

Understanding the ASD's Cloud Computing Security for Tenants in the Context of AWS

June 2017



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

Introduction	1
AWS Shared Responsibility approach to Managing Cloud Security	2
What does the shared responsibility model mean for the security of customer content?	3
Understanding ASD Cloud Computing Security for Tenants in the Context of AWS	4
General Risk Mitigations	4
IaaS Risk Mitigations	27
PaaS Risk Mitigations	38
SaaS Risk Mitigations	40
Further Reading	41
Document Revisions	42

Introduction

The Australian Signals Directorate (ASD) publishes the *Cloud Computing Security for Tenants* paper to provide guidance on how an organisations’ cyber security team, cloud architects and business representatives can work together to perform a risk assessment and use cloud services securely. The paper highlights the shared responsibility that organisations (referred to as Tenants) share with the cloud service providers (CSP) to design a solution that uses security best practices.

This document addresses each risk identified in the Cloud Computing Security for Tenants paper, and describes the AWS services and features that you can use to mitigate those risks.

Important: You should understand and acknowledge that that the risks discussed in this document cover only part of your responsibilities for securing your cloud solution. For more information about the [AWS Shared Responsibility Model](#), see [AWS Shared Responsibility Approach to Managing Cloud Security](#) below.

AWS provides you with a wide range of security functionality to protect your data in accordance with ASD’s Information Security Manual (ISM) controls, agency guidelines and policies. We are continually iterating on the security tools we provide our customers, and regularly release enhancements to existing security functionality. AWS has assessed ASD’s ISM controls against the following services:

- **Amazon Elastic Compute Cloud (Amazon EC2)** – Amazon EC2 provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. For more information, go [here](#).
- **Amazon Simple Storage Service (S3)** – Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. For more information, go [here](#).

- **Amazon Virtual Private Cloud (VPC)** – Amazon VPC provides the ability for you to provision a logically isolated section of AWS where you can launch AWS resources in a virtual network that you define. For more information, go [here](#).
- **Amazon Elastic Block Store (EBS)** – Amazon EBS provides highly available, highly reliable, predictable storage volumes that can be attached to a running Amazon EC2 instance and exposed as a device within the instance. For more information, go [here](#).

Important: AWS provides many services in addition to those listed above. If you would like to use a service not listed above, you should evaluate your workloads for suitability. Contact [AWS Sales and Business Development](#) for a detailed discussion of security controls and risk acceptance considerations.

Our global whitepapers have recommendations for securing your data that are just as applicable to Australian government workloads on AWS. For a complete list of our security and compliance whitepapers, see the [AWS Whitepapers website](#).

Our [AWS Compliance](#) website contains more specific discussions of security, AWS Risk and Compliance practices, certifications, and reports.

If you need answers to questions that are not covered in the above resources, you can contact your account manager directly.

AWS Shared Responsibility approach to Managing Cloud Security

When you move your IT infrastructure to AWS, you will adopt a model of shared responsibility between you and AWS (as shown in Figure 1). This shared model helps relieve your operational burden because AWS operates, manages, and controls the IT components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate.

As part of the shared model, you are responsible for managing the guest operating system (including updates and security patches to the guest operating system) and associated application software, as well as the configuration of the AWS provided security group firewall and other security-related features. You will also generally connect to the AWS environment through services that you acquire from third parties (for example, internet service providers). As AWS does not provide these connections, they are part of your area of responsibility. You should consider the security of these connections and the security responsibilities of such third parties in relation to your systems.

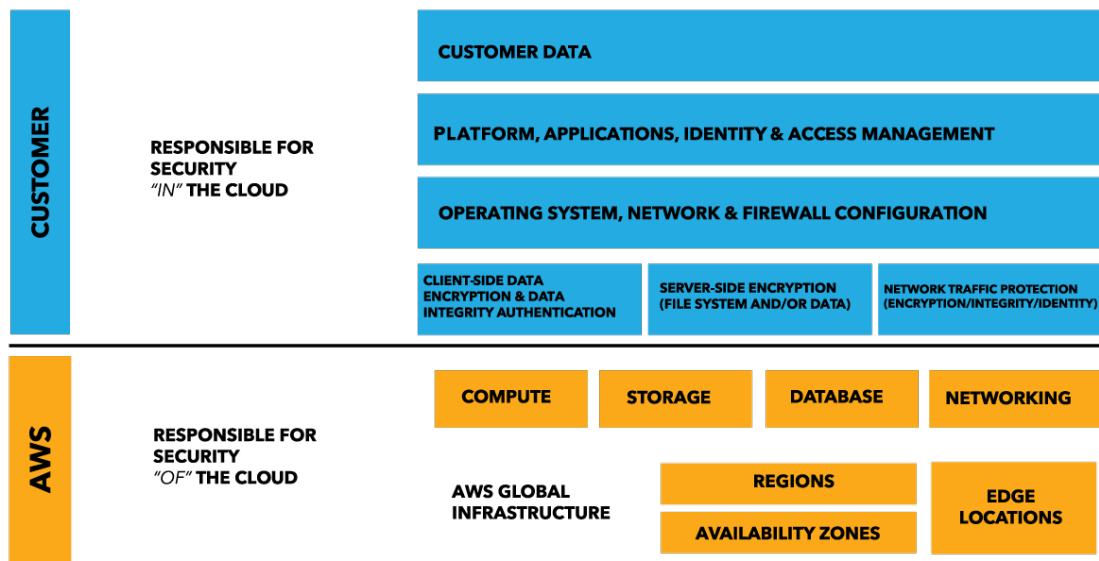


Figure 1: The AWS Shared Responsibility Model

What does the shared responsibility model mean for the security of customer content?

When evaluating the security of a cloud solution, it is important for you to understand and distinguish between:

- Security measures that AWS implements and operates – “security *of* the cloud”
- Security measures that you implement and operate, related to the security of your content and applications that make use of AWS services – “security *in* the cloud”.

While AWS manages the security *of* the cloud, security *in* the cloud is your customer responsibility, as you retain control of what security you choose to implement to protect your own content, platform, applications, systems and networks – no differently than you would for applications in an on-site data centre.

Understanding ASD Cloud Computing Security for Tenants in the Context of AWS

The following sections describe the AWS compliance and AWS offerings that can help you, as the Tenant, mitigate the risks identified in the Cloud Computing Security for Tenants paper.

General Risk Mitigations

1 – General

Requirement

Use a cloud service that has been assessed, certified and accredited against the ISM at the appropriate classification level, addressing mitigations in the document Cloud Computing Security for Cloud Service Providers.

AWS Response

An independent IRAP assessor examined the controls of in-scope AWS services’ people, process, and technology to ensure they address the needs of the ISM. AWS has been certified for Unclassified DLM (UD) workloads by the Australian Signals Directorate (ASD) as the Certification authority and is an inaugural member of the [ASD Certified Cloud Services List \(CCSL\)](#).

2 – General

Requirement

Implement security governance involving senior management directing and coordinating security-related activities including robust change management, as well as having technically skilled staff in defined security roles.

AWS Response

AWS customers are required to maintain adequate governance over the entire IT control environment regardless of how IT is deployed. This is true for both on premise and cloud deployments. Leading practices include:

- Develop an understanding of required compliance objectives and requirements (from relevant sources)
- Establish a control environment that meets those objectives and requirements
- Understand the validation required based on the organization’s risk tolerance.
- Verify the operating effectiveness of their control environment.

AWS provides options to apply various types of controls and verification methods.

