Australian Government
Australian Signals Directorate

Secure Administration

First published: September 2015
Last updated: October 2021

Introduction

Privileged access allows administrators to perform their duties such as establishing and making changes to key servers, networking devices, user workstations and user accounts. Privileged access or credentials are often seen as the 'keys to the kingdom' as they allow the bearers to have access and control over many different assets within a network.

Privileged access is often a key goal of malicious actors. They can use privileged access to:

- propagate malware to multiple workstations and servers
- add new user accounts, including privileged accounts
- bypass controls for applications, databases and file servers
- implement configuration changes to make future access easier.

Given the scale and complexity of enterprise networks, it is reasonable to assume that at least one standard user account and workstation within an organisation's internet-connected network could be compromised by malicious actors. As administrator accounts often have unrestricted access to critical resources, this publication focuses on protecting sensitive accounts and resources from malicious actors who have gained a presence on a network.

This publication is designed to complement and expand upon the guidance contained within the *Information Security Manual* (ISM).

Secure administration and the cloud

The primary intent of this publication is to secure the administration of traditional corporate network assets such as domain controllers and application servers as well as the infrastructure used for the administration of these assets.

Administration of cloud-based infrastructure, systems and applications brings different challenges and may require a different approach. As such, not all controls within this publication may be directly applicable to the administration of cloud assets and may require assessment and adjustment before being applied to infrastructure used for cloud administration.

Throughout the document, the controls will contain guidance on applying the recommendation within a cloud environment.

Rationale for implementing secure administration

The goals of malicious actors, whatever they may be, are far easier to achieve when privileged access on a network or system has been attained. Assuming malicious actors have not been able to achieve their goals with the access granted by their initial foothold, gaining privileged access to a network or system may be an extremely high priority. This is for a number of reasons:

- **Wider malware propagation.** Privileged accounts, especially built-in *Administrator* accounts and groups such as Domain Administrators, have privileged access to a wide range of systems. If malicious actors compromise such an account, they can more easily propagate malware to multiple workstations and servers. This increases the chance of accessing the information they want while also making it harder to remediate the cyber intrusion.
- **More intrusive access.** Privileged access will also provide malicious actors greater access to systems. For example, privileged access allows malicious actors to utilise more advanced tools and techniques, such as credential theft, in order to move laterally within a network.
- **Access to sensitive data stores.** Privileged access may allow malicious actors direct access to a sensitive data store, such as a database, without having to access the data store through regular means where illegitimate access may be audited and detected.
- **Insight into detection and remediation activities.** If malicious actors have sufficient access, they may gain insight into any detection and remediation activities being performed. This may enable malicious actors to avoid complete detection, or hamper the effectiveness of mitigation activities. This may necessitate the creation of secure administration assets prior to any significant remediation work proceeding.

Organisations which implement secure administration strategies can increase their resilience to a successful cyber intrusion by preventing or greatly delaying malicious actors from gaining access to privileged accounts and credentials. Furthermore, any cyber security incident response activities and remediation work will be more effective and agile.

Lessons from a cyber security incident response team

The Australian Signals Directorate has responded to compromises of government networks. In the majority of compromises, malicious actors had managed to acquire significant privileged access, such as Domain Administrator access, which allowed them to more efficiently steal, modify and damage data; move laterally within the network to gain access to additional resources and remaining undiscovered on the network for longer.

A well-known cyber security incident where privileged accounts were compromised and used to further the goals of malicious actors is the 2014 compromise of Sony Pictures Entertainment. A large amount of data, including intellectual property that directly affected the company's operations, was exfiltrated. In addition, sensitive emails were publicly disclosed that severely damaged the company's reputation. Finally, malware was installed on workstations that wiped hard drives which resulted in the loss of a large amount of corporate information.

While also being a critical control, the implementation, or lack of, a secure administration environment can make a significant difference to the effectiveness of investigation and remediation activities carried out in response to a cyber security incident. This is particularly the case when malicious actors already have privileged access, or it is suspected that it is highly possible due to a poor network security posture.

