Common Perimeter Exposures and What You Can Do About Them

Certain types of attacks are straightforward and easy to understand. Email phishing tries to trick users into opening a link or entering their credentials. Distributed denial of service attacks flood websites with so much traffic that the real communications can't get through.

Other network attacks are more difficult to understand. The details of the exposure can have a big influence on the ease of the attack and the nature of the impacts. In this paper, we'll highlight five of the most common perimeter exposures and discuss how you can address them.

Background

Protecting your attack surface should start with a data-driven approach. Lots of good information exists detailing how adversaries search for vulnerable systems across the internet. The most important takeaway is that while the internet is a large place, it is actually very small from an attacker's perspective. In less than an hour, anyone can have a list of every single exposed system of a given type across the entire internet.

In one experiment, Cortex® Xpanse™ set up honeypots to see who scans the internet. Figure 1 shows how different actors have different strategies to get data. Some only scan on HTTPS (Web Only), and some alternate their scans across ports throughout the week to get more data (Alternating). Others "hammer" the internet, scanning everyone on everything all of the time (Hammer).

This gives them a lot of data but can also be noisy and lead to getting block listed. With so many people scanning the internet, someone is always waiting for a new device to become exposed.

Xpanse also analyzed what attackers are looking for. Figure 2 shows a time period around the WannaCry and NotPetya outbreaks. Nearly 50% of all scan traffic was looking for vulnerable server message block (SMB) instances on open ports on the spreading mechanism for those exploits. Telnet is also a popular target, followed by database exposures.

Attackers don't just focus on registered ranges to find exposed devices; they target the cloud as well. Many organizations don't know where their assets reside in cloud environments, leading to more accidental misconfigurations and exposures.

When a system is left out in the open, it's only a matter of time before someone finds it. The impacts can include network intrusion, business disruption, data loss or theft, and reputation damage. Some of the most well-known breaches of the last few years have occurred via these exposed systems.

Exposures

Some of the most common exposures are also the simplest. Attackers don't need zero-day exploits when insecure systems are left open to the public. Prioritizing the following exposures can go a long way to making sure your network and data stay secure.

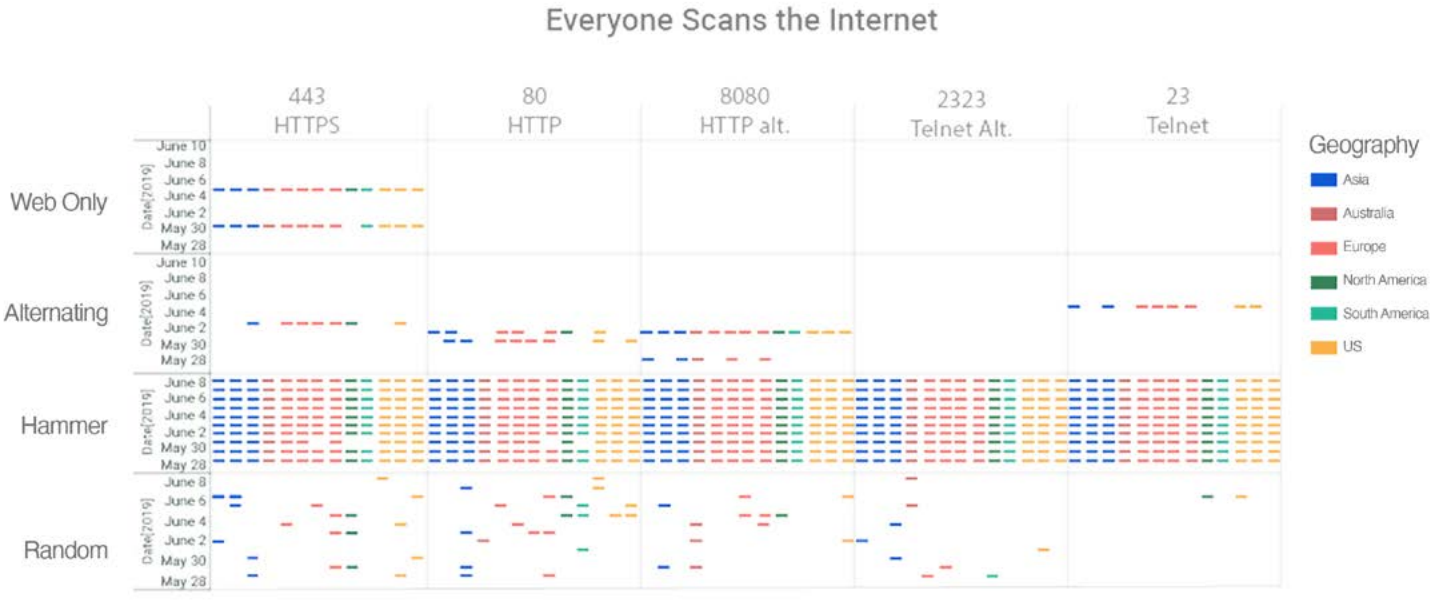

Figure 1: Attackers' internet scanning strategies

Remote Desktop Protocol

The Problem

Remote Desktop Protocol (RDP) is a Microsoft-based service that allows users to remotely connect to and control a computer. It provides a screen share that makes it look as if you are sitting in front of that computer's monitor.

