of attackers and sophistication of attack technique varies only with the importance of the target system. To thwart these potential attackers, your organization needs to ensure that you have the right levels of security in place to prevent systems and networks being compromised. This may require multiple security techniques deployed across multiple layers to provide depth of data security needed for your business. While it may be impossible to protect against every possible attack that may occur, defense-in-depth strategy may be the best option to successfully protect your organization against the most likely types of attacks. This is where a defense in depth combined with zero trust systematically applied to your security architecture comes into play.

Defense-in-depth is the practice of using layers of security measures including firewalls, secured gateways, authentication and intrusion detection systems, to protect internal networks not only from external attacks. This provides backup levels of security in case other security measures fail. The goal is to keep threat actors outside of your network perimeter. So, if one defensive measure is compromised, another set of defense mechanisms can detect and prevent a breach attempt. By including redundancies and using security defenses across solutions, your organization can aim to close gaps in security and thwart potential attacks.

A layered approach to security can be applied to all levels of IT systems ranging from individual user laptops to thousands on servers in the data centers and extending beyond to the cloud environments. Note that not only does defense -in-depth protects your organizations sensitive data assets from external threats, the practices applied also protect your assets from insider threats.

Elements of defense in depth

New security products continue emerging to protect networks and systems. Here are some of the more common security elements found in a Defense-in-Depth strategy:

Network Security Controls

The first line of defense when securing a network is the analysis of network traffic. Firewalls prevent access to and from unauthorized networks and will allow or block traffic based on a set of security rules. Intrusion protection systems often work in tandem with a firewall to identify potential security threats and respond to them quickly.

Antivirus Software

Antivirus software is critical to protecting against viruses and malware. However, many variants often rely heavily upon signature-based detection. While these solutions offer strong protection against malicious software, signature-based products can be exploited by intelligent cybercriminals. For this reason, it is wise to use an antivirus solution that includes heuristic features that scan for suspicious patterns and activity.

Analyzing Data Integrity

Integrity of data, while in storage or in transit, is critical for your business. Data integrity can be managed in several ways. These include:

Enterprise Data Security for US Europe and Asia

- Using checksums associate with stored fields. This is a mathematical representation of the size of the file and any changes in the checksum while in storage or in transit can flag a loss of data integrity.
- File provenance, the source of its original location, the frequency of its use, checks against a known list of viruses and other malicious code.
- If an incoming file is completely unique to the system, it may be marked as suspicious.
- Data integrity solutions can also check the source IP address to ensure it is from a known and trusted source.

Behavioral Analysis

File and network behaviors often provide insight while a breach is in progress or has occurred. If behavioral analysis shows unexpected activity related to the use of the data it could indicate a firewall or intrusion protection activation. Behavioral analysis picks up the change from the norm and can either send alerts or execute automatic controls that prevent a breach from continuing any further. For this to work effectively, organizations need to set a baseline for "normal" behavior.

Defense-in-Depth vs. Zero Trust

Zero Trust is based on the concept of “never trust, always verify.” This cybersecurity framework is based on the idea that no entity – inside or outside of a network – should be implicitly trusted. Everything and everyone need to prove who they are and that they have the right permissions to access the target resource. In short, they need to undergo continual authentication, verification and authorization to keep access to applications and data.

What is the difference between Zero Trust vs defense-in-depth?

The main difference between Zero Trust and defense-in-depth security strategies is that Zero Trust never implicitly trusts users inside or outside network perimeters. Whereas with a defense-in-depth security strategy, users within a network are usually implicitly trusted but controls are applied nonetheless to ensure they have the proper credentials to access data.

While Zero Trust aims to prevent attackers from the outset, the defense-in-depth strategy simply aims to delay the attack by increasing the number of barriers an attacker must overcome to get within the security perimeter.

Benefits of Defense-in-Depth

The main cited benefit of a defense-in-depth approach is through layered security. If one part of an enterprise fails, then other networks and platforms are kept safe. To build it effectively, you must consider your strengths and weaknesses, and add layers where a breach is most likely. This redundancy helps to ensure that a single point of failure won't be the sole factor in a successful breach. By layering perimeter defense with a defense-in-depth configuration, threat actors must clear multiple hedgerows which slows down threat advancement and adds complications to

When you design your Defense-in-Depth security architecture, you need to keep some key considerations front and center as part of your design effort. The architecture will be addressing a continually growing attack surface and will need the flexibility to manage that. You will also need to manage your users bringing their own devices and potentially applications that must be administered effectively. With the many layers, threat identification and management can become tedious and complicated as you address total visibility across your organization. Integrating a wide range of utilities and tools with their own ecosystems can require a lot of attention on an ongoing basis.

Best Practice for Protecting Workstations

With users, especially those working remotely, management of their devices and access rights and permissions is an ongoing challenge. The security architecture must address all aspects of user identification and access management. Key considerations for your security architecture include the following.

Remove Local Admin Rights

Microsoft Windows, macOS and Linux administrator accounts open the door for threat actors. These privileged accounts are used to install and update workstation software, configure system settings and manage user accounts. Bad actors often leverage them to disable antivirus software or disaster recovery tools, and to launch ransomware and other types of malware. You need to carefully evaluate what rights they should have and minimize rights to what they really need.

Enforce Least Privilege

Users often have a legitimate need to perform an action requiring administrative privileges. Best-of-breed endpoint privilege managers let you elevate user permissions to perform certain specified tasks, based on policy, without requiring

end-user action or help-desk intervention. By enforcing the principle of least privilege — giving end users and applications the bare minimum set of privileges needed to perform their jobs — you can mitigate workstation vulnerabilities and strengthen your security posture. Consider using a jump server to elevate user accounts to privileged accounts. Leading endpoint privilege managers integrate with privileged access management (PAM) solutions to safeguard administrative passwords and strengthen workstation security.

Institute Application Control Policies

Application control is fundamental for protecting against malware and other threats. Many endpoint privilege managers support allow listing and deny listing functionality to explicitly permit or block known applications. The difference between a good application control and a great one is the ability to create comprehensive rules for applications that can be allowed and how they can be used.

Secure and Rotate Workstation Admin Passwords

Administrative accounts are a favorite target for threat actors. Even if you remove local admin rights from users, compromised admin accounts can still be used to launch attacks. Good admin password hygiene is essential, and policy-based governance is the best way to ensure good hygiene. Password must be rotated (changed) on a quarterly basis. With a PAM solution you can safely store passwords in a secure digital vault to prevent theft and rotate them automatically.

Deceive Threat Actors and Detect Them Early in the Attack Chain

Best-in-class endpoint privilege managers support privilege deception functionality to help you lure and throw off would-be attackers. You can create fake “honeypot” privileged accounts with alluring user names like “admin2” and intentionally weak passwords like “admin” to reel in and block adversaries at the point of entry. The endpoint privilege manager will generate an alert to let you know when you’ve lured an attacker.

Proactively Monitor Privileged Workstation Activity

Many endpoint privilege managers provide monitoring and reporting capabilities to help you track privileged workstation activity, identify suspicious behavior, and streamline compliance audits and forensics investigations. Most include canned reports that let you easily identify which workstations and applications are protected, as well as the specific privilege elevation policies and business rules being enforced. Leading endpoint privilege managers generate detailed event

messages whenever a policy is invoked and include dashboards and tools for observing and exploring policy-driven management.

Security Architecture

A cyber security architecture is the foundation of an organization's defense against cyber threats and ensures that all components of its IT infrastructure are protected. It is a set of security principles, methods and models designed to align to your objectives and help keep your organization safe from cyber threats. Security architecture translates the business requirements to executable security requirements. All of your data centers and cloud environments must be secured by your security architecture.

Infrastructure security is the practice of protecting critical systems and assets against physical and cyber threats. From an IT standpoint, this typically includes hardware and software assets such as end-user devices, data center resources, networking systems, and cloud resources. A security architect needs to understand the network, firewalls, defenses, detection systems, and many other factors that contribute to the security of your system.

In the previous section we discussed the important of Defense-in-Depth. The diagram below shows the layers of a Defense-in-Depth Security Architecture.

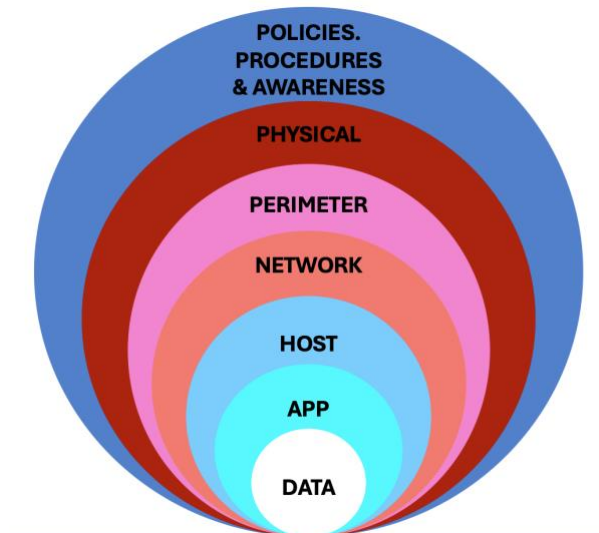


Figure 11-1: Defense-in-Depth Security Architecture

The main goal of infrastructure security is to reduce the level of risks the organization faces. The layers described in the diagram ensure that the security is managed at every layer of the architecture starting at the high-level with policies and procedures, followed by the physical infrastructure, the perimeter security at the perimeter boundary, the network security within the perimeter, the individual hosts (and in the case of the cloud, the virtual hosts), the applications and finally the data. Every one of these layers needs to have its own individual security.

Goals of Security Architecture

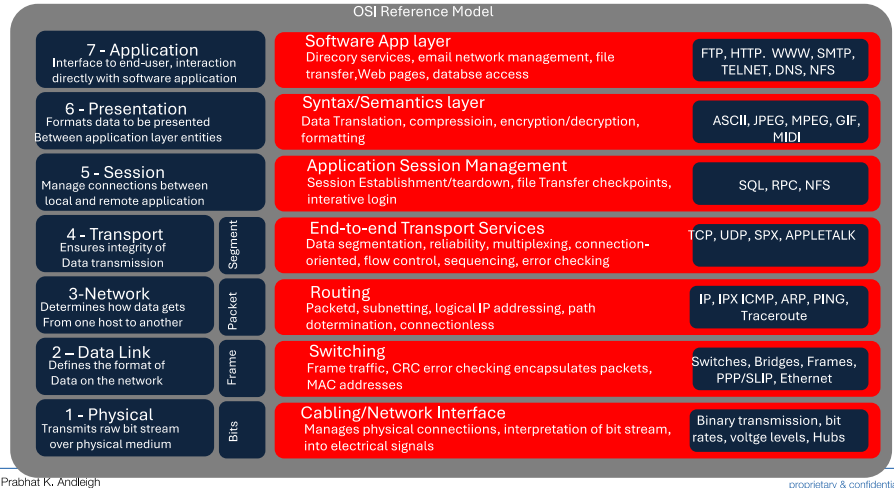
The following lists the key components and the role they play in the security architecture.

- A security service is a processing or communication service that improves the security of the data processing systems by protecting the flow of information. These services include authentication, access control, data confidentiality, data integrity, nonrepudiation, and availability.
- Authentication verifies that the user attempting to access a system is who he or she claims to be. This may be accomplished through a combination of passwords, usernames, tokens, biometrics, and so on.
- Access control restricts access to a system or information to authorized individuals. It may be based on user credentials, user location, user roles, or other criteria.
- Data confidentiality maintains the privacy of information while it is being transmitted. This is usually accomplished by encryption, using secure file-sharing software, and implementing confidentiality policies.
- Data integrity ensures that data is not altered or modified in transit or storage. Additionally, it ensures that the data received by the receiver is from a trusted source.
- Non-repudiation prevents users from denying that they sent or received a specific message.

Open Systems Interconnection (OSI) Architecture Framework

The components of your standard IT infrastructure can be broken down into the following three categories: hardware, software, and networking. While these represent the pillars of more traditional infrastructure, some of the same components are still used in cloud infrastructure.

Enterprise Data Security for US Europe and Asia



Source: OSI Organization

Figure 11-2: OSI Framework

The OSI framework described in the diagram above describes an important aspect of security architecture – how information flows across networks and how systems communicate. The seven layers of the OSI framework shows on the right the data formats and security protocols used at each layer. This forms the foundation of protecting data as it is communicated across networked systems. While describing these individual protocols is outside the scope for this book, you will do well to study the formats and protocols and how and understand where they are used.

Components of IT Infrastructure

Another important aspect to understand about security architecture are the components that are the building blocks of the security architecture. These are described in the following:

- **Hardware.** This includes servers, virtual hosts, other computers, network devices, storage systems, and peripheral devices.
- **Software** including all system software, vendor software, utilities, and applications.
- **Networks** and all network components such as routers, interface cards and gateways., and all network software.

Enterprise Data Security for US Europe and Asia

- Data Centers including on-premises, data centers managed by a hosting service provider, and cloud data centers managed by a cloud service provider.
- Cloud Services include all cloud native services software provided by the cloud service provider.
- Security Systems includes all security software hosted by security vendors, installed on your organization's systems, whether acquired from vendors or developed in-house.
- IT Service Management (ITSM) as an ongoing service to manage the IT infrastructure.

As you can see, this is a fairly comprehensive list that needs to be an integral part of the security architecture design and development.

Security Architecture Frameworks

Security Architecture Frameworks provide security architects a structure and a set of guidelines called frameworks to describe security architecture. A security architecture framework is a set of consistent guidelines and principles for implementing different levels of business' security architecture. The following are the three most common frameworks.

