Segmentation and Shared Active Directory of a Cardholder Data Environment

This is something that many QSAs (Qualified Security Assessors) are unable to agree on and I've been involved in many discussions for and against utilising a shared Microsoft Windows Active Directory (AD) infrastructure. As I've already written about, this is something that in the past I have not allowed because when I used to do penetration testing, the ability to compromise the whole Windows Domain from simply gaining administrative access to a domain joined host was usually trivial. That said, one could argue that this was due to the configuration of AD and the security of the Windows systems supporting it.

From a PCI DSS perspective, the PCI SSC guidance "Information Supplement: Guidance for PCI DSS Scoping and Network Segmentation" (https://www.pcisecuritystandards.org/documents/Guidance-PCI-DSS-Scoping-and-Segmentation_v1_1.pdf) (*Referred to as the PCI SSC Scoping Guidance herein*) includes the use of Microsoft AD as a 'Shared Service', meaning that the same Microsoft AD can be used to service the in-scope Cardholder Data Environment (CDE) and the not-in-scope environments. This still finds me somewhat uneasy and I would still highly recommend designing a separate Microsoft AD environment to support the CDE and non-CDE environment as sharing a Microsoft AD environment will introduce unnecessary risks to the CDE and the cardholder data (CHD) being handled. Mis-configurations over time to a shared Microsoft AD will only increase the chances of a compromise, therefore this is an important risk that must be managed well to ensure the environment stays secure.

That said, I have set off on this path to identify if there is a way to configure Microsoft AD and it's supporting Windows infrastructure in a way that will provide increased protection against compromised untrusted domain joined hosts from leading to a complete compromise of Microsoft AD and the Domain Controllers (DCs). This article will discuss the approach and findings to this research in the hope of providing further guidance for hardening Microsoft AD to reduce the risks of this type of deployment if an organisation is adamant on deploying Microsoft AD in this manner.

Objective

I've discussed my approach to 'Scoping and Segmentation' in a previous post (here) (https://pciramblings.com/2018/09/10/my-scoping-and-segmentation-methodology/), if you haven't already read it I would recommend you do so. As a recap, I discuss an approach which was originally introduced within the "Open PCI DSS Scoping Toolkit" (https://www.isaca.org/Groups/Professional-English/pci-compliance/GroupDocuments/OpenPCIScopingToolkit.pdf) (*Referred to as the Open Scoping Toolkit herein*) and which maps to a similar concept within the PCI SSC Scoping Guidance. Both documents discuss the concept of 3 categories of network segments of different security boundaries. These are defined as follows;

- Page 10 of PCI SSC guidance document discusses three categories of systems as follows;
    - CDE Systems,
    - Connected-to or Security-impacting Systems (Shared Service),
    - Out-of-Scope Systems.
- Open PCI DSS Scoping Toolkit as follows;
    - CAT 1 (i.e. CDE Systems),
    - CAT 2 (i.e. Connected-to or Security-impacting Systems),
    - CAT 3 (i.e. Out-of-Scope Systems)

To make the post easier to read, I'll be using the terms CAT 1, CAT 2 and CAT 3, however these are synonymous to the three categories (CDE Systems, Connected-to/Security-impacting systems and Out-of-Scope Systems) identified within the PCI SSC Scoping Guidance as detailed above.

The key concept within these two documents is that the CAT 1 and CAT 3 segments should never have direct connectivity. "Shared Services" can operate within the CAT 2 environment, which can serve both the CAT 1 and CAT 3 segments. CAT 1 and CAT 2 are in scope for PCI DSS assessment activities, whilst, typically the CAT 3 segment remains out-of-scope for PCI DSS assessment activities.

The PCI SSC guidance document includes Microsoft Active Directory as a "Shared Service" which can service both the CAT 1 and the CAT 3 segments, which, as I've already discussed, I have always felt uncomfortable with. The objective of this research and this article is to assume that the CAT 3 segment and all 'not-in-scope' (or untrusted) domain joined hosts within this segment have been compromised and a threat actor has "Administrative" control of one of (or all of) them. Administrative access may give the threat actor access through Firewalls / Access Control Lists (ACLs) providing segmentation between the various network segment security boundaries. Therefore, the goal is to identify a configuration of Active Directory that protects the CAT 2 and CAT 1 Active Directory infrastructure against a domain level attack, limiting the exposure to just CAT 3 devices.

The following paragraph on Page 6 of the PCI SSC Scoping Guidance says;

"The intent of segmentation is to prevent out-of-scope systems from being able to communicate with systems in the CDE or impact the security of the CDE. Segmentation is typically achieved by technologies and process controls that enforce separation between the CDE and out-of-scope systems. When properly implemented, a segmented (out-of-scope) system component could not impact the security of the CDE, even if an attacker obtained administrative access on that out-of-scope system."

Therefore, my objective is aligning to this premise of obtaining administrative access to the CAT 3 system(s), proving little to no risk to the environment, namely Active Directory, thus maintaining CAT 3 'out-of-scope'.

I'll be looking at techniques to do this without the need to implement mitigations on the untrusted CAT 3 devices themselves. Since we are trying to keep these hosts out-of-scope, if we start to implement mitigation controls on them, we are then bringing them into scope for some requirements which could then snowball into the whole of CAT 3 pretty much implementing all applicable PCI DSS controls thereby negating the point of them being not-in-scope.

Architecture

Due to the scoping and segmentation concepts discussed, a lot of organisations and QSAs will accept a design which is architecturally implemented in this manner, see the simplistic diagram below.