Strong customer compliance and governance might include the following basic approach:

1. Review information available from AWS together with other information to understand as much of the entire IT environment as possible, and then document all compliance requirements.
2. Design and implement control objectives to meet the enterprise compliance requirements.
3. Identify and document controls owned by outside parties.
4. Verify that all control objectives are met and all key controls are designed and operating effectively.

Approaching compliance governance in this manner will help you gain a better understanding of your control environment, and will help you clearly delineate the verification activities that you need to perform.

You can run nearly anything on AWS that you would run on premise, including websites, applications, databases, mobile apps, email campaigns, distributed data analysis, media storage, and private networks. AWS provides services that are designed to work together so that you can build complete solutions. An often overlooked benefit of migrating workloads to AWS is the ability to achieve a

higher level of security, at scale, by utilizing the many governance-enabling features offered. For the same reasons that delivering infrastructure in the cloud has benefits over on-premise delivery, cloud-based governance offers a lower cost of entry, easier operations, and improved agility by providing more oversight, security control, and central automation.

The [Governance at Scale](#) whitepaper describes how you can achieve a high level of governance of your IT resources using AWS.

3 – General

Requirement

Implement and annually test an incident response plan covering data spills, electronic discovery, and how to obtain and analyse evidence, e.g. time-synchronised logs, hard disk images, memory snapshots and metadata.

AWS Response

AWS recognizes the importance of customers implementing and testing an incident response plan. Using AWS, you can requisition compute power, storage, and other services in minutes and have the flexibility to choose the development plan or programming model that makes the most sense for the problems you’re trying to solve. You pay only for what you use, with no up-front expenses or long-term commitments, making AWS a cost-effective way to deliver applications plus conduct incident response tests and simulations in realistic environments. [This](#) presentation from AWS re:Invent 2015 conference provides further details on incident response simulation on AWS.

The AWS platform includes a range of monitoring services that can be leveraged as part of your incident detection and response capability some. In-scope services include the following:

- [CloudWatch](#)
- [CloudWatch Logs](#)
- [CloudWatch Events](#)
- [Cloudtrail](#)
- [Trusted Advisor](#)
- [Elastic Load Balancer Logs](#)
- [S3 logs](#)
- [Cloudfront logs](#)
- [VPC Flow Logs](#)

- [Simple Notification Service](#)
- [Lambda](#)

4 – General

Requirement

Use ASD-approved cryptographic controls to protect data in transit between the Tenant and the CSP, e.g. application layer TLS or IPsec VPN with approved algorithms, key length, and key management.

AWS Response

AWS allows customers to use their own encryption mechanisms for nearly all the services, including S3, EBS, and EC2. IPsec tunnels to VPC are also encrypted. Customers may also use third-party encryption technologies. In addition, customers can leverage AWS Key Management Systems (KMS) to create and control encryption keys (refer to <https://aws.amazon.com/kms/>). All of the AWS APIs are available via TLS-protected endpoints which provide server authentication.

AWS cryptographic processes are reviewed by independent third party auditors for our continued compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.

For Tenants leveraging the Amazon Elastic Load Balancer in their solutions, it has security features relevant to this mitigation. Elastic Load Balancing has all the advantages of an on-premises load balancer, plus several security benefits:

- Takes over the encryption and decryption work from the Amazon EC2 instances and manages it centrally on the load balancer
- Offers clients a single point of contact, and can also serve as the first line of defense against attacks on your network
- When used in an Amazon VPC, supports creation and management of security groups associated with your Elastic Load Balancing to provide additional networking and security options
- Supports end-to-end traffic encryption using TLS (previously SSL) on those networks that use secure HTTP (HTTPS) connections. When TLS is used, the TLS server certificate used to terminate client connections can be managed centrally on the load balancer, rather than on every individual instance.

HTTPS/TLS uses a long-term secret key to generate a short-term session key to be used between the server and the browser to create the ciphered (encrypted)

message. Amazon Elastic Load Balancing configures your load balancer with a pre-defined cipher set that is used for TLS negotiation when a connection is established between a client and your load balancer. The pre-defined cipher set provides compatibility with a broad range of clients and uses strong cryptographic algorithms. However, some customers may have requirements for allowing only specific ciphers and protocols (such as PCI, SOX, etc.) from clients to ensure that standards are met. In these cases, Amazon Elastic Load Balancing provides options for selecting different configurations for TLS protocols and ciphers. You can choose to enable or disable the ciphers depending on your specific requirements.

To help ensure the use of newer and stronger cipher suites when establishing a secure connection, you can configure the load balancer to have the final say in the cipher suite selection during the client-server negotiation. When the Server Order Preference option is selected, the load balancer will select a cipher suite based on the server’s prioritization of cipher suites rather than the client’s. This gives you more control over the level of security that clients use to connect to your load balancer.

For even greater communication privacy, Amazon Elastic Load Balancer allows the use of Perfect Forward Secrecy, which uses session keys that are ephemeral and not stored anywhere. This prevents the decoding of captured data, even if the secret long-term key itself is compromised.

5 – General

Requirement

Use ASD-approved cryptographic controls to protect data at rest on storage media in transit via post/courier between the tenant and the CSP when transferring data as part of on-boarding or off-boarding.

AWS Response

[Snowball](#) is a petabyte-scale data transport solution that uses secure appliances to transfer large amounts of data into and out of the AWS cloud. Using Snowball addresses common challenges with large-scale data transfers, including high network costs, long transfer times, and security concerns. Transferring data with Snowball is simple, fast, secure, and can be as little as one-fifth the cost of high-speed Internet.

Snowball encrypts all data with AES-256-bit encryption. You manage your encryption keys by using the AWS Key Management Service (AWS KMS). Your keys are never sent to or stored on the appliance. Further details on the [AWS KMS](#) are available in this [paper](#).

In addition to using a tamper-resistant enclosure, Snowball uses an industry-standard Trusted Platform Module (TPM) with a dedicated processor designed to detect any unauthorized modifications to the hardware, firmware, or software. AWS inspects every appliance for any signs of tampering and to verify that no changes were detected by the TPM.

When the data transfer job has been processed and verified, AWS performs a software erasure of the Snowball appliance that follows the National Institute of Standards and Technology (NIST) guidelines for media sanitization.

Snowball uses an innovative, E Ink shipping label designed to ensure the appliance is automatically sent to the correct AWS facility and which also helps in tracking. When you have completed your data transfer job, you can track it by using Amazon SNS, text messages, and the console.

6 – General

Requirement

Use a corporately approved and secured computer, multi-factor authentication, a strong passphrase, least access privileges and encrypted network traffic to administer (and, if appropriate, access) the cloud service.

AWS Response

All of the AWS APIs are available via TLS-protected endpoints that provide server authentication. For more information on our region end points, go [here](#).

AWS requires that all API requests be signed—using a cryptographic hash function. If you use any of the AWS SDKs to generate requests, the digital signature calculation is done for you; otherwise, you can have your application calculate it and include it in your REST or Query requests by following the directions in our documentation.

Not only does the signing process help protect message integrity by preventing tampering with the request while it is in transit, it also helps protect against potential replay attacks. A request must reach AWS within 15 minutes of the time stamp in the request. Otherwise, AWS denies the request.

The most recent version of the digital signature calculation process is Signature Version 4, which calculates the signature using the HMAC-SHA256 protocol

AWS Identity and Access Management (IAM) enables you to securely control access to AWS services and resources for your users. Using IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources.