When confronted with a significant network compromise, one of the common remediation activities a cyber security incident response team will recommend is to establish a secure administration environment. Because of this, resources which would otherwise be focused on remediation activities may be spent on establishing a secure administration environment. This can lead to a delay in the full remediation of a network, increasing the consequences associated with a compromise, as malicious actors have more time to propagate within the network resulting in the theft of sensitive data.

Elements of secure administration

The table below provides an outline of the key controls which underpin secure administration within an organisation. Each control is expanded upon in greater detail in later sections of this publication.

While implementing each control in isolation would provide some benefit (some more than others), these controls are most effective when implemented as a package. If the entire package of controls cannot be implemented in one step, the preferred order of implementation is to start at the top of the table and work down.

Most of the controls below can be achieved using a combination of standard tools in modern operating systems and logical changes to an organisation and their corporate network.

Element	Description
Privileged access control	Controlling access to privileged accounts is a fundamental control that will protect privileged accounts from misuse. The access control methodology will encompass the concepts of 'least privilege' and 'need to have' as well as processes and procedures for managing service accounts and staff movements.
Multi-factor authentication	Implementing additional factors of authentication beyond usernames and passphrases, such as physical tokens or smartcards, can help protect critical assets. If malicious actors were to compromise credentials for privileged accounts, as all administrative actions would first need to go through some form of multi-factor authentication, the consequences can be greatly reduced.
Privileged workstations	The use of a known secure environment for administrative tasks can result in a lesser risk of the network being compromised due to the implementation of additional controls.
Logging and auditing	Automated generation, collection and analysis of security and administrative related events from workstations, servers, network devices and jump boxes will enable detection of compromises and attempted compromises. This will enable organisations to respond more quickly, reducing the implications of a compromise.
Network segmentation and segregation	Segmenting a network into logical zones such as differing security domains, and further segregating these logical networks by restricting the types of data that flow from one zone to another, restricts lateral movement. This will prevent malicious actors from gaining access to additional resources.
Jump boxes	A jump box is a hardened remote access server, commonly utilising Microsoft's Remote Desktop Services or Secure Shell (SSH) software. It acts as a stepping point for administrators accessing critical systems with all administrative actions performed via a jump box.

Privileged access control

Privileged accounts are often targeted by malicious actors due to their access across the breadth of an organisation's network and systems. Restricting the use of built-in administrator groups and accounts and delegating privileged permissions in accordance with the principles of least privilege is an effective way to reduce the impact and spread of malicious actors' access during a cyber intrusion.

To assist in restricting the use of privileged accounts, the following procedural and technical controls should be implemented:

- Ensure that unique identifiable accounts are linked to individual users and they are authenticated every time privileged access is granted on a system. This will ensure accountability and attribution of all actions.
- Restrict access for privileged accounts by issuing administrators a standard user account in addition to separate privileged and unprivileged administrator accounts for administrative purposes. Separate user and administrator accounts will provide a logical separation of administrative and user tasks while the use of privileged and unprivileged administrator accounts will provide another layer of abstraction to further protect privileged account credentials. The unprivileged administrator account is used to control remote access to a jump box and corporate workstations and servers while all actual administration tasks are performed with the privileged administrator account using 'runas', 'sudo' and remote administration management tools.
- Privileged administrator accounts should not be used to run ongoing tasks on a system, including with 'runas' or by escalating privileges via Windows User Account Control. It is possible to extract passphrase hashes of processes running with UAC and 'runas'; however, the hashes are cleared following process termination. In cases where services or applications require additional privileges indefinitely or over a long period of time, specific service accounts with the minimum required permissions should be used instead.

- Enforcing role-based delegation of privileges for privileged administrator accounts will reduce the exposure of an account in the event of compromise by malicious actors. As functional levels have progressed for Microsoft Active Directory, built-in administrator groups have been provided with the goal of encouraging administrative role-based delegation. While these groups provide a useful starting point, it is likely further customisation will be required based on any organisational structures in place.
- Enforcing strong passphrase management for privileged and unprivileged administrator accounts will reduce the risk of passphrase guessing and brute-force attacks against administrator accounts.
- Disabling local administrator accounts will reduce the risks associated with credential theft.
- Enforcing the use of Microsoft Window's Secure Desktop for entry of all privileged account credentials where possible. This will prevent key loggers from capturing credentials when administrators use their unprivileged administrator account to elevate to their privileged administrator account to perform an administrative action. This setting can be set in Group Policy (User Account Control: Switch to the secure desktop when prompting for elevation).
- On recent Microsoft Windows releases, consider utilising the Protected Users group for privileged accounts if supported by the environment. This will enforce strong authentication protocols and Kerberos settings.
- Do not allow service accounts to be members of any built-in administrator groups. This will reduce the risks associated with credential theft. Where possible, managed service accounts should be used with delegation of specific privileges as required.

