LMNTRIX
BE THE HUNTER | NOT THE PREY

SECURING THE CLOUD: A COMPREHENSIVE GUIDE TO CLOUD COMPUTING & SECURITY

Navigating Cloud Deployment Models, Security Concerns, and Advanced Protection Strategies

WHITEPAPER 2023

LMNTRIX USA
333City Blvd West
18th Floor, Suite
1805 Orange,
CA 92868
+1.888.958.4555

LMNTRIX UK
200 Brook Drive,
Green Park
Reading, RG2 6UB
+44.808.164.9442

LMNTRIX AUSTRALIA
Level 32, 101 Miller
Street, North Sydney
NSW 2060
+61.288.805.198

lmntrix.com

LMNTRIX SINGAPORE
60 Kaki Bukit Place
#05-19
Eunos TechPark
+65 3159 0639

LMNTRIX HONG KONG
14F, Manning House,
38-48 Queen's Road
Central , Hong Kong
+852.580.885.33

LMNTRIX INDIA
VR Bengaluru, Level 5
ITPL Main Rd,
Devasandra Industrial
Estate Bengaluru,
Karnataka v560048
Email: sales@lmntrix.com
+91-22-49712788

EXECUTIVE **SUMMARY**

Organizations using cloud services or in the process of migrating to a cloud-based solution for their business systems need awareness of the security implications of adopting the cloud for information processing.

The cloud offers organizations significant benefits with its on-demand, highly scalable services, typically hosted on shared infrastructure and accessible via a network. In the case of public cloud solutions, connectivity is via the Internet, providing access from anywhere worldwide. Typically cloud-based services offer customers storage and processing solutions, applications, platforms, and infrastructure to suit their needs.

The increasing adoption of cloud services has produced a corresponding increase in the sophistication and prevalence of threats against these services. The rising levels of security risk resulting from using cloud-based solutions have created the need for businesses to implement adequate security solutions. These must cover the entire cloud-based element of the corporate environment to identify and manage risks and accommodate vulnerabilities and weaknesses inherent in cloud services.

Cloud-based Software-as-a-Service offers customers a pre-built cloud service that solves a specific business problem and includes built-in security controls to protect customer information. Cloud-based Platform-as-a-Service, Infrastructure-as-a-Service, and serverless solutions offer customers flexible, configurable components for building services to solve business problems. However, the customer is responsible for implementing and managing the security controls for these latter services.

To implement adequate security for cloud services, you must thoroughly understand the balance of responsibilities between you and the service provider, and the separation measures a service provides.

Information security in cloud services is most effective when introduced as a fully integrated solution at the start of the deployment process. Attempting to add security controls to a live deployed service as a last-minute bolt-on can have unpredictable and expensive consequences.

Tools and techniques for managing cloud security are available to support organizations. These start with specialist services with defined but limited scopes, such as Cloud Security Posture Management and Cloud Infrastructure Entitlement Management. A more comprehensive Cloud Detection and Response offers broad protection.

This white paper explores the issues around implementing adequate security controls into cloud-based corporate systems and options for protecting these systems using security solutions. The aim is to help you determine your cyber security protection requirements for your cloud-based services and select the appropriate managed security solution for business needs

CONTENTS

CLOUD **COMPUTING**

Introduction to Cloud Computing

Overview

Cloud computing refers to computer resources located remotely from the user and accessed via a network. These resources typically comprise processing capability, storage, and applications. The general principle of cloud computing is that it allows users access to resources without needing to install any software on their device, making the cloud platform independent and technology-agnostic.

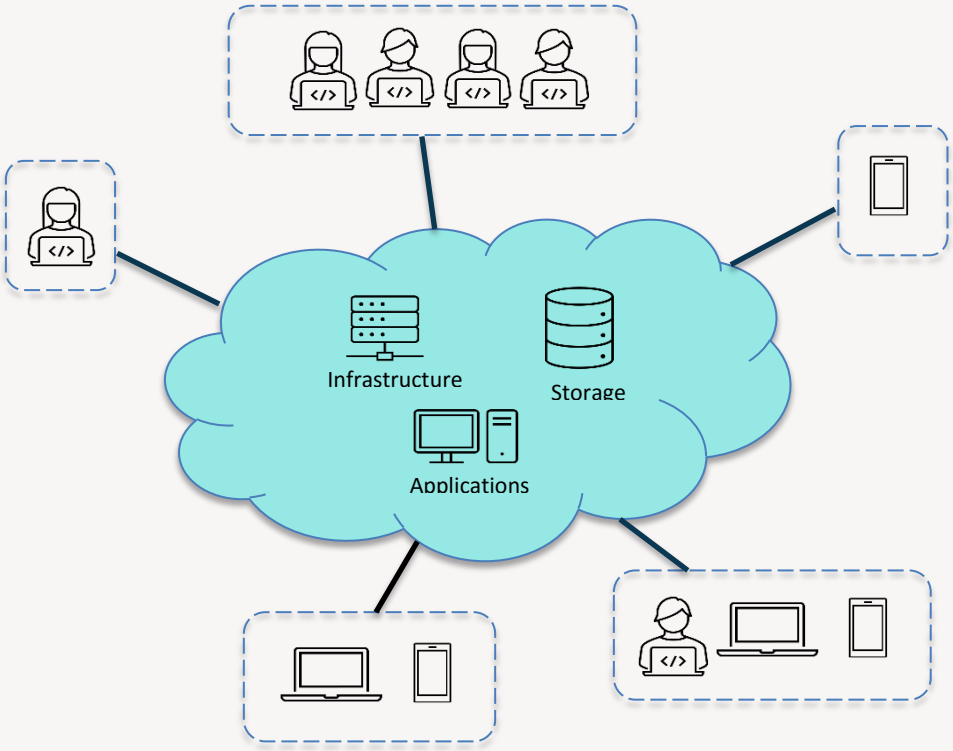

Figure 1 – Cloud Computing Model

The user, an associate, or a third party may own these computer resources. The network access may be a private connection or a public network, such as a telephone system or the Internet. Depending on the user's needs, many resources and choices for connectivity are available.

Cloud computing makes applications and resources available to users worldwide, using a supported device and appropriate network connectivity to access cloud services.

Concentrating services into a cloud-based architecture and making these available to a broad user base offers economic benefits for users.

Terminology

Virtualization technology allows multiple users to share a single physical application or resource instance as if it were a dedicated service for each user. A logical name points to a physical resource, providing local access to each user. Multi-tenant environments employ virtual isolation techniques, allowing users to customize their virtual resources without affecting the underlying physical or other users' virtual resources. A hypervisor is a low-level function that manages the virtualization technology, deploying virtual machines for each user.

Service-Oriented Architecture (SOA) technology allows applications to exchange information irrespective of the underlying technology of each application without the need for changes to any applications to facilitate complete vendor, product, and technology-agnostic compatibility.

Grid computing technology enables cloud computing systems to couple heterogeneous and geographically dispersed resources into an integrated service where the user requires no awareness of the distributed nature of the cloud architecture.

Utility computing principles allow users to access cloud computing resources on demand as a metered service using the principles of the pay-as-you-go model.

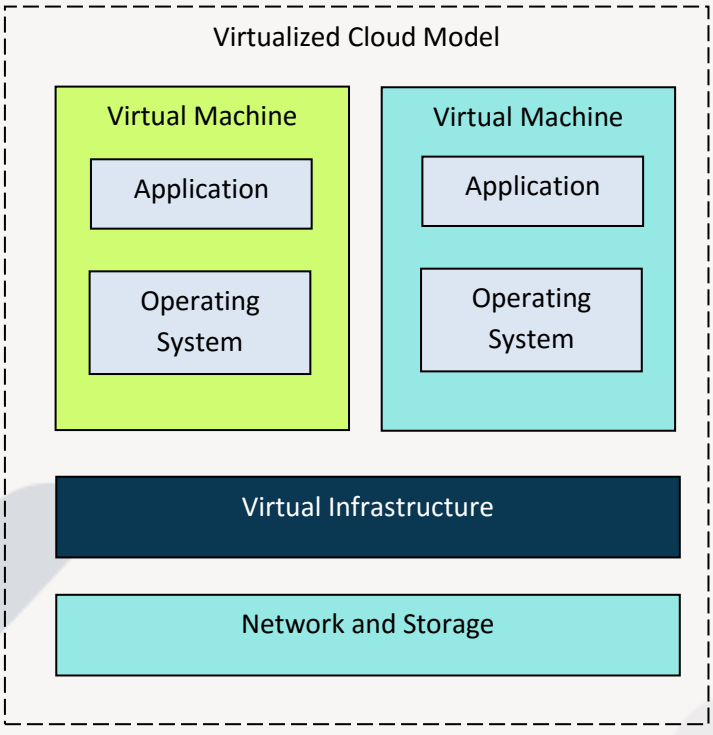

Figure 2 – Virtualized Cloud Model

Architecture

A cloud computing service comprises two loosely coupled elements. The front-end element includes the applications and interface functions resident on the user's infrastructure that access and manage the cloud computing resources. These include the processing capability provided by servers, storage capability, a network interface which manages communications routing and protocols, deployment applications for resources, and management applications to configure, monitor, and maintain the resources.

The back-end elements comprise the cloud-based resources. The front-end and back-end elements communicate over the network that links the user to the cloud environment.

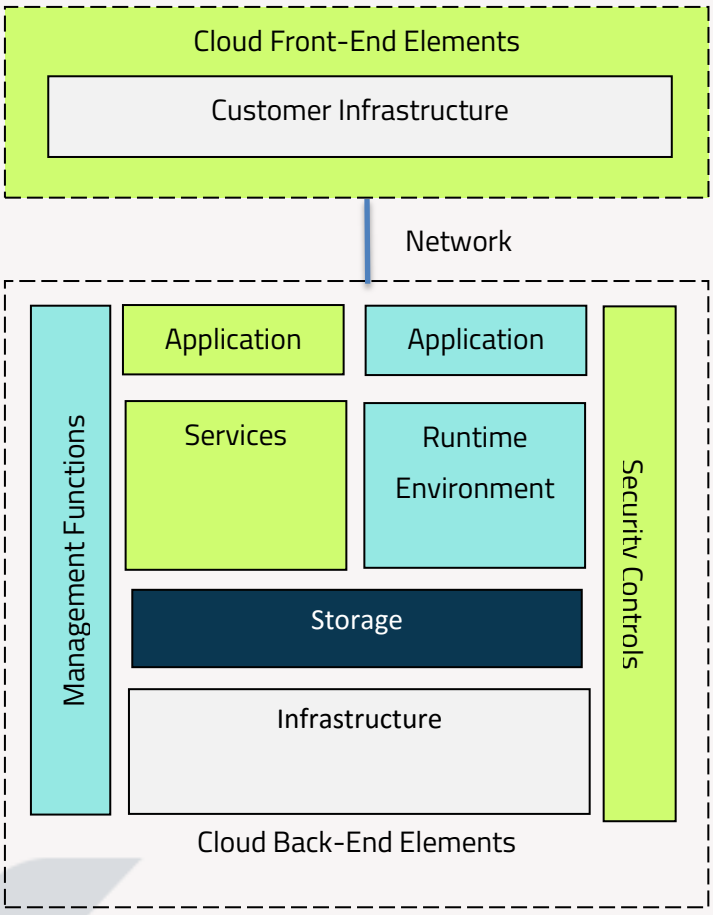

Figure 3 – Cloud Architecture

Cloud Deployment Models

The cloud deployment model defines who has access to the cloud resources, and four main types of models are used.

Public Cloud

The public cloud deployment model allows anyone to access its resources subject to signing up for the terms and conditions, including subscription-based payment for the resources. Major cloud service providers, including Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure (Azure), use this model.

The key benefit of this deployment model is that it offers a cost-effective solution using economies of scale by sharing resources across a broad user base. Internet connectivity allows cloud infrastructure to be located anywhere worldwide, taking advantage of local pricing differences.

Public clouds typically use multiple geographical instances to provide resilience to failures and deliver high availability using grid computing technology.

Public clouds can be integrated with other cloud deployment models to offer the flexibility of hybrid cloud solutions.

Public clouds support utility computing principles for pay-per-use pricing models with high scalability thanks to the large cloud computing resource pool available to users looking to scale up deployments.

The disadvantages of this deployment model are limitations to the customization of resources in a shared environment and more significant security challenges when protecting the confidentiality and integrity of data in a public cloud environment.

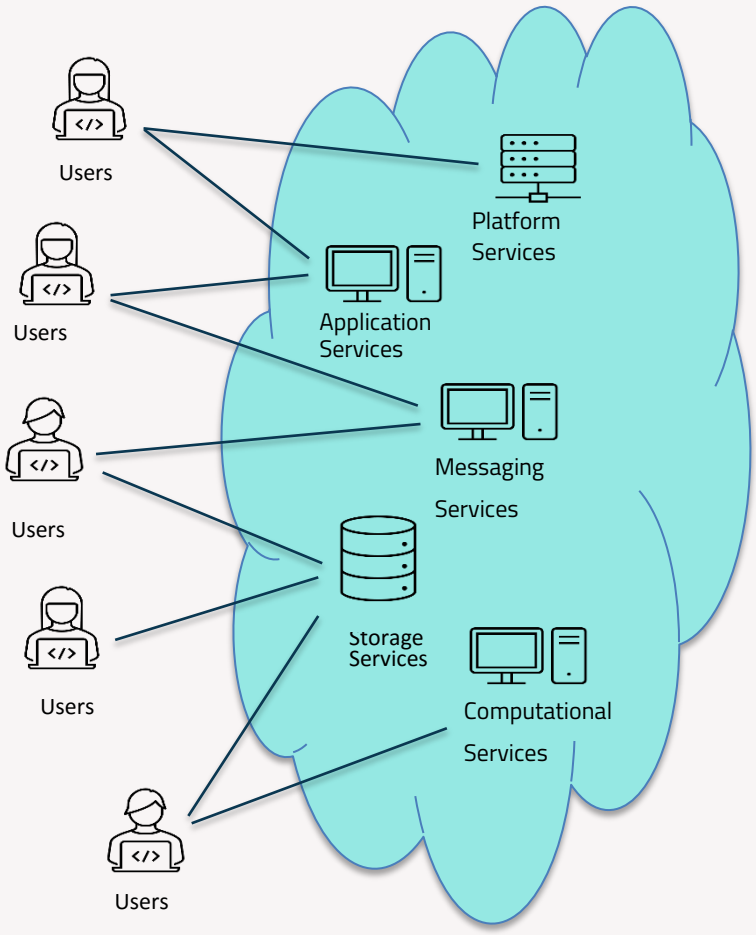

Figure 4 – Public Cloud Model

Community Cloud

The community cloud deployment model restricts access to resources to specific organizations within a particular community. Operation and management may be performed by the organization or a third-party service provider.

The main advantage is that it offers the collaborating organizations the benefits of private cloud deployment with cost savings through the sharing of resources. The organizations benefit from better security than the public, similar in robustness to a private cloud deployment, subject to the nature of the trust relationships between the community organizations. The main security challenge is partitioning information in the pooled cloud storage resources to prevent other organizations within the community from compromising its confidentiality or integrity.

