
# SCS Azure Security Checklist

## Baseline Security Controls

**Version 0.3**
**22 August 2017**

# Table of Contents

# 1 Purpose

This document provides guidance to enterprises on security controls that must be implemented when developing or deploying applications on Azure.

It details a list of prescriptive actions that can be taken for immediate risk reduction and also seeks to raise architectural design concepts that should be considered when securing workloads moving into Azure. As the product landscape of Azure continues to evolve, so will the security model. This document will continue to be updated as new security controls and services are made available.

Project teams are expected to use this document in conjunction with artefacts like the Statement of Applicability (SoA) that detail the security objectives that a project must meet. It does not supersede the need for security controls that fall outside the domain of Azure, such as following a secure software development lifecycle (S-SDLC) or performing code scanning and regular penetration testing.

SCS is continually improving its security services and operational capabilities within cloud and this document is intended to communicate a number of initiatives that projects may be able to leverage in the short to medium term. This is intended to promote design decisions that enable decoupling from native Azure or avoiding implementation of complex and siloed solutions that may be provided as an enterprise service in the near future.

The information contained within this document includes data that has been collated from a number of sources including from the Cloud Security Alliance (CSA), Azure Cloud Security Blog and Homepage (Azure Security Centre) and the Centre for Internet Security (CIS).

# 2 Target Audience

The target audience of this document is any organisation looking to leverage the benefits of Azure while needing to easily and consistently meet enterprise security policies and standards.

# 3 Structure

This document begins by describing some of the threats to cloud service utilisation and how these relate to the enterprise. It also includes a high-level overview of Cloud Security Principles and the ‘Shared Responsibility Model’ employed by Azure. The idea behind this introduction is

to solidify the idea that security of cloud is of major importance for enterprises and must be considered from the very beginning and at every stage of the Plan, Build, Run approach. Security is everyone's responsibility.

The Security Checklist is broken down into sections that consist of a summary of controls for easy referencability and follows with further detailed implementation steps.

Each section consists of a summary of security requirements that in the opinion of SCS MUST be met in order to have a secure deployment, recommendations that SHOULD be implemented by organisations wanting to extend the security posture of their environment, or FUTURE design considerations when considering enterprise security services delivered further down the line.

| List of Requirements | | |
|---|---|---|
| **x.x.x** | **This is a mandatory requirement that must be met for all projects to deploy into Azure regardless of maturity** | **MUST** |
| **x.x.x** | **This is a recommendation that extends baseline security and should be considered for teams that are more mature in their use of Azure** | **SHOULD** |
| **x.x.x** | **This is architectural guidance on services or initiatives that are on SCS's Cloud Security Enterprise Roadmap and should be considered when designing applications for Azure** | **FUTURE** |

# 4 Inherent Security of Azure

According to Azure:

*"Azure provides a secure platform on which you can build your solutions. We also provide services and technologies to make your solutions on Azure more secure. Because of the many options available to you, many of you have voiced an interest in what Microsoft recommends as best practices and patterns for improving security."*

The recurring message from Azure is that security is their first and most important consideration when developing services for customers and which they aim to consistently apply across the product portfolio. The advantage of which is a robust, supported and well integrated security model available for customer use.

A downside to the many inherent Azure security controls is their technical nature that require individual configuration and implementation for an account, Azure Virtual Network (AVN) or environment. This creates a problem where teams are forced to 're-invent the wheel' and design security solutions for each project they deliver. This can lead to an overly complex control

environment that has to be supported by teams possibly unfamiliar with security. Gartner captured this concept in one of its latest Top 10 Predictions: "Through 2020, 95 percent of cloud security failures will be the customer's fault."

Secure Cloud Services is attempting to address this issue by embarking on a number of Cloud control initiatives including the development of a security reference architecture for Infrastructure-as-a-Service (IaaS). The project will deliver reusable services and design patterns that projects and teams can leverage to improve standardisation, simplification, reuse and drive down security costs.

# 5 Threats to Enterprise use of Cloud

While there are many security concerns in the Cloud, the CSA has published a list of the top 12 listed in order of risk:

## 5.1 Data Breach

A data breach is an incident in which sensitive, protected or confidential information is released, viewed, stolen or used by an individual who is not authorised to do so. A data breach may be the primary objective of a targeted attack or may simply be the result of human error, application vulnerabilities or poor security practices.

**Impact on your organisation**
Data breach is of significant concern to the enterprise as it potentially holds billions of records with personally identifiable information (PII). Industry estimates related to PII breach costs stand at $154 per record in 2015.

## 5.2 Insufficient Identity, Credential and Access Management

Data breaches and enabling of attacks can occur because of a lack of scalable identity access management systems, failure to use multifactor authentication, weak password use, and a lack of ongoing automated rotation of cryptographic keys, passwords and certificates.

**Impact on your organisation**
Illustrating this problem is the historically disparate Shadow IT environments within customer organisations that have sprung up across various areas of their business. Some staff take it upon themselves to manage the delivery of their own applications into IaaS environments such as Azure, but consequently then have to manage the non-functional aspects such as security and operations. This can lead to an environment where management of credentials and access is far from what security policy stipulates, creating considerable risk.

## 5.3 Insecure Interfaces and APIs

Cloud computing providers expose a set of software user interfaces (UI's) or application programming interfaces (API's) that customers use to manage and interact with cloud services. Provisioning, management, orchestration and monitoring are all performed with these interfaces. The security and availability of general cloud services is dependent on the security of these basic API's. From authentication and access control to encryption and activity monitoring, these interfaces must be designed to protect against both accidental and malicious attempts to circumvent policy.

