


---

# Lesson 11: Cloud and Zero Trust





---


# What is Cloud Computing?

- Cloud Computing is a general term used to describe a new class of network based computing that takes place over the Internet, a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform). Using the Internet for communication and transport provides hardware, software and networking services to clients
- These platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API (Applications Programming Interface).

---

# What is cloud computing?

- Cloud computing means storing and accessing data and programs over the Internet instead of your computer's hard drive. – PC Mag
- Cloud computing is the delivery of computing services—servers, storage, databases, networking, software, analytics, and more—over the Internet (“the cloud”). – Microsoft
- Cloud computing is the on-demand delivery of compute power, database storage, applications, and other IT resources through a cloud services platform via the internet with pay-as-you-go pricing. –Amazon



---

# 2021 AWS Cloud Security Report

- 95% of cybersecurity professionals confirm they are extremely to moderately concerned about public cloud security.
- #1 concern remains Misconfiguration of the cloud platform (71%). Exfiltration of sensitive data came in second (59%). Insecure APIs (54%) rounded out the top 3 concerns.
- 58% use periodic vulnerability and compliance reports as the primary method of communication with system owners about security and compliance issues needing remediation. This is followed by automatically opened tickets (47%) using tools such as Jira, ServiceNow, etc.
- >40% of organizations embrace hybrid cloud (44%) and multi-cloud deployments (43%) for planned redundancy because of commitments to legacy applications in traditional data centers. Single cloud deployments (11%) continue to diminish in importance.
- 90% of organizations use more than two cloud providers.
- 66% of organizations prioritize cost-effectiveness as a leading criterion when selecting a cloud security provider, followed by scalability (52%), ease of deployment (51%), and tools that can be deployed with automation (48%).



# NIST

The National Institute of Standards and Technology (NIST) defines cloud computing as “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.”

---

# Cloud Characteristics

- In all clouds, someone else is providing the physical machines
- You aren't concerned about power, bandwidth, maintenance, physical security, or (sometimes) scaling
- You only pay for what you use
- Cloud computing uses servers distributed geographically. In some cases, the servers are in other countries. This brings the benefit of fault tolerance, but has some security concerns:
- Privacy laws in different regions
- Ensuring that a customer's data is segregated from other customers data is the primary data protection issue

---

# Essential Cloud Characteristics



**ON-DEMAND SELF-SERVICE:** CUSTOMERS CAN PROVISION COMPUTING CAPABILITIES.




**BROAD NETWORK ACCESS:** RESOURCES ARE AVAILABLE OVER THE NETWORK THROUGH STANDARD MECHANISMS.



**RESOURCE POOLING:** THE PROVIDER'S COMPUTING RESOURCES ARE POOLED TO SERVE MULTIPLE CONSUMERS USING A MULTI-TENANT MODEL.



**RAPID ELASTICITY:** CAPABILITIES CAN BE RAPIDLY AND ELASTICALLY PROVISIONED, PREFERABLY AUTOMATICALLY.



**MEASURED SERVICE:** RESOURCE-USAGE IS MONITORED AND AUTOMATICALLY CONTROLLED AND OPTIMIZED. THE ORGANIZATION PROVIDES TRANSPARENCY FOR BOTH ITSELF AND THE CUSTOMER OF THE UTILIZED SERVICE.

---

# Cloud Types

- **Private cloud**

The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premises

- **Public cloud**

Large cloud infrastructure is made available to the general public and is owned by an organization selling cloud services

- **Community cloud**

The (usually private) cloud infrastructure is shared by multiple organizations that share the same goal.

- **Hybrid cloud**

The Hybrid cloud is a composition of two or more cloud deployment models (private, community or public) that are bound together to enable data and application portability. Most common applications consist of cloud bursting (handling temporary peaks) and load-balancing between clouds.

---

# Basic Cloud Concepts



In all clouds, someone else is providing the physical machines



You aren't concerned about power, bandwidth, maintenance, physical security, or (sometimes) scaling



You only pay for what you use

---

---

# Multi-cloud



Multiple different cloud vendors are used heterogeneously. This mitigates dependency on a single vendor. Cloud assets (applications, virtual servers, etc.) are hosted across multiple different public clouds. One can also include private clouds in the architecture.



Poly cloud is similar, but in this case the different public clouds are being utilized not for flexibility and redundancy, but rather for specific services each provider offers.

---



---

# HPC Cloud

- HPC Cloud
- This is the use of cloud services for high performance computing. Such HPC applications would normally require clusters of computers or a supercomputer. There are several companies including Amazon Web Services that offer HPC cloud computing.



---

# Virtualization

- Started in 1967 with the IBM CP-40
- Virtual machine (VM) software is a program that emulates a physical machine
- A VM should behave as if it were an independent physical machine. You see this on desktops with Oracle Virtual Box and VMWare workstations

---

# Virtual Systems

Virtual Machine

Software as a Service (SaaS)

Platform as a Service (PaaS)

Infrastructure as a Service (IaaS)

Desktop as a Service (DaaS)

Information Technology Management as a Service (ITMaaS).