Pass-the-Hash attacks

Pass-the-Hash (PtH) is a very common and effective method malicious actors use to move laterally within an organisation's network with their ultimate goal being to gain access to sensitive data.

One common technique to perform a PtH attack is for malicious actors to expose the authentication credentials of user accounts that have been previously authenticated on a compromised workstation. These credentials are then passed directly around the network. The credentials are usually hashed passwords, although there may be cases where plaintext passwords or single sign on and multi-factor tokens are stolen.

Environments where local administrator accounts share the same password throughout the environment and where domain and enterprise administrator accounts are used for regular workstation support are particularly susceptible to this attack.

By following the previous controls for privileged access control, leveraging PtH attacks become far more difficult. In particular, the separation of privileged and unprivileged administrator accounts and the use of remote administration management tools from a jump box reduce the number of privileged administrator account hashes being stored on workstations and servers.

Role-based delegation of administrative privileges further strengthens the security posture of an organisation by reducing the implications of PtH attacks as much as possible by minimising the exposure of networks even in the event of a privileged account hash being compromised.

Other controls throughout this publication further mitigate PtH attacks and should be employed where possible.

Considerations when administering a cloud environment

There are no specific issues to be considered when applying the above guidance to a cloud environment.

Multi-factor authentication

By implementing an additional authentication factor, such as a hardware authentication token, the consequences associated with the theft of passphrases can be greatly reduced. When implementing multi-factor authentication (MFA) in a secure administration context, the primary consideration is where MFA takes place. The three possible locations are the privileged workstation, the jump box or the target asset. The following table outlines the advantages and disadvantages associated with the three different locations.

Privileged Workstation	Jump Box	Target Asset
Advantages ▪ Ease of integration with MFA product. ▪ MFA can also cover workstation login.	**Advantages** ▪ Ease of integration with MFA product. ▪ MFA only needs to be deployed to a centralised location.	**Advantages** ▪ MFA process occurs on the target asset at time of privileged action.
Disadvantages ▪ Multiple deployment locations ▪ MFA process not performed on the target asset. ▪ MFA process may not be performed temporally close to the privileged action.	**Disadvantages** ▪ MFA process not performed on the target asset. ▪ MFA process may not be performed temporally close to the privileged action.	**Disadvantages** ▪ Varied target assets on different platforms must support integration with MFA. ▪ Multiple deployment locations.
RECOMMENDED	**RECOMMENDED**	**RECOMMENDED**

Considerations when administering a cloud environment

The key decision when implementing MFA for the administration of a cloud-based service is determining where to implement the MFA: within an organisation's local network, within the cloud service or both.

If implementing MFA within the cloud service, the following issues should be considered:

- Many cloud service providers offer a MFA service which can be used to enhance the authentication process when accessing cloud assets, such as accessing an administrator web console or protecting an application programming interface (API).
- Organisations may establish their own dedicated MFA server within a virtual private cloud to enhance the authentication process to the various cloud servers hosted within the virtual private cloud.
- An existing local MFA server may also provide coverage of cloud assets.
- An MFA process could be enforced for local jump boxes which are then used to access cloud services. This provides little benefit if these cloud services can be administered by hosts other than the jump boxes.
- If a virtual private cloud is accessed via a virtual private network (VPN), the existing local MFA server could extend its coverage to include the cloud-based assets.

Privileged workstations

A key goal of malicious actors is to compromise the workstation of an administrator or a user with privileged access to one or more systems. This compromise can occur due to common intrusion methods such as phishing emails, drive-by downloads or through lateral movement when malicious actors are already within the network.

