BeyondTrust

Adapting to the 2021 Telecommunications (Security) Act Regulations and Security Requirements

Understanding the TSA Code of Practice *and* How BeyondTrust Can Help

TABLE OF CONTENTS

New Legislation Means Big Changes

Even those outside of the telecommunications world will have heard of the government's instruction to remove providers such as Huawei from UK communications networks. What you may not know is that this forms part of a much broader effort to secure our digital communications.

In response to increasing reliance upon and greater threats to telecommunications networks across the UK , the government has amended the 2003 Communications Act with the Telecommunications (Security) Act 2021 ('the TSA'[1]). Within this act there are not only high-level security goals as to how telecommunications companies should be hardening themselves (defined in the Electronic Communications (Security Measures) Regulations 2022[2], but also specific technical guidance through the Telecommunications Security Code of Practice[3] on how this can be achieved. The Code of Practice covers the concepts and principles that should be borne in mind when implementing changes to comply with the TSA, finally breaking them down into a numbered list of measures tied to specific deadlines.

Throughout this article, we will discuss key points described in the TSA, the Code of Practice, and how BeyondTrust solutions can be leveraged to meet these technical requirements, highlighting specific measures where relevant.

[1]https://www.legislation.gov.uk/uksi/2022/933/pdfs/uksiem_20220933_en.pdf
[2]https://www.legislation.gov.uk/uksi/2022/933/contents/made
[3]https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1120531/E02781980_Telecommunications_Security_CoP_Accessible.pdf

Understanding Tiers and Timelines

Public telecoms providers will have different timeframes to implement the measures detailed in the Code of Practice depending on their size; this is laid out via a tiering system, depicted below. [4]

- **Tier 1 –** public telecoms providers with relevant turnover in the relevant period of £1bn or more.
- **Tier 2 –** public telecoms providers with relevant turnover in the relevant period of more than or equal to £50m but less than £1bn.
- **Tier 3 –** public telecoms providers whose relevant turnover in the relevant period is less than £50m, but who are not micro-entities.
- (Relevant period represents a 12-month reporting cycle)

The implementation deadlines for the measures described in the Code of Practice are set out in phases, the simplest of which need to be addressed for Tier 1 providers by the 31 March 2024 and Tier 2 providers by 31 March 2025. Providers will have longer to implement the more complex measures.

Below is a table of each of the deadlines and how they correspond to the groups of measures and tiers. For Tier 3 providers, the measures are voluntary and will not be enforced.

[4]Pg 7 Code of Practice – Explanation of terms

Deadline per Provider Tier	Implementation Requirements
Completed by 31 March 2024 (Tier 1 providers) **Completed by 31 March 2025** (Tier 2 providers)	• Overarching security measures • Management plane 1 • Signalling plane 1 • Third-party supplier measures 1 • Supporting business processes
Completed by 31 March 2025	• Management plane 2 • Signalling plane 2 • Third-party supplier measures 2 • Customer premises equipment
Completed by 31 March 2024 (Tier 1 providers) **Completed by 31 March 2025** (Tier 2 providers) **Completed by 31 March 2027** (All provider Tiers)	• Third-party supplier measures 3
Completed by 31 March 2027	• Management plane 3 • Signalling plane 3 • Virtualisation 1 • Third-party supplier measures 4 • Network Oversight Functions • Monitoring and analysis 1
Completed by 31 March 2028	• Management plane 4 • Signalling plane 4 • Virtualisation 2 • Monitoring and analysis 2 • Retaining national resilience and capability

Requirements: Asked or Told?

There are significant penalties detailed in the TSA for companies who the regulator deems to have not implemented appropriate controls after the deadlines have passed. These penalties could be up to 10% of turnover, or in the case of a continuing issue, £100,000 per day. With such severe penalties at stake, these requirements are not ones to ignore.

When specific technical guidance is provided, we often assume it must be followed to the letter. In the Code of Practice, this is not the case; instead, the government recognises there is more than one way to achieve the same result. They do specify that when companies decide to deviate from the suggestions made in the Code of Practice, they must justify these deviations and demonstrate that they are at least equally secure. What this means, practically, is that with the proper solutions, companies can take a flexible approach.