## 5.4 System Vulnerabilities

System vulnerabilities are exploitable bugs in programs that attackers can use to infiltrate a computer system for the purpose of stealing data, taking control of the system or disrupting service operations. Vulnerabilities within the components of the operating system – kernel, system libraries and application tools – put the security of all services and data at significant risk.

## 5.5 Account Hijacking

Account or service hijacking is not new. Attack methods such as phishing, fraud and exploitation of software vulnerabilities still achieve results. Credentials and passwords are often reused, which amplifies the impact of such attacks. Cloud solutions add a new threat to the landscape. If an attacker gains access to your credentials, they can eavesdrop on your activities and transactions, manipulate data, return falsified information and redirect your clients to illegitimate sites. Your account or service instances may become a new base for attackers. From here, they may leverage the power of your reputation to launch subsequent attacks.

## 5.6 Malicious Insiders

The risk caused by malicious insiders has been debated in the security industry. While the level of threat is left to debate, the fact that insider threat is a real adversary is not. CERN defines an insider threat as follows:

"A malicious insider threat to an organization is a current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems."

## 5.7 Advanced Persistent Threats

Advanced Persistent Threats (APTs) are a parasitical form of cyber attack that is usually orchestrated by well funded nation states or organised crime gangs and infiltrates systems to establish a foothold in the computing infrastructure of target companies from which they smuggle data and intellectual property. APTs pursue their goals stealthily over extended periods of time, often adapting to the security measures intended to defend against them. Spearphishing, direct hacking systems, delivering attack code through USB devices, penetration through partner networks and use of unsecured or third-party networks are common points of entry for APT's. Once in place, APTs can move laterally through data center networks and blend in with normal network traffic to achieve their objectives.

## 5.8 Data Loss

Data stored in the cloud can be lost for reasons other than malicious attacks. An accidental deletion by the cloud service provider, or worse, a physical catastrophe such as a fire or earthquake, can lead to the permanent loss of customer data unless the provider or cloud consumer takes adequate measures to back up data, following best practices in business continuity and disaster recovery – as well as daily data backup and possibly off-site storage. Furthermore, the burden of avoiding data loss does not fall solely on the provider's shoulders. If a customer encrypts his or her data before uploading it to the cloud but loses the encryption key, the data will be lost as well.

## 5.9 Insufficient Due Diligence

When executives create business strategies, cloud technologies and Cloud Service Providers (CSP's) must be considered. Developing a good roadmap and checklist for due diligence when evaluating technologies and CSP's is essential for the greatest chance of success. An organization that rushes to adopt cloud technologies and choose CSP's without performing due diligence exposes itself to a myriad of commercial, financial, technical, legal and compliance risks that jeopardize its success. This applies whether the company is considering moving to the cloud or merging with or acquiring a company that has moved to the cloud or is considering doing so.

## 5.10 Abuse and Nefarious Use of Cloud Services

Poorly secured cloud service deployments, free cloud service trials and fraudulent account sign-ups via payment instrument fraud expose cloud computing models such as IaaS, PaaS, and SaaS to malicious attacks. Malicious actors may leverage cloud computing resources to target users, organizations or other cloud providers. Examples of misuse of cloud service-based resources include launching DDoS attacks, email spam and phishing campaigns; "mining" for

digital currency; large-scale automated click fraud; brute-force compute attacks of stolen credential databases; and hosting of malicious or pirated content.

## 5.11 Denial of Service

Denial-of-service (DoS) attacks are attacks meant to prevent users of a service from being able to access their data or their applications. By forcing the targeted cloud service to consume inordinate amounts of finite system resources such as processor power, memory, disk space or network bandwidth, the attacker—or attackers, as is the case in distributed denial-of-service (DDoS) attacks—causes an intolerable system slowdown and leaves all legitimate service users confused and angry as to why the service is not responding.

## 5.12 Shared Technology Vulnerabilities

Cloud service providers deliver their services scalably by sharing infrastructure, platforms or applications. Cloud technology divides the “as a Service” offering without substantially changing the off-the-shelf hardware/software—sometimes at the expense of security. Underlying components (e.g., CPU caches, GPUs, etc.) that comprise the infrastructure supporting cloud services deployment may not have been designed to offer strong isolation properties for a multitenant architecture (IaaS), re-deployable platforms (PaaS) or multi-customer applications (SaaS). This can lead to shared technology vulnerabilities that can potentially be exploited in all delivery models. A defense in-depth strategy is recommended and should include compute, storage, network, application and user security enforcement and monitoring, whether the service model is IaaS, PaaS, or SaaS. The key is that a single vulnerability or misconfiguration can lead to a compromise across an entire provider’s cloud.

# 6 Principles

- Security is everyone’s responsibility
- Design for resiliency
- Expect the worst
- Enable logging and maintain environment visibility through monitoring
- Keep infrastructure and applications up-to-date
- Grant least privilege - only authorised and monitored access to production resources

# 7 Azure Shared Responsibility Model

The following concepts should be considered when securing systems and data on Azure:

- Security measures that Azure implements and operates – "security of the cloud"

- Security measures that the customer implements and operates, related to the security of cloud customer content and applications that make use of Azure services – "security in the cloud"

While Azure manages security of the cloud, security in the cloud is the responsibility of the cloud customer. The customer retains control of what security they choose to implement to protect their content, platform, applications, systems and networks, no differently than they would for applications in an onsite data center.