- TOGAF
- SABSA
- OSA

TOGAF Framework

TOGAF, or The Open Group Architecture Framework, helps determine which problems need to be solved within the security infrastructure in a business. Its primary focus is on the organization's goal and scope, as well as the preliminary phases of security architecture. TOGAF does not, however, give specific guidance on ways to address security issues.

SABSA Framework

SABSA, or the Sherwood Applied Business Security Architecture, is a policy-driven framework. It helps define the critical questions that security architecture can only answer: what, why, when, and who. The goal of SABSA is to ensure that after the design of security services, they are then delivered and supported as an

integral part of the enterprise's IT management. One downside, however, is that SABSA doesn't get into specifics regarding technical implementation.

OSA Framework

The Open Security Architecture (OSA) is a framework related to technical and functional security controls. OSA offers a comprehensive overview of crucial security components, principles, issues, and concepts that underlie architectural decisions involved in designing effective security architectures. However, OSA can only be used if the security architecture has already been designed.

Enterprise information Systems Architecture (EISA)

EISA is the framework for planning and implementing security measures for enterprise data. It provides the basis for understanding the information security goals of the organization and ensures that the right controls are in place to meet those goals. While EISA is a comprehensive approach to information security, it is more broadly concerned with optimizing business security. It also provides a systematic approach to managing and assessing risk, as well as it is a framework for designing, implementing, and maintaining information security solutions. Enterprise Information Security Architecture is an approach to security that is based on best practices and includes both technical and non-technical controls. It represents how information security is practiced within the organization and provides the basis for information security policies and procedures.

Cloud Security Architecture

Data stored in the cloud also needs protection from unauthorized access, malicious attacks, and other potential threats. Cloud security architecture is the combination of strategies, policies, and controls used to protect the cloud-based data that organizations store and process. This includes the physical, network, and host security controls for data in the cloud. Cloud security architecture is a critical component of any expanding business because of the increasing dependence on cloud computing for data storage and processing.

Cloud service models are classified into three major categories: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). Each of these service models has unique security requirements that can be addressed by different security architectures.

Enterprise Data Security for US Europe and Asia

When it comes to cloud security, AWS, GCP and Azure have a wide range of built-in security features and tools to help you secure your cloud data. All of them follow a shared responsibility model where the responsibility for security is divided between the cloud provider and the cloud customer. They also offer built-in compliance tools that can audit your cloud resources and recommend appropriate security best practices to help you secure your data and meet your compliance requirements. Compliance solutions on all three platforms support the majority of the major compliance standards including ISO 27001, PCI, DSS, and many others.

All three, AWS, GCP and Azure also provide Identity and Access Management (IAM) services, encryption in transit and at rest, and firewall rules, and VPN services. In addition to these built-in security features, each cloud platform also includes a marketplace where users can purchase third-party vendor applications to satisfy specific security needs.

Security of Physical Premises

We design and build our own data centers with multiple layers of physical security. Access to these data centers is tightly controlled by our own IT staff. We use multiple physical security layers to protect our data center floors. We use biometric identification, metal detection, cameras, vehicle barriers, and laser-based intrusion detection systems. For more information.

Data Centers managed by the cloud service providers have similar but much more intensive security procedures due to the scale and the range of customer data hosted on their systems.

A cloud service provider's data center (such as Google or Amazon) consists of tens of thousands to millions of servers connected to a local network. The CSPs design the server boards and the networking equipment and vet the component vendors that they work with and choose components checked for security. The CSPs work with vendors to audit and validate the security properties that are provided by the components. CSPs may also design custom security that they may deploy on servers, devices, and peripherals. These chips let the CSP identify and authenticate legitimate CSP devices at the hardware level and serve as hardware roots of trust.

Secure boot stack and machine identity

The CSP servers use various technologies to ensure that they boot the correct software stack. Cryptographic signatures for low-level components like the baseboard management controller (BMC), BIOS, bootloader, kernel, and base operating system image ensure boot level security. These signatures can be validated during each boot or update cycle. The first integrity check for the servers

uses a hardware root of trust. The components are completely CSP controlled, built, and hardened with integrity attestation. With each new generation of hardware, the CSPs improve security.

Secure Service Deployment

All cloud-native services provided by the CSP are the application binaries that their developers or business partners write and run on their infrastructure. To handle the required scale of the workload, thousands of machines might be running binaries of the same service.

The infrastructure does not assume any trust between the services that are running on the infrastructure. This trust model as we have seen earlier in this chapter is referred to as a *zero-trust security model*. A zero-trust security model means that no devices or users are trusted by default, whether they are inside or outside of the network.

Multi-Tenant Security

Because the infrastructure is designed to be multi-tenant, data from a large number of the CSP customers may share the same host and even the same hypervisor. Refer to the section on hypervisors attacks to see how such data is protected.

Inter-service Access Management

The owner of a service can use access-management features provided by the infrastructure to specify exactly which other services can communicate with the service. For example, a service can restrict incoming RPCs solely to an allowed list of other services. That service can be configured with the allowed list of the service identities, and the infrastructure automatically enforces this access restriction. Enforcement includes audit logging, justifications, and unilateral access restriction (for engineer requests, for example).

Benefits of Security Architecture

Clearly, strong security architecture leads to fewer security breaches. With modern technology, an organization is required to have a security architecture framework to protect vital information. This drastically reduces the threats associated with an attacker successfully breaching an organization's systems. Among the many benefits of security architecture is that it can translate each unique requirement into executable strategies and develop a risk-free environment for a business while aligning with the latest security standards and business needs.

Chapter 12 - Data Security Architecture

Data security is crucial for organizations to protect their business information assets including all confidential, highly confidential personal (PI) data and personally Identifiable (PII) data, collectively called sensitive information or sensitive data. All major financial and large business entities that collect their employees' and/or customers' sensitive data are likely to be required to comply with regulations imposed by local, state, or federal or central government agencies as well as by industry consortiums. Furthermore, documenting and ensuring appropriate data security policies and governance are crucial for maintaining customer trust, mitigating financial losses, and safeguarding competitive advantage.

A well thought out and designed Data Security Architecture plays an essential role in meeting the data security criteria laid out by the regulators and industry consortiums.

Cybersecurity vs. Data Security

Cybersecurity has achieved a high level of recognition due to major data breaches and exposures of data. The focus on cybersecurity does not often lead to an equally broad understanding of data security and data protection. Cybersecurity involves data security, but it also describes the way an organization protects its digital networks, programs, devices, systems, servers, and other online assets. In fact, cyber security is one component of data security, but it gets the most attention due to rapidly evolving and increasingly sophisticated cyber threats. Malware, hacking, and internal errors are the leading causes of cyber breaches, so it makes sense to incorporate cyber security planning to determine how to mitigate these risks effectively.

Data Security, on the other hand is, at a broad level, primarily concerned with protecting digital data, while cybersecurity is, as just mentioned, concerned with protecting computer systems, networks, and other digital assets.

However, one cannot consider cybersecurity independently from data protection. Together, they provide a powerful set of complementary measures and tools to protect your organization's sensitive data and the privacy of the data subjects whose data your organization is collecting and processing.

In addition to cybersecurity, Data security includes measures that are put in place to protect private data, especially sensitive data. So, data security focuses on the data

itself and how it needs to be and is protected. This type of security is concerned with the protection of data from accidental or purposeful unauthorized revisions. Additionally, data security is designed to use physical security, logical controls, cyber security, administrative controls, and other protocols in place for cybersecurity to protect your data.

Data Security Architecture, therefore, focuses on architecture, designs, processes, and operational controls that focus on data security.

Significant Aspects of Data Security Architecture

Developing data security architecture requires a clear understanding of the significant aspects of data security that drive the approach for the selections of technologies and their use in the architectural layers.

Protecting the Data Assets

At a very basic data level, data protection entails the following basic data protections:

- **Data encryption** — Encoding the data so it requires a key to unlock and read. The key itself needs to be protected.
- **Data masking or tokenization** — Masking specified areas of data or replacing it with tokens so that it is not visible or readable by unauthorized users and only authorized users can view it.
- **Data erasure** — Ensuring that obsolete data is completely removed so that sensitive information contained in it is not inadvertently exposed to unauthorized users.
- **Data backup** — Creating copies of data often using a data backup plan that meets the RTO (return to operation) and RPO (Recovery Point Objective) criteria for data recovery if the original is lost or locked through a cyberattack.

While all of these data protection mechanisms are important, data encryption and data masking have really significant implications for data security architecture.

Four Elements of Data Security

The four essential principles for effective security on any site, whether it's a small independent business with a single site, or a large multinational corporation with hundreds of locations and where data is stored on-premises or in public clouds are:

Enterprise Data Security for US Europe and Asia

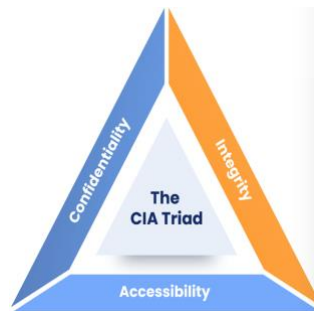
- Protection – of data from unauthorized exposure, corruption, or exfiltration.
- Detection – of any attempt for unauthorized access to data with underlying bad intention.
- Verification – of any attempts for unauthorized changes to data and if any data corruption occurred.
- Reaction – steps taken to mitigate the attack and restoring of data for normal operations.

These essential principles determine a wide swath of decisions for data security architectures on access controls, technologies, and layered patterns of the architecture.

The Three Areas of Data Security

The CIA triad refers to an information security model made up of the three main components: confidentiality, integrity, and availability. Each component represents a fundamental objective of information security. Let's look more closely at the CIA triad to understand the data security impact.

- **Confidentiality:** Confidentiality in this context means that the data is only available to authorized parties so that it has not been compromised or modified in an unauthorized manner, and that it has not been disclosed to unauthorized users. Organizing access to data on a need-to-know basis ensures confidentiality; it means that the data can be accessed only by users who have a need to access the data and are authorized to view it and use it. A breach of confidentiality may take place through a cyber-attack or social engineering.
- **Integrity:** Data integrity refers to the certainty that the data is not tampered with or degraded during or after submission. It is the certainty that the data has not been subject to unauthorized modification, either intentional or unintentional. There are two points during the transmission process during which the integrity could be compromised: during the upload or transmission of data or during the storage of the document in the database or collection.
- **Availability:** It implies that the data is available to authorized users when it is needed. For a system to demonstrate availability, it must have properly functioning computing systems, security controls and communication channels. Systems defined as critical (power generation, medical equipment, safety



systems) often have extreme requirements related to availability. These systems must be resilient against cyber threats, and have safeguards against power outages, hardware failures and other events that might impact the system availability.

What are the types of data security? Some of the most common types of data security, which organizations should look to combine to ensure they have the best possible strategy, include encryption, data erasure, data masking and tokenization, and data resiliency.

Data Security Reference Architecture

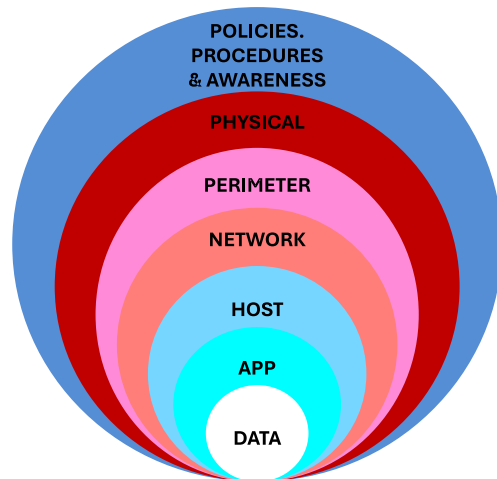
With an understanding of data security requirements, we are now ready to look at what we mean by data security architecture. Let's start with the important elements of data security.

Data architecture describes how data is managed--from collection through to transformation, distribution, and consumption. It sets the blueprint for data and the way it flows through data storage systems. It is foundational to data processing operations and artificial intelligence (AI) applications. Data Security Architecture determines how the data is protected through its lifecycle of the events described for the data architecture.

In general cybersecurity architecture combines security software and appliance solutions, providing the infrastructure for protecting an organization from a broad range of cyber-attacks. Your cybersecurity architecture should be able to adapt to the evolving cyber threat landscape as your organization engages in business. If your organization collects and uses sensitive data and stores private data, its use within as well as outside your security perimeter exposes it to internal as well as external. With cyber threats existing inside and outside the security perimeter, it has become essential for your security architecture to be based on a zero-trust framework. The diagram below shows the zero-trust layered security architecture.

Zero Trust Best Practices

- No implicit trust.
- Use dynamic and adaptive policies.
- Identify and authenticate all users, devices and sessions. Use MFA.
- For every request, verify the device, network flow, and access request.
- Use end-point detection and response. Use Firewalls.
- Use least privilege for all accesses
- Enforce RBAC, ABAC controls.
- Implement micro-segmentation for the networks.
- Maintain continuous monitoring and use behavior analytics.
- Classify all data and securely manage all data storage and access.



Prabhat K. Andleigh

proprietary & confidential

Figure 12-2: Zero trust Based Defense-in-Depth Data Security Architecture

Your organization's data architecture guides how the data is collected, integrated, enhanced, stored, and delivered to the businesspeople who use it to do their jobs. It helps make data available, accurate, and complete so it can be used for business decision-making. This determines the layers and the services at each layer of your data security architecture.

A security architect evaluates the security of your organization's systems for vulnerabilities. They perform penetration tests, risk analyses, and ethical hacks on LANs, WANs, and VPNs. To determine the efficacy and efficiency of routers, firewalls, and comparable systems, they also evaluate these systems and fix any misconfigurations that pose an exposure risk.