For this reason, the language used in the Code of Practice is worth exploring in further depth. Doing so gives us insight into where the government sees room for flexibility and where suggested controls are instead closer to mandates.

Explanation of Terms:

Where the term **'reduce'** is used in the regulations, it is expected that the provider will reduce risks as much as possible.

The terms **'shall'**, **'should'**, and **'may'** have been defined in relation to the guidance given in the remainder of the Code of Practice. This is to distinguish between where the government believes there is likely to be only one acceptable way of implementing the specific measure, and those which have potential alternatives.

Shall:
The use of the word **'shall'** indicates where government guidance states there is likely only one viable technical solution to secure the network or service in line with the regulations. We would not expect these technical solutions to vary as a result of different network configurations or business structures.

Should:
Where the word **'should'** is used in the guidance, the government views the solution provided as being the best way to implement the measures in the majority of cases. However, there are known alternatives that providers can deploy that may attain a satisfactory security outcome, depending on their network or service configuration and business structure.

May:
The use of the word **'may'** in the guidance indicates that providers are likely to have multiple options, all of which could deliver a satisfactory solution. There are likely to be differences between providers in their implementation.

When implementing controls to respond to the TSA, we should keep our eyes on those 'shall' components when we're looking to see how much flexibility providers are given.

Where BeyondTrust Can Help

The scope of the TSA is broad. For brevity's sake, we'll be highlighting a few notable areas where BeyondTrust solutions can ease the journey to compliance.

The Basics

Before we get into the deeper technical requirements, the first round of measures in 'Management Phase 1' are worth touching on.

These measures represent easy wins for those using BeyondTrust technologies. They include fundamental capabilities of our Privileged Access Session Management (PASM) tools. For example, M2.02 and M2.03 ("All privileged access shall be logged." and "Privileged access shall be via secure, encrypted and authenticated protocols whenever technically viable.")

Out of the box, BeyondTrust's session management products not only offer access logging, but also screen recording and a searchable index of events that took place during the session.

M2.05 requires the removal of default passwords.

When using our vaulting capabilities, passwords can be stored and injected in a way where they are never exposed to the user. Going a step further, both the **Privileged Remote Access** and **Password Safe** products offer credential management capabilities for a range of different credential types and systems, dramatically limiting the exposure risk of such accounts.

As you can see from the timeline table (pg. 5), there isn't long before these measures need to be in place. Providers who have not already covered them may want to consider what can be done off the shelf, as opposed to embarking upon a custom solution.

Privileged Access Workstations

One of the primary considerations from a Privileged Access Management (PAM) perspective is the proposal of the use of Privileged Access Workstations (PAWs) and the avoidance of a 'browse-up' architecture which the NCSC describes as one of its six security architecture anti-patterns.[5]

This guidance has the potential to significantly impact the logistics of providing privileged access because it requires dedicated devices. This architecture has been proposed in part because the threat model that applies to public telecoms providers in the UK is one which includes state-sponsored actors, and thus requires the most stringent controls.

A 'browse-up' architecture exists when a secure system is accessed from a less secure system. This would be the case if a corporate or personal device were used to access the management plane of a provider. These devices are less secure because of the more relaxed controls that are necessary for people to perform their day-to-day work (such as accessing emails or browsing the internet).

[5]https://www.ncsc.gov.uk/whitepaper/security-architecture-anti-patterns#section_3

