Microsoft

Microsoft Cybersecurity Architect

Exam Ref SC-100

Yuri Diogenes
Sarah Young
Mark Simos
Gladys Rodriguez

Microsoft

Exam Ref SC-100
Microsoft Cybersecurity Architect

Yuri Diogenes
Sarah Young
Mark Simos
Gladys Rodriguez

- Use Azure Key Vault for certificate and secrets management. Also, you can encrypt `.pfx` files, passwords, authentication keys, storage account keys, and data encryption keys by using keys protected by hardware security modules.
- Use the Sensitive Operations Report workbook to identify modifications to applications' service principals' credentials or authentication methods.
- Look for opportunities to convert explicit credentials (such as connection strings and API keys) to use managed identities.
- Standardize on using managed identities where possible.
- Review and document the service interconnections and the permissions setup across them and use this information when performing audits or configuration reviews.
- Review organization processes and policies that may describe how identities should be used and the systems they should interconnect to reduce employee friction and confusion on how the service should be configured.
- Ensure that the monitoring systems and alerts are interconnected and that the required organization teams, including the SOC, have access to the required insights.
- If you are considering enabling sign-in frequency, make sure there is no other conditional access policy enabling the configurable token lifetime feature for any of the users for which you are enabling sign-in frequency. Microsoft does not support this configuration because the configurable token lifetime feature was retired in January 2021.
- Do not overuse the sign-in frequency >every time setting. Only use it for specific business needs because it can affect productivity and increase the risk of users approving authentications they might not have initiated because they have grown numb to the prompts.
- Disable legacy authentication using a phased approach. First, you might want to disable basic authentication on a per-protocol basis using Exchange Online authentication policies. Then you might want to extend the protection using conditional access.
- Use a single identity provider that supports all platforms, including operating systems, cloud providers, and third-party services.
- For older tenants that explicitly disabled the CAE preview or only explicitly enabled some users using the old experience, CAE is not enabled by default. Consider enabling it. You do this by selecting **Azure Active Directory > Security > Continuous Access Evaluation** and selecting **Migrate**. This process will create a new conditional access policy named **CA Policy Created From CAE Settings**. For newer tenants or customers that explicitly enabled CAE for all users using the old experience, CAE is automatically enabled, and no migration is needed.
- Be aware that not all services support CAE, though Microsoft is working on extending that support for both Microsoft cloud and third-party services.
- Update your organization code to use CAE.

Skill 3-4: Recommend an authorization strategy

The previous skill focused on describing authentication approaches to ensure the identity is known and trusted. For this skill, the focus will be extending the validation to check the trust factors of that identity and ensuring that the right authorizations are provided. However, we must first start with provisioning the required access for the right authorizations to take place. After all, authentication is the act of granting permissions to perform certain required tasks.

Configuring access to support authorization

Different authorization methods can ensure the resource has the required access. People often plan for authorization to be driven mostly by user accounts. However, as mentioned earlier, in order to get the best security and Zero Trust strategy, all identities and network capabilities should work in harmony to properly authorize the required access. Because of this, the type of access controls and verifications needed to properly authorize a session will be very diverse and will include checking

- **What devices are authorized to communicate with another device** An SQL server may only need to communicate with the identity provider and supporting websites. In addition, most on-premises user workstations do not need to communicate with other workstations. They only need to communicate with servers/applications. This means you can use micro-segmentation capabilities to only allow the required traffic requirements. Some capabilities that can be used include implementing host firewall policies through Microsoft Endpoint Manager, using Azure Network Security Group (NSG), and using Azure Firewall to limit open ports and traffic paths. Be aware that this strategy might not work for mobile workloads, but it can reduce risk on server workloads and specialized endpoints that do not move across environments. For example, on-premises Privileged Access Workstations (PAW) can be used to administer IaaS and on-premises server workloads. (See "Skill 3-7: Design security strategy for privileged-role access to infrastructure, including identity-based firewall rules and Azure PIM," later in this chapter.)
- **What roles and group membership should a user belong to, and whether the access should be time-bound** Group membership helps provide common permissions to a group of users. Some memberships might be permanent, while others, like privileged groups, might not need to be if a just-in-time (JIT) service is used to grant time-bounded granular elevated privileged access in real-time. In addition, membership of a group may be controlled dynamically based on the value of an attribute defined within the identity object—for example, the user's title or department or portions of the device's name.
- **What data may need to be accessed and when** Permissions can be assigned to control data access. However, other capabilities, such as Microsoft Data Loss Prevention (DLP) and Microsoft Sensitive Labels, use attributes defined within the data to drive access control. In addition, retention labels and policies can control the time a file or data may be available.