Security Architecture Framework

A security architecture framework is a set of consistent guidelines and principles for implementing different levels of an enterprise security architecture. Organizations often combine elements of each of these standard frameworks to build the design of their cybersecurity architecture.

The Security Architecture and Design documentation describes fundamental logical hardware, operating system, and software security components and how to use those components to design, architect, and evaluate secure computer systems.

Understanding these fundamental issues is critical for an information security professional.

Security-First Architecture

A security-first architecture that embraces the mindset that everything should be scrutinized and authenticated, must implement cutting-edge cybersecurity protocols, and operate on a philosophy of transparency where risks at each layer are visible and mitigated.

Security Architecture Components

Firewalls, antivirus and protection from malware software, threat intelligence platforms, and other security tools and applications that defend the organization's network are thus included as components of security architectures. A strong security architecture combines three elements: people, processes, and tools and describes how they interact and the services the processes and tools provide.

Common Data Security Architecture

The Common Data Security Architecture (CDSA) defines the infrastructure for a comprehensive set of security services to address the needs of individual users and the business enterprise. CDSA) consists of a set of layered security services and a cryptographic framework that provides an infrastructure for creating cross-platform, interoperable, security-enabled applications for client-server on-premises and cloud environments.

CDSA is an extensible architecture that provides mechanisms to manage add-in security service modules. These modules can provide cryptographic services and certificate services for use in building secure applications. The CDSA platform shows the five basic layers of the Common Data Security Architecture: Applications, System Security Services, the Common Security Services Manager, and Security Add-in Modules. The Common Security Services Manager (CSSM) is the core of CDSA. It provides a means for applications to directly access security services through the CSSM security API, or to indirectly access security services via layered security services and tools implemented over the CSSM API. CSSM manages the add-in security modules and re-directs application calls through the CSSM API to the selected add-in modules that will service the request.

This four-layer architecture defines five categories of basic add-in modules for security services to meet the security needs of all applications. CSSM also supports the dynamic inclusion of APIs for new categories of security services, required by

some applications. These elective services are dynamically, and transparently added to a running CSSM environment when required by an application. Called Elective Services, they are required by only a subset of security aware applications. When an elective service is needed a module manager for that category of service can be transparently attached to the system followed by the requested add-in service module. Once attached to the system, the elective module manager is a peer with all other CSSM module managers. Applications interact uniformly with add-in modules of all types. The following diagram describes the CDSA layers and services.

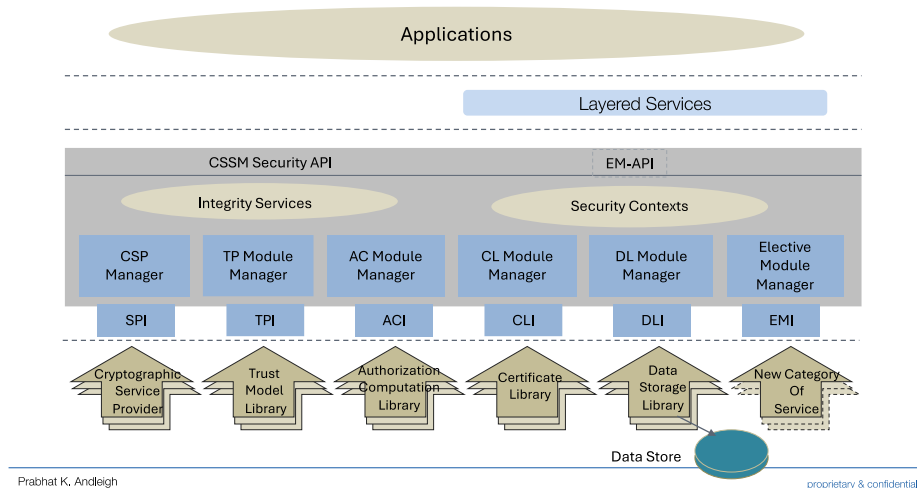


Figure 12-3: The Layered CDS Architecture and Service Classes

Basic Categories of Security Services

The CDSA layered diagram shows the following five basic categories of security services modules and an optional new category of services.:

- Cryptographic Service Providers (CSP) -- perform cryptographic operations including encryption, decryption, digital signaturing, key pair generation, random number generation, and key exchange.
- Trust Policy Modules (TPM) -- implement policies defined by authorities, institutions, and applications, such as your Corporate Information Technology Group (as a certificate authority). Each trust policy module embodies the

semantics of a trust environment based on digital credentials. A certificate is a form of digital credential. Applications may use a digital certificate as an identity credential and/or an authorization credential.

- Certificate Library Modules (CLM) -- provide format-specific, syntactic manipulation of memory-resident digital certificates and certificate revocation lists.
- Data Storage Library Modules (DLM) -- provide persistent storage for certificates, certificate revocation lists, and other security-related objects.
- Authorization Computation Modules (ACM) – evaluate user privileges and provide access to authorized users.

All of these modules are add-ins to the core system.

Data Security Architecture Based on CSDA

Applications dynamically select the modules used to provide security services. These add-in modules can be provided by independent software and hardware vendors. A single add-in module can provide services in multiple categories of services. A standalone registry system called the Module Directory Services (MDS) provides applications with information about the service modules available for use by applications.

The majority of the CSSM APIs support service operations. Service operations are functions that perform a security operation, such as encrypting data, adding a certificate to a certificate revocation list, or verifying that a certificate is trusted and/or authorized to perform some action. Service providers can require caller authentication before providing services. Application authentication is based on signed manifest credentials associated with the application.

Service modules can leverage other service modules in the implementation of their own services. Service modules acquire attach handles to other modules by:

- Receiving additional module handles from an invoking application.
- Selecting and attach additional service module directly.

To prevent stealth attacks, CSSM performs secure linkage checks on function invocation.

Modules can also provide services beyond those defined by the CSSM API. Module-specific operations are enabled in the API through pass-through functions whose behavior and use are defined by the add-in module developer. (For example, a CSP implementing digital signatures with a fragmented private key can make this service available as a passthrough.) The passthrough is viewed as a proving ground for potential additions to the CSSM APIs.

CSSM core services support:

- Module management.
- Security context management.
- System integrity services.

The module management functions are used by applications and by add-in modules to support runtime access to security service modules.

Security context management provides runtime caching of user-specific, cryptographic context information. Multi-step cryptographic operations, such as staged hashing, require multiple calls to a CSP.

CSSM, security services modules, and optionally applications, check the identity and integrity of components of CDSA. Components that can be checked include: add-in service modules, CSSM itself, and in the future applications that use CSSM.

In summary, the direct services provided by CSSM through its API calls include:

- Runtime management and access to all security service modules.
- Runtime management and access to elective module managers providing new security services.
- Caching of context information for cryptographic operations.
- Call-back functions used by add-in modules and CSSM to interact with an application process.
- Notification services to inform add-in modules of selected actions taken by an application.
- Management support for concurrent security operations.

Selecting CDSA Components

A single system can host multiple instances of CSSM. These instances can be distinct versions of CSSM or multiple copies of the same instance of CSSM. Applications can select which instance of CSSM to use at compile-time or at runtime, depending on how the CSSM is deployed. The dynamic components of a CDSA configuration require some level of compatibility to interoperate correctly. Three pieces of information form the basis for determining compatibility and interoperability among CSSM, service modules, EMMs, and applications:

- a globally unique identification (GUID), which distinguishes the component and its manufacturer.
- major and minor version numbers, which further distinguishes the supported APIs, feature set, and bug fixes of the component.

The Module Directory Service (MDS), is a standalone service outside of CDSA, which implements a database describing CDSA components available from the local platform. Applications and CDSA components can query MDS to obtain the compatibility information and numerous other attributes describing features of the CDSA components. This information can be used as the basis for selecting appropriate and compatible components at runtime.

Every CDSA component must have a unique identification GUID. Not all CDSA applications are required to have an identifying GUID, but it is highly recommended. MDS uses the GUID as the primary database key for locating information about the CDSA component. Specification of the version numbers is optional but believed to be of value as an augmentation to the distinguished name for an executable CDSA component.

When components are selected at runtime, applications use the MDS query functions to select components based on GUID or based on other properties of the component (such as the algorithms or features provided by the component).

If the application is selecting a CSSM at runtime, the application is responsible for loading the selected CSSM using services provided by the platform-specific environment. The application is also responsible for any load-specific initialization, such as symbol resolution. Once a CSSM has been loaded with the application,

Applications also use MDS to identify a service module providing the features and services required by the application.

Building a Data Security Architecture

As we have seen, the CSDA and other common security reference architectures used in financial institutions and large organizations subject to regulatory inspections and standardization includes governance, discovery and classification, vulnerability management, encryption, monitoring, data loss prevention, auditing, and more.

Data Security Architecture is the description and visualization of how data security controls and the related countermeasures for cyberattacks are positioned and how they relate to the overall systems and security architecture.

Use of Layered Technologies for Data Security Architecture

Your organization needs a single access point from which to monitor and control files wherever they're stored. This means full audit capability and data protection for content at end points, in transit, and in storage.

Securing User Accounts

The first line of defense in file security is to specify user access. This starts with built-in controls like password strength, rotation, and two-factor authentication. Typically, this authentication is via some compliment of AD/LDAP/SSO (Active Directory, LDAP, and Single Sign-on). It is important to implement controls such as customize account expiry, user account management, and policy definition and enforcement. Any access to such controls must require privilege access management (PAM) approval.

Securing User Devices

User devices, typically with a myriad of third-party applications are a persistent source of cyber-attack risks. For BYOD (bring your own device), it is important to set access restrictions based on device type or operating system along with access controls such as passcodes and device locks that prevent unauthorized users from using the device. The ability to remotely wipe lost or stolen devices help ensure that business-critical content never falls into the wrong hands.

Securing Data Centers

Data centers are risk at multiple levels and need to be protected at each of these levels. Here is a sampling of common approaches to data center security.

- Physical security against bad actors infiltrating into the facility.
- Brute force attacks on the network perimeter to achieve Denial-of-Service that overwhelms the network capacity of the servers.
- Implement zero trust security for access control and network segmentation.
- Use automation to prevent misconfigurations that allow bad actors access to key user authentication areas that give them control of the network.
- Prevent denial-of-service attacks by automatically terminating suspected sessions and preventing restart of sessions without fresh checks.

Data centers should be designed to provide complete confidentiality, integrity and availability for data at rest. They must be SSAE-16 Type II compliant. with 24-hour surveillance and biometric access controls.

Logging and Monitoring of Events

Data forensics is a specialized capability that allows tracking when users accessed the systems and the data files they accessed. It is useful to find the potential bad actors, the method of the attack, and it can also drive the mitigation approach to prevent such attacks in the future.

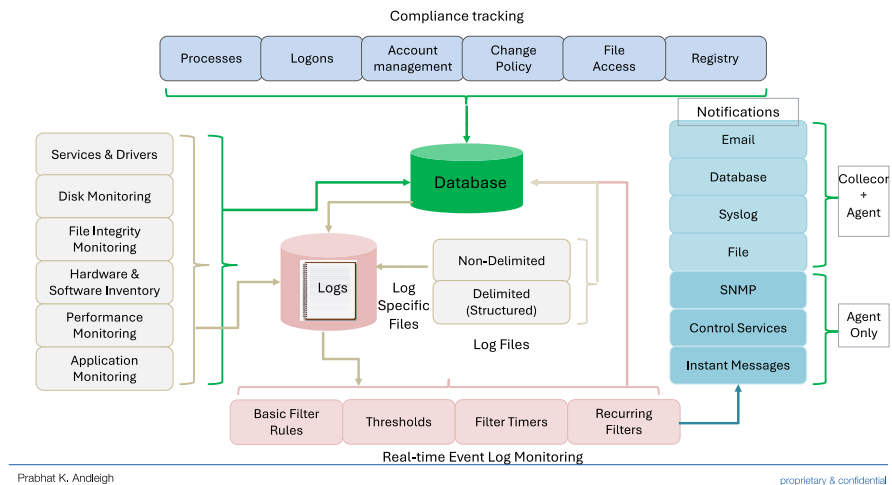


Figure 12-4: Logging and Monitoring Architecture

Logging of critical operations and monitoring operations that can be the target of an attack is an important tool for protecting data and being able to trace the history of data access transactions.

The States of Data - Securing Content

Your organization, depending on the nature of the markets it serves, needs to follow encryption standards set by NIST and other standards groups. 256-bit AES file encryption is a very common requirement and provides customers with unique encryption keys. Effective enterprise key management should be implemented for complete control of storage and use of encryption keys.

Data in the cloud security lifecycle can take on one of four forms: data in motion, data in use, data at rest, and data in stasis.

The confidentiality of data is typically provided through encryption technologies. Encryption is a broad topic. You can choose from many algorithms and strengths. Different encryption types exist for different storage options:

- Disk encryption of hard disks.
- Image encryption of snapshots.
- File encryption of files in file systems.
- Volume encryption of data in databases.
- Application encryption of data that is written to object storage.

Meanwhile, secure key management and strong identity and access management play

Securing Files in Transit

Data in motion is data as it traverses the internal and external networks in packets. To protect it, you need provide encrypted data communication pathways by using transport layer security (TLS). The most common secure communications options for cloud providers and clients are either through a direct connection from the client to the cloud or through a secured VPN over the internet.

Files being transferred from one location to another such as among user device, your organization's data centers on-premises and in cloud environments, and customer data centers are exposed to what is known as man-in-the-middle attacks. In this kind of attack, the file is intercepted in transit and modified or exfiltrated or corrupted to make it unusable. Specialized file transfer protocols and use of encryptions must be used by your organization for data protection.