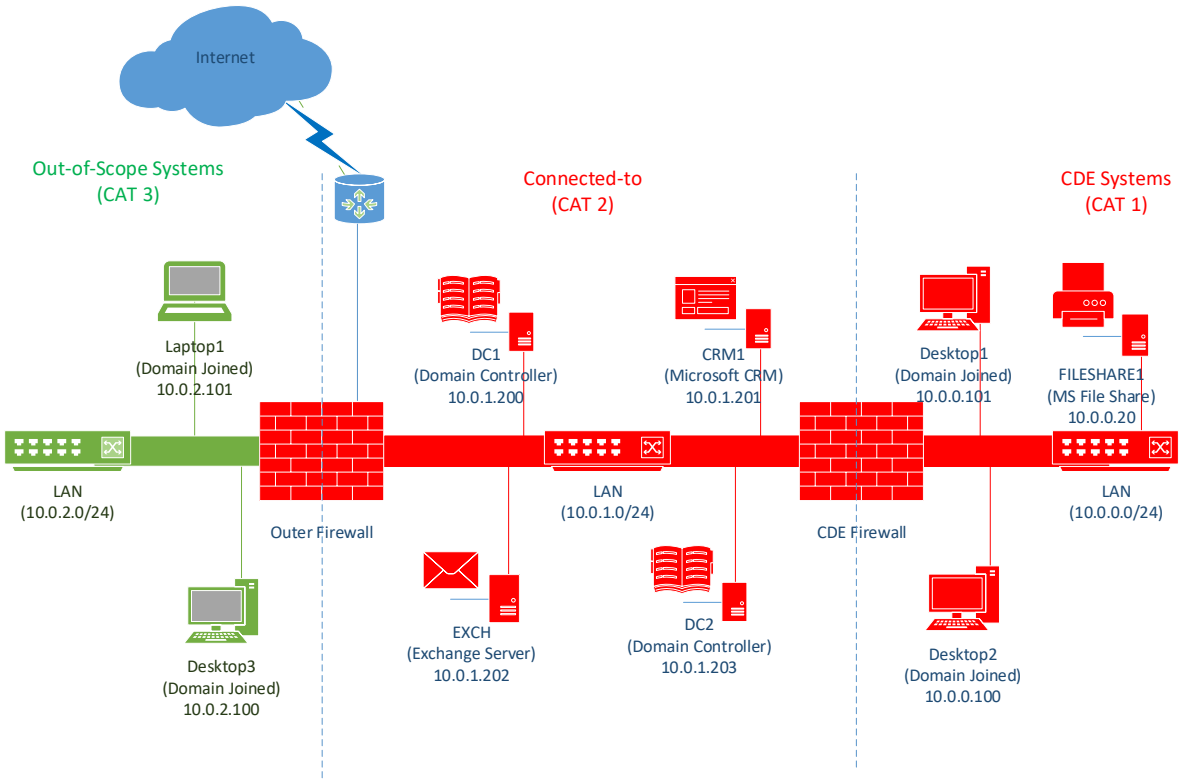

Simplistic Diagram

Domain Controllers could be deployed in one or many various locations across CAT 1, CAT 2 and CAT 3. Regardless of their deployments, providing Microsoft Active Directory Sites and Services is configured correctly, you can still restrict which DCs clients will utilise for authentication to limit blanket access through firewalls from many Domain Joined hosts.

For Microsoft Windows and Microsoft AD to function correctly, multiple services operating within a Windows infrastructure have various port requirements. The following table lists various port requirements which allow connectivity between Windows workstations and DCs and for intra-DC communications. These ports are used to support services such as, but not limited to; Active Directory, Distributed File System Namespaces, Distributed File System Replication, Distributed Transaction Coordinator and Group Policy. For a breakdown of ports and services, see the following Microsoft article; "Service Overview and Network Port Requirements for Windows" (https://support.microsoft.com/en-gb/help/832017/service-overview-and-network-port-requirements-for-windows);

Source Port	Protocol	Application Protocol	Destination Ports
1024-65535	TCP/UDP	DNS	53
1024-65535	TCP/UDP	Kerberos	88
1024-65535	TCP	RPC (Remote Procedure Call)	135

1024-65535	TCP/UDP	NetBIOS Name Server	137
1024-65535	UDP	NetBIOS Datagram Service	138
1024-65535	TCP/UDP	NetBIOS Session Services	139
1024-65535	TCP/UDP	LDAP (Lightweight Directory Access Protocol)	389
1024-65535	TCP	SMB (Server Message Block)	445
1024-65535	TCP	Microsoft Active Directory Global Catalog	3268
1024-65535	TCP	Microsoft Active Directory Global Catalog with LDAP/SSL	3269
1024-65535	TCP	RPC	5722**
1024-65535	TCP	Active Directory Web Services (ADWS) / Active Directory Management Gateway Services	9389
1024-65535	TCP/UDP	Dynamic Port Ranges	49152-65535*

* Windows 2000, XP and Windows Server 2003 utilise a smaller range of dynamic ports; 1025 to 5000

** Windows 2008/2008 R2 servers only

So, these firewall ports will be included between CAT 3 and CAT 2 and between CAT 1 and CAT 2 in the diagram above. These ports will likely need to be in both directions.

As you can see from the table, the port requirements are quite significant with some of these ports introducing many additional security concerns to the environment, especially SMB, RPC and NetBIOS ports which can allow threat actors to map out the Windows environment, brute-force user credentials or can expose vulnerabilities as these services are often the target of threat actors looking to find 0-day vulnerabilities.

Attacks/Risks

There are many attacks and mis-configurations that can lead to a compromise of a Microsoft AD infrastructure. Due to the nature of Microsoft AD and the way in which it controls user authentication and utilises trusts-relationships between other domains, forests and computers, there are inherent risks that can often be introduced. Some of the most common issues include;

Over privileged accounts: A secure design of Microsoft AD, or any other authentication system usually includes a least privilege concept. This ensures users' accounts are configured with the absolute minimum privileges required for the user to carry out their job function. It is often the case that Microsoft AD accounts are configured with more privileges than needed and within larger environments, you can often see 100's or even 1000's of accounts like this. This increases the attack surface area and the probability of compromising an over privileged account that can be used to gain full administrative control over the entire Microsoft AD infrastructure.

Privileged accounts with weak credentials: Linked to the above, there is a risk of privileged accounts having weak passwords that can be easily guessed and therefore compromised. If a threat actor can gain access to a privileged account, this could provide an attack vector that leads to a compromise of the entire Microsoft AD infrastructure.

Password reuse: This can pose a real problem within Microsoft AD environments. Often, passwords will be re-used across multiple accounts and across multiple devices. If not managed correctly, it only takes a compromise of a single system to then yield a privileged account password that is used across many devices or either the entire environment. An attacker can leverage this to then compromise the entire Microsoft AD infrastructure.

Abuse of Delegation Tokens: An issue within a Microsoft AD environment is that delegation tokens can often be cached on domain joined servers and workstations. This becomes a massive issue if there is inadequate patching which exposes domain joined machines to privileged access. Finding a domain joined machine with a domain administrative delegation token provides an attacker with administrative privileges across the whole domain. With a delegation token, it allows the computer or service to present the credentials of that user via their delegation token. Therefore, if a delegation token of a domain administrator is located, the delegation token can be used as if you were that domain administrator.

Pass-the-hash attacks: With pass-the-hash attacks, if the user's password hash is obtained (for example, by dumping a local machines password database), Windows will allow a user to present the hash of their password instead of the actual password. This negates the need of an attacker cracking a user's password since they can simply just present the hash obtained, authenticating as that user.

SMB Relay attacks: This is an old problem dating back to 1997 in a paper by Dominique Brezinski called "A weakness in CIFS Authentication" (http://diswww.mit.edu/menelaus.mit.edu/BoS/23). This man-in-the-middle attack allows an attacker to authenticate to a domain resource as a valid user. The attack works by waiting for the start of a user's SMB authentication process, then routing this traffic through the attacker's computer, modifying the request so it is coming from the attacker but with the user's credentials, thus authenticating to the resource as that user. A good write up of this can be found here. (https://cqureacademy.com/blog/penetration-testing/smb-relay-attack)

Kerberoasting attacks: This attack relies on having a valid user account to request service tickets for service accounts via their Service Principal Name (SPN) values. Active Directory will return an encrypted ticket, which is encrypted using the (NT Lan Manager) NTLM hash of that account. This can then be used to perform an off-line brute force attack to obtain the service account password in plain text.