RDP is a useful program, but it was never intended to be public-facing. Unfortunately, users commonly open their computers up to the world on RDP to get business done, or mistakenly expose their laptops when their endpoint firewall is misconfigured.

RDP is one of the most common entry points for ransomware. Several variants specifically target RDP, including SamSam, CRYSIS, and LowLevel04. Attackers who identify an exposed RDP instance have access to a login screen, where they can launch a brute-force password-guessing attack. If the password is easily guessed, the attacker gains access to a user's account.

LabCorp experienced a significant business disruption incident via RDP ransomware, impacting 7,000 systems and 1,900 servers.[1] The impact can be even greater for other organizations depending on which system is exposed and how it spreads internally.

What You Can Do

Organizations can protect themselves with several layered controls:

- ✓ Don't permit RDP to be public-facing. Block at the port/protocol level via the machine build and the endpoint firewall.
- ✓ Enforce login attempt lockouts.
- ✓ Consider using multi-factor authentication, even for user workstations.

RDP exposures can be difficult for organizations to detect on their own because they frequently occur when IT assets are misconfigured and outside of core networks that are regularly monitored by IT staff, such as when employees work from coffee shops and conferences. As a result, only a whole-internet approach to RDP discovery can ensure that organizations have no such services exposed.

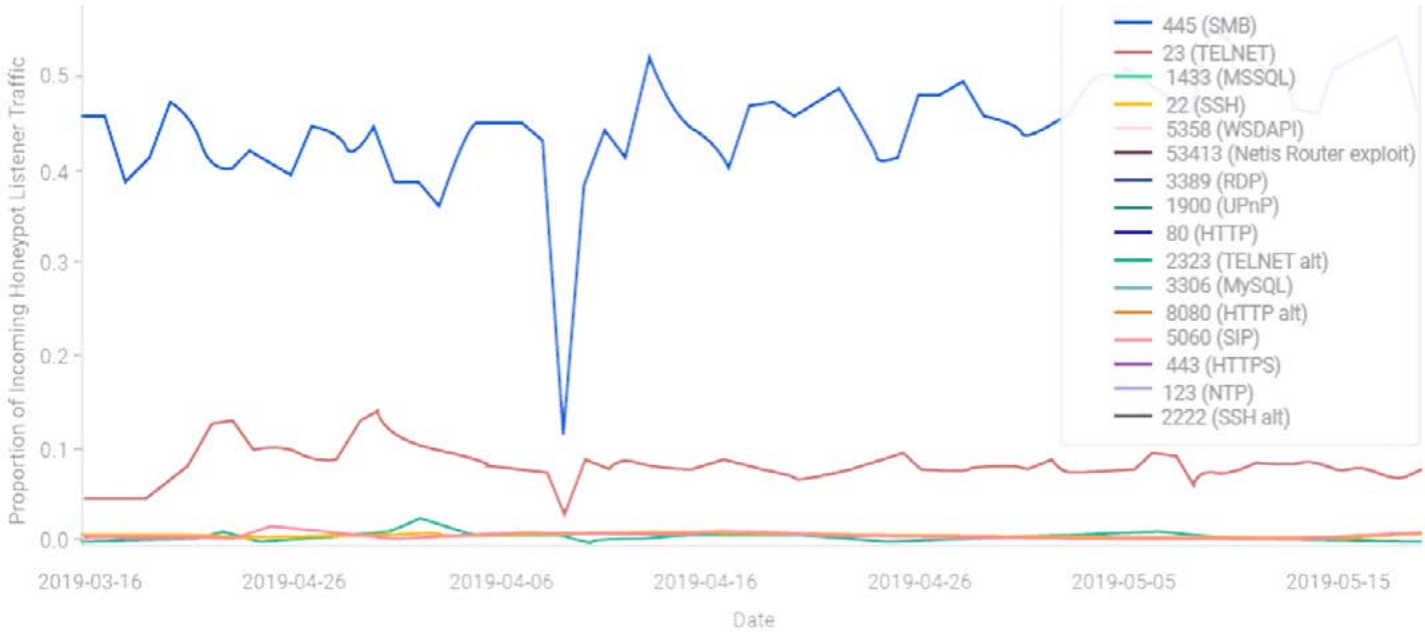

Figure 2: Internet scanner traffic—relative volumes by port

1. "LabCorp still recovering from ransomware, won't say if it's SamSam," Healthcare IT News, July 20, 2018, https://www.healthcareitnews.com/news/labcorp-still-recovering-ransomware-wont-say-if-its-samsam.

Telnet

The Problem

Telnet is an unencrypted remote access protocol. It is sort of like RDP without the desktop interface, using terminal commands instead. Telnet is an extremely old protocol that has no real requirement to be used today.

It is unencrypted, meaning any passwords or commands entered over it are sent in cleartext and can be intercepted by an attacker. Even Telnet's replacement (SSH) is not permitted to be public-facing at most organizations these days.