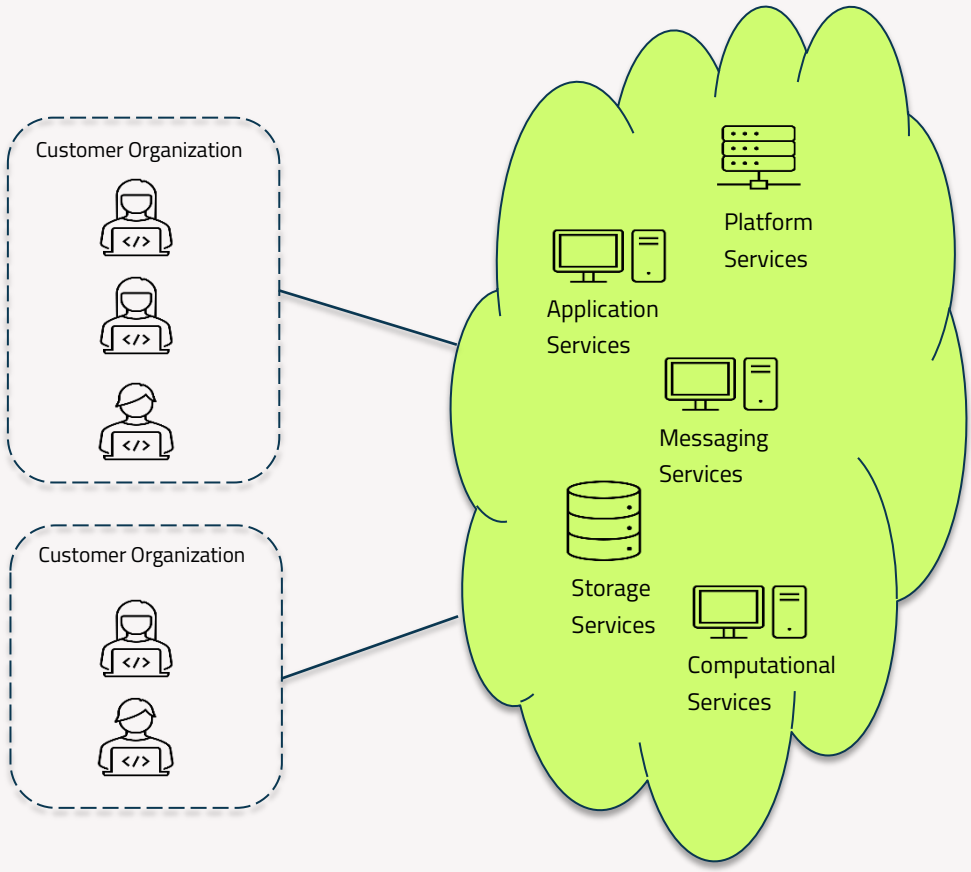

Figure 5 – Community Cloud Model

Private Cloud

The private cloud deployment model restricts access to its resources to a specific organization, either operated by that organization or provided as an exclusive service by a third-party operator.

The key benefit of this deployment model is that it offers high levels of security for information and applications, protecting confidentiality. In addition, this deployment model provides more control and customization of resources for the user.

The disadvantages of this deployment model are higher costs compared to a public cloud deployment, though this can be offset with better operating efficiencies due to operating fewer resources. However, the fewer resources themselves lead to limitations on scalability without the capital expenditure needed to acquire and deploy additional resources to meet increased demand.

Another disadvantage of this deployment model is the more significant connectivity challenges in enabling secure access.

Hybrid Cloud

The hybrid cloud deployment model combines the public cloud deployment model for non-critical resources with the private cloud deployment model for critical resources.

It combines the scalability benefits and cost savings of public cloud deployment with the more robust security available in private cloud deployment.

The disadvantage is that managing public and private cloud connectivity is more complex, and it restricts how this can be practically implemented across different cloud providers.

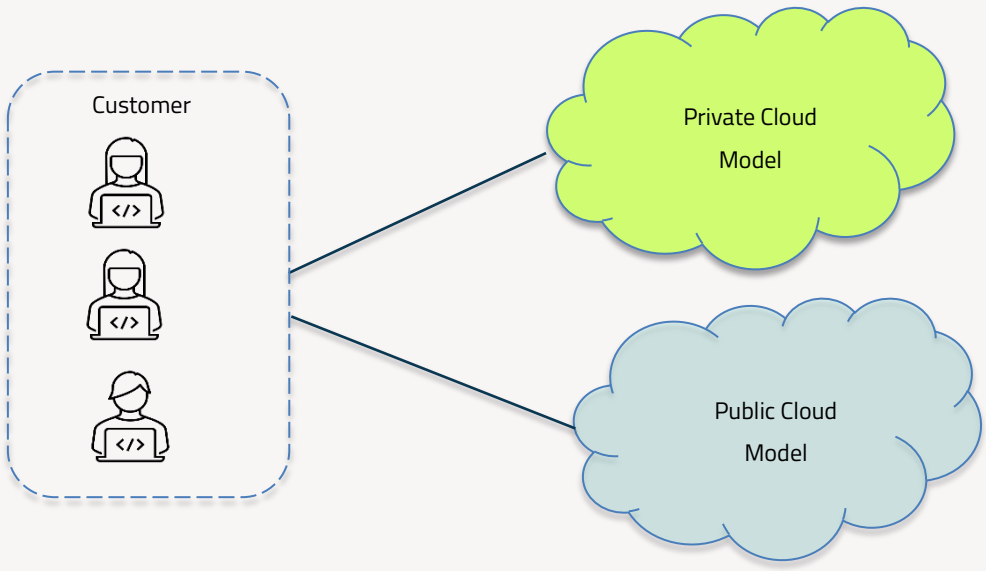

Figure 6 – Hybrid Cloud Model

Cloud Service Models

The cloud service model defines what cloud resources a user may access, and three main types of models are used.

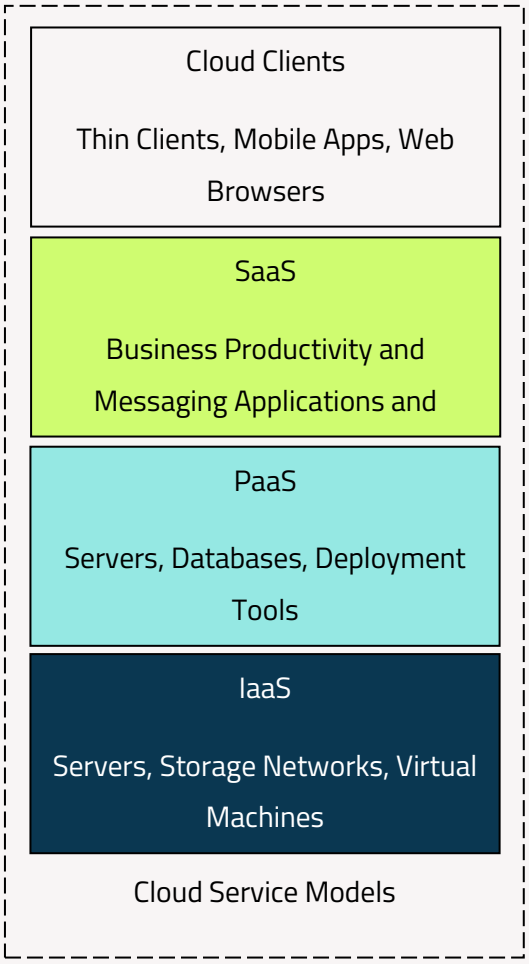

Figure 7 – Cloud Service Models

Infrastructure-As-A-Service

The Infrastructure-as-a-Service (IaaS) model allows the user to access the physical resources within the cloud computing facility, such as processing machines and storage devices. This access can be to physical devices or virtual instances. For example, processing machines include computational processing units, load balancers, system monitoring, and perimeter access controls such as firewalls. Virtual machines are typically offered with preinstalled operating systems and applications.

The user has the freedom to select and scale their allocated resources at will on a PAYG (pay-as-you-go) lease basis. In addition, administrative access to virtual resources allows users to deploy and run any application or service they need.

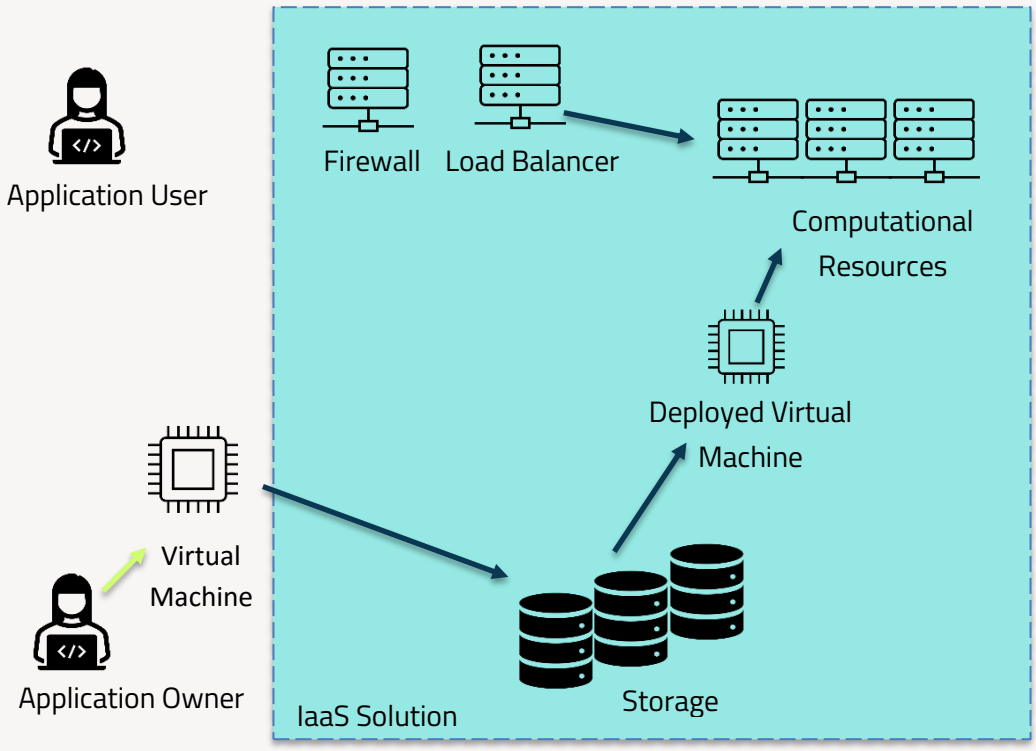

Figure 8 – IaaS Model

The IaaS model allows users to have control of their computing resources with the flexibility and scalability of a cloud-based service, including bandwidth, replication, and geolocation configuration. This allows the configuration of high-dependability systems and backup solutions.

Users can issue administrative commands to virtual machines to manage applications, services, and storage. Virtualization also ensures the interoperability of legacy applications. However, users should be aware that security vulnerabilities in legacy applications will impact the security across their cloud-based systems. This risk is compounded by users' procedural and technical challenges in patching and updating virtual machines to the latest secure build standard.

Virtual storage solutions can also introduce risks to information confidentiality when the data stored on physical media is not deleted when the associated virtual storage is released. This has the potential to allow another user to view the data when they set up a new virtual storage area that uses the same locations on the physical media.

The downside of IaaS is that service availability is directly linked to network availability. Any loss of connectivity due to technical failure or malicious action such as a DDoS attack will lead to a loss of access to cloud-based services and applications.

Platform-As-A-Service

The Platform-as-a-Service (PaaS) model allows the user to access the runtime environment of the cloud computing facility, such as development, provisioning, and deployment tools. This service can range from basic development platforms for browser-based environments to point-and-click tools accessed using APIs for creating and configuring web applications without development knowledge.

The PaaS model simplifies the integration of applications developed within the same environment and offers web service interfaces for integration with applications outside of this environment. It also supports the entire development lifecycle, including workflow processes, review and approval procedures, and deployment mechanisms.

The PaaS model benefits businesses that need to develop bespoke web applications with lower costs, both directly from acquiring resources and their administration and maintenance, and reduces the skillset required, which lowers the barriers to the internal development of applications. In addition, PaaS offers flexibility and scalability to development. However, it constrains the user to a specific toolset, restricting the ability to switch PaaS providers.

The availability and security of applications depend on the PaaS provider's network connectivity. Bandwidth constraints and the need to protect data in transit using cryptography techniques can cause development responsiveness issues and process bottlenecks.

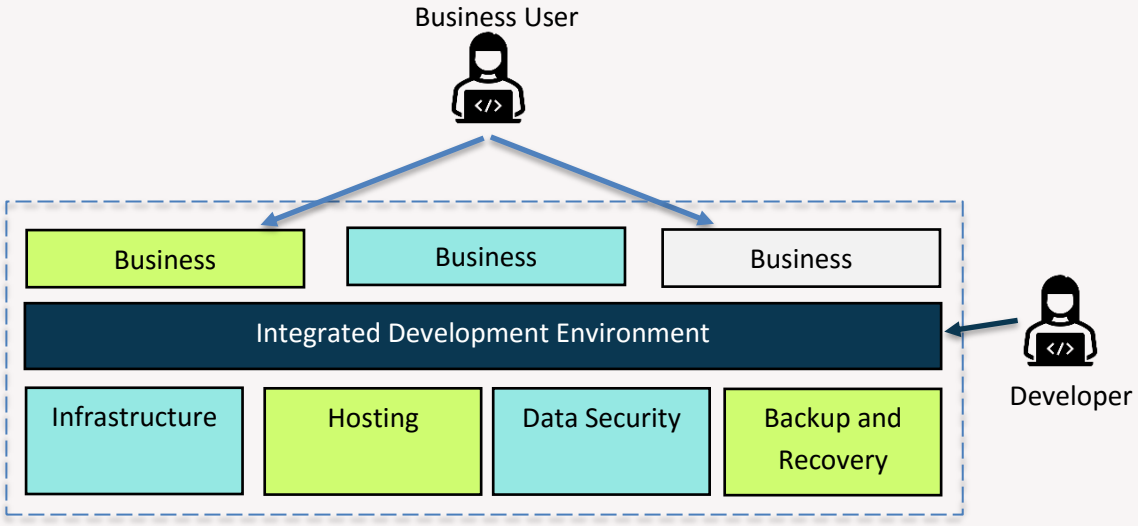

Figure 9 – PaaS Model

Software-As-A-Service

The Software-as-a-Service (SaaS) model allows a business to offer end users access to software applications installed in a cloud computing facility without the need to install the applications on their systems. In addition, this approach ensures all users have access to a common build standard. These applications can be provided as a standard service, such as Microsoft Office 365, or as a fully configurable service, such as a Customer Relationship Management (CRM) application.

The SaaS model allows access on demand to applications over networks hosted and maintained by the service provider, reducing the administrative overhead for the application. In addition, the application can be easily scaled to match changes in the user base.

Cloud deployment of SaaS allows multi-tenant access to applications to deliver economies of scale to reduce licensing costs.

The SaaS model supports various licensing models, including subscription-based, usage-based, or a hybrid combination.

Application availability depends on network connectivity to the SaaS provider, while application security robustness will depend on browser and network security controls. Migration between SaaS providers can also be complex, potentially exposing user information to a specific service provider.

Bandwidth constraints and the need to protect data in transit using cryptography techniques can cause development responsiveness issues and process bottlenecks.

Benefits of Cloud Computing

Cloud-based computing services offer several benefits over premise-based alternatives.