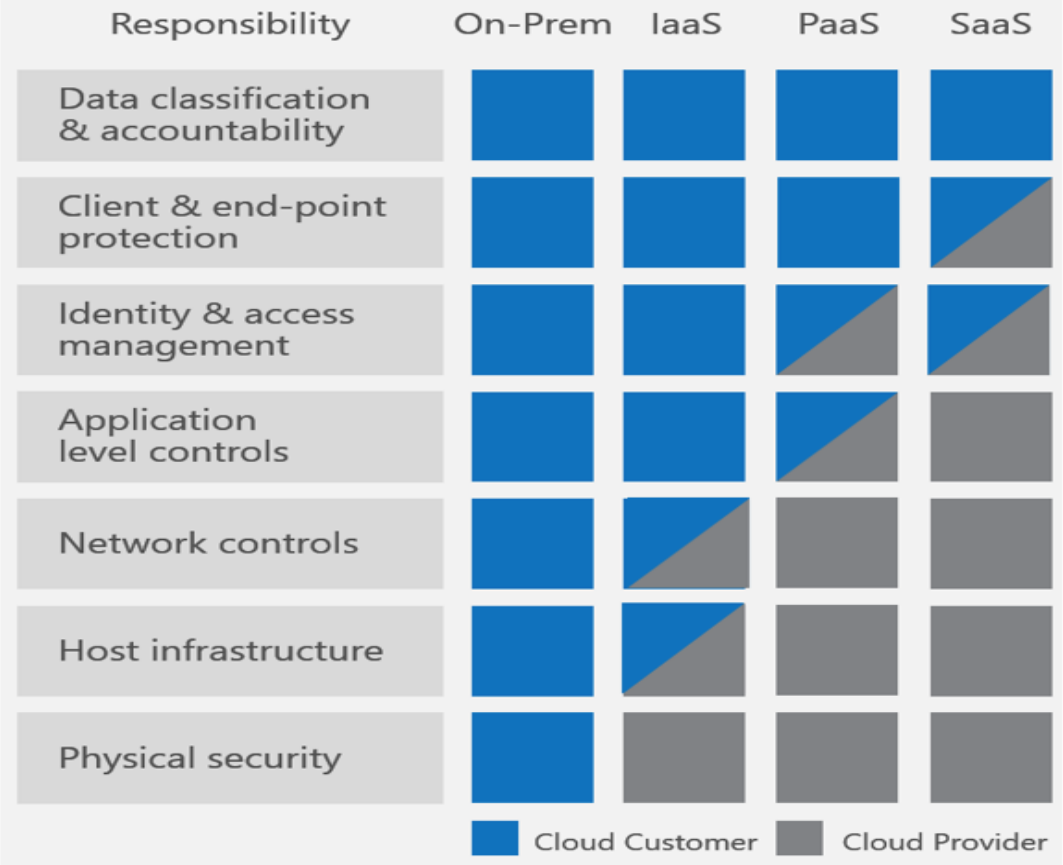


# 8 Security Checklist

## 8.1 Governance

| Governance | | |
|---|---|---|
| 8.1.1 | Project has been assessed by security department | MUST |
| 8.1.2 | Project has completed a Privacy Impact Assessment | MUST |
| 8.1.3 | Project has documented Data Classification Levels (DCL) for assets | MUST |

| | | |
|---|---|---|
| **8.1.4** | **An owner has been assigned to receive alerts** | **MUST** |
| **8.1.5** | **Azure account must be integrated with corporate alerting tools (e.g. CloudCheckr)** | **FUTURE** |
| **8.1.6** | **A read-only security account has been created for audit purposes** | **MUST** |
| **8.1.7** | **An administrator account has been created for emergency access** | **MUST** |

## 8.1.1 Project has been through Security triage

MUST All projects must go through security assessment which should be overseen by the office of the CISO

## 8.1.2 Project has completed a Privacy Impact Assessment

MUST This Privacy Impact Assessment (PIA) is aimed at new projects that involve personal data which impact the enterprise. The purpose of this assessment is to help mitigate the privacy risks that are involved when project teams carry out new projects, including creating systems that involve 'personal data'.

## 8.1.3 Project has documented Data Classification Levels (DCL) for assets

MUST An asset owner must follow the Data Classification Policy and assess systems for Data Classification Level. This assessment is used to determine the commensurate level of controls needed to protect a particular resource. The DCL is also needed for asset tagging within Azure for compliance purposes.
Azure tagging guidelines can be found here:
https://docs.microsoft.com/en-us/azure/azure-resource-manager/resource-group-using-tags

## 8.1.4 An owner has been assigned to receive alerts

MUST While many controls within Azure are protective in nature, some require monitoring and alerting when a breach of policy has been detected. These alerts should be sent to an active distribution list, or at a minimum a member of staff, responsible for that Azure environment or account. Usually this would be the Data Owner.

## 8.1.5 Azure account must be integrated with corporate alerting tools (e.g. CloudCheckr)

FUTURE CloudCheckr is used by SCS to gather metrics on inventory, usage and security across Azure accounts. All accounts should be onboarded to central tooling such as CloudCheckr.

