What to Verify for CPCSC Level 1 Boundary Protection

Understanding Boundary Protection Under ITSP.10.171

When a Canadian defence supplier works toward CPCSC Level 1 compliance, the boundary between its internal network and any external network becomes a central focus. The Controlled Goods Program and the IT Security Standard for Canadian Defence Suppliers (ITSP.10.171) require that the system boundary be clearly defined and actively defended. For Level 1, the emphasis is on practical, verifiable controls that prevent unauthorized access and limit exposure. This article walks through the key evidence points an assessor will look for, so you can prepare your documentation and configuration with confidence.

Defining the System Boundary and Managed Interfaces

The first piece of evidence is a clear statement of where your system boundary lies. You need a diagram or a written description that shows every network connection crossing that boundary. Each connection must pass through a managed interface - typically a firewall or a router with access control lists. The assessor will want to see that no direct, unmanaged path exists between the internal network and the internet or any other untrusted network. Your boundary documentation should list every managed interface, its role, and how it enforces the organization's security policy.

Firewall Rule Set and Default Deny Logic

At the heart of boundary protection is the firewall rule set. For Level 1, the rule set must follow a default deny posture: all traffic is blocked unless a specific rule explicitly allows it. Evidence for this includes a copy of the active rule set, annotated to show which rules are in effect and why. You should also document the review process for these rules - who approves changes, how often the rule set is audited, and how stale or unused rules are removed. A rule set that has not been reviewed in months will raise questions about whether the boundary is still secure.

Network Segmentation and Internal Boundaries

Boundary protection is not only about the edge of your network. Internal segmentation matters too. The standard expects that sensitive systems - such as those handling Controlled Unclassified Information (CUI) or supporting critical functions - are placed on separate network segments. Evidence for segmentation can include VLAN configurations, firewall rules between segments, or router ACLs that limit east-west traffic. The goal is to show that a compromise on

one segment does not automatically give an attacker access to all systems. Your documentation should explain the segmentation scheme and how it maps to data classification or system criticality.

Patch Window and Patch Report Evidence

Boundary devices - firewalls, routers, intrusion detection systems - must be kept up to date with security patches. An assessor will ask for a patch window: the maximum time allowed between a patch release and its deployment on boundary devices. This window should be documented in your security policy. You will also need a patch report that shows the current patch status of each boundary device, including the version number and the date the last patch was applied. If any device falls outside the defined window, you need a risk acceptance or a plan to bring it back into compliance. The patch report is a concrete piece of technical evidence for Level 1 controls that demonstrates your organization's discipline in maintaining the boundary.

Endpoint Console and Malware Blocking on Boundary Systems

Boundary devices themselves are endpoints that can be compromised. You should show that each managed interface has an endpoint console - a centralized management tool - that enforces security settings and reports on the device's health. The console should log configuration changes, alert on anomalies, and allow remote verification of the device's state. Additionally, any boundary device that runs a general-purpose operating system must have malware blocking software installed and active. Evidence for this includes a screenshot or report from the endpoint console showing that malware protection is enabled, definitions are current, and no active infections are present. This combination of centralized management and active defense closes common gaps that attackers exploit.

