‹packt›

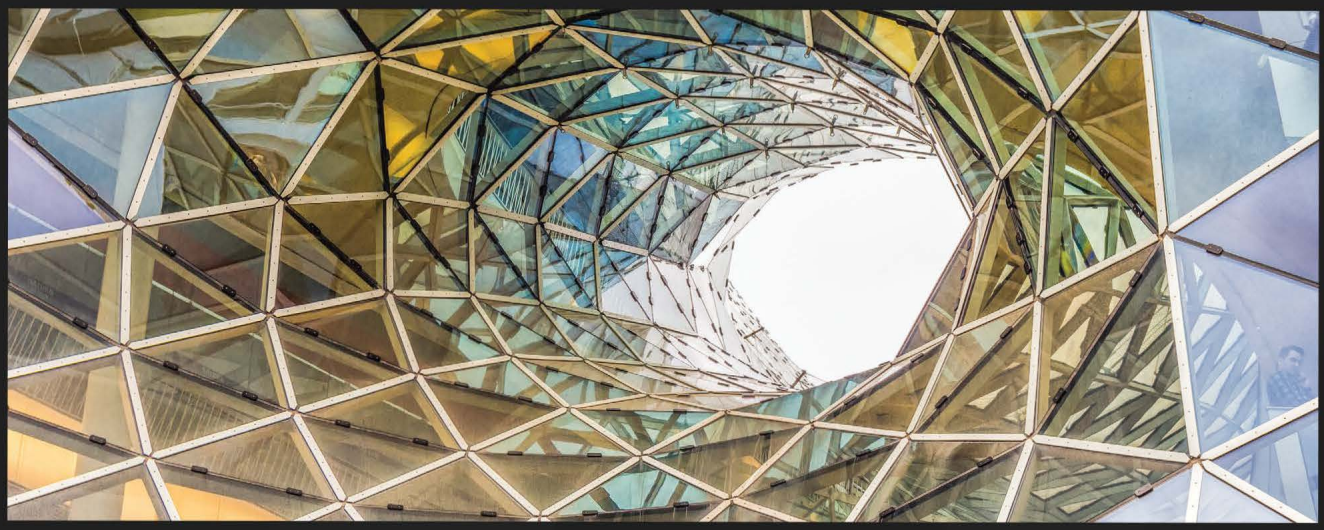

1ST EDITION

Microsoft Cybersecurity Architect Exam Ref SC-100

Get certified with ease while learning how to develop highly effective cybersecurity strategies

DWAYNE NATWICK

Foreword by Rod Trent, Security Cloud Advocate, Microsoft

Microsoft Cybersecurity Architect Exam Ref SC-100

Get certified with ease while learning how to develop highly effective cybersecurity strategies

Dwayne Natwick

‹packt›

BIRMINGHAM—MUMBAI

Microsoft Cybersecurity Architect Exam Ref SC-100

Associate Group Product Manager: Mohd Riyan Khan
Publishing Product Manager: Mohd Riyan Khan
Senior Editor: Divya Vijayan
Technical Editor: Rajat Sharma
Copy Editor: Safis Editing
Project Coordinator: Ashwin Kharwa
Proofreader: Safis Editing
Indexer: Tejal Daruwale Soni
Production Designer: Shankar Kalbhor
Marketing Coordinator: Marylou De Mello and Ankita Bhonsle

First published: January 2023

Production reference: 1091222

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham
B3 2PB, UK.

ISBN 978-1-80324-239-2

www.packt.com

Contributors

About the author

Dwayne Natwick is the Global Principal Cloud Security Lead at Atos, a multi-cloud GSI. He has been working in IT, security design, and architecture for over 30 years. His love of teaching led him to become a **Microsoft Certified Trainer** (**MCT**) Regional Lead and a Microsoft **Most Valuable Professional** (**MVP**).

Dwayne has a master's degree in Business IT from Walsh College, the CISSP and CCSP certifications from ISC2, and 18 Microsoft certifications, including Identity and Access Administrator, Azure Security Engineer, and Microsoft 365 Security Administrator. Dwayne can be found providing and sharing information on social media, at industry conferences, on his blog site, and on his YouTube channel.

Originally from Maryland, Dwayne currently resides in Michigan with his wife and three children.

About the reviewers

Dan Gora is currently a Lead Cloud Security Architect at Cloudreach, an Atos company. He has over a decade of experience in cybersecurity consulting. He has a broad set of experience in guiding highly regulated industries in securing their cloud transformation journey, adopting approaches such as DevSecOps and Zero Trust.

Dan is also involved in open source communities, where he is an OWASP City Chapter Lead and contributor to the Cloud Security Alliance. Dan has a master's degree in Secure Software Engineering and certifications including CISSP and CSSLP from ISC2, CCSK from CSA, as well as multiple Microsoft and AWS certifications.

Originally from Germany, Dan currently resides in Edinburgh, Scotland, with his partner.

Frank Grimberg has over 35 years of experience in delivering information systems solutions. His expertise spans cybersecurity, Azure, and web application development. His current focus is on providing CISO, cybersecurity architect, and **Digital Forensics Incident Response** (**DFIR**) services. He is the author of the Microsoft Official Technical curriculum for the SC-200 Microsoft Security Operations Analyst Learn content. His industry certifications include **GIAC Certified Forensic Analyst** (**GCFA**) and **Offensive Security Certified Professional** (**OSCP**). His Microsoft certifications include Cybersecurity Architect, Azure Solutions Architect, Azure Security Engineer, Security Operations Analyst, Identity and Access Administrator, and Microsoft Certified Trainer.

4

Designing an Identity Security Strategy

The previous chapter discussed the design for an identity security strategy and how to evaluate a strategy for security operations while utilizing the concept of Zero Trust. This chapter will discuss how to design an identity security strategy for cloud-native, hybrid, and multi-cloud identity and access management infrastructures. This will include understanding the design criteria and recommendations of a Zero Trust strategy for identity and access management for internal tenants, external customers and partners, and hybrid architectures.

In this chapter, we are going to cover the following main topics:

- Zero Trust for identity and access management
- Designing a strategy for access to cloud resources
- Recommending an identity store (tenants, B2B, B2C)
- Recommending an authentication and authorization strategy
- Designing a strategy for **conditional access** (**CA**)
- Designing a strategy for role assignment and delegation
- Designing a security strategy for privileged role access
- Designing a security strategy for privileged activities
- Case study – designing a Zero Trust architecture

Zero Trust for identity and access management

In the previous chapter, you learned about the design for an identity security strategy and how to evaluate a strategy for security operations. Managing and monitoring potential threats and vulnerabilities with security operations is aided by a Zero Trust strategy. This is increasingly true when determining

a strategy for recommendations to design a secure identity and access architecture. Let's review some of the foundational elements of Zero Trust and how they relate to securing identities.