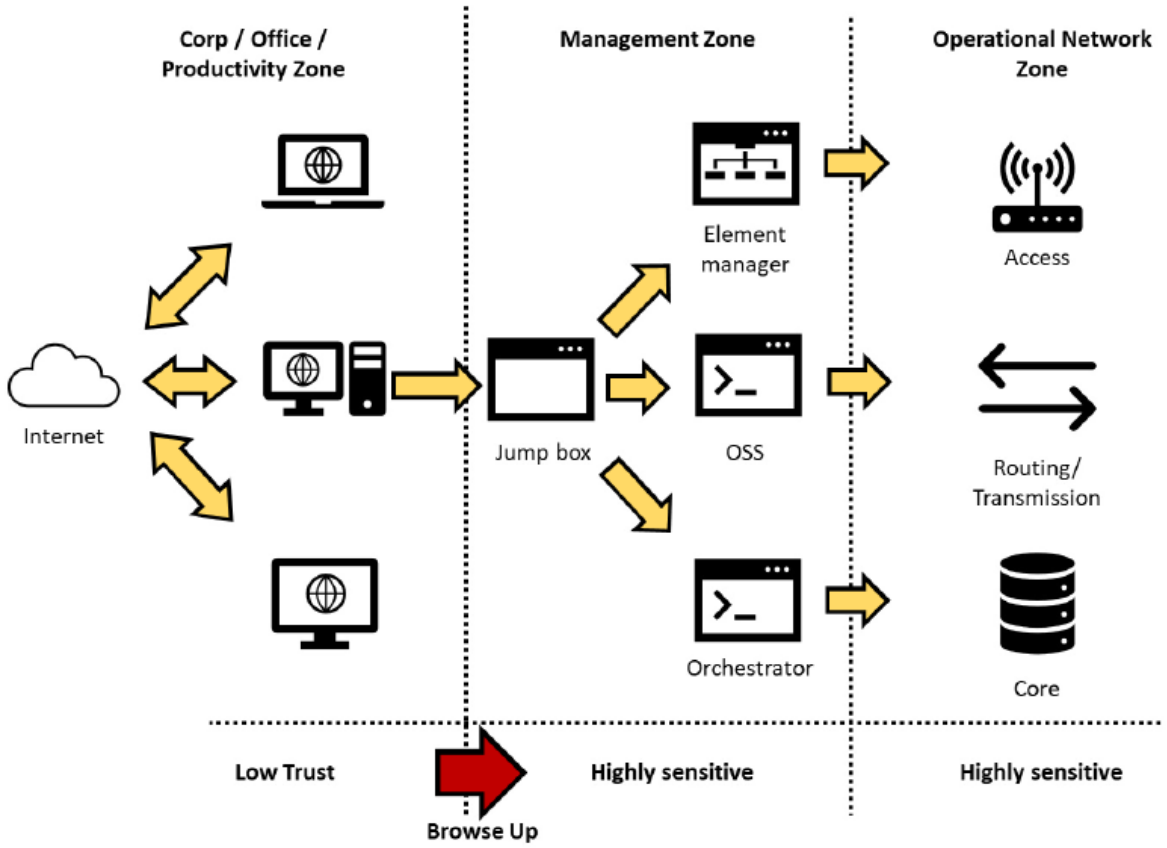

Figure 1: Browse-Up Architecture[6]

The government has been clear that the use of a secure connection mechanism (including the use of jump hosts) from non-PAW devices does not mitigate the risks associated with a browse-up architecture and does not constitute an effective control on its own.

PAWs must be segregated from corporate and personal services and activities, existing as a dedicated asset. If we examine section 2.17 of the Code of Practice, we will note that this falls under the 'shall' verb; the government does not anticipate alternatives to this model.

[6]Pg18 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1120531/E02781980_Telecommunications_Security_CoP_Accessible.pdf

2.17
Workstations dealing with general office productivity tools and external access to external services over the internet **shall** be logically or physically separate from those with any access to the management plane. Any administrative users who previously performed these functions via a single device will need to operate differently to protect their network.

Looking at two of the third-phase management plane measures, we see how strict this requirement is.

M11.24
A PAW shall only have access to the internet to the extent it is needed to carry out changes to security critical functions, and such access shall be secured (e.g., via VPN).

M11.25
The PAW shall only have access to internal-only business systems (e.g., not corporate email).

The need for this measure was challenged during the consultation period for the TSA. Several alternatives were proposed, and while the government has repeated statements that other mechanisms may be used if justified, they have stated that dedicated PAWs should be considered best practice.[7]

A browse-up architecture lowers the security of the more secure system to that of the less secure system. Using PAWs, providers are aligning with the "clean source principle[8]," the fundamental concept of which is described by Microsoft in their privileged access guidance as follows.

[7]https://www.gov.uk/government/consultations/proposal-for-new-telecoms-security-regulations-and-code-of-practice/outcome/proposals-for-new-telecoms-security-regulations-and-code-of-practice-government-response-to-public-consultation#fnref:8
[8]https://learn.microsoft.com/en-us/security/compass/privileged-access-success-criteria#clean-source-principle

"Any subject in control of an object is a security dependency of that object. If an adversary can control anything in control of a target object, they can control that target object. Because of this threat, you must ensure that the assurances for all security dependencies are at or above the desired security level of the object itself."