- Cloud-hosted services and applications are accessible from anywhere globally, constrained only by access to the network on which the cloud computing resources are connected.
- Cloud-hosted services and applications make it easier for users to implement effective information-sharing and collaborative working practices to improve productivity and engagement.
- Centralized hosting of services and applications ensures all users access identical build standards and configurations, improving organizational security posture.
- Centralized hosting of information ensures all users access a coherent data set with standard security controls and backup solutions to protect confidentiality and integrity and manage business continuity and disaster recovery.
- Duplication and redundancy built into cloud-hosted solutions enhance services, applications, information availability, and resilience.
- Distribution of cloud-hosted solutions across geographic regions can minimize the impact of major disasters such as earthquakes or conflicts on service availability.
- The use of third-party cloud-based computing services eliminates the need for consumers of the services to maintain systems and applications.
- Cloud-hosted services can be deployed quickly and scaled on-demand to manage usage fluctuations for short-term demand variations across the workday and long-term business growth and operations changes.
- Cloud-hosted services offer real-time operational metrics for informed decision-making based on business process insights and service usage.
- Cloud-hosted services replace the short-term high capital expenditure associated with infrastructure acquisition with lower licensing costs expended over the long term, aiding cash flow management.
- Concentrating physical infrastructure away from individual businesses into centralized cloud hosting facilities can aid in the implementation of energy efficiency measures to reduce the overall environmental impact.

CLOUD **SECURITY**

Cloud Security Overview

Cloud security is the collection of processes, procedures, and technologies that manage the threats to a cloud-based service. These threats may be internal to the cloud service, internal to the users of the cloud service, or external threats targeting the user, the cloud service, and the network that connects them. Cloud security covers all types of deployment models and all levels of service.

Adopting a cloud computing service for businesses requires a balance of productivity and security. Overly restrictive controls can restrict business processes to the point where they are too inefficient to function. In contrast, insufficient rules can lead to unacceptable risks to business services and information that can compromise the organization's viability.

A key challenge in cloud security is managing security across responsibility boundaries to ensure that no gaps exist where the security responsibilities of the cloud service provider end and the user's security responsibilities begin.

Cloud Security Concerns

The critical security issues for adopting cloud computing include:

Information Protection

Cloud-based services have simplified information sharing as a core benefit, with public cloud services, in particular, focusing on offering users facilities for information sharing with other parties using simple processes, including direct links to data or indirect access via email or other messaging technology. Such links often allow access to anyone possessing the link, with no authentication process to confirm the user is authorized to use the link.

While this easy information sharing supports collaborative working practices, it creates significant vulnerabilities that can result in information being inadvertently or maliciously shared with unauthorized parties. Tools and techniques are available for attackers to search the Internet for accessible cloud-based data repositories with little or no access controls and data-sharing links.

Failure to adequately protect information can have significant consequences for a business. Violation of data protection regulations, including the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accessibility Act (HIPAA), and the EU's General Data Protection Regulation (GDPR), can result in severe financial penalties. Data breaches can also lead to significant reputational damage, harming consumer trust in brands and impacting business competitiveness.

Data Sovereignty

A key benefit of cloud computing is that data and services can be located in geographically distributed data centers to maximize resilience from significant and widespread incidents and natural disasters. Public cloud solutions, in particular, can be located almost anywhere worldwide to take advantage of local economic benefits and resource availability regarding power and environmental management. However, for some regulated industries, the location of data can create significant challenges.

For example, an organization collecting, processing, and holding personal information related to identifiable EU citizens will encounter significant challenges if the information is processed or stored in cloud-based systems outside the EU. As a result, data sovereignty becomes a substantial issue in this event, requiring affected organizations to have a complete understanding of where their cloud-based services reside across the whole data processing chain for the entire life of that data.

Another complication is that data held in different jurisdictions will be subject to varying laws regarding the holding organization's responsibilities. The disclosure of data to law enforcement and national security officials affected data privacy.

Incident Response

Incident response procedures for cloud-based services will significantly differ from those for traditional on-premises infrastructure, where security teams have complete visibility of networks and endpoints within the boundaries of their information systems. As a result, the security team will be able to lock down any compromised services or devices to manage an incident and initiate remediation actions.

Incident response for cloud-based infrastructure is more challenging due to the partial visibility across the affected environment and limited control of the infrastructure, platform, and applications.

Access Control

Weak or compromised access credentials for cloud-based services can have a more severe impact than equivalent traditional on-premises services due to the reduced visibility of user actions within a third-party managed infrastructure.

Phishing attacks are evolving to make use of weaknesses in user awareness around access to cloud services. For example, links to cloud based document-sharing resources can be easily imitated by attackers to steal access credentials unknowingly.

Where a business handles regulated information, there will typically be legislative and regulatory compliance requirements to manage access. For example, PCI DSS and HIPAA require organizations to demonstrate that access to sensitive information is restricted to authorized users with legitimate access needs. Cloud-based security controls are required to physically or logically partition such information from other business information in such a manner as to be able to demonstrate compliance with regulatory requirements and industry best practices.

Where cloud users have limited visibility and control over cloud services, demonstrating regulatory compliance and meeting governance requirements can be technologically and financially challenging.

Cloud Security Boundaries

The perimeters of cloud service models define the security boundaries of cloud computing solutions. These perimeters establish the security responsibilities of the service provider and the consumer of the services. These boundaries can be seen in the Cloud Security Alliance (CSA) stack model.

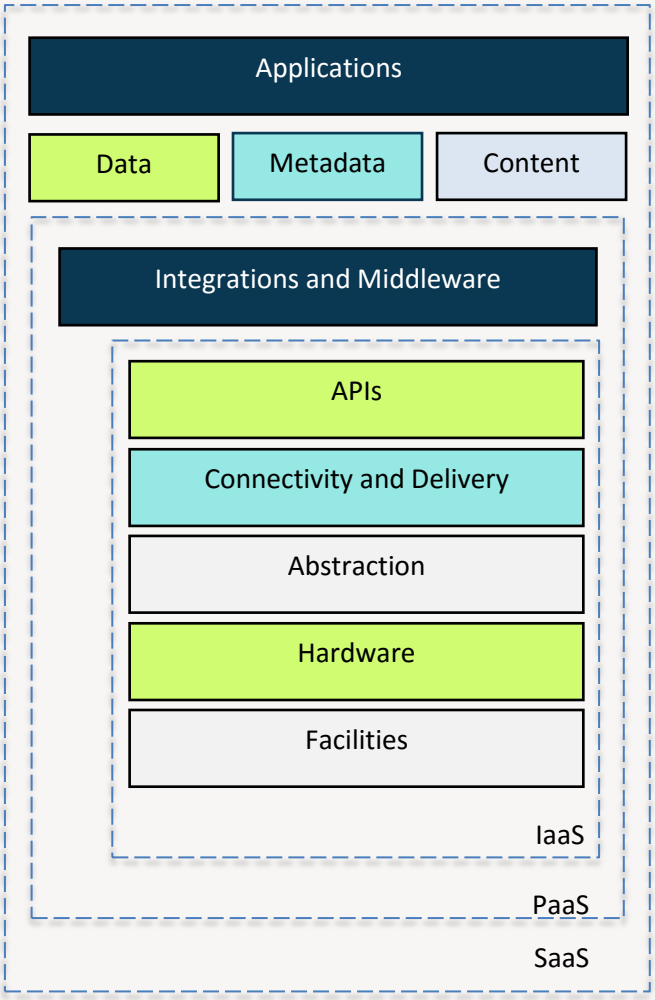

Figure 10 – Cloud Security Boundaries

The nested security boundaries have an accumulative effect. The perimeter around IaaS solutions provides the baseline security controls for the cloud security solution. The next PaaS layer will inherit the IaaS security capabilities and concerns and apply an additional layer of security controls. The final SaaS layer will then inherit the PaaS security capabilities and considerations and apply an outer layer of security controls.

Cloud Security Threats

Figure 11 – Cloud Threats

The critical threats to cloud-based services include the following:

Misconfiguration

Cloud infrastructure can be complex to set up and manage, which creates a high risk of misconfiguration.

Poor Authentication Controls

Controlling access to cloud services is significantly more complex compared to internal networks, where the administrator has full control and complete visibility, which creates a risk of misconfigured or misapplied authentication controls.

External Data Breaches

The scale of cloud computing environments and their multitenant features increases the risk of data loss from cloud based resources. In addition, cloud-based services can be vulnerable due to the reliance on the service provider to maintain secure networks and resources and the customer's responsibilities for patching and securing infrastructure and platform services.

Api Insecurity

Insecure APIs are an increasingly attractive and low-risk attack vector for attackers to gain unauthorized access to cloud resources.

Account Hijacking

Accounts for cloud-based resources can be vulnerable to compromise through phishing attacks and other social engineering techniques.

Cloud Security Requirements

Understanding the security requirements for protecting cloud-based services and information is critical to implementing robust security controls. Cloud computing service providers typically manage the security of the infrastructure and elements of platform and software services but are not usually responsible for information security.

Cloud service providers will typically follow security best practices and comply with recognized standards for information security for the physical environment and their equipment. However, as part of the service agreement, all users of cloud-based services will accept responsibility for protecting data, applications, and workloads in the cloud environment. Cloud service providers have limited visibility into the configuration, access, and usage of a user's cloud services, limiting their protection. Responsibility will lie with the user for operational security.

The security threat landscape has constantly evolved, and threats against cloud computing providers are becoming more sophisticated due to the benefits to an attacker of a successful attack on a cloud service. Organizations consuming cloud services do not face significant governance and compliance risks due to the advanced threats to these services.

While cloud services offer benefits in resilience and recovery, successful cloud adoption will require the user to implement robust security controls to counter credible threats identified and assessed using risk management processes. All cloud deployment models will typically face the same threats, with the type of model affecting the risk levels. For example, the threat from a lone hacker will be lower for a private cloud deployment compared with a public cloud deployment.

Information Security

A key consideration in cloud computing solutions is the reliance on network connectivity between the service provider and the consumer. Typically, these networks are managed by third-party service providers or, in the case of public cloud models, use the publicly accessible Internet. Information flows between the user and the service are vulnerable to interception or, disruption, potentially compromising their confidentiality, integrity, or availability. Protecting data in transit across networks is a critical element of cloud security.

Data Confidentiality

Message encryption techniques will protect data confidentiality as it transits across a network. Robust encryption algorithms with an adequate key length can prevent intercepted messages from being decoded using available technology within a practical timeframe.

Data Integrity

Message encryption practices can protect data integrity using hashing techniques that indicate if the contents of an intercepted message are altered in transit.

Data Availability

Duplication of cloud services and multiple network routing options ensure that no single failure can prevent message traffic between the user and the cloud service outside of major disruptive events such as a DDoS attack on the user's systems, or network access equipment.

Brokered cloud storage access systems provide a mechanism to manage access to data stored in a cloud-based service. Isolating data from direct network access offers an element of information security management by controlling access using authentication and authorization mechanisms. The brokered access system places a proxy-broker pairing between the user and the stored data to manage the access process.

The proxy manages access between the user and the broker element, while the broker manages access between the proxy and the stored data. When the user requires access to cloud-stored data, they submit an access request to the external facing interface of the proxy service. The proxy then forwards the authenticated request to the broker service, which retrieves the data from the cloud

storage. The broker service then transfers the data to the internal interface of the proxy service, which in turn transmits the data via its external-facing interface to the user.

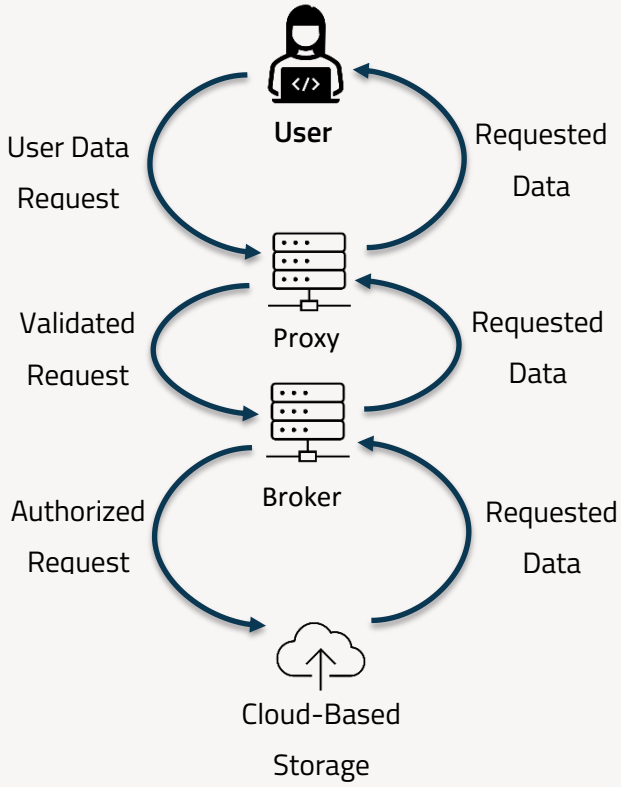

Figure 12 – Cloud Storage Access Mechanism

Cloud Security Challenges

Access Management

The access mechanisms for cloud services are accessible to anyone accessing the network that connects to these services. Where networks are managed by third-party organizations such as telecommunications providers or the Internet, collecting access records for security auditing and monitoring can be limited. This broad accessibility and limited visibility create challenges in authenticating users and tracking access-related event data.

The network-based challenges also impact the ability to control the pathways users may follow to gain network access, making it challenging to prevent communications traffic to and from cloud services from passing over insecure or compromised networks. This challenge is compounded by the global access that cloud computing solutions offer, which may limit both visibility of message integrity and security control technology. For example, some countries do not permit encryption techniques for protecting message confidentiality and integrity due to legislative and regulatory restrictions.

Where the cloud-based services are accessible from the Internet, the risk of an attacker gaining undetected unauthorized access using compromised credentials or exploitation of misconfigured access controls is significantly increased.

Access management challenges include device authentication management issues where the role and identity-based access management processes can allow authenticated users to access cloud-based services using personal or otherwise insecure devices.

Shared Resources

Cloud computing environments commonly employ multitenancy hosting to maximize cost efficiencies. Typically, only user-owned and managed private cloud services will operate on a single tenancy basis, with services dedicated to a single organization. However, this multitenancy mode of operation creates several challenges.

Firstly, cloud resources released by a tenant scaling down its services may be taken up by another tenant scaling up existing or deploying new services. Therefore, there is a risk that information pertinent to the original tenant may inadvertently become accessible to the other tenant.

Secondly, a cyber security attack on the cloud-based services of one tenant may result in the compromise of the services of all tenants in the shared environment, be that at the infrastructure, platform, or application level.

Thirdly, a malicious actor may become a legitimate tenant of the cloud services with the intention of leveraging their legitimate access to launch attacks on the other tenants within the shared environment.

Compliance And Governance

Legislative and regulatory compliance management will rely on information from the cloud services provider, which can be challenging to obtain in a suitable format and timely manner to support corporate compliance and governance processes.

While cloud service providers have compliance responsibilities, overall legislative and regulatory compliance for information security and privacy ultimately rests with the owner of the information, irrespective of whether it resides on-premises or in the cloud. However, the owner will be reliant on the third-party solution provider for elements of managing and monitoring compliance and providing evidence for governance processes.