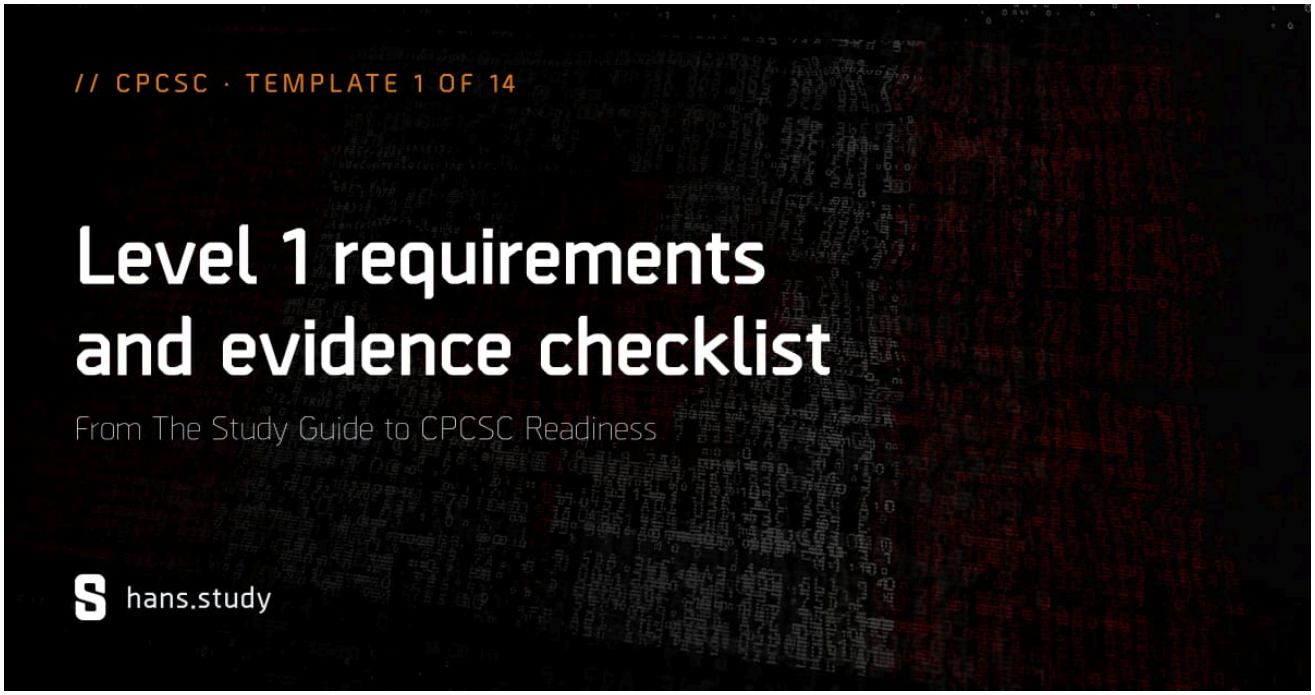

Logging and Monitoring of Boundary Traffic

To prove that the boundary is actually protected, you need logs. The assessor will look for evidence that traffic crossing the boundary is logged consistently. This includes connection logs from the firewall, authentication logs from remote access services, and any alerts from intrusion detection or prevention systems. Your logging policy should define what is logged, how long logs are retained (typically at least one year for Level 1), and how logs are protected from tampering. A log review process - whether manual or automated - must be in place to spot suspicious activity. Without logs, you cannot demonstrate that the boundary controls are working as intended.

Configuration Management and Change Control

Boundary devices change over time: rules are added, interfaces are reconfigured, firmware is updated. Each change is a potential source of misconfiguration. The assessor will want to see a change control process for boundary devices. Evidence includes a change request form, an approval record, and a post-change verification that the device still enforces the security policy. Configuration backups are also important. You should keep a known-good baseline configuration for each device and compare it against the running configuration periodically. A configuration drift report that shows differences between the baseline and the live state is strong evidence that you are in control.

Testing the Boundary: Vulnerability Scans and Penetration Tests

A static configuration is not enough; you must test that the boundary holds. For Level 1, periodic vulnerability scans of the boundary devices and the network segments behind them are expected. You should retain scan reports that show no critical or high-severity vulnerabilities on the boundary. If vulnerabilities are found, you need a remediation plan and a timeline. Penetration testing is not always required at Level 1, but if you have done one, it adds weight to your evidence. The key is to show that you have actively looked for weaknesses and addressed them. Scan reports, along with the remediation records, form another layer of **technical evidence for Level 1 controls**.

Putting It All Together: A Practical Evidence Package

When you assemble your evidence for a CPCSC Level 1 assessment, think of it as a story: this is my boundary, here is how I enforce the rules, here is how I keep my devices current and healthy, and here is how I verify everything works. Each piece of evidence should be easy to find and clearly labelled. A common mistake is to provide raw logs or configuration dumps without any explanation. Instead, include a brief narrative that maps each piece of evidence to a specific ITSP.10.171 requirement. For example, next to your firewall rule set, note which requirement it supports and how the rule set meets the intent. This makes the assessor's job easier and reduces the chance of follow-up questions. Preparing boundary protection evidence for CPCSC Level 1 is largely about being thorough and organized. The controls themselves are not exotic; they are the basics of network security done right. The challenge is documenting them in a way that leaves no doubt. By verifying each of the areas above - boundary definition, rule set, segmentation, patching, endpoint management, logging, change control, and testing - you build a complete and defensible evidence package.