To get started using IAM, go to the AWS Management Console and get started with these IAM Best Practices.

You can set a password policy on your AWS account to specify complexity requirements and mandatory rotation periods for your IAM users' passwords. You can use a password policy to do these things:

- Set a minimum password length.
- Require specific character types, including uppercase letters, lowercase letters, numbers, and non-alphanumeric characters. Be sure to remind your users that passwords are case sensitive.
- Allow all IAM users to change their own passwords.

Note: When you allow your IAM users to change their own passwords, IAM automatically allows them to view the password policy. IAM users need permission to view the account's password policy in order to create a password that complies with the policy.

- Require IAM users to change their password after a specified period of time (enable password expiration).
- Prevent IAM users from reusing previous passwords.
- Force IAM users to contact an account administrator when the user has allowed his or her password to expire.

AWS Multi-Factor Authentication (MFA) is a simple best practice that adds an extra layer of protection on top of your user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password (the first factor—what they know), as well as for an authentication code from their AWS MFA device (the second factor—what they have). Taken together, these multiple factors provide increased security for your AWS account settings and resources.

You can enable MFA for your AWS account and for individual IAM users you have created under your account. MFA can be also be used to control access to AWS service APIs.

After you've obtained a supported hardware or virtual MFA device, AWS does not charge any additional fees for using MFA.

If you already manage user identities and MFA outside of AWS, you can use IAM identity providers instead of creating IAM users in your AWS account. With an identity provider (IdP), you can manage your user identities outside of AWS and give these external user identities permissions to use AWS resources in your account. This is useful if your organization already has its own identity system, such as a corporate user directory. It is also useful if you are creating a mobile app or web application that requires access to AWS resources.

To use an IdP, you create an IAM identity provider entity to establish a trust relationship between your AWS account and the IdP. IAM supports IdPs that are compatible with [OpenID Connect \(OIDC\)](#) or [SAML 2.0 \(Security Assertion Markup Language 2.0\)](#).

The following services are relevant in the enforcing use of corporate controlled computers:

- A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. When you launch an instance in a VPC, you can assign the instance to up to five security groups. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC could be assigned to a different set of security groups. For each security group, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound

traffic. For example, you could restrict access to SSH and RDP ports to only your approved corporate IP ranges.

- A *network access control list (ACL)* is a recommended layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC. For more information about the differences between security groups and network ACLs, see [Comparison of Security Groups and Network ACLs](#).
- Permissions let you specify access to AWS resources. Permissions are granted to IAM entities (users, groups, and roles) and by default these entities start with no permissions. In other words, IAM entities can do nothing in AWS until you grant them your desired permissions. To give entities permissions, you can attach a policy that specifies the type of access, the actions that can be performed, and the resources on which the actions can be performed. In addition, you can specify any conditions that must be set for access to be allowed or denied. To assign permissions to a user, group, role, or resource, you create a policy that lets you specify:
 - **Actions** – Which AWS actions you allow. For example, you might allow a user to call the Amazon S3 ListBucket action. Any actions that you don't expressly allow are denied.
 - **Resources** – Which AWS resources you allow the action on. For example, what Amazon S3 buckets will you allow the user to perform the ListBucket action on? Users cannot access any resources that you do not explicitly grant permissions to.
 - **Effect** – Whether to allow or deny access. Because access is denied by default, you typically write policies where the effect is to allow.
 - **Conditions** – Which conditions must be present for the policy to take effect. For example, you might allow access only to the specific S3 buckets if the user is connecting from a specific IP range or has used multi-factor authentication at login.

For an example of this policy, go [here](#).

7 – General

Requirement

Protect authentication credentials, e.g. avoid exposing Application Programming Interface (API) authentication keys placed on insecure computers or in the source code of software that is accessible to unauthorised third parties.

AWS Response

When you access AWS programmatically, you use an access key to verify your identity and the identity of your applications. An access key consists of an access key ID (something like AKIAIOSFODNN7EXAMPLE) and a secret access key (something like wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY).

Anyone who has your access key has the same level of access to your AWS resources that you do. Consequently, AWS goes to significant lengths to protect your access keys, and, in keeping with our shared-responsibility model, you should as well.

The following steps can help you protect access keys. For general background, see [AWS Security Credentials](#).

Note: Your organization may have different security requirements and policies than those described in this topic. The suggestions provided here are intended to be general guidelines.

- [Remove \(or Don't Generate\) a Root Account Access Key](#)

One of the best ways to protect your account is to not have an access key for your root account. Unless you must have a root access key (which is very rare), it is best not to generate one. Instead, the recommended best practice is to create one or more AWS Identity and Access Management (IAM) users, give them the necessary permissions, and use IAM users for everyday interaction with AWS.

- [Use Temporary Security Credentials \(IAM Roles\) Instead of Long-Term Access Keys](#)

In many scenarios, you don't need a long-term access key that never expires (as you have with an IAM user). Instead, you can create IAM

roles and generate temporary security credentials. Temporary security credentials consist of an access key ID and a secret access key, but they also include a security token that indicates when the credentials expire.

- [Manage IAM User Access Keys Properly](#)

If you do need to create access keys for programmatic access to AWS, create an IAM user and grant that user only the permissions he or she needs. Then generate an access key for that user. For details, see [Managing Access Keys for IAM Users](#) in *IAM User Guide*.

Observe these precautions when using access keys:

- Don't embed access keys directly into code.
- Use different access keys for different applications.
- Rotate access keys periodically.
- Remove unused access keys.
- Configure multifactor authentication for your most sensitive operations.

- [More Resources](#)

You can also leverage AWS Trusted Advisor checks as part of your Security monitoring. AWS Trusted Advisor provides best practices in four categories:

- Cost Optimization
- Security
- Fault Tolerance
- Performance Improvement

The complete list of over 50 Trusted Advisor checks available with business and enterprise support plans can be used to monitor and improve the deployment of Amazon EC2, Elastic Load Balancing, Amazon EBS, Amazon S3, Auto Scaling, AWS Identity and Access Management, Amazon RDS, Amazon Redshift, Amazon Route 53, CloudFront, and CloudTrail. You can view the overall status of your AWS resources and savings estimations on the Trusted Advisor dashboard.

One of the Trusted Advisor checks is for exposed Access Keys. This checks popular code repositories for access keys that have been exposed to the public and for irregular Amazon Elastic Compute Cloud (Amazon EC2) usage that could be the result of a compromised access key. An access key consists of an access key ID and the corresponding secret access key. Exposed access keys pose a security risk to your account and other users, could lead to excessive charges from unauthorized activity or abuse, and violate the AWS Customer Agreement. If your access key is exposed, take immediate action to secure your account. To additionally protect your account from excessive charges, AWS temporarily limits your ability to create some AWS resources. This does not make your account secure; it only partially limits the unauthorized usage for which you could be charged. Note: This check does not guarantee the identification of exposed access keys or compromised EC2 instances. You are ultimately responsible for the safety and security of your access keys and AWS resources.

8 – General

Requirement

Obtain and promptly analyse detailed time-synchronised logs and real-time alerts for the Tenant’s cloud service accounts used to access, and especially to administer, the cloud service.

AWS Response

AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service.

With CloudTrail, you can get a history of AWS API calls for your account, including API calls made via the AWS Management Console, AWS SDKs, command line tools, and higher-level AWS services (such as AWS CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing.