Once malicious actors have access to a privileged workstation, malicious actors can perform some form of credential theft (e.g. keylogging, hash dumping) to gain access to additional privileged account credentials. Therefore, it is essential to reduce the likelihood that a workstation used for administration purposes will be compromised. This can be achieved through the use of dedicated and hardened privileged workstations.

Dedicated privileged workstations

Privileged users utilising a low-privilege user account for high risk activities such as email or web browsing reduces the risk of privileged account credentials being compromised. However, if a user was to elevate their privileges (e.g. to a local administrator or SYSTEM account) after being compromised by malicious actors, their privileged administrator account credentials could be at risk.

Providing a physically separate workstation that is used exclusively to carry out privileged activities (a privileged workstation) will greatly reduce the risk of privileged administrator accounts being compromised.

Privileged workstations should not have internet or email access. Additionally, the workstations should be logically separated from other zones within the network, and only allowed to communicate with assets which require management (ideally through jump boxes), or infrastructure required for the proper management of workstations and servers such as internally managed patching servers. Communication between this dedicated privileged workstation and with less trusted zones, such as regular user workstations, should be prevented.

The diagram below shows the different use cases for a dedicated privileged workstation and a regular user workstation.

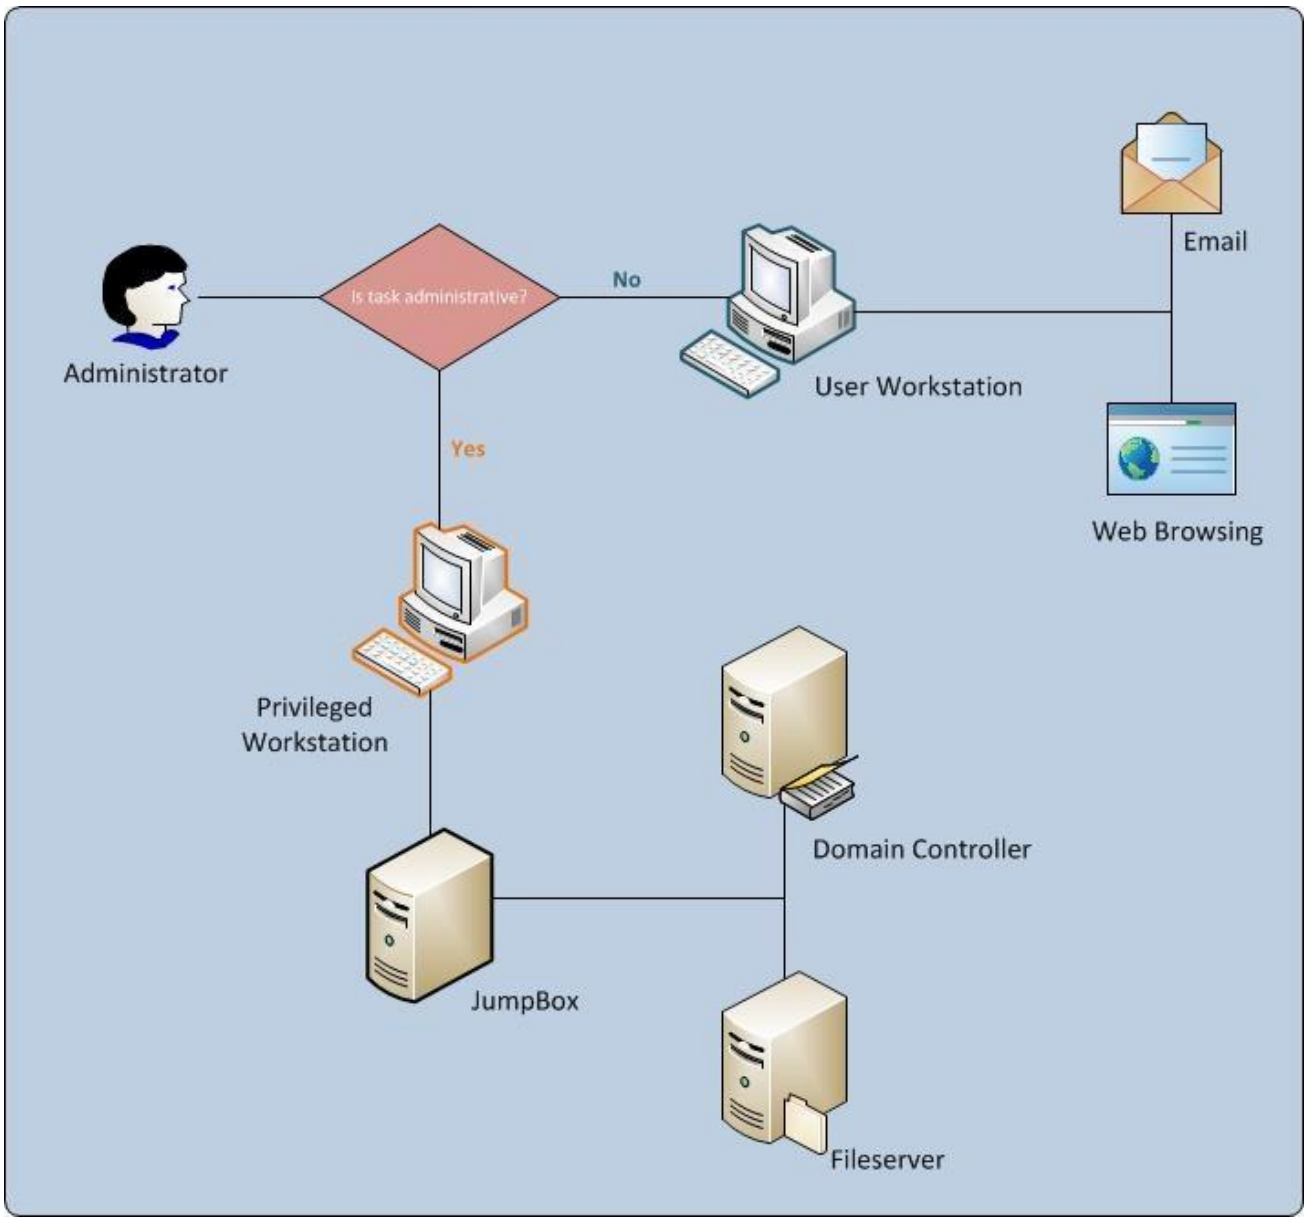

The use of virtualisation to achieve dedicated workstations

The use of virtualisation does not yield the same security benefit as using a dedicated physically separate workstation. However, a risk-based approach may determine that utilising a virtualisation solution is sufficient.

Care should be taken to ensure separation of physical and virtual machines to prevent leakage across the virtualisation boundary. For example, drag and drop integration, clipboard integration and shared folders should all be disabled where possible when using a virtualised solution.