The government has alluded to further guidance being provided by the NCSC as to how PAW solutions can be built, but until we see otherwise, it seems clear that at least some degree of physical separation is non-negotiable.[9]

Managed Service Provider and Third-Party Administrators

Supply chain attacks have long been a worry for organisations trying to secure themselves, and the TSA shares the concern that lax controls at third-party organisations should not expose public telecoms infrastructure to threat. As a result, we see that third parties accessing the management plane of a provider must conform to the same security standards as the provider. Described in section 2.26 of the Code of Practice is as follows:

2.26

Managed service providers (MSPs) or third-party administrators (3PAs) are prize targets for attackers, as they will often have privileged access to multiple networks. Because of this, where these third parties have access to the management plane, they shall have to meet the same security principles as those employed by public telecoms providers themselves, and ideally shall use the same methods.

[9]https://www.gov.uk/government/consultations/proposal-for-new-telecoms-security-regulations-and-code-of-practice/outcome/proposals-for-new-telecoms-security-regulations-and-code-of-practice-government-response-to-public-consultation#fnref:8

We see in section 2.28 that it is the responsibility of the provider to oblige the third party to conform to these standards, and providers should be able to audit to confirm this. Section 2.28 goes into further detail.

2.28

To ensure that security controls are applied correctly, it will be essential for public telecoms providers to have contractual arrangements in place which oblige third party administrators to undertake this activity. It will also be necessary to have robust powers of audit to permit spot checks and ongoing monitoring of security governance arrangements. Public telecoms providers shall ensure they are able to fully control and monitor access by third parties into their management plane independently of the third party.

For MSPs who may be particularly concerned around increasing costs when deploying separate PAWs for each engineer and customer combo, clarification has been added to the Code of Practice that PAWs may make use of virtualisation and be used to access more than one secure system provided appropriate segregation is in place.[10] [11]

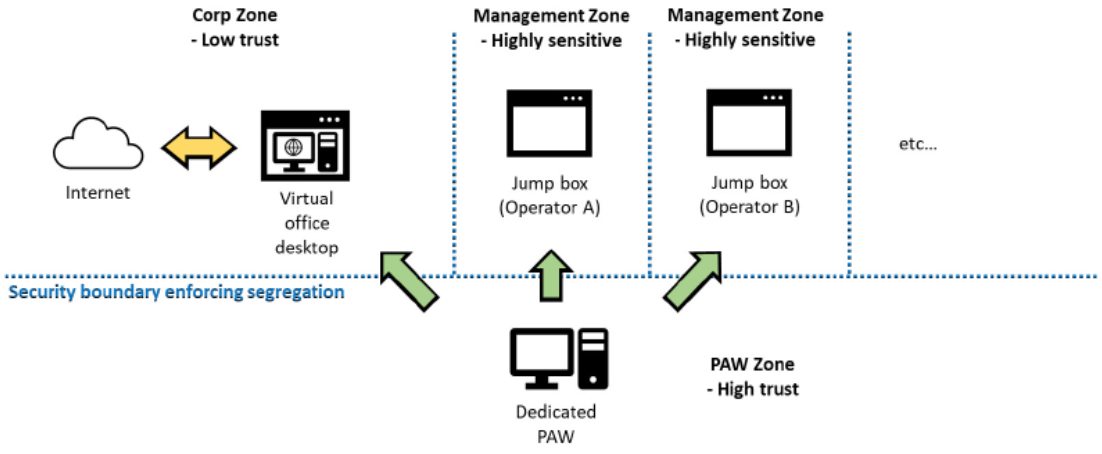

Figure 2: Third-party administrator secure access to multiple providers.[12]

[10]Pg 14 Code of Practice sec 2.27

[11]https://www.gov.uk/government/consultations/proposal-for-new-telecoms-security-regulations-and-code-of-practice/outcome/proposals-for-new-telecoms-security-regulations-and-code-of-practice-government-response-to-public-consultation#fnref:8

[12]Pg21 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1120531/E02781980_Telecommunications_Security_CoP_Accessible.pdf

2.27
Given a trusted device, 3PAs can access securely segregated management systems for multiple providers, as shown in Figure 3. [above] Critically, such an approach must maintain the security and integrity of the PAW, and segregation between each provider's management environment.