To maintain the integrity of your log data, it is important to carefully manage access around the generation and storage of your log files. The ability to view or modify your log data should be restricted to authorized users. A common log-related challenge for on premise environments is the ability to demonstrate to regulators that access to log data is restricted to authorized users. This control can be time-consuming and complicated to demonstrate effectively because most on premise environments do not have a single logging solution or consistent logging security across all systems.

With AWS CloudTrail, access to Amazon S3 log files is centrally controlled in AWS, which allows you to easily control access to your log files and help demonstrate the integrity and confidentiality of your log data.

This [paper](#) provides an overview of common compliance requirements related to logging and details how AWS CloudTrail features can help satisfy these requirements.

9 – General

Requirement

Obtain and promptly analyse detailed time-synchronised logs and real-time alerts generated by the cloud service used by the tenant e.g. operating system, web server and application logs.

AWS Response

You can execute Continuous Monitoring of logical controls on your own systems. You assume the responsibility and management of the guest operating system (including updates and security patches), other associated application software as well as the configuration of the AWS provided security group firewall. In addition to the monitoring services that AWS provides, you can leverage most OS level and application monitoring tools that you have used in traditional on premise deployments.

[Amazon CloudWatch](#) is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. Amazon CloudWatch can monitor AWS resources as well as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch

to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

[CloudWatch Logs](#) lets you monitor and troubleshoot your systems and applications using your existing system, application, and custom log files. With CloudWatch Logs, you can monitor your logs, in near real-time, for specific phrases, values or patterns (metrics). For example, you could set an alarm on the number of errors that occur in your system logs or view graphs of web request latencies from your application logs. You can view the original log data to see the source of the problem if needed. Log data can be stored and accessed for as long as you need using highly durable, low-cost storage so you don’t have to worry about filling up hard drives.

You can use Amazon CloudWatch Logs to monitor, store, and access your log files from Amazon Elastic Compute Cloud (Amazon EC2) instances, AWS CloudTrail, or other sources. You can then retrieve the associated log data from CloudWatch Logs using the Amazon CloudWatch console, the [CloudWatch Logs commands in the AWS CLI](#), the [CloudWatch Logs API](#), or the [CloudWatch Logs SDK](#).

You can use CloudWatch Logs to:

- Monitor Logs from Amazon EC2 Instances in Real-time
- Monitor AWS CloudTrail Logged Events
- Archive Log Data

10 – General

Requirement

Avoid providing the CSP with account credentials (or the ability to authorise access) to sensitive systems outside of the CSP’s cloud such as systems on the tenant’s corporate network.

AWS Response

AWS does not request that you disclose your customer passwords in order to provide the services or support. AWS provides infrastructure, and you manage everything else, including the operating system, the network configuration, and

the installed applications. You control your own guest operating systems, software and applications.

When you launch an instance, you should specify the name of the key pair you plan to use to connect to the instance. If you don't specify the name of an existing key pair when you launch an instance, you won't be able to connect to the instance. When you connect to the instance, you must specify the private key that corresponds to the key pair you specified when you launched the instance. Amazon EC2 doesn't keep a copy of your private key; therefore, if you lose a private key, there is no way to recover it.

11 – General

Requirement

Use multi-tenancy mechanisms provided by the CSP e.g. to separate the tenant’s web application and network traffic from other tenants, use the CSP’s hypervisor virtualisation instead of web server software virtual hosting.

AWS Response

Amazon EC2’s simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon’s proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change.

The AWS environment is a virtualized, multi-tenant environment. Customer can also select dedicated Amazon EC2 instances, which are single tenant. AWS has implemented security management processes, PCI controls, and other security controls designed to isolate each customer from other customers. AWS systems are designed to prevent you from accessing physical hosts or instances not assigned to you by filtering through the virtualization software. This architecture has been validated by an independent PCI Qualified Security Assessor (QSA) and was found to be in compliance with all requirements of PCI DSS version 3.1 published in April 2015.

Note: AWS also has single-tenancy options. Dedicated Instances are Amazon EC2 instances launched within your Amazon Virtual Private Cloud (Amazon VPC) that run hardware dedicated to a single customer. Dedicated Instances

let you take full advantage of the benefits of Amazon VPC and the AWS cloud while isolating your Amazon EC2 compute instances at the hardware level.

12 – General

Requirement

Perform up-to-date encrypted backups in a format avoiding CSP lock-in, stored offline at the tenant’s premises or at a second CSP requiring multi-factor authentication to modify/delete data. Annually test the recovery process.

AWS Response

You retain control and ownership of your content and it is your responsibility to manage your data backup plans. You can export your EC2 instance image (an EC2 instance image in AWS is referred to as an Amazon Machine Image AMI) and use it on premise or at another provider (subject to software licensing restrictions). For more information, see [Introduction to AWS Security Processes](#).

AWS supports several methods for loading and retrieving data, including: the public Internet; a direct network connection with AWS Direct Connect; the AWS Import/Export service where AWS will import data into S3; and for backups of application data, the AWS Storage Gateway helps you backup your data to AWS.

AWS allows you to move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport.

AWS allows you to perform your own backups to tapes using your own tape backup service provider. However, a tape backup is not a service provided by AWS. Amazon S3 service is designed to drive the likelihood of data loss to near zero percent and the durability equivalent of multi-site copies of data objects is achieved through data storage redundancy.

Amazon S3 provides a highly durable storage infrastructure. Objects are redundantly stored on multiple devices across multiple facilities in an Amazon S3 Region. Once stored, Amazon S3 maintains the durability of objects by quickly detecting and repairing any lost redundancy. Amazon S3 also regularly

verifies the integrity of data stored using checksums. If corruption is detected, it is repaired using redundant data. Data stored in S3 is designed to provide 99.999999999% durability and 99.99% availability of objects over a given year.

AWS allows you to use your own encryption mechanisms to encrypt backups for nearly all the services, including S3, EBS, and EC2. IPsec tunnels to VPC are also encrypted. Amazon S3 also offers you Server Side Encryption as an option. You can also use third-party encryption technologies.

The AWS CloudHSM service allows you to protect your encryption keys within HSMs designed and validated to government standards for secure key management. You can securely generate, store, and manage the cryptographic keys used for data encryption such that they are accessible only by you. AWS CloudHSM helps you comply with strict key management requirements without sacrificing application performance.

AWS Key Management Service (KMS) is a managed service that makes it easy for you to create and control the encryption keys used to encrypt your data, and uses Hardware Security Modules (HSMs) to protect the security of your keys. AWS Key Management Service is integrated with several other AWS services to help you protect your data you store with these services. AWS Key Management Service is also integrated with AWS CloudTrail to provide you with logs of all key usage to help meet your regulatory and compliance needs.

13 – General

Requirement

Contractually retain legal ownership of tenant data. Perform a due diligence review of the CSP’s contract and financial viability as part of assessing privacy and legal risks.

AWS Response

You retain control and ownership of your data.

AWS only uses your content to maintain or provide the AWS services that you have selected or to comply with the law or a binding, legal government request. AWS treats all customer content the same and has no insight as to what type of content that you choose to store in AWS. AWS simply makes available the

compute, storage, database and networking services that you select. See <https://aws.amazon.com/agreement/> for further information.

AWS errs on the side of protecting your privacy and is vigilant in determining which law enforcement requests we must comply with. AWS does not hesitate to challenge orders from law enforcement if we think the orders lack a solid basis.

Further legal information is available at this site <https://aws.amazon.com/legal/>.

14 – General

Requirement

Implement adequately high bandwidth, low latency, reliable network connectivity between the tenant (including the tenant’s remote users) and the cloud service to meet the tenant’s availability requirements.