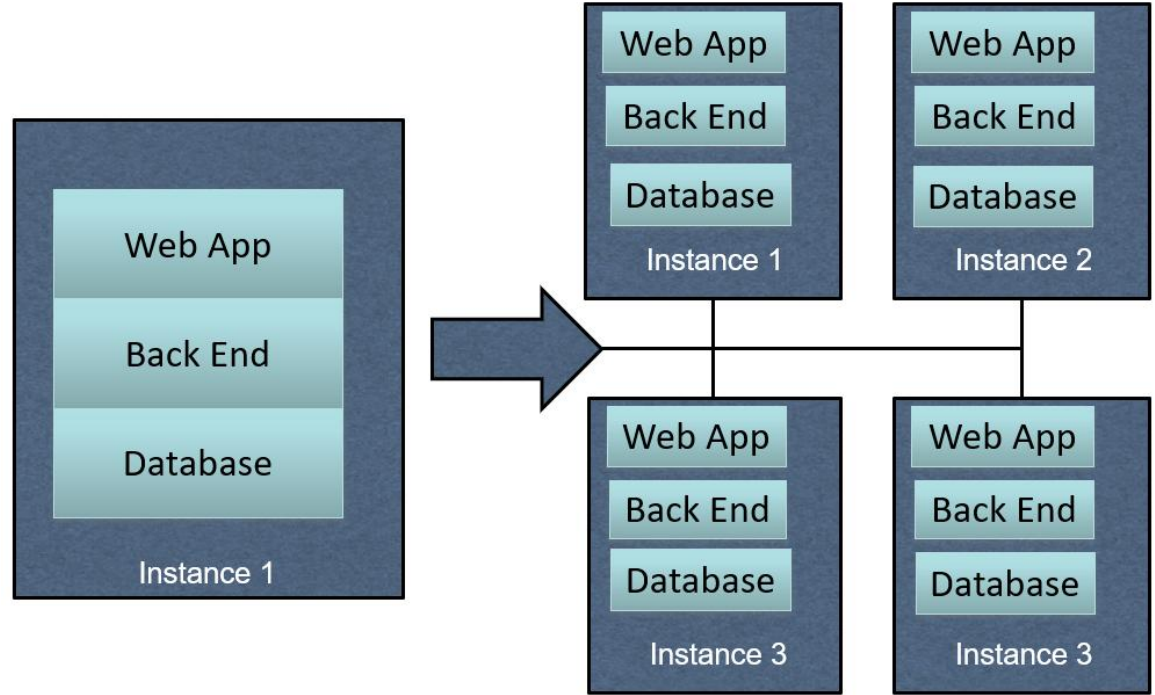
---



# IAAS

- You can't tell if you are on a cloud machine or not
- From the perspective of the software (or an administrator), a cloud machine is identical to a physical machine. The process is:
  - Determine your operating system
  - Determine how much computing you need
  - Find an instance in your cloud provider library of machines
  - Start that instance
- Automatically scale up / down machines as needed

# IAAS



---

# PaaS



CPU resources scale up or down as needed



No need to spin up new machines, manage load balancing, etc.



There are several types of PaaS, including public, private and hybrid.



There are variations such as Communications Platform as a Service (CPaaS) and Mobile Platform as a Service (mPaaS).

---

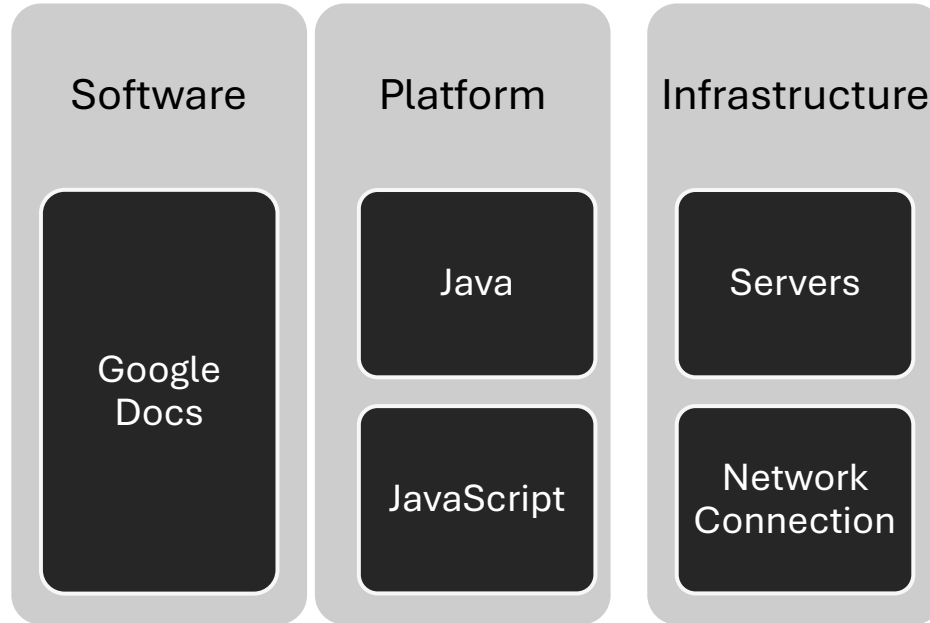
---

# SaaS

- Basically, renting an application instead of setting it up on your own server.
- Usually, users access SaaS apps via some thin client, often web browsers. There are a wide range of applications available in this fashion. The applications are provided by application service providers (ASP). There are subsets of SaaS such as DbaaS (database as a service)
- There are two main variations of SaaS:
  - Vertical SaaS
    - Software which is for a specific industry such as healthcare, finance, etc.
  - Horizontal SaaS
    - The products which focus on a particularly category of software such as software development, sales, etc., but is not for a particular industry
    - OpenSaaS refers to software as a service (SaaS) based on open source code.