### 8.1.6 A read-only security account has been created for audit purposes

MUST In case of security investigation or audit, security requires the creation of an account with a read-only role. This role must provide access to audit Azure infrastructure but not the data within databases or access to hosts. This must be enabled at all times and will be managed in accordance with enterprise Privileged User Management Standard. For Azure purposes, a Reader role would fit this description as per table below:

| Role | Edit security policy | Apply security recommendations for a resource | Remediate or Dismiss alerts | View alerts across a subscription | View alerts for a specific resource |
|---|---|---|---|---|---|
| Subscription Owner | X | X | X | X | X |
| Subscription Contributor | X | X | X | X | X |
| Resource Group Owner | -- | X | -- | -- | X |
| Resource Group Contributor | -- | X | -- | -- | X |
| Reader | -- | -- | -- | X | X |

Table 1

### 8.1.7 An administrator account has been created for emergency access

MUST In case of a major security incident, security requires the creation of an account with administrative privileges This account should be a subscription owner or contributor. The account must be provided access to ALL Azure infrastructure including data within databases. This must be enabled at all times and will be enrolled within CyberArk (or similar) and managed in accordance with enterprise Privileged User Management Standard. It is expected that this account will only be used in case of emergency and the account owner would be notified of its use.
For Azure purposes, a Subscription Contributor role would fit this description

## 8.2 Azure Specific Best Practices

| Azure Specific Security Best Practices | | |
|---|---|---|
| **8.2.1** | **Deploy all resources through Resource Manager** | **MUST** |
| **8.2.2** | **Segregate environments with DevTest Labs** | **MUST** |

| 8.2.3 | Activate Azure Security Center | MUST |
|---|---|---|

## 8.2.1 Deploy all resources through Resource Manager

MUST A resource group is a container for resources that share a common lifecycle. The Resource Manager deployment model provides several benefits:

- Deploy, manage, and monitor all the services for your solution as a group, rather than handling these services individually.
- Repeatedly deploy your solution throughout its lifecycle and have confidence your resources are deployed in a consistent state.
- Apply access control to all resources in your resource group, and those policies are automatically applied when new resources are added to the resource group.
- Apply tags to resources to logically organize all the resources in your subscription.
- Use JavaScript Object Notation (JSON) to define the infrastructure for your solution. The JSON file is known as a Resource Manager template.
- Define the dependencies between resources so they are deployed in the correct order.

For information about the deployment models, see Understanding deployment models.

## 8.2.2 Segregate environments with DevTest Labs

SHOULD Using Azure for labs and development environments enables organizations to gain agility in testing and development by taking away the delays that hardware procurement introduces. Unfortunately, a lack of familiarity with Azure or a desire to help expedite its adoption might lead the administrator to be overly permissive with rights assignment. This risk might unintentionally expose the organization to internal attacks. Some users might be granted a lot more access than they should have.

The Azure DevTest Labs service uses Azure Role-Based Access Control (RBAC). By using RBAC, you can segregate duties within your team into roles that grant only the level of access necessary for users to do their jobs. RBAC comes with predefined roles (owner, lab user, and contributor). You can even use these roles to assign rights to external partners and greatly simplify collaboration.

Because DevTest Labs uses RBAC, it's possible to create additional, custom roles. DevTest Labs not only simplifies the management of permissions, it simplifies the process of getting environments provisioned. It also helps you deal with other typical challenges of teams that are working on development and test environments. It requires some preparation, but in the long term, it will make things easier for your team.

Azure DevTest Labs features include:

- Administrative control over the options available to users.
- The administrator can centrally manage things like allowed VM sizes, maximum number of VMs, and when VMs are started and shut down.
- Automation of lab environment creation.
- Cost tracking.
- Simplified distribution of VMs for temporary collaborative work.
- Self-service that enables users to provision their labs by using templates.
- Managing and limiting consumption.

### 8.2.3 Activate Azure Security Center

MUST Security Center helps prevent, detect, and respond to threats with increased visibility into and control over the security of Azure resources. It provides integrated security monitoring and policy management across Azure subscriptions, helps detect threats that might otherwise go unnoticed, and works with a broad ecosystem of security solutions.

**Key Capabilities**
Security Center delivers easy-to-use and effective threat prevention, detection, and response capabilities that are built in to Azure. Key capabilities are:

| Stage | Capability |
|---|---|
| Prevent | Monitors the security state of Azure resources |
| Prevent | Defines policies for Azure subscriptions and resource groups based on enterprise security requirements, the types of applications used, and the sensitivity of the organisation’s data |
| Prevent | Uses policy-driven security recommendations to guide service owners through the process of implementing needed controls |
| Prevent | Rapidly deploys security services and appliances from Microsoft and partners |
| Detect | Automatically collects and analyzes security data from Azure resources, the network, and partner solutions like anti-malware programs and firewalls |
| Detect | Leverages global threat intelligence from Microsoft products and services, the Microsoft Digital Crimes Unit (DCU), the Microsoft Security Response Center (MSRC), and external feeds |

| | |
|---|---|
| Detect | Applies advanced analytics, including machine learning and behavioural analysis |
| Respond | Provides prioritized security incidents/alerts |
| Respond | Offers insights into the source of the attack and impacted resources |
| Respond | Suggests ways to stop the current attack and help prevent future attacks |

Security Center is enabled with Azure subscription. Access Security Center from the Azure portal.

See Getting started with Azure Security Center to guide you through the security-monitoring and policy-management components of Security Center.

## 8.3 General Azure Security Configuration

| General Azure Security Configuration | | |
|---|---|---|
| **8.3.1** | **Use stand-alone hardened workstation for accessing Subscription Owner/Subscription Contributor account** | **MUST** |
| **8.3.2** | **Use Azure logon restrictions to constrain source IP addresses for accessing administrative tools and audit access requests** | **MUST** |
| **8.3.3** | **Deploy management gateway** | **SHOULD** |
| **8.3.4** | **Ensure administrator account credentials are stored within enterprise Privileged User Management system** | **MUST** |
| **8.3.5** | **Ensure no High Security alerts listed in Azure Advisor/Security Center** | **MUST** |
| **8.3.6** | **Security contact details for Azure subscription must be set to security generated unique email address** | **MUST** |
| **8.3.7** | **Azure resources used for production systems must be separate from resources used for pre-production systems (see DevTest Labs)** | **MUST** |
| **8.3.8** | **Install client-side management certificates for all workstations with administrator access** | **SHOULD** |

**Note:** Client systems that have administrator access to infrastructure components should be subjected to the strictest possible policy to reduce security risks. Examples:

- Security policies can include Group Policy settings that deny open Internet access from the device and use of a restrictive firewall configuration.
- Use Internet Protocol security (IPsec) VPNs if direct access is needed.
- Configure separate management and development Active Directory domains.
- Isolate and filter management workstation network traffic.
- Use anti-malware software.
- Implement multi-factor authentication to reduce the risk of stolen credentials.