Telnet is so old and insecure that it is one of the most consistently scanned protocols over time. It's often used to infect devices for use in a botnet and has led to some of the largest internet-scale attacks ever seen. In addition to the risk of plaintext credentials submitted to a Telnet server being sniffed over a network, Telnet software implementations were not designed with security in mind and frequently contain exploitable bugs, making Telnet an easily exploited vector from which malicious actors can perform lateral movement from the internet to internal networks.

What You Can Do

Organizations should never have Telnet (or any other unencrypted remote access service) reachable from the public internet. If a legacy Telnet service must be run, make it accessible only on a segmented internal network and require strong authentication from other devices to connect to that network.

As Telnet is frequently associated with legacy systems, don't focus discovery and detection efforts only on networks where normal day-to-day IT operations occur, but also on lesser-known parts of your network, such as ISP-provided equipment for your regional offices.

Systems like these are often overlooked in asset management systems, and frequently contain misconfigurations, making them ideal targets for malicious actors to gain access to your network without your knowledge.

Exposed Databases

The Problem

With the ease and cheapness of cloud services, organizations are using basic computing resources at a growing rate. Unfortunately, this also results in millions of insecure databases. It's easier than ever for a developer to pull out a credit card and spin up a database in IP space that isn't owned or monitored by his or her organization. Test databases are supposed to use test data, but production data is often used instead. With no other controls or protections in place, huge volumes of data can become exposed in minutes.

Database servers should generally never be connected to the public internet. Securing them has always been notoriously difficult, and even allowing the database to be queryable can be risky. Database security is very difficult to get right, which makes it a huge target for hackers. Instead of needing to penetrate a network and find data, database servers prepackage all of the valuable data for the attackers. Dozens of data breaches were caused by exposed database servers in cloud environments in 2018.

What You Can Do

Organizations should never have database servers accessible from the public internet. A robust discovery program looking for SQL databases such as MSSQL, MySQL, Postgres, MongoDB, Elasticsearch, and memcached is critical. It's important to look for what the attackers are already searching for.

Exposed Engineering Systems

The Problem

Developers and engineers tend to focus on getting things done, sometimes at the expense of security. Verifying that policy is being followed is hard enough on a core network, but when employees can spin up their own infrastructure outside of your corporate environment, it can be nearly impossible to detect.

Exposed engineering systems can pose many risks. Development, staging, and QA environments are often exposed briefly for testing and are supposed to be taken offline after a few hours. They often aren't included in vulnerability scanning or patch management because they are assumed to be specially managed. However, these assets are often forgotten and then end up persisting for months or even years. Since they aren't actively monitored or managed, they may become vulnerable as new exploits are released.

In addition, these systems are usually not provisioned or configured to production security standards in the first place, and may contain sensitive data that was only intended for testing and integration purposes, but was then accidentally left on the machines, putting that data at risk of compromise.

What You Can Do

Organizations need routine discovery and monitoring of everything they have connected to the internet. Taking an outside-in view, just like an attacker does, can uncover systems that look normal from the inside but look like a big bullseye from the outside. Test, development, and staging environments should reside behind the firewall and only be exposed for brief testing periods required.

Certificate Issues

The Problem

Certificate health is one of the basic tenets of good cyber hygiene. Despite this, certificate issues remain incredibly common across large and small organizations. Certificates that are expired, self-signed, or that use a deprecated signature algorithm all provide a road map for attackers to find unmonitored and potentially unprotected systems.

On its own, certificate health is sometimes considered merely a "hygiene issue" by some IT staff—not an ideal state of affairs, but also not the most urgent of tasks to fix. While this may be the case in certain circumstances, there are many ways in which poor certificate health is a leading indicator (to both IT staff and potential attackers) of larger issues on a given system.

For example, an expired certificate is not necessarily a large security risk in itself, though it can in a number of cases be linked to other more concerning security issues. If a certificate has been expired for multiple years, it is a strong indication that the system the certificate is on likely isn't being actively managed. If the certificate has been expired for two years and hasn't been updated, then security patches probably haven't been applied, and the software probably hasn't been updated either. Attackers use certificate issues as indications that a device might be insecure.

What You Can Do

Anything that is expired, self-signed, or that uses deprecated algorithms should almost certainly be both known to IT staff and not part of a business-critical system.

Users often can't even access those systems because of browser warnings. If it doesn't have a healthy certificate, it likely doesn't need to be public-facing and should be taken offline.

Conclusion

The cause of many data breaches and hacks these days is not state-of-the-art exploits or zero days. Instead, basic security lapses are driving many of the largest incidents. Keeping up on basic hygiene and minimizing your attack surface area are some of the most impactful things you can do to keep your organization secure.

Xpanse discovers unknown devices and what they're talking to. We specialize in tracking down rogue devices across the internet so that they can be remediated. With Xpanse, you can keep a constant eye on your entire globally distributed network to make sure that nothing slips through the cracks. Own your perimeter before someone else does.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com

cortex-xpanse-Common-Perimeter-Exposures-and-What-You-Can-Do-About-Them-wp-05132021