It's important to consider how to extend encrypted communications into and within the cloud. For example, when you use Kubernetes or Red Hat® OpenShift®, you can use a service mesh such as Istio, which provides secure network connectivity within the complexity of a Kubernetes space. Other providers, such as VMware® services, offer secure extensions into the cloud that allow application mobility and infrastructure hybridity. For fuller discussion about protecting data in motion.

Securing Data at Rest

Data at rest is data that is in a data store on an operational system but that isn't being accessed and used at that time. The type of encryption that is required to provide confidentiality depends on what type of data it is and what kind of data

Enterprise Data Security for US Europe and Asia

store is in use. The solution architecture which determines the requirements for the type of data for data security provides more details about encryption and key management. Your sensitive data at rest should always be fully protected.

Encrypting data at rest protects files and databases at rest so that only authorized users with encryption keys and decrypt the data and read it. Protecting the encryption keys is important from hackers getting access to the encryption keys.

Full disk and file encryptions are the most common ways of protecting data at rest. Field level encryption has also been used for databases in very special cases but that method adds significant processing overhead.

Securing Data in Use

Data in use refers to data that is being used by your workforce, vendors or customers through applications that may process the data while accessing it. managing their access permissions so that they can access only the data they are authorized to use.

The data in this state is the most vulnerable to attacks as well as inadvertent misuse. Protection of data in use is achieved through robust controls for authentication of users and critical roles through the data security lifecycle as well.

The protection of data in use is a rapidly evolving field focused on protecting data in use by an application accessing and processing the data. This is achieved by performing computations in a hardware-based trusted execution environment (TEE). TEEs contain their own keys to ensure that no one else can access the data that is in the isolated environment of the TEE. Most of the major cloud platforms offer confidential computing on their clouds. It's important to look for confidential computing components that secure all places where data might be used, including secure access service edge (SASE). For more information about confidential computing,

Data in stasis

Data in stasis is data that is on either a physical disk or a virtual image, neither of which are up and running and the data is, in fact maybe no longer required. However, this data may contain confidential or highly-confidential sensitive data. The data must be protected even if it is destined to be removed soon. If Protecting data in stasis requires f due to regulatory reasons, it needs to remain stored unused for long periods of time, it should still be fully protected. Encryption is the most common approach to protect data: full disk encryption for protecting all data. In addition, file encryption should be used for data files. Encrypting the disk or image

prevents unauthorized people from spinning it up and accessing data on the disk. However, after the disk or image is running, the data is available for read or write and the service and applications are available for use. Therefore, data in stasis requires extra encryption architectures for complete data security.

Data Security Architecture Best Practices

As hackers get smarter, your organization as a business also needs to get smarter. You need to document the data security architecture best practices and ensure that they are adhered to in the initial architecture and design as well as in any follow-on changes to the architecture and design, especially, due to introduction of new tools and services.

Your organization must maintain compliance with the standards and policies that meet your regulatory requirements or industry norms and standards, and that ensure your reputation with your customers.

Architecting and Designing Data Security Architecture

A number of steps are required for developing the architecture and design for data security. Typically this would be incremental to the Security Architecture approach for the architecture and design for your application, irrespective of whether is intended for deployment to on-premises environments, cloud environments or both. The following lists the key steps for adding data security to your security architecture for your application.

1. Determining what data needs to be protected and its location.
2. Determining the users who need to access it and their roles.
3. Crafting the role design and service account requirements.
4. Decide on encryption technologies and key management system (KMS).
5. Designing data at rest strategy for on-premises and cloud environments.
6. Identify data to be moved and the locations participating in the moves.
7. Documenting the data security architecture diagram.

Determine What Data Needs to be Encrypted

Not all data needs to be encrypted. In this step you need to decide what data requires protection and how strong that protection must be. This depends as we have seen earlier if the data is public which vs. sensitive data) data (confidential,

highly-confidential, and personally identifiable data). You need to follow the data classification standard (within your organization or from standards setting bodies such as NIST).

Determine Users Who Need to Access the Data

List all of the users who will be accessing the data based on the business unit they belong to and their specific job functions. You can create groups of users who perform similar functions so that they can be allocated permissions as part of a group. Users can belong to multiple groups as per their job role.

Entitlements, Role Design and Service Account Requirements

The entitlements process is the analysis that is used to support the concept of least privilege. Based on your business processes and your data flow, decide which actors, functions, and locations are possible, and then within them, which are allowed by business role. If a role needs to know the data or operate on the data, it can be allowed. These sets of allowed functions by role constitute your definition of least privilege. The controls for least privilege are typically implemented by using identity and access management (IAM) system.

For the user listed in the previous step, determine the type of access they need to the data. Are they the data owners, do they need editing and delete permissions, should they be limited to a viewer role only with limited permissions that would prevent them from writing new data or deleting existing data. They are then provided those permissions.

Service accounts are accounts that are created to provide elevated privileges to users performing actions on production data. This is limited to key senior professionals in the team.

Determine Encryption Technologies and Key Management Systems

The Federal Information Security Modernization Act (FISMA) defined the framework of guidelines and security standards for to protect United States government information and operations. FIPS 140-2 under FISMA is the Federal Information Processing Standard that governs security requirements for the use of cryptographic systems. These are described in detail in Chapter 9.

Your design for the data security architecture needs to incorporate modules at the appropriate layers to manage encryption and decryption of data in the various states where it is applied.

Design Data-at-Rest Strategies for On-premises and Cloud Environments

This is an important step in the protection of sensitive data rest. After identifying the data that needs to be protected in the earlier step, you need to determine where, when and how to apply the encryption technologies and the Key Management Systems.

Identify Storage Location of Data to be Moved and Protection in Transit

You need to know the type of data and the location of data to know which compliance requirements apply. In addition, your risk assessment indicates which data needs to be protected to reduce the business risk to your enterprise. The compliance requirements and the data risk determine what needs to be encrypted and at what strength.

Most organizations don't know where all of their data is. But this is a very important step to ensure that all data is protected as needed. Automated discovery and classification tools can make that task easier.

The data inventory needs to include locations, storage types, file systems, database and version, type of data, and the protected elements in the data. It's also useful to know the encryption that is in use for each data store, the key management system that holds the keys, and the hardware security module (HSM), if applicable.

Sensitive Data Management Best Practices. Many organizations don't know where sensitive data is in their environment, nor do they track who the stakeholders are viewing and modifying that data. Most current privacy and compliance regulations require you to understand what data is sensitive, why the data is sensitive, and who has a stake before you can decide on the appropriate level of protection.

Your compliance requirements must be able to provide an understanding of what data needs to be protected in transit and what constitutes an appropriate level of protection in transit. This determines the technologies and utilities used for data transit among on-premises and cloud environments.

Data Security Architecture to Facilitate Operations

After your data security policies are defined, your data is known and understood, and your roles are defined by using least privilege, you have completed the critical first step in developing the data security architecture that will be used for architecture governance.

The next phase is to determine which processes to use to implement those architectural constructs and usage policies and guidelines. The processes can be manual or automated. They must be documented alongside your policies. A critical element is to ensure that the processes are assigned to roles and that the people in the roles are fully trained on the processes. The processes constitute the set of controls that you selected to address your data security risk.

Many organizations find it helpful to use automation to implement these controls. Infrastructure-as-Code (IaC) for example using a Terraform service provide automation for creating on-premises and cloud environments. This approach provides critical business efficiency and reduces the risks that are associated with human error due to misconfigurations. This approach also provides consistency in the creation of the infrastructure. You can use the NIST Cybersecurity Framework for guidance.

Architecting for Data Security Audits

The final piece of governance is to understand whether your policies are being enforced. Audits are used to review each control and to see evidence of compliance. Audits can vary in complexity depending on the size and operations of your organization. Organizations that are subject to external audits typically complete internal audits in preparation for the external audits.

Findings from audits must be addressed as part of this governance step, so you must have a feedback loop to ensure that changes are implemented to close out those findings. Larger organizations with complex operations often rely on automation to implement the processes, alert on compliance failures, and provide evidence of compliance for audits. This data security architecture strives to present the required functional elements to support the data security lifecycle and data security governance.

Drawing the Security Architecture Diagram?

Finally draw the diagrams depicting your security modules and how they are used, and present the system by doing the following:

- Clearly identify what's in scope and out-of-scope for the architecture.
- Show the high level, software, and hardware components.
- Show all interconnections and data flows between components.
- Label all communication lines and interconnections.

Enterprise Data Security for US Europe and Asia

Where needed, provide the details to describe the reasons for technology selections, the configurations for the technologies, and the inputs and outputs for each software module.

Chapter 13 - Business Continuity & Disaster Recovery

Disasters come in many forms such as, earthquakes, storms and winter freezes that cause loss of power, wars and hostilities of various kinds that affect your operations preventing people from accessing data centers, cyber-attacks such as data lockdowns for ransomware and even a disgruntled employee causing your data center disruption.

But business must continue through these to prevent financial loss and loss of reputation. Furthermore, a disaster that causes loss of your corporate business data can be catastrophic for your business. As a member of the IT staff, it is incumbent upon you to prevent any loss of business or data due to a natural or human made disaster.

To prepare for any such disaster you must address two concepts simultaneously: how to keep the business running without loss and also how to recover from any kind of data loss.

BCDR or Business Continuity and Disaster Recovery is a set of practices that bring together people, technology and processes that you need to employ to keep your business running at an acceptable level and recover the business operations to normal levels at the earliest during and after the disaster strikes.

So, business continuity focuses on keeping the lights on and keep the business open and operating in an acceptable capacity, while disaster recovery focuses on getting operations fully back to the level before the disaster struck. Your organization needs to incorporate business continuity and disaster as integral parts of your IT strategy. It requires defining a BCDR plan that spells out the resources and people essential for business continuity and disaster recovery and the processes used by each person in managing the resources.

Business Continuity

To ensure business continuity, your organization needs a documented plan that focuses on how your organization maintains critical business operations during and after a disaster. Note that it is important to keep the business running at acceptable levels even while a disaster is in progress. Your BCDR plan must

include every aspect of your organization: your employees at all levels and roles, communication channels within your organization and outside your organizations that are critical to your business, office buildings and data centers that your staff works in, your IT infrastructure including on-premises data centers and hosted data centers and cloud operations, your key business partners who may be affected by the disaster and any operational slowdowns, and so on.

The BCDR comprises specific actions and pre-determined responsibilities for all personnel responsible for maintaining your applications, data security and your IT infrastructure, that must be taken when disaster strikes for ensuring that operations can continue at acceptable levels.

Disaster Recovery

While business continuity is concerned with keeping the lights on and running operations at acceptable levels during a disaster, disaster recovery (DR) is a part of a business continuity plan concerned with the actions to be taken after the disaster is over and the business can return to normal operations. The disaster recovery part of the business continuity is the process, policies and procedures related to assessing the damage caused by the disaster and preparing for recovery of technology infrastructure, systems and applications to the state they were in before the disaster struck and which are vital to an organization after a disaster or outage.

A disaster recovery plan is focused on not only the recovery of the infrastructure but also the full restoration of important IT applications and data after a disaster or catastrophe. DR focuses on minimizing downtime as well as the impact of a disaster by ensuring vital support systems are up and running as quickly as possible with minimal loss of data.

DR is dependent on the backups of the information describing the IT infrastructure and the applications and data storage that are critical to the recovery processes to ensure that the recovery is complete and has returned all systems, applications and data to the state they were in before the disaster. This required well-planned and comprehensive processes to ensure full recovery.

Recovery Point Objective) & Recovery Time Objective)

Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are two very important terms for business continuity and disaster recovery.

RPO, or Recovery Point Objective, refers to how much data loss your application can tolerate without disrupting the business or incurring financial loss or loss of reputation. Examples of RPOs include time between data backups for business

financial data/banking transactions, customer relationship management databases, and patient records. Business units may each be able to continue to function if they lose data from within a specified period (the RPO). Subsequent data recovery effort attempt to recreate the data back to normal operations.

RTO or Recovery Time Objective and is a measure of how quickly after an outage an application must be available again. For example, suspending business operations on a retail sales site for any duration of time can result in financial loss. Similarly, unavailability of on-line banking services creates a level of suspicion and discomfort in the minds of the bank’s customers. We have seen the wide-scale impacts of temporary outages of outages operations of popular websites.

The diagram below shows the relationship between RPO and RTO.

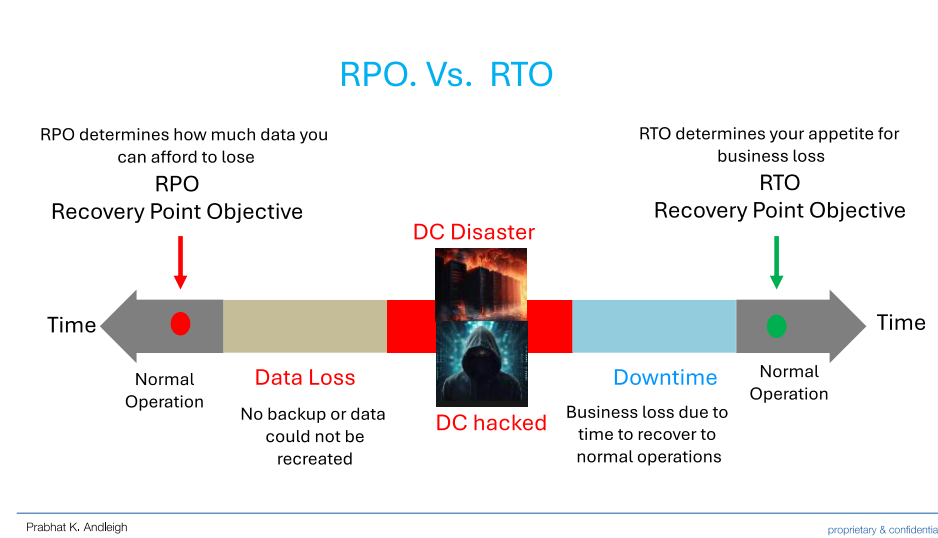


Figure 13-1: RPO and RTO Definition

To summarize, RTO is the goal your organization sets for the maximum length of time it should take to restore normal operations following an outage or data loss. RPO is your goal for the maximum amount of data the organization can tolerate losing. The diagram shows the RPO and RTO relevant to the timing of an event and how the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are established and managed.

