Standards In Depth

Lesson 10

NIST Special Publications

- SP 800-12: An Introduction to computer Security. A handbook. Chapter 4 is common threats
- SP 800-14: Generally Accepted Principles and Practices for Securing Information Technology Systems
- SP 800-18: Guide for Developing Security Plans for Federal Information Systems
- SP 800-27:Engineering Principles for Information Technology Security (A Baseline for Achieving Security)
- SP 800-30:Risk Management Guide for Information Technology Systems (this has been superseded)
- SP 800-61: Computer Security Incident Handling Guide
- SP 800-64:Security Considerations in the System Development Life Cycle
- SP 800-100: Information Security Handbook: A Guide for Managers

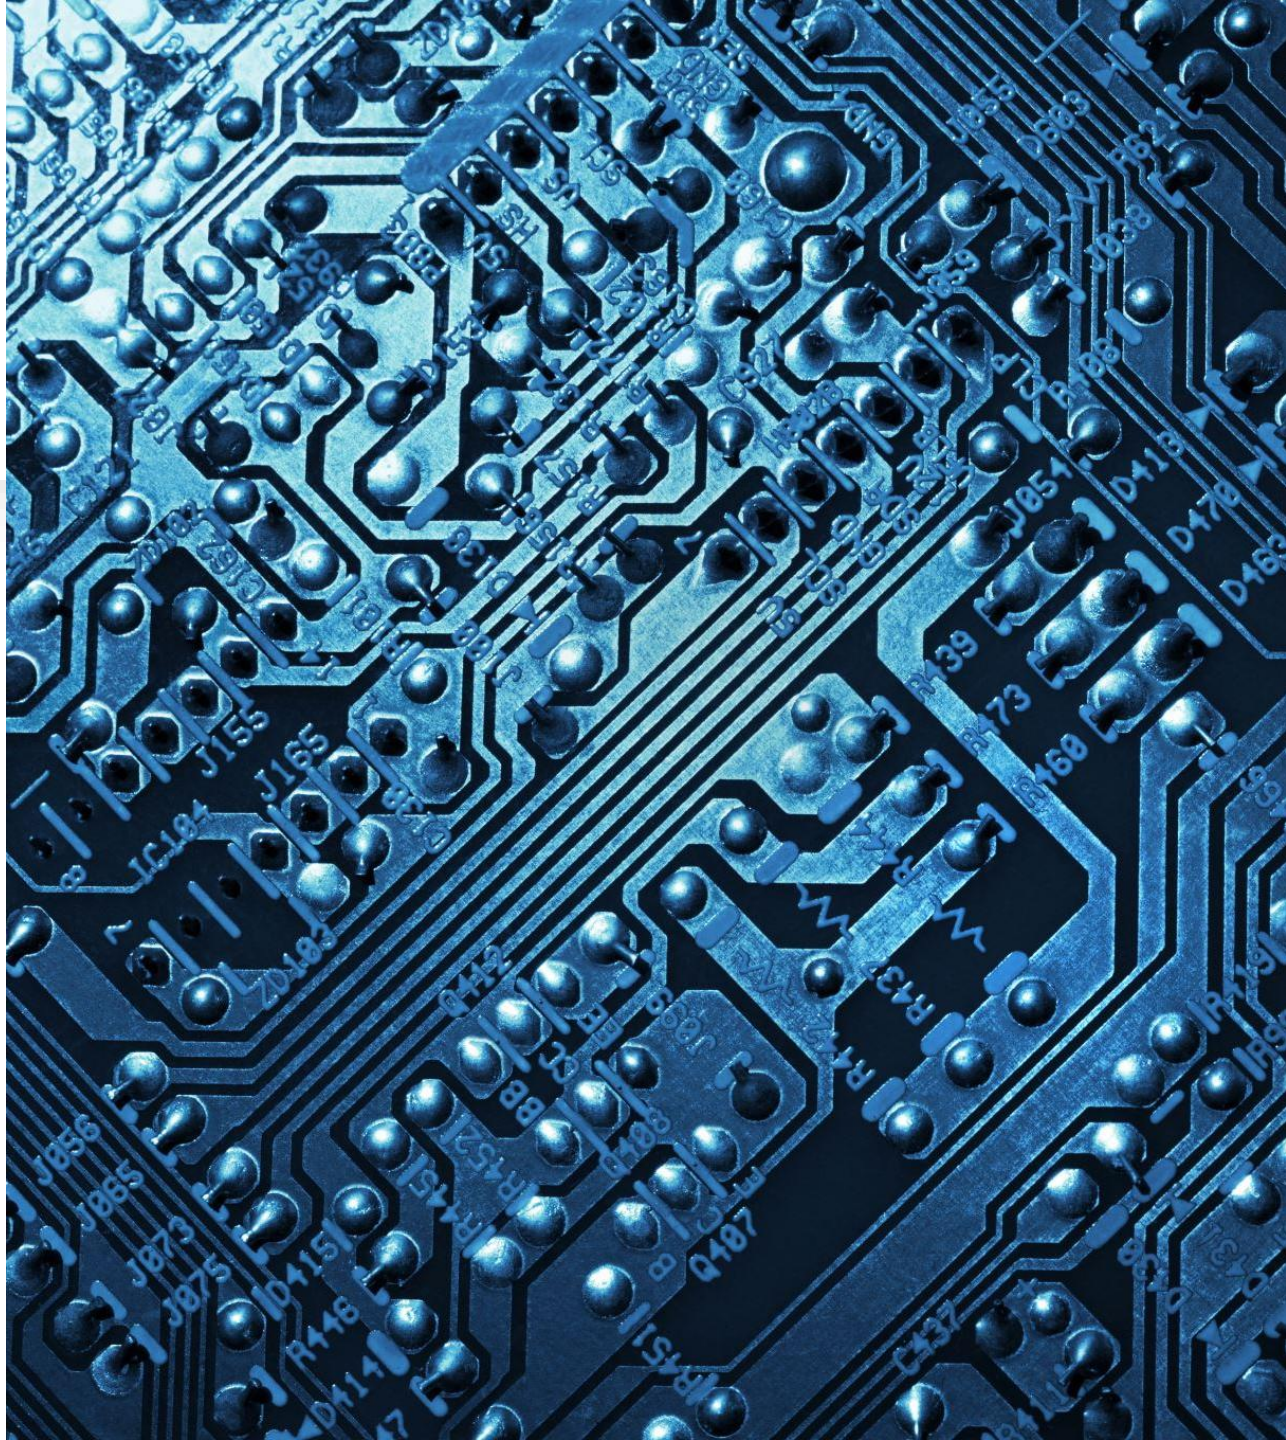

ISO

- ISO/IEC 15408: The Common Criteria for Information Technology Security Evaluation
- ISO/IEC 25000: Systems and Software Engineering
- ISO/IEC 27000: Information technology — Security Technology
- ISO/IEC 27001: Information Security Management
- ISO/IEC 27005: Risk Management
- ISO/IEC 27006: Accredited Certification Standard
- ISO/IEC 28000: Specification for security management systems for the supply chain

ISO/IEC 27000 family

- New family accepted in the spring 2005
  - Different numbers for BS 7799 based standards (ISO/IEC 17799 vice versa ISO/IEC 24743)
- Harmonization of two concepts
  - ISO/IEC 13335 – academic approach
  - BS 7799 – pragmatic approach
- ISO/IEC 27000 family criteria
  - Provide direct support or detailed guidance and interpretation for the implementation of the PDCA processes and requirements of defined in ISO/IEC 27001 (e.g. risk assessment, identification of assets, ISMS effectiveness)
  - Address conformity assessments or sector-specific requirements for ISMS
  - Contribute and add value to ISO/IEC 27001 of the PDCA processes
  - Specification of a relationship to ISO/IEC 27001
  - Standards excluded from this ISMS family are those that only address the implementation of controls from ISO/IEC 27002 – see WG4 activities

ISO/IEC 27001 Requirements

- General requirements (4.1)
- Establishing and managing the ISMS (4.2)
- Documentation requirements (4.3)
- Management commitment (5.1)
- Resource management (5.2)
- Internal ISMS audits (6)
- Management review (7)
- Continual improvement (8.1)
- Corrective action (8.2)
- Preventive action (8.3)
- Annex A – Control objectives and controls

(A total of 35 Control Objectives and 114 Controls are grouped under 14 main categories as listed out in Table A.1 of ISO/IEC 27001)

Certification of ISMS to ISO/IEC 27001

Certification is an attestation issued by **a third-party body**, through a **formal conformity assessment process**, that specified requirements (e.g., ISO/IEC 27001) are fulfilled.