As discussed within Microsoft's "Best Practices for Securing Active Directory" (https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/best-practices-for-securing-active-directory), avenues to compromise also includes; outdated applications and operating systems, missing patches and gaps in anti-virus and anti-malware deployments. These elements are not specific to this discussion; therefore, these other attack vectors are not discussed further within this document which will purely focus on specific configuration and techniques to help secure Microsoft AD deployments.

Targeted Accounts

Within Microsoft AD, there are three **main** privileged groups that are targeted by threat actors, EA (Enterprise Admins), DA (Domain Admins) and BA (Built-in Administrators) groups, DNSAdmin is also targeted which is discussed later. Users of these groups are the target of threat actors since these groups of users have privileged access to the Active Directory. On workstations and member servers, the local Administrators group is also targeted as this group provides the attack with privileged access to the workstation/server.

Additionally, VIP (Very Important Person) domain accounts are often targeted due to the types of data they have access to.

The remainder of this document discusses one possible design and various best practice configuration techniques which will help mitigate some of the threats to Microsoft AD helping to secure Microsoft AD as a 'Shared Service'.

One Possible Design Scenario

My issue with the simplistic design above is that domain controllers are located within CAT 2, therefore there are a lot of '*undesired*' ports required to allow functionality between the untrusted domain joined machines in CAT 3 and the Microsoft AD services in CAT 2. The following design can somewhat simplify this since a second DC (or multiple for fault tolerance) can be included within CAT3, however this DC would still be in scope for PCI DSS assessment activities. Additionally, since there is no firewall between the CAT 3 system components and the DC housed within CAT 3, this exposes the DC to additional threats. For this reason, it is recommended that the local software firewall be utilised to provide additional firewalling functionality (which would come into scope for PCI DSS Requirement 1).

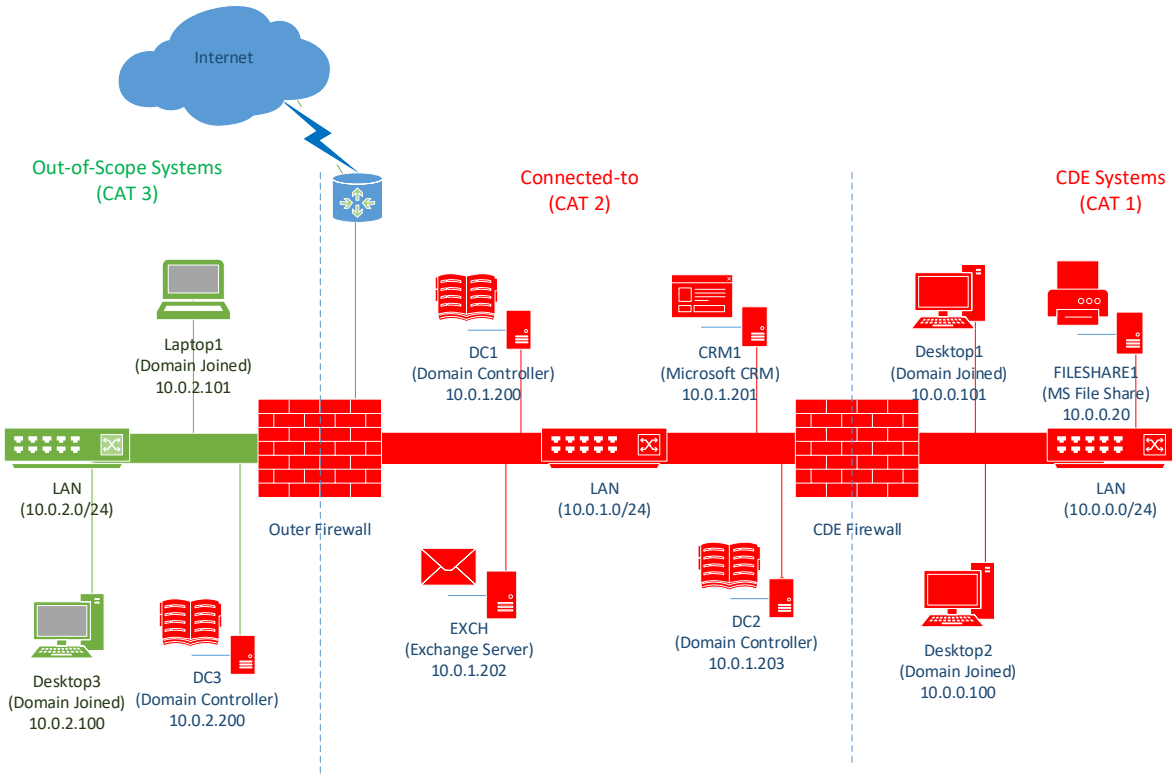

Example Design Scenario Diagram

Specific Configuration and Best Practice Techniques

This section provides some best practice and configuration specifics that can (**and most definitely should**) be used to harden the Microsoft AD environment. The narrative of the discussions within this section will be based on the example design scenario above. The following sub-sections will discuss some specific configurations, techniques and areas of AD that can help secure the environment.

Microsoft Active Directory Sites and Services

With this example design, traffic from the untrusted domain joined machines within CAT 3 and Microsoft AD can be limited to DC3. This can be done by configuring Microsoft AD Sites and Services correctly. Once setup correctly, the firewall rules can be locked down to block traffic between domain joined machines and CAT 2, only allowing DC3 into CAT 2.

N.B. *The design diagram does identify additional shared services, such as Microsoft CRM and Exchange Server. Access to these resources can be restricted on the firewall or front-end servers (Microsoft Exchange) could be used within CAT 3.*

Separate Accounts

Privileged access accounts should not be used for general productivity (day-to-day) tasks. Privileged accounts should only be accessed when required to perform a specific task which requires elevated privileged and should only be used to complete said task.

There are a lot of Windows vulnerabilities that require elevated privileges to work, therefore only standard user accounts should be used for general day-to-day activities. As per the BeyondTrust Annual Microsoft Vulnerabilities Report (2019) (Available Here; https://www.beyondtrust.com/blog/entry/the-annual-microsoft-vulnerabilities-report-just-released), 81% of crucial vulnerabilities discovered in 2018 would be mitigated by simply not running as an Administrator.

Additionally, as already discussed in the Attacks/Risks section above, limiting use of these accounts will reduce the attack surface and the chances a threat actor will gain access to a privileged account.

Securing Privileged Accounts

This section discusses the various techniques that can help with securing privileged accounts.

Privileged Account Hardening/Best Practice

Microsoft has released many extremely useful articles discussing how privileged user credentials can be secured to protect against many of the threats they face. This section will provide information on some (not all) of the recommendations as follows:

- Privileged accounts should not be used to browse the Internet or receive emails. In fact, Domain Controllers and other sensitive systems, including Privileged Access Workstations (PAWs); which are discussed later, should not have access to the Internet. Restrictions on web browsers and email clients can be placed on these systems and firewalls can be configured to block Internet access.
- Password re-use should not be in use within the environment. Password re-use is where the same password is used across many accounts. In fact, this is a breach of one of PCI DSS requirements 8.5.b which says “8.5.b Examine authentication policies and procedures to verify that use of group and shared IDs and/or passwords or other authentication methods are explicitly prohibited.”. LAPS can help with this which is discussed later, however it still comes with its complications so the easiest way to overcome this problem is to simply disable Administrator accounts across all servers and workstations (booting into Safemode will still allow disabled Administrator accounts to be used). LAPS coupled with disabling the Administrator account will provide coverage against password reuse and the ability to utilise the Administrator account remotely.