No mathematical formulae exist to compute RTO or RPO values. They are strictly numeric time values set by your organization. For example, an RTO for a fairly

critical server might be one hour, whereas the RPO for less-than-critical data transaction files might be 24 hours, and might also support the use of backup storage devices.

A business creates a backup plan that uses traditional backups, and may conduct scheduled backups once or twice a day depending on the appetite for potential data loss. [time](#)

Since RTO and RPO are not directly related, RPO does not need to be less than RTO or vice-versa – you could have an RTO of 24 hours and an RPO of one hour, or an RTO of two hours and an RPO of 12 hours. For example, an e-commerce site may need to be online 4 hours after a disruption, so RTO is four hours.

For important applications that are not mission critical (tier-2 applications), the RTO is typically four hours and the RPO is two hours. For all other applications (tier-3 applications), a typical RTO is 8 to 24 hours and RPO is four hours.

BCDR Plan

As we saw earlier, BCDR stands for business continuity and disaster recovery. It refers to a collection of processes an organization can use to help it recover from a disaster and continue its normal operations. BCDR plans can help organizations prepare for unexpected natural disasters like floods.

A Disaster Recovery Plan (DRP) is a plan describing how to restore services after a disaster. It's usually used for businesses that operate during certain hours, like retail stores and restaurants. A BCP is a broad overview of what's needed to restore services after a disruption. It is a comprehensive plan that outlines the steps that an organization needs to take to ensure that critical business functions continue to operate during and after the event. A BCP typically includes a detailed analysis of potential threats, risk assessments, and strategies for risk mitigation.

A Business Continuity Plan (BCP) is a document that outlines how an organization will continue to operate during an unplanned event, such as a natural disaster, cyber-attack, or other disruptive event. It is a comprehensive plan that outlines the steps that an organization needs to take to ensure that critical business functions continue to operate during and after the event. A BCP typically includes a detailed analysis of potential threats, risk assessments, and strategies for risk mitigation. It also outlines communication plans, employee safety measures, and the steps that will be taken to restore normal business operations.

The goal of a Business Continuity Plan is to ensure that an organization can continue to operate, even during challenging and unexpected situations, while minimizing disruptions to critical business functions, and protecting the safety and well-being of employees and customers.

BCDR Planning

The fundamental goal of BCDR planning is not only to provide data recovery but also to minimize the effects of a crisis on business operations and enable an organization to get back to normal quickly in the aftermath of a disaster.

Listed below are five goals that you can use to fortify your BCDR plans.

1. *Assess the state of business:* Assessing the current state of an organization can help identify the threats and set priorities for remediation efforts. The plan should be updated routinely to account for changes to things such as personnel or systems.
2. *Find weaknesses and provide solutions:* The risks should be constantly evaluated to identify any gap that could potentially disrupt business operations and jeopardize BCDR strategies. It is important to acknowledge risks and address gaps uncovered during routine assessment of the BCDR plan.
3. *Review and test the plan:* Review the BCDR plan on at least a yearly basis to ensure it remains up to date and covers all aspects of the business for rapid recovery. There are several ways to test your plan, from tabletop simulations to full cut-over. Depending on your environment and the resources available, you may use one or several testing methods throughout the course of your evaluation.
4. *Identify location for data storage:* Identifying where critical business data and assets are being stored is one of the crucial objectives of BCDR planning. This will help disaster recovery personnel to start the recovery process even if the designated IT professionals are unavailable.
5. *Know the disaster recovery teams:* Knowing recovery personnel, their roles and how they can be reached during an emergency is another important goal of a BCDR plan. Communicate roles and responsibilities to all key stakeholders and keep this documentation accessible to employees and updated regularly.

Importance for Your Organization Have a BCDR Plan

A business continuity and disaster recovery plan helps organizations prepare for potentially disruptive events. It enhances an organization's ability to continue

business operations with little or no disruption and minimizes the risk in the event of a natural or man-made disaster.

Organizations without a BCDR plan may find it challenging to survive or recover from a major disaster. In fact, the effects of large-scale disasters can shut down operations. A large majority of companies without a DR plan that suffer a major disaster are out of business within 12 months. A BCDR plan is like an insurance policy for your organization. BCDR programs help your organization to reduce overall risk, get back up and running after an outage or disruption, mitigate the risk of data loss and protect against reputational damage.

Improving Business Resilience

Business resilience describes an organization's ability to respond and adapt quickly to disruptions or significant, unplanned changes that could threaten its operations, people, assets, brand, or reputation. By making your organization resilient, you are better equipped to manage unplanned disastrous events and emerge unscathed or even stronger by adopting new processes and tools to better manage such events in the future.

Your organization can build business resilience through innovation and investment, new tools, products or services; technology that helps your organization become more efficient and adaptable to rapid changes.

Best Practices for Improving Business Resilience

The well managed organizations follow many of the best practices listed below. They train their managers and staffs to follow these best practices on a regular basis.

Assess Your Business Operations

The day-to-day running of your organization will involve many different processes and systems. You may have already invested in technology such as a cutting phone system and IVR (Interactive Voice Response) for your organization. If you don't secure your data, however, this investment may not provide the anticipated returns.

You need to evaluate all processes used in your and determine which processes are crucial for your business to run? Do you have areas where failure could result in public liability, financial loss or loss of reputation in case of failure due to equipment damage or personnel injury. Each process relies on data, so in examining your processes, you will also need to examine your data and if it is at risk. By identifying the most important aspects of your organization, you can better prepare for the worse-case scenario.

It's important that you can ensure synergy between different departments. All teams should have a Certified Data Recovery Professional who has the competence to manage all systems used in your organization. You should also have a documented Corporate Data Recovery Plan. The plan should, at a minimum, address the following:

- If a certain operation failed, would day-to-day work be able to continue?
- Are you handling data that could put customers at risk if exposed?
- Are appropriate measures in place to deal with any disasters that occur?
- What are the cost implications of disruption?

Identify Your Risks

A risk assessment is a crucial part of your corporate data recovery plan for protecting data. It offers an opportunity to anticipate future disasters that could occur and their effects on your business. There is a wide range of potential risks that could affect both operations and data. They could be the result of human error, system failure, or criminal activity. A systematic review of your operations and processes can help identify things that could go wrong. Your plan should identify each risk and describe the tools and processes that must be used to mitigate the risk.

Assess the Risks and Develop Mitigations

Not all risks are equal, however. Some will have greater impacts than others. Some will impact output, others have a cost impact, and some will result in data corruption or loss.

The next step is, therefore, to assess the impact of each risk. You'll also need to consider the likelihood of each risk occurring. It's important that you score and document both the impact and likelihood of each risk. It's the combination of these scores that will enable you to focus attention where it's needed. A risk that has both a high impact and likelihood will need strong management attention.

Once your organization has identified and categorized the risks, your organization needs to determine what steps can you take to reduce the likelihood of a risk materializing. Determine what measures can be put in place to minimize the effects of a disaster. The actions needed to mitigate various aspects of a disaster or different disasters will vary greatly depending on the nature of the risk. This could even involve upgrading your security software to protect against data breaches or changing processes and assigning new resources..

Backup of Systems and Data

If you have carried out risk planning properly, and created a backup plan your organization will be able to recover provided you have backups of any data lost during an event. You need to consider backup strategies on-premises as well as in the cloud environments.

More and more businesses are moving to cloud processing. Your cloud service provider (CSP) can be expected to have good recovery practices and processes. If their processing system fails, due to a natural disaster such as a flood or loss of power including their backup power, they should be able to switch to another facility, sometimes without any impact in processing. You might not even be aware that there has been a problem. You may want to consider alternate data backup approach, such as another CSP just, in case their backup systems do not

The same thing applies to your data if a database is corrupted. If you don't use a third-party supplier, the costs of replicating systems and databases, and keeping them up to date, can be very high.

Ensure Maximum Protection

A third-party expert audit is an excellent approach to determine if you have the best security software for your needs, proper firewalls to protect from unwanted intruders, anti-virus and phishing software, encryption, and virus scanning. It can all get a bit overwhelming. The security and backup systems efficiency for organization's data is of crucial importance. The best disaster recovery systems in the world won't undo the impact on your customers and your reputation if your data has been lost or is stolen. So you need to ensure that you have the right protection from the outset. Once this is in place, ensure your systems have the latest

patches and upgrades. New threats emerge all the time. You should have a patching policy that keeps you up to date.

Another very important aspect of ensuring security is to make sure your staff is trained to employ best practices in password management and to recognize bogus phone calls, phishing emails, and other threats. In short, security awareness should be instilled into the culture of your business.

Ownership

Your corporate data recovery plan should define individual responsibility so that it is clear to the entire team who amongst them is responsible for every system and aspect of the recovery plan. Those are the people who should know the critical areas and understand what protection is in place. Importantly, they will also understand any weaknesses that have been identified and are responsible for signing off any mitigation plans to address these.

Forensics is an important part of any recovery plan to know what went wrong and at what time. A proper logging of all activities performed by the recovery team on a day-to-day basis and regularly scheduled reviews and reports of the activities updated by the owners of these critical processes will allow your team to identify any shortfalls in the processes and responsibilities and what needs to be done to fix it.

Policy and Procedures

It's important that the key aspects of your business are documented. Business Continuity policy and procedures are no exception. You might have expert staff who know exactly what to do when problems occur. But what if they are on sick leave the day a risk impacts? Or what if they leave on short notice and suddenly that experience is lost?

The answer is to ensure that you have disaster recovery and data protection policies and procedures in place. These should describe the detailed steps (both the 'who' and the 'what') to be taken in the event of problems.

Keep on Top of Business Continuity

Once you've put effort into establishing a business continuity and disaster recovery plan in place and augmented it with a corporate data recovery plan, it is important to audit it on a regular basis to ensure it remains capable of addressing new emerging threats and risks. Updating risks and risk mitigation actions regularly

Enterprise Data Security for US Europe and Asia

is crucial for ensuring that your organization is always ready to address disastrous events.

Most organizations run desk simulations as well as live exercises to test corporate disaster recovery response team readiness and capability. Such planning and operational exercises mimic steps to be taken in the event of a disaster. Simulated event for live exercise can involve senior management including even Board of Directors to role play key decisions in case of a real disaster.

Chapter 14 – Data Security Controls

A *security control* is any type of safeguard or countermeasure your organization uses to detect, counteract, or mitigate security risks to physical property and data centers, computer systems and servers within the data centers and your corporate offices as those used by your workforce, all information assets, and data across all work environments including cloud environments.

The security controls include the processes your organization has in place to protect from cyber-attacks that attempt to exploit network vulnerabilities and data hacks performed by hackers. The cybersecurity controls your organization implements should be designed to detect and manage the threats to your network and data. Security Controls can be technical, administrative, or physical. Firewalls, intrusion detection systems (IDS), encryption, and identification and authentication mechanisms are examples of technical controls. Policies and standards governance are examples of administrative controls. Examples of physical controls include data center perimeter fencing, locks, guards, access control cards, biometric access control systems, surveillance cameras, and intrusion detection sensors.

In this chapter we will focus on the technical and administrative controls. These controls are guided by the CIA Triad, *confidentiality*, *integrity* and *availability* that was reviewed in Chapter 2.

The security controls implemented by your organization should be comprehensive, consolidated and collaborative.

Security controls can be *detective* or *preventive*. Detective controls identify and report on any intrusion or hacking attempts. Preventive controls take it one step further and are designed to be implemented before a threat is enacted to reduce, mitigate, or prevent the impact of such an attempted intrusion or attack. These controls besides being technical such as firewalls, encryption, and IDS and IDP systems, can also be administrative to include policies, standards, processes, and governance.

Security controls can also be *corrective* in that your organization designs them to report on and correct procedural issues, and fix issues and inadequacies of technical controls either manually or in an automated manner.

Some common approaches very well understood by most organizations include the following:

- Establish strong passwords. Implementing strong passwords is the easiest thing you can do to strengthen your security.
- Put up a strong firewall.
- Install antivirus protection.
- Update your programs regularly.
- Secure your laptops.
- Secure your mobile phones.
- Backup regularly.
- Monitor diligently.
- Email, IM and Web Surfing.
- Workforce training on understanding and using controls.

As is obvious, security controls are an important mechanism used to ensure the security of all data and more importantly, sensitive data.

Security controls can be applied at the application level as discussed in the chapter 10 on application security or overall, to the entire network and operating environments. Our focus in this chapter is on standards-based security controls.

Standards That Define Security Controls

Chapter 2 described in detail the international standards and as a reader you will do well to study those before trying to understand the controls. Rather than describe the standards in detail, in this chapter we will review the security controls based on these standards. The two main standards used internationally and in the US include the ISO 27002:2022 and NIST 853 R4.

Most organizations develop their own internal policies and standards derived from these international standards.

ISO Standards that Define Controls

The ISO/IEC 27002 requires organizations to implement controls that meet its standards for an information security management system. A “control” in the ISO standard is defined as a measure that modifies or maintains risk. An information

Enterprise Data Security for US Europe and Asia

security policy, for example, can only maintain risk, whereas compliance with the information security policy can modify risk. Moreover, some controls describe the same generic measure in different risk contexts.

ISO 27002:2022 has 93 controls in the 2022 edition. These security controls are categorized into four control “themes.”

