



2021 Cortex Xpanse Attack Surface Threat Report

Lessons in Attack Surface Management from Leading
Global Enterprises

May 2021 | Estimated reading time: 15 minutes

Every time a new security vulnerability surfaces, a frenzied race kicks off between attackers scanning the internet to identify vulnerable systems and defenders scrambling to implement patches and other mitigations to protect their networks.

Computing has become so inexpensive that a would-be attacker need only spend about \$10 to rent cloud computing power to do an imprecise scan of the entire internet for vulnerable systems. We know from the surge in successful attacks that adversaries are regularly winning the race and finding vulnerable assets before defenders can patch new vulnerabilities. It's hard to ignore the increasingly common firsthand experiences with breaches disrupting

our digital lives, as well as the continuous flow of news reports chronicling the surge in cyber extortion.

To help enterprises gain ground in this battle, the Palo Alto Networks Cortex® Xpanse™ research team studied the public-facing internet attack surface of some of the world's largest businesses. From January to March, we monitored scans of 50 million IP addresses associated with 50 global enterprises to understand how quickly adversaries can identify vulnerable systems for fast exploitation.

Nearly one in three vulnerabilities we uncovered were due to issues with the widely used Remote Desktop Protocol (RDP), use of which has surged since the beginning of 2020 as enterprises expedited moves to the cloud to support remote workers during the COVID-19 pandemic. This is troubling because RDP can provide direct admin access to servers, making it one of the most common gateways for ransomware attacks. They represent low-hanging fruit for attackers, but there is reason for optimism: most of the vulnerabilities we discovered can be easily patched.

Here are our key findings.

- **Adversaries Are at Work 24/7**

Adversaries work around the clock to find vulnerable systems on enterprise networks that are exposed on the open internet. Exposure of enterprise systems has expanded dramatically over the past year to support remote workers. On a typical day, attackers conducted a new scan once every hour, whereas global enterprises can take weeks.

- **Adversaries Rush to Exploit New Vulnerabilities**

As soon as new vulnerabilities are announced, adversaries rush to take advantage. Scans began within 15 minutes after Common Vulnerabilities and Exposures (CVE) announcements were released between January and March. Attackers worked faster for the Microsoft Exchange Server zero-days, launching scans within five minutes of Microsoft's March 2nd announcement.

- **Vulnerable Systems Are Widespread**

Xpanse discovered that global enterprises found new serious vulnerabilities every 12 hours, or twice daily. Issues included insecure remote access (RDP, Telnet, SNMP, VNC, etc.), database servers, and exposures to zero-day vulnerabilities in products such as Microsoft Exchange Server and F5 load balancers. Experiencing one issue every 12 hours highlights the ephemeral nature of today's IT

infrastructure, where not only infrastructure changes but so does the vulnerability footprint.

- **RDP Amounted to One-Third of All Security Issues**

Remote Desktop Protocol accounted for about one-third of overall security issues (32%). Other commonly exposed vulnerabilities included misconfigured database servers, exposure to high-profile zero-day vulnerabilities from vendors such as Microsoft and F5, along with insecure remote access through Telnet, Simple Network Management Protocol (SNMP), Virtual Network Computing (VNC), and other protocols. Many of these high-risk exposures can provide direct admin access, if exploited. In most cases, these vulnerabilities can be patched easily, yet they represent low-hanging fruit for attackers.

- **Cloud Comprised the Most Critical Security Concerns**

Cloud footprints were responsible for 79% of the most critical security issues we found in global enterprises. This highlights how the speed and nature of cloud computing drives risk in modern infrastructure, especially considering how quickly cloud environments have grown over the past year as enterprises moved computing off-premises to enable the surge in remote work during the COVID-19 pandemic.

1. Key Findings

Mean Time to Inventory

Mean time to inventory is a race between attackers and security teams, and attackers have a significant edge.

The mean time to inventory (MTTI) is far faster for threat actors than for enterprise security teams. Our analysis found that threat actors scan to inventory vulnerable internet assets once per hour and even more frequently—in 15 minutes or less—following CVE disclosures. Meanwhile, global enterprises need 12 hours, on average, to find vulnerable systems, and that assumes the enterprise knows about all assets on its network.