# Example SaaS: Google

## Docs



---

# Variations

- There are numerous variations such as
- Security as a service (SECaaS or SaaS)
- Knowledge as a service (KaaS)
- data as a service (DaaS)
- Mobile backend as a service (MBaaS)
- Artificial intelligence as a service (AlaaS)
- Content as a Service (CaaS)
- These are typically specialized variation of PaaS, IaaS , SaaS.

---

# Distributed Systems Issues

- Synchronization: multiple clocks (difficult to agree on exact time)
- Concurrency: multiple simultaneous accesses potentially conflicting.
- Failures: high probability of failures (too many components).  
Complex failure modes (single, multiple simultaneous, network partition, ...)
- Consensus: difficult to reach consensus (odds includes failures, lack of synchronization, ...)

---

# Fog Computing

An architecture that uses edge devices for processing. Sometimes called fogging or fog networking. This approach is seen in cloud systems. There are two aspects to Fog computing. The control plane and the data plane. Fog networking is often used in IoT.

National Institute of Standards and Technology in March 2018 released NIST Special Publication 500-325, Fog Computing Conceptual Model, that defines fog computing as a horizontal, physical or virtual resource paradigm that resides between smart end-devices and traditional cloud computing or data center

---

---

# Future Trends

Cloud robotics: integrating robotics with cloud storage and processing. This will allow the robot itself to have less processing power. There is interest in medical, industrial, and domestic robots integrating with the cloud.

Improving Edge Computing: In January 2017, RECAP, an EU-funded project, was launched to advance cloud and edge computing technology.

Robotics as a service: The idea of integration robotics, the web, and cloud computer so that robots can be utilized as a service. Sometimes termed RaaS.

---

# Cloud Threats On the Rise

A 2019 report found that cloud infrastructure attacks are on the rise



An unauthenticated command execution vulnerability in Apache Hadoop through ResourceManager REST API



A Redis remote command execution bug



CVE-2016-3088, an ActiveMQ arbitrary file execution flaw.



<https://www.helpnetsecurity.com/2019/01/25/cloud-infrastructure-attacks/>

---

---

# Cloud Vulnerabilities

- CVE-2020-8960: Western Digital mycloud.com before Web Version 2.2.0-134 allows XSS.
- CVE-2020-3154: A vulnerability in the web UI of Cisco Cloud Web Security (CWS) could allow an authenticated, remote attacker to execute arbitrary SQL queries. The vulnerability exists because the web-based management interface improperly validates SQL values. An authenticated attacker could exploit this vulnerability sending malicious requests to the affected device. An exploit could allow the attacker to modify values on or return values from the underlying database.

---

# 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)

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

---

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 and Privacy in Public Cloud Computing



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 Computing

## 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

Table 1: Security and Privacy Issues and Recommendations

Areas	Recommendations
Governance	<p>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.</p> <p>Put in place audit mechanisms and tools to ensure organizational practices are followed throughout the system lifecycle.</p>
Compliance	<p>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.</p> <p>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.</p> <p>Ensure that the cloud provider's electronic discovery capabilities and processes do not compromise the privacy or security of data and applications.</p>
Trust	<p>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.</p> <p>Establish clear, exclusive ownership rights over data.</p> <p>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.</p> <p>Continuously monitor the security state of the information system to support on-going risk management decisions.</p>

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

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	<p>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.</p> <p>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.</p> <p>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.</p>
Availability	<p>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.</p> <p>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 reinstated in a timely and organized manner.</p>
Incident Response	<p>Understand the contract provisions and procedures for incident response and ensure that they meet the requirements of the organization.</p> <p>Ensure that the cloud provider has a transparent response process in place and sufficient mechanisms to share information during and after an incident.</p> <p>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.</p>

## **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 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)



---

# 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](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.



# NSA Guidance



- - [https://media.defense.gov/2020/Jan/22/2002237484/-1/-1/0/CSI-MITIGATING-CLOUD-VULNERABILITIES\\_20200121.PDF](https://media.defense.gov/2020/Jan/22/2002237484/-1/-1/0/CSI-MITIGATING-CLOUD-VULNERABILITIES_20200121.PDF)

## Cloud Threat Actors

Threat actors may target the same types of weaknesses in both cloud and traditional system architectures. This section focuses on cloud-specific activities, but administrators should be aware that traditional tactics still apply. For example, an unpatched web application in the cloud bears similar risk of compromise as one served from an on-premises network. The following threat actors are relevant to cloud computing:

### *Malicious CSP Administrators*

- Leverage privileged credentials or position to access, modify, or destroy information stored on the cloud platform;
- Leverage privileged credentials or position to modify the cloud platform in order to gain access to networks connected to or consuming cloud resources;

### *Malicious Customer Cloud Administrators*

- Leverage privileged credentials to access, modify, or destroy information stored on the cloud platform;

### *Cyber Criminals and/or Nation State-Sponsored Actors*

- Leverage a weakness in the cloud architecture or configuration to obtain sensitive data or consume cloud resources at the victim's expense;
- Exploit weak cloud-based authentication mechanisms to obtain user credentials (e.g., password spray attacks);
- Leverage compromised credentials or incorrect access privileges to gain access to cloud resources;
- Gain privileged access to the cloud environment to compromise tenant resources;
- Leverage the trust relationship between an organization's networks and cloud resources to pivot from clouds into protected networks or vice versa;

### *Untrained or Neglectful Customer Cloud Administrators*

- Expose sensitive data or cloud resources unintentionally.