- People (8 controls)
- Organizational (37 controls)
- Technological (34 controls)
- Physical (14 controls)

Furthermore, ISO 27001:2022 controls have five types of attributes so that they’re easier to categorize. You can see the complete list of these 93 controls with their descriptions at the website:

<https://www.iso27001security.com/html/27002.html>

The 93 controls are each tagged with one or more values from each of 5 ‘attributes’ so they can be grouped, selected, or filtered in other ways too:

- Control type: preventive, detective and corrective.
- Information security properties: confidentiality, integrity and availability.
- Cybersecurity concepts: identify, protect, detect, respond and recover.
- Operational capabilities: continuity, application security, governance, asset management, information security protection assurance, human resource security, physical security, system and network security, secure configuration, identity and access management, threat and vulnerability management, continuity, supplier relationships security, legal and compliance, and information security event management.
- Security domains: governance and ecosystem, protection, defense and resilience.

The table below lists all 93 ISO controls.

#	ID	Name	Type
1	5.1	Policies for information security	Preventive
2	5.2	Information security roles and responsibilities	Preventive
3	5.3	Segregation of duties	Preventive
4	5.4	Management responsibilities	Preventive

Enterprise Data Security for US Europe and Asia

5	5.5	Contact with authorities	Preventive
6	5.6	Contact with special interest groups	Preventive
7	5.7	Threat intelligence	Preventive
8	5.8	Information security in project management	Preventive
9	5.9	Inventory of information and other associated assets	Preventive
10	5.10	Acceptable use of information and other associated assets	Preventive
11	5.11	Return of assets	Preventive
12	5.12	Classification of information	Preventive
13	5.13	Labelling of information	Preventive
14	5.14	Information transfer	Preventive
15	5.15	Access control	Preventive
16	5.16	Identity management	Preventive
17	5.17	Authentication information	Preventive
18	5.18	Access rights	Preventive
19	5.19	Information security in supplier relationships	Preventive
20	5.20	Addressing information security within supplier agreements	Preventive
21	5.21	Managing information security in the ICT supply chain	Preventive
22	5.22	Monitoring, review and change management of supplier services	Preventive
23	5.23	Information security for use of cloud services	Preventive
24	5.24	Information security incident management planning and preparation	Corrective
25	5.25	Assessment and decision on information security events	Detective
26	5.26	Response to information security incidents	Corrective
27	5.27	Learning from information security incidents	Preventive
28	5.28	Collection of evidence	Corrective
29	5.29	Information security during disruption	Preventive

Enterprise Data Security for US Europe and Asia

31	5.31	Legal, statutory, regulatory and contractual requirements	Preventive
32	5.32	Intellectual property rights	Preventive
33	5.33	Protection of records	Preventive
34	5.34	Privacy and protection of PII	Preventive
35	5.35	Independent review of information security	Preventive
36	5.36	Compliance with policies, rules and standards for information security	Preventive
37	5.37	Documented operating procedures	Preventive
38	6.1	Screening	Preventive
39	6.2	Terms and conditions of employment	Preventive
40	6.3	Information security awareness, education and training	Preventive
41	6.4	Disciplinary process	Preventive
42	6.5	Responsibilities after termination or change of employment	Preventive
43	6.6	Confidentiality or non-disclosure agreements	Preventive
44	6.7	Remote working	
45	6.8	Information security event reporting	Detective
46	7.1	Physical security perimeters	Preventive
47	7.2	Physical entry	Preventive
48	7.3	Securing offices, rooms and facilities	Preventive
49	7.4	Physical security monitoring	Preventive
50	7.5	Protecting against physical and environmental threats	Preventive
51	7.6	Working in secure areas	Preventive
52	7.7	Clear desk and clear screen	Preventive
53	7.8	Equipment siting and protection	Preventive
54	7.9	Security of assets off-premises	Preventive
55	7.10	Storage media	Preventive
56	7.11	Supporting utilities	Preventive
57	7.12	Cabling security	Preventive

Enterprise Data Security for US Europe and Asia

58	7.13	Equipment maintenance	Preventive
59	7.14	Secure disposal or re-use of equipment	Preventive
60	8.1	User endpoint devices	Preventive
61	8.2	Privileged access rights	Preventive
62	8.3	Information access restriction	Preventive
63	8.4	Access to source code	Preventive
64	8.5	Secure authentication	Preventive
65	8.6	Capacity management	Preventive
66	8.7	Protection against malware	Preventive
67	8.8	Management of technical vulnerabilities	Preventive
68	8.9	Configuration management	Preventive
69	8.10	Information deletion	Preventive
71	8.12	Data leakage prevention	Preventive
72	8.13	Information backup	Corrective
73	8.14	Redundancy of information processing facilities	Preventive
74	8.15	Logging	Detective
75	8.16	Monitoring activities	Detective
76	8.17	Clock synchronization	Detective
77	8.18	Use of privileged utility programs	Preventive
78	8.19	Installation of software on operational systems	Preventive
79	8.20	Networks security	Preventive
80	8.21	Security of network services	Preventive
81	8.22	Segregation of networks	Preventive
82	8.23	Web filtering	Preventive
83	8.24	Use of cryptography	Preventive
84	8.25	Secure development life cycle	Preventive
85	8.26	Application security requirements	Preventive
86	8.27	Secure system architecture and engineering principles	Preventive
87	8.28	Secure coding	Preventive
88	8.29	Security testing in development and acceptance	Preventive

Enterprise Data Security for US Europe and Asia

89	8.30	Outsourced development	Preventive
90	8.31	Separation of development, test and production environments	Preventive
91	8.32	Change management	Preventive
92	8.33	Test information	Preventive
93	8.34	Protection of information systems during audit testing	Preventive

Table 14-1: ISO 27002:2022 Controls

NIST Standards that define Controls

A wide range of security controls are required to ensure security of data.

NIST controls are meant to enhance an organization's cybersecurity program, risk posture, information protection, and security standards. NIST 800-53 is mandatory for federal agencies, but any other organization can use the standard to improve their own security program. NIST defines over 1,000 security controls.

Five Principles of NIST

The five principles of NIST are: Identify, Protect, Detect, respond and Recover.

Identify. Make a list of all equipment, software, and data you use. This should include all data center servers and end-user devices such as laptops, smartphones, tablets, and point-of-sale devices.

Protect. Your organization should control who can log on to your extended network, and use your organization issued devices as well as BYOD devices. NIST also requires your organization to use appropriate security software to protect your data at rest as well as in transit. All of your security related software must be updated and kept current to ensure that all security patches have been applied. A complete set of policies and standards must be defined, and they should also specify how data is destroyed after use. All workforce members must be made aware of and trained to identify and manage risks.

Detect. Monitor your infrastructure on a continual basis to detect all access to your systems and software and prevent unauthorized use. Your security staff must actively monitor any alerts and alarms caused by unauthorized access attempts.

Respond. Your organization should have formal plans for addressing all potential risk events or emergencies that may put your data assets at risk. This plan should prescribe how to keep the business up and running and how to recover from a

Enterprise Data Security for US Europe and Asia

disaster. Some events may require regulatory responses such as workforce, customers, law enforcement, and regulatory agency notifications. Every event or attack should be investigated and contained. The response plan along with policies and standards should be updated to reflect lessons to prevent future events.

Recover. All equipment that and software that could be impacted by similar potential future events should be updated. It is always very beneficial to inform your workforce and customers about the events, how they were mitigated and what actions have been taken to prevent such attacks and event in the future.

NIST Families

The following table describes the Security and Privacy Control families for NIST. NIST document source: Security and Privacy Controls for Information Systems and Organizations.

<https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf>

ID	FAMILY	ID	FAMILY
AC	Access Control	PE	Physical and Environmental Protection
AT	Awareness and Training	PL	Planning
AU	Audit and Accountability	PM	Program Management
CA	Assessment, Authorization, and Monitoring	PS	Personnel Security
CM	Configuration Management	PT	PII Processing and Transparency
CP	Contingency Planning	RA	Risk Assessment
IA	Identification and Authentication	SA	System and Services Acquisition
IR	Incident Response	SC	System and Communications Protection
MA	Maintenance	SI	System and Information Integrity
MP	Media Protection	SR	Supply Chain Risk Management

Table 14-2: NIST Families

While it may be useful to review the NIST standard to fully understand the full scope, from the data security perspective, the families of interest from the perspective of data security include AC, AT, AU, CA, CM, IA, IR, RA, and SI.

Family	Description
AC	<ul style="list-style-type: none"> • (AC-1) Policy and Procedures • (AC-2) Account Management • (AC-3) Access Enforcement • (AC-4) Information Flow Enforcement • (AC-5) Separation of Duties

Enterprise Data Security for US Europe and Asia

	<ul style="list-style-type: none">• (AC-6) Least Privilege• (AC-7) Unsuccessful Login Attempts• (AC-8) System Use Identification• (AC-9) Previous Login Notification,• AC-10) Concurrent Session Control• (AC-11) Device Lock• (AC-12) Session Termination• (AC-13) Supervision and Review – Access Control• (AC-14) Permitted Actions without Identification or Authentication• (AC-15) Automated Marking• (AC-16) Security and Privacy Attributes• (AC-17) Remote Access• (AC-18) Wireless Access• (AC-19) Access Control for Mobile Devices• (AC-20) Use of External Systems• (AC-21) Information Sharing• (AC-22) Publicly Accessible Content• (AC-23) Data Mining Protection• (AC-24) Access Control Decisions• (AC-25) Reference Monitor
AT	<ul style="list-style-type: none">• (AT-1) Policy and Procedures• (AT-2) Literacy Training and Awareness• (AT-3) Role-based Training• (AT-4) Training Records• (AT-5) Contacts with Security Groups and Associations• (AT-6) Training Feedback
AU	<ul style="list-style-type: none">• (AU-1) Policy and Procedures• (AU-2) Event logging• (AU-3) Content and Audit Records• (AU-4) Audit Log Storage Capacity• (AU-5) Response and Audit logging Features• (AU-6) Audit Record Review, Analysis and Reporting• (AU-7) Audit Record Reduction and Report Generation• (AU-8) Time Stamps• (AU-9) Protection of Audit Information• (AU-10) Non-repudiation

Enterprise Data Security for US Europe and Asia

	<ul style="list-style-type: none"> • (AU-11) Audit Record Retention • (AU-12) Audit Record Generation • (AU-13) Monitoring and Information Disclosure • (AU-14) Session Audit • (AU-15) Alternate Audit Logging Capability • (AU-16) Cross-organizational Audit logging
CA	<ul style="list-style-type: none"> • (CA-1) Policy and Procedures • (CA-2) Control Assessments • (CA-3) information Exchange • (CA-4) Security Certification • (CA-5) Plan of Action and Milestones • (CA-6) Authorization • (CA-7) Continuous Monitoring • (CA-8) Penetration Testing • (CA-9) internal System Connections
CM	<ul style="list-style-type: none"> • (CM-1) Policy and Procedures • CM-2) Baseline Configuration • (CM-3) Configuration Change Control • (CM-4) Impact Analysis • (CM-5) Access Restrictions for Change • (CM-6) Configuration Setting • (CM-7) Least Functionality • (CM-8) System Component Inventory • (CM-9) Configuration Management Plan • (CM-10) Software usage Restrictions • (CM-11) User-installed Software • (CM-12) information Location • (CM-13) Data Action Mapping • (CM-14) Signed Components
IA	<ul style="list-style-type: none"> • (IA-1) Policy and Procedures • (IA-2) Identification and Authentication (Organizational Users) • (IA-3) Device identification and Authentication • (IA-4) Identifier Management • (IA-5) Authentication Management • (IA-6) Authentication Feedback • (IA-7) Cryptographic Module Authentication

Enterprise Data Security for US Europe and Asia

	<ul style="list-style-type: none"> • (IA-8) Identification and Authentication (Non-organizational users) • (IA-9) Service Identification and Authentication • (IA-10) Adaptive Authentication • (IA-11) Re-authentication • (IA-12) Identity Procedure
IR	<ul style="list-style-type: none"> • (IR-1) Policy and Procedures • (IR-2) Incident Response Training • (IR-3) Incident Response Testing • (IR-4) incident Handling • (IR-5) incident Monitoring • (IR-6) Incident Reporting • (IR-7) Incident Response Assistance • (IR-8) incident Response Plan • (IR-9) Information Spillage Response • (IR-10) Integrated information Security Analysis Team
RA	<ul style="list-style-type: none"> • (RA-1) Policy and Procedures • (RA-2) Security Categorization • (RA-3) Risk Assessment • (RA-4) Risk Assessment Update • (RA-5) Vulnerability Monitoring Scanning • (RA-6) Technical Surveillance Countermeasures Survey • (RA-7) Risk Response • (RA-8) Risk Impact Assessments • (RA-9) Criticality Analysis • (RA-10) Threat Hunting
SI	<ul style="list-style-type: none"> • (SI-1) Policies and Procedures • (SI-2) Flaw remediation • (SI-3) Malicious Code Protection • (SI-4) System Monitoring • (SI-5) Security Alerts, Advisories, and Directives • (SI-6) Security Privacy Protection Verification • (SI-7) Software, Firmware and Information Integrity • (SI-8) Spam Protection • (SI-9) information Input Restrictions • (SI-10) Information Input Validation • (SI-11) Error Handling

Enterprise Data Security for US Europe and Asia

- (SI-12) Information Management and Retention
- (SI-13) Predictable Failure Prevention
- (SI-14) Non-persistence
- (SI-15) Information Output Filtering
- (SI-16) Memory Protection
- (SI-17) Fail-safe Procedures
- (SI-18) Personally Identifiable information Quality Operations
- (SI-19) De-Identification
- (SI-20) Tainting
- (SI-21) Information Refresh
- (SI-22) Information Diversity
- (SI-23) Information Fragmentation

Table 14-3: NIST Control applicable to Data Security

Given this vast number of clauses for the NIST standards, one way to address them is to create your own standards documents that cover the domains and list the clauses that need to be addressed by creating controls applicable to your business domain and your organization.