Let's look at two CVEs to understand the spectrum. Scanning started 15 minutes after the release of a CVE for a vulnerability that enabled remote access to products from a maker of 'prosumer' networking devices. By contrast, we saw large-scale scanning begin just 5 minutes after the high-profile disclosure of Microsoft Exchange Server and Outlook Web Access vulnerabilities.

The vulnerability management system most enterprises follow is not designed to cope with the modern reality of a growing attack surface. The director of Dartmouth College's Institute for Security, Technology, and Society, V.S. Subrahmanian, recently warned that "Cybersecurity Needs a New Alert System," in a Wall Street Journal [op-ed](#) that details flaws in the communication process for patches from vendors to customers. This, however, is only half the story.

The process of vulnerability management within security teams is also flawed. Like antivirus systems, scanners rely

on a database of known CVEs—making them only as good as the latest update. This means you wait hours or days for an updated CVE profile, which consequently translates to a longer overall patching process that takes days. Worse, vulnerability scanners query only known devices to see what is exposed. What about unknown assets?

For things enterprises don't know about, third parties—usually quarterly [penetration tests](#) or red teaming—perform partial asset enumeration to find and test infrastructure. (The Palo Alto Networks Unit 42 incident response team typically commences most breach investigations with external scans.) Typically, discovery of assets happens just once per quarter and uses a patchwork of scripts and programs the pentesters have put together to find some of the infrastructure that is potentially vulnerable. Their methods are rarely comprehensive, however, and regularly fail to find all vulnerable infrastructure of a given organization. These issues, coupled with a rapid shift to digital transformation and inherently ephemeral systems on the scale of hours and minutes, results in enterprise security teams significantly lagging behind attackers in both having an inventory of assets and knowing if those assets are vulnerable.

The Attack Surface Scanning Industry

The ease of scanning gave rise to a cottage industry of analysts and criminals who scan for vulnerabilities and infrastructure—especially in the age of ransomware. In the past five years, attackers have perfected techniques that scale at speed. To identify new targets, scanners just need a target—usually a list of IPs or a specific vulnerability. For attackers not using online tools, many scanners are open source. With a simple trip to GitHub, attackers need merely download a scanner, deploy onto infrastructure, and off they go.

Mean Time to Inventory Defined

MTTI measures the time required for organizations to perform a full external asset inventory that assigns ownership to drive classification and protection based on value. MTTI becomes especially critical during CVE announcements. There is an expected increase in attackers scanning for vulnerable services immediately after the release of a new critical CVE. Furthermore, there will be an additional increase once proof-of-concept code is released to the general public. Unfortunately, both of these cycles will usually occur before most organizations have completed their own first pass of an inventory scan. Today, organizations use metrics to gauge cybersecurity effectiveness, and typical yardsticks often include dwell time, mean time to detect, or mean time to respond. However, these measurements, inherently reactive in nature, focus only on known assets.

High-Priority Security Issues

Global enterprises average one serious security issue every 12 hours.

Our scans identified an average of two high-priority internet-exposed security issues each day at the enterprises we studied. They included insecure remote access (RDP, Telnet, SNMP, VNC, etc.), database servers, and exposures to zero-day vulnerabilities in products such as Microsoft Exchange Server and F5 load balancers. Experiencing one issue every 12 hours highlights the ephemeral nature of today's IT infrastructure, where not only infrastructure changes but so does the vulnerability footprint. Tracking an ever-changing landscape is an impossible task for humans and requires an automated approach.

This rate of security issues arising is driven by the fact that asset leak is inevitable. In many organizations, focusing on the big picture means 'little' things get forgotten, and the more unknown and obscure IT will slip through.

The speed of business today often demands a mission-first, security-second attitude, which makes it all but inevitable that even the most technically proficient companies will have vulnerable portions of their internet attack surface. Being too forward-leaning and operationalizing development environments can also introduce even bigger vulnerabilities.