Zero Trust is an integrated approach to securing access with adaptive controls and continuous verification across your entire digital estate. Everything from the user's identity to the application's hosting environment verifies the request and prevents a breach. To limit the impact of potential breaches, we apply segmentation policies, employ the principle of least privilege access, and use analytics to help detect and respond quickly.

Zero Trust is a security strategy. It is not a product or a service, but an approach to designing, adopting, and implementing defined, verified security principles.

Simplified, Zero Trust can be boiled down to three basic tenets, as defined by Microsoft:

- **Verify explicitly**: In other words, authentication and authorization should be reverified when a user or device is determined to have changed location, when a device has changed compliance or health, or if there is a behavioral anomaly detected. These should all be triggers for further identity verification.
- **Use least privilege access**: Limiting user access with just-in-time access and only the level of access or assigned role that is required to perform their defined tasks to protect both data and productivity.
- **Assume breach**: Minimize the blast radius for breaches and employ security strategies to prevent lateral movement.

Microsoft Zero Trust is currently defined by six pillars of coverage to which each of these principles can apply:

- **Identity**: Identity is the first pillar of Zero Trust, much like a cloud defense-in-depth strategy. Access should only be granted to people, devices, and resources that are authorized and required to complete a task.
- **Endpoints**: Assessing the security compliance of device endpoints is the next area to which Zero Trust applies, including **Internet of Things** (**IoT**) devices.
- **Applications**: The entry point of access to an application should use Zero Trust oversight. Access to resources that are required for the application such as databases or data sources should also require additional verification.
- **Network**: Zero Trust protections at the network layer for accessing resources are crucial – especially those within the corporate perimeter. This includes utilizing network and web application firewalls that provide deep packet inspection and application protection, respectively, to verify the integrity of the information that is coming into the network.

- **Infrastructure**: Applying Zero Trust principles to the infrastructure hosting your data on-premises and in the cloud is a step just beyond the Network pillar. This can be done with the overall hardening of operating systems, containers, or microservices that are part of the physical or virtual infrastructure to decrease the potential attack surface. Protecting access to ports through just-in-time access and bastion hosts can also decrease the likelihood of a security breach.
- **Data**: Protecting the integrity of your data for customers, people, and the overall business can be done through utilizing key management systems, such as Azure Key Vault, data loss prevention and data governance with Purview, and vulnerability scanning and threat detection on database platforms.

These six pillars create an end-to-end model for enforcing Zero Trust and decreasing the attack surface and the number of entry points for attackers or channels open to leak sensitive information. Optimizing identity and access management as the first line of defense in this strategy ensures that users are who they say they are at every access attempt, and regularly reaffirms their trustworthiness.

The next section will expand upon how the Zero Trust approach should be used to design a strategy for access to resources within cloud and hybrid environments in relation to identity and access.

Designing a strategy for access to cloud resources

As we determine how access to cloud resources will be handled, let's think about the evolution of identity and access architectures. Identity and access prior to cloud technologies were typically handled from application to application. Developers of on-premises applications would have their user database accessed from within the application. Every application would have a database of users and passwords. Windows AD user databases could be tied into these applications and could also be integrated into some of these applications. However, full single sign-on capabilities were not widely available.

As cloud technologies became available and **software-as-a-service** (**SaaS**) applications became more widely used by companies, the need to manage and govern identity and access expanded beyond the on-premises architecture. This expansion decreased the amount of control that IT departments had to protect identities and access behind network technologies, such as firewalls, **virtual local area networks** (**VLANs**), and **virtual private networks** (**VPNs**).

Modern authentication requires identity and access to be managed across on-premises networks, SaaS applications (Microsoft 365, Google Workspace, formerly called G-Suite, and so on), and cloud providers such as Microsoft Azure, Amazon Web Services, and Google Cloud Platform. These environments have data that includes **personally identifiable information** (**PII**), business-critical information, and intellectual property that needs to be protected. The identity and access architecture of these environments must take a Zero Trust approach to verify users and devices explicitly before authorizing access to resources. *Figure 4.1* provides a diagram that shows the relationship between the various resources to access, the access control, and privileged access for the administration of resources:

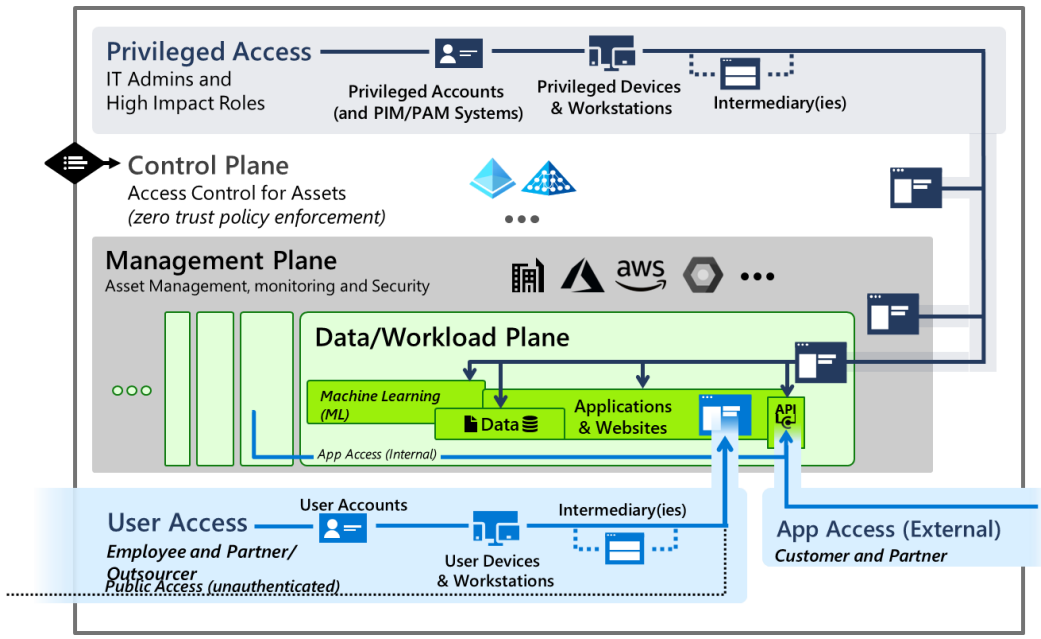

Figure 4.1: Identity and access control for secure access

Modern authentication has the following requirements for the access controls for resources and the parameters for privileged administrative access:

- Maintain a secure level of access to all resources by explicitly validating the trust of users and devices during access requests, using all available data and telemetry.
- Ensure that security assurances are applied consistently and seamlessly across the environments, including on-premises, SaaS, and cloud providers.
- The architecture should be comprehensive while enforcing an access policy as close to the resources and access entry points as possible.
- Controls that are in place should be identity-centric. They should prioritize the use of identity and related controls when available since the physical infrastructure is no longer a responsibility when working in cloud environments.