---

# Cloudbleed

- Cloudbleed
- This was a bug in Cloudflare's reverse proxy servers that cause their edge servers to send back confidential information from their memory buffer. It was discovered in February 2017.
- Basically it was a buffer overrun that revealed data that should have been confidential. Data included HTTP cookies and authentication tokens.

---

# OWASP Cloud-Native Application Security Top 10

- 
- CNAS-1: Insecure cloud, container or orchestration configuration
  - CNAS-2: Injection flaws (app layer, cloud events, cloud services)
  - CNAS-3: Improper authentication & authorization
  - CNAS-4: CI/CD pipeline & software supply chain flaws
  - CNAS-4: CI/CD pipeline & software supply chain flaws
  - CNAS-6: Over-permissive or insecure network policies
  - CNAS-7: Using components with known vulnerabilities
  - CNAS-8: Improper assets management
  - CNAS-9: Inadequate 'compute' resource quota limits
  - CNAS-10: Ineffective logging & monitoring (e.g. runtime activity)

---

# OWASP Top 10 Cloud Security Risks (Draft)

- 
- R1. Accountability & Data Risk
  - R2. User Identity Federation
  - R3. Legal & Regulatory Compliance
  - R4. Business Continuity & Resiliency
  - R5. User Privacy & Secondary Usage of Data
  - R6. Service & Data Integration
  - R7. Multi-tenancy & Physical Security
  - R8. Incidence Analysis & Forensics
  - R9. Infrastructure Security
  - R10. Non-production Environment Exposure
- 
- [https://owasp.org/www-pdf-archive/OWASP\\_Cloud\\_Top\\_10.pdf](https://owasp.org/www-pdf-archive/OWASP_Cloud_Top_10.pdf)

---

# Checkpoint Cloud Security

---

Misconfiguration

---

Unauthorized Access

---

Insecure Interfaces/APIs

---

Hijacking of Accounts

---

Lack of Visibility

---

External Sharing of Data

---

Insider threats

---

Cyberattacks

---

Denial of Service Attacks

---

<https://www.checkpoint.com/cyber-hub/cloud-security/what-is-cloud-security/top-cloud-security-issues-threats-and-concerns/>

---

---

# Critical Security Areas in Cloud Computing (CSA)



## Governing in the Cloud

Governance and Enterprise Risk Management

Legal and Electronic Discovery

Compliance and Audit

Information Lifecycle Management

Portability and Interoperability



## Operating in the Cloud

Traditional Security, Business Continuity, and Disaster Recovery

Data Center Operations

Incident Response, Notification, and Remediation

Application Security

Encryption and Key Management

Identity and Access Management

Virtualization

# Cloud Security Alliance - Guidance

The Cloud Security Alliance's 13 Critical Areas Of Focus for Cloud:

1. Architecture & Framework	
<i>Governing the Cloud</i>	<i>Operating the Cloud</i>
2. Governance & Risk Mgmt.	8. Traditional BCM, DR
3. Legal & Electronic Discovery	9. Datacenter Operations
5. Compliance & Audit	10. Incident Response
6. Information Lifecycle Mgmt.	11. Application Security
7. Portability & Interoperability	12. Encryption & Key Mgmt.
	13. Identity & Access Mgmt.

---

# Top 10 Customer Issues Eroding Cloud Confidence (from CSA)

Government regulations keeping pace with the market
Exit strategies
International data privacy
Legal issues
Contract lock in
Data ownership and custodian responsibilities
Longevity of suppliers
Integration of cloud with internal systems
Credibility of suppliers
Testing and assurance

---

# Risks

Risk 1: Resource Exhaustion

Risk 2: Customer Isolation Failure

Risk 3: Management Interface Compromise

Risk 4: Interception of Data in Transmission

Risk 5: Data leakage on Upload/Download, Intra-cloud

---

# Risks



- 
- Risk 6: Insecure or Ineffective Deletion of Data
  - Risk 7: Distributed Denial of Service (DDoS)
  - Risk 8: Economic Denial of Service
  - Risk 9: Loss or Compromise of Encryption Keys
  - Risk 10: Malicious Probes or Scans



# Risks

- Risk 11: Compromise of Service Engine/Hypervisor\*
- Risk 12: Conflicts between customer hardening procedures and cloud environment
- Risk 13: Subpoena and E-Discovery\*
- Risk 14: Risk from Changes of Jurisdiction\*
- Risk 15: Licensing Risks\*

---

# Risks

Risk 16: Network Failure

Risk 17: Networking Management

Risk 18: Modification of Network Traffic

Risk 19: Privilege Escalation\*

Risk 20: Social Engineering Attacks

---

---

# Risks

Risk 21: Loss or Compromise of Operation Logs

Risk 22: Loss or compromise of Security Logs

Risk 23: Backups Lost or Stolen


Risk 23: Unauthorized Access to Premises, Including Physical Access to Machines and Other Facilities