On the flip side, the ephemeral nature of today's infrastructure also means issues disappear. Issues go away because cloud instances get turned off, infrastructure becomes disconnected from the public internet because of an updated firewall rule, legacy systems get reconfigured, or the security issue is remediated. Despite these possibilities, IT should not be sitting back and waiting for issues to disappear.

Attack Surface Rankings

RDP is the most common security issue attackers can find on the global enterprise attack surface.

Remote Desktop Protocol was the most common security issue we found among the global enterprises we studied, representing 32% of overall security issues (see figure 1). RDP's top spot is particularly worrisome because it's a top gateway for ransomware. Our analysis found constant RDP scanning for port 3389—reserved for RDP. The Unit 42 incident response team has seen in its investigations that scanning is followed by brute forcing credentials or basic credential cracking tools. Worse, in the remote work environment, connecting from a personal device means it's out of the security team's control. This gap means most companies don't have the right controls, and without visibility, attackers have the luxury of time to find and exploit RDP.

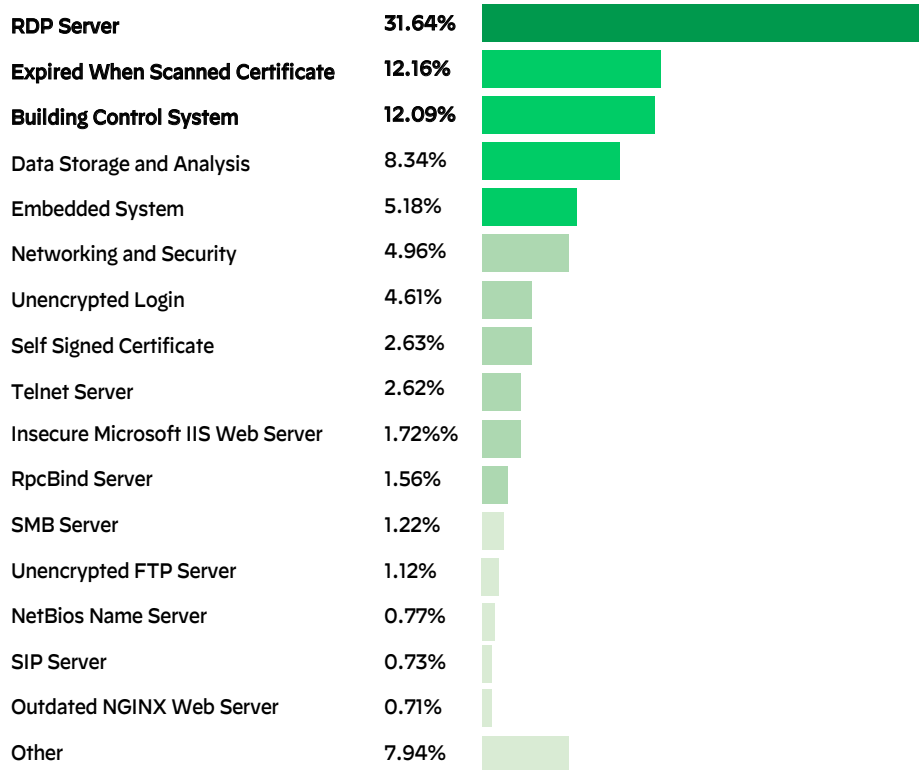


Figure 1: Top security issues based on prevalence

Many of the security issues we found and included in our top 10 Attack Surface Management list amount to basic hygiene. This means attackers don't need to be clever. They just have to find the issues. No matter how sophisticated you make your applications, with basic vulnerabilities, enterprises are still at risk. Further, exposure of applications that should never be public facing is a problem. Applications like Jenkins, Grafana, and even Tableau can expose proprietary information and represent an attractive pivot point for attackers. As a result, they are generally not intended to be exposed to the public internet.

Cloud Security Risks

For global enterprises, cloud-based security issues comprises 79% of observed exposures compared to 21% for on-prem, highlighting the risk incurred from digital transformation

Organizations are moving to the cloud, and it is too easy for employees to spin up a cloud instance outside of normal IT processes. Across cloud infrastructure providers like Amazon

Web Services (AWS), Microsoft Azure, Google Cloud, Oracle, Rackspace, and more, our findings show that organizations experience nearly four times the total number of critical issues for cloud infrastructure than they do for on-premises environments (see figure 2).¹ Several factors contribute to cloud's significant risk.