The configuration of virtualised desktops can take many forms. The following table is ranked from most secure to least secure and encompasses the most likely approaches.

Configuration	Advantages	Disadvantages	Notes
Remotely virtualised privileged and user workstations (preferred)	Separation of the workstations managed by the hypervisor and possibly physical hardware. Improved control over workstation. No physical access to workstations so controls can't be circumvented.	Relatively difficult to use. Possibly more expensive as a thin client will also be required. Availability in the event of service outage.	A thin client will be required to allow access to the virtualised desktops. The privileged workstation should be non-persistent.
Physical privileged workstation and virtualised user workstation	Relatively hard to escape from the virtualised user workstation back to the privileged workstation.	Privileged workstation must be persistent.	The user workstation should be remotely virtualised in order to prevent console access to the virtual machine.
Physical user workstation and virtualised privileged workstation (not preferred)	Logical separation of the workstations. Privileged workstation can be non-persistent.	The workstation can be used to pivot into the privileged workstation leveraging Pass-the-Hash, key logging or saved passphrase functionality.	The privileged workstation should be remotely virtualised in order to prevent console access to the virtual machine.

Hardening privileged workstations

Regardless of whether dedicated privileged workstations are employed, workstations that are used by privileged users should be subject to additional controls. At a minimum, the following procedural and technical controls should be implemented for privileged workstations:

- Ensure administrators log in with an unprivileged administrator account.
- Implement host-based firewall rules restricting both inbound and outbound traffic to only traffic required to perform administration functions. For example, restrict outbound remote access services (e.g. Remote Desktop Services (RDP), PowerShell Remoting and Secure Shell (SSH)) to only the jump box.
- Implement a strict removable storage policy. This will prevent malware from using storage devices to bypass other separation controls in place.
- Implement full disk encryption to protect the integrity of workstations as well as the confidentiality of any sensitive data.