Consolidating access resources and eliminating unmanaged endpoints also simplifies management tasks.

Also, using a hardened workstation configuration for administering your Azure cloud services, Virtual Machines, and applications can help avoid numerous risks and threats that can come from remotely managing critical IT infrastructure. Both Azure and Windows provide mechanisms that can be employed to help protect and control communications, authentication, and client behavior.

### 8.3.1 Use stand-alone hardened workstation for accessing Subscription Owner/Subscription Contributor account

MUST With a stand-alone hardened workstation, administrators have a PC or laptop that they use for administrative tasks and another, separate PC or laptop for non-administrative tasks. A workstation dedicated to managing your Azure services does not need other applications installed. Additionally, using workstations that support a Trusted Platform Module (TPM) or similar hardware-level cryptography technology aids in device authentication and prevention of certain attacks.

In the stand-alone hardened workstation scenario (shown below), the local instance of Windows Firewall (or a non-Microsoft client firewall) is configured to block inbound connections, such as RDP. The administrator can log on to the hardened workstation and start an RDP session that connects to Azure after establishing a VPN connect with an Azure Virtual Network, but cannot log on to a corporate PC and use RDP to connect to the hardened workstation itself.

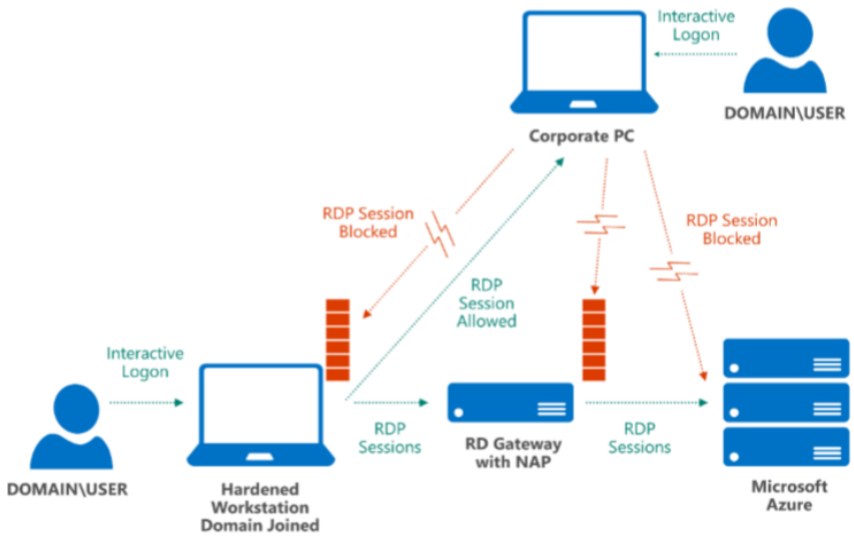


### 8.3.2 Use Azure logon restrictions to constrain source IP addresses for accessing administrative tools and audit access requests.

MUST To help Azure restrict management access (workstations and/or applications), you MUST restrict management access to company office address ranges. This can be used in conjunction with a VPN solution to enable remote administration.

### 8.3.3 Configure management gateway for Azure account access

SHOULD To centralise all administrative access and simplify monitoring and logging, deploy a dedicated Remote Desktop Gateway (RD Gateway) server in the corporate on-premises network, connected to the Azure environment.

A Remote Desktop Gateway is a policy-based RDP proxy service that enforces security requirements. Implementing RD Gateway together with Windows Server Network Access Protection (NAP) helps ensure that only clients that meet specific security health criteria established by Active Directory Domain Services (AD DS) Group Policy objects (GPOs) can connect. In addition:

- Provision an Azure management certificate on the RD Gateway so that it is the only host allowed to access the Azure portal.
- Join the RD Gateway to the same management domain as the administrator workstations. This is necessary when you are using a site-to-site IPsec VPN or ExpressRoute within a domain that has a one-way trust to Azure AD, or if you are federating credentials between your on-premises AD DS instance and Azure AD.
- Configure a client connection authorization policy to let the RD Gateway verify that the client machine name is valid (domain joined) and allowed to access the Azure portal.
- Use IPsec for Azure VPN to further protect management traffic from eavesdropping and token theft, or consider an isolated Internet link via Azure ExpressRoute.
- Enable multi-factor authentication (via Azure Multi-Factor Authentication) or smart-card authentication for administrators who log on through RD Gateway.
- Configure source IP address restrictions or Network Security Groups in Azure to minimise the number of permitted management endpoints.

### 8.3.4 Ensure Subscription Owner account credentials are stored within enterprise Privileged User Management system

MUST Storing privileged account credentials in a secure privileged user management system ensures they can only be accessed by following approved and audited authorisation processes and that credentials are rotated once no longer needed.

### 8.3.5 Ensure no High Security alerts reported in Azure Advisor

MUST Azure Advisor/Security Center are Azure tools that provide best practice checks for cost optimisation, security, fault tolerance and performance improvements. Most of the security checks are for controls listed within this document. All alerts rated as High must be remediated.