- **Cloud Is Harder to Manage Because It's Easy to Deploy**
Employees can set up in any cloud provider, oblivious to corporate policies that state otherwise. For example, DevOps often enables developer self-service to stand up cloud infrastructure. Further, most enterprises don't have a clear source of truth on which providers are being used and whether those providers safeguard data to an extent that meets internal legal and security demands. The COVID-19 pandemic accelerated the growth of cloud, which, most likely, won't revert to old-school IT anytime soon. In fact, cloud spending rose 37% to \$29 billion during the first quarter of 2020. According to Gartner, cloud spending rose to 19% in 2020, even as IT spending fell 8%.²

1. Cloud issues are modeled by domain, so figures in this report are not overinflated by IPs changing regularly.
2. Katie Costello and Meghan Rimol, "Gartner Says Global IT Spending to Decline 8% in 2020 Due to Impact of COVID-19," Gartner, May 13, 2020, <https://www.gartner.com/en/newsroom/press-releases/2020-05-13-gartner-says-global-it-spending-to-decline-8-percent-in-2020-due-to-impact-of-covid19>.

The Attack Surface Management Top 10

The ASM Top 10 serves as a guide to help security teams identify attack surface exposures ranked on potential risk posed to enterprises. The list is rank ordered based on two main tenets. First, certain things should never be on the internet because, if exploited, attackers can move laterally or expose sensitive business processes and data. These include:

- Inherently bad protocols
- Exposures that relate to the control plane of your network (e.g., admin portals, VPNs)

The second tenet: Older assets that were secure in the past may have since become vulnerable.

The ASM Top 10 includes:

1. Remote access services (e.g., RDP, VNC TeamViewer)
2. Insecure file sharing/exchange services (e.g., SMB, NetBIOS)
3. Unpatched systems vulnerable to public exploit and end-of-life (EOL) systems
4. IT admin system portals
5. Sensitive business operation applications (e.g., Jenkins, Grafana, Tableau)
6. Unencrypted logins and text protocols (e.g., Telnet, SMTP, FTP)
7. Directly exposed Internet of Things (IoT) devices
8. Weak and insecure/deprecated crypto
9. Exposed development infrastructure
10. Insecure or abandoned marketing portals (which tend to run on Adobe Flash)

To learn more about the ASM Top 10, go to www.asmtop10.com.

• Cloud Constantly Changes

Previous Cortex Xpanse research shows that, on average, companies add 3.5 new publicly accessible cloud services per day—nearly 1,300 per year.³ At the high end, one customer added 693 new publicly accessible cloud services in a single day. Amid these rapid growth rates, many cloud providers have default settings configured when connected to the public internet, and ease of use keeps it that way. Although cloud providers are getting better at security, the problem persists. Poorly provisioned cloud and on-premise might both be exposed on the internet, and enterprises are at risk in either case. However, a well-provisioned cloud is by definition on the internet, whereas well-provisioned on-premise might indeed be air-gapped.

• CSP Security May Not Suffice

Relying on only what the baked-in security cloud service providers (CSPs) include can be insufficient. CSP tooling can provide basic vulnerability scanning and cloud security posture management capabilities, but it's just the basics. For enterprises, it doesn't provide the visibility or full-stack security that you would need to be cloud native.

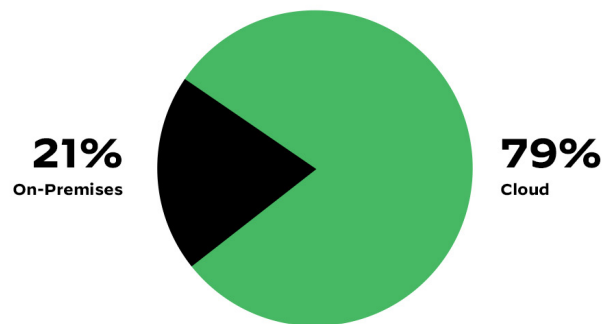


Figure 2: Observed cloud-based vs. on-prem security exposures for global enterprises

2: Conclusion and Recommendations

The Internet Is Small

In the past, it would take months for a team of researchers to scan the entire internet and determine its structure and topology. Moreover, internet scanning was the realm of nation-state actors. Worse, as recently as 2012, identifying every device on the internet in a reasonable amount of time required illegal methods (like building your own botnet, for example).

