CYBER ESSENTIALS PLUS

Preparation Guide

Test Case 1: Remote Vulnerability Assessment

What an internet attacker can see and reach. How the external scan works, why the April 2026 scope change widens it, and how to close your exposure.

Classification: Internal Use

Version 1.0

Test procedures follow the NCSC Cyber Essentials Plus Test Specification v3.2 (April 2025). Control language and the April 2026 marking changes follow the Requirements for IT Infrastructure v3.3 and the Danzell question set, in force for assessment accounts created from 27 April 2026. Remediation is shown on Microsoft platforms (Entra, Microsoft 365, Defender, Intune, Windows), which is where most applicants operate.

Ridgeline Cyber ridgelinecyber.com

Licence, Copyright, and Disclaimer

This page may be removed before the guide is issued internally.

Copyright

Crown copyright material

Licence

Disclaimer

This document is preparation guidance, not a certification, and is not legal advice. Only a licensed Certification Body can award Cyber Essentials Plus. [Draft pending legal review.]

Ridgeline Cyber · Northgate Engineering Ltd, Company No. 04827311 · ridgelinecyber.com

Contents

1. What this test checks

Test Case 1 looks at your organisation from the outside, the way an opportunist attacker does. The assessor runs a vulnerability scan against every internet-facing IP address in your scope, including Infrastructure-as-a-Service, across a recommended set of TCP and UDP ports, and reviews every internet-accessible service it finds. The purpose is narrow and blunt: can a low-skill internet-based attacker hack into your systems using typical, widely available methods?

This is the one test that is not sampled. It covers your whole external footprint, so a single forgotten service on a single public address is enough to fail it.

The rule for this test

Every internet-accessible service discovered on your in-scope addresses is assessed. A single failing service fails the whole test case.

Where you use dynamic IP addresses, scope can be defined by DNS entries instead. Take care that services like carrier-grade NAT do not send assessment traffic to the wrong destination.

2. The April 2026 scope change widens your footprint

The v3.3 requirements quietly removed two words from the scope conditions, and the effect is that more of your estate is now exposed to this test. Previously a device was in scope if it could accept incoming connections from untrusted internet hosts, or allow user-initiated outbound connections. The words 'untrusted' and 'user-initiated' have gone.

In practice this means a server that only accepts connections from a trusted, whitelisted source, or one that only makes automated outbound connections such as reaching a vendor update service, can now fall in scope where before it could be excluded. A back-end server behind a proxy, or a system that only phones home to a trusted endpoint, may now need the same scrutiny as a public web server.

Re-check your external footprint against the new wording

Do not assume your scope from last year still holds. Enumerate everything that accepts any inbound internet connection or makes any outbound internet connection, not only the obviously public services, and confirm what is genuinely in scope before you scan.

3. Run the test yourself

Mirror the assessor by scanning your own external footprint first with an unauthenticated external scan, using a tool of the kind approved for the scheme, such as a PCI Approved Scanning Vendor product like Tenable, Qualys, or Rapid7. Do it in three steps.

1. Enumerate every public IP address you own or use, including static and dynamic addresses, the public IPs of your Azure and other cloud IaaS, and any address a third party exposes on your behalf. Confirm DNS records resolve where you expect.
2. Scan all of them across the recommended TCP and UDP port range, not just the few ports you think are open.

3. Review every service the scan reports. For each, decide whether a low-skill attacker could exploit it: an unpatched service, an exposed management interface, weak or legacy protocols, or a service that should not be on the internet at all.

Enumerate the full footprint, not the obvious one

The gap that fails this test is almost always a forgotten address, not a known one. Your external footprint is wider than the firewall in the server room. It includes the public IPs of cloud load balancers, application gateways, and platform services such as Azure App Service; addresses a managed service provider exposes for you; and anything published from a branch office or a piece of operational technology. Pull the list from your cloud portals and your asset inventory together, and reconcile it against what the scan actually finds, because the scan will find addresses you forgot you had.

Scan across the full recommended set of TCP and UDP ports rather than the handful you expect to be open. The point of the exercise is to find the service you did not know was listening, on the port you were not watching, which is exactly what an opportunist attacker sweeps for.

4. Close your exposure on Microsoft platforms

Stop exposing management interfaces

The most common and most serious finding is a remote administration interface open to the internet: RDP, SSH, or a web admin panel. Do not publish these. For Azure virtual machines, remove public RDP and SSH and reach them through Azure Bastion, and enable Just-in-Time VM access in Microsoft Defender for Cloud so ports open only briefly, only to approved sources, and only on request. Put administrative access behind a VPN or Entra-governed access rather than on the open internet.

Control what is published at the boundary

Use network security groups and Azure Firewall, or your perimeter firewall, to allow only the inbound services that genuinely must face the internet, and deny everything else by default. Every service you do expose must be supported, patched within the 14-day window, and configured to current standards, with legacy protocols and weak cipher suites disabled. Treat the boundary as an allow-list, not a deny-list.

Watch the cloud-native ways things become public without anyone opening a firewall port. A storage account left open to anonymous access, an App Service without access restrictions, a database with its public endpoint enabled, or a management plane exposed to the internet are all findings here. Use Microsoft Defender for Cloud's secure score and recommendations to surface internet-exposed resources across the subscription, and disable public network access on platform services that do not need it, reaching them over private endpoints instead.

Patch and harden what remains

For the services that must stay public, the same patching discipline as Test Case 2 applies: a vendor fix available for more than 14 days that addresses a critical or high vulnerability fails you here too if the service is reachable. Harden them with secure configuration baselines and remove default content, default accounts, and unnecessary features.

Remember how this test is scored: every internet-accessible service is judged individually, and a single service that an attacker could exploit fails the whole test case. There is no sampling to hide behind and no averaging across a mostly-clean estate. Treat each exposed service as a pass-or-fail decision in its own right, and the safest position for any service you cannot justify keeping public is to take it off the internet entirely.

5. Common failures, and how to clear them

What the scan finds	Why it fails	How to clear it
Exposed RDP or SSH	Remote admin reachable by an internet attacker.	Remove public access; use Azure Bastion and Just-in-Time access, or VPN.
Exposed web admin or management panel	Administrative surface on the internet.	Restrict to internal or VPN; put behind Entra-governed access.
Unpatched internet-facing service	Exploitable by low-skill methods.	Patch inside 14 days; remove the service if not needed.
Legacy TLS or weak ciphers	Known-weak configuration.	Disable deprecated protocols; apply a current TLS baseline.
Newly in-scope back-end via proxy or outbound	v3.3 scope change pulled it in.	Re-scope, then scan and harden it like any public service.

6. Evidence to capture

Record proof against this test case in the Evidence Capture Pack:

- Your own external scan report showing no exploitable internet-accessible services, dated close to the audit.
- Your list of all in-scope public IP addresses and DNS entries, reconciled against the asset inventory.
- Firewall and network security group rules showing a default-deny boundary and only justified services exposed.
- Evidence that remote administration is not internet-facing: Bastion, Just-in-Time access, or VPN configuration.