### 8.3.6 Security contact details for Azure subscription must be set to security generated unique email address

MUST Azure Security Center will recommend that you provide security contact details for your Azure subscription if you haven't already. This information will be used by Microsoft to contact you if the Microsoft Security Response Center (MSRC) discovers that your customer data has been accessed by an unlawful or unauthorized party. MSRC performs select security monitoring of the Azure network and infrastructure and receives threat intelligence and abuse complaints from third parties.

An email notification is sent on the first daily occurrence of an alert and only for high severity alerts. Email preferences can only be configured for subscription policies. Resource groups within a subscription will inherit these settings.

Security department will provide all Azure accounts with a unique email addresses that redirects to the Security Operations Centre (SOC).

### 8.3.7 Azure resources used for production systems must be separate from resources used for pre-production systems

MUST Production environments must be separate from test, QA, UAT, staging and any other pre-production environments. AVNs must be split into production and pre-production AVNs. Resources including, but not limited to, VM instances, databases, Azure Storage and Key Vault keys must not be shared between production and pre-production environments. This can be achieved using DevTest labs (see above).

### 8.3.8 Install client-side management certificates for all workstations with administrator access

SHOULD To help Azure identify management clients (workstations and/or applications), you SHOULD configure both SMAPI (via customer-developed tools such as Windows PowerShell cmdlets) and the Azure portal to require client-side management certificates to be installed, in addition to SSL certificates.

## 8.4 Identity and Access Management

| IdAM | | |
|---|---|---|
| 8.4.1 | Use multifactor authentication for all users that have admin access | MUST |
| 8.4.2 | Enroll all privileged host and database accounts within PUM solution (e.g. CyberArk) | FUTURE |
| 8.4.3 | Integrate Azure AD with corporate identity and authentication service | MUST |
| 8.4.4 | Account credentials unused for 60 days should be disabled | MUST |
| 8.4.5 | Password length and complexity should be enforced and comply with enterprise password policy and standards | MUST |
| 8.4.6 | Passwords should expire within 60 days | MUST |
| 8.4.7 | Azure AD policies should only be applied to groups and roles | MUST |

### 8.4.1 Use Multi-Factor Authentication (MFA) for all users that have admin access

MUST Some applications or services that you deploy into Azure may have their own authentication mechanisms for both end-user and administrator access, whereas others take full advantage of Azure AD. For accounts that manage your Azure subscription and for accounts that can sign in to virtual machines, using Multi-Factor Authentication gives you a much greater level of security than using only a password.

### 8.4.2 Enroll all privileged host and database accounts within Privileged User Management solution

FUTURE All accounts that could potentially enable user access to host systems and databases must be enrolled within CyberArk (or similar) to ensure such access is authorised and audited. This reduces the risk of unauthorised changes to non-production and production systems.

### 8.4.3 Integrate Azure AD with corporate identity and authentication service

MUST Identity and access management should be handled centrally to ease management of accounts and reduce the risks associated with individual users and credentials

### 8.4.4 Account credentials unused for 60 days should be disabled

MUST Azure AD users can access Azure resources using different types of credentials, such as passwords or access keys. It is recommended that all credentials that have been unused in 90

or greater days be removed or deactivated. *If using ForgeRock for centralised authentication, this requirement is not applicable.*

### 8.4.5 Password length and complexity should be enforced and comply with secure password policy and standards

MUST Passwords should be complex so that the effectiveness of dictionary and brute force attacks are minimised.

### 8.4.6 Passwords should expire within 60 days

MUST Reducing the password lifetime increases account resiliency against brute force login attempts.

### 8.4.7 AD policies should only be applied to groups and roles

MUST Assigning privileges at the group or role level reduces the complexity of access management as the number of users grow. Reducing access management complexity may in-turn reduce opportunity for a principal to inadvertently receive or retain excessive privileges.

## 8.5 Network Security

| Network Security | | |
|---|---|---|
| **8.5.1** | **Enable Network Security Groups (NSGs)** | **MUST** |
| **8.5.2** | **Ensure only ports 80 and 443 are permitted from the internet and only into networks designated as public** | **MUST** |
| **8.5.3** | **No network security groups should allow ingress from */* port 22** | **MUST** |
| **8.5.4** | **No network security groups should allow ingress from */* port 3389** | **MUST** |
| **8.5.5** | **Configure User Defined Routes** | **MUST** |
| **8.5.6** | **NSGs must not contain unused service definitions** | **MUST** |
| **8.5.7** | **Applications must be logically segregated by resource groups** | **SHOULD** |
| **8.5.8** | **AVNs should be split into public and private network zones** | **MUST** |
| **8.5.9** | **Private networks must not be internet routable** | **MUST** |
| **8.5.10** | **Traffic from the internet must terminate in a public zone** | **MUST** |
| **8.5.11** | **All traffic outbound to the internet must traverse a NAT or proxy that logs connection data** | **MUST** |
| **8.5.12** | **Any connectivity from Azure to on premise infrastructure must be via managed VPN using enforced tunneling or Azure ExpressRoute** | **MUST** |

| | | |
|---|---|---|
| **8.5.13** | **Configure NSG Diagnostic Logs to OMS Analytics or event hubs and store for 3 months** | **SHOULD** |
| **8.5.14** | **NSGs must be approved by security before migration into production** | **FUTURE** |

### 8.5.1 Enable Network Security Groups

MUST Network Security Groups (NSGs) can be used to control network segmentation. NSGs are simple stateful packet inspection devices that use the 5-tuple (the source IP, source port, destination IP, destination port, and layer 4 protocol) approach to create allow/deny rules for network traffic. You can allow or deny traffic to and from single IP address, to and from multiple IP addresses or even to and from entire subnets.