Coordinating the various environments to access adds additional complexity to the enforcement of Zero Trust. Traditional on-premises identity and access directories do not have **single sign-on** (**SSO**) integration with cloud providers and applications. In addition, these on-premises networks do not have visibility of identity threats within these expanded environments.

The architect of the identity and access requires a holistic review of the integration points in the hybrid environment. *Figure 4.2* shows the expanded view of the enterprise.

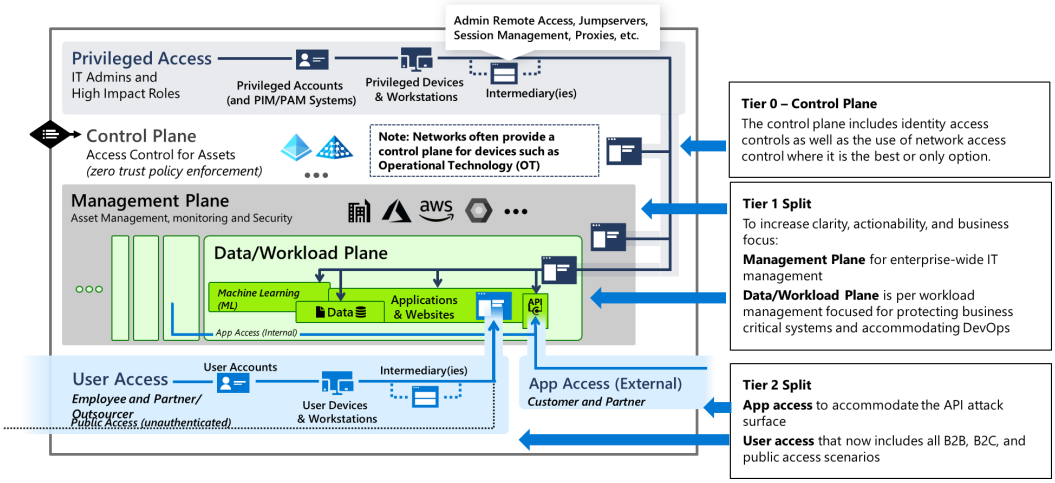

Figure 4.2: Identity and access model for the enterprise

The various means of user and app access need to be considered when determining the verification points for a Zero Trust framework. Taking the steps to protect these access points with Zero Trust will be determined as follows:

- Integration and federation are created between cloud identities and on-premises identity databases for a single management plane. To be able to enforce Zero Trust across the enterprise, you need to have a strong modern authentication architecture.
- **Azure Active Directory** (**Azure AD**) provides a method of enabling strong authentication along with a point of integration for endpoint security to multiple types of devices. Azure AD can also be the primary point of enforcing policies to verify and guarantee least-privilege access across the hybrid environment.
- Azure AD provides CA to enforce policy-based decision points for access to cloud and on-premises applications based on user identity, the access environment, device health, and potential risk. These policies provide explicit verification of identity at the point and time of access. CA policies provide a gate to verify access and identify areas for the remediation of potential vulnerabilities and risks. Planning and creating a strategy for CA will be discussed in the *Designing a strategy for CA* section of this chapter.
- The visibility of user and device activity through analytics can provide operational insights into the entire identity and access infrastructure. Logging within Azure AD provides reporting on authentication, authorization, and provisioning activities across users and devices. These logs can be used in Azure AD for monitoring and reporting as well as integrated into SIEM solutions, such as Microsoft Sentinel, for advanced hunting of potential identity attacks on the MITRE framework.

- Identity protection and identity governance monitor and manage the access privileges of identities. Identity protection determines the potential risks to users and sign-ins. Identity governance reviews the roles of users and manages their continued access through access reviews. Microsoft's machine learning capabilities can analyze various threat signals from security solutions to improve the overall detection, protection, and response to user and device identity and access.
- Identity and access architectures that span across on-premises and cloud providers can be integrated into Azure AD for a complete view and more effective management and governance. This can be accomplished with **Microsoft Defender for Identity** (**MDI**) and Microsoft Defender for Cloud Apps.
- MDI provides data on risky behavior for users that are accessing on-premises resources and applications, through either traditional or modern authentication methods. Microsoft Defender for Cloud Apps provides and discovers third-party cloud applications that are being accessed on the company network. This discovery allows you to verify and approve access to compliant applications.
- Utilizing Azure AD for identity and access provides a cloud-based identity and access management service for all Microsoft cloud solutions and an integration point for on-premises and cloud-based identity providers.
- Azure AD can provide SSO authentication, CA, verification with MFA, and password-less authentication. This adds new capabilities of provisioning users through automation, and provides additional enterprise-enabled features for protecting and automating identity protection and governance.

The next section will discuss recommendations for an identity store and how Azure AD can be used as an integration point for multiple identity providers for cloud and on-premises applications.

Recommending an identity store

As stated in the previous section, you need to understand the resources needed for accessing and designing an identity and access architecture with Zero Trust. The landscape of identity and access expands beyond the company member users. You need to understand the users that will be guests in your tenant, the companies that partner with your company for collaboration, and the customers that will be accessing your company resources, such as an e-commerce website or registration pages for events.

These scenarios pose potential risks to the organization as the identity and access environment expands beyond your own tenant. The foundational security policies and techniques that are used to protect member users should be maintained by anyone that is accessing the organization's resources. Users' identities, business partners, and customers should be protected through the security capabilities available within Azure AD.

Overall, users want flexibility when accessing resources. Azure AD provides these capabilities through foundational security features, such as **multi-factor authentication** (**MFA**). Expanded features to provide SSO for members, guests, and customers can be provided and allow access to applications that are published and verified through Azure AD. The following sections will discuss the various options for providing this flexibility in Azure AD.