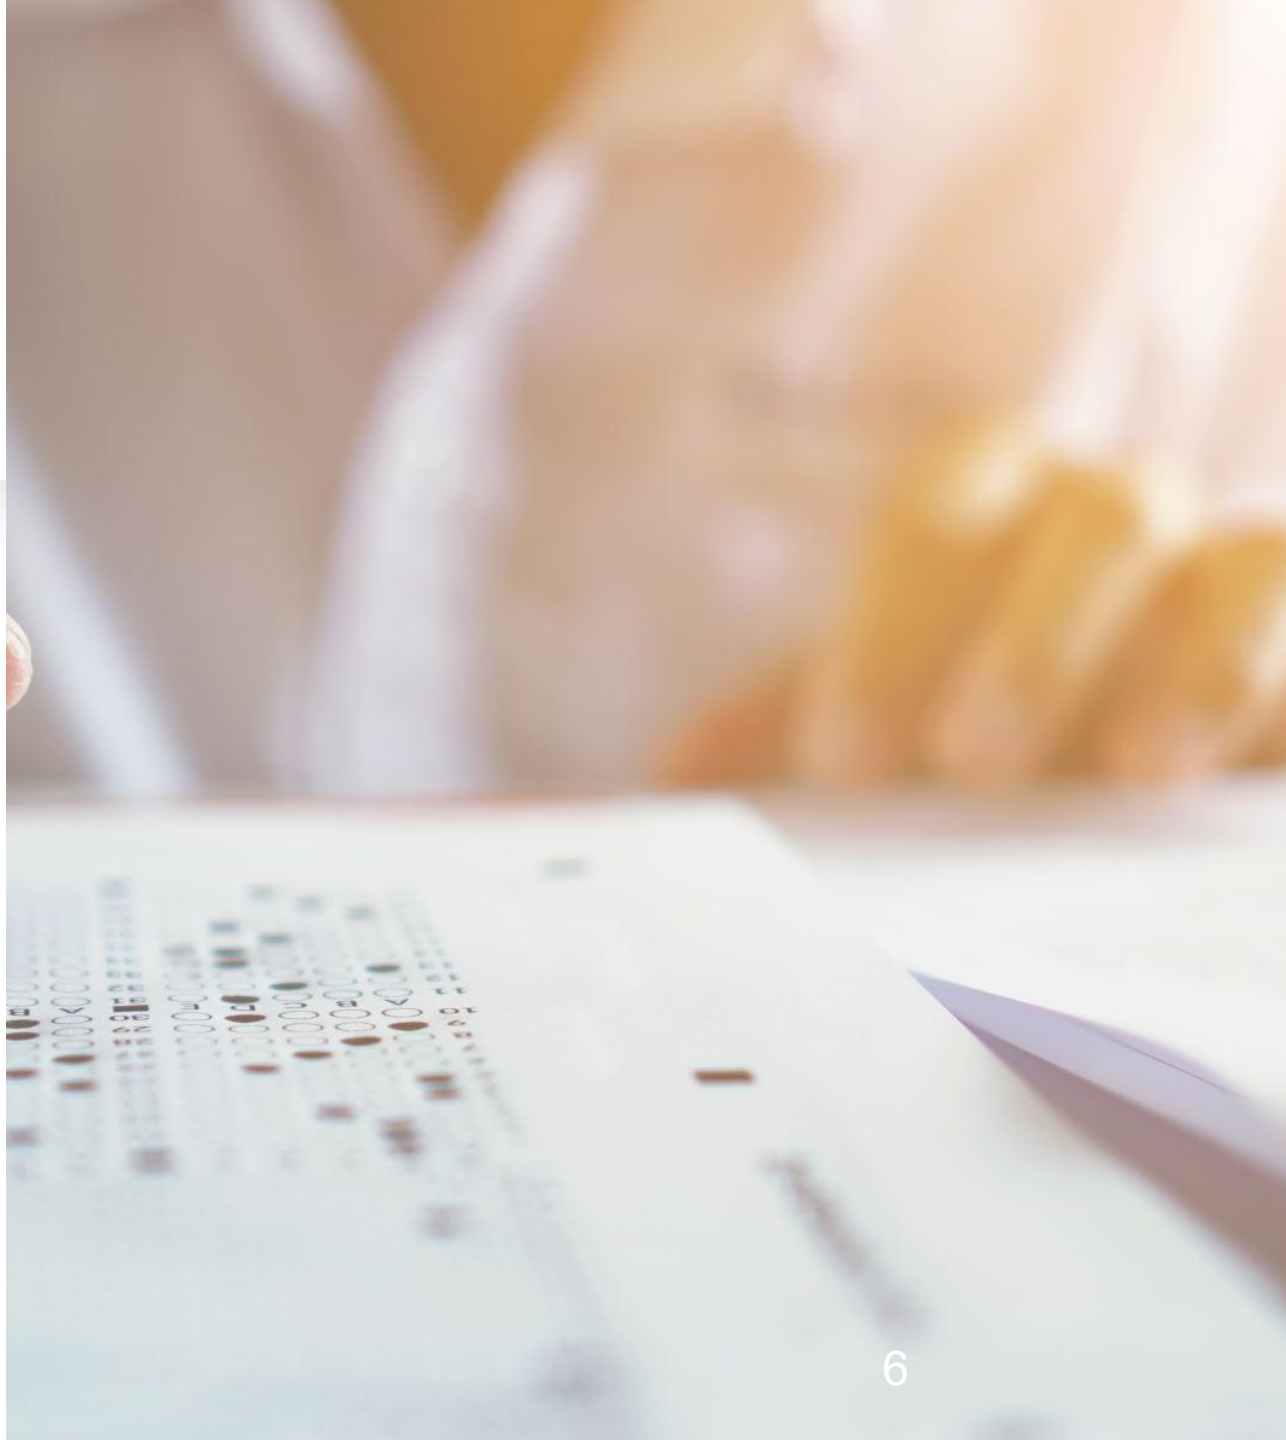

What is accreditation ?

- According to ISO/IEC 17000:2004 Conformity assessment – Vocabulary and general principles:
  - “Accreditation” – Issuance of conformance statement by a third party (i.e. accreditation body)

Cryptography and Security Mechanisms

- ISO/IEC 7064 Data processing – Check character systems
    - Published 2003
- ISO/IEC 9796 Digital signature schemes giving message recovery
    - 3 parts published 2002 – 2006, under revision
- ISO/IEC 9797 Message authentication codes (MACs)
    - 2 parts published 1999 – 2002, under revision, 3rd part is upcoming
- ISO/IEC 9798Entity authentication
    - 6 parts published 1997 – 2005
- ISO/IEC 10116 Modes of operation for an n-bit block cipher algorithm
    - Published 2006
- ISO/IEC 10118 Hash-functions
    - 4 parts published 1998 – 2004 (2006), under revision
- ISO/IEC 11770 Key management
    - 4 parts published 1996 – 2006, under revision