Using NSGs for network access control between subnets enables you to put resources that belong to the same security zone or role in their own subnets.

To learn more about Network Security Groups and how you can use them to logically segment your Azure Virtual Networks, please read the article What is a Network Security Group.

### 8.5.2 Ensure only ports 80 and 443 are permitted from the internet and only into networks designated as public

MUST Traffic that originates from an untrusted zone (internet) into a semi-trusted zone (public) is only permitted on ports 80 and 443.
If a project requires additional inbound ports to be opened, a request for an exception needs to be made.

### 8.5.3 No network security groups should allow ingress from */* port 22

MUST Removing unfettered connectivity to remote console services, such as SSH, reduces a server's exposure to risk. If host management is needed via SSH, restrict access to company network ranges only.

### 8.5.4 No network security groups should allow ingress from */* port 3389

MUST Removing unfettered connectivity to remote console services, such as RDP, reduces a server's exposure to risk. If host management is needed via SSH, restrict access to company network ranges only.

### 8.5.5 Configure User Defined Routes

MUST There is a collection of system routes that are enabled by default that allow every virtual machine on an Azure Virtual Network to connect to any other virtual machine on the same Azure Virtual Network, even if the other virtual machines are on different subnets. These default

routes allow virtual machines on the same Azure Virtual Network to initiate connections with each other, and with the Internet (for outbound communications to the Internet only).

You can learn more about User Defined Routes and how to configure them by reading the article What are User Defined Routes and IP Forwarding.

### 8.5.6 Network security groups must not contain unused service definitions

MUST Network security groups should not contain rules for services that are not actively being used across hosts. This limits the attack surface of resources and reduces risk.

### 8.5.7 Applications must be logically segregated by resource groups

SHOULD Assets that logically form part of an application or platform should be grouped together within separate Resource Groups

### 8.5.8 AVNs should be split into public and private network zones

MUST In accordance with the enterprise Firewall Standard, zoning should be enforced by the creation of a public network zone that is internet routable and a private network zone that is not internet routable.

### 8.5.9 Private networks must not be internet routable

MUST Assets located in a private network zone must only be accessed from untrusted networks, such as the internet, by terminating in a semi-trusted zone where traffic can be authenticated or authorised. This provides a defense-in-depth approach to network security and limits the risks associated with exposing internal assets to the internet.

### 8.5.10 Traffic from the internet must terminate in a public zone

MUST Traffic from the internet must not terminate in private networks, but must first terminate in an internet routable zone.

### 8.5.11 All traffic outbound to the internet that originates in a private network must traverse a NAT or proxy that logs connection data

MUST Traffic that originates in a trusted zone (private) must traverse a NAT or proxy within the semi-trusted zone (public) and this traffic must to logged. This provides visibility into network traffic that attempts to access the internet and can be used to detect anomalous traffic or insight during security workflows.

**Note:** If using a tailored NAT host, the easiest way to log network connections is to ensure NSG Diagnostic Logs are configured on any internally facing network interfaces of the NAT server(s) and log to OMS Log Analytics, Event Hubs or Azure Storage. Information on how this can be configured can be found on the Azure website:

https://github.com/Microsoft/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring-overview-of-diagnostic-logs.md

### 8.5.12 Enforced Tunneling through VPN must be used for connectivity between Azure and on-premise estate

MUST Any network connections from Azure to on-premise must traverse a VPN. This can be implemented via the use of Enforced Tunneling

The default routes for an Azure Virtual Network allow virtual machines to initiate traffic to the Internet. This represents a security risk, as these outbound connections could increase the attack surface of a virtual machine and be leveraged by attackers. For this reason, you MUST enable forced tunneling on your virtual machines when you have cross-premises connectivity between your Azure Virtual Network and your on-premises network.

If you do not have a cross premises connection, you MUST use Network Security Groups or Azure virtual network security appliances to control outbound connections to the Internet from your Azure Virtual Machines.

To learn more about forced tunneling and how to enable it, please read the article Configure Forced Tunneling using PowerShell and Azure Resource Manager.

### 8.5.13 Configure NSG Diagnostic Logs to OMS and store for 3 months

SHOULD NSG Diagnostic Logging is a feature that enables the capturing of information about the IP traffic going to and from network interfaces in a AVN. Ensure NSG Diagnostic Logs are configured and log to OMS Log Analytics, Event Hubs or Azure Storage. Information on how this can be configured can be found on the Azure website:
https://github.com/Microsoft/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring-overview-of-diagnostic-logs.md

**Note:** Logging at a network level is a key function for any network security scenario. In Azure, you can log information obtained for Network Security Groups to get network level logging information. With NSG logging, you get information from:

- Activity logs – these logs are used to view all operations submitted to your Azure subscriptions. These logs are enabled by default and can be used within the Azure portal. They were previously known as "Audit logs" or "Operational Logs".
- Event logs – these logs provide information about what NSG rules were applied.
- Counter logs – these logs let you know how many times each NSG rule was applied to deny or allow traffic.

You can also use Microsoft Power BI, a powerful data visualization tool, to view and analyze these logs.

Learn more:

- Log Analytics for Network Security Groups (NSGs)

### 8.5.14 NSGs must be approved by security department before migration into production

FUTURE As enterprises continue to migrate products into the cloud and benefit from the inherent automation found within Azure and other Cloud Service Providers (CSPs), security should be aiming to move from a legacy 'gates' based culture to one of 'security self-service' and notifications. This will attempt to address historic delays caused by firewall rule approval and implementation being sought just before go-live, in favour of pre-approved security group templates that can be applied on the fly.

The goal is that security will provide services to detect and report back any network policy violations to teams in development/testing environments so remediation can be undertaken before applications are deployed into production. Policy violation within production environments will be raised as security incidents and managed accordingly.