Azure AD tenant synchronization with SCIM

For large companies that partner and need to federate with your tenant, the **System for Cross-Domain Identity Management** (**SCIM**) protocol can be used. SCIM is an open source protocol that can be configured with a user management API to synchronize and automatically provision users and groups to an application in Azure AD. *Figure 4.3* provides a diagram of how this works across multiple identity providers:

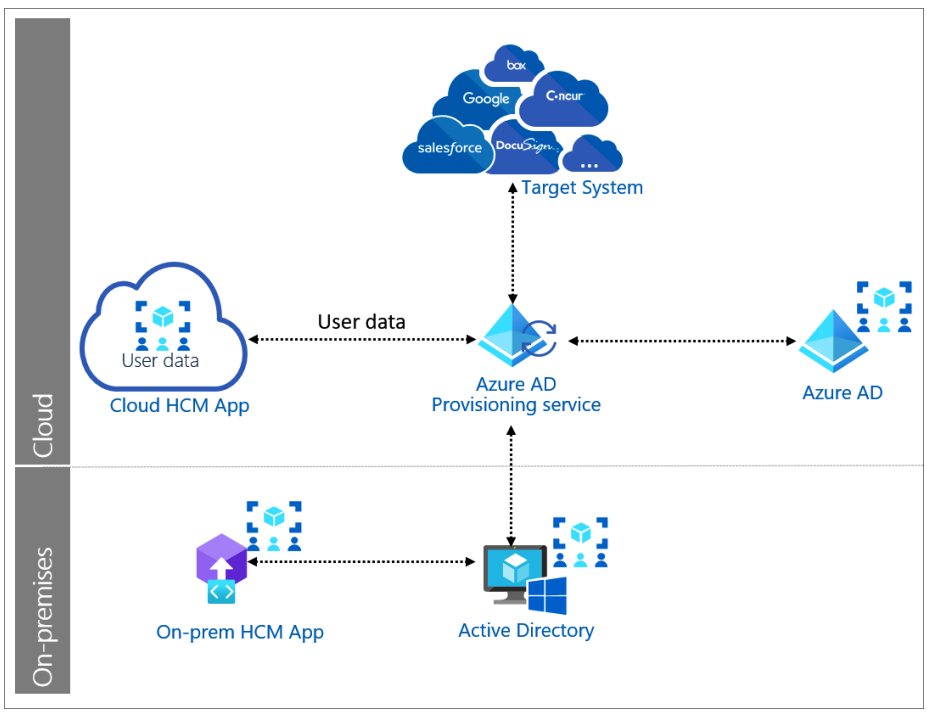

Figure 4.3: SCIM Azure AD provisioning diagram

For more information on how to use SCIM in Azure AD, please go to the following link: `https://docs.microsoft.com/azure/active-directory/fundamentals/sync-scim`

SCIM synchronization with Azure AD is a good option for integrating companies that need cross-platform access to resources in a merger and acquisition scenario. For other partner relationships between two companies, **business-to-business** (**B2B**) access can be used.

B2B

B2B guests are best described as a partnership relationship between users within two separate companies that need to collaborate on a project. These B2B relationships may be created through business mergers and acquisitions, project needs, or support relationships.

Within these B2B relationships, an external company may bring its own Azure AD and Microsoft 365 licenses for collaboration. Otherwise, these licenses can be assigned if external users do not come from an Azure AD tenant with Microsoft 365 licensing. External company collaboration settings provide these users with an SSO experience for both business tenants. If an external user does not use these licenses, then the username and password can be used from another identity, and that user can be assigned licenses from within the invited tenant.

B2B provides the flexibility of access to resources for business collaborations. For customers that would like flexibility in completing forms or shopping on e-commerce sites, Azure AD can be used with registered applications through the configuration of **business-to-consumer** (**B2C**) access.

B2C

The B2C relationship is an external account used for customers that are accessing applications and resources within the tenant. An example of this would be using LinkedIn, Facebook, or Google credentials to log in and access an account on a shopping site. This provides convenience to customers by not requiring them to create another username and password. These identity provider relationships enable B2C authentication relationships.

When architecting identity and access with Zero Trust, the more that you can integrate the security features of Azure AD into the methods for accessing resources within your tenant, the better. The methods described in this section provide these capabilities to enforce Zero Trust features, such as MFA and CA.

In the next section, you will learn more about choosing a strategy for authentication.

Recommending an authentication and authorization strategy

In current organizational infrastructures, authentication and authorization to resources are not limited to cloud-only users. Many companies have applications that are still in on-premises data centers that users require access to. This provides additional challenges to enforcing the modern authentication techniques for Zero Trust. When using Azure AD for authentication and authorization to cloud resources, you should also determine the proper techniques for users to access on-premises resources. Azure AD Connect provides this capability for SSO to cloud and on-premises resources, but you need to determine the best method for your company to synchronize and manage these hybrid users.

Hybrid identity infrastructure

The term hybrid identity is meant to signify that the company has users that use on-premises resources, and users that use cloud-native resources. Within this hybrid identity infrastructure, there is going to be an on-premises Windows AD domain controller that is used to manage the on-premises users, and Azure AD, which manages the cloud-native users, both members and guests. This infrastructure coincides with companies that have a *hybrid cloud* approach. Many companies have Windows AD domain controllers in place today. Azure AD Connect provides a hybrid infrastructure connection to Azure AD.

Azure AD Connect is a software solution that is installed within the on-premises infrastructure and configured to synchronize users and groups to Azure AD. Azure AD Connect simplifies the management of these users and groups by providing ways that an identity and access administrator can manage users in one interface and have the changes updated in near real time.

Since there are structural differences in how **Active Directory Domain Services** (**AD DS**) and Azure AD are built, Azure AD Connect provides the means to create a consistent user and administrator experience for identity and access management.

There are some prerequisites and aspects that are out of scope within Azure AD Connect that you should understand. The installation prerequisites can be found at this link: `https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-install-prerequisites`. When planning to implement Azure AD Connect, the following information should be understood and planned accordingly:

- Azure AD Connect synchronizes users and groups, not devices or applications. There are ways to co-manage devices and support on-premises application access within Azure AD. Both of these will be covered in later chapters.
- Azure AD Connect synchronizes a single AD DS forest per Azure AD tenant. If there are multiple forests, then multiple tenants will be required in Azure AD.

There are additional considerations in planning for AD DS and Azure AD synchronization with Azure AD Connect, but these will be covered within the use cases for each synchronization type.