By implementing these controls, in conjunction with the network architecture controls mentioned below, an organisation can reduce the ability of malicious actors to target and compromise privileged workstations.

Considerations when administering a cloud environment

One of the key controls for privileged workstations is enforcing a lack of internet connectivity. When administering an internet-based cloud environment, such as public cloud services, this recommendation may require some changes depending on the secure administration environment in place.

It is desirable that no internet connectivity be provided for privileged workstations. However, for the purposes of managing cloud resources exceptions can be made:

- A local jump server used solely for administering cloud-based assets could be used. This jump box would have internet connectivity, but be restricted to only accessing the domains or IP addresses required for accessing the necessary cloud services and other critical services such as patch servers.

- The cloud-based environment could be configured to only accept administration traffic from an organisation's public IP address range to prevent unauthorised access.
- A dedicated physical link may provide additional security for management connections to cloud services.
- In cases where a jump box is not feasible and privileged workstations must connect directly to cloud-based infrastructure, restrict workstations to only access the domains or IP addresses required for accessing the necessary cloud services.
- If the only cloud service being administered is a private virtual network (e.g. Amazon's Virtual Private Cloud or Microsoft Azure's Virtual Network) accessed via a VPN, then an organisation should determine whether internet connectivity restrictions should apply. The use of a local jump server is still recommended.

Logging and auditing

Log collection and analysis for the purposes of auditing is a core element of an organisation's defence of their networks. A comprehensive log collection and audit plan is a key method of detecting when security and administrative events of significance occur, allowing for detection of cyber intrusions and determining appropriate cyber security incident response activities.

When implementing a secure administration approach using a jump box, logging and auditing becomes easier as all administrative access should pass through the jump box. While the jump box is the most obvious target for logging and auditing activities, logging and associated activities should still be performed within the secure administration environment and servers to ensure that all key events are subjected to appropriate auditing.

All log events should be captured and stored in a centralised, backed-up location. Centralisation of event logs can allow easier log collation and correlation which aids the identification of suspicious events or anomalies. In addition, care should be taken to ensure time synchronisation across the environment (and by extension the logs) in order to allow for proper log analysis.

Considerations when administering a cloud environment

Cloud services may create their own logging events and files which may be accessed by an organisation for auditing purposes. Some cloud providers may even offer a logging cloud service which allows aggregation and storage of all activity and events across multiple cloud services.

If possible, all logs should be retrieved and incorporated into a centralised logging environment so a comprehensive picture of administrative events can be developed.

Management logs for cloud services should be regularly audited.

Network segmentation and segregation

Designing a network architecture based on security zones is an extremely effective method of restricting malicious actors' ability to move laterally within a network and gain access to administrative resources.

Controls for network segmentation and segregation in relation to a secure administration environment are:

- Create zones to logically separate different security domains. Examples of different zones may be user workstation zone, privileged workstation zone and asset zone.
- Restrict traffic flows between the zones. Especially for more sensitive security domains such as the privileged workstation zone and the asset zone, as well as sensitive management traffic types such as RDP and SSH.
- Management traffic should only flow between jump boxes and managed assets. When jump boxes are implemented the only assets which should be able to administer managed assets are the jump boxes.

Some examples of traffic restrictions that should be put in place are outlined in the tables below. By default, traffic between zones should be denied. These examples are designed around an implementation which utilises jump boxes.

Allowed traffic

Source	Destination	Traffic Type
Privileged workstation	Jump box	Management traffic (e.g. RDP, SSH)
Jump box	Managed assets	Management traffic (e.g. RDP, SSH)
User workstation	Managed assets	Business traffic (e.g. web, mail, business applications, application portals, Active Directory)

Blocked traffic

Source	Destination	Traffic Type	Explanation
User workstation	Managed assets	Management traffic (e.g. RDP, SSH)	User workstations should not be performing administration activities.
User workstation	Jump box	Any traffic type	User workstations should not interact with IT administration assets.
User workstation	Privileged workstation	Any traffic type	User workstations should not interact with IT administration assets.
Privileged workstation	Managed assets	Any traffic type	Privileged actions should be performed via the jump box.