- Limit the number of accounts within privileged domain groups, or better yet, add accounts to the "Privileged Groups" only when needed and only for the duration of when that user requires said privileges. This should be supported by a strong approvals process for granting these privileges.
- Authentication Policy Silos can be used to limit where account can be used. This can effectively be used to limit privileged and "sensitive" accounts to where they can be used within the domain. This can be used to limit high-value account usage to high-value assets, thus limiting the risks of account compromises of these high-value accounts by enforcing their usage to high-value assets only (i.e., can limit accounts for us in CAT 1 only). See the "Authentication Policies and Authentication Policy Silos" article for further reading. https://docs.microsoft.com/en-us/windows-server/security/credentials-protection-and-management/authentication-policies-and-authentication-policy-silos
- Limit the number of accounts with direct administrative control over workstations and servers. Many organisations will limit access to privileged domain groups but will provide user accounts with local administrative rights to a lot of the estate's workstations and servers. These accounts would still be targeted by attackers and could be leveraged to compromise the AD estate.
- LAPS can be used to provide just-in-time local admin passwords. With LAPS, unique random passwords are generated for administrative user accounts across all workstations and servers the LAPS client has been deployed to. The passwords are then stored within Active Directory, which are protected by ACLs (Access Control Lists) which can be configured to a select number of system administrators. Ideally, the Administrative account should be disabled instead of relying on LAPS, due to additional complications as detailed below.
  N.B. This can cause an issue with PCI DSS requirements because potentially, multiple users will have access to the administrative passwords. A suitable compensating control would need to be implemented to work with LAPS due to this issue.
- MFA using Smart cards or certificate-based authentication mechanisms can and should be used to further protect privileged accounts.
- A Microsoft article discusses activities that increase the likelihood of compromise (https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/manage/component-updates/executive-summary). Within this article, it suggests that logging on to unsecured computers with privileged accounts can increase the risk of a compromise. The issue is that the security posture of the unsecured domain joined computers is unknown, therefore utilising an AD account that has elevated privileges across the domain is a bad idea as the credentials could be compromised. To perform administrative functions across these machines, the recommendation would be to create specific domain user accounts **(Not Domain Admin)** designed for the purpose of performing administrative functions on these unsecured computers. Add these users to an Active Directory Group (Unsecure Comp Local Admins) and add this group into the local Administrators groups on each machine within CAT 3. This will give these users the ability to perform local administrative functions. Domain user accounts should only be added to this group when they need to perform remote administrative activities so this group should remain empty for the majority of the time.

  Care still needs to be taken when logging on directly or utilising logon types which will leave reusable credentials on the destination host (e.g. RUNAS authentication). As already mentioned, it would be recommended that these accounts utilise MFA (all privileged accounts would be recommended). The "Administrative Tools and Logon Types" section of

the "Active Directory administrative tier model; https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged-access-reference-material. " Microsoft document details which logon types leave reusable credentials on a destination host. Again, it is important that the membership of the "Unsecure Comp Local Admins" group is kept empty until needed. To further secure these accounts, Active Directory Users and Computers can be used to restrict authentication of these user accounts to only work for the unsecured computers. This can be done by configuring the "Logon Workstations" for the user account profile (this can be found under Account → Log On To). You need to specify the NetBIOS machine name or DNS computer name of each machine the user account can logon to. This can get tiresome; therefore, this can be done via PowerShell by providing a list of users. Note: This property is limited to 64 computers.

Workstations and Member Servers Local Privileged Accounts

The following best practices are recommended for privileged accounts on domain joined workstations and servers;

- Limit the exposure of compromised Administrator accounts by limiting where they can logon;
  Via Group Policy, this can be done by adding **Administrator** to the following objects; Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments:
  - Deny access to this computer from the network
  - Deny log on as a batch job
  - Deny log on as a service
  - Deny log on through Remote Desktop Servers
- The following "Securing Privileged Access" (https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged-access) Microsoft article recommends the use of just-in-time passwords (i.e. LAPS) to perform administrative functions on workstations which are in less secure environments (i.e. CAT 3), so as not to expose a domain admin account. A contradictory recommendation within the follows "Avenues-to-Compromise" (https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/avenues-to-compromise) article also suggests disabling the built-in administrator account, which is disabled by default in current versions of Windows Workstation (from Windows Vista). Under normal operation of the Windows host, a disabled administrator accounts cannot be used. This, coupled with LAPS, will therefore further secure the environment since should a host get compromised and the attacker obtain the local administrator password hash, it will be a randomly generated password. A disabled administrator account can only be used if the machine is booted up into safe mode, so for Disaster Recovery purposes, an authorised user can access the LAPS passwords within Active Directory and recover the system through safe mode.

Privileged Accounts and Groups in Active Directory

The following best practices should be applied to privileged accounts and groups in Microsoft AD;

- For BA accounts, the following controls should be implemented;
  - Enable the "Account is sensitive and cannot be delegated" flag on the account. This can be done within the user's properties within Active Directory Users and Computers.

    - Enable the "Smart card is required for interactive logon" attribute on the account. Windows will reset the account's password to a 120-character random value. This makes this account unusable, although the password could be reset if required.
    - Within a Group Policy object that is linked to all workstations and servers, add this BA account into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
        - Deny access to this computer from the network
        - Deny log on as a batch job
        - Deny log on as a service
        - Deny log on through Remote Desktop Services
    - Within a Group Policy object that is linked to domain controllers, add this BA account into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
        - Deny access to this computer from the network
        - Deny log on as a batch job
        - Deny log on as a service
        - Deny log on through Remote Desktop Services

      (Note: This does not stop the Administrator account from being used at the console, or in disaster recovery scenarios)

    - Disable the BA account.
    - Enable auditing of the BA account to highlight any changes to the configuration of the account.
- For the EA group, the following controls should be implemented;
    - On a day-to-day basis, the EA group should not have any members, apart from possibly the Forest Root domain's administrator, which should be secured as above.
    - Within a Group Policy object that is linked to all workstations and servers, add this EA group into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
        - Deny access to this computer from the network
        - Deny log on as a batch job
        - Deny log on as a service
        - Deny log on through Remote Desktop Services
    - Enable auditing of the EA group to highlight any changes to the configuration of the account.
- For the DA group, the following controls should be implemented;
    - On a day-to-day basis, the DA group should not have any members and should only be used in build or disaster recovery scenarios. The only possible exception to this is possibly the domain's administrator account, which should be secured as above.
    - Within a Group Policy object that is linked to all workstations and servers, add this DA group into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
        - Deny access to this computer from the network
        - Deny log on as a batch job
        - Deny log on as a service
        - Deny log on through Remote Desktop Services

      N.B. If jump servers are used to administer DCs and Microsoft AD, ensure these are within an Organisational Unit (OU) that restricts this GPO from being applied.