BeyondTrust's **Privileged Remote Access** can provide logical separation in a way which does not hinder productivity or negatively impact user experience. Each access console instance is permanently tied to the instance it has been deployed from; it cannot be reassociated with a new or different instance of **Privileged Remote Access**.

Through this immutable bonding of appliances and consoles, a third party can create a single device (a PAW) that uses multiple instances of the **Access Console** to allow access into a provider's management zones without fear of improper segregation between those providers. It is particularly important to make this as easy as possible; as we mentioned, providers will have to contractually oblige their service providers to conform to these requirements, and a cumbersome solution is unlikely to make those negotiations straightforward (Code of Practice, Section 2.28 pg.13).

Access Console:
The client used to open sessions through **Privileged Remote Access**. Available as a thick, web, or a mobile client. For those particularly attached to their own tools, it can even enable Bring-Your-Own-Tool functionality.

Newly-released Infrastructure Access Mode for the console streamlines the user experience even further when using native tools.

The Consequences of an Unclean Source

The requirement for physically separate privileged access infrastructure may feel to some like a difficult pill to swallow, but it is there for good reason.

Password vaulting vendor LastPass recently disclosed details around an August 2022 attack which showed exactly how risky browse-up architectures can be. The attack against one of their senior DevOps engineers exposed encrypted customer vault data and has subsequently sparked a fresh debate about BYOD and the use of personal devices.

An initial foothold was established after the use of compromised third-party software on a personal device enabled an attacker to perform remote code execution and deploy a keylogger. Because this personal device was also being used to access privileged data, it exposed the wider environment to attack. From that point, lateral movement into the privileged environment (a cloud storage instance) was possible through a combination of the keylogging and exploitation of MFA push spamming. Push spamming was a popular and effective breach tactic throughout 2022, with Microsoft, Cisco, and Uber all falling victim.[13]

We've written about this topic in further detail **here.**

The specific mechanisms of this attack illustrate why we must be cautious of trust in MFA alone, and by extension, any control, as a fool-proof way to ensure authorisation and authentication when an originating device is compromised.

[13]https://www.bleepingcomputer.com/news/security/mfa-fatigue-hackers-new-favorite-tactic-in-high-profile-breaches/

While MFA vendors are implementing features to limit the effectiveness of MFA push spamming, attacks will never stop evolving. When we consider a threat model which includes state-sponsored actors, there is little that can compare to physical segregation in terms of security.

Securing the PAW: Privilege Management

As part of maintaining a PAW as a clean source, BeyondTrust recommends not only limiting the activities that should be carried out (in line with the Code of Practice), but also limiting the activities that can be carried out. *Removing local administrative rights from users, controlling elevation, and application allow-listing* are powerful controls that can address measures such as M11.29, described below.

> **M11.29**
> A PAW shall prevent the execution of unauthorised code such as binaries or macros within documents.

BeyondTrust's **Privilege Management for Windows & Mac** and **Privilege Management** for Unix & Linux enable least privilege to be enforced in a manner that does not present an obstacle to useability or flexibility. Instead of granting local administrative rights to users, or relying on a support team for necessary elevations, these solutions can elevate applications and tasks, not users, on-demand per a role-specific policy.

[14]https://www.beyondtrust.com/resources/whitepapers/microsoft-vulnerability-report

People often consider the implementation of application control to be an enormous undertaking. BeyondTrust's Quick Start policies provide out-of-the-box templates based on years of implementation experience. With matching criteria that can be as broad or as narrow as you need, implementation becomes an incredibly achievable goal.

Controlling Access from the PAW: Privileged Remote Access

We see below a diagram from the Code of Practice, illustrating the desired browse-down architecture wherein a sensitive system is accessed only from a system in a zone of high trust (equal to that of the environment it is accessing). The PAW is the starting point for securing access in this architecture, but not the end of this story. Providers still need a mechanism to give appropriate, controlled, and recorded access to the highly sensitive zones from the PAW.