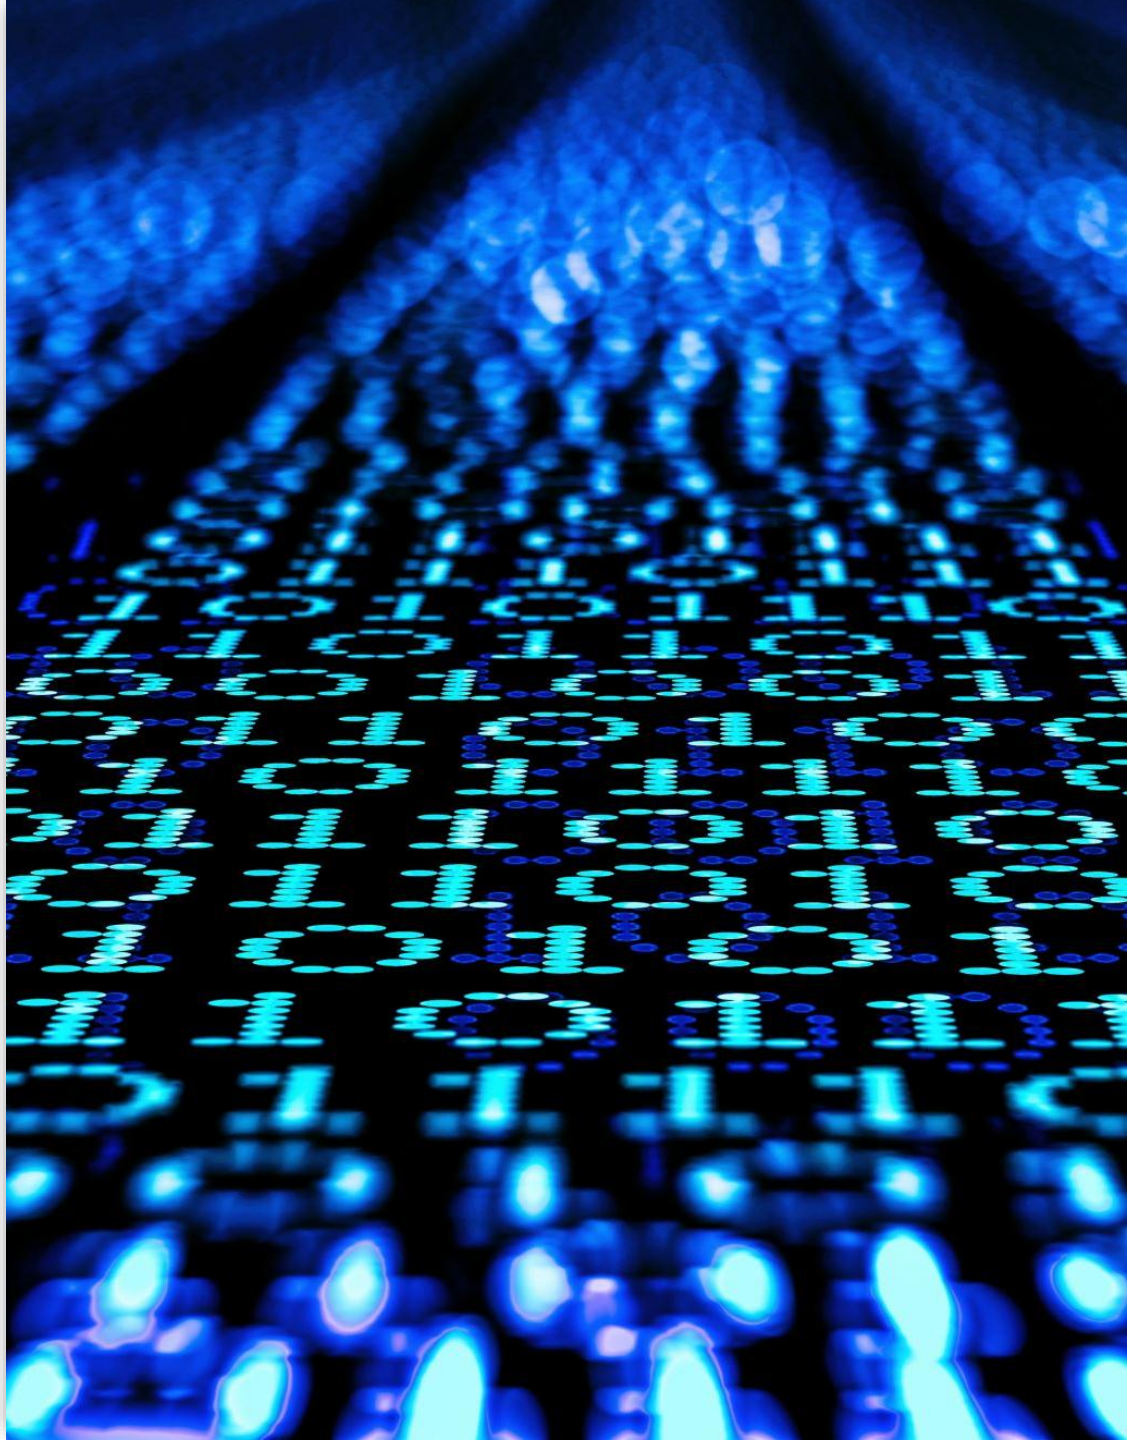

Cryptography and Security Mechanisms

- ISO/IEC 13888 Non-repudiation
    - 3 parts published 1997 – 2004, under revision
- ISO/IEC 14888 digital signatures with appendix
    - 3 parts published 1998 – 2006, under revision
- ISO/IEC 1594 Cryptographic techniques based on elliptic curves
    - 4 parts published 2002 – 2004, under revision
- ISO/IEC 18031 Random bit generation
    - Published 2005
- ISO/IEC 18032 Prime number generation
    - Published 2005
- ISO/IEC 18033Encryption algorithms
    - 4 parts published in 2005 – 2006
- ISO/IEC 19772 Data encapsulation mechanisms
    - Upcoming

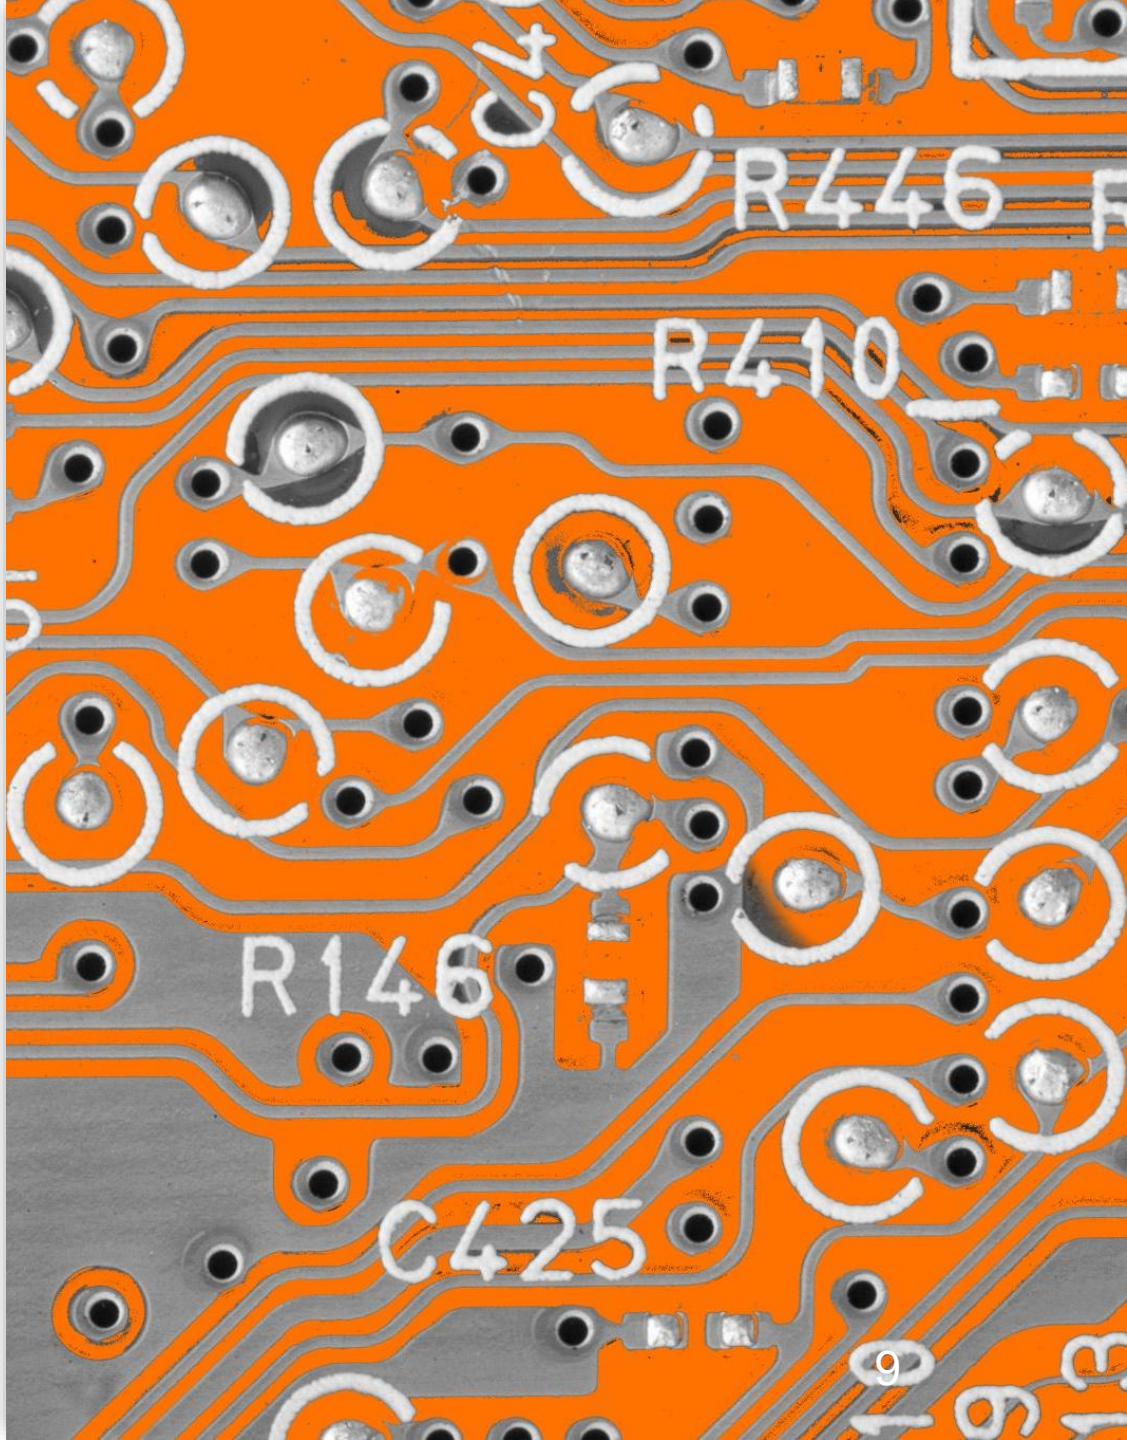

Security Evaluation and Assessment

Common criteria etc.

- ISO/IEC 15408:2005
  Evaluation criteria for IT Security
- ISO/IEC 18045:2005
  Methodology for IT security evaluation
- ISO/IEC TR 19791:2006
  Security assessment of operational systems
- ISO/IEC 15292:2003
  Protection profile registration procedures
- ISO/IEC TR 15446:2004
  Guide on the production of protection profiles and security targets

- ISO/IEC 19790:2006
  Security requirements for cryptographic modules
- ISO/IEC 24759
  Test requirements for cryptographic modules
- ISO/IEC 21827:2003
  Systems Security Engineering – Capability Maturity Model (SSE-CMM)
- ISO/IEC 15443
  A framework for IT security assurance

NIST SP 800-207

1. All data sources and computing services are considered resources. A network may be composed of multiple classes of devices. A network may also have small footprint devices that send data to aggregators/storage, software as a service (SaaS), systems sending instructions to actuators, and other functions. Also, an enterprise may decide to classify personally owned devices as resources if they can access enterprise-owned resources.
2. All communication is secured regardless of network location. Network location alone does not imply trust. Access requests from assets located on enterprise-owned network infrastructure (e.g., inside a legacy network perimeter) must meet the same security requirements as access requests and communication from any other nonenterprise-owned network. In other words, trust should not be automatically granted based on the device being on enterprise network infrastructure. All communication should be done in the most secure manner available, protect confidentiality and integrity, and provide source authentication.
3. Access to individual enterprise resources is granted on a per-session basis. Trust in the requester is evaluated before the access is granted. Access should also be granted with the least privileges needed to complete the task. This could mean only “sometime recently” for this particular transaction and may not occur directly before initiating a session or performing a transaction with a resource. However, authentication and authorization to one resource will not automatically grant access to a different resource.

NIST SP 800-207

. Access to resources is determined by dynamic policy—including the observable state of client identity, application/service, and the requesting asset—and may include other behavioral and environmental attributes. An organization protects resources by defining what resources it has, who its members are (or ability to authenticate users from a federated community), and what access to resources those members need. For zero trust, client identity can include the user account (or service identity) and any associated attributes assigned by the enterprise to that account or artifacts to authenticate automated tasks. Requesting asset state can include device characteristics such as software versions installed, network location, time/date of request, previously observed behavior, and installed credentials. Behavioral attributes include, but not limited to, automated subject analytics, device analytics, and measured deviations from observed usage patterns. Policy is the set of access rules based on attributes that an organization assigns to a subject, data asset, or application. Environmental attributes may include such factors as requestor. network location, time, reported active attacks, etc. These rules and attributes are based on the needs of the business process and acceptable level of risk. Resource access and action permission policies can vary based on the sensitivity of the resource/data. Least privilege principles are applied to restrict both visibility and accessibility

NIST SP 800-207

1. The entire enterprise private network is not considered an implicit trust zone.
2. No resource is inherently trusted.
3. Remote enterprise subjects and assets cannot fully trust their local network connection.