## 8.6 Audit and Monitoring

| Audit and Monitoring | | |
|---|---|---|
| **8.6.1** | **Log access to Azure APIs by enabling log integration in all regions** | **MUST** |
| **8.6.2** | **Deploy VM Agent on all VMs** | **MUST** |
| **8.6.3** | **Ensure data collection is enabled for all VMs** | **MUST** |
| **8.6.4** | **Integrate Azure Security Center alerts with Azure log integration** | **MUST** |
| **8.6.5** | **Install Endpoint Protection to Windows VMs** | **MUST** |
| **8.6.6** | **Ensure any master keys within Azure Key Management Service are rotated** | **MUST** |
| **8.6.7** | **Deploy Vulnerability Assessment Agents on VM instances** | **MUST** |
| **8.6.8** | **Configure Activity Monitor alarms for basic security alerting** | **MUST** |
| **8.6.9** | **Extend Application Insights alarms for additional security alerting** | **SHOULD** |
| **8.6.10** | **Feed Application Insights logs into Security Operations Centre (SOC)** | **FUTURE** |

### 8.6.1 Log access to Azure APIs by enabling log integration in all regions

MUST The Azure API call history produced by Azure Activity Log enables security analysis, resource change tracking, and compliance auditing. Additionally, sending these logs to Operations Management Suite provides a single ‘pane-of-glass’.

### 8.6.2 Deploy VM Agent on all VMs

MUST The VM Agent must be installed on virtual machines (VMs) in order to enable data collection. Azure Security Center enables you to see which VMs require the VM Agent and will recommend that you enable the VM Agent on those VMs.
The VM Agent is installed by default for VMs that are deployed from the Azure Marketplace. The article VM Agent and Extensions – Part 2 provides information on how to install the VM Agent.

### 8.6.3 Ensure data collection is enabled for all VMs

MUST Security Center collects data from your virtual machines (VMs) to assess their security state, provide security recommendations, and alert you to threats. When you first access Security Center, data collection is enabled on all VMs in your subscription. Ensure data collection has been enabled. More information on how to set this up can be found on the Azure website:
https://github.com/Microsoft/azure-docs/blob/master/articles/monitoring-and-diagnostics/monitoring-overview-activity-logs.md

### 8.6.4 Integrate Azure Security Center alerts with Azure log integration

MUST Azure log integration currently supports the integration of:

- Azure VM logs
- Azure Audit Logs
- Azure Security Center alerts

Documentation is provided by Azure on how this can be configured:
https://docs.microsoft.com/en-gb/azure/security-center/security-center-integrating-alerts-with-log-integration

### 8.6.5 Install Endpoint Protection to Windows VMs

MUST Azure Security Center will recommend that you provision an anti-malware program to your Azure virtual machines (VMs) if anti-malware is not already enabled. This recommendation applies to Windows VMs only. Presently, this recommendation checks for the presence of either Windows Defender or Trend Micro Deep Security. Additional endpoint protection solutions are planned to be added in the future.

### 8.6.6 Ensure any master keys within Azure Key Management Service are rotated

MUST Rotating encryption keys helps reduce the potential impact of a compromised key as data encrypted with a new key cannot be accessed with a previous key that may have been exposed. This requirement covers Azure Key Vault as well as Azure Trust Centre.

### 8.6.7 Deploy Vulnerability Assessment Agents on VM instances

MUST By default Azure VM instances do not log to a central collection point. Installing the Application Insights Log Agent on VM instances allows for system level OS logs or custom application logs to be fed to Application Insights for storage or analysis. More information can be found here:
https://docs.microsoft.com/en-gb/azure/security-center/security-center-vulnerability-assessment-recommendations

### 8.6.8 Configure alarms for basic security alerting

MUST There are many external tools that can be configured to consume Activity logs and provide alerting (such as Splunk) however OMS is the native Azure monitoring tool for security, performance and health metrics. At a minimum it should be configured for basic security alerting and alarm on the following events:

- Production OS login
- Production infrastructure changes
- Production network security group changes
- Production network (ACLS, gateways, route tables)
- AD policy changes
- Azure portal access without MFA
- Azure portal authentication failures

### 8.6.9 Extend alarms for additional security alerting

SHOULD Additional events that should be alerted on include:

- Security Center configuration
- Customer key state changes
- Azure Storage policy changes
- Azure NSG policy changes

### 8.6.10 Feed Security Center logs into Security Operations Centre (SOC)

FUTURE Implementing a next-generation Security Information and Events Management platform will provide advanced analytics, forensics and incident response capabilities. This platform may be architected on systems such as Splunk which can natively consume Security

Operations Centre data, so by ensuring that projects feed all security information to the Security Operations Center by design, this can be integrated easily once delivered.

## 8.7 Compliance

| Compliance | | |
|---|---|---|
| **8.7.1** | **DCL3 and DCL4 data must be encrypted in transit and at rest** | **MUST** |

### 8.7.1 DCL3 and DCL4 data must be encrypted in transit and at rest

MUST All projects must encrypt any information that is of Data Classification Level 3 or above in Azure. Azure provides native encryption for Azure Storage (Blobs, Tables, Files), Azure SQL and Activity Monitor.

# 9 Supplementary Information

### 9.1 Checklist version:

**v.0.3**

# 10 List of Abbreviations

| Term | Description |
|---|---|
| AVN | Azure Virtual Network |
| DCL | Data Classification Level |
| IPS | Intrusion Prevention System |
| NAT | Network Address Translation |
| NSG | Network Security Group |
| WAF | Web Application Firewall |