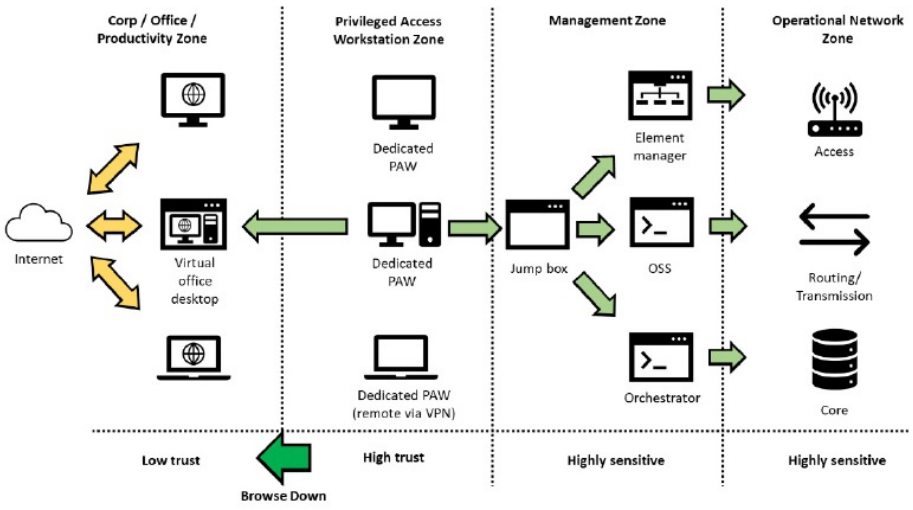

Figure 3: Browse-Down Architecture[15]

[15]Pg20 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1120531/E02781980_Telecommunications_Security_CoP_Accessible.pdf

As depicted in Figure 3, the user is in the high trust zone and can remote to a low trust zone to perform corporate activities if required, the reverse of what we would see in a browse-up architecture.

If we dive into the measures, we see a host of requirements around things such as credential storage M6.01, vault encryption M11.03, RBAC and non-repudiation M6.02, time-bound, ticket linked access M11.05, session recording M11.13 and approval flows M15.10, to highlight a few.

Such a range of capabilities will create a significant burden on providers if they choose to implement home-grown solutions, with so many moving parts and the amount configuration errors, human errors and support overhead also increases.

Deploying an enterprise solution means that providers can have confidence that the technology they are using is tried and tested, will continue to be developed, and will be fully supported in the event of any issues.

This is where BeyondTrust's **Privileged Remote Access** comes in, purpose built to provide role-based access control, auditing, and credential management when connecting to secure environments. There's no steep learning curve when onboarding users, even with external vendors where pushback can present a challenge.

Below is an example of an architecture using an on-prem deployment of **Privileged Remote Access** to grant access from a PAW which has been secured using **Privilege Management.**

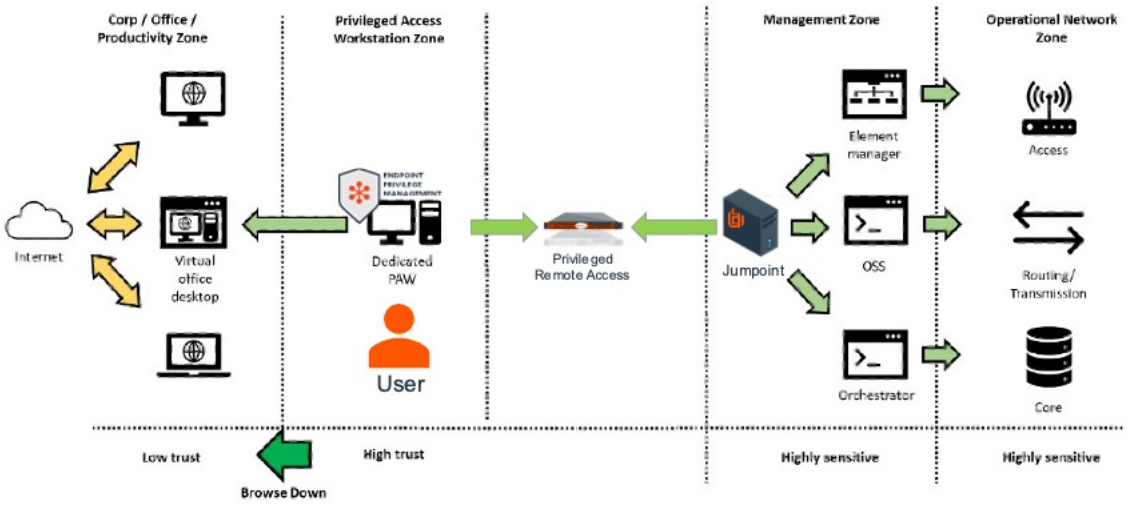

Figure 4: Architecture with Privileged Remote Access Implemented[16]