- **What explicit permissions are applied, and which service controls it** Microsoft Office 365 applications like SharePoint and Teams can be controlled using their own permissioning system or Azure AD specialized groups (Microsoft 365 group). In addition, Azure AD entitlement management can automate the provisioning and revocation of access based on group membership or attributes defined within the user object. To ensure that proper authorizations are in place, there must be an understanding of the different provisioning services available and how they overwrite each other to reduce the friction that these services may otherwise create. This means that processes must be defined and made available for administrators to follow to reduce conflicts that may be created. In addition, audits should be planned to review the different access provided by these services and ensure that excessive permissions are not being granted.
- **Who should you share data with** DLP can be used to control what data should be shared based on the content of that data. For example, the organization may define that no file containing a Social Security Number will be shared externally. Microsoft Defender for Cloud Apps can use policies defined by DLP and Information Protection to search for sensitive data that may be shared externally through the different Software as a Service (SaaS) applications used by the organization. Then, Defender for Cloud Apps can take the proper steps to either remove the data or set the proper permissions so the data is not exposed externally. Microsoft Purview Information Barriers is another example of how you can control shared data; this compliance solution prevents two-way communication and collaboration between groups and users meeting certain criteria. For example, finance personnel might be unable to communicate and share information with certain groups within their organization, such as marketing.

All these capabilities enforce least-access/privileged access, which, in turn, ensures that the proper authorization to resources is provided. By incorporating all access controls from network, privileged, identity, external, endpoint, and adaptive access, you can implement a defense-in-depth approach, where each layer supports the others and helps contain threats. More detailed recommendations will be discussed in subsequent sections.

Network access

Nowadays, many organizations are migrating to the cloud. However, many still have network management infrastructure (such as IaaS and on-premises). To protect anywhere, security must have a cloud-delivered capability and be interconnected, so insights can be shared across services. This is because isolated signals can create blind spots that will prevent analysts from determining if malicious events have occurred and cannot verify if the attackers have used those blind spots to jump to other environments in the organization.

Secure Access Service Edge (SASE) is a framework that tries to interconnect signals from cloud-native security technologies such as Secure Web Gateway, Cloud Access Security Broker (CASB), Firewall as a Service (FWaaS), Remote Brower Isolation, encryption, and software-defined WAN (SDWAN). This interconnection of signals helps you gain end-to-end visibility of the environment, consistent policy enforcement, control of sensitive data, and so on.

Besides the capabilities mentioned, there are other functionalities you should consider using a threat protection solution to uncover network misconfiguration and improve the environment's security posture. For example, Defender for Cloud provides capabilities to control network access, such as:

- Adaptive network hardening helps filter traffic to and from resources by using network security groups (NSGs). It also provides hardening alerts and recommended policy rules.
- Network map provides a graphical view of the connections between your virtual endpoints and subnets to provide insights that help harden the environment. Figure 3-6 shows how Microsoft Defender for Cloud reports the network-related recommendations.

MORE INFO **A list of Defender for Cloud network alerts can be seen at *https://aka.ms/NetworkAlerts*. Also, the network map helps uncover unwanted connections, enabling you to better isolate workloads and subnets.**