Risk 25: Theft of Computer Equipment.\*

---

# Minimize Loss of Control: Monitoring

- Cloud consumer needs situational awareness for critical applications
  - When underlying components fail, what is the effect of the failure to the mission logic
  - What recovery measures can be taken (by provider and consumer)
- Requires an application-specific run-time monitoring and management tool for the consumer
  - The cloud consumer and cloud provider have different views of the system
  - Enable both the provider and tenants to monitor the components in the cloud that are under their control





---

# Minimize Multi-tenancy

- Can't force the provider to accept less tenants
- Can try to increase isolation between tenants
  - Strong isolation techniques
  - VM Side channel attacks
  - QoS requirements need to be met
  - Policy specification
- Use SLAs to enforce trusted behavior

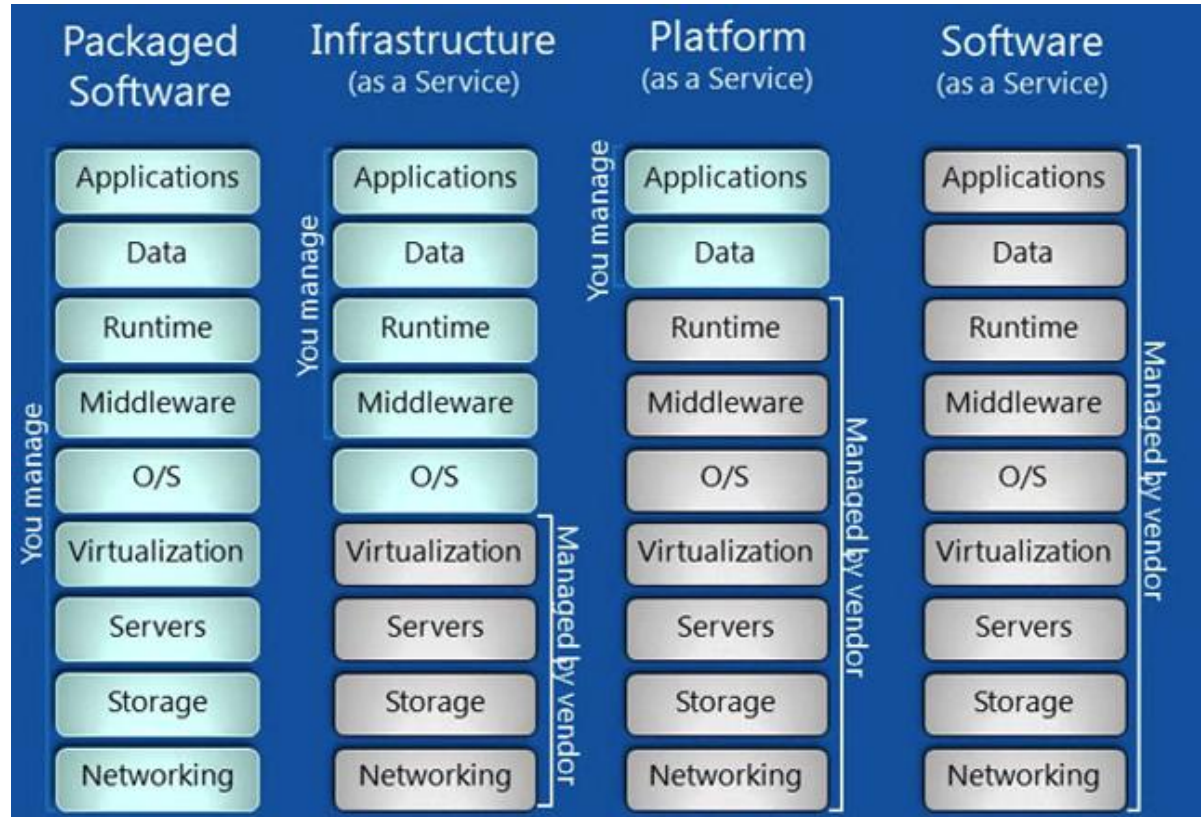
---

# Cloud Security Challenges

Data dispersal and international privacy laws

- EU Data Protection Directive and U.S. Safe Harbor program
- Exposure of data to foreign government and data subpoenas
- Proprietary cloud vendor implementations can't be examined
- Loss of physical control
- Possibility for massive outages
- Encryption needs for cloud computing
  - Encrypting access to the cloud resource control interface
  - Encrypting administrative access to virtual instances
  - Encrypting access to applications
  - Encrypting application data at rest

# Services Delivered in each Model



---

# Part II Zero Trust



---

# What is Zero Trust?

- Zero Trust is a security framework requiring all users, whether in or outside the organization's network, to be authenticated, authorized, and continuously validated for security configuration and posture before being granted or keeping access to applications and data. **Zero Trust assumes that there is no traditional network edge**; networks can be local, in the cloud, or a combination or hybrid with resources anywhere as well as workers in any location.
- Zero Trust is a framework for securing infrastructure and data for today's modern digital transformation. It uniquely addresses the modern challenges of today's business, including securing remote workers, hybrid cloud environments, and ransomware threats. While many vendors have tried to create their own definitions of Zero Trust, there are a number of standards from recognized organizations that can help you align Zero Trust with your organization
- -<https://www.crowdstrike.com/cybersecurity-101/zero-trust-security/>