Privileged Remote Access is purpose built with zero trust as a priority. Sessions are brokered via connections coming out of the target environment, not into it, allowing you to block inbound traffic. This greatly reduces the attack surface and makes it almost impossible for a threat actor to perform enumeration by port scanning.

Each session brokered through the appliance is video recorded, with events within the session also logged and indexed. This session data is available on the appliance in an un-editable format for up to 90 days, and it can be moved to an external database for long-term archiving.

Out of the box, the solution offers strong authentication through TOTP, but also the ability to hand off authentication to an organisation's existing identity providers to maintain a single source of truth.

Through these features, providers can easily make sure the people who they want are getting access to what they want, and nothing else.

[16]Modified from pg20 https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1120531/E02781980_Telecommunications_Security_CoP_Accessible.pdf

Assumed Compromise and Infrastructure-as-Code

One of the concepts the Code of Practice brings up is assumed compromise: the idea that providers should always work on the basis that they have an attacker in their environment on the basis that, and that active steps should be taken to limit their dwell time and opportunity for lateral movement.

A highly effective way to deal with an unknown, dwelling attacker is to regularly restore to known-good images or configurations, thereby purging any mechanisms for persistence that may have been deployed but not detected.

We see discussion of the use of infrastructure-as-code and automation in section 2.69 of the code from the perspective of building scaling infrastructure, with measure M15.03 (described below) defining the government's expectation in terms of rebuilds.

> **M15.03**
> Any workstations or functions (e.g., jump boxes) through which it is possible to make administrative changes to network oversight functions shall be rebuilt from an up-to-date known-good software state on a yearly basis. This applies to the workstation or function's operating systems and above.

Regular restoration of infrastructure from a playbook ensures specific knowledge of an attack is not required for proactive measures. This move to a more ephemeral infrastructure model not only increases security, but substantially simplifies disaster recovery processes.

BeyondTrust's **Privileged Remote Access** has been rapidly expanding capabilities and streamlining user experience to serve infrastructure-as-code use cases.

Recent releases have seen UI changes adding a lightweight Access Console mode, the publication of API cookbooks that allow administrators to provision into IaaS vendors and their Privileged Remote Access instance simultaneously, and the development of an API CLI tool to support development and automation.

Soon, we'll be seeing a direct Terraform integration; features such as ephemeral agents that allow us to seamlessly mesh with ephemeral infrastructure and integrate with containerisation technologies such as Kubernetes.

Imagine the benefits of being able to spin up the infrastructure required to jump into a secure environment, from a gold build, on-demand. With these capabilities, not only do rebuilds become trivial, but they can be used to provision on a just-in-time basis, which significantly reduces the attack surface and opportunity for lateral movement.

Keep your eyes open for an increasing catalogue of integrations throughout this year.

Read More About Infrastructure Access with Privileged Remote Access.

Conclusion

Whether you're on a compliance journey of your own, or simply looking to assist your customers with theirs, BeyondTrust looks forward to supporting you with the next, most impactful step you can take toward securing privileged access.

ABOUT BEYONDTRUST

BeyondTrust is the worldwide leader in intelligent identity and access security, empowering organizations to protect identities, stop threats, and deliver dynamic access to empower and secure a work-from-anywhere world. Our integrated products and platform offer the industry's most advanced privileged access management (PAM) solution, enabling organizations to quickly shrink their attack surface across traditional, cloud and hybrid environments.

With a heritage of innovation and a staunch commitment to customers, BeyondTrust solutions are easy to deploy, manage, and scale as businesses evolve. We are trusted by 20,000 customers, including 75 of the Fortune 100, and a global partner network. Learn more at **beyondtrust.com**