Considerations when administering a cloud environment

The controls listed above can be applied to administering an Infrastructure as a Service (IaaS)-based network, particularly if a virtual private cloud is also implemented.

Depending on the functionality of a Software as a Service (SaaS) or Platform as a Service (PaaS) offering, it may be impossible or impractical to implement some of the controls listed above. For example, organisations may be limited to just controlling IP addresses from which a cloud asset can be administered.

Jump boxes

A jump box can enhance a secure administration environment by providing three main benefits:

- Single point of origin for administration connections. A jump box, or specific jump box zone, provides a single point of origin for administration connections, such as RDP or SSH, destined for critical assets. This can help reduce administrative overhead and technical complexity, thereby reducing the chance of configuration mistakes.
- Easier detection of anomalous or suspicious behaviour. If a jump box and administration zone have been configured it can be easier to identify anomalous behaviour. For example, regular users and workstations should not be administering critical servers and applications. As such, attempted RDP connections originating from a user workstation zone destined for an asset zone or jump box would be suspicious.
- Enables an effective and easier deployment of multi-factor authentication. With all administrative actions now performed via a jump box, it can be easier to deploy multi-factor authentication by centralising it at the jump box.

Without a jump box it may be more difficult to ensure that all assets can support the multi-factor implementation used by an organisation.

The configuration of a jump box is critical to realising the benefits listed above. Minor configuration mistakes can lead to a jump box becoming a security liability as it has wide ranging access across the breadth of an organisation's systems and is likely to be targeted.

In addition to application control, patching operating systems and applications, restricting administrative privileges, and other general hardening practices, the following procedural and technical controls should be implemented for a jump box:

- Implement strong device or computer authentication in addition to user authentication to ensure only trusted workstations are able to connect to the jump box. For example, the issuance of computer certificates to administrative workstations and the jump box combined with IP Security (IPSec) policy settings allowing remote network access to the jump box by only workstations with trusted certificates.
- Prevent direct internet access to and from the jump box.
- Only allow access to the jump box from the secure administration environment. Block traffic from all other sources to help protect the jump box from malicious actors who already have a presence on the network.
- Prevent privileged accounts from logging on to the jump box remotely. Administrators should connect to the jump box with an unprivileged administrator account and then elevate their privileges as required.
- Perform thorough auditing with automated analysis of all activities carried out on the jump box including logins, logouts, program execution, outbound connections and failed login attempts.
- Implement a host-based intrusion prevention and detection system (HIPS/HIDS) to identify and prevent malicious and suspicious activities.
- Investigate system crashes, unusual activity and unusual network traffic as a priority to determine if they are indicators of compromise.

The diagram below demonstrates the process flow of administration utilising a jump box.