AWS Response

You can choose your network path to AWS facilities, including multiple VPN endpoints in each AWS Region. In addition, AWS Direct Connect makes it easy to establish a dedicated network connection from your premises to AWS. Using AWS Direct Connect, you can establish private connectivity between AWS and your data center, office, or colocation environment, which in many cases can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than Internet-based connections. Refer to AWS Overview of Security Processes Whitepaper for additional details - available at <http://aws.amazon.com/security>.

AWS Direct Connect lets you establish a dedicated network connection between your network and one of the AWS Direct Connect locations. Using industry standard 802.1q VLANs, this dedicated connection can be partitioned into multiple virtual interfaces. This allows you to use the same connection to access public resources such as objects stored in Amazon S3 using public IP address space, and private resources such as Amazon EC2 instances running within an [Amazon Virtual Private Cloud \(VPC\)](#) using private IP space, while maintaining network separation between the public and private environments. Virtual interfaces can be reconfigured at any time to meet your changing needs.

Network latency over the Internet can vary given that the Internet is constantly changing how data gets from point A to B. With AWS Direct Connect, you choose the data that utilizes the dedicated connection and how that data is routed which can provide a more consistent network experience over Internet-based connections.

AWS Direct Connect makes it easy to scale your connection to meet your needs. AWS Direct Connect provides 1 Gbps and 10 Gbps connections, and you can easily provision multiple connections if you need more capacity. You can also use AWS Direct Connect instead of establishing a VPN connection over the Internet to your Amazon VPC, avoiding the need to utilize VPN hardware that frequently can’t support data transfer rates above 4 Gbps.

15 – General

Requirement

Use a cloud service that meets the tenant’s availability requirements. Assess the Service Level Agreement penalties, and the number, severity, recency and transparency of the CSP’s scheduled and unscheduled outages.

AWS Response

AWS commits to high levels of availability in its service level agreements (SLAs). For example, Amazon EC2 commits to annual uptime percentage of at least 99.95% during the service year. Amazon S3 commits to monthly uptime percentage of at least 99.9% Service credits are provided in the case these availability metrics are not met. See <https://aws.amazon.com/legal/service-level-agreements/> .

For many services, AWS can perform regular maintenance and system patching without rendering the service unavailable or requiring reboots. AWS’ own maintenance and system patching generally do not impact you. You control maintenance of the instances themselves.

AWS publishes our most up-to-the-minute information on service availability on the Service Health Dashboard. Amazon Web Services keeps a running log of all service interruptions that we publish for the past year. Refer to <http://status.aws.amazon.com>.

You should architect your AWS usage to take advantage of multiple Regions and Availability Zones. Distributing applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes including natural disasters or system failures.

AWS utilizes automated monitoring systems to provide a high level of service performance and availability. Proactive monitoring is available through a variety of online tools both for internal and external use. Systems within AWS are extensively instrumented to monitor key operational metrics. Alarms are configured to notify operations and management personnel when early warning thresholds are crossed on key operational metrics. An on-call schedule is used such that personnel are always available to respond to operational issues. This includes a pager system so alarms are quickly and reliably communicated to operations personnel. AWS Network Management is regularly reviewed by independent third-party auditors as a part of AWS ongoing compliance with SOC, PCI DSS, ISO 27001 and FedRAMP.

16 – General

Requirement

Develop and annually test a disaster recovery and business continuity plan to meet the tenant’s availability requirements e.g. where feasible for simple architectures, temporarily use cloud services from an alternative CSP.

AWS Response

You retain control and ownership of your data. AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move your data traffic away from the affected area. AWS SOC reports provides further details. ISO 27001 standard Annex A, domain 15 provides additional details. AWS has been validated and certified by an independent auditor to confirm alignment with ISO 27001 certification.

Using AWS, you can enable faster disaster recovery of your critical IT systems without incurring the infrastructure expense of a second physical site. The AWS cloud supports many popular disaster recovery (DR) architectures from “pilot light” environments that are ready to scale up at a moment’s notice to “hot standby” environments that enable rapid failover. For more information about

Disaster Recovery on AWS see the [Disaster Recovery](#) website and [Disaster Recovery](#) whitepaper.

AWS provides you with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures. You can place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. Each Availability Zone is designed as an independent failure zone. In case of failure, automated processes move customer data traffic away from the affected area.

AWS data centers incorporate physical protection to mitigate against environmental risks. AWS’ physical protection against environmental risks has been validated by an independent auditor and has been certified as being in alignment with ISO 27002 best practices. Refer to ISO 27001 standard, Annex A domain 9 and the AWS SOC 1 Type II report for additional information.

You retain control and ownership of your content and it is your responsibility to manage your data backup plans. You move data as needed on and off AWS storage. AWS Import/Export service for S3 accelerates moving large amounts of data into and out of AWS using portable storage devices for transport.

VM Import/Export enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances and export them back to your on-premises environment. This offering allows you to leverage your existing investments in the virtual machines that you have built to meet your IT security, configuration management, and compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. You can also export imported instances back to your on-premises virtualization infrastructure, allowing you to deploy workloads across your IT infrastructure.

VM Import/Export is available at no additional charge beyond standard usage charges for Amazon EC2 and Amazon S3. See <https://aws.amazon.com/ec2/vm-import/> for further information.

17 – General

Requirement

Manage the cost of a genuine spike in demand or denial of service via contractual spending limits, denial of service mitigation services and judicious use of the CSP’s infrastructure capacity e.g. limits on automated scaling.

AWS Response

To help guarantee availability of AWS resources, as well as minimize billing risk for new customers, AWS maintains service limits for each account. Some service limits are raised automatically as you build a history with AWS, though most AWS services require that you request limit increases manually.

For a list of the default limits for each service, as well as how to request a service limit increase, see [AWS Service Limits](#).

Note: Most limits are specific to a particular AWS region, so if your use case requires higher limits in multiple regions, file separate limit increase requests for each region you plan to use.

To avoid exceeding service limits while building or scaling your application, you can use the AWS Trusted Advisor [Service Limits](#) check to monitor some limits. For a list of limits that are included in the Trusted Advisor check, see [Service Limits Check Questions](#).

EC2 has a service-specific limits dashboard that can help you manage your instance, EBS, and Elastic IP limits. For more information about EC2’s Limits dashboard, see [Amazon EC2 Service Limits](#).

For more information about service limits, go [here](#).

You can also monitor your AWS costs by using CloudWatch. With CloudWatch, you can create billing alerts that notify you when your usage of your services exceeds thresholds that you define. You specify these threshold amounts when you create the billing alerts. When your usage exceeds these amounts, AWS sends you an email notification. You can also sign up to receive notifications when AWS prices change. For more information, go [here](#).

Cost Explorer is a free tool that you can use to view graphs of your costs (also known as spend data) for up to the last 13 months, and forecast how much you are likely to spend for the next three months. You can use Cost Explorer to see patterns in how much you spend on AWS resources over time, identify areas that need further inquiry, and see trends that you can use to understand your costs. You can also specify time ranges for the data you want to see, and you can view time data by day or by month.

For example, you can use Cost Explorer to see which service you use the most, which Availability Zone (AZ) most of your traffic is in, which linked account uses AWS the most, and more. For more information, go [here](#).