Service Configuration

Configuring cloud services can be challenging for organizations without the in-house expertise to correctly configure services to maximize productivity while being sufficiently secure to counter threats by mitigating security risks.

Misconfiguration of cloud services is one of the leading root causes of compromise by creating exploitable vulnerabilities in the cloud infrastructure, platform, and applications. Cloud services are often, by default, configured to maximize usability and flexibility and promote information sharing. This usually includes features that allow the sharing of data for collaborative working with anyone with access to a shared link to that data, irrespective of if they are an authorized user with explicitly defined access rights, bypassing access security controls. Configuring services to be secure will often restrict access to information and impose user constraints.

Configuring cloud services security controls is also challenging due to partitioning protective technology between the user-managed elements and the third-party-managed elements of the cloud computing infrastructure. Consequently, the user of the services does not always have complete visibility of the infrastructure's security controls and is reliant on the cloud service provider for the correct configuration, monitoring, and updating of protective controls upon which their security posture is dependent.

Migrating services between different cloud service providers can also create configuration challenges where differences in the third-party-provided security controls can impact user-managed configuration controls.

Third-Party Application Vulnerabilities

Cloud service providers commonly provide users access via application programming interfaces (APIs) to manage connectivity. These APIs will be third-party-provided services with unknown security integrity. Any weaknesses or errors due to flaws in their development or implementation will create exploitable vulnerabilities that external attackers can use to compromise the cloud-based services to which the API provides users access.

Connectivity Dependency

Access to cloud-based services is dependent on maintaining network access to those services. Networks can be configured to provide redundancy so that no single technical failure can cause a total loss of connectivity. However, connectivity of any public network access, such as the Internet, will always be vulnerable to DDoS attacks. There is a continuous capability battle between technological

controls that counter DDoS attacks and the threat actors leveraging broader and more dispersed networks of compromised devices from which to launch ever larger scale attacks.

A successful DDoS attack can deny users access to cloud services, including business-critical information and applications. Additionally, in multi-tenant environments, an attack can simultaneously compromise the availability across multiple tenants. This represents a significant incentive for attackers looking to use the attack to either compromise systems for further attack or gain immediate financial benefit for blackmailing the affected businesses.

Cloud Security Controls

Overview Of Cloud Security Controls

Cloud security controls protect cloud environments by mitigating system weaknesses to reduce the likelihood and impact of attacks.

Technology-based preventive controls counter system vulnerabilities and prevent unauthorized access using multi-factor authentication, boundary firewalls, and endpoint protection techniques. These can be complemented with human factor-focused preventive controls, including training and awareness to minimize exploitable weaknesses.

Threat detection controls are necessary to complement preventive controls by providing a separate independent protective layer to mitigate weaknesses in the preventive controls. System, endpoint, and network security monitoring solutions provide information using services such as a Security Information and Event Management (SIEM) solution or the more capable Security Operations Center (SOC) solution.

Threat response controls provide the mechanism to reduce the impact of attacks identified by the threat detection controls. A mix of technical, physical, and procedural measures is necessary to reverse the damage and restore compromised systems to an operational state.

Overview Of Cloud Security Controls

Cloud security controls must provide comprehensive coverage across all cloud services, including infrastructure, platform, and application levels. The controls should counter all credible threats and reduce the associated risks to an acceptable level. The controls must also be proportionate, providing the required coverage without adversely impacting service usability and business productivity.

The challenge for users of cloud-based services is fully understanding which controls are the responsibility of the service provider and which are the responsibility of the user. This understanding allows the user to ensure all controls across the entire cloud infrastructure, its platforms, and applications are in place and effective as a cohesive suite of protective measures. This includes confirming the service provider's controls are correctly configured and compatible with the user controls to deliver comprehensive, integrated coverage.

Cloud-based services should include a shared responsibility model that details the security responsibilities of the service provider and user to minimize the risk of gaps in protective measures due to misunderstanding or ambiguities.

Cloud control frameworks are available as part of international standards such as the ISO/IEC 27000 series of information security standards and from organizations such as the Cloud Security Alliance (CSA). These frameworks help users gain a comprehensive understanding of their security requirements and responsibilities and which controls they should implement to manage security threats correctly.

ISO/IEC 27017 Security Controls for Cloud Services

International standard ISO/IEC 27002 provides a code of practice for information security controls for cloud services to supplement the ISO/IEC 27001 standard for information security. The 27000 series of standards also include in ISO/IEC 27017 a set of enhanced controls for cloud services. These controls clarify the roles and responsibilities of the service providers and the consumers of the services. The goal is to achieve the same security posture in cloud services as the on-premises certified information management systems certified to the ISO/IEC 27001 standard.

The additional security controls for cloud services are:

- Definition of the security responsibilities of the cloud service provider and the cloud customer
- The management of information assets on termination of a cloud service
- Cloud environment administration
- Cloud environment activity monitoring
- Virtual environment protection and separation controls
- Virtual machine configuration
- Alignment of virtual and cloud network environments

CSA Cloud Controls Framework

The CSA is a valuable source of advice for cloud computing environment security best practices. This guidance includes a Cloud Controls Matrix (CCM) designed to help secure cloud services by detailing controls and defining the responsibilities of the cloud provider and consumer, along with any shared responsibilities.

The CCM provides control and implementation guidance that covers high-level controls to address organizational best practices and cloud provider agnostic technical controls. The CCM encompasses 197 controls across 17 domains. These controls map to other cloud controls standards and frameworks, including ISO/IEC 27001 and PCI DSS, to simplify implementation and support certification activities.

Audit and Assurance

Users of cloud services should have documented policies and procedures supported by independent auditing to ensure appropriate process assurance and governance, including defining key security responsibilities within the organization.

Application and Interface Security

The actions necessary for the planning, deployment, and ongoing support of application security controls that satisfy security and compliance requirements. This includes vulnerability identification and remediation and the segregation of duties to minimize the impact of insider threats.

Business Continuity Management and Operational Resilience

The alignment of Business Continuity Management (BCM) and Disaster Recovery (DR) processes with cloud services.

Change Control and Configuration Management

A risk-based approval process for changes to cloud services, their gateways, and network connectivity. This includes controls to detect or prevent unapproved or unintentional modifications that deviate from an approved baseline and create unacceptable risks using cloud-native guardrails.

Compliance and Governance

Establishing roles, responsibilities, and accountability for information security within the cloud environment.

Cryptography, Encryption, and Key Management

The enforcement of minimum encryption standards, including algorithm and key length configuration and key exchange protocols. Controls should protect data in transit to and from the cloud services and

at rest within the cloud environment from unauthorized access. This includes controls for migrating algorithms and replacing keys in response to vulnerabilities or compromises.

Datacenter Security

Identifying, registering, and managing information assets within the cloud infrastructure and inclusion within the risk management processes.

Data Security and Privacy Lifecycle Management

The identification, classification, protection, and management of critical information in compliance with legislative and regulatory requirements. Controls include data loss prevention tools and services to ensure the security of regulated data in the cloud environment.

Endpoint Management

Controls to protect cloud services from compromised endpoint devices with direct access to cloud infrastructure.

Human Resources

Processes for identity verification, background checks for privileged users, defining roles and responsibilities, and managing security awareness and training programs.

Identity and Access Management

The processes to protect service user identities and associated access credentials and impose minimum security requirements, including complexity controls, multi-factor options, time and geolocation restrictions, role-based privilege management, and segregation of duties. Includes the management of digital identities for all users of cloud services to enable active monitoring and access restrictions as necessary in response to abnormal or otherwise suspicious behavior.

Interoperability and Portability

The processes for application interface communications and data format commonality with technology and provider agnosticism to facilitate migration between service providers.

Infrastructure and Virtualization Security

The processes to manage capacity and resource planning to meet the organization's confidentiality, integrity, and availability requirements of services, applications, and data within the cloud environment.

Logging and Monitoring

A logging collection and analysis strategy for cloud-based services with processes for the collation and secure persistent tamperproof storage of log data.

Security Incident Management, E-Discovery, and Cloud Forensics

Incident management processes are required for cloud-based services that manage investigation and evidence collection for potentially compromised services without physical access to the cloud infrastructure.

Supply Chain Management, Transparency, and Accountability

The shared responsibility model's management process includes responsibilities as understood and defined within contracts and service-level agreements.

Threat and Vulnerability Management

The processes for identifying and mitigating This includes procedures for replacing vulnerable services with the redeployment of a patched version.

Cloud Security Controls

Cloud security assurance allows users of cloud services to assess the robustness of security controls employed by cloud service providers. Assurance programs will provide a measure of the cloud service provider's ability to deliver cloud services with the appropriate levels of security.

For cloud service providers, certification and accreditation to recognized standards will demonstrate compliance with industry cloud best practices, standards, legislation, and regulations.

For cloud service consumers, certification and accreditation to recognized standards will demonstrate their commitment to cloud security internally to staff and externally to customers and suppliers. It can also verify compliance with relevant regulatory and legislative standards for regulated businesses.

A cloud security assurance program will provide an organization with a framework for assessing and improving its cloud security practices against best practices, regulations, and standards. This supports processes for reviewing and maintaining security posture, management of emerging threats, and continuous improvement.

CLOUD SECURITY **POSTURE MANAGEMENT**

Definition

Overview

Traditional cloud security solutions focus on detecting threats from external attacks and malicious activities within the cloud environment. However, they do not address protection against inadvertent misconfigurations leading to data leakage.

Cloud Security Posture Management (CSPM) automates identifying and remedying all risks within a cloud infrastructure, encompassing IaaS and PaaS services. CSPM facilitates risk management, incident response, and compliance management for cloud services that cover deliberate attacks and accidental issues.

CSPM is an evolution of the Cloud Infrastructure Security Posture Assessment (CISPA) technique for reporting vulnerabilities in cloud services. CSPM introduces automation to the process to deliver faster and more comprehensive protection.

The principle of CSPM is the continuous monitoring of risk levels in the cloud environment to identify new vulnerabilities and emerging or evolving known threats that increase the risk to a level where measures are required to reduce or eliminate that risk. This proactive threat detection allows for identifying and rectifying security weaknesses before threat actors discover and exploit them.

The purpose of CSPM is to manage the dynamic network connectivity of cloud environments within a third-party managed infrastructure where traditional boundary security controls can't be applied and manual processes cannot provide responsive and scalable protection. Additionally, CSPM helps manage the reduced visibility inherent in decentralized cloud-based solutions

Features

CSPM provides cloud service users with automatic discovery and visibility into the cloud infrastructure and their security configuration on deployment. Additionally, it offers centralized management of security policies across deployments in single and multi-cloud environments. This enables the identification of misconfiguration, change management monitoring, and oversight of cloud security controls.

CSPM minimizes the risk of misconfiguration-related vulnerabilities by automating the assessment of configuration settings against benchmarks and best practices to identify and remediate real-time

risks, including unauthorized modification. This protects against inadvertent errors in setting and changing the configuration and malicious interference in settings as part of an attack's kill chain.

CSPM can proactively detect threats within the cloud environment using a risk-based identification and management approach. It uses the perceived likelihood of vulnerability exploitation to monitor those areas with the most significant risk levels to reduce the security alert landscape to a manageable scope. CSPM continuously monitors the high-risk areas of the cloud environment for indicators of unauthorized, suspicious, or malicious activities and access to cloud resources in real time.

CSPM can be used in pre-deployment environments to detect and remediate risks due to misconfiguration or poor development practices before any service or application is deployed into the operational cloud environment. It can be integrated into DevOps (Development-Operations) processes to provide a DevSecOps (Development-Security-Operations) capability to improve the security posture of development activities without compromising the benefits of the DevOps pipeline process.

Risk Coverage

The typical security risks that CSPM solutions can detect and resolve include the following:

- Service misconfiguration issues
- Missing access controls exposing data or services
- Access control permission errors or weaknesses
- Access control authentication issues
- Absent or weak encryption mechanisms
- Legal and regulatory compliance issues

CSPM Benefits

Key Benefit Overview

CSPM covers security threats from both deliberate and accidental sources. Most cloud security solutions focus on deliberate attacks from external threat actors and internal attacks from malicious insiders. However, CSPM also covers security threats from unintentional actions such as misconfigurations or incorrectly executed configuration changes.

CSPM provides administrators with unified visibility across multi-cloud environments, with a coherent, integrated single source of truth for all cloud services. This eliminates the risk of undetected attacks

where information from multiple vendors displayed across multiple consoles creates blind spots or misinterpretation of data.

CSPM's unified visibility also supports the integration of system alerts into a single coherent data set where duplicate or consequential alerts can be filtered to reduce workload and the accompanying risk of alert fatigue with the benefit of improving the effectiveness of security monitoring processes.

CSPM automation combined with the continuous monitoring of the cloud environment improves threat detection and response reaction times with the benefit of reducing the risk of threats being exploited before they are identified and remediated. This has the additional benefit of automatically detecting and resolving compliance issues with security policies to support governance and maintain its security posture.

Benefits over CASB

Cloud Access Security Brokers (CASB) use security policy enforcement to protect cloud-based data and services from threats. It primarily extends the visibility of on-premises systems into the cloud environment. Their essential functions include:

- Identity and access management
- Authentication management
- Data sharing controls and loss prevention
- Configuration audit
- Malware detection
- User and entity behavior analytics
- Risk assessment, compliance, and governance support
- Data encryption

CASB solutions protect cloud-based services from attacks by APTs, malware, malicious insiders, and the exploitation of compromised access credentials. They offer risk assessment, e-discovery, and auditing functions to support compliance activities and incident investigation. They also enhance the protection of information confidentiality and integrity using data encryption techniques.

CASB solutions manage security policy enforcement at the interface between the cloud-services user and the services providers. In addition, CSPM offers benefits in managing enforced security policies and continuous compliance monitoring.

Benefits over SSPM

SaaS Security Posture Management (SSPM) solutions provide continuous cloud application monitoring services. They primarily manage the security controls around integrating SaaS with existing business infrastructure. Their essential functions include:

- Monitoring of on-demand applications
- Security policy compliance
- Protection of operating systems, hypervisor, and infrastructure
- Network traffic monitoring

SSPM solutions protect existing systems against threats from misconfiguration and authentication issues in cloud-based applications. They also enhance the visibility and monitoring of the applications to improve security posture.

SSPM solutions offer narrow cloud-based application security controls. CSPM offers additional benefits in covering all cloud services, including protecting platform and infrastructure services with comprehensive protective measures and intelligent automated processes.

Benefits over CWPP

Cloud Workload Protection Platform (CWPP) services are agent-based server protection solutions that operate on a scalable workload management solution that focuses on maintaining cloud-based business operations in a threat environment. They primarily manage visibility across multiple cloud environments to deliver cohesive security controls. Their essential functions include:

- Workload configuration and behavior monitoring
- Cloud application management
- Event log monitoring and management
- Host-based Intrusion Detection System
- Malware scanning
- Network protection using segmentation, firewalling, and traffic monitoring
- System integrity and vulnerability management

CWPP solutions support the hardening of cloud-based services and the management of vulnerabilities using micro-segmentation techniques to separate security controls from the underlying infrastructure and extend protection and visibility in multi-cloud environments.