There are three options when configuring Azure AD Connect for synchronization:

- **Password hash synchronization** (**PHS**)
- Pass-through synchronization
- Federation with **Active Directory Federated Services** (**AD FS**) synchronization

Each of these options has its own unique uses. The following subsections will detail these uses and options to consider when choosing one to use within your company.

PHS

PHS is the easiest to configure and is the default option within Azure AD Connect's express setup. PHS maintains both the on-premises and cloud identities of users. This takes place by providing on-premises user identities to Azure AD along with an encrypted hash of their passwords. This allows users to sign into on-premises and cloud applications with the same authentication credentials.

PHS is a good option when a company has a single on-premises domain and is moving quickly to a cloud-native infrastructure. PHS is not for companies with complex authentication and password requirements within an on-premises AD.

As previously stated, PHS maintains authentication credentials on-premises and in Azure AD. Therefore, PHS can have users authenticate to cloud applications through Azure AD, while passing authentication responsibilities to on-premises applications to on-premises AD. The benefit here is that if the connection fails in Azure AD Connect between the on-premises AD DS and Azure AD, users are still able to authenticate to their cloud applications and remain partially productive. *Figure 4.4* provides a visual diagram of how this workflow is handled:

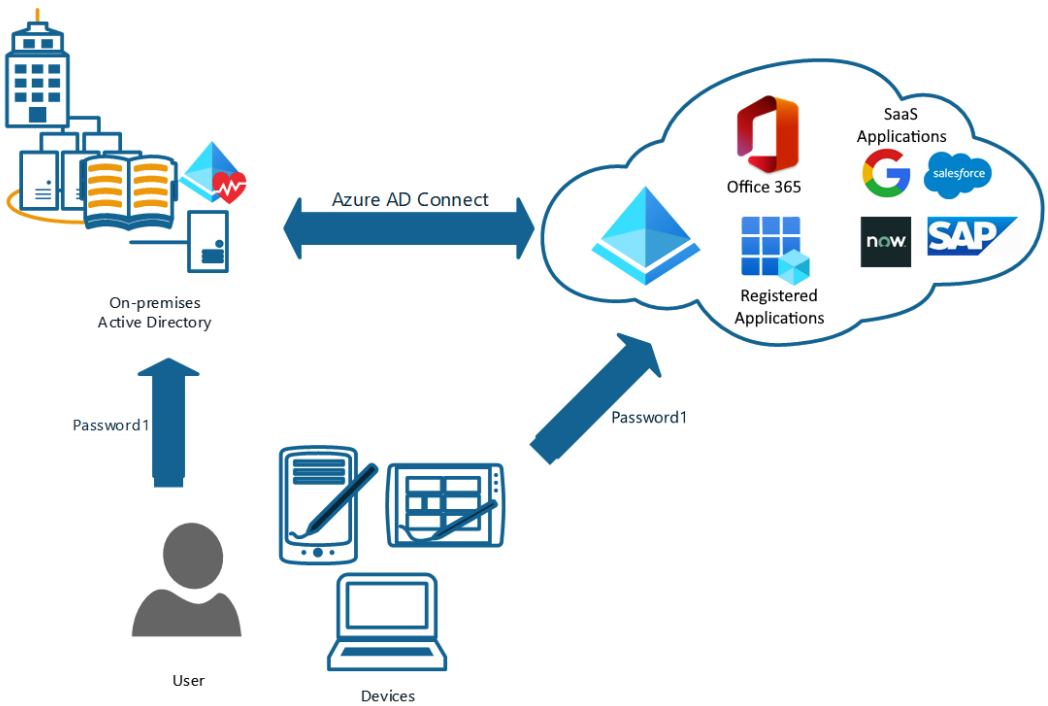

Figure 4.4: An overview of PHS

Additional information can be found at this link: `https://docs.microsoft.com/en-us/azure/active-directory/hybrid/whatis-phs`

This configuration for Azure AD Connect is the least complex of the three options and should be preferred for a cloud-native authentication approach. The next sections will provide information about pass-through synchronization and AD FS synchronization.

Pass-through synchronization

The next hybrid identity synchronization option is pass-through synchronization. Unlike PHS, which allows user identities to be authenticated in either the on-premises AD or Azure AD, pass-through synchronization requires all users to authenticate to the on-premises AD.

In this configuration, if the Azure AD Connect connection between Azure AD and on-premises AD were to become disconnected, no users would be able to authenticate to on-premises or cloud resources. Therefore, it is important to actively monitor this connection and build resiliency in the architecture. *Figure 4.5* shows a diagram of how pass-through synchronization functions and how you can build resiliency with redundancy in **pass-through agents** (**PTA**) and a backup domain controller:

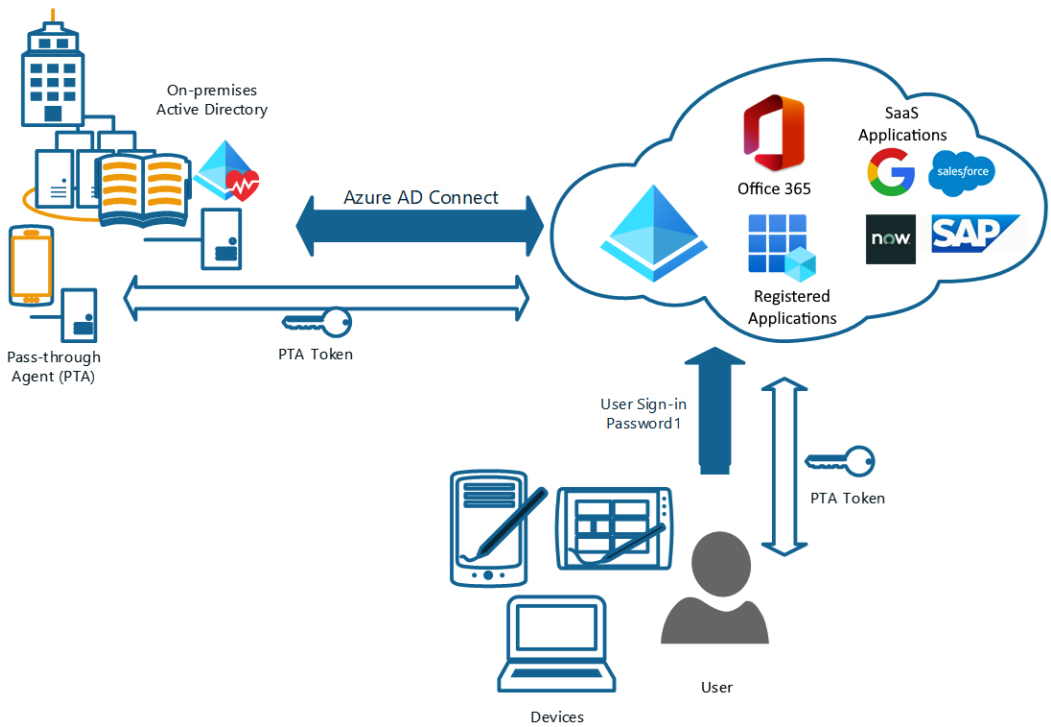

Figure 4.5: An overview of pass-through synchronization

There are some reasons to utilize pass-through synchronization. An organization may require authentication parameters and limits that only allow users to access resources during certain times. Such rules can only be configured currently on an AD domain controller. With pass-through