Chapter 15 - Data Security Governance

The first rule of data security is to keep data secure. businesses have leaned on enterprise Data Loss Prevention (DLP) solutions, considering them fundamental to data security. Sensitive data sprawls across cloud folders, websites, company laptops and personal smartphones. This complexity adds up to a lot of risk, including potential breaches and the looming specter of regulatory non-compliance. Organizations deploy products like Cloud Application Security Brokers (CASB) or Secure Web Gateways (SWG), each with their own separate DLP. Managing these disparate mechanisms for cybersecurity and more specifically data security requires careful orchestration of access controls and access management.

Data Governance is the overall management of the availability, usability, integrity, and security of data used in an organization. It's a collection of practices and processes to ensure the formal management of data assets.

While data governance focuses on managing data as an enterprise asset, data security prioritizes protection against threats. A holistic approach that combines robust governance and security measures is crucial for unlocking the full potential of data while ensuring its confidentiality and integrity. So, data security governance is the overall orchestration and management, data classifications, secure storage, transfers and use of the data, access controls applied to access management, and compliance with corporate data security standards.

Data Governance Framework

Your organization should have a fully detailed and documented Data Governance Framework that aligns with the following:

1. Policies and standards framework
2. Data Security Governance framework
3. Architecture and Design framework
4. Deployment Framework -- On-Prem and Cloud

These frameworks govern all aspects of data creation, data storage, data transfers and data use. Because different teams are responsible for each of these areas, the framework documents must be specific to each function. For example, the central architecture team under the office of the chief technical officer (CTO) documents the framework for architecture and design governance.

Policies and Standards Framework

The Policies and Standards development team spells out the content and process for governance risk and compliance (GRC). All of these frameworks contribute to data security.

Your organization must follow the standards and controls described in previous chapters in establishing and documenting the frameworks and assigning tasks based on the frameworks to the various teams including but not limited to the central architecture team under the office of the CTO, the GRC team, the internal audit team, the infrastructure and operations teams, and cybersecurity governance teams.

Data Security Governance Framework

Our focus in this chapter is Data Security Governance (DSG). Data security governance spans several different functions and domains within the organization and includes all of the frameworks described in the previous section. DSG must be applied to a few specific dimensions. These include among others:

1. Data Classification
2. Data Quality Management
3. Data Life Cycle Management
4. Data Access and Entitlements Governance

Let's review in detail each of these data security governance dimensions.

Data Classification Governance

Data classification, as we have seen in previous chapters, is important to identify the PI (personally identifiable) data, the highly confidential data, and the confidential data (collectively sensitive data) which your organization has an obligation to protect from unauthorized exposure and use.

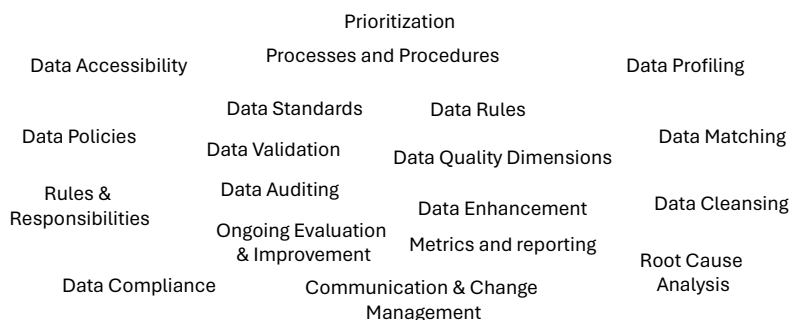
Your organization needs to have internal standards, typically derived from the worldwide regulatory standards (as applicable) described in Chapter 2. These can include ISO, NIST, GDPR, Privacy Act of India and others. These standards should be documented and socialized across your organization so that all departments understand their roles in ensuring that the standards are complied with at all times from the point the data is created until it is destroyed. We have also defined the roles of the various entities such as the data creator, owner, steward and so on.

Data classification governance permeates through all of the frameworks as data moves through the various stages of creation, storage, transfers and use by applications across various on-premises and cloud environments used by your organization. So, every framework must have a section dedicated to data classification governance with references to applicable policies and standards.

Data Quality Governance

Data quality is used to describe the degree to which data is accurate, complete, timely and consistent with business requirements rules. It requires that data has not been altered in an unauthorized manner during its lifetime and the data integrity has been maintained at all times. Several measures and processes need to be employed to ensure the accuracy, completeness, timeliness, and consistency of data.

The management of these measures is a function of data quality governance. Data quality governance is the exercise of authority, control, and shared decision-making over the management of data assets. The diagram below describes the relationship between data quality and the governance of data quality.



Data Governance

Data Quality

Prabhat K. Andleigh

proprietary & confidential

Figure 15-1: Data Governance and Data Quality Overlap

Data Quality Guarantees are an expectation of your customers and are essential for your organization to ensure that you have the trust of your customers. Data quality guarantees can be assured if your organization has a data quality standard and a data quality policy that defines how the standards are met. The policies and

associated governance documentation must define the processes that affect the data throughout its lifecycle and the roles of all members of your workforce who handle the data in some manner. A good measure of the effectiveness of your data quality governance is tracking metrics and key performance indicators and, most importantly, adjusting the governance processes, procedures, and roles if the metrics and KPIs do not meet the goals established already.

Data Life Cycle Management

Data is one of the most valuable assets for your organization, but it also comes with various risks throughout its lifecycle. From creation until it is deleted or destroyed from all storage areas, data can be exposed to threats such as loss, theft, corruption, misuse, or breach or exposure.

Data lifecycle management (DLM) is the process of safeguarding data appropriately throughout its existence from the time it was created until it is destroyed and no longer exists in any form of storage.

The primary goal of data lifecycle management is to ensure a seamless flow of information throughout its lifecycle and make it available to authorized users and applications, and allow them to process and share the data under controlled circumstances. The data lifecycle is governed by your organizations management teams in various departments as defined by the data lifecycle governance documentation.

The Lifecycle Management Data Phase consists of six phases that cover the various stages of the data. The diagram below describes these six phases along with the governance implications for these six phases known by the acronym – CSU-SAD (Collect, Store, Use, Share, Archive, Destroy).

Enterprise Data Security for US Europe and Asia

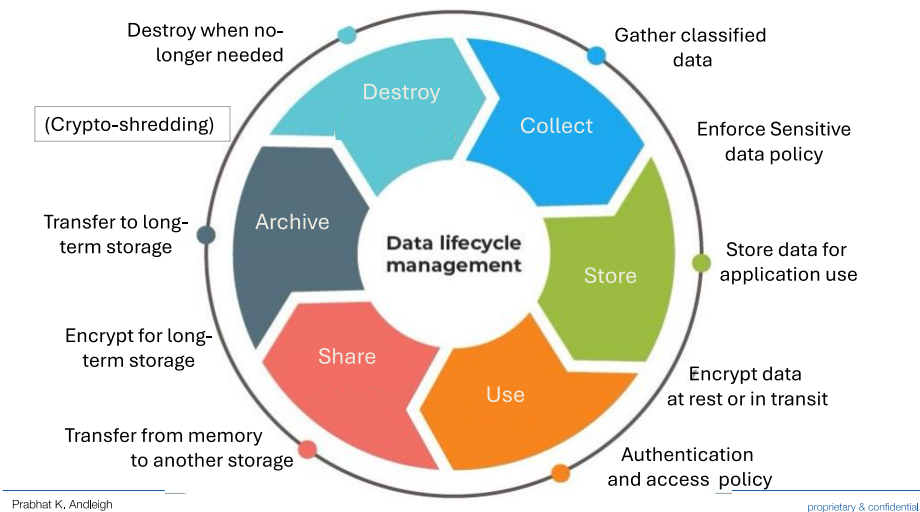


Figure 15-2: Data Lifecycle Management

- **Collect** -- Data collection process of gathering information in an established systematic fashion where the creator of the information defines the classification of the information. This can be performed by an individual or an application.
- **Store** -- Data storage refers to temporary or permanent storage of data organized as objects or data fields within files. The storage can be on disk drives or SSDs on premises or in cloud environments.
- **Use** -- Data access and usage by an individual or an application is the on-demand ability to retrieve, modify, copy, or move data from IT systems as an authorized user or application. It allows them to access different repositories perform their business function by accessing and manipulating the data.
- **Share** -- Data sharing or transfers refers to copying data from a storage device to memory or another storage device within the on-premises data center or to cloud environments or to third party organizations' storage systems.
- **Archive** -- Data can be stored on a long-term basis either for business or regulatory purposes. Archived data must be protected in the same manner as data stored for use with access controls and security protections.
- **Destroy** - Data destruction is the process of destroying data stored on hard disks and other forms of electronic media so that it is completely unreadable before it is deleted. Crypto-shredding is a common approach to data destruction.

Note that each of these stages of data lifecycle has associated with it governance requirements as described in the diagram above.

Data Access and Entitlements Governance

Data Access and Entitlements Governance helps you enforce data access rules and policies via data discovery, data classification, entitlements definition and access management. This includes understanding the privileges and permissions assigned to users and applications associated with the data access, with the goal of allowing access based on a least privileged model. This ensures that all data is used and managed in compliance with laws, regulatory requirements, industry standards your organizations ternal policies and standards.

Data Access Governance (DAG) is a market segment that focuses on identifying and addressing the malicious and non-malicious threats that can come from unauthorized access to structured sensitive data as well as valuable unstructured data. As an important function of DAG, your organization's policy management and visibility to wherever your organization stores and uses data should be governed through a single, powerful data access and entitlement-controlled security solution. You should create the policies and standards for access management and deploy them throughout your organization for seamless operations.

As we have seen, data and security controls refer to the tactics, policies, and procedures that your organization will use to meet your data governance and data management objectives. They are the rules and systems that your organization will rely on to ensure that only authorized users can access their data, ensuring its security and integrity.

Data level entitlements are rules that dictate whether or not a user can view or download or use otherwise certain sets of data that they have been specifically permitted access to. An entitlement can be a role, responsibility, or group membership. For example, a user granted the Sales Analyst role on a target system, then that user can use that entitlement to access and generate sales-related reports from the target systems.

Note that Authorization has a Prescriptive behavior whereas Entitlement has Descriptive behavior. Authorization controls access to the data resource while entitlement defines what access the user is allowed to exercise. Access management is the process of managing access to the resources.

Architecture and Design framework

We reviewed the architecture and design frameworks in depth in Chapter 12. Your organization must document your own architecture and framework that all development teams are required to follow and comply with the guardrails and guidelines specified in the framework. Most importantly, your architecture framework must ensure that your product are designed with data security as a built-in component of the architecture and as part of the modules and their interconnections.

Deployment Framework -- On-Prem and Cloud

On completion of development, applications are deployed on premises and or the public clouds. To ensure smooth, consistent, and secure deployment, you need to have a document that describes the deployment checklists and playbooks. They should describe the deployment, change management, monitoring, and emergency response procedures.

The Deployment framework should describe the following:

- The CI/CD pipelines and configuring the deployment environments.
- Storage and management of deployment images and use of IaC for deployments.
- Checklists and playbooks for the Site Reliability Engineers (SREs) who will be managing the production deployments.
- Deployment best practices for the production teams.

The policies and standards associated with application deployment should be fully socialized with the deployment teams and documented in the checklists.

Benefits of Data Security Governance

Establishing and rigorously following Data Security Governance is paramount for your organization due to the following benefits. These may be essential for you to not only meet any regulatory demands on your organization but also to ensure the trust of your customers.

- **Risk Mitigation:** By adopting DSG, your organization can identify potential security, privacy, and compliance risks associated with your data assets, and mitigate them. This may be essential to meet regulatory demands and establish customer trust.

Enterprise Data Security for US Europe and Asia

- *Business Continuity*: Effective DSG ensures data remains available and accessible to authorized users through an event that may include downtime caused by system or platform failures or loss of certain sections of operating environments.
- *Regulatory Compliance*: DSG helps organizations adhere to relevant data protection and compliance regulations, such as GDPR and SOX, and comply with data sovereignty requirements, thereby avoiding fines and reputational damage.
- *Data Privacy Protection*: The framework established by DSG includes privacy policies and controls that safeguard sensitive data, ensuring it is used appropriately and individual privacy rights are respected.
- *Safeguarding Mission-Critical Assets*: With data proliferation, protecting crown jewel datasets is paramount. DSG helps prevent unauthorized access and theft of your critical sensitive data and business information.
- *Cost Savings*: The costs associated with not implementing an effective DSG program can be substantial. While there is an obvious cost to establishing good governance and associated technology that drives this to fulfillment, the costs associated with regulatory fines, customer distrust and loss of sales and loss of reputation can be substantial.

Suffice it to say that these benefits far outweigh the costs and the resources required to establish and maintain effective DSG.