3. "Quantifying the Velocity of Digital Transformation," Expanse, December 1, 2020, <https://expance.co/blog/quantifying-the-velocity-of-digital-transformation/>.

Everything changed in 2013. New scanning algorithms allowed global internet scanning at a rate of 1,500 times faster than previous methods. While it used to take weeks or months to scan the global internet, it takes less than 45 minutes today to communicate with every public-facing IP address in IPv4 space (4.3 billion IPs) on one port-protocol pair. To put this in perspective, this is about the time it takes to watch one episode of [The Expanse](#).

What happened in the past few years to make global-scale internet scanning feasible?

Computing and Infrastructure Costs Dropped

Previously, scanning the internet required a room full of supercomputers working 24/7—restricted to well-funded nation-state threat actors. Then, the cloud came, lowering the barrier to entry for scanning as the ease of spinning up infrastructure became negligible. The risk of conducting malicious scanning activity dropped drastically in a few years, going from the loss of significant leased infrastructure to a “cease and desist” from a CSP that was disposable to the attacker’s end goals.

Algorithms Randomized the IP Addresses Scanned

Scanning IP addresses sequentially can look like a distributed denial of service attack because traditional on-premises IP addresses are issued to individuals and organizations in blocks. New scanning techniques used an algorithm that randomized the order of IPs scanned, while also allowing scanners to ignore who was scanned. Additionally, this approach allowed for the use of distributed scanning architectures and reduced computing resources needed since tracking the order of scanning targets was no longer necessary.

Packet Lag Parallelized the Handshake Process

When two devices on the internet communicate, there is always some delay as packets of information travel through the system. At global scale, this lag adds up fast. New ways to keep track of connections allowed scanners to send more packets at once, effectively bundling the delay time into one brief waiting period. This made bulk scanning of millions of devices almost as fast as making a single connection to just one of those devices.

Given these advancements, enterprises must think harder about the attack surface. In particular, digital transformation has turned enterprises inside out, creating numerous and frequently insecure backdoors into their network in the form of abandoned, rogue, or misconfigured assets. The advancements in scanning technology made these backdoors easier to find and fundamentally changed how we think about the internet and gathering information on it—especially for hackers, who, by definition, are innovators and early adopters.

As the Internet Shrinks, the Enterprise Attack Surface Grows

It is well established that digital transformation had been underway for some time only to be accelerated by the COVID-19 pandemic. Not as well known, however, has been the impact on cybersecurity.

In the past, when the majority of infrastructure lived on-premises, security and IT had a good (though not great) idea of asset inventories with established processes for change control. Today, however, a growing group of assets outside core subnets and data centers is ephemeral, leading to network sprawl and making it difficult for security operations centers to understand how well their networks are configured.

For example, while engineers needing an internet-facing development server would have previously applied for a corporate IP space and undergone the corresponding security conversations, any engineering lead today with a corporate credit card can spin up an AWS server. Worse yet, the cost of that server is trivial enough that it could easily be forgotten as a recurring charge on the corporate account—and a persistent ever-degrading vulnerability on the corporate attack surface. The result? Security can’t stay on top of an ever-changing enterprise as marketing, finance, and other departments bypass change control.

On top of internal drivers, external factors such as mergers and acquisitions (M&A), supply chain and IoT also bypass change control. M&A activity, for example, has [reached all-time highs](#) during the COVID-19 pandemic. While security is doing a good job, the task of protecting a static IP address is now obsolete as infrastructure changes so quickly.

Meanwhile, old-school security change control processes continue with monthly vulnerability reviews or quarterly pen-testing. As noted previously, attackers scan hourly; and for significant CVEs, attackers ramp up in just minutes. Concerns about digital transformation introducing security gaps not only proved grounded but also understated the impact.