synchronization, you can utilize modern authentication features with Azure AD, such as MFA and **self-service password reset** (**SSPR**). However, you will need to enable the password writeback feature within Azure AD Connect to use SSPR. Password writeback will allow the Azure AD password to be written to the on-premises AD. If this is not enabled, the AD password will always take precedence.

In order to have a resilient architecture, it is recommended that at least two of these PTAs are installed on member servers in the on-premises infrastructure. Microsoft's recommendation is that up to four PTAs should be deployed. The next section will discuss the third and final synchronization type within hybrid identity infrastructures, AD FS.

Federation with AD FS synchronization

Federation with AD FS synchronization is the most complex of the three Azure AD Connect synchronization types. AD FS requires additional infrastructure in place to support the authentication process. In comparison, PHS and pass-through authentication can be installed directly on the on-premises domain controller in many cases. *Figure 4.6* provides an overview that shows the complexity of the infrastructure and necessary components:

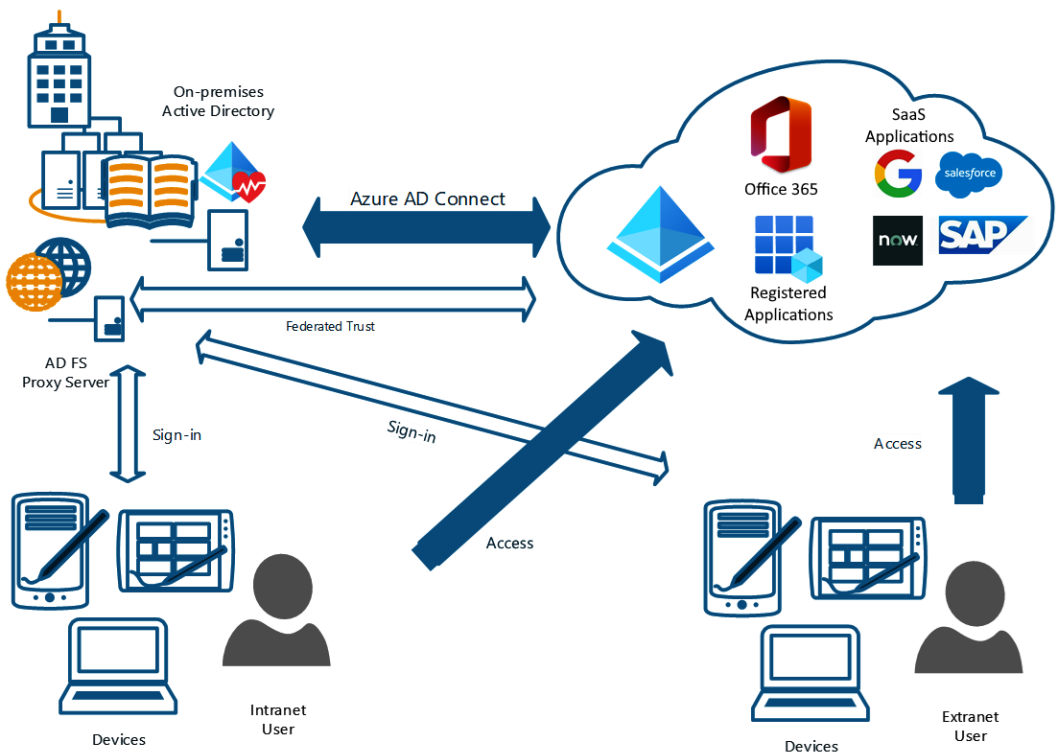

Figure 4.6: An overview of AD FS synchronization

AD FS synchronization is utilized in complex AD infrastructures where there are multiple domains, and third-party MFA solutions or smart cards are utilized. For additional information on the configuration of AD FS synchronization, you can read more here: `https://docs.microsoft.com/en-us/azure/active-directory/hybrid/how-to-connect-fed-management`.

PHS and pass-through synchronization are the more widely discussed of the three Azure AD Connect options. Federation with AD FS is necessary when utilizing third-party solutions that are not native to Microsoft, such as third-party MFA.

The next section will discuss additional methods for secure authorization that should be considered when designing identity and access security with Zero Trust.

Secure authorization methods

Once you have determined the method that you will be using for authentication within your hybrid infrastructure, you should consider how to authorize users to access resources. The synchronization methods for authentication allow you to utilize Azure AD security methods for authorization. These methods include the following:

- **Security group membership**: Planning for users to be assigned roles and authorized access based on group membership decreases the management overhead for identity and access administrators
- **Role-based access control**: Authorization for access to resources should be based on the role of the user and the level of access should be specific to the tasks that they need to perform

Properly assigning these authorization methods allows for a more secure environment for managing and monitoring access to resources.

The next section will discuss how to use CA for the dynamic enforcement of Zero Trust.

Designing a strategy for CA

CA policies enforce additional verification actions based on a signal that a user or device may be potentially compromised. The foundation of CA policies is the Zero Trust methodology. Azure AD CA analyzes signals such as user, device, and location to automate decisions and enforce organizational access policies for the resource. CA policies allow you to prompt users for MFA when needed for security and stay out of the user's way when not needed.