Enterprise Data Security for US Europe and Asia

- Access Controls --*, 74
- Access Controls responsibility, 104
- Active Directory (AD), 157
- Administrative accounts, 196
- Antivirus Software, 193
- Asia-Pacific Economic Cooperation (APEC) Privacy Framework, 28
- Asymmetric Encryption, 131
 - Diffie Hellman (DH) key exchange, 131
- Asymmetric encryption
 - Digital signature, 131
- Audit Plan, 47
- Australia Privacy Principles, 24
- Australian Privacy Act 1988. *See* International Laws
- Authenticity, 173
- AWS (Amazon Web Services),. *See* public clouds
- AWS IAM, 159

- BCDR
 - Business Continuity Plan (BCP), 228
 - Disaster Recovery Plan (DRP), 228
 - Plan, 226
 - Recovery Point Objective (RPO), 226
 - Recovery Time Objective (RTO), 227
- Behavioral Analysis, 194
- Business Continuity (BC), 225
- Business intelligence, 117
- Business resilience
 - Best practices for improving, 230

- C S U – S A D, 77
- CCPA - California Consumer Privacy Act. *See* International Laws

- CDSA (Common Data Security Architecture), 211
- Certification framework, 48
- Chain of Custody, 76
- Cloud Architecture
 - Intercloud, 122
 - Multicloud, 121
- Cloud Audits, 75
- Cloud deployment models, 64
 - Community cloud, 65
 - Hybrid cloud, 65
 - Private cloud, 64
 - Public cloud, 64
- Cloud Environments, 59
 - SANDBOX, 59
- Cloud management plane, 69
- Cloud Segmentation, 67
- Cloud service models, 61
 - IaaS, 61
 - PaaS, 63
 - Private IaaS, 62
 - SaaS, 63
- Cloud storage, 70
- COBIT5
 - Control Objectives for Information and Related Technologies, 51
- Compensating controls*, 73
- Container security, 70
- Cryptography
 - Asymmetric Encryption, 131
 - crypto-shredding, 127
 - Hardware security module (HSM), 128
 - Private Keys, 126
 - Public Key Infrastructure (PKI), 126
 - Symmetric Key Encryption, 130

Enterprise Data Security for US Europe and Asia

- Trusted platform modules (TPMs), 128
- CSA Treacherous Twelve, 182
- CSU-SAD
 - Collect, Store, Use, Share, Archive, Destroy, 250
- Customer Data Protection**, 89
- Cybersecurity, 14, 206
- Cybersecurity architecture, 197
 - Goals, 198
- Cybersecurity Architecture
 - Benefits, 204
 - Frameworks, 200
- Cybersecurity architecture framework, 210

- Data Access and Entitlements
 - Governance, 252
- Data Access Governance (DAG), 252
- Data Alteration
 - All-or-Nothing Transform with Reed-Solomon (AONT-RS), 139
 - Bit splitting, 140
 - Data Anonymization, 138
 - Data diddling, 136
 - Data masking, 136
 - Ddigital signature, 140
 - for Data protection, 136
 - SSMS (Secret Sharing Made Short), 139
 - Tokenization, 138
- Data Alteration Treatments for Data Protection, 136
- Data architecture, 113
 - Data integration, 117
- Data Architecture
 - Spanning, 121
 - Types of, 114
- Data at rest, 90
 - Securig, 218
- data breaches, 20

- Data classification
 - Governance, 248
- Data Classification, 80
- Data Classification by Type, 80
- Data confidentiality, 107
- Data Controller. *See* GDPR
- Data Destruction
 - Crypto-shredding, 140
- Data Discovery Tools, 90
- data encryption, 15
- Data encryption, 124
 - DES, 124
- Data forensics, 217
- Data Governance, 247
 - Framework, 247
- Data in motion*
 - Securing, 218
- Data in Motion, 91
- Data in stasis*
 - Securing, 219
- Data in use*
 - Securing, 219
- Data in Use, 92
- Data Integrity, 193
- Data lake, 118
- Data Lakehouse, 119
- Data lifecycle management (DLM), 250
- Data loss prevention, 16
- Data Loss Prevention (DLP)*, 19, 247
- Data Management
 - Best practices, 108
- Data minimization**. *See* GDPR
- Data model
 - Enterprise data model, 115
 - Entity-relationship, 116
 - Hierarchical, 115
 - Network, 116
 - NoSQL, 116
 - Object-oriented, 116
 - Relational, 116
- Data Processing. *See* GDPR

Enterprise Data Security for US Europe and Asia

- Data Processor. *See* Law
 - Data protection responsibility
 - Business partners, 98
 - Data Custodian, 100
 - Data owner, 99
 - Data Protection
 - Data backups, 95
 - Data loss prevention (DLP), 94
 - Encryption, 95
 - Best practices, 94
 - Data erasure, 95
 - Data resiliency, 95
 - Firewalls, 95
 - Framework, 94
 - Key elements, 93
 - Data Protection Laws, 20
 - Data protection Responsibility
 - Data Aggregator, 101
 - Data protection responsibility
 - Data Controller, 99
 - Data protection responsibility
 - Cloud Service Provider (CSP), 98
 - Data Processor, 100
 - Data Steward, 100
 - Data User, 101
 - Hosting Service Provider, 98
 - Data protection responsibility
 - Data Subject, 101
 - Data protection responsibility
 - The business on-premises IT, 97
 - Data quality
 - Governance, 249
 - Data Risk Assessment, 92
 - Data risk management, 147
 - Data security, 15
 - Reference Architecture, 209
 - Data Security, 206
 - Availability, 209
 - CIA Triad elements, 208
 - Confidentiality, 208
 - Four elements of, 208
 - Integrity, 208
 - Data Security Architectu
 - Layered technologies, 216
 - Data security architecture, 207
 - Data Security Architecture
 - Architecting and Designing, 220
 - Best practices, 220
 - for Audits, 223
 - Data Security Governance
 - Benefits, 253
 - Data Security Governance (DSG), 248
 - Data Security Posture Management, 83
 - Data source, 111
 - Data Subject. *See* GDPR
 - Database
 - Definition, 112
 - Database as a Service (DBaaS), 113
 - Defense in Depth, 71, 192
 - Defense-in-Depth
 - Benefits, 195
 - Department of Regulatory
 - Agencies (DORA), 27
 - Digital signature, 131
 - Disaster recovery, 60
 - Disaster Recovery (DR), 226
 - DLP, 16
 - Doctrine of Silver Platter*, 21
 - DSPM, 83
- EISA (Enterprise information Systems Architecture), 201
 - Encryption, 219
 - ALE (Application Level Encryption), 132
 - Format-preserving Encryption, 134
 - Homomorphic Encryption, 133
 - Self-decrypting storage disks, 135
 - Transparent Encryption, 135
 - Enterprise Risk Management (ERM), 72
 - EU Directive. *See* Laws

Enterprise Data Security for US Europe and Asia

- Exploit, 175
- Exploit Classification, 176
- FISMA - Federal Information Security Modernization Act 2014, 38
- forensic evidence, 21
- FTC (Federal Trade Commission), 23
- Full disk encryption, 219
- GCP
 - Workforce Identity Federation, 159
- GCP (Google Cloud Platform). *See* public clouds
- GCP IAM, 158
- GDPR, 27
- General Data Protection Regulation (GDPR). *See* International Laws
- GLBA – Gramm-Leach-Bliley Act (Financial Services Modernization Act of 1999). *See* International Laws
- Global standards, 37
 - COBIT5, 37
 - GDPR, 37
 - HIPAA, 37
 - ISO/IEC 27017, 37
 - NIST-RMF, 37
 - PCI-DSS, 37
- Global Standards
 - CSA CCM, 37
- Governance Risk and Compliance (GRC), 248
- Governance, Risk, and Compliance (GRC), 147
- HIPAA - Health Insurance Portability and Accountability Act. *See* International Laws
- Hybrid Cloud, 119
 - Data architecture and implementation, 119
- Hypervisor attacks
 - Hyperjacking, 186
 - Rootkits, 187
- Hypervisors, 21, 170
 - Type 1 hypervisors, 21
- Hypervisors
 - Attacks, 185
- IAM
 - Across the Enterprise, 156
 - Anomaly detection, 146
 - B2B identity, 146
 - Customer Identity and Access Management (CIAM), 146
 - Federated Identity., 146
 - FID (Federated Identity and Directory Service), 161
 - Identity and Access Management, 145
 - Multi-factor authentication (MFA), 146
 - Policy guidelines, 149
 - Policy Management, 148
 - Single Sign-On (SSO), 146
 - Workforce identity, 145
- Identity and Access Management (IAM), 142
 - Roles, 145
 - Workforce identity, 145
- Identity Governance and Administration (IGA), 164
- IGA Governance Model, 165
- II, 23
- India Digital Personal Data Protection Act 2023 (DPDPA). *See* International Laws
- Information Technology Infrastructure Library (ITIL), 53
- Infrastructure as a Service (IaaS)*, 61
- Integrity, 173
- Integrity checks, 21
- International Laws

Enterprise Data Security for US Europe and Asia

- GLBA – Gramm-Leach-Bliley Act (Financial Services Modernization Act of 1999), 25
- International Laws, 22
 - Australia Privacy Principles, 24
 - Australian Privacy Act 1988, 28
 - CCPA - California Consumer Privacy Act, 26
 - General Data Protection Regulation (GDPR), 22
 - HIPAA - Health Insurance Portability and Accountability Act, 25
 - India Digital Personal Data Protection Act 2023 (DPDPA), 24
 - Japan APPI, 23
 - Personal Data Protection Act, 29
 - The Personal Information Protection and Electronic Documents Act (PIPEDA), 27
- International Standards
 - ENISA, 38
 - ISO, 38
 - ISO/IEC, 38
- Internet Service Provider (ISP), 58
- Intrusion detection systems (IDS), 235
- ISO
 - ISO 27018, 24
 - ISO/IEC 27002, 236
 - ISO/IEC standards, 42
- IT infrastructure, 17
- IT Infrastructure Components, 200
- ITIL 4, 53
- Japan APPI. *See* International Laws
- key management, 15
- Least Privilege Enforce, 195
- Litigation hold. *See* Litigation Hold
- Litigation Hold Notice*, 21
- Logging, 66
- Managed Service Provider (MSP), 58
- Master data, 117
- Microsoft Azure. *See* public clouds
- Monitoring, 66
- MRA (Matters Requiring Attention), 75
- Multitenancy, 60
- Multi-Tenant Security, 204
- Network Security Controls, 193
- NIST
 - Families, 242
 - Five Principles, 241
- NIST 800-53, 241
- NIST Standards, 40
- Nonrepudiation, 174
- OECD – Organization for Economic Cooperation and Development, 36
- Office of the Comptroller of the Currency, 46
- Open Systems Interconnection (OSI) Architecture Framework, 199
- Operational Data store, 111
- OSA (Open Security Architecture), 201
- OWASP 10, 177
- Payment Card Industry Data Security Standard (PCI-DSS), 45
- Penetration testing, 175
 - Pivoting, 176
- Personal Data. *See* GDPR
- Personal Data Protection Act. *See* International Laws
- PI - Personal Information, 87
- PI - Personal information, 86
- PII - Personally identifiable information, 86
- Platform as a Service (PaaS)*, 61

Enterprise Data Security for US Europe and Asia

- Policies and Standards
 - Framework, 248
- Policy and Procedures, 233
- Potecting data
 - In motion, 91
- Preponderance of Evidence Principle*, 20
- Preventive Controls*, 73
- Protecting data
 - At rest, 90
 - in transit, 91
- public clouds, 17
- Public Key Infrastructure (PKI), 126
- Purpose limitation.** *See* GDPR

- QoS (Quality of Service), 75

- Regulatory Framework, 19
 - NIST, 19
- regulatory standards, 20
- Resource Optimization, 53
- Risk assessment, 231
- Risk Management, 47
- Risk optimization, 53
- RUM (Real User Monitoring) tools, 171

- SABSA (Sherwood Applied Business Security Architecture), 201
- Secure Hashing Algorithm, 131
 - SHA 256, 131
- Secure Service Deployment, 203
- Security control, 235
- Security of Physical Premises, 202
- Security Operations Center (SOC) reports, 50
- Security policies, 15
- Security Services
 - Basic Categories, 212
- Security-first architecture, 211
- sensitive data, 20
- Sensitive information, 86
- Sensitive Information, 88

- Service Level Agreement (SLA), 75
- Software as a service (SaaS)*, 61
- SOX – Sarbanes Oxley Act 2002, 26, *See* Laws
- Spoliation.* *See* Litigation Hold Notice
- SSAE SOC Reports, 51
- Statement on Standards for Attestation Engagements # 18, 50
- Storage limitation.** *See* GDPR
- Storage security management, 107
- STRIDE, 171
- Symmetric Key Encryption, 130
- System of Engagement, 110
- System of Record, 109
- System of Reference, 109

- The Cloud Security Alliance (CSA) STAR, 48
- The Personal Information Protection and Electronic Documents Act (PIPEDA), 27
- Third-party data transfers, 110
- Threat Actors, 196
- Threat Management, 18
- Threats
 - Denial of Service (DoS), 170
 - SQL Injection, 170
 - Zero-day vulnerability, 174
- TOGAF (The Open Group Architecture Framework), 201
- tokenization, 15
- Transport Layer Security (TLS), 129
- Trasnprt Layer Seucirty (TLS) TLS 1.3, 130
- Type 2 hypervisors. *See* Hypervisor

- US Department of Commerce, 23
- US Law
 - Case Law, 33
 - Civil Law, 33
 - Common Law Privacy Act, 35

Enterprise Data Security for US Europe and Asia

- Contract Law, 35
- Criminal Law, 33
- US Laws
 - Federal Information Security Modernization Act 2014, 38
- US Standards, 37
 - FedRAMP, 38, 39
 - FIPS 140-2, 37, 38
 - NIST, 38
 - PCI-DSS, 38
- Virtual Machine (VM), 169
- Virtual machines, 71
- Virtualization Security Vulnerabilities, 189
- VPN (e.g. IPSec Gateway), 140
- Workforce (Employee) Data Protection, 89**
- Zero Trust, 194
 - Best practices, 210
- Zero-day vulnerability, 174