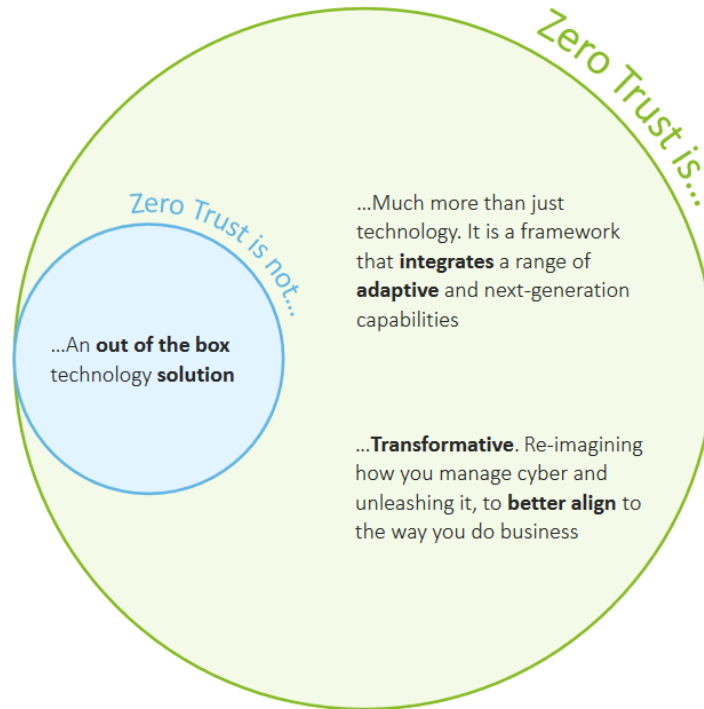
---

# What is Zero Trust?

- Many sources trace the term “zero trust” to the doctoral thesis of Stephen Marsh, published in 1994.
- The term ‘zero trust model’ is often attributed to Forrester Research analyst John Kindervag in 2009. In fact, many sources ignore Marsh’s dissertation and instead attribute Zero Trust entirely to Forrester Research. However, you can find Marsh’s dissertation at <https://dspace.stir.ac.uk/handle/1893/2010#.YvActBzMI9E> which has since been cited over 2200 times.

---

**Zero Trust – Deloitte**  
**<https://www2.deloitte.com/content/dam/Deloitte/de/Documents/risk/deloitte-cyber-zero-trust.pdf>**



# Zero Trust – Oracle <https://www.oracle.com/security/what-is-zero-trust/>



All data sources and computing services are considered resources.



All communication is secure regardless of network location; network location does not imply trust.



Access to individual enterprise resources is granted on a per-connection basis; trust in the requester is evaluated before the access is granted.



Access to resources is determined by policy, including the observable state of user identity and the requesting system, and may include other behavioral attributes.



The enterprise ensures all owned and associated systems are in the most secure state possible and monitors systems to ensure that they remain in the most secure state possible.




User authentication is dynamic and strictly enforced before access is allowed; this is a constant cycle of access, scanning and assessing threats, adapting, and continually authenticating.

# Zero Trust Principles




## Verify Explicitly

Always authenticate and authorize based on all available data points, including user identity, location, device health, data classification, and anomalies.



## Least Privilege

Minimize user access with Just-In-Time and Just-Enough Access (JIT/JEA), risk-based adaptive policies, and data protection which protects data and productivity.




## Assume Breach

Minimize scope of breach damage and prevent lateral movement by segmenting access via network, user, devices and application awareness. Verify all sessions are encrypted end to end. Use analytics to get visibility and drive threat detection.

# Zero Trust Access Control Strategy

Never Trust. Always verify.



## Signal

*to make an informed decision*

### Device Risk

- Device Management
- Threat Detection
- and more...

### User Risk

- Multi-factor Authentication
- Behavior Analytics
- and more...

## Decision

*based on organization's policy*

**Apply to inbound requests**

**Re-evaluate during session**

## Enforcement

*of policy across resources*



**Modern Applications**

**SaaS Applications**

**Legacy Applications**

**And more...**

# Zero Trust Access Control Paradigm

	 Network	 Identity
Control Plane	Apply Zero Trust Policy to <i>network connections</i>	Apply Zero Trust Policy to <i>access requests</i>
Industry Proponents	Network Security Vendors	Identity Vendors
Overall Effect	<b>Micro segmentation</b> enhances existing network perimeter by shrinking "trusted network" to each server / IP address.	<b>Dual Perimeter</b> – Adds an identity perimeter where "inside" is defined by authentication and authorization. <i>Coexists with network perimeter</i>
Applicability/Scope	<b>Limited to networks</b> controlled by customer. Doesn't protect modern SaaS and PaaS assets. <i>Micro segmentation approach varies by vendor</i>	<b>Applies to all assets</b> – <ul style="list-style-type: none"> <li>• Natively protects modern cloud assets</li> <li>• Protects legacy intranet assets via proxy</li> </ul>
Differentiation	<b>Scope of assets</b> where zero trust is enforced  <b>Threat Intelligence</b> signal Integration	Integration of <b>Behavior Analytics (UEBA)</b> risk signal  Use of <b>ML</b> across large datasets decisions  <div style="border: 1px solid black; border-radius: 15px; padding: 5px; background-color: #008080; color: white;">             Microsoft focuses on protecting modern and legacy assets as well as integration of ML, UEBA, and massive diverse threat intelligence           </div>
Common Components	Evaluate trust signals for Devices & User Identities with per application policy	