A downside of CWPP is that this solution requires the installation of an agent on every cloud-based service it protects. These agents must be deployed and maintained, adding an operational overhead to the business. These agents can also potentially adversely impact system performance. Where legacy

systems cannot support the installation of an agent, this system will fall outside the scope of the CWPP protection and create a blind spot in the business environment. CWPP can also not identify risks associated with the lateral movement of attacks across services, a critical element of advanced attack vectors.

Another issue for CWPP is they cannot assess the impact of security issues to enable effective prioritization of alerts due to limited visibility of the context for the services they protect. This creates the need for the security team to triage alerts to determine remediation importance and schedule manually.

CWPP solutions offer unified workload protection using traditional on-premises security controls adapted for a cloud environment. CSPM offers additional benefits in being purpose-built for cloud services with capabilities that extend beyond workloads to the service control plane and include intelligent automated processes with guided remediation procedures.

CSPM Tools and Platforms

CSPM tools autonomously manage and mitigate risk across an organization's entire cloud-based environment to enhance security posture while reducing resourcing requirements.

CSPM tools provide the necessary visibility to manage and mitigate security risks across an organization's cloud environment. They employ continuous monitoring to provide threat detection and protection. They also support service hardening, compliance monitoring, and remediation workflows. Cloud-based workflows that are non-compliant with security policies are reported and prioritized for remediation.

CSPM tools should provide the following capabilities:

- The monitoring and management of IaaS, PaaS, SaaS, and solutions across the entire business ecosystem, including cloud-based and on-premises environments. This includes all cloud deployment models and multi-cloud environments
- The autonomous identification and remediation of misconfigurations in cloud-based services, including unauthorized or erroneous changes
- The continuous monitoring of security policy compliance across all cloud-based services
- The continuous monitoring of changes to regulatory compliance mandates, including HIPAA, PCI DSS, and GDPR, and reporting of impact on compliance of cloud-based services
- Risk assessment of cloud-based services using baseline frameworks and external standards to assess security posture.

CSPM Deployment

Implementing effective CSPM requires careful planning and preparation to maximize the benefits of deployment. The following best practices will aid successful deployment:

Proportionate Risk Appetite

All businesses face numerous credible threats that can lead to an enormous list of risks. Attempting to manage too many risks can quickly overwhelm resources and prevent the effort from being directed to more immediate security tasks such as detecting attacks and managing ongoing incidents. CSPM should focus on addressing that subset of risks that exceed the risk appetite for the business. This risk level should be initially set high to focus on critical security threats and subsequently reduced once the critical risks are being actively managed to a proportionate level so all unacceptable and intolerable risks are managed.

Prioritized Security Alerts

Security teams managing cloud-based services, particularly in public cloud deployments, will face significantly large numbers of security alerts which can easily overwhelm the team. Alert fatigue can lead to alerts relating to severe incidents, including ongoing attacks, being lost in the noise of more mundane security policy violation alerts. Security alerts should be prioritized to allow security analysts to concentrate on those relating to critical cloud-based services.

Security by Design

Development pipelines should include security controls as an integrated component of the design lifecycle to prevent applications that continuously use and access new resources from being vulnerable to exploitable weaknesses. In addition, the development process should ensure that any misconfiguration issues are identified and remediated before the developed application is deployed into a live environment.

Benchmarked Security Posture

Monitoring cloud security posture can provide meaningful metrics measured against relevant and appropriate cloud-specific benchmarks. The Center for Internet Security (CIS) produces benchmarks for vendor product families. These benchmarks offer prescriptive configuration recommendations based on industry best practices against which an organization can compare its security policies to gain insight into the effectiveness of its security posture.

Continuous Improvement Processes

The security threat landscape continuously evolves in parallel with the security controls available to counter threats. Therefore, the security posture of a business needs to progress alongside these factors by incorporating continuous improvement practices, including the inclusion of the latest standards of security posture benchmarks.

CLOUD INFRASTRUCTURE **ENTITLEMENT MANAGEMENT**

Overview

Cloud infrastructure entitlement management (CIEM) solutions manage user identities and their privileges in a cloud environment. CIEM aims to identify threats to cloud-based services from accounts with unnecessarily high access privileges across different cloud-based services and multi-cloud environments.

This identifies and mitigates risks associated with access to cloud-based services and supports identity and access management processes by improving visibility in the cloud environment. It also makes the enforcement of security policies based on the principle of least privilege simpler to implement.

The principle of least privilege is an information security concept where a person or service is granted the minimum levels of access necessary for that entity to perform its assigned duties. This may be an ordinary user accessing applications and data in the cloud, an administrator managing the configuration of applications and data in the cloud, or a cloud-based application accessing another application or data.

Applying the principle of least privilege is an information security best practice for managing and securing privileged access to critical assets in a cloud environment to reduce risks of accidental misuse or malicious compromise without disrupting the business needs of the assets.

Adopting a least privilege policy reduces the attack surface of the cloud environment by limiting the assignment of super-user and administrator privileges. This reduces the opportunity for an attacker to exploit privileged credentials as part of its kill chain. This also reduces the ability of malware to move laterally across infected systems by leveraging elevated privileges on endpoints.

CIEM solutions enable businesses to implement on-demand, just-in-time privilege elevation in response to a legitimate business need and, more critically, remove those privileges as soon as they are no longer required to eliminate privilege creep.

CIEM solutions allow security teams to centrally manage identities and privileges for cloud services and monitor compliance with security policies that enforce the principle of least-privileged access across cloud infrastructure and resources. CIEM uses advanced analytics underpinned by machine learning techniques to monitor user and entity behavior and assess privilege assignments instead of the traditional use of generic rules and conditions.

CIEM solutions also support automated security policy adherence through monitoring and auditing, allowing demonstration of compliance with regulatory requirements and governance processes.

Importance of Cloud Permissions

Cloud permissions allow system administrators to control access to cloud-based services to facilitate business processes while complying with security requirements. The challenge in multi-cloud environments is that each cloud provider operates their own IAM processes with divergent capabilities and differing principles and constructs for setting access permissions. This divergence also impacts the migration of cloud-based services between different service providers.

Traditional IAM permission models for on-premise systems cannot be readily translated into cloud services where services and data are hosted on equipment within highly distributed locations not owned and operated by the service user. Models built using the principles of privileged and non-privileged access do not directly map to dynamic cloud-based services. Using shared responsibility models can also increase the complexity of securely managing access.

Typically the cloud provider provides directory services and user authentication, authorization, and auditing functions for the infrastructure and services. The user is then responsible for adding access controls and monitoring to form a comprehensive and cohesive solution to manage cloud permissions.

Multi-cloud environments will host services accessed by enormous numbers of users across large numbers of clients, creating visibility and complexity challenges. In addition, each user or entity that accesses cloud-based services will have access rights for each service. Hence, a large number of users (U) accessing a large number of services (S) will potentially create U multiplied by S different permissions that need to be monitored, audited, and managed.

The dynamic nature of cloud services, with the instantiation of services and applications on demand and their disposal when no longer needed, can result in cloud permissions requiring frequent changes as the cloud environment changes. Containers can be spun up and down continuously, which creates an administrative burden for managing access privileges assignment and tracking.

In a typical cloud environment, the IAM permissions will be required to manage access and administration of the following elements:

- Multi-factor authentication configuration and control.
- Single sign-on and role-based access control.
- Cloud storage service access.
- Cloud data service access, including virtualization, databases, and network services.

- Cloud resources access includes Virtual Machine servers, Kubernetes containers, and serverless infrastructure.
- Cloud administrative account access, including service and financial management consoles and security administration functions.

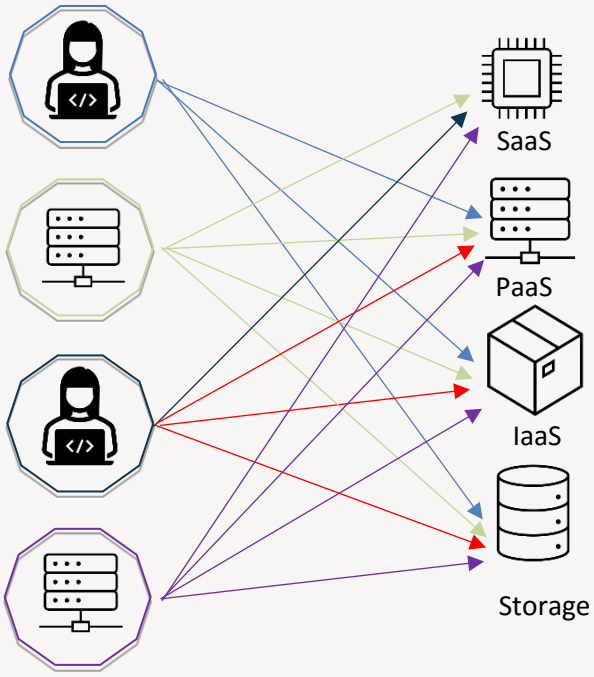

Figure 13 – Cloud Permissions Visualized

The use of manual cloud IAM processes can significantly increase the risk of misconfiguration when assigning or revoking cloud permissions with reduced visibility, potentially hiding errors from review and audit processes. Manual following procedures to manage consent will also be time-consuming and resource intensive compared with operating automated processes. Any delay can be critical in a dynamic cloud environment, reducing productivity or creating vulnerabilities.

Manual processes also tend to increase the use of static or infrequently updated authentication credentials, which increases the risk of security breaches from compromised credentials. They can also result in over-permissioned accounts and users with excessive cloud entitlements where permissions are not removed once their business need has ended.

Benefits of CIEM

Identity and access management solutions are available to manage access controls on traditional on-premises infrastructure and in static self-hosted environments. However, the move to dynamic cloud-

based services, including platforms and infrastructure with restricted visibility, creates access management challenges.

Cloud service providers typically offer identity and access management services for their specific services, but these will not integrate across a multi-cloud environment provided by different providers. This creates the problem where large organizations attempt to manage account privileges separately for each cloud environment without the ability to visualize which user accounts have which rights across the entire infrastructure surface. This creates the risk that an account may have unnecessary elevated privileges across multiple cloud environments, allowing an attacker to extend an attack on one cloud service to numerous diverse cloud services.

CIEM solutions allow businesses to centrally manage user identities and their privileges across a multi-cloud environment with complete transparency to manage access risks. The key benefits of CIEM include the following:

- Visibility of privileged access rights across the multi-cloud environments offering the ability to gain insight into access-related risks and compliance with security policies
- Enhances identity and access management service with continuous monitoring of access activity to support identification of deviations from security policies for least privilege principles and removal of inactive accounts
- Provides automatic detection and remediation of security issues, including the use of compromised accounts, suspicious activity, and malicious insider activity based on monitoring of user behavior and comparison with baseline activity levels
- Provides monitoring and auditing of privilege assignment and use across the cloud environment to support compliance monitoring and adherence to policies, accreditation, certification, and regulatory requirements.

Implementing CIEM

Overview

CIEM controls aim to protect cloud environments from breaches from access control-related vulnerabilities that result from users or entities being assigned overly permissive credentials or the application of poor credential hygiene practices.

A challenge in the cloud environment is that permissions can be inherited from multiple sources, which, combined with visibility limitations, makes the process of obtaining complete visibility of all permissions for each user and entity complex. Permissions can be directly assigned or obtained from resources, access control lists, or assigned roles in cloud services.

The critical issue with excessive cloud permissions is the opportunity for an attacker to use compromised credentials for one cloud service to extend control across the entire cloud infrastructure. In addition, such access levels make it easier for attackers to gain complete control over an organization's information systems through the cloud presence.

Manual implementation of least privilege access principles across an organization is resource intensive due to continuous monitoring and adjusting permissions in a dynamic environment. In addition, rapid changes to the instantiation and retirement of cloud services can create compliance challenges and increase the risk of misconfiguration of access controls.

CIEM Solution Features

CIEM solutions should monitor and manage permissions across the cloud environment, with complete visibility across multi-cloud ecosystems, to produce a single source of truth for assigned access credentials.

The product should be capable of detecting unnecessary or unused permissions, incorrect privilege levels, and vulnerabilities in access credentials. The product should also be capable of identifying users and entities with permissions that create risk levels that exceed the organization's risk appetite to allow review and reassignment of duties as necessary to reduce risks to an acceptable level.

Continuous monitoring of compliance with security policies and benchmark security posture should employ user and entity analytical-based behavioral monitoring techniques to detect indicators of compromise and support audit and governance processes.

The solution should also be capable of automatically resolving issues and applying for permissions in compliance with the principle of least privilege access and in accordance with best practices continuously in a highly dynamic multi-cloud environment.

The solution should also be capable of integrating with the cloud service provider's identity management services to support cloud-agnostic single sign-on functionality. This will enable transparent migration of cloud services across provider environments that do not compromise access control security posture.

Permissions Management Use Cases

CIEM permissions management solutions cover three critical use cases: discovery, remediation, and monitoring. Following each use case as a sequential series will provide comprehensive insight into cloud permissions.

Discovery

In the discovery stage, CIEM will assess permission risks using gap analysis techniques to evaluate the difference between assigned and actively used permissions. This will highlight unused and unnecessary permissions accumulated due to permission creep. It also provides visibility of the permissions-related risk across the entire cloud infrastructure for all user and entity identities, actions, and resources.

The results of the discovery stage can be quantified using a Permission Creep Index (PCI) metric. This represents a measure of the number of unused or excessive permissions across the cloud environment and the risk these issues create based on the permission levels and the impact of misuse.

Remediation

In the remediation stage, CIEM will right-size permissions based on user and entity behavior by revoking unused permissions. New access permissions for cloud services can be granted on demand using automated just-in-time workflows that integrate authorization processes.

Monitoring

In the monitoring stage, CIEM will detect anomalous or suspicious user and entity behavior to generate security alerts supported with detailed forensic event capture and reporting for incident response investigations and response.

CIEM will also manage compliance to least privilege access principles and support implementing a zero-trust security strategy. This has the benefit of unifying access control policies across multi-cloud environments, including IaaS platforms, to deliver consistent and coherent security policies for IAM.

Permissions Management Case Study

IAM provides the infrastructure necessary to manage authentication and authorization for cloud services. In the example of an AWS account, the IAM infrastructure supports the following access control process:

- Users or entities requiring access submit their sign-in credentials to the AWS service
- A trusted principal matches the credentials to authenticate these. A principal is a user or entity with a root user AWS account or IAM privileges to sign in and make requests to AWS
- The principal requests access to the relevant cloud resources
- Access is then granted to the user or entity by the principal in response to an authorization request

For example, a request by a user to access a specific service will result in the intermediatory principle issuing a request to the service on their behalf for it to provide the user with access. The service then verifies that the user is authorized to use its service and what security policies govern the access. Only if the user is allowed access and the request adheres to the policies in force for that access will the service grant access to the user via the principal.

AWS applies the least privilege access principle by denying all access requests by default and only permitting access if explicitly allowed by the permissions policy, using identity-based or resource-based access controls.

For example, an AWS EC2 instance that requires access to an AWS S3 bucket to retrieve configuration data will only be granted read access to the specific S3 bucket that contains the requested data via an IAM principal. The EC2 instance will be denied access to change the data and will be unable to access any other data in any other S3 bucket. Furthermore, once the data is retrieved, the EC2 instance will no longer have access to the S3 bucket.