- Enable auditing of the DA group to highlight any changes to the configuration of the account.

***N.B**. By default the Domain Admins group is added as a member of the local Administrators group on workstations and servers when they get added to the domain. This should not be modified as it can affect supportability and disaster recovery options.*

- For the local Administrators group, the following controls should be implemented;
  - On a day-to-day basis, the DA group should not have any members and should only be used in build or disaster recovery scenarios. The only possible exception to this is possibly the domain's administrator account, which should be secured as above.
  - Within a Group Policy object that is linked to all workstations and servers, add this local Administrators group into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
    - Deny access to this computer from the network
    - Deny log on as a batch job
    - Deny log on as a service
  - Within a Group Policy object that is linked to all domain controllers (usually the Domain Controllers OU), add this local Administrators group into the following rights in Computer Configuration\Policies\Windows Settings\Security Settings\Local Policies\User Rights Assignments;
    - Access this computer from the network
    - Allow logon locally
    - Allow logon through Remote Desktop Services
  - Enable auditing of the local Administrators group to highlight any changes to the configuration of the account.

***N.B**. Be careful do not implement Deny logon locally or Deny logon through Remote Desktop Services for the local Administrators group as this could cause issues.*

N.B. All additional PCI DSS Requirements for user accounts should also be implemented as per the PCI DSS V3.2.1; i.e. Requirements 7 & 8.

Multi-Factor Authentication should be used for all "Administrative" accounts.

<u>Role Based Access Controls (RBAC)</u>

As discussed in the previous section, in most cases, membership of the EA, DA and BA privileged groups should not have any members in them. So, how do system administrators carry out admin tasks? This is where RBAC comes in. This can be achieved via native tools and interfaces, tooling created by the company or third-party products. The idea of RBAC is to create roles within Active Directory and delegating the necessary rights and permissions to that group, allowing its members to perform day-to-day admin tasks. By utilising this approach, the organisation can create roles with specific rights and permissions, ensuring that the group doesn't have too much privilege.

This can be challenging and will rely on the organisation having clearly defined roles and responsibilities for the various groups of administrators within the IT teams. To ensure each role is configured with the least privileges required, for each role, the organisation should identify;

For each role you define, you should identify:

1. Which tasks members of the role perform on a day-to-day basis and which tasks are less frequently performed.
2. On which systems and in which applications members of a role should be granted rights and permissions.
3. Which users should be granted membership in a role.
4. How management of role memberships will be performed.

N.B. *If highly privileged rights/permissions are required, then separate accounts should be used, with MFA enabled on these accounts. All the guidance above for "Privileged" accounts should be followed for these highly privileged roles.*

Microsoft AD does have some predefined security groups that may be appropriate for specific job roles, for example; the DNS Admins group. The DNS Admin group has the necessary rights and permissions to maintain and manage DNS zones and records, however care of this group needs to be taken since DNS is heavily utilised for the correct operation of AD so a malicious threat actor could compromise the correct operation of AD by modifying DNS or can redirect traffic flows.

Membership within the privileged groups; EA, DA and BA should be temporary and immediate revoked once the user has carried out their task.

Service Account Hardening

Many system administrators will configure service accounts with excessive permissions. These accounts will often be configured as members of some of the highly privileged groups such as Domain Admins or Schema Admins. Each service account's function needs to be fully understood so the minimum level of privileges can be configured for the account to operate correctly. A compromise of a privileged account could allow an attacker to compromise the accounts of every user or service account utilising this information to compromise other systems within the environment.

Ensuring all service accounts are limited to privileges required for the service to perform its functions ensures that a compromise of the account doesn't give the attacker too much functionality within the Microsoft AD environment.

Another issue with service accounts is managing passwords. Typically, a service account will have a strong password configured and securely stored within the environment for when additional services are installed that requires that service account. This causes a security concern and causes issues with PCI DSS since multiple users know the service account's password. Something added in Windows Server 2012 was gMSA (Group Managed Service Accounts) which utilises the Key Distribution Service to generate new passwords in line with the ManagedPasswordInterval setting. This can be used for scheduled tasks, IIS Application Pools, and other applications that support gMSA.

Where service accounts are used, gMSA should be investigated to determine if this feature can be used. If gMSA cannot be used, strong passwords should be configured along with a documented process ensuring that each service account has its password changed at an acceptable pre-defined interval and the supporting applications are re-configured to reflect this. An individual shouldn't have access to the password, which should be under dual-control (i.e. two administrators who both

set half the password). There should be no reason to store this as processes should be in place to simply reset the password and application should this be required.

Additional hardening should be carried out; such as, but not limited to:

- Enable the "Account is sensitive and cannot be delegated" flag on the account. This can be done within the user's properties within Active Directory Users and Computers.
- If the service account is used on a handful of computers, restrict the account to these specific computers within Microsoft Active Directory for Users and Computers.
- Enable auditing of the service account to highlight any changes to the configuration of the account.
- Ensure a service account is created for each specific function. Don't share service accounts that have similar privilege requirements.

The "Windows server 2012: Group Managed Service Accounts" (https://techcommunity.microsoft.com/t5/Core-Infrastructure-and-Security/Windows-Server-2012-Group-Managed-Service-Accounts/ba-p/255910) article discusses how gMSA's can be configured for scheduled tasks.

This should provide some protection against Kerberoasting attacks.

Read-Only Domain Controllers

Read-Only Domain Controllers (RODC) are ideal for locations which are less secure or present a greater risk to the asset. This RODC type is designed so no changes can be made to AD and, that replication traffic is in-bound only.

By default, credential caching is not enabled on RODCs, therefore no passwords are stored. This behaviour can be changed if required by permitting specific groups to cache passwords. Due to the physical risks of RODCs, this would not be advisable.

Another feature of the RDOC is the ability to filter specific object attributes to stop replication of these attributes to the RDOC. This would be recommended if utilising LAPS (https://www.microsoft.com/en-us/download/details.aspx?id=46899) to ensure the administrator account passwords are not replicated to these RODCs.

There are other features of RODC; such as SYSVOL Modification Isolation, Administrator Role Separation and Read-only DNS Server which help to protect the environment. These are discussed further within the following paper (Beyond the MCSE: Active Directory for the Security Professional; https://www.blackhat.com/docs/us-16/materials/us-16-Metcalf-Beyond-The-MCSE-Active-Directory-For-The-Security-Professional-wp.pdf) that was presented at Black Hat and should be configured.

Finally, where RODCs are used in less secure remote locations, BitLocker Drive Encryption should be used to protect the operating system against physical attacks. That said, since this 'critical' piece of infrastructure is within these remote locations, PCI DSS Requirements 9.1 to 9.4 would also be included for these physical locations.