# Microsoft's Recommended Zero Trust Priorities



1. **Align segmentation strategy & teams** by unifying network, identity, app, etc. into a single enterprise segmentation strategy (as you migrate to Azure)



2. **Build identity-based perimeter** to protect modern *and* legacy enterprise assets



3. **Refine network perimeter** using micro segmentation (if required for residual risk)

---

# 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

# NIST

4. 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

The enterprise monitors and measures the integrity and security posture of all owned and associated assets. No asset is inherently trusted. The enterprise evaluates the security posture of the asset when evaluating a resource request. An enterprise implementing a ZTA should establish a continuous diagnostics and mitigation (CDM) or similar system to monitor the state of devices and applications and should apply patches/fixes as needed. Assets that are discovered to be subverted, have known vulnerabilities, and/or are not managed by the enterprise may be treated differently (including denial of all connections to enterprise resources) than devices owned by or associated with the enterprise that are deemed to be in their most secure state. This may also apply to associated devices (e.g., personally owned devices) that may be allowed to access some resources but not others. This, too, requires a robust monitoring and reporting system in place to provide actionable data about the current state of enterprise resources.

6. All resource authentication and authorization are dynamic and strictly enforced before access is allowed. This is a constant cycle of obtaining access, scanning and assessing threats, adapting, and continually reevaluating trust in ongoing communication. An enterprise implementing a ZTA would be expected to have Identity, Credential, and Access Management (ICAM) and asset management systems in place. This includes the use of multifactor authentication (MFA) for access to some or all enterprise resources. Continual monitoring with possible reauthentication and reauthorization occurs throughout user transactions, as defined and enforced by policy (e.g., time-based, new resource requested, resource modification, anomalous subject activity detected) that strives to achieve a balance of security, availability, usability, and cost-efficiency.

7. The enterprise collects as much information as possible about the current state of assets, network infrastructure and communications and uses it to improve its security posture. An enterprise should collect data about asset security posture, network traffic and access requests, process that data, and use any insight gained to improve policy creation and enforcement. This data can also be used to provide context for access requests from subjects (see Section 3.3.1).



---

# NIST SP 800-207

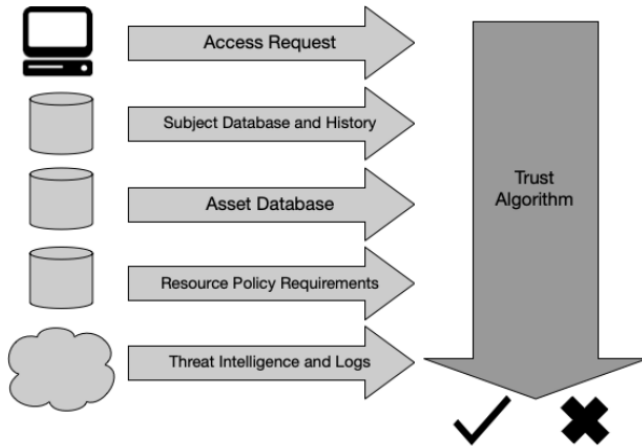


Figure 7: Trust Algorithm Input

- 3.1.1 ZTA Using Enhanced Identity Governance
- 3.1.2 ZTA Using Micro-Segmentation
- 3.2.4 Device Application Sandboxing
- 4.3 Enterprise with Contracted Services and/or Nonemployee Access

---

# 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.

---

# NIST SP 800-205 Network Requirements for Zero Trust

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).

---

# NIST SP 800-205 Network Requirements for Zero Trust

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.

---

# 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

- 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.

•[https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model\\_Draft.pdf](https://www.cisa.gov/sites/default/files/publications/CISA%20Zero%20Trust%20Maturity%20Model_Draft.pdf)