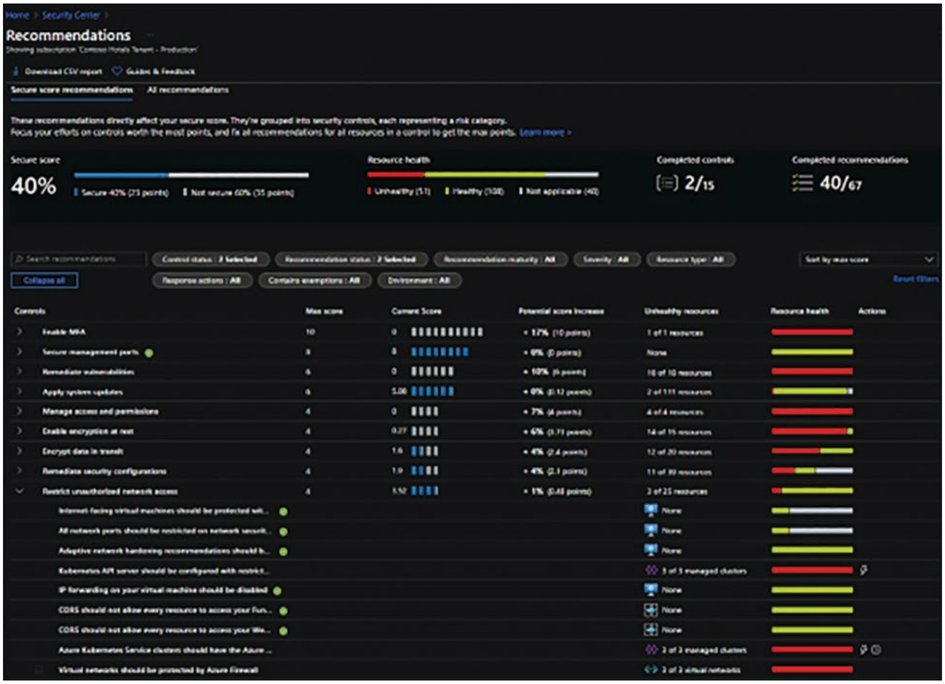
Home > Security Center >

Recommendations

Showing subscription 'Contoso Hotels Tenant - Production'

Download CSV report | Guides & Feedback

Secure score recommendations | All recommendations

These recommendations directly affect your secure score. They're grouped into security controls, each representing a risk category.
Focus your efforts on controls worth the most points, and fix all recommendations for all resources in a control to get the max points. Learn more >

Secure score: 40% — Secure 40% (23 points) | Not secure 60% (35 points)

Resource health: Unhealthy (51) | Healthy (108) | Not applicable (40)

Completed controls: 2/15

Completed recommendations: 40/67

Search recommendations | Control status : 2 Selected | Recommendation status : 2 Selected | Recommendation maturity : All | Severity : All | Resource type : All | Sort by max score

Collapse all | Response actions : All | Contains exemptions : All | Environment : All | Reset filters

Controls	Max score	Current Score	Potential score increase	Unhealthy resources	Resource health	Actions
Enable MFA	10	0	+ 17% (10 points)	1 of 1 resources		
Secure management ports ●	8	8	+ 0% (0 points)	None		
Remediate vulnerabilities	6	0	+ 10% (6 points)	10 of 10 resources		
Apply system updates	6	5.88	+ 0% (0.12 points)	2 of 111 resources		
Manage access and permissions	4	0	+ 7% (4 points)	4 of 4 resources		
Enable encryption at rest	4	0.27	+ 6% (3.73 points)	14 of 15 resources		
Encrypt data in transit	4	1.6	+ 4% (2.4 points)	12 of 20 resources		
Remediate security configurations	4	1.9	+ 4% (2.1 points)	11 of 30 resources		
Restrict unauthorized network access	4	3.52	+ 1% (0.48 points)	2 of 25 resources		
Internet-facing virtual machines should be protected wit... ●				None		
All network ports should be restricted on network securit... ●				None		
Adaptive network hardening recommendations should b... ●				None		
Kubernetes API server should be configured with restrict...				3 of 3 managed clusters		
IP forwarding on your virtual machine should be disabled ●				None		
CORS should not allow every resource to access your Fun... ●				None		
CORS should not allow every resource to access your We... ●				None		
Azure Kubernetes Service clusters should have the Azure ...				2 of 2 managed clusters		
Virtual networks should be protected by Azure Firewall				2 of 2 virtual networks		

FIGURE 3-6 Defender for Cloud Recommendations

Defender for IoT provides insights through the Device Map, such as the capabilities of Defender for Cloud Network Map, which can help define segmentation rules to isolate devices. In addition, the user interface provides capabilities to search by IP, MAC, protocols, known applications, and so on, which can help uncover unexpected connections and enable continuous configuration improvements in the environment. Defender for IoT interconnects with

Defender for Cloud, Defender for Endpoint, Sentinel, and others to help provide an end-to-end view of the environment. Figure 3-7 shows the Defender for IoT Device Map.