Just in Time (JIT) Administration

JIT (https://docs.microsoft.com/en-gb/windows-server/identity/securing-privileged-access/securing-privileged-access?redirectedfrom=MSDN), also referred to as just in time privileges utilises Microsoft solutions to provide administrative privileges when required for a specific task and for a specified

time period. Introduced in Windows Server 2016, this is performed by leveraging Microsoft Identity Manager and Privileged Access Manager (PAM) (https://docs.microsoft.com/en-gb/microsoft-identity-manager/pam/privileged-identity-management-for-active-directory-domain-services). PAM builds upon this JIT least privilege principle ensuring administrators don't have administrative privileges for a prolong period. This builds on a PowerShell toolkit defining a set of commands for performing privileged activities known as Just Enough Administration (JEA) (Discussed here from a TechEd 2014 Talk; https://channel9.msdn.com/Events/TechEd/NorthAmerica/2014/DCIM-B362)

PAM works by separating privileged accounts to a BASTION domain from an existing Active Directory environment. For an account to be provided with privileged access, this needs to be requested and then approved before the privileged account is given the necessary permission via a foreign principal group (shadow principals) within the BASTION forest which mirror objects within the existing Active Directory. Workflows are leveraged for users to request temporary access to privileged groups within the main active directory which is achieved via shadow principal users and groups within the BASTION forest, which have a time-to-live value so the shadow principle through these shadow principle users and groups.

PAM has four setup and operation steps as follows;

1. ***Prepare***: This preparation step is concerned with identifying the groups within the existing Active Directory that has significant privileges which need to then be created within the BASTION domain.
2. ***Protect***: Designed to protect the authentication of the users requesting JIT administration using; for example, Multi-Factor Authentication (MFA).
3. ***Operate***: This step is responsible for granting administrative privileges by adding the user account to the privileged group within the BASTION domain, for a set period of time. This will only be fulfilled if the authentication requirements are met, and the request gets approved.
4. ***Monitor***: This whole activity is audited with alerting and report functionality included by PAM.

Privileged Access Workstations (PAW)

Aka Secure Administrative Hosts, are hosts dedicated to providing a secure platform to allow privileged accounts to undertake administrative tasks within AD and the supporting environment. These hosts should not be permitted to access the Internet and should not run unnecessary applications such as, but not limited to, email applications, web browsers or productivity software (for example Microsoft Office).

These systems should never be accessed from a less-trusted system and the privileged accounts used when operating on these hosts shouldn't rely on a single authentication factor. Physical security of these devices should be considered and therefore PCI DSS Requirements 9.1 to 9.4 should be included wherever these hosts are located.

Smart cards should be used for privileged accounts used to access the PAW. This can be done by enforcing the use of smart cards on the PAW through Group Policy through the following GPO setting;

- Computer Configuration\Policies\Windows Settings\Local Policies\Security Options\Interactive logon: Require smart card

The accounts permitted to access the PAW should be limited through the following GPO settings;

- Computer Configuration\Policies\Windows Settings\Local Policies\Security Settings\Local Policies\User Rights Assignment

A secure baseline of the PAW should be created utilising tools such as; but not limited to, the Microsoft Security Configuration Wizard, Microsoft Security Compliance Manager and other third-party tools; for example, CIS (Centre for Internet Security) baselines.

AppLocker or a third-party application restriction software should be deployed to restrict which scripts, tools and applications are permitted to run.

RDP restrictions should be configured to limited who can connect to the Remote Desktop Service.

PAWs shouldn't have Internet connectivity and should be properly patched inline with PCI DSS Requirement 6.2.

Finally, care needs to be taken if the PAWs are running on a virtualisation infrastructure since additional security controls and hardening of the virtualisation infrastructure would need to be carried out.

SMB Message Signing

SMB Message Signing helps to prevent attacks relaying NTLM authentication messages, therefore stopping SMB Relaying discussed above. This can be done through GPO by enabling the;

- SMB Servers
    - Microsoft network server: Digitally sign communications (always)
- SMB Clients
    - Microsoft network client: Digitally sign communications (always)

N.B. As discussed in the following blog article "Security Advisory: Critical Vulnerabilities in NTLM Allow Remote Code Execution and Cloud Resources Compromise" (https://blog.preempt.com/security-advisory-critical-vulnerabilities-in-ntlm), a recent vulnerability allows SMB Signing to be bypassed, therefore the June 2019 Microsoft patch for CVE-2019-1040 and CVE-2019-1019 needs to be applied.

Domain Controllers and Member Servers

Regardless of their location within the environment, Security Features on Domain Controllers and Member Servers should not be disabled. The Windows Firewall automatically configures itself based on the roles or features installed on it, additionally User Access Control (UAC) should not be disabled on any server unless a specific scenario requires it as per Microsoft Support article 2526083 (https://support.microsoft.com/en-gb/help/2526083/disabling-user-account-control-uac-on-windows-server).

As already discussed, access to the Internet from Domain Controllers and Members Servers should be blocked and a robust patching policy / process needs to be implemented. Also, as discussed under PAWS, a secure baseline should be created utilising tools such as; but not limited to, the Microsoft Security Configuration Wizard, Microsoft Security Compliance Manager and other third-party tools.

As highlighted in one of the linked Microsoft articles, care needs to be taken when security baselines are transposed from older operating systems to newer operating systems as some of the newer

security settings may end up being misconfigured. When installing newer servers or raising the functional levels of the domain/forest, the newer security features should always be reviewed, creating a new baseline for these systems.

RDP restrictions should be in place on all Domain Controllers and Member Servers. This should be configured using a combination of user rights settings allowing specific users RDP access and through the correct implementation of WFAS (Windows Firewall with Advanced Security) configuration to limit access to authorised clients such as PAWs. This should be configured through GPO to ensure a consistent configuration and to ensure the settings are re-applied, should they get changed.

Remote Desktop Admin Mode

This should be configured across the entire estate. With this setting enabled, the transmission of reusable credentials to the remote systems are prevented. This will prevent administrative credentials from being harvested should the remote system be compromised. How to configure this on the remote systems and how to launch remote desktop in RestrictedAdmin mode is explained within this "Remote Desktop Services: Enable Restricted Admin mode" article. (https://social.technet.microsoft.com/wiki/contents/articles/32905.remote-desktop-services-enable-restricted-admin-mode.aspx)

Alerting / Logging

As with any environment, adequate logging and alerting should be configured to help identify compromises, hacking attempts and to assist in any forensics work to help identified specific details of the attack/compromise. The Microsoft "Audit Policy Recommendations" https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/audit-policy-recommendations article provides details of what logging should be enabled.