Within the Cost Explorer tool, a budget is a way to plan your costs (also known as spend data), and to track how close your costs are to exceeding your budgeted amount. Budgets use data from Cost Explorer to provide you with a quick way to see your estimated charges from AWS, and to see how much your predicted usage will accrue in charges by the end of the month. Budgets also compare the estimated charges to the amount that you want to spend, and lets you see how much of your budget has been spent. Budgets are updated every 24 hours. Budgets track your unblended costs and subscriptions, but do not track refunds. AWS does not use your forecasts to create a budget for you.

You can create budgets for different types of cost. For example, you can create a budget to see how much you are spending on a particular service, or how often you call a particular API operation. Budgets use the same data filters as Cost Explorer. For more information, go [here](#).

Auto Scaling helps you maintain application availability and allows you to scale your [Amazon EC2](#) capacity up or down automatically according to conditions you define. You can use Auto Scaling to help ensure that you are running your desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs. Auto Scaling is well suited both to applications that have stable demand patterns or that experience hourly, daily, or weekly variability in usage. You can specify the maximum number of instances in each Auto Scaling group, and Auto Scaling ensures that your group never goes above this size.

Distributed denial of service (DDoS) attacks are sometimes used by malicious actors in an attempt to flood a network, system, or application with more traffic, connections, or requests than it can handle. Not surprisingly, customers often ask us how we can help them protect their applications against these types of attacks. To help you optimize for availability, AWS provides best practices that allow you to use the scale of AWS to build a DDoS-resilient architecture.

IaaS Risk Mitigations

1 – IaaS

Requirement

Securely configure, harden and maintain VMs with host based security controls e.g. firewall, intrusion prevention system, logging, antivirus software, and prompt patching of all software that the tenant is responsible for.

AWS Response

You retain control of their own guest operating systems, software and applications and are responsible for performing vulnerability scans and patching of your own systems.

Regularly patch, update, and secure the operating system and applications on your instance. For more information about updating Amazon Linux, see [Managing Software on Your Linux Instance](#). For more information about updating your Windows instance, see [Updating Your Windows Instance in the Amazon EC2 User Guide for Microsoft Windows Instances](#).

Amazon EC2 provides a complete firewall solution; this mandatory inbound firewall is configured in a default deny-all mode and Amazon EC2 customers must explicitly open the ports needed to allow inbound traffic. The traffic may be restricted by protocol, by service port, as well as by source IP address (individual IP or Classless Inter-Domain Routing (CIDR) block). AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic.

AWS [WAF](#) is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you

control over which traffic to allow or block to your web applications by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. Also, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

With AWS WAF you pay only for what you use. AWS WAF pricing is based on how many rules you deploy and how many web requests your web application receives. There are no upfront commitments.

This paper provides AWS best practices for DDoS resiliency
https://d0.awsstatic.com/whitepapers/DDoS_White_Paper_June2015.pdf.

AWS [Elastic Beanstalk](#) is an easy-to-use service for deploying and scaling web applications and services developed with Java, [.NET](#), PHP, Node.js, Python, Ruby, Go, and [Docker](#) on familiar servers such as Apache, Nginx, Passenger, and [IIS](#).

Elastic Beanstalk regularly releases updates for the Linux and Windows Server based [platforms](#) that run applications on an Elastic Beanstalk environment. A platform consists of a software component (an AMI running a specific version of an OS, tools, and Elastic Beanstalk-specific scripts), and a configuration component (the default [settings](#) applied to environments created with the platform). New platform versions provide updates to existing software components and support for new features and configuration options.

With managed platform updates, you can configure your environment to automatically upgrade to the latest version of a platform during a scheduled maintenance window. Your application remains in service during the update process with no reduction in capacity.

You can configure your environment to automatically apply patch version updates, or both patch and minor version updates. Managed platform updates don't support major version updates, which may introduce changes that are backwards incompatible.

2 – IaaS

Requirement

Use a corporately approved and secured computer to administer VMs requiring access from the tenant’s IP address, encrypted traffic, and a SSH/RDP PKI key pair protected with a strong passphrase.

AWS Response

Amazon VPC offers a wide range of tools that give you more control over your AWS infrastructure. Within a VPC, you can define your own network topology by defining subnets and routing tables, and you can restrict access at the subnet level with network ACLs and at the resource level with VPC security groups. You can isolate your resources from the Internet and connect them to your own data center through a VPN. You can assign elastic IP addresses to some instances and connect them to the public Internet through an Internet gateway, while keeping the rest of your infrastructure in private subnets. VPC makes it easier to protect your AWS resources while you keep the benefits of AWS with regards to flexibility, scalability, elasticity, performance, availability, and the pay-as-you-use pricing model.

You can add or remove rules for a security group (also referred to as authorizing or revoking inbound or outbound access). A rule applies either to inbound traffic (ingress) or outbound traffic (egress). You can grant access to a specific CIDR range, or to another security group in your VPC or in a peer VPC (requires a VPC peering connection). For example by leveraging part of your organisation’s public IP address range, you could limit inbound SSH and RDP access to be allowed only from your network (via the VPC Internet Gateway). Similarly if a VPN or Direct Connect connection to the VPC is in place you could limit SSH and RDP access to only a section of your organisation’s private IP range.

You can connect your VPC to remote networks by using a VPN connection. The following are some of the connectivity options available to you.

- AWS Hardware VPN (VPC VPG)
- AWS Direct Connect
- Software VPN

Amazon EC2 uses public–key cryptography to encrypt and decrypt login information. Public–key cryptography uses a public key to encrypt a piece of data, such as a password, then the recipient uses the private key to decrypt the data. The public and private keys are known as a key pair.

To log in to your instance, you must create a key pair, specify the name of the key pair when you launch the instance, and provide the private key when you connect to the instance. Linux instances have no password, and you use a key pair to log in using SSH. With Windows instances, you use a key pair to obtain the administrator password and then log in using RDP.

You can use Amazon EC2 to create your key pair, this will create a 2048-bit SSH-2 RSA keys. For more information, see [Creating Your Key Pair Using Amazon EC2](#).

Alternatively, you could use a third-party tool and then import the public key to Amazon EC2. For more information, see [Importing Your Own Key Pair to Amazon EC2](#).

Amazon EC2 stores the public key only, and you store the private key. Anyone who possesses your private key can decrypt your login information, so it's important that you store your private keys in a secure place.

Amazon EC2 accepts the following formats:

- OpenSSH public key format (the format in `~/.ssh/authorized_keys`)
- Base64 encoded DER format
- SSH public key file format as specified in [RFC4716](#)

Amazon EC2 does not accept DSA keys. Make sure your key generator is set up to create RSA keys. Supported lengths: 1024, 2048, and 4096.

3 – IaaS

Requirement

Only use VM template images provided by trusted sources, to help avoid the accidental or deliberate presence of malware and backdoor user accounts.

Protect the tenant’s VM template images from unauthorised changes.

AWS Response

An Amazon Machine Image (AMI) provides the information required to launch an instance, which is a virtual server in the cloud. You specify an AMI when you launch an instance, and you can launch as many instances from the AMI as you need. You can also launch instances from as many different AMIs as you need.

You can customize the instance that you launch from a public AMI and then save that configuration as a custom AMI for your own use. Instances that you launch from your AMI use all the customizations that you've made. You can also use custom AMI instances with AWS CloudFormation. AWS CloudFormation gives developers and systems administrators an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

After you create an AMI, you can keep it private so that only you can use it, or you can share it with a specified list of AWS accounts. You can also make your custom AMI public so that the community can use it. Building a safe, secure, usable AMI for public consumption is a fairly straightforward process, if you follow a few simple guidelines. For information about how to create and use shared AMIs, see [Shared AMIs](#).