In reality, digital transformation has realigned the risk equilibrium in the attacker’s favor. Most tools in IT and security’s arsenal—namely asset and vulnerability management—focus on evaluation but not discovery. In other words, these tools manage known assets while remaining blind to unknown ones. Worse yet, the common methods of discovering unknown assets—such as pen-testing—take place on a quarterly basis (see figure 3).

These programs should start with the basics:

- **Global internet visibility:** Implement a system of record to track every asset, system, and service you own that is on the public internet, including across all major CSPs and dynamically leased (commercial and residential) ISP space using comprehensive indexing, spanning common and often misconfigured port/protocols (i.e., not limited to the old perspective of only tracking HTTP and HTTPS websites). M&A, supply chain and IoT also bypass change control. M&A activity, for example, has reached all-time highs during the COVID-19 pandemic. While security is doing a good job, the task of protecting a static IP address is now obsolete as infrastructure changes so quickly.
- **In-depth attribution:** Detect systems and services belonging to your organization using a full protocol handshake to verify details about a specific service running at a given IP address. By fusing this information with a number of public and proprietary datasets, match the full and correct set of internet-facing systems and services back to a specific organization.

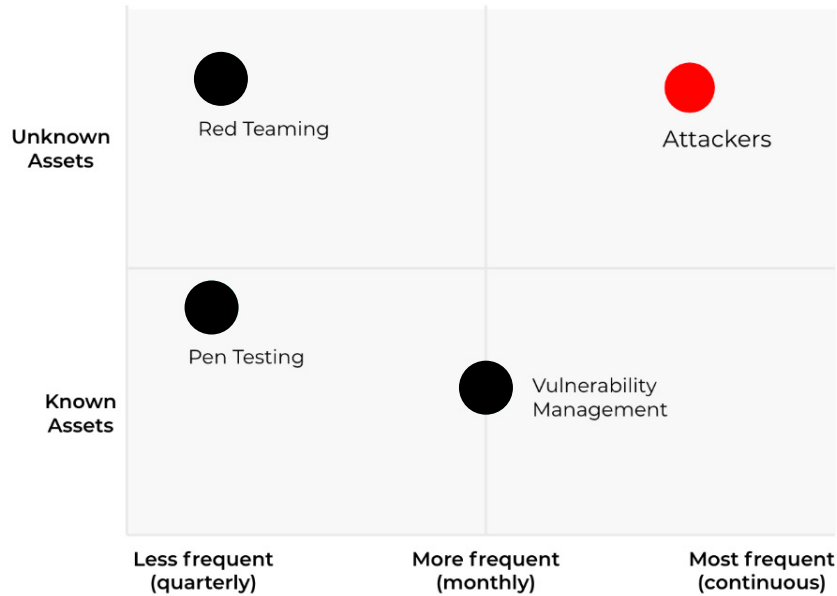


Figure 3: IT security toolset for attack surface evaluation

Methodology

Using the externally available attack surface from global enterprises, Xpanse researchers examined and interpreted data to help defenders understand the attack surface in order to:

- Quantify and remediate externally facing vulnerabilities.
- Provide security teams with attack surface benchmark metrics.
- Optimize threat modeling.
- Convey the threat landscape to technical and nontechnical audiences.
- Deploy proactive security measures.

Xpanse operates a proprietary platform that continuously collects more than one petabyte per day of information related to all systems on the public internet to ascertain how attackers view potential targets. We fuse this information to discover cybersecurity risks present on the networks of the world’s largest and most complex organizations, which no one else can find. Our technology helps our customers see the world through the eyes of highly sophisticated attackers.

For this report, we looked at the attack surface and threat data coming from 50 global enterprises, including a subset of the Fortune 500, covering around 50 million IP addresses from Q1 2021 (January 2021 – March 2021) and representing 1% of total, global IPv4 space.

About Cortex Xpanse

Xpanse, a global internet collection and attribution platform, empowers CISOs to continuously discover, evaluate, and mitigate their external attack surface. Today, Xpanse customers collectively represent 12% of the overall IPv4 internet and include leading Fortune 500 companies as well as both US government organizations and military branches.