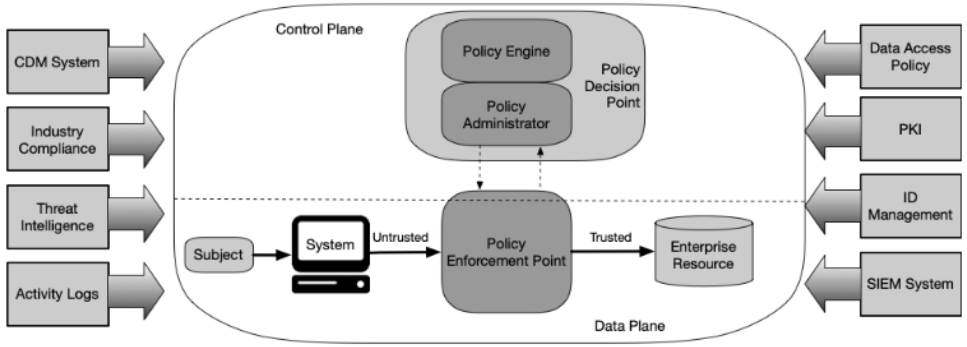

Figure 2: Core Zero Trust Logical Components

NIST SP 800-207

Zero Trust Architecture

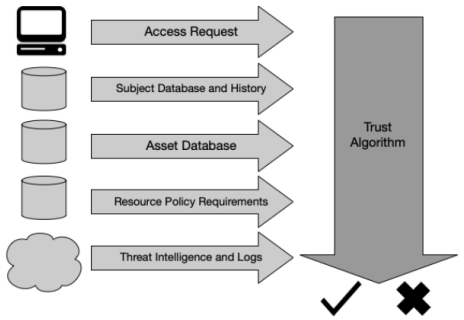

Figure 7: Trust Algorithm Input

NIST SP 800-207

Zero Trust Architecture

- **Policy engine (PE):** This component is responsible for the ultimate decision to grant access to a resource for a given subject. The PE uses enterprise policy as well as input from external sources (e.g., CDM systems, threat intelligence services described below) as input to a trust algorithm (see Section 3.3 for more details) to grant, deny, or revoke access to the resource. The PE is paired with the policy administrator component. The policy engine makes and logs the decision (as approved, or denied), and the policy administrator executes the decision.
- **Policy administrator (PA):** This component is responsible for establishing and/or shutting down the communication path between a subject and a resource (via commands to relevant PEPs). It would generate any session-specific authentication and authentication token or credential used by a client to access an enterprise resource. It is closely tied to the PE and relies on its decision to ultimately allow or deny a session. If the session is authorized and the request authenticated, the PA configures the PEP to allow the session to start. If the session is denied (or a previous approval is countermanded), the PA signals to the PEP to shut down the connection. Some implementations may treat the PE and PA as a single service; here, it is divided into its two logical components. The PA communicates with the PEP when creating the communication path. This communication is done via the control plane.
- **Policy enforcement point (PEP):** This system is responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource. The PEP communicates with the PA to forward requests and/or receive policy updates from the PA. This is a single logical component in ZTA but may be broken into two different components: the client (e.g., agent on a laptop) and resource side (e.g., gateway component in front of resource that controls access) or a single portal component that acts as a gatekeeper for communication paths. Beyond the PEP is the trust zone (see Section 2) hosting the enterprise resource.

NIST SP 800-205 Network Requirements for Zero Trust

- 1. Enterprise assets have basic network connectivity. The local area network (LAN), enterprise controlled or not, provides basic routing and infrastructure (e.g., DNS). The remote enterprise asset may not necessarily use all infrastructure services.
  2. The enterprise must be able to distinguish between what assets are owned or managed by the enterprise and the devices' current security posture. This is determined by enterprise-issued credentials and not using information that cannot be authenticated information (e.g., network MAC addresses that can be spoofed).
  3. The enterprise can observe all network traffic. The enterprise records packets seen on the data plane, even if it is not be able to perform application layer inspection (i.e., OSI layer 7) on all packets. The enterprise filters out metadata about the connection (e.g., destination, time, device identity) to dynamically update policies and inform the PE as it evaluates access requests

NIST SP 800-205 Network Requirements for Zero Trust

- 4. Enterprise resources should not be reachable without accessing a PEP (Policy Enforcement Point). Enterprise resources do not accept arbitrary incoming connections from the internet. Resources accept custom-configured connections only after a client has been authenticated and authorized. These communication paths are set up by the PEP. Resources may not even be discoverable without accessing a PEP. This prevents attackers from identifying targets via scanning and/or launching DoS attacks against resources located behind PEPs. Note that not all resources should be hidden this way; some network infrastructure components (e.g., DNS servers) must be accessible.
- 5. The data plane and control plane are logically separate. The policy engine, policy administrator, and PEPs communicate on a network that is logically separate and not directly accessible by enterprise assets and resources. The data plane is used for application/service data traffic. The policy engine, policy administrator, and PEPs use the control plane to communicate and manage communication paths between assets. The PEPs must be able to send and receive messages from both the data and control planes.

- 6. Enterprise assets can reach the PEP component. Enterprise subjects must be able to access the PEP component to gain access to resources. This could take the form of a web portal, network device, or software agent on the enterprise asset that enables the connection.
  7. The PEP is the only component that accesses the policy administrator as part of a
  business flow. Each PEP operating on the enterprise network has a connection to the policy administrator to establish communication paths from clients to resources. All enterprise business process traffic passes through one or more PEPs.
  8. Remote enterprise assets should be able to access enterprise resources without
  needing to traverse enterprise network infrastructure first. For example, a remote
  subject should not be required to use a link back to the enterprise network (i.e., virtual private network [VPN]) to access services utilized by the enterprise and hosted by a public cloud provider (e.g., email).

- 9. The infrastructure used to support the ZTA access decision process should be made scalable to account for changes in process load. The PE(s), PA(s), and PEPs used in a ZTA become the key components in any business process. Delay or inability to reach a PEP (or inability of the PEPs to reach the PA/PE) negatively impacts the ability to perform the workflow. An enterprise implementing a ZTA needs to provision the components for the expected workload or be able to rapidly scale the infrastructure to handle increased usage when needed.
  10. Enterprise assets may not be able to reach certain PEPs due to policy or observable factors. For example, there may be a policy stating that mobile assets may not be able to reach certain resources if the requesting asset is located outside of the enterprise's home country. These factors could be based on location (geolocation or network location), device type, or other criteria.

Zero Trust and Government Services Organizations

- There is no single technology, product, or service that can achieve the goals of implementing a ZTA. A truly effective ZTA incorporates technologies that:

- Authenticate, monitor, and validate user identities and trustworthiness.
- Identify, monitor, and manage devices and other endpoints on a network.
- Control and manage access to and data flows within networks.
- Secure and accredit applications within a technology stack.
- Automate security monitoring and connect tools across information systems.
- Analyze user behavior and other data to observe real-time events and proactively orient network defenses.
- Support IPv4 and IPv6.

OMG M-22-09 MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

The strategic goals set forth in this memorandum align with CISA's five pillars:
1. Identity: Agency staff use enterprise-managed identities to access the applications they use in their work. Phishing-resistant MFA protects those personnel from sophisticated online attacks.
2. Devices: The Federal Government has a complete inventory of every device it operates and authorizes for Government use, and can prevent, detect, and respond to incidents on those devices.
3. Networks: Agencies encrypt all DNS requests and HTTP traffic within their environment, and begin executing a plan to break down their perimeters into isolated environments.
4. Applications and Workloads: Agencies treat all applications as internet-connected, routinely subject their applications to rigorous empirical testing, and welcome external vulnerability reports.
5. Data: Agencies are on a clear, shared path to deploy protections that make use of thorough data categorization. Agencies are taking advantage of cloud security services to monitor access to their sensitive data, and have implemented enterprise-wide logging and information sharing.

https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf

OMG M-22-09 MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

Agencies must employ centralized identity management systems for agency users that can be integrated into applications and common platforms.
2. Agencies must use strong MFA throughout their enterprise.
• MFA must be enforced at the application layer, instead of the network layer.
• For agency staff, contractors, and partners, phishing-resistant MFA is required.
• For public users, phishing-resistant MFA must be an option.
• Password policies must not require use of special characters or regular rotation.
3. When authorizing users to access resources, agencies must consider at least one device-level signal alongside identity information about the authenticated user.

https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf

OMG M-22-09 MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

- Agencies must resolve DNS queries using encrypted DNS wherever it is technically supported.
- CISA's Protective DNS program will support encrypted DNS requests.

2. Agencies must enforce HTTPS for all web and application program interface (API) traffic in their environment.
- Agencies must work with CISA to "preload" their .gov domains into web browsers as only accessible over HTTPS.

3. CISA will work with FedRAMP to evaluate viable Government-wide solutions for encrypted email in transit and to make resulting recommendations to OMB.
4. Agencies must develop a zero trust architecture plan that describes the agency's approach to environmental isolation in consultation with CISA and submit it to OMB as part of their zero trust implementation plan.

- https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf

CISA (Cybersecurity & Infrastructure Security Agency) Zero Trust Maturity Model

The Zero Trust Maturity Model represents a gradient of implementation across five distinct pillars, where minor advancements can be made over time toward optimization. The pillars, depicted in Figure 1, include Identity, Device, Network, Application Workload, and Data. Each pillar also includes general details regarding Visibility and Analytics, Automation and Orchestration, and Governance. This maturity model is one of many paths to support the transition to zero trust.