You also control the updating and patching of your guest OS, including security updates. Amazon-provided Windows and Linux-based AMIs are updated regularly with the latest patches, so if you do not need to preserve data or customizations on your running Amazon AMI instances, you can simply relaunch new instances with the latest updated AMI. In addition, updates are provided for the Amazon Linux AMI via the Amazon Linux yum repositories.

VM Import/Export enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances and export them back to your on-premises environment. This offering allows you to leverage your existing investments in the virtual machines that you have built to meet your IT security, configuration management, and compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. You can also export imported instances back to your on-premises virtualization infrastructure, allowing you to deploy workloads across your IT infrastructure.

VM Import/Export is available at no additional charge beyond standard usage charges for Amazon EC2 and Amazon S3.

The Center for Internet Security, Inc. (CIS) is a 501c3 nonprofit organization focused on enhancing the cyber security readiness and response of public and private sector entities, with a commitment to excellence through collaboration. CIS provides resources that help partners achieve security goals through expert guidance and cost-effective solutions. CIS provide preconfigured AMI’s on the AWS Marketplace here: https://aws.amazon.com/marketplace/seller-profile/ref=dtl_pcp_sold_by?ie=UTF8&id=6b3b0dc2-c6f4-487b-8f29-9edba5f39eed

[Amazon Inspector](#) is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector automatically assessed applications for vulnerabilities or deviations from best practices. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of security.

4 – IaaS

Requirement

Implement network segmentation and segregation e.g. n-tier architecture, using host based firewalls and CSP’s network access controls to limit inbound and outbound VM network connectivity to only required ports/protocols.

AWS Response

Amazon Virtual Private Cloud (Amazon VPC) lets you provision a logically isolated section of the Amazon Web Services (AWS) cloud where you can launch AWS resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways.

You can easily customize the network configuration for your Amazon Virtual Private Cloud. For example, you can create a public-facing subnet for your web servers that has access to the Internet, and place your backend systems such as databases or application servers in a private-facing subnet with no Internet access. You can leverage multiple layers of security, including security groups and network access control lists, to help control access to Amazon EC2 instances in each subnet.

Additionally, you can create a Hardware Virtual Private Network (VPN) connection between your corporate datacenter and your VPC and leverage the AWS cloud as an extension of your corporate datacenter.

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. When you launch an instance in a VPC, you can assign the instance to up to five security groups. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC could be assigned to a different set of security groups. If you don't specify a particular group at launch time, the instance is automatically assigned to the default security group for the VPC.

For each security group, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic. This section describes the basic things you need to know about security groups for your VPC and their rules.

The default state is to deny all incoming traffic, and you should plan carefully what you will open when building and securing your applications. Well-informed traffic management and security design are still required on a per instance basis. AWS further encourages you to apply additional per-instance filters with host-based firewalls such as IPtables or the Windows Firewall and VPNs. This can restrict both inbound and outbound traffic.

A network access control list (ACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC. For more information about the differences between security groups and network ACLs, see [Comparison of Security Groups and Network ACLs](#).

5 – IaaS

Requirement

Utilise secure programming practices for software developed by the tenant

AWS Response

It is your responsibility to use secure programming practices.

AWS’s development process for AWS infrastructure and services follows secure software development best practices, which include formal design reviews by the AWS Security Team, threat modeling, and completion of a risk assessment. Static code analysis tools are run as a part of the standard build process, and all deployed software undergoes recurring penetration testing performed by carefully selected industry experts. Our security risk assessment reviews begin during the design phase and the engagement lasts through launch to ongoing operations.

This [whitepaper](#) describes how Amazon Web Services (AWS) adds value in the various phases of the software development cycle, with specific focus on development and test. For the development phase, it shows how to use AWS for managing version control; it describes project management tools, the build process, and environments hosted on AWS; and it illustrates best practices. For the test phase, it describes how to manage test environments and run various kinds of tests, including load testing, acceptance testing, fault tolerance testing, etc. AWS provides unique advantages in each of these scenarios and phases, allowing you to pick and choose the ones most appropriate for your software development project. The intended audiences for this paper are project managers, developers, testers, systems architects, or anyone involved in software production activities.

With AWS, your development and test teams can have their own resources, scaled according to their own needs. Provisioning complex environments or platforms composed of multiple instances can be done easily using AWS CloudFormation stacks or some of the other automation techniques described. In large organizations comprising multiple teams, it is a good practice to create an internal role or service responsible for centralizing and managing IT resources running on AWS. This role typically consists of:

- Promoting internal development and test practices described here
- Developing and maintaining template AMIs and template AWS CloudFormation stacks with the different tools and platforms used in your organization
- Collecting resource requests from project teams, and provisioning resources on AWS according to your organization’s policies, including network configuration (e.g., Amazon VPC), security configurations (e.g., Security Groups and IAM credentials)

- Monitoring resource usage and charges using Amazon CloudWatch, and allocating these to team budgets

While you can use the AWS Management Console to achieve the tasks above, you might want to develop your own internal provisioning and management portal for a tighter integration with internal processes. You can do this by using one of the AWS SDKs, which allow programmatic access to resources running on AWS.

6 – IaaS

Requirement

Architect to meet availability requirements e.g. minimal single points of failure, data replication, automated failover, multiple availability zones, geographically separate data centres and real-time availability monitoring.

AWS Response

AWS provides you with the capability to implement a robust continuity plan, including the utilization of frequent server instance back-ups, data redundancy replication, and multi-region/availability zone deployment architectures.

The [AWS Well-Architected Framework](#) whitepaper describes how you can assess and improve your cloud-based architectures to better understand the business impact of your design decisions. Included in the paper are the four general design principles, as well as specific best practices and guidance in four conceptual areas (security, reliability, performance efficiency, and cost optimization). These four areas are defined as the pillars of the Well-Architected Framework.

AWS provides you with the flexibility to place instances and store data within multiple geographic regions as well as across multiple Availability Zones within each region. You should architect your AWS usage to take advantage of multiple Regions and Availability Zones.

The [Architecting for the Cloud](#) whitepaper Amazon CloudWatch is a monitoring service for AWS cloud resources and the applications you run on AWS. You can use Amazon CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. Amazon CloudWatch can monitor AWS resources such as Amazon EC2 instances, Amazon DynamoDB tables, and Amazon RDS DB instances, as well

as custom metrics generated by your applications and services, and any log files your applications generate. You can use Amazon CloudWatch to gain system-wide visibility into resource utilization, application performance, and operational health. You can use these insights to react and keep your application running smoothly.

This [whitepaper](#) is intended for solutions architects and developers who are building solutions that will be deployed on Amazon Web Services (AWS). It provides architectural patterns and advice on how to design systems that are secure, reliable, high performing, and cost efficient. It includes a discussion on how to take advantage of attributes that are specific to the dynamic nature of cloud computing (elasticity, infrastructure automation, etc.). In addition, this whitepaper also covers general patterns, explaining how these are evolving and how they are applied in the context of cloud computing.

7 – IaaS

Requirement

If high availability is required, implement clustering and load balancing, a Content Delivery Network for public web content, automated scaling with an adequate maximum scale value, and real-time availability monitoring.

AWS Response

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances in the cloud. It enables you to achieve greater levels of fault tolerance in your applications, seamlessly providing the required amount of load balancing capacity needed to distribute application traffic.