In AWS, the IAM principal can be a user, a role, or a group. When any of these identities are created, they have no default or inherited permissions. Instead, all permissions are explicitly assigned using policies. For example, identity-based policies permit IAM principals to grant defined permissions to a specific identity. In addition, AWS provides fixed-managed policies from which customer-managed policies can be built to support policy creation.

The AWS policy statements are formed from the following elements:

- Effect: Defines if the policy will allow or deny the defined action.
- Action: A definition of the specific actions that the policy governs when specified conditions are satisfied. For example, the s3:CreateBucket actions an IAM Principal to grant an AWS S3 service permission for its API to request the creation of an S3 bucket.

- NotAction: A counterpart of the Action element that defines the specific actions that the policy will deny when all other actions are permitted by default when specified conditions are satisfied.
- Resource: A definition of the specific resources the policy applies to using the AWS Resource Name (ARN) format.
- NotResource: A counterpart of the Resource element that specifies which specific resources the policy does not apply to when all other resources apply by default.
- Condition: An expression that defines which condition keys and values in the policy must match against keys and values in the request context sent by the IAM principal. Condition keys can be specific to a resource or be global for the AWS service, such as aws:CurrentTime which can be applied to grant access when specified date and time values are met.

AWS policy statements can be created easily using the IAM console integrated into the AWS Management Console. A create policy option on the policies navigation pane opens a visual editor which allows the selection of specific services and the addition of permissions. In addition, policy statements can be manually added or imported from any existing managed policies.

Various tools and techniques are available to support the creation of AWS policy statements. The AWS Lambda service provides a convenient method. AWS Lambda is a serverless, event-driven computational service that enables users to execute instructions virtually for any cloud-based application or service without the overhead of provisioning or managing servers. The majority of AWS services and SaaS applications support the Lambda function.

CLOUD DETECTION **& RESPONSE**

Overview

Introduction

Cloud detection and response (CDR) extends traditional detection and response strategies across the cloud environment. It enables cohesive monitoring for security threats across a multi-cloud ecosystem encompassing dynamic IaaS, PaaS, and SaaS resources and supports the incident investigation and response processes.

CDR solutions follow the identify, protect, detect, respond, and recover lifecycle, following information security best practices to detect and mitigate real-time threats in the cloud environment.

CDR is gaining importance as more organizations embrace multi-cloud solutions while attacks from advanced persistent threats on cloud services increase in number and sophistication. The challenge is managing the volume of alerts that using dynamic, scalable services in a public cloud environment creates.

Threat Detection

Threat detection is a critical capability in the process of detecting and remediating attacks. However, traditional threat detection tools employ a "one-size fits all" approach that can generate alerts for every anomaly occurring in the cloud environment. This can overwhelm security analyst resources, create alert fatigue, and result in the volume of benign abnormalities hiding indicators of attack.

A CDR solution should consider the cloud environment along with user and service behavior across this environment and use this information to correlate alert data and identify actual threats.

The solution should provide observability into attack propagation by identifying and monitoring anomalous behavior in runtime across the operating environment for the complete workflow lifecycles. Additionally, the solution should correlate multiple suspicious events into a cohesive and understandable set to provide observability of the attack manifestation. This helps reduce the attack surface area, mitigate data exfiltration, eliminate account compromise, and minimize business disruption.

Operations

CDR solutions enhance capabilities compared with existing cloud security tools for posture management. Machine learning-based technology monitors system operations to establish a baseline of typical behavior within the cloud environment based on observed user and entity behavior.

Metadata and system logs allow the tool to gain a precise outlook on normal activities to create a contextualized baseline from which abnormal or suspicious behavior patterns can be identified and attacks identified and alerted.

Examples of abnormal and suspicious behavior include:

- APIs accessed at unusual times or days or by an unexpected user or entity
- Services accessed by a specific user at an unusual time for that user
- Machine invocations at unusual times or days or by an unexpected user or entity
- Stored information accessed by an unexpected user or entity
- An abnormally large volume of stored information accessed by a user or entity

These unusual events are labeled as malicious behavior indicators (MBIS).

CDR Benefits

CDR solutions offer continuous monitoring for threats in a multi-cloud environment to detect anomalous behavior that indicates an attack is in progress. Typically, CSPM solutions scan for vulnerabilities, misconfigurations, and other static events within the dynamic cloud environment at fixed points in time. However, this creates a window of opportunity between vulnerabilities emerging and these being detected and subsequently mitigated. In addition, the CSPM solution will have no visibility of any attack launched to exploit the vulnerability before it has been resolved.

This lack of behavioral analysis of the cloud environment limits observability for any initiated attacks. Furthermore, it creates a gap in the runtime protection required to prevent the exploitation of cloud-based services.

CDR solutions plug this gap by providing observability within the runtime to provide immediate detection when multiple correlated events occur. In addition, they allow security analysts to identify how an attack started and provide visibility of how it traverses through public cloud infrastructure to understand how the attack can be halted and the cloud environment remediated and restored.

The key benefits of CDR solutions are:

- Automated correlation and processing of event information from multi-cloud environments to reduce analyst workload and the associated risk of alert fatigue
- The use of intelligent machine learning-based technology to correlate event information to discover the attack sequence and deduce the actual threat
- The generation of graphical attack storylines to show attack propagation along the end-to-end threat chain and reduce the forensic analysis burden on security specialists
- The correlation of multiple malicious behavior indicators into a coherent attack story initiates an incident response only once the attack risk level exceeds the risk appetite of the organization to reduce workload by reducing the number of alerts and false positives to allow them to focus on critical incidents
- Provide threat detection for suspicious activity and across dynamic contexts using runtime behavioral monitoring, including:
    - Monitoring of geographic location on first-time API usage
    - Detection of abnormal behavior for nonhuman identities, including machine roles, function roles, cross-account roles, and other role types
    - Detection of abnormal North-South network communications
    - Detection of abnormal East-West network communications
    - Detection of irregular access to S3 storage
- Provide the ability to detect and correlate anomalous events from large data sets associated with dynamic, scalable multi-cloud environments of large enterprise accounts.

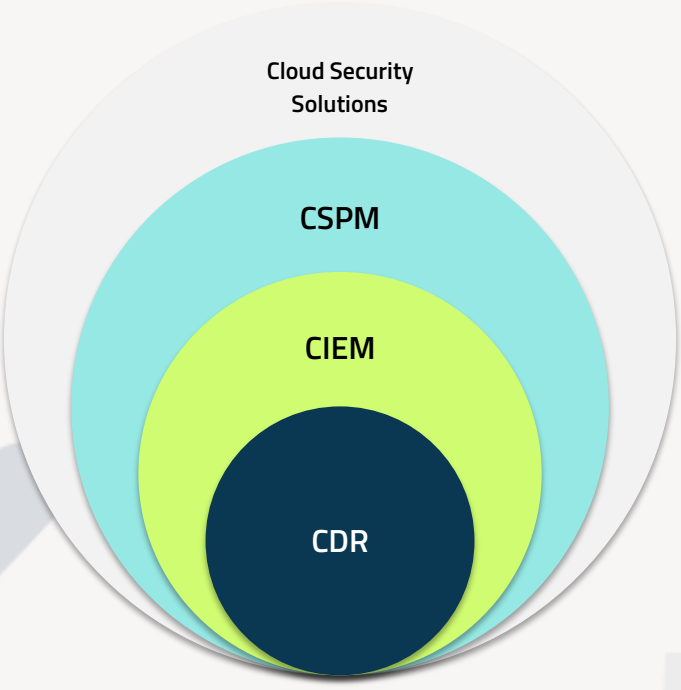

Figure 14 – Cloud Security Solutions

CDR Tools and Platforms

CDR tool operation revolves around detecting malicious behavior indicators (MBIS). Machine learning and historical rules-based detectors continuously analyze user and entity activities to identify actions that deviate from baseline behavior or increase risk above an acceptable level. Identifying any anomalous or suspicious event will generate an MBI and trigger an investigation.

Tools that correlate individual MBI into a coherent sequence of unexpected behavior support investigation activities and provide enhanced visibility into an attacker's actions within the cloud environment. These sequences also make identifying false positive alerts easier and quicker for analysts to concentrate resources on true positive ones.

The MBIs sequences should include details of any correlation with the tactics and techniques of the MITRE ATT&CK framework to allow analysts to more readily and speedily understand how the attack is being undertaken and what resources are being exploited to support faster and more effective response.

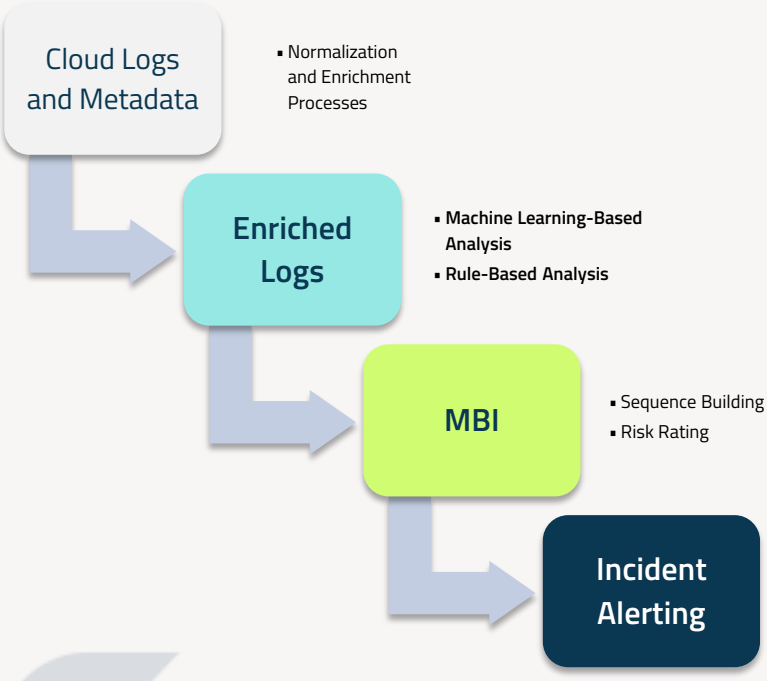

Figure 15 – CDR Detection Process

CDR tools ingest metadata and log information from cloud infrastructure to identify abnormal behavior and create a timeline of a possible incident.

Metadata is the descriptive inventory data for the cloud services consumed at any time. This includes data collected from the activities of users and groups, workloads information, managed services usage, PaaS status, network definitions, and permissions management as examples. Each type of data is collected at a periodicity specific to the nature of the data to balance the collection of a comprehensive data set against the storage and processing constraints of big data.

The metadata is collected via the public cloud interfaces, including command line interfaces (CLI) and API, and stored in a CDR-managed database. This collated data is available for use by suites of security solutions, including CSPM and CIEM services, to provide integrated security monitoring. In addition, the data has applications in threat and misconfiguration detection, permission hardening, public exposure monitoring, and general visibility purposes.

Log data is the transactional event-based telemetry available for the cloud services, including activity logs, flow logs, storage logs, and DNS Query logs. The telemetries provide granular visibility to critical events within the cloud environment, including east-west and north-south network transactions, API activity, data access, and identity administration. In addition, analysis of log data supports threat detection and monitoring for over-permission.

The collected and collated metadata and log information can be aggregated, enriched, and analyzed using machine learning-based processes within the CDR platform to identify abnormal behavior and recommend remediation strategies. The critical information security services that CDR tools support include the following:

- Misconfiguration detection and resolution are performed following any significant system change and as a frequently run process to minimize vulnerability exploitation risks
- Permission hardening to implement the least privilege principle by monitoring for over-permission risks as a frequently run process
- Rule-based behavioral monitoring at each level of the cloud environment, including IaaS, PaaS, and SaaS, to generate MBI. Rules can be a simple binary test for a condition or a complex algorithm with temporal and context-specific parameters. Process execution frequency will reflect the rate at which behavioral telemetry is updated and balanced against the time at risk between a malicious event occurring and its detection
- Machine learning-based detection of anomalous behavior using advanced complex alert criteria based on learned and predictive processes. Process execution frequency will reflect the rate at which behavioral telemetry is updated in line with rule-based detection processes
- Sequence-building processes correlate MBIs with known current attack sequences or into new attack sequences when new MBIs are detected

- Alert prioritization processes assess attack sequences to determine the risk level and deliver prioritized reporting to security teams depending on the type of incident. These can range from the identification of misconfiguration or public exposure vulnerabilities to detected threats from ongoing attacks

CDR tools should be capable of integrating into an organization's information systems, including messaging systems, productivity tools, development environment management tools, and cloud service administration consoles.

CDR Deployment

CDR Deployment Principals

CDR deployment should protect cloud-based information and services by detecting and blocking the exploitation of vulnerabilities and evasive attacks in the cloud environment. It extends traditional threat monitoring and hunting processes across cloud resources using techniques that scale with services, so the security solution seamlessly keeps pace with the provisioning of new workloads.

CDR deployment should be capable of building and reporting MBI timelines across the entire cloud ecosystem by integrating activity and audit logs with metadata from cloud services and applications with endpoint and network traffic data. In addition, the CDR solution should block attacks targeting cloud services while gathering context information for detection and response across cloud environments.

The CDR deployment model uses elements of traditional security solutions, including endpoint detection and response (EDR) and network detection and response (NDR) in a cloud-centric application. EDR solutions focus on on-premises workload security, while NDR solutions focus on network activity. CDR deployment uses these techniques within the cloud environment using cloud provider services and APIs

CDR Deployment Challenges

The challenge for CDR is that traditional security solutions, such as EDR, primarily focus on physical and virtual endpoints and servers. The techniques used by EDR cannot be readily applied to equivalent cloud-based workloads, which can exist as dynamically instantiated virtual machines, containers, or serverless applications. The cloud-based services offer attackers a significantly different attack surface and exploitable weaknesses than traditional on-premises infrastructure. CDR leverages large-scale event processing and automated risk analysis using machine learning and analytics to manage the scale of the security challenge.

CDR solutions cover all aspects of cloud security visibility to implement detection and response for all workloads, storage mechanisms, and information flows across the entire cloud environment. In addition, CDR protection covers virtual machines, serverless workloads, and container-based services using frictionless automated deployment and scaling technology across all cloud service and deployment models.

A benefit for CDR solutions is access to the software-defined cloud infrastructure with its APIs that facilitate the automation of detection and response capabilities. This allows CDR solutions to continuously evaluate workload events and network traffic, detecting MBI and creating alerts to trigger automated response actions.

CDR solutions can quarantine suspicious and anomalous workloads, adapt network access controls, isolate cloud-based assets, and retire and redeploy new instances of workloads using securely managed images to manage threats.

Cloud Incident Response

Cloud incident response follows the same principles as traditional security response processes but with subtle differences that significantly impact effectiveness.

The incident response includes the plans, processes, and controls that help organizations prepare for, detect, analyze, and recover from an incident. The challenge for cloud-based security controls is that not all elements of the cloud services are within the control of the user of the services, which impacts incident response.

The shared responsibility model sets out which security elements of the cloud services are the responsibility of the user and which will remain the responsibility of the cloud services provider.

The impact is best demonstrated by considering the difference between SaaS and IaaS. A security incident due to vulnerability within a SaaS product will be mainly invisible to the users due to the application's lack of visibility and telemetry. Users will only see the consequences of an attack where it impacts the availability of the service or compromises the data that the service can access. In these cases, the cloud services provider must have responsibility for incident response.