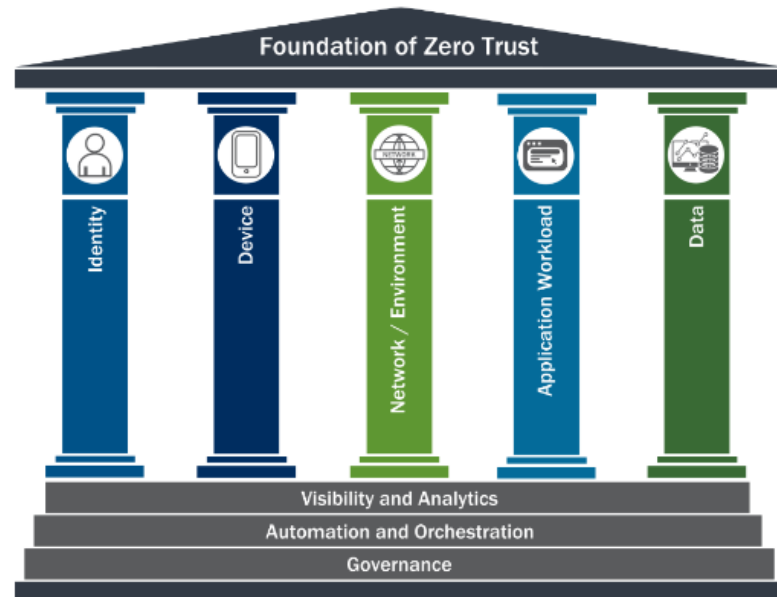


Figure 1: Foundation of Zero Trust<sup>7</sup>

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

	Identity	Device	Network / Environment	Application Workload	Data
Traditional	<ul style="list-style-type: none"> <li>• Password or multifactor authentication (MFA)</li> <li>• Limited risk assessment</li> </ul>	<ul style="list-style-type: none"> <li>• Limited visibility into compliance</li> <li>• Simple inventory</li> </ul>	<ul style="list-style-type: none"> <li>• Large macro-segmentation</li> <li>• Minimal internal or external traffic encryption</li> </ul>	<ul style="list-style-type: none"> <li>• Access based on local authorization</li> <li>• Minimal integration with workflow</li> <li>• Some cloud accessibility</li> </ul>	<ul style="list-style-type: none"> <li>• Not well inventoried</li> <li>• Static control</li> <li>• Unencrypted</li> </ul>
Advanced	<ul style="list-style-type: none"> <li>• MFA</li> <li>• Some identity federation with cloud and on-premises systems</li> </ul>	<ul style="list-style-type: none"> <li>• Compliance enforcement employed</li> <li>• Data access depends on device posture on first access</li> </ul>	<ul style="list-style-type: none"> <li>• Defined by ingress/egress micro-perimeters</li> <li>• Basic analytics</li> </ul>	<ul style="list-style-type: none"> <li>• Access based on centralized authentication</li> <li>• Basic integration into application workflow</li> </ul>	<ul style="list-style-type: none"> <li>• Least privilege controls</li> <li>• Data stored in cloud or remote environments are encrypted at rest</li> </ul>
Optimal	<ul style="list-style-type: none"> <li>• Continuous validation</li> <li>• Real time machine learning analysis</li> </ul>	<ul style="list-style-type: none"> <li>• Constant device security monitor and validation</li> <li>• Data access depends on real-time risk analytics</li> </ul>	<ul style="list-style-type: none"> <li>• Fully distributed ingress/egress micro-perimeters</li> <li>• Machine learning-based threat protection</li> <li>• All traffic is encrypted</li> </ul>	<ul style="list-style-type: none"> <li>• Access is authorized continuously</li> <li>• Strong integration into application workflow</li> </ul>	<ul style="list-style-type: none"> <li>• Dynamic support</li> <li>• All data is encrypted</li> </ul>

Figure 2: High-Level Zero Trust Maturity Model

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

Function	Traditional	Advanced	Optimal
<b>Authentication</b>	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.
<b>Identity Stores</b>	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.
<b>Risk Assessment</b>	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.
<b>Visibility and Analytics Capability</b>	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).
<b>Automation and Orchestration Capability</b>	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.
<b>Governance Capability</b>	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.

---

# **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 Architecture

[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)

---

- **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.

---

# DoD Zero Trust Reference Architecture

[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)

- 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.



---

# DoD Zero Trust Reference Architecture

[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)

## 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.
-

# DoD Zero Trust Reference Architecture

[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)

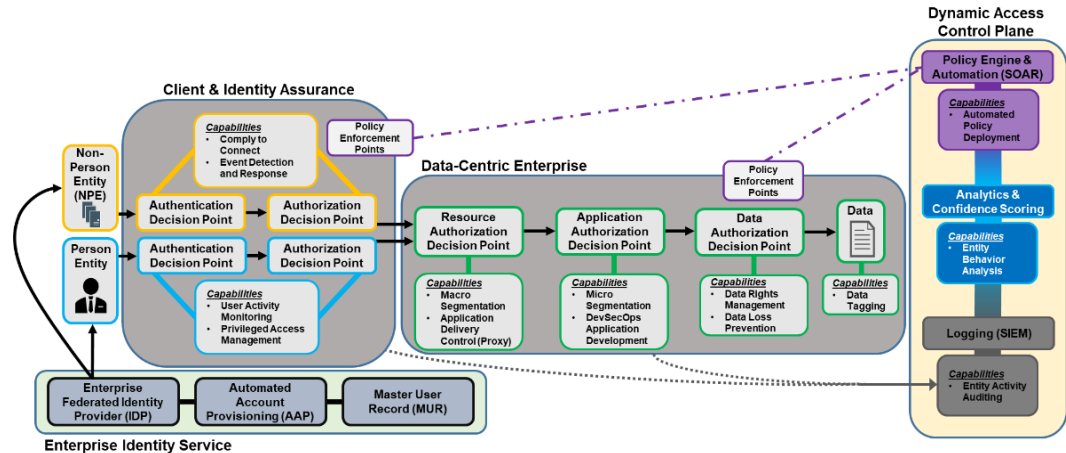


Figure 2: High-Level Operational Concept (OV-1)

---

# Risk Management and Zero Trust



## Risk mitigation

Can you minimize the risk?



## Risk Transference

Can you transfer the risk to some other entity (like an insurance carrier)



## Risk avoidance

Can you avoid the risk?



## Risk acceptance

Can you accept the risk?  
Basically, does it cost more to avoid

- Or mitigate than a breach would cost?