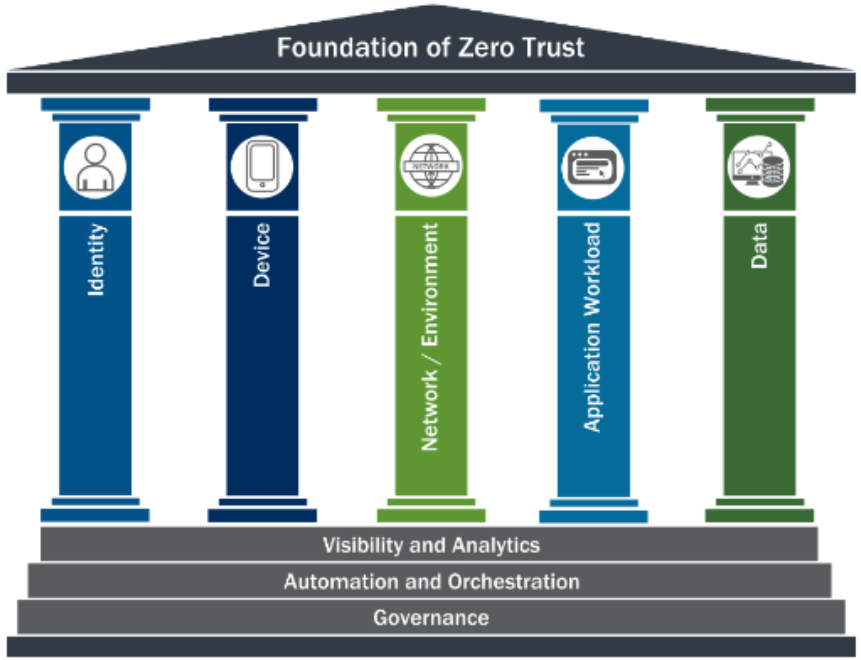

Figure 1: Foundation of Zero Trust[7]

CISA (Cybersecurity & Infrastructure Security Agency

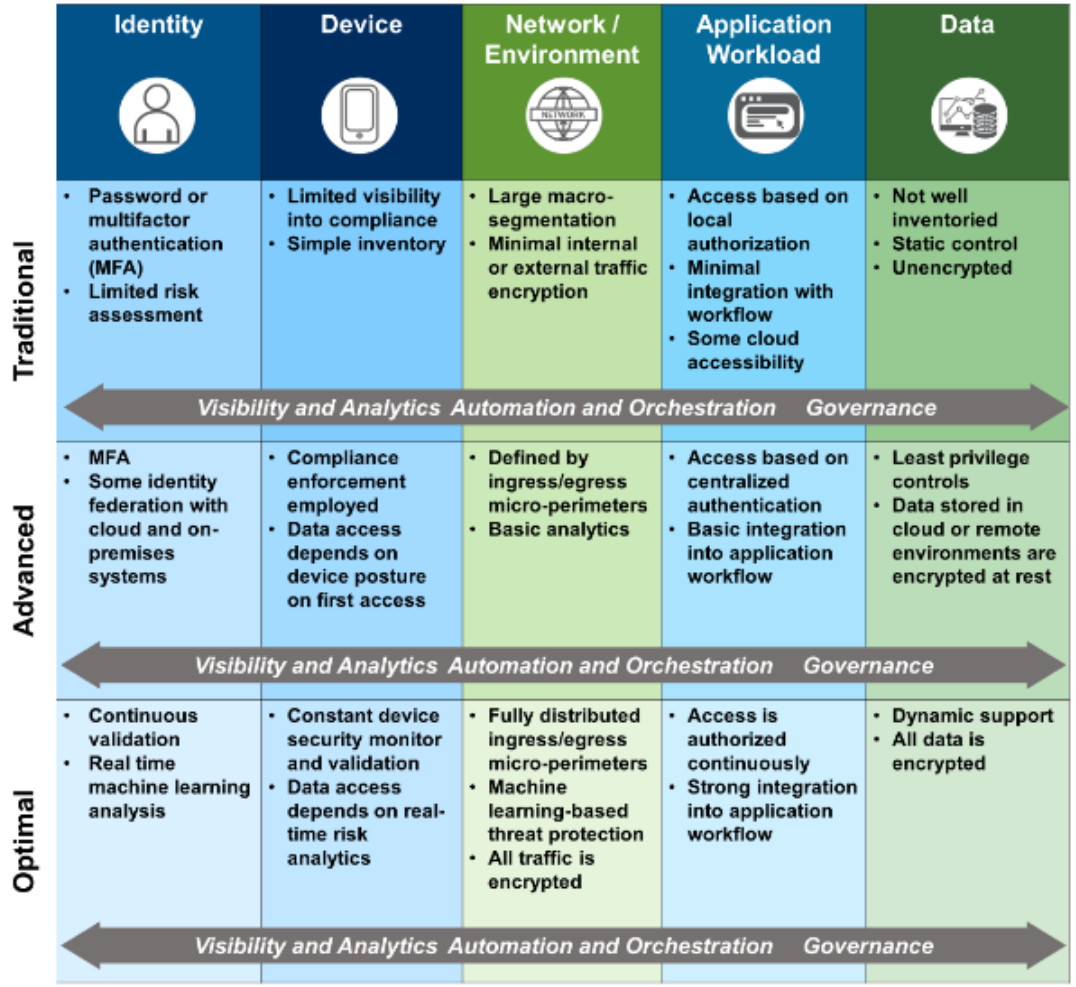

Figure 2: High-Level Zero Trust Maturity Model

https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf

CISA (Cybersecurity & Infrastructure Security

Function	Traditional	Advanced	Optimal
Authentication	Agency authenticates identity using either passwords or multi-factor authentication (MFA).	Agency authenticates identity using MFA.	Agency continuously validates identity, not just when access is initially granted.
Identity Stores	Agency only uses on-premises identity providers.	Agency federates some identity with cloud and on-premises systems.	Agency has global identity awareness across cloud and on-premises environments.
Risk Assessment	Agency makes limited determinations for identity risk.	Agency determines identity risk based on simple analytics and static rules.	Agency analyzes user behavior in real time with machine learning algorithms to determine risk and deliver ongoing protection.
Visibility and Analytics Capability	Agency segments user activity visibility with basic and static attributes.	Agency aggregates user activity visibility with basic attributes and then analyzes and reports for manual refinement.	Agency centralizes user visibility with high fidelity attributes and user and entity behavior analytics (UEBA).
Automation and Orchestration Capability	Agency manually administers and orchestrates (replicates) identity and credentials.	Agency uses basic automated orchestration to federate identity and permit administration across identity stores.	Agency fully orchestrates the identity lifecycle Dynamic user profiling, dynamic identity and group membership, just-in-time and just-enough access controls are implemented.
Governance Capability	Agency manually audits identities and permissions after initial provisioning using static technical enforcement of credential policies (e.g., complexity, reuse, length, clipping, MFA, etc.).	Agency uses policy-based automated access revocation. There are no shared accounts.	Agency fully automates technical enforcement of policies. Agency updates policies to reflect new orchestration options.

https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf

NIST SPECIAL PUBLICATION 1800-35B

Implementing a Zero Trust Architecture

- Authentication and periodic reauthentication of the requesting user's identity
- Authentication and periodic reauthentication of the requesting endpoint
- Authentication and periodic reauthentication of the endpoint that is hosting the resource being accessed

In addition, the following capabilities are also considered highly desirable:

- Verification and periodic reverification of the requesting endpoint's health
- Verification and periodic reverification of the health of the endpoint that is hosting the resource being accessed

DoD Zero Trust Reference

- **Defense Enterprise Identity, Credential, and Access Management (ICAM):** which includes Identity Provider (IDP), Automatic Account Provisioning (AAP) and a Master User Record (MUR), identifies and manages the roles, access privileges, and the circumstances in which users are granted or denied privileges.
    - IDP: A system that performs direct authentication and optionally can provide authorization data on behalf of one or more information systems. This system also provides authentication for NPE's.
    - AAP: Provides identity governance services such as user entitlement management, business role auditing and enforcement and account provisions and deprovisioning based on identity data produced during DOD people-centric activities such as on and off-boarding, continuous vetting, talent management and readiness training.
    - MUR: Enables DOD-wide knowledge, audit, and data rollup reporting of who has access to what system or applications. MUR will also provide support in identifying insider and external threats.
- **Client and Identity Assurance:**
    - Authentication Decision Point: This evaluates the identity of the user, NPE, and or device as access is attempted to applications and data. Devices may also be evaluated as to whether they are managed or unmanaged. Additional use cases for non-user NPE and user assisted NPE are available in the ICAM Reference Design.
    - Authorization Decision Point: A system entity that makes authorization decisions for entities that request such access decisions. It examines requests to access resources and compares them to the policy that applies to all requests for accessing that resource to determine whether specific access should be granted to the requester who issued the request under consideration. The client and device authorizations are the first stage in conditional access to resources, applications, and ultimately the data.

https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf

https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf

DoD Zero Trust Reference Architecture

- Capabilities:
  - Macro Segmentation - Macro-segmentation, the concept of dividing a network into smaller, controlled segments with different attributes, can be achieved through the application of additional hardware or VLANs.
  - Application Delivery Control (Proxy) - An application delivery controller is a device that is typically placed in a data center between the firewall and one or more application servers (an area known as the DMZ). Application delivery controllers primarily perform application acceleration and handle enterprise-level load balancing between servers. Earlier generations of Application Delivery Controllers can handle a variety of tasks including, but not limited to, content-caching, SSL offload and acceleration services, data compression as well some intrusion prevention services.

https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf

DoD Zero Trust Reference Architecture

Capabilities:

- Micro segmentation - This is the practice of creating logical network zones to isolate segments. These segments are secured by enabling granular access control, whereby users, applications, workloads, and devices are segmented based on logical attributes. This also provides an advantage over traditional perimeter security, as the smaller segments present a reduced attack surface (for malicious personas). In a Zero Trust Architecture, security settings can be applied to different types of traffic, creating policies that limit network and application flows between workloads to those that are explicitly permitted. Segmentation Gateways and API access decision points can limit access on a per identity basis to explicitly allowed API invocations, with allowance granularity down to the "verb" level.
- DevSecOps Application Development – DevSecOps is a set of software development practices that combines software development (Dev), security (Sec), and information technology operations (Ops) to secure the outcome and shorten the development lifecycle. Software features, patches, and fixes occur more frequently and in an automated fashion. Security is applied at all phases of the software lifecycle. Adoption of DevSecOps applies to application development and production environments equally
- Data Authorization Decision Point: Data owners use Data Reference Architecture to apply tagging to data via orchestrator or DLP/DRP Servers.

https://dodcio.defense.gov/Portals/0/Documents/Library/(U)ZT_RA_v1.1(U)_Mar21.pdf

DoD Zero Trust Reference Architecture

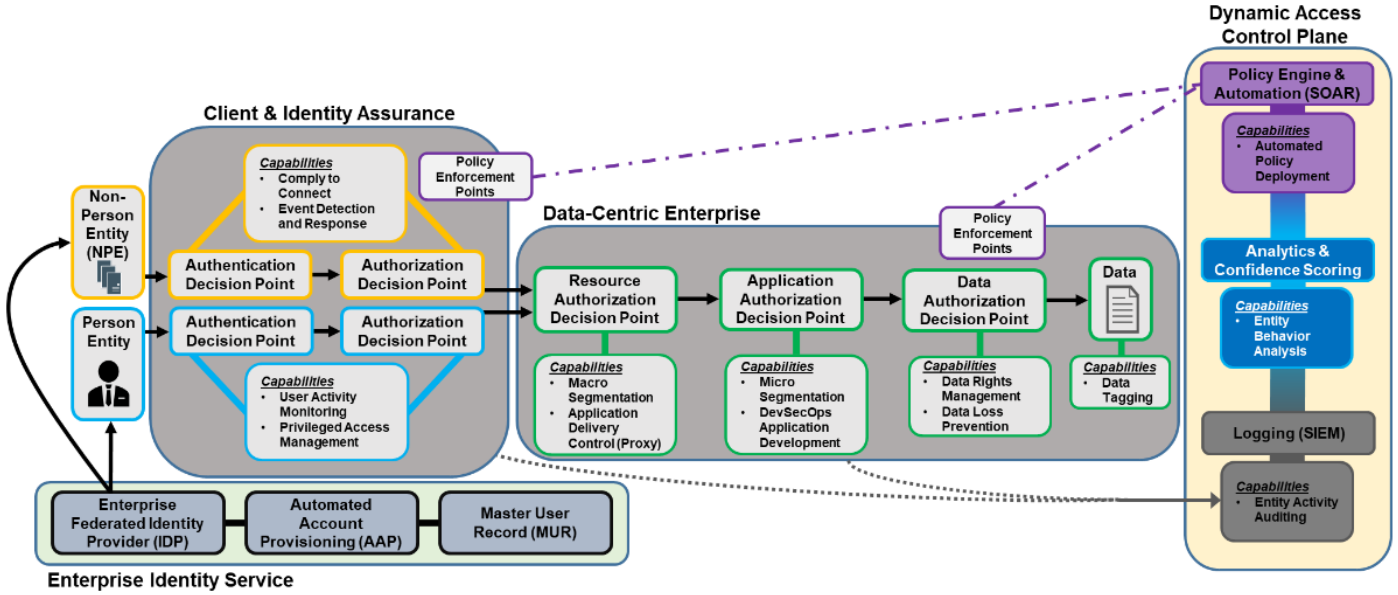

Figure 2: High-Level Operational Concept (OV-1)

2.1 RISK MANAGEMENT PROCESS

Risk assessment is a key component of a holistic, organization-wide *risk management process* defined in NIST Special Publication 800-39, *Managing Information Security Risk: Organizati Mission, and Information System View*. Risk management processes include: (i) framing risk; (assessing risk; (iii) responding to risk; and (iv) monitoring risk. Figure 1 illustrates the four ste in the risk management process—including the risk assessment step and the information and communications flows necessary to make the process work effectively.[13]

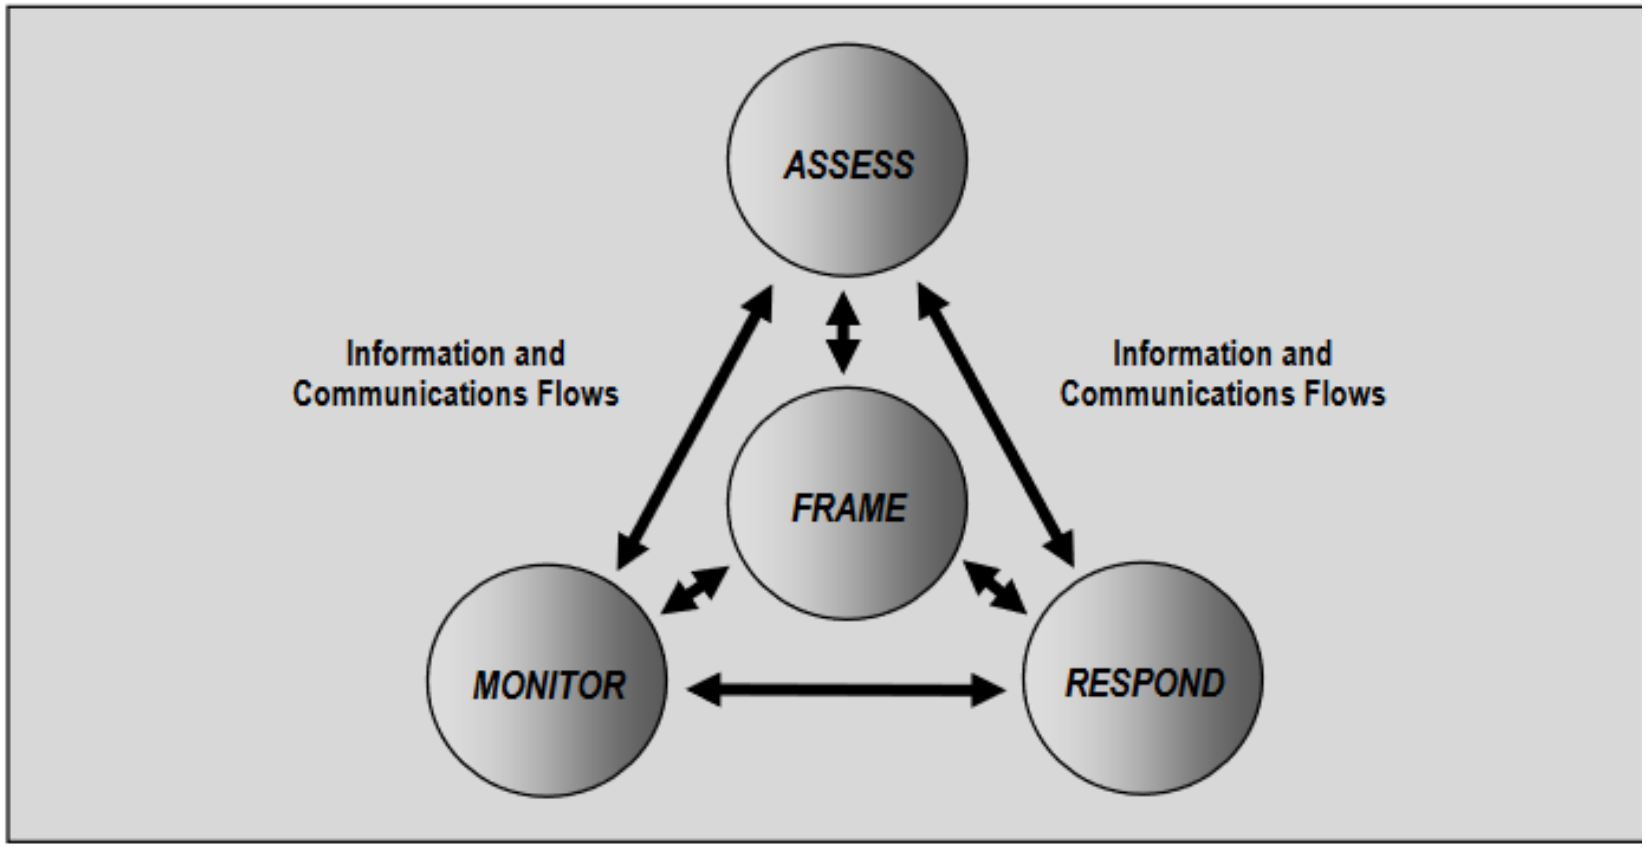

NIST 800-30 is the U.S. standard for how to conduct risk assessments. The standard is divided into three chapters. The first is just introductory information such as the target audience and purpose of the standard. Chapter 2 discusses the risk management process and concepts. Chapter 3 provides a process for conducting a risk assessment.