On the other hand, the user of an IaaS product will have direct control of assets and objects managed by that service and will have visibility of any attack. By contrast, the cloud service provider will have minimal visibility of the user instance of the IaaS. In these cases, the service user must have responsibility for incident response.

Incident Response Frameworks

Cloud incident response frameworks are available to guide effective CDR deployment and typically encompass four critical phases as follows:

Preparation

The incident preparation phase covers implementing security controls, including technology-based security solutions, developing and reviewing policies and procedures, undertaking staff training and awareness, and producing a comprehensive set of playbooks. This element of incident response should deliver all the resources required by the security team to handle any cloud-related incident.

Detection

The detection phase covers monitoring the cloud environment for potential indicators of attack and other MBI and tracking possible precursors, such as new vulnerabilities and attack vectors, before appropriate mitigations have been implemented. CDR solutions allow automation of the detection phase to support agile security teams.

Analysis

Data from the detection phase is triaged and assessed to determine if prioritized alerts from the CDR solution represent a valid indicator of attack or other incident types that requires a response. The timeline information generated by the CDR solution allows the security team to make an evidence-based decision on initiating an incident response.

Response

The response phase covers containment, eradication, and recovery from an attack. The containment process focuses on halting the spread of an attack and preventing its impact on business systems from increasing. In the cloud environment, this element can include moving business processes to a different instance of the cloud environment or isolating and quarantining affected assets and maintaining operations within the existing cloud environment.

The eradication process involves removing the incident's root cause from the affected assets, such as disabling a compromised account or removing and redeploying a new instance of a malware-infected workload.

The recovery process covers the steps necessary to resume normal business operations in the cloud environment.

Post-Incident Review

This post-incident review phase covers the lessons learned element of incident recovery as part of a continuous improvement program to resolve any issues in the planning, procedures, or playbooks used for the incident response.

CDR Deployment Best Practices

Develop a clear and well-defined security strategy before deploying a CDR solution. This process includes fully understanding the credible threats to the cloud environment and maintaining this as the threat landscape evolves and the cloud environment scales.

Create an incident response plan covering all cloud-related assets, including infrastructure, platform, applications, and data. The incident response process should include clear and well-defined procedures for responding to security threats, including breach notification processes where regulated or sensitive information is held or processed by cloud-based services.

Metadata and log information used by CDR solutions should be protected from malicious deletion or modification using secure write-once storage mechanisms. Secure retention and recovery are vital where the data used by the security solution resides in the same cloud environment that the solution is protecting. This approach will prevent attackers from hiding evidence of their presence and activities in a compromised cloud environment by altering log files to delete evidence of their actions.

Ensure cloud guardrail services and cloud-wide logging features offered by the cloud service provider are enabled to maximize visibility and monitoring capabilities before implementing a CDR solution.

Ensure that the security team responsible for managing the cloud environment's security has the necessary skills and expertise to detect and respond to threats in the cloud effectively. This includes continuous training programs to maintain currency and awareness of evolving threats.

BEST PRACTICES **FOR CLOUD SECURITY**

Overview

Cloud security can appear daunting for organizations migrating their on-premises information systems into the cloud environment. While the majority of the underlying security principles remain unchanged, implementing security controls can be significantly different and more challenging. However, the following best practices provide a foundation for implementing adequate cloud security.

Identity and Access Management

Adopting a Zero Trust approach to identity and access management will protect against compromising access credentials, one of cloud services' most common attack vectors. This approach assumes that credentials have been compromised and applies protections to counter this attack.

All identities and devices accessing cloud services are verified at every interaction with services, and the least-privilege access principle is applied. Conditional access policies can reduce the authentication overhead to limit any identity verification to high-risk processes such as access to sensitive information or critical business workloads. All activities are monitored and logged, and suspicious or abnormal behavior is identified and analyzed to deduce if an attack is in progress.

Multi-factor authentication should be enforced for all user accounts using independent factors that cannot be compromised with a single action. For example, an authenticator app on a mobile device should not be verified using an email or SMS message sent to the same device to prevent theft of the device, compromising all elements of the multi-factor authentication process.

Security Posture Management

Automated processes and tools that identify and manage risks across the cloud infrastructure are necessary to ensure an organization establishes and maintains an adequate security posture in a dynamic cloud environment. Infrastructure, where services can be deployed and scaled in seconds, will face a changing threat landscape and the risks associated with the services.

Security Control Strategy

Cloud-based infrastructure, platforms, applications, and data should be protected using comprehensive and coherent security controls that provide layered protection. As a result, the failure

of any one control will not compromise the assets it protects. In addition, the defense-in-depth principle should be applied to protect identities, hosts, networks, and data stores.

Security controls need to reflect the shared responsibility model of the cloud services and ensure complete coverage of all elements that fall under the user's responsibility. This is critical for IaaS solutions, where the user responsibilities are the greatest.

Encryption of all cloud-based data at rest and in transit is essential in all cloud deployments, though it's most critical in the multi-tenant environments of public cloud deployments.

Threat Protection

Threat detection should be adequate for all elements of a cloud ecosystem, including virtual machines, containers, databases, storage, and other resources the cloud service provider provides. In addition, threat detection should be proactive to counter advanced persistent threats and scalable to keep in step with the dynamic nature of the cloud environment.

Threat intelligence services should inform operational security processes to protect, detect and respond to threats to ensure that new and evolving threats can be countered and previously unknown threats within the cloud environment can be detected at the earliest opportunity.

Network Security

Distributed denial-of-service (DDoS) is a critical threat to cloud services with its ability to deny the availability of the entire cloud environment by targeting network connectivity. Therefore, network security controls should protect the cloud environment from malicious traffic targeting the application and network layers to maintain availability.

Identity and access management processes provide the first layer of defense for network access but need to be supported with additional controls such as firewall technology to provide strength in depth. Flat network structures within the cloud environment also enable attackers to move laterally across services once a service has been compromised. Virtual networking, subnet provisioning, and IP addressing techniques will enhance security posture while micro-segmented networks support zero-trust networking principles to contain threats.

Principles

The following fundamental principles underpin cloud security and provide a mechanism to assess the security posture of cloud environments.

Protection of Data in Transit

All data should be adequately protected against alteration or eavesdropping as it transits networks connecting the user to the cloud deployment and between services within the cloud environment. Protection can employ a combination of encryption, authentication, and network-level protections.

Data Protection and Resilience

All cloud-stored data and the cloud services that hold and process this data should be protected against the data's loss, damage, or unauthorized disclosure. In addition, protection should meet all necessary requirements when data is subject to regulatory compliance.

Customer Separation

Effective security boundaries for applications, data storage, and network management should ensure no malicious or compromised cloud service customer can access or affect another customer's service or data.

Governance Frameworks

Any cloud service provider should have a security governance framework that coordinates and directs its management of the service and information within it.

Operational Security

All cloud services should be operated and managed securely to prevent, detect and prevent attacks using a combination of vulnerability management, protective monitoring, configuration, change management, and incident management.

Personnel Security

No cloud service provider personnel should have access to customer data and services without a legitimate need. In the event of such a need, they should be able to demonstrate personal trustworthiness and be subject to technical measures that audit and constrain their actions.

Secure Development

All cloud services should be designed, developed, and deployed to minimize and mitigate security threats, including a robust software development lifecycle that uses an automated and audited integration and deployment pipeline.

Supply Chain Security

Third parties with access to customer data or the services, or where these services depend on a third party, should comply with the same security standards as the cloud service provider.

User Management

Cloud services should include tools to enable the secure management of service users that prevent unauthorized access and alteration of resources, applications, and data.

Identity and Authentication

Only securely authenticated and authorized identities should be granted access to service interfaces.

External Interfaces

All external APIs, web consoles, and command line interfaces should be protected.

Service Administration

The design, implementation, and management of administration systems for cloud services should follow enterprise best practices and protect adequate protection against attack.

Audit and Alerting

The cloud service should provide audit information and security alerts following the detection of an attack with sufficient details to allow investigation and incident response.

Secure Service Use

The cloud service provider should ensure customer meet their data protection responsibilities using the principle of security by design and by default for provided services.

CLOUD SECURITY **CASE STUDIES**

Compromised AWS Credentials

Any user or entity accessing a cloud service requires valid credentials to access the cloud infrastructure hosting the service. These access credentials are used to authenticate the identity of the user or entity. They also verify that they are authorized to use the service in the manner they are attempting. For example, a user may have a valid identity but may try to modify the information where they have read-only access. In this case, they case the authentication but fail the authorization checks.

In this case study, an attacker was able to launch a phishing attack by stealing the access credentials of an AWS Lambda entity. Additionally, an attacker was able to extract the Lambda function's environment variables and export them to use in their attack (AWS_ACCESS_KEY_ID, AWS_SECRET_ACCESS_KEY, AWS_SESSION_TOKEN).

The attack began with using the GetCallerIdentity command. This command provides information about the entity the credentials are associated with. From the response, the attacker can gain additional information, such as the account ID and the credentials type that was stolen. They cannot, however, determine anything about the privileges associated with the identity.

The credentials allowed the attacker to execute API calls by impersonating the Lambda function, enumerating different services in the cloud environment. Most API calls were blocked as the authorization checks failed because the Lambda function was not permitted to make these calls. However, the Lambda function had permission to access and instantiate the AWS Simple Email Service (SES). This provides the attacker the ability to execute VerifyEmailIdentity and UpdateAccountSendingEnabled commands. Consequently, the attacker was able to use the SES to launch a phishing attack from within the boundary of the affected organization's systems and bypass the perimeter protective mechanisms.

The attacker could also hide evidence of their activities by deleting the SES identity using the DeleteIdentity command. However, the attack was identified as the Lambda calls were associated with an IP address outside the affected organization's cloud environment, which provided an indicator of compromise for security analysts to investigate. Additionally, applying security controls to restrict access to authorized IP address ranges would have mitigated this vulnerability.

Compromised GCP Credentials

In this case study, an attacker was able to compromise the credentials for a GCP App Engine service account (SA). The App Engine is a fully managed GCP Cloud serverless platform that uses a token for the service account credentials. When the user creates an App Engine instance, the cloud provider creates a default SA and attaches it to the newly created App Engine.

User credentials are vulnerable when stored in insecure locations or when users follow poor security practices and select easily guessed or brute-forced passwords. In this case study, the stolen SA credentials were for the default account that had been left active and unchanged. This poor security practice was compounded as this was a highly privileged nature of the SA role.

An attacker was able to use the default credentials to launch an attack that resulted in the creation of a significant number of high-core CPU virtual machines (VMs) that were used to run crypto mining applications. While this attack did not adversely affect the victim organization's cloud-based services, it did result in a financial impact due to the costs of the deployed applications, which consumed significant processing resources for the time they were undetected.

The attack began with privilege escalation, which added a compute.admin role into the IAM policy for the compromised SA account. This allowed the attacker to modify the firewall rules to add a new subnet for crypto mining applications with all network restrictions removed so the VMs could undertake unrestricted operations within the cloud environment. The attacker was also able to create additional SA keys to counter any detection and revocation of the exploited default key.

One of the key indicators of compromise beyond the unexpected resource usage fees was the IP addresses used to access the App Engine SA, including active TOR exit nodes, outside the affected organization's cloud environment.

AWS Log4Shell Hot Patches

In this case study, we look at Log4Shell, which resulted in one of the most significant vulnerabilities for cloud-based services in recent times. This story starts with the original high-risk vulnerability being identified in the globally used Log4Shell logging function, which resulted in a rush to resolve the issue before attackers could exploit this remote code execution (RCE) vulnerability.

AWS produced hot patch solutions covering different environments as part of the remediation process. These open-source releases were intended as short-term solutions to mitigate the risk until the Log4Shell application could be corrected and released to users.

- A hot patch Daemonset was released for Kubernetes clusters.
- A hot patch solution named Hotdog was bundled as a set of OCI hooks for Bottlerocket hosts.
- A hot patch service was bundled in an RPM package and automatically installed with AWS Linux JDK (Java) packages.

These hot patches were not exclusive to AWS environments, covering the majority of affected cloud and on-premises environments, including Kubernetes and ECS clusters, Fargate containers, and standalone servers.

Hot patching is the process of injecting a small section of modified code into a running application to fix a localized problem in the application without stopping and restarting the system.

However, security researchers found that the hot patches contained a high-risk vulnerability that could be exploited for container escape and privilege escalation. Following the installation of the hot patch, any new container had the ability to compromise the underlying host. In some cases, existing containers also inherited this capability and could be exploited for malicious use.

Containers are employed to create security boundaries between applications running on the same machine. Container escape vulnerability will allow an attacker to extend the attack across containers and move laterally to neighboring services. For Kubernetes clusters, container escape vulnerabilities can allow an attacker to compromise the entire cluster.

An additional impact was that an attacker could use an unprivileged process to escalate privileges and gain complete control over their underlying server. The result was that exploitable processes could perform root code execution.

This case study highlights the impact of unintentional consequential effects of mitigating security vulnerabilities. The urgency for Log4Shell users to mitigate the initial high-risk vulnerability led to the deployment of hot patches at scale into the live environment without testing. Including a new high-risk vulnerability in the hot patches puts the container environments at risk. Once the updated Log4Shell application was released, users were not incentivized to remove the hot patch, maintaining the presence of the vulnerability.

AWS has now mitigated these vulnerabilities and released a fix for each solution.

LMNTRIX **CLOUD XDR**

Overview

LMNTRIX XDR Cloud Delivers Multi-Vector Managed Detection and Response for Multi-Cloud Workloads

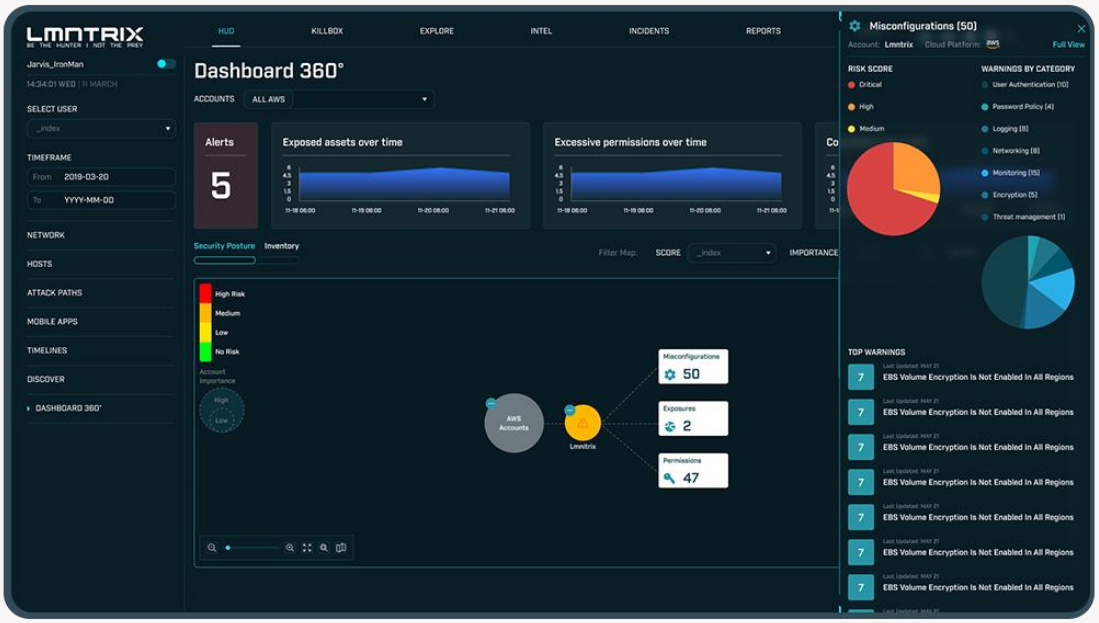

Figure 16 - LMNTRIX CLOUD XDR

Security Operations Center (SOC) and Cloud SecOps teams are overwhelmed with investigating and closing every "alert", most of which are not actual threats or incidents. LMNTRIX XDR CLOUD contextualizes the cloud, application, and user behavior in your environment and creates an attack storyline to identify actual threats. By focusing on actual threats in the runtime, LMNTRIX improves productivity and morale of the SOC and reduces business risk.