Achieve higher levels of fault tolerance for your applications by using Elastic Load Balancing to automatically route traffic across multiple instances and multiple Availability Zones. Elastic Load Balancing ensures that only healthy Amazon EC2 instances receive traffic by detecting unhealthy instances and rerouting traffic across the remaining healthy instances. If all of your EC2 instances in one Availability Zone are unhealthy, and you have set up EC2 instances in multiple Availability Zones, Elastic Load Balancing will route traffic to your healthy EC2 instances in those other zones.

[Auto Scaling](#) helps you maintain application availability and allows you to scale your [Amazon EC2](#) capacity up or down automatically according to conditions you define. You can use Auto Scaling to help ensure that you are running your desired number of Amazon EC2 instances. Auto Scaling can also automatically increase the number of Amazon EC2 instances during demand spikes to maintain performance and decrease capacity during lulls to reduce costs. Auto Scaling is well suited both to applications that have stable demand patterns or that experience hourly, daily, or weekly variability in usage.

Whether you are running one Amazon EC2 instance or thousands, you can use Auto Scaling to detect impaired Amazon EC2 instances and unhealthy applications, and replace the instances without your intervention. This ensures that your application is getting the compute capacity that you expect.

[Amazon Route 53](#) is a highly available and scalable cloud Domain Name System (DNS) web service. You can use Amazon Route 53 health checking and DNS failover features to enhance the availability of the applications running behind Elastic Load Balancers. Route 53 will fail away from a load balancer if there are no healthy EC2 instances registered with the load balancer or if the load balancer itself is unhealthy.

Using Route 53 DNS failover, you can run applications in multiple AWS regions and designate alternate load balancers for failover across regions. In the event that your application is unresponsive, Route 53 will remove the unavailable load balancer endpoint from service and direct traffic to an alternate load balancer in another region. To get started with Route 53 failover for Elastic Load Balancing, visit the [Elastic Load Balancing Developer Guide](#) and the [Amazon Route 53 Developer Guide](#).

[Amazon CloudFront](#) is a global content delivery network (CDN) service. It integrates with other Amazon Web Services products to give developers and businesses an easy way to distribute content to end users with low latency, high data transfer speeds, and no minimum usage commitments. The service automatically responds as demand increases or decreases without any intervention from you. Amazon CloudFront also uses multiple layers of caching at each edge location and collapses simultaneous requests for the same object before contacting your origin server. These optimizations further help reduce the need to scale your origin infrastructure as your website becomes more popular.

Amazon CloudFront is built using Amazon’s highly reliable infrastructure. The distributed nature of edge locations used by Amazon CloudFront automatically routes end users to the closest available location as required by network conditions. Origin requests from the edge locations to AWS origin servers (e.g., Amazon EC2, Amazon S3, etc.) are carried over network paths that Amazon constantly monitors and optimizes for both availability and performance.

[AWS WAF](#) is a web application firewall that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow or block to your web applications by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting, and rules that are designed for your specific application. New rules can be deployed within minutes, letting you respond quickly to changing traffic patterns. Also, AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of web security rules.

With AWS WAF you pay only for what you use. AWS WAF pricing is based on how many rules you deploy and how many web requests your web application receives. There are no upfront commitments.

PaaS Risk Mitigations

1 – PaaS

Requirement

Securely configure and promptly patch all software that the tenant is responsible for.

AWS Response

While AWS provides a managed service, you are responsible for setting up and managing network controls, such as firewall rules, and for managing platform-level identity and access management separately from IAM

AWS is responsible for patching systems supporting the delivery of service to customers. This is done as required per AWS policy and in accordance with ISO 27001, NIST, and PCI requirements. AWS manages the underlying

infrastructure and foundation services, the operating system and the application platform.

Elastic Beanstalk regularly releases [platform updates](#) to provide fixes, software updates and new features. With managed platform updates, you can configure your environment to automatically upgrade to the latest version of a platform during a scheduled [maintenance window](#). Your application remains in service during the update process with no reduction in capacity.

You can configure your environment to automatically apply [patch version updates](#), or both patch and minor version updates. Managed platform updates don't support major version updates, which may introduce changes that are backwards incompatible.

When you enable managed platform updates, you can also configure AWS Elastic Beanstalk to replace all instances in your environment during the maintenance window, even if a platform update isn't available. Replacing all instances in your environment is helpful if your application encounters bugs or memory issues when running for a long period.

2 – PaaS

Requirement

Utilise secure programming practices for software developed by the tenant

AWS Response

Covered in 5 IaaS.

3 – PaaS

Requirement

Architect to meet availability requirements e.g. minimal single points of failure, data replication, automated failover, multiple availability zones, geographically separate data centres and real-time availability monitoring.

AWS Response

Covered in 6 IaaS.

4 – PaaS

Requirement

If high availability is required, implement clustering and load balancing, a Content Delivery Network for public web content, automated scaling with an adequate maximum scale value, and real-time availability monitoring.

AWS Response

Covered in 7 IaaS.

SaaS Risk Mitigations

1 – SaaS

Requirement

Use security controls specific to the cloud service e.g. tokenisation to replace sensitive data with non-sensitive data, or ASD-approved encryption of data (not requiring processing) and avoid exposing the decryption key.

AWS Response

AWS provides specific SOC controls to address the threat of inappropriate access, and the public certification and compliance initiatives covered in this document address efforts to prevent inappropriate access. All certifications and third-party attestations evaluate logical access preventative and detective controls. In addition, periodic risk assessments focus on how access is controlled and monitored.

AWS allows you to implement your own security architecture. For more information about server and network security, see the [AWS security whitepaper](#).

All data stored by AWS on behalf of you has strong tenant isolation security and control capabilities. You retain control and ownership of your data, thus it is your responsibility to choose to encrypt the data. AWS allows you to use your own encryption mechanisms for nearly all of the AWS services, including S3, EBS, and EC2. IPsec tunnels to VPC are also encrypted. In addition, you can leverage AWS Key Management Systems (KMS) to create and control encryption keys using 256-bit AES envelope encryption (refer to <https://aws.amazon.com/kms/>).

2 – SaaS

Requirement

If high availability is required, where possible and appropriate, implement additional cloud services providing layered denial of service mitigation, where these cloud services might be provided by third party CSPs.

AWS Response

Covered in 7 IaaS.

Additionally, the [AWS Best Practices for DDoS Resiliency](#) whitepaper provides guidance on how you can improve the resiliency of your applications running on Amazon Web Services (AWS) against Distributed Denial of Service attacks. The paper provides an overview of Distributed Denial of Service attacks, techniques that can help maintain availability, and reference architectures to provide architectural guidance with the goal of improving your resiliency.

Further Reading

For additional help, see the following sources:

- AWS Security Page: <http://aws.amazon.com/security>
- AWS Compliance Page: <http://aws.amazon.com/compliance>
- AWS IRAP Page: <http://aws.amazon.com/compliance/irap/>
- Overview of AWS Security Processes:
http://d0.awsstatic.com/whitepapers/Security/AWS_Security_Whitepaper.pdf
- AWS Risk and Compliance Whitepaper:
https://d0.awsstatic.com/whitepapers/compliance/AWS_Risk_and_Compliance_Whitepaper.pdf
- AWS Security Best Practices:
<https://d0.awsstatic.com/whitepapers/aws-security-best-practices.pdf>
- KMS Cryptographic Details:
<https://d0.awsstatic.com/whitepapers/KMS-Cryptographic-Details.pdf>

Document Revisions

Date	Description
June 2017	Initial publication.