NIST 800-30

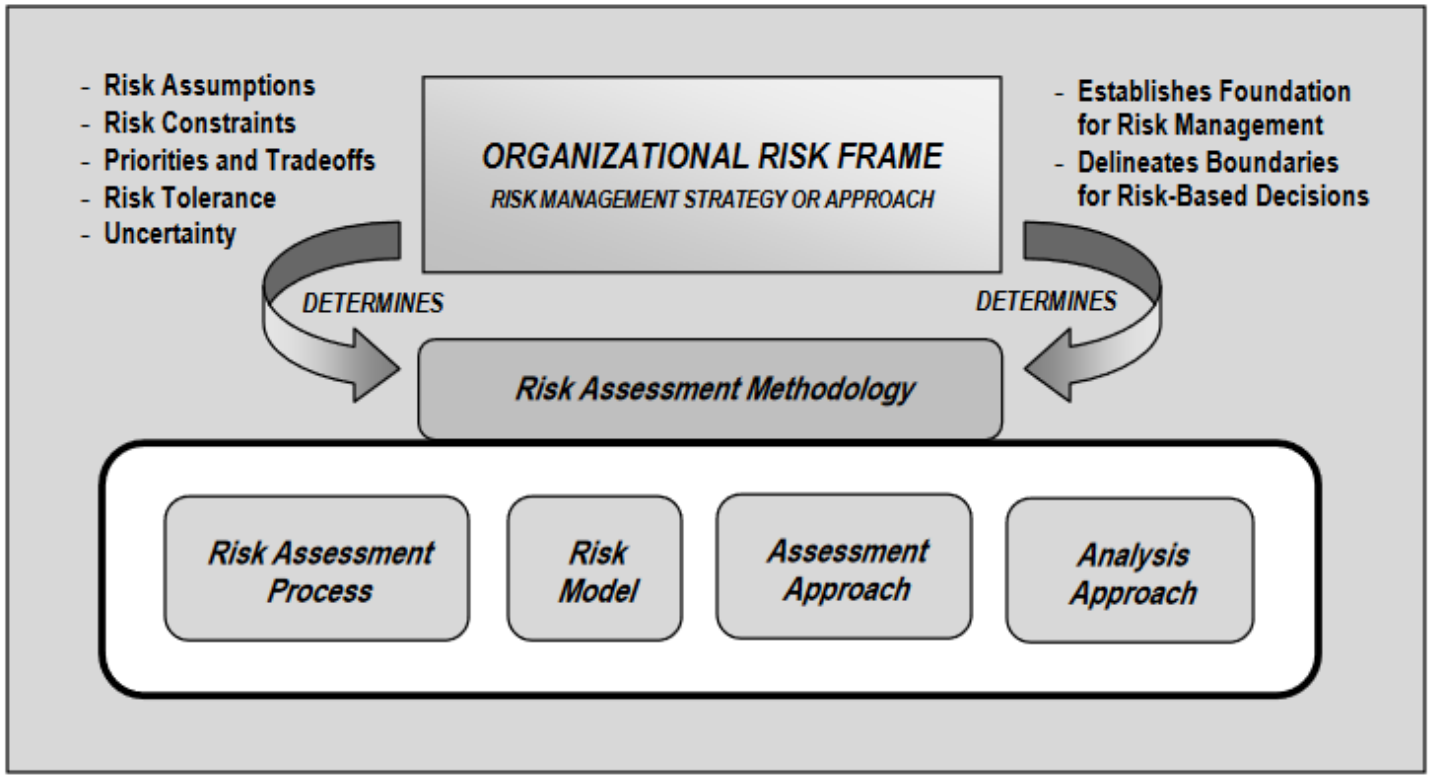

FIGURE 2: RELATIONSHIP AMONG RISK FRAMING COMPONENTS

NIST 800-30

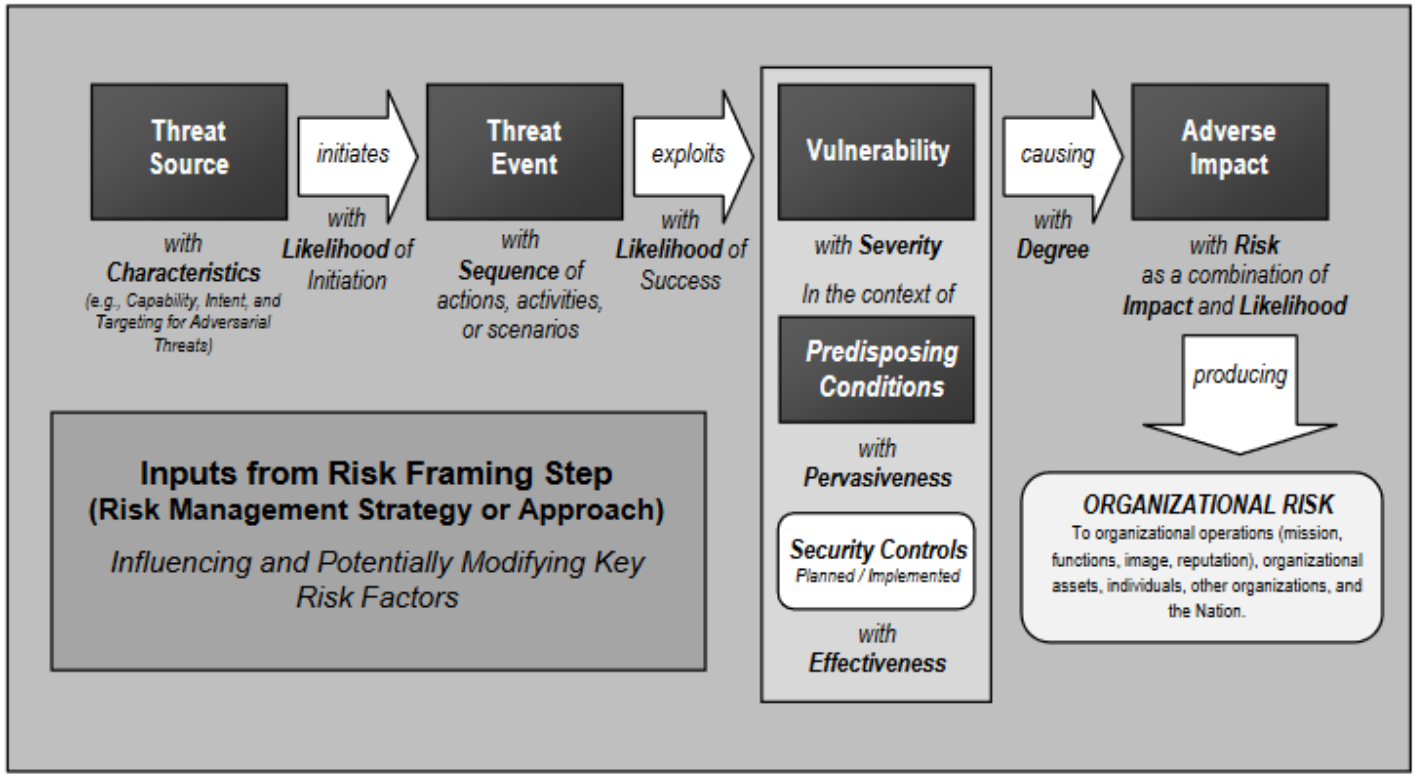

FIGURE 3: GENERIC RISK MODEL WITH KEY RISK FACTORS

ISO 27017

- ISO 27017 is guidance for cloud security. It does apply the guidance of ISO 27002 to the cloud, but then adds 7 new controls.
  - CLD.6.3.1: Agreement on shared or divided security responsibilities between the customer and cloud provider
  - CLD.8.1.5: Addresses how assets are returned or removed from the cloud when the contract is terminated
  - CLD.9.5.1: This control states that the cloud provider must separate the customers virtual environment from other customers or outside parties.
  - CLD.9.5.2: This control states that the customer and the cloud provider both must ensure the virtual machines are hardened.
  - CLD.12.1.5: It is solely the customer's responsibility to define and manage administrative operations.
  - CLD.12.4.5: The cloud providers capabilities must enable the customer to monitor their own cloud environment.
  - CLD.13.1.4: The virtual network environment must be configured so that it least meets the security policies of the physical environment.

ISO 27018

ISO 27018 is closely related to ISO 27017. ISO 27018 defines privacy requirements in a cloud environment. Particularly how the customer and cloud provider must protect personally identifiable information (PII)

FedRAMP

The Federal Risk and Authorization Management Program (FedRAMP) is a government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud products and services. Third-party assessment organizations (3PAOs) play a critical role in the FedRAMP security assessment process, as they are the independent assessment organizations that verify cloud providers' security implementations and provide the overall risk posture of a cloud environment for a security authorization decision

https://www.fedramp.gov/

NSA Guidance

- The NSA offers guidance on cloud security https://media.defense.gov/2020/Jan/22/2002237484/-1/-1/0/CSI-MITIGATING-CLOUD-VULNERABILITIES_20200121.PDF
- While not a base component of cloud architectures, encryption and key management (KM) form a critical aspect of protecting information in the cloud.
- While CSPs are generally responsible for detecting threats to the underlying cloud platform, customers bear the responsibility of detecting threats to their own cloud resources.
- Incident Response: CSPs are uniquely positioned to respond to incidents internal to the cloud infrastructure and bear responsibility for doing so. Incidents internal to customer cloud environments are generally the customer's responsibility, but CSPs may provide support to incident response teams.
- Patching/Updating: CSPs are responsible for ensuring that their cloud offerings are secure and rapidly patch software within their purview but usually do not patch software managed by the customer (e.g., operating systems in IaaS offerings). Because of this, customers should vigilantly deploy patches to mitigate software vulnerabilities in the cloud. In some cases CSPs offer managed solutions in which they perform operating system patching as well.

NIST Special Publication 800-144, Guidelines on Security

NIST Special Publication 800-144, **Guidelines on Security and Privacy in Public Cloud Computing**, December 2011
NIST Special Publication 800-145, **NIST Definition of Cloud Computing**, September 2011
NIST Special Publication 800-146, **Cloud Computing Synopsis and Recommendations**, May 2012
NIST Special Publication 500-291, **NIST Cloud Computing Standards Roadmap**, July 2011
NIST Special Publication 500-292, **NIST Cloud Computing Reference Architecture**, September 2011
NIST Special Publication 500-299, **NIST Cloud Computing Security Reference Architecture (Draft)**

NIST Special Publication 800-144, Guidelines on Security

This standard emphasizes the importance of the service level agreement (section 3.1).