The LMNTRIX XDR Cloud combines Cloud Security Posture Management (CSPM) with Identity Threat Detection and Response (ITDR) and our runtime unique Cloud Threat Detection and Response (CDR) to monitor and analyze all threat vectors including excess privileges and identity risks, exposed assets, and behavioral anomalies to identify and respond to real-time risk.

- Reduce the cloud attack surface by understanding all permissions and their use so unused permissions can be safely eliminated without impacting productivity

- Centralized runtime analysis and observability identifies risky behaviors to stop threats before they become a newsworthy incident
- Respond to actual incidents with contextualized models built to alert on atypical behaviors for your specific cloud, application, and users.
- Resolve issues fast with a full understanding of exactly how threat actors penetrated the environment with attack sequences
- Support internal and external compliance initiatives with automated review
- Identify privilege escalation and other malicious uses of unsuspecting identities

Runtime Observability

While most cloud security tools analyze static points in time, the LMNTRIX XDR Cloud continuously monitors and analyzes configurations, permissions, and behavior in the environment providing constant vigilance and real-time observability. Particularly important for addressing misconfigurations that may take time to address or simply cannot be addressed without taking down the environment, this level of monitoring reduces the misconfiguration risks your organization is forced to tolerate.

Identifying real alerts with MBIs and the Attack Sequence and eliminate AlertFatigue.

LMNTRIX leverages malicious behavior indicators (MBIs) sequenced together to indicate an actual threat. MBIs are behaviors and activities that are detected from the metadata and logs we collect from the cloud environment. Over a period of time, a string of MBIs can reveal an attack sequence. Our Machine Learning (ML) will score these activities and create alerts for the LMNTRIX team to investigate and validate

The MBIs are sequenced into a storyline, which is then scored. The sequence is continuously evaluated as new MBIs are added to the sequence and scored. Once the score of the sequence is higher than 7, you get a realert and you know you need to investigate. This ensures you respond to the riskiest alerts first.

The LMNTRIX XDR attack sequence provides a complete overview to your team of how the attacker got in and then moved around the organization. Your team, regardless of experience level, can easily identify the vulnerabilities and take steps to close these gaps in real-time.

The LMNTRIX XDR AI and ML are different, really.

LMNTRIX leverages advanced machine learning (ML) techniques and artificial intelligence (AI) to build models for ongoing behavioral analysis of the runtime for more accurate threat detection. Models are trained and tuned every day on customer data at a global level. The output of the models are reviewed by security experts to ensure models are accurate. For more information, check out our blog, "The Science Behind Our Security".

Eliminating Excessive Permissions

Access to cloud resources is granted through permissions and in order to ensure employee productivity is not impacted, many cloud teams will grant permissions broadly. Employees may only use one or two of these permissions to access the information they need, meaning the other permissions assigned to them are unused. This expands the attack surface and provides additional tools for hackers to use to penetrate your cloud.LMNTRIX XDR CLOUD automatically detects all the permissions that are assigned to users/groups/roles and analyzes their usage. This information is presented to your team so they can revoke permissions if needed and reduce the attack surface.

Effectively Implement Security Best Practices

Managing security processes to ensure that your cloud security framework adheres to best practices is not an easy task. LMNTRIX helps ensure compliance with one-click reporting across a variety of common industry standards, with detailed visual reports on where you are successful and where you need to do some work. An additional layer of protection is delivered with custom governance enforcement via a query language for custom rules.

LMNTRIX's XDR Cloud continuously monitors and alerts on potentially dangerous misconfigurations such as public exposure of assets,authentication misconfigurations, password policy, logging, networking, monitoring, and encryption. Alerts deliver granular details including the affected assets, users, and the compliance rules being violated.

User-Defined Automated Response

Automated response within LMNTRIX XDR CLOUD addresses compliance/governance issues (aka misconfigurations), public exposures, and realert. You define rules within the LMNTRIX XDR CLOUD platform and only execute them if all the criteria are met.

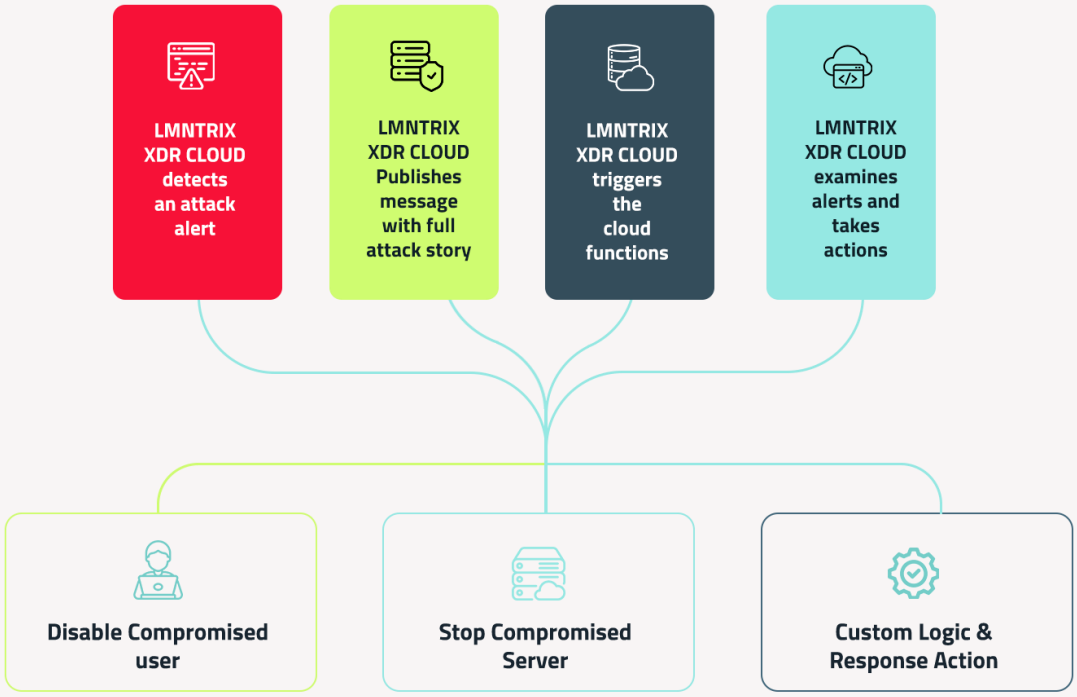

Figure 17 - User-Defined Automated Response

Security analysts create remediation rules with specific configurations to resolve various issues such as public exposures or misconfigurations. Auto-remediation is an optional capability, and is rule based allowing you to decide if and when to execute remediation actions.

A Complete Picture Of Cloud Risk

LMNTRIX XDR CLOUD takes a comprehensive approach to threat detection and response. It analyzes the configuration of your environment and all enabled permissions to fully assess the attack surface. It then analyzes the runtime within the environment, so your team can see which misconfigurations orpermissions are compromised and are being exploited – shutting down these vulnerabilities is where the your team needs to start. This is the level of security analysis that LMNTRIX XDR CLOUD provides. Understanding where your vulnerabilities are, clarifies how and when your vulnerabilities are being exploited.

LMNTRIX XDR CLOUD not only provides a complete analysis of static configurations, it reveals risky behavior in the runtime so you can stop attacks in their tracks fast. AI and ML driven models create the right context for your business for your cloud, applications, and users, so that alerts are actual threats and not just anomalies or one-offs. The overall productivity of the team is greatly improved as they are focusing on real alerts, and not chasing random activities

ABOUT **LMNTRIX**

Overview

Your cloud security solution should provide peace of mind that adopting cloud-based services will not compromise your sensitive information, business processes, or other corporate systems. Typically, the difference between preventing a cyber-attack and falling victim is understanding how and where to seek compromise indicators. Even the most advanced attackers leave traces of their presence, so an effective defense must be vigilant and ever-adaptive in response to changes in attacker tactics. A critical element of constantly evolving threats in this age is a detailed view of an organization's potential attack surface, including all cloud-based aspects.

Unfortunately, adopting cloud-based services limits the visibility of activities and places limitations on log collection. This prevents traditional security solutions developed for on-premises environments from collecting and analyzing data in a manner that allows rapid attack detection when faced with today's advanced threat actors.

LMNTRIX has reimagined cybersecurity, once again turning the tables in favor of the defenders. We have cut out the bloat of SIEM, log analysis, false positives, and associated alert fatigue and created new methods for confounding even the most advanced attackers. We combine deep expertise with cutting-edge technology, leading intelligence, and advanced analytics to detect and investigate threats with great speed, accuracy, and focus. We believe that in a time of continuous compromise, you need a continuous response – not an incident response.

As a company, we stand in defiance of the unwanted human presence within corporate networks by attacking the root of the problem—the adversary's ability to gain entry and remain undetected. Our real-time hunt operations identify signs of planned and active attacks and take action to neutralize them, forming the basis of our comprehensive Active Defense approach to limiting security exposure.

We are a partner who becomes an extension of your internal team, can augment your MSSP, or be a full-service SOC as a cloud services security solution.

LMNTRIX Active Defense

LMNTRIX Active Defense is a three-tier outcome-based solution (The industry refers to it as Managed Detection & Response (MDR) and our platform as Extended Detection & Response (XDR).

1) LMNTRIX XDR (AWS Data Lake and Platform)
2) LMNTRIX TECHNOLOGY STACK (Deployed deep within Customer Networks)
3) LMNTRIX CYBER DEFENSE CENTRE (Security Analyst Driven).

LMNTRIX XDR natively unifies Machine and Underground Intelligence, NGAV, EDR, NDR, UEBA, and Deception Everywhere with completely automated attack validation, investigation, containment, and remediation on a single, intuitive platform. Backed by a 24/7 Managed Detection and Response service at no extra cost, LMNTRIX provides comprehensive protection of the environment for even the smallest security teams. In addition, it is a single investigative platform for insights into threats on enterprise, cloud, hybrid, and industrial control systems (ICS) networks.

The LMNTRIX XDR delivers unique advantages over current network security solutions. It is a holistic and multi-vector platform with an unlimited retention window of full-fidelity network traffic, innovative security visualizations, and the ease and cost-savings of an on-demand deployment model.

LMNTRIX XDR is based on multiple detective, responsive, and predictive capabilities that integrate and share information to build a security protection system that is more adaptive and intelligent than any one element. The constant exchange of intelligence between the Active Defense components and the broader cybersecurity community enables LMNTRIX to keep abreast of the TTP of the most persistent, well-resourced, and skilled attack groups.

Figure 18 - LMNTRIX XDR

LMNTRIX Tech Stack

The LMNTRIX Tech Stack is a powerful, proprietary threat detection stack embedded within the client environment behind existing controls. TECHNOLOGY STACK comprises multiple detective systems, combining contextual threat intelligence and correlation, static-file analysis, user and entity behavior analytics (UEBA), and anomaly detection techniques to find threats in real time. In addition, it eliminates alert fatigue, determining which alerts to escalate through multi-platform consensus.

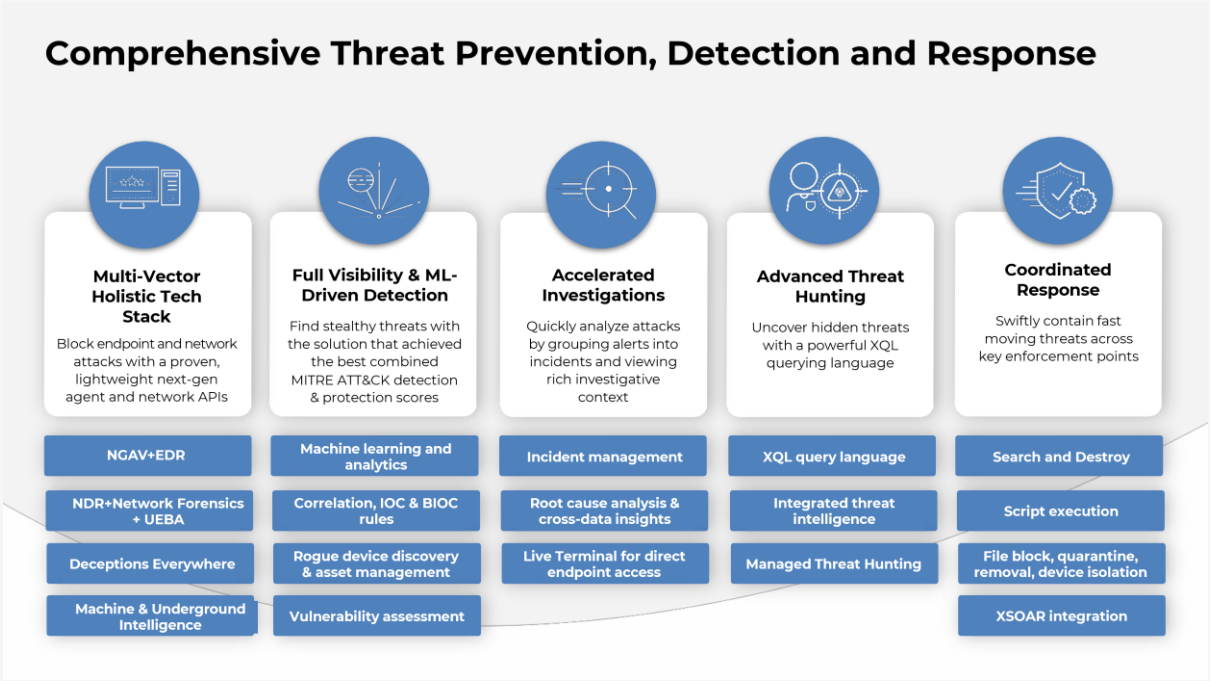

Figure 19 - LMNTRIX XDR Features

LMNTRIX Cyber Defense Centers

LMNTRIX employs a global network of Cyber Defense Centers (CDC) comprising trained and certified security analysts and provides constant vigilance and on-demand analysis of your cloud-based digital assets and networks. Our intrusion analysts actively probe and monitor your networks and endpoints 24x7, using the latest intelligence and proprietary methodologies to look for signs of compromise. When a suspected breach is detected, the team performs an in-depth analysis of potentially affected systems to confirm the breach. Additionally, when data theft or lateral movement is imminent, our endpoint containment feature makes immediate action possible by quarantining affected hosts, whether on or off your corporate network. This significantly reduces or eliminates the consequences of a breach.

Figure 20 - LMNTRIX Cyber Defense Centre