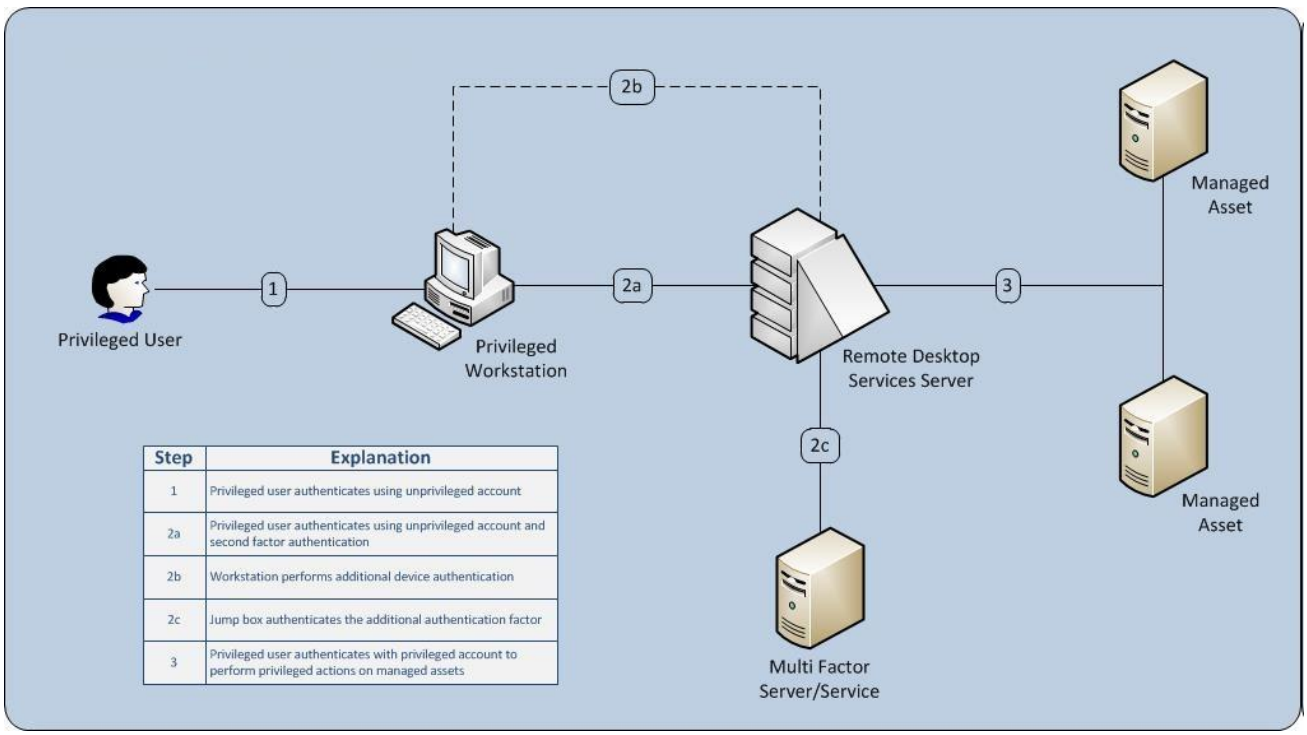

Step	Explanation
1	Privileged user authenticates using unprivileged account
2a	Privileged user authenticates using unprivileged account and second factor authentication
2b	Workstation performs additional device authentication
2c	Jump box authenticates the additional authentication factor
3	Privileged user authenticates with privileged account to perform privileged actions on managed assets

Considerations when administering a cloud environment

Jump boxes can be implemented in a variety of ways when administering a cloud environment. The key consideration will be the type of cloud service being administered.

- **Software as a Service.** As a cloud-based jump box is likely not possible with a SaaS offering, a local cloud-specific jump box can be used. Such a jump box should only be used for administering cloud services and not local assets since it will likely have to be able to connect to the internet.
- **Platform as a Service.** PaaS will have a similar approach as SaaS since organisations will likely be unable to establish a cloud-based jump box.
- **Infrastructure as a Service.** In addition to the local jump box approach outlined for SaaS and PaaS offerings, a cloud-based jump box approach could be used for IaaS services. In this scenario a cloud server is used as a dedicated jump box. Administrators would first connect to the cloud-based jump box, and then use it to administer other cloud-based servers.

Further information

The *Information Security Manual* is a cyber security framework that organisations can apply to protect their systems and data from cyber threats. The advice in the *Strategies to Mitigate Cyber Security Incidents*, along with its Essential Eight, complements this framework.

If your network is based on Microsoft Windows and Active Directory, the following two Microsoft articles contain advice relevant to many parts of this publication:

- *Mitigating Pass-The-Hash (PtH) Attacks and Other Credential Theft, Version 1 and 2*
- *Best Practices for Securing Active Directory*.

Further information on cloud computing is available in the *Cloud Computing Security for Executives* and *Cloud Computing Security for Tenants* publications.

Further information on implementing network segmentation and segregation is available in the *Implementing Network Segmentation and Segregation* publication.

Further information on the benefits of multi-factor authentication, and implementation considerations, is available in the *Implementing Multi-Factor Authentication* publication.

Further information on the Protected Users group is available in Microsoft's *Protected Users Security Group* article.

Further information on Managed Service Accounts is available in Microsoft's *Managed Service Accounts* article.

Contact details

If you have any questions regarding this guidance you can write to us or call us on 1300 CYBER1 (1300 292 371).