As you will notice in *Figure 4.7*, the policies that we determine for our company are what then enforce these CA requirements from signal to decision to enforcement:

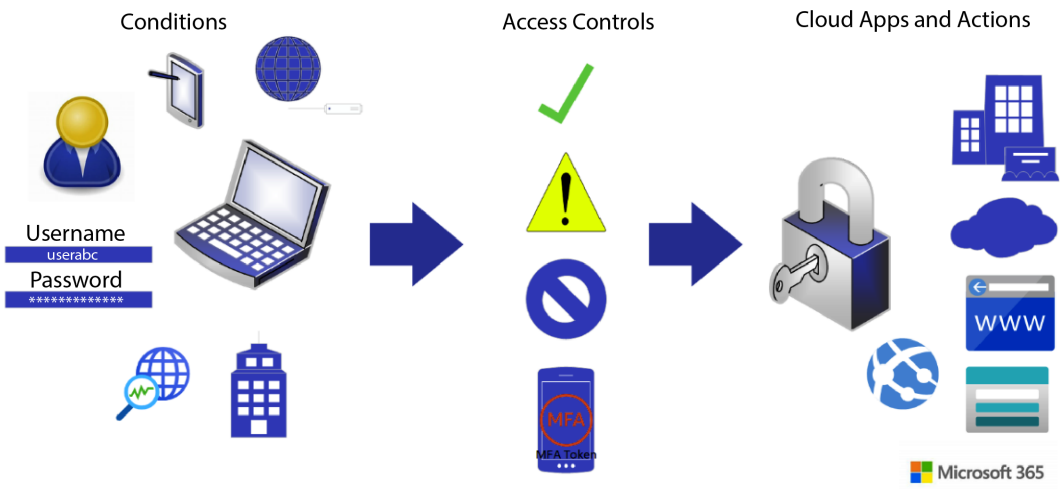

Figure 4.7: CA workflow

The planning and creation of CA policies should be a foundation of access policy enforcement in Zero Trust. In addition, you should have a set of active and fallback policies to start your deployment. You need to have a proper plan and understand how conditional access policies would potentially affect the user experience. There is a balance that a company should attempt to maintain between the enforcement of policies to secure and protect data. This includes the ability for a user to have access to the applications and data that they need to be effective at their required tasks.

Before we can create a CA policy, we need to meet the prerequisites. There are a couple of areas that we need to address to implement this solution, that is, licensing and security defaults. For licensing, CA policy features are available with an Azure AD Premium P1-level license. This level of Azure AD licensing is included with Microsoft 365 Business Premium, Office 365 E3/A3, Microsoft 365 E3/A3, Office 365 E5/A5, and Microsoft 365 E5/A5. These licenses must be assigned to the users that we are attempting to enforce CA policies for.

The full list of licensing requirements can be found at this link: `https://docs.microsoft.com/azure/active-directory/conditional-access/overview`

In addition to having the proper licenses, we will be required to turn off the Azure AD security defaults. Security defaults are a feature that is turned on when we create our Azure AD tenant. This provides a baseline level of protection to require. Examples include when users enroll in MFA, enforce MFA for administrators, and block the use of legacy authentication for apps. Security defaults will be discussed further in *Chapter 7*. To be able to implement CA, navigate back to **Security Defaults** and turn them off. When you do this, there will be a list of reasons that appear and you will see the selection for using CA policies.

Once we have the proper licensing assigned and security defaults turned off, we can begin planning for CA policies. Some commonly used CA policies can be found on the Microsoft Docs website at this link: `https://docs.microsoft.com/azure/active-directory/conditional-access/plan-conditional-access`.

The key to planning for CA is to understand the following:

- The groups of users that access company applications and data
- The devices that they are using to access company applications and data
- The locations where they may be accessing company applications and data
- The applications that are being used to access the company data

Once a CA policy is created, you should test it in report-only mode using the **What-If** feature to verify the policy is working correctly against users, devices, and applications. What-If testing will also show other policies that may overlap and cause a conflict with user access.

As we continue to design and plan access to resources, the assignment of roles and delegation of roles to resources in a Zero Trust architecture is needed. The next section will discuss how to design a strategy for role assignments and delegation to access resources.

Designing a strategy for role assignment and delegation

Managing role assignments and how users are delegated access to resources can be a source of concern for security and governance. To alleviate the issue of users requesting access and being assigned a role on a one-to-one basis, planning a strategy for role assignments by groups of users creates a more manageable environment.

Security groups can be created for departments or project groups that include dynamic assignments. These groups can be assigned the roles defined by management and supervisor stakeholders for members of the groups. When a user is then manually or dynamically assigned to that group, they inherit that role. When users are removed from the group, the role no longer is available to them. This allows you to better manage and govern the access levels and roles that users in your company have to resources. The differences between Microsoft 365 groups and security groups can be found at this link: `https://learn.microsoft.com/en-us/microsoft-365/community/all-about-groups`.

This becomes very important when you are planning for the privileged roles of administrators in your tenant. The next section will discuss the planning of this strategy for these elevated user privileges.

Designing a security strategy for privileged role access

Designing a secure strategy for privileged roles will allow you to protect and defend identity and access by utilizing the concept of Zero Trust and the principle of least privilege to assign authorization for

administrator accounts. You should have a clear strategy with defined job tasks for every administrator user account to plan for the proper assignment of these roles. This strategy should include meeting with stakeholders and discussing the roles that each department member requires to complete their job tasks. In addition, you should be monitoring the activity of these accounts and verifying the continued requirement for users to have these privileged access roles.

To enforce the concept of Zero Trust, you have the capability to assign CA policies to these accounts. To address and protect privileged assignments, Azure AD provides **Privileged Identity Management** (**PIM**) within the Identity Governance solutions.

Azure AD PIM

PIM provides just-in-time privileged access to users. Since users are only provided active administrator roles for a short window of time, this reduces the attack surface and potential for these user accounts to cause exposure to privileged access in an attack. PIM provides an approval and justification process for activating privileged role assignments, which includes notifications when a role is activated and an audit trail of these activations.

PIM requires an Azure AD Premium P2 license. To assign PIM to member accounts, each user must have this license. However, for guest users that require privileged access with PIM, five guests can be assigned PIM roles for each Azure AD Premium P2 license that you have in your tenant.

When planning for Azure AD PIM, you should consider the following options:

- Roles should be assigned utilizing the principles of least privilege. The assignment should use a role that provides the authorized level of access that is necessary to perform the task and no more.
- PIM should be used for just-in-time access to roles and should only be active for a specified and limited amount of time. Administrator roles should not be permanently assigned; they should be eligible by request.
- Administrator accounts should have MFA enabled and enforced without exceptions. These accounts should also be cloud-native accounts and not accounts that are synchronized with Windows AD.
- Recurring access reviews should be used to verify that access to the assigned roles is still necessary for the users and groups.
- Global administrators within the company should be limited to five (5) users. This protects the attack surface in the case of a compromised account and also protects against users that may leave the company with elevated privileges.