NIST 800-144 discusses governance as a security issue (section 4.1)

Virtual Network Protection is also emphasized (section 4.4)
Authentication is addressed (section 4.5). Many cloud providers are using SAML (We will discuss SAML in some depth later in this workshop)

NIST Special Publication 800-144, Guidelines on Security and Privacy in Public Cloud

- SLA should cover:

- Personnel requirements, including clearances, roles, and responsibilities
- Regulatory requirements
- Service availability
- Problem reporting, review, and resolution
- Information handling and disclosure agreements and procedures
- Physical and logical access controls
- Network access control, connectivity, and filtering
- Data protection
- System configuration and patch management
- Backup and recovery
- Data retention and sanitization
- Security and vulnerability scanning
- Risk management
- Incident reporting, handling, and response
- Continuity of operations
- Resource management
- Certification and accreditation
- Assurance levels
- Independent auditing of services

NIST Special Publication 800-144, Guidelines on Security and Privacy in Public Cloud Computing

Recommendations

Table 1: Security and Privacy Issues and Recommendations

Areas	Recommendations
Governance	Extend organizational practices pertaining to the policies, procedures, and standards used for application development and service provisioning in the cloud, as well as the design, implementation, testing, use, and monitoring of deployed or engaged services. Put in place audit mechanisms and tools to ensure organizational practices are followed throughout the system lifecycle.
Compliance	Understand the various types of laws and regulations that impose security and privacy obligations on the organization and potentially impact cloud computing initiatives, particularly those involving data location, privacy and security controls, records management, and electronic discovery requirements. Review and assess the cloud provider's offerings with respect to the organizational requirements to be met and ensure that the contract terms adequately meet the requirements. Ensure that the cloud provider's electronic discovery capabilities and processes do not compromise the privacy or security of data and applications.
Trust	Ensure that service arrangements have sufficient means to allow visibility into the security and privacy controls and processes employed by the cloud provider, and their performance over time. Establish clear, exclusive ownership rights over data. Institute a risk management program that is flexible enough to adapt to the constantly evolving and shifting risk landscape for the lifecycle of the system. Continuously monitor the security state of the information system to support on-going risk management decisions.

NIST Special Publication 800-144, Guidelines on Security and Privacy in Public Cloud Computing

Recommendations

Areas	Recommendations
Architecture	Understand the underlying technologies that the cloud provider uses to provision services, including the implications that the technical controls involved have on the security and privacy of the system, over the full system lifecycle and across all system components.
Identity and Access Management	Ensure that adequate safeguards are in place to secure authentication, authorization, and other identity and access management functions, and are suitable for the organization.
Software Isolation	Understand virtualization and other logical isolation techniques that the cloud provider employs in its multi-tenant software architecture, and assess the risks involved for the organization.
Data Protection	Evaluate the suitability of the cloud provider's data management solutions for the organizational data concerned and the ability to control access to data, to secure data while at rest, in transit, and in use, and to sanitize data. Take into consideration the risk of collating organizational data with that of other organizations whose threat profiles are high or whose data collectively represent significant concentrated value. Fully understand and weigh the risks involved in cryptographic key management with the facilities available in the cloud environment and the processes established by the cloud provider.
Availability	Understand the contract provisions and procedures for availability, data backup and recovery, and disaster recovery, and ensure that they meet the organization's continuity and contingency planning requirements. Ensure that during an intermediate or prolonged disruption or a serious disaster, critical operations can be immediately resumed, and that all operations can be eventually reinstituted in a timely and organized manner.
Incident Response	Understand the contract provisions and procedures for incident response and ensure that they meet the requirements of the organization. Ensure that the cloud provider has a transparent response process in place and sufficient mechanisms to share information during and after an incident. Ensure that the organization can respond to incidents in a coordinated fashion with the cloud provider in accordance with their respective roles and responsibilities for the computing environment.

NIST Special Publication 800-144, Guidelines on Security and Privacy in Public Cloud Computing

The standard lists specific concerns:

Inadequate Policies and Practices.

Weak Confidentiality and Integrity Sureties.

Weak Availability Sureties.

Principal-Agent Problem

Attenuation of Expertise.

ISO Standards

- ISO/IEC 42001 AI management systems
- ISO/IEC 23894 AI Guidance on risk management
- ISO/IEC 23053 Framework for AI Systems Using ML
- ISO/IEC DIS 12792 Information technology — Artificial intelligence — Transparency taxonomy of AI systems

- https://webstore.ansi.org/industry/software/artificial-intelligence

NIST Standards

NIST SP 800-218A Secure Software Development Practices for Generative AI and Dual-Use Foundation Models: An SSDF Community Profile
https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-218A.ipd.pdf

NIST Special Publication 1270: Toward a Standard for Identifying and Managing Bias in Artificial Intelligence
https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf

NIST AI 100-3: The Language of Trustworthy AI
https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-3.pdf

NIST AI 600-1: AI RMF Generative AI Profile
https://airc.nist.gov/docs/NIST.AI.600-1.GenAI-Profile.ipd.pdf

NIST AI 100-1 Artificial Intelligence Risk Management Framework (AI RMF 1.0)

- "The AI RMF is intended to be practical, to adapt to the AI landscape as AI technologies continue to develop, and to be operationalized by organizations in varying degrees and capacities so society can benefit from AI while also being protected from its potential harms"

- Part 1 discusses how organizations can frame the risks related to AI and describes the intended audience
- Part 2 comprises the "Core" of the Framework.

- https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

MITRE ATLAS Framework

Adversarial Threat Landscape for AI Systems (ATLAS)

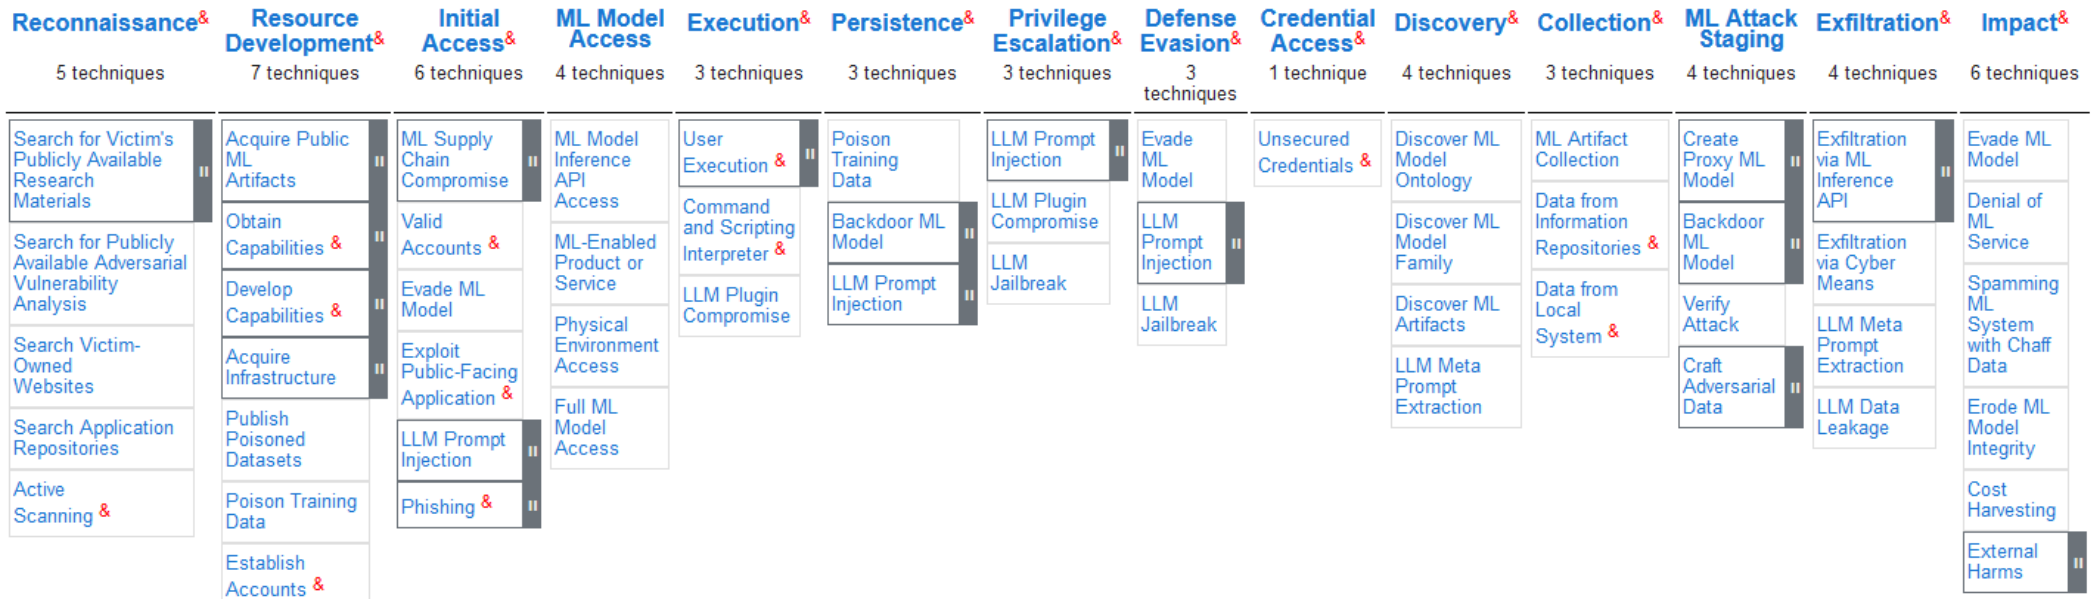

https://atlas.mitre.org/matrices/ATLAS

NSA AI Security Guidelines

- Secure Design
  - Model the threats to your system
  - Design your system for security as well as functionality and performance
  - Consider security benefits and trade-offs when selecting your AI model
- Secure Development
  - Secure your supply chain
  - Identify, track and protect your assets
  - Document your data, models, and prompts
- Secure Deployment
  - Secure your infrastructure
  - Protect your model continuously.
  - Develop incident management procedures

- https://media.defense.gov/2023/Nov/27/2003346994/-1/-1/0/GUIDELINES-FOR-SECURE-AI-SYSTEM-DEVELOPMENT.PDF