Additional General Hardening

This section is in no way a comprehensive list of all best practice security hardening required to design a secure Microsoft Active Directory environment. Industry-recognised hardening resources and tools such as the Microsoft Security Configuration Wizard, Microsoft Security Compliance Manager and other third-party tools should be utilised to ensure a robust hardening configuration strategy is used. Some hardening examples include;

- Enforce NTLMv2 within the environment (or investigate replacing it with a more secure authentication protocol such as Kerberos and restricting the use of NTLM within the environment. The following "Restricting NTLM usage" https://docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2008-R2-and-2008/jj865676(v=ws.10) article provides additional information)
- Enforce LDAP/S Signing (Microsoft plan to enforce a security update early 2020 to enabled LDAP channel binding and LDAP signing hardening changes; see https://support.microsoft.com/en-us/help/4520412/2020-ldap-channel-binding-and-ldap-signing-requirement-for-windows)
- Enforce EPA (Enhanced Protection for Authentication)
- Group Policy Preferences is a feature introduced in Windows 2008. In 2014, Microsoft released the encryption key used by this feature to encrypt credentials that can be stored within Group Policy for scheduling tasks, mapping drives, etc… Although in 2014 Microsoft released a patch for this, the old credentials can still be left lingering within the SYSVOL share or cached on local machines. It is important to ensure these are not located within the

environment. GrimHacker released a tool called “GP3Finder” (https://grimhacker.com/2015/04/10/gp3finder-group-policy-preference-password-finder/) that can be used to help identify this issue

- On Member Servers, set the “Computer Configuration\Windows Settings\Security Settings\Account Policies\Local Policies\Interactive logon: Number of previous logons to cache (in case domain controller is not available)” to 0. This will ensure users credentials are no longer cached on Member Servers which is set to 10 by default.
- On all systems, disabled the “Store password using reversible encryption” option in Group Policy. This is within the following setting; “Computer Configuration\Windows Settings\Security Settings\Account Policies\Password Policy: Store passwords using reversible encryption”
- Prevent WDigest passwords being stored in memory. A patch was released to resolve this issue back in 2014, however a registry key setting is also required; “HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\SecurityProviders\WDigest: UseLogonCredential” set to 0
  See the following Microsoft Security Advisory “Update to improve credentials protection and management” (https://support.microsoft.com/en-gb/help/2871997/microsoft-security-advisory-update-to-improve-credentials-protection-a)
- Disable NetBIOS and LLMNR, which can be done by as follows:
  - For NetBIOS, use the “001 Microsoft Disable Netbios Option” and then configure the clients to accept the DHCP server to determine the NetBIOS behaviour. This can be done through the following registry key, which can be configured through Group Policy;
    HKLM:SYSTEM\SYSTEM\CurrentControlSet\Services\NetBT\Parameters\Interfaces\tcpip{*}:NetbiosOptions -Value 2
    N.B. The * denotes a unique hexadecimal number
  - For LLMNR, this can be done via GPO by setting enabling “Turn off Multicast Name Resolution” option within Computer Configuration\Administrative Templates\Network\DNS Client within Group Policy

Regular Review of Controls

Since there is a lot of the security configuration, techniques and controls used to help secure Microsoft AD for us as a “Shared Service”, this configuration must be regularly reviewed against the baseline. Privileged Groups/Accounts should be monitored to identify any unauthorised changes which could indicate a potential compromise of the environment, or a simply mis-configuration by an administrator.

My opinion would be that this review needs to be carried out at least quarterly to help ensure no mis-configurations compromising the security of the design and implementation is introduced and that all security controls, configurations and techniques are maintained. Staff should also be allowed time to keep abreast of new security techniques designed to further secure Microsoft AD and Windows hosts to maintain a high security baseline across this type of deployment.

Penetration Testing

There is no denying that the use of Microsoft Active Directory as a “Shared Service” is highly risky and requires a lot of technical configuration and discipline to best practices to help secure the environment. A breakdown of just one of these technical controls or a discipline to best practices implemented can result in a compromise of the secure environment and breach of cardholder data.

To ensure the security best practices and technical configuration is maintained, it is important that the environment is adequately scoped for Penetration Testing activities. Testing should be increased within these environments to validate the security and segmentation across the environment with penetration testing recommended to be a minimum of every 3 to 6 months depending upon the volumes of cardholder data that could be at risk through a compromise.

MFA

With the introduction of PCI DSS Version 3.2.1, all non-console administrative access needs to support Multi-Factor Authentication (MFA). The guidance within the PCI DSS Requirements and Security Assessment Procedures document does highlight that where the CDE is segmented, MFA can be configured at the network level for entry into the CDE, it doesn't then need to be configured on each system. However, with this type of setup, I would suggest that for all non-console access where administrative users are being used, MFA be enabled to protect against replay attacks. Additionally, it is also worth pointing out that for some authentication; SMB access, it is unlikely that MFA can be configured, therefore it is best to ensure administrative accounts are not used on any authentication mechanisms that does not support MFA.

Domain Controller PCI DSS Requirements

When the DC is located within the CAT 3 zone, it will still be in scope for all applicable requirements. This section gives an idea of what PCI DSS Version 3.2.1 requirements will **NOT** likely be applicable; however, this list can still depend upon the environment setup and the QSAs discretion.

N.B. This section only covers the applicability of PCI DSS Requirement to the Domain Controller in the scenario highlighted at the beginning of the post, i.e. operating within the CAT 3 zone. It doesn't include the organisations overall PCI DSS requirement. It doesn't include the Jump Box or PAWS that may be located within the CAT 3 zone used to access the DC within its site location.

Requirement 1: Install and maintain a firewall configuration to protect cardholder data

This requirement is applicable since the WFAS should be configured to reduce the attack surface of the DCs. If other "in-scope" devices; for example, a Jump Box/PAWS, are located at remote sites where the RODCs are located, a hardware firewall may be used to protect all these in scope devices. This will very much depend upon the environment since the use of the WFAS can simplify the firewall configuration since Windows can pre-configure this based on the roles/components installed.

The following requirements have been identified as potentially not being applicable;

Requirement	Discussion
1.1.4	DMZ wouldn't be applicable for the CAT 3 Domain Controller.
1.2.2	Router isn't applicable for the CAT 3 Domain Controller.
1.3	Public access to the Domain Controller wouldn't be applicable for the CAT 3 Domain Controller.
1.4	Portable computing devices aren't applicable for the CAT 3 Domain Controller.

Requirement 2: Do not use vendor-supplied defaults for system passwords and other security parameters

Hardening of the DCs will be in scope to ensure they are not susceptible to compromise. Only the following requirements have been identified as potentially not being applicable;

Requirement	Discussion
2.1.1	Wireless wouldn't be applicable for the CAT 3 Domain Controller.
2.6	Shared Hosting shouldn't be applicable for the CAT 3 Domain Controller.

As per requirement 2.2, configuration standards need to be in place for **ALL** system components (which include DCs). Suitable hardening needs to be carried out, this should be aimed at protecting the systems from attack. Industry-accepted system hardening standards exist and must be used to help meet this requirement. My recommendation would be to incorporate as many security hardening standards and sources as possible to ensure DCs (and other system components) are securely configured.

Requirement 3: Protect stored cardholder data

This requirement wouldn't be applicable since the DCs in CAT 3 wouldn't be storing CHD.

Requirement 4: Encrypt transmission of cardholder data across open, public networks

This requirement wouldn't be applicable since the transmission of CHD wouldn't touch the DCs.

Requirement 5: Maintain a Vulnerability Management Program

The only requirement that wouldn't be applicable within Requirement 5 would be Requirement 5.1.2, since the DCs would be considered as commonly affected by malicious software and therefore would need anti-virus to be installed.