To protect against a potential lockout and ensure access is still available in a potential emergency situation, you should configure at least two emergency access or "break-glass" accounts. These accounts are accounts of high privilege with access at the level of a global administrator. These accounts are not protected with MFA, so they can gain access quickly to resources when other administrator accounts

cannot gain access. The use of these accounts should be limited to this scenario and the credentials should be locked away until the time that they are absolutely needed.

Break-glass accounts are member accounts that are tied directly to the Azure AD tenant. Therefore, they can be utilized in situations where federated identity providers are being used for authentication and there is an outage to that identity provider. Other use cases would be that the global administrator has lost access to their MFA device to verify their identity, a global administrator has left the company and it is necessary to delete that account, and a storm has taken down cellular services and you cannot verify the identity with MFA.

Additional information on emergency access or "break-glass" accounts can be found at this link: `https://docs.microsoft.com/en-us/azure/active-directory/roles/security-emergency-access`

The next section will provide additional ways to manage and govern the activities of privileged users and their access to resources.

Designing a security strategy for privileged activities

In terms of access life cycle, you should consider the access life cycle of your member users and your guest users, and especially your privileged users. These should be handled differently as the life cycle of member users is based on their employment within the company and the access that is required for the department or team that they belong to. Guest users are provided access based on a partnership and external collaboration trust relationship.

Privileged access reviews

Privileged user access should be regularly reviewed in a similar manner. Since these are elevated access assignments, the review of these should be done on a consistent basis as identified by the company. Unused and unnecessarily privileged assignments should be removed. Automated removal should also be configured for users that are no longer with the company or have changed departments within the company. In the next section, you will learn how entitlement management can be used to govern access to resources for internal and external users.

Entitlement management (aka permission management)

Entitlement management provides this governance through the creation of catalogs and access packages that you can build for these groups of users. Entitlement management is found under **Identity Governance** within Azure AD. In entitlement management, the catalogs define the groups and teams, applications, and SharePoint sites within Identity Governance. Creating a catalog does not establish access to these catalogs. You must go through the creation of access packages to approve and allow access to these catalogs.

Before creating catalogs and access packages, you should plan and determine how these are going to be used within your company. Entitlement management can be a helpful tool for companies that have projects that utilize internal and external users, departments that utilize different and specialized resources that other departments don't require access to, and branch and global offices that have their own users, groups, and partners.

Being in charge of Identity Governance, it is important to work with stakeholders to plan these catalogs and access packages, as well as to determine how often they will be reviewed for continued use and access. Proper planning with these stakeholders will allow them to quickly provide users access to the resources that are required for a given project or department once they are onboarded.

Like privileged access reviews, entitlement management can use access reviews to verify continued membership to access packages.

The next section will discuss how to bring together the use of privileged access management, access reviews, and role permissions to administer identity and access across your tenant.

Cloud tenant administration

The previous sections discussed many of the strategies that can be used to enforce and verify user access through Zero Trust. This is a multi-staged strategy across your tenant to plan and implement. As a cybersecurity architect, you should work with the various stakeholders to properly plan and implement these strategies in stages.

These stages can be broken down and prioritized as follows:

- *Stage 1*: Critical items that we recommend you do right away. This stage should be executed in the first 24-48 hours of creating your cloud tenant. Steps within this stage include the identification of privileged roles that should use PIM and the creation of break-glass accounts.
- *Stage 2*: Mitigate the most frequently used attack techniques. In this stage, you should spend 2-4 weeks addressing the common attacks, such as brute-force and phishing attacks. You should determine a proper hybrid identity synchronization strategy and set up identity protection to recognize potential identity attacks. The incident response process should begin to be developed with the assignment of owners.
- *Stage 3*: Build visibility and full control of administrator activity. The next 1-3 months should begin to see the full control of PIM for administrators, access reviews, and utilizing standards, such as those of the **National Institute of Standards and Technology** (**NIST**), for recommendations on incident response procedures.
- *Stage 4*: Continue building defenses to further harden your security platform. This stage is ongoing for reviewing and improving security controls and validating response plans. In this stage, you should determine ways to manage devices with endpoint management, and you should have all company-owned devices compliant and Azure AD joined.

Additional information on these stages can be found at this link: `https://docs.microsoft.com/en-us/azure/active-directory/roles/security-planning`

The next section will provide a case study that summarizes the Zero Trust architecture discussed in *Chapters 2*, *3*, and *4*.

Case study – designing a Zero Trust architecture

In this section, you will be given a company scenario and asked to complete a number of tasks to meet the requirements of the Zero Trust architecture.

Company ABC has concerns across its Azure, on-premises, and SaaS applications architecture. They have come to you for assistance in addressing their security concerns. They want you to provide suggestions on how they can use the security capabilities within Azure, Azure AD, and Microsoft 365 to enforce Zero Trust methodologies across the company's technology infrastructure.

The areas of concern and requirements include the following:

- The utilization of strong modern authentication techniques for all users, including cloud-native and on-premises directory users
- Guest users should be only allowed to access resources that are assigned to them, and they should be regularly reviewed
- Administrative users should have just-in-time access to privileged resources and all access should be justified and audited
- When accessing applications that contain sensitive information, users should be required to verify their identity
- All user identities should be protected from common attacks
- Users that are accessing company resources from potentially dangerous locations should be forced to re-authenticate
- Devices should be verified with proper security patches before accessing company resources
- Users should be analyzed for anomalous behavior and potential threats to identities
- Network resources with sensitive information should not be accessible through publicly available connections
- All activity and event data should be logged and can be reviewed
- Reports can be generated for executive reviews and incident handling

Suggested responses to this case study will be provided in *Chapter 11, Case Study Responses and Final Assessment/Mock Exam*.

Summary

In this chapter, we discussed the design and strategy for creating a Zero Trust architecture for identity and access. This included an overview of Zero Trust for identity and access management and how to design a strategy for access to cloud resources. We then learned ways to recommend an identity store for hybrid and guest access and recommend an authentication and authorization strategy.

Finally, we learned about the various strategies for designing CA policies, determining role assignments and delegation, handling privileged role access, and reviewing and governing privileged activities. We then wrapped up the chapter with a case study to provide design and architecture suggestions for Zero Trust for users, devices, and networks.

In the next chapter, you will learn how to design a strategy for regulatory compliance.