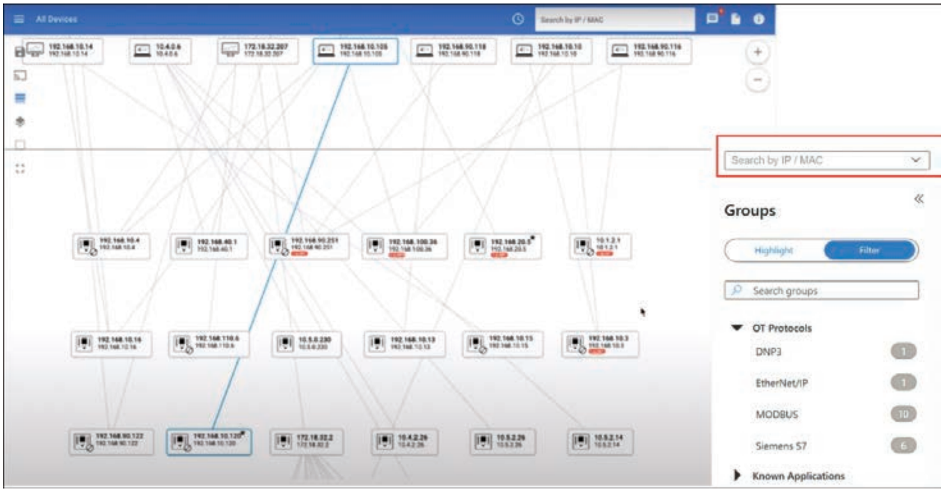

FIGURE 3-7 Defender for IoT Device Map and network-related filtering capabilities

There are many other network-related capabilities from Microsoft. However, this section focuses on network capabilities that help enforce least-access to ensure proper authorizations occur. You can create a comprehensive authorization strategy by incorporating these approaches with other access controls described in the following sections. In addition, the configuration baseline created by these approaches helps provide information supporting audits and/or investigations of unexpected configurations in the environment.

Privileged access

When starting an attack, cybercriminals target privileged accounts to quickly spread and attain wide access across the environment. Therefore, it's important to make it difficult for attackers to find sources that can provide privileged access and carry the attack further: Make sure attackers cannot see who the members of the privileged groups are. Make sure attackers cannot use well-known and easy-to-use attack vectors to compromise users (email and Internet browsing). And make sure those identities never have excessive permissions.

> ***MORE INFO*** **Because of the extensive number of privileged access solutions provided by Microsoft, "Skill 3-7: Design security strategy for privileged-role access to infrastructure including identity-based firewall rules and Azure PIM," describes these capabilities.**

Identity access

Security group membership, role-based access control, security namespaces and permissions, feature flags, and access levels are some of the capabilities that can help the proper authorizations take place. To ensure proper authentication takes place and the least-access principle is applied, some configurations that you should consider when architecting your authentication strategy include the following:

- Security namespaces in Azure DevOps is a functionality that stores access control lists (ACL) in tokens to determine the level of access that different entities must perform over a resource.
- Feature Flags in Azure DevOps deployments is a technique that uses conditional logic to control the visibility of certain application functionality for users, for example, deploying functionality only to a limited audience for testing or paying higher subscription fees or others.
- Access Levels grants or restricts access to select web portal features.
- Microsoft Office utilizes its own Microsoft 365 roles and Azure AD Microsoft 365 groups to control access to Microsoft Office resources.
- SaaS applications such as ServiceNow, WorkDay, and SuccessFactors use their own access management systems in conjunction with Azure AD or other identity providers to control access to their services.

As seen, user access is controlled in many areas of the environment. To control all the access levels, organizations should dedicate personnel for auditing all these different accesses and develop guidance for employees to follow when granting and revoking access. To help with the organization's governance needs, Microsoft has released many different governance solutions, such as Azure AD Entitlement Management, which controls permission granting and revoking.

When planning, it is also important to try to automate and simplify the environment's permissioning for the environment. Consider these questions:

- Should you use dynamic groups that populate membership based on some value within the user object (such as department) and then use those groups to provide access within the applications?
- Do the applications support using dynamic groups to assign permissions?
- Should you use services like Entitlement Management for granting and revoking access?

Some organizations work on projects that have citizenship, education, certification, or other employee requirements. They might need employees to sign yearly agreements and take yearly training for them to be able to work on those organization projects. This means that you should consider whether information should be populated from an HR or personnel system to drive automated workloads around employee screening, clearance, training, and other requirements. And you should consider the type of access needed once that information is verified.

External access

External access is more than the access from your organization users to your internal or DMZ resources. Now you must deal with:

- External collaborations (B2B and B2C) as described as part of "Skill 3-2: Recommend an identity store (tenants, B2B, B2C, and hybrid)"