Requirement 6: Develop and maintain secure systems and applications

Within Requirement 6, only Requirements 6.1, 6.2 and 6.4 would likely be applicable for the in-scope DCs.

Requirement 7: Restrict access to cardholder data by business need to know

All Requirement 7 is likely in scope for the CAT 3 DCs.

Requirement 8: Identify and authenticate access to system components

All Requirement 8 is likely in scope for the CAT 3 DCs. The majority of this will be covered by settings which are configured within the Active Directory domain which would be replicated to the DCs; therefore, already be covered by the entire PCI DSS program.

Requirement 9: Restrict physical access to cardholder data

Only the requirements from Requirement 9.1 to Requirement 9.4 (inclusive) would likely be included in scope for PCI DSS assessment activities. These cover the physical location of where the DCs are physically located.

Requirement 10: Track and monitor all access to network resources and cardholder data

All Requirement 10 is likely in scope for the CAT 3 DCs.

Requirement 11: Regularly test security systems and processes

The only exception to all Requirement 11 being in scope is Requirement 11.2.2 for External ASV scanning.

Requirement 11.1 would be applicable for the CAT 3 DCs and the location where the CAT 3 Domain Controllers are located.

Requirement 12: Maintain a policy that addresses information security for all personnel

I would expect all of Requirement 12 to be in scope for PCI DSS. 12.3 should include acceptable use of some of the privileged accounts, limited their use on certain devices.

My Final Ramblings

Reviewing the PCI SSC released the PCI SSC Scoping guidance, the use of Microsoft Active Directory as a 'Shared Service' is included. For this reason, many organisations have started to adopt this same strategy within their architectural designs for the CDE and non-CDE environments. In my honest opinion, this is still dangerous since one mis-configuration of the myriad of security controls designed to secure this setup can potentially lead to a compromise of the environment. It is generally much safer to separate the non-CDE and CDE Microsoft Active Directories.

Although the above techniques should help organisations configure Microsoft Active Directory securely to support its use within a PCI DSS Cardholder Data Environment (CDE). Organisations still need to be careful of scope creep within the CAT 3 environment. The premise of determining if a system component is in scope or not, depends on the risk that system has to the environment should it be compromised.

Your specific environment should be reviewed by a QSA to determine what the scope is and only a QSA would be able to sign this off. Although the customer is responsible for defining the scope of the environment, the QSA will validate it during the assessment. This is a requirement within the front section of the Report on Compliance (ROC) so a QSA would still have to sign it off.

This is just one type of Microsoft Active Directory design and doesn't include the use of multiple domains within the forest or multiple inter-connected forests. There are specific attacks against domain trusts which would need to also be considered should a Microsoft Active Directory forest with multiple Active Directory domains be in use. There are some articles discussing this at the end of this blog.

If the entity being assessed is deploying a shared Active Directory, the specific hardening and secure design concepts being followed should be adequately documented so the QSA is able to validate the secure design and review specific configuration settings before signing it off as adequately segmented and secured.

I would also like to make a very specific point on segmentation testing and penetration testing. If a shared Active Directory is deployed, the scope of the internal penetration testing and segmentation testing **MUST** include an exhaustive test of all CAT 3 environments, **INCLUDING** system components within them. This **MUST** include domain joined hosts, even if they are not determined in scope for the PCI DSS Assessment. The aim of the penetration testing and/or segmentation testing should be

to compromise CAT 3 hosts to see if the penetration tester can leverage the host to compromise in-scope CAT 3 systems or the CAT 2/1 environments. This should be fully documented within the organisations penetration testing methodology which is a requirement under PCI DSS Requirement 11.3.

Planning for Compromise

The various Microsoft articles linked within this blog post include a section called "Planning for compromise" (https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/planning-for-compromise). This is an interesting read and organisations should plan for when a Microsoft Active Directory gets compromised because it may require a full rebuild of the environment and therefore should be included within the organisation's disaster recovery plan. A compromise of a system makes that system untrustworthy, which means it will need to be rebuilt before it is considered trustworthy within the environment. If the compromise impacts the Domain Controllers within the environment, due to the replication between all Domain Controllers, all Domain Controllers within the environment become untrustworthy, therefore requiring a complete rebuild of the environment. Therefore, it is important that this needs to be planned for.

Some of the articles used during research

I don't usually include a section like this within my blog posts, however I do feel that some of the resources I used in the writing of this post may be useful to others.

1. A Guide to Attacking Domain Trusts: http://www.harmj0y.net/blog/redteaming/a-guide-to-attacking-domain-trusts/
2. Active Directory administrative model: https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged-access-reference-material
3. Active Directory forest trusts part 1 - How does SID filtering work?: https://dirkjanm.io/active-directory-forest-trusts-part-one-how-does-sid-filtering-work/
4. Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide: https://technet.microsoft.com/library/dd378897.aspx
5. Best Practices for Security Active Directory: https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/best-practices-for-securing-active-directory
6. Beyond Domain Admins – Domain Controller & AD Administration: https://adsecurity.org/?p=3700
7. Beyond the MCSE: Active Directory for the Security Professional: https://www.blackhat.com/docs/us-16/materials/us-16-Metcalf-Beyond-The-MCSE-Active-Directory-For-The-Security-Professional-wp.pdf
8. Designing the Site Topology: https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/designing-the-site-topology
9. How Active Directory Authentication Works: http://www.rebeladmin.com/2018/06/active-directory-authentication-works/
10. How to configure a firewall that resides between a Windows Domain Controller and a NetApp controller: https://kb.netapp.com/app/answers/answer_view/a_id/1030571/~/how-to-configure-a-firewall-that-resides-between-a-windows-domain-controller
11. Microsoft Attack Surface Analyzer: https://github.com/Microsoft/AttackSurfaceAnalyzer/wiki
12. Protecting Privileged Domain Accounts: Safeguarding Access Tokens: https://digital-forensics.sans.org/blog/2012/03/21/protecting-privileged-domain-accounts-access-tokens?fbclid=IwAR3IBEmPT_WRUuPLScrjudvrceS7jw1YF_rZITfVZGsrYfevH7L1s4pglXk

13. Remote Desktop Services: Enable Restricted Admin mode: https://social.technet.microsoft.com/wiki/contents/articles/32905.remote-desktop-services-enable-restricted-admin-mode.aspx
14. Security Advisory: Critical Vulnerabilities in NTLM Allow Remote Code Execution and Cloud Resources Compromise: https://blog.preempt.com/security-advisory-critical-vulnerabilities-in-ntlm
15. Service overview and network port requirements for Windows: https://support.microsoft.com/en-gb/help/832017/service-overview-and-network-port-requirements-for-windows
16. Windows 10 Device Guard and Credential Guard Demystified: https://blogs.technet.microsoft.com/ash/2016/03/02/windows-10-device-guard-and-credential-guard-demystified/
17. Windows Server 2012: Group Managed Service Accounts: https://techcommunity.microsoft.com/t5/Core-Infrastructure-and-Security/Windows-Server-2012-Group-Managed-Service-Accounts/ba-p/255910