paloalto NETWORKS | CORTEX XPANSE BY PALO ALTO NETWORKS

LESSONS IN ATTACK SURFACE MANAGEMENT BASED ON OBSERVABLE DATA

2022 CORTEX XPANSE ATTACK SURFACE THREAT REPORT

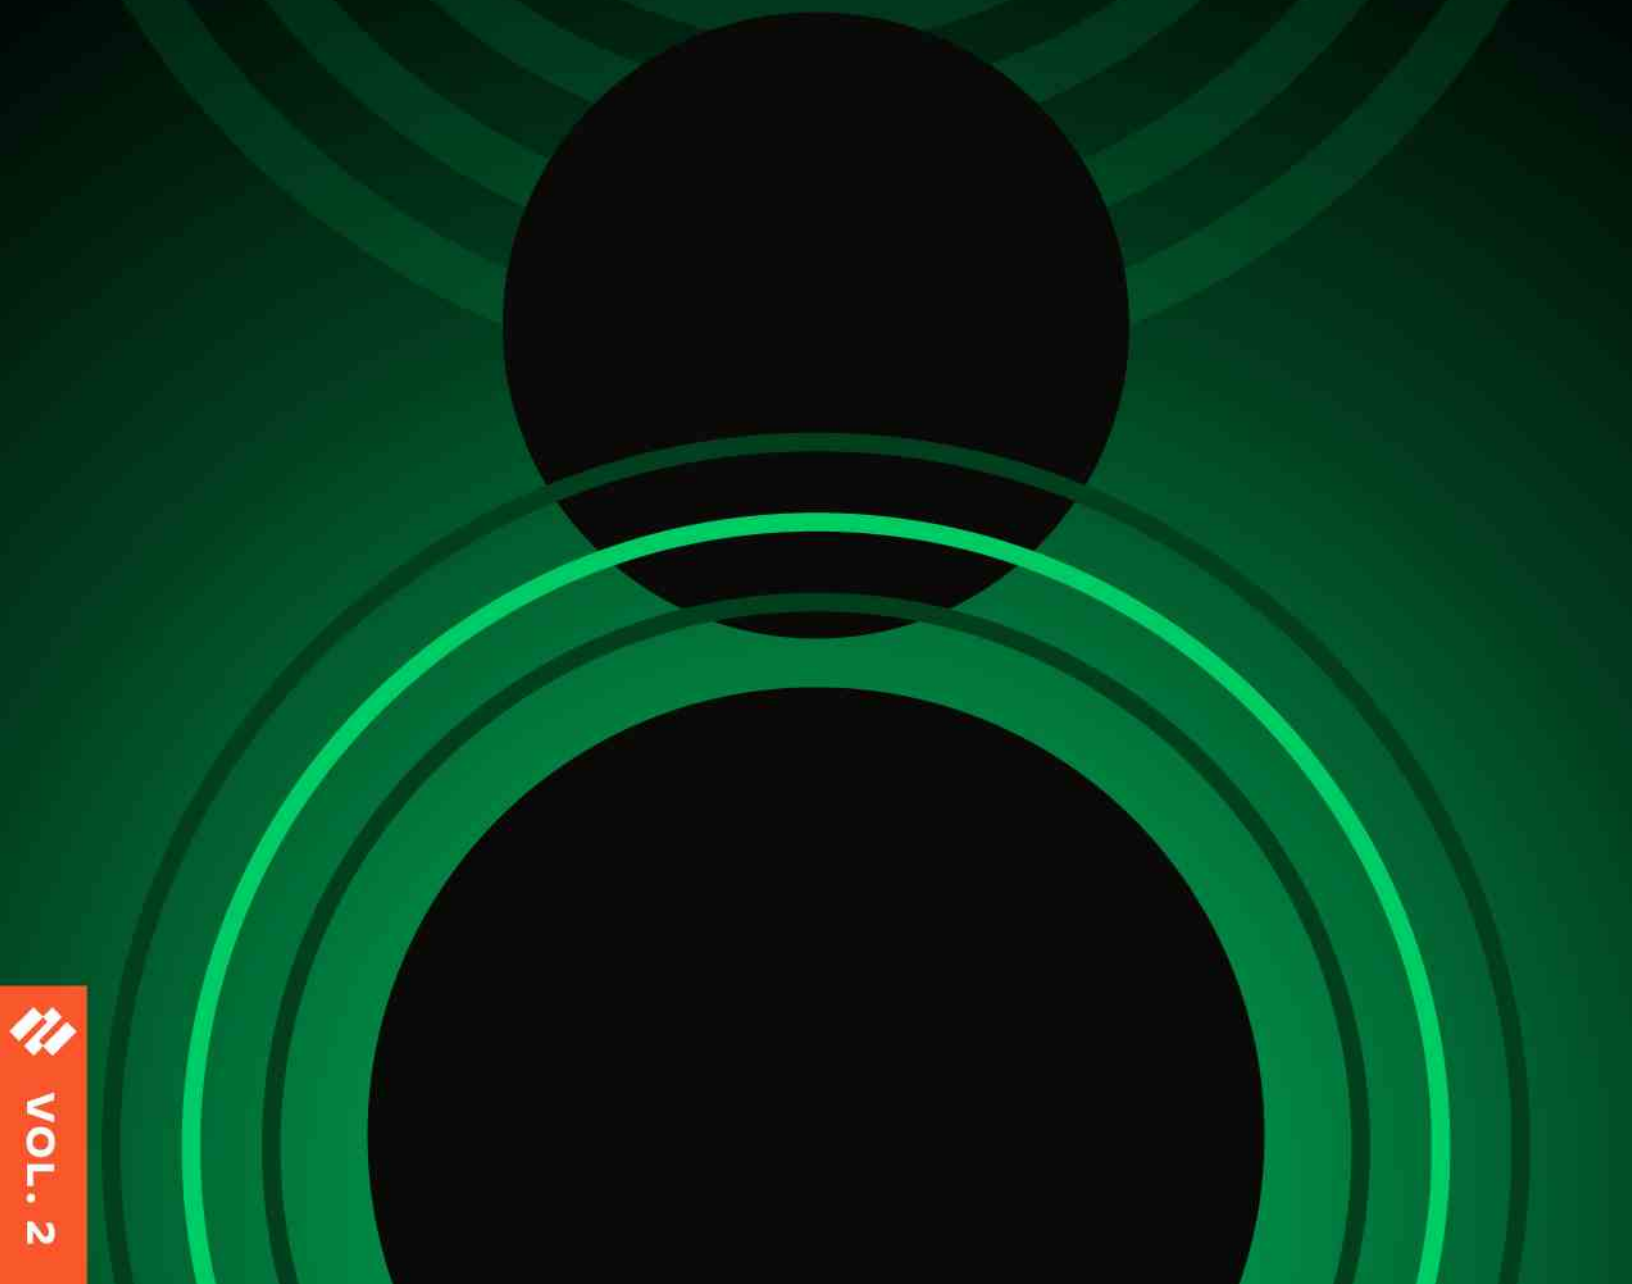

VOL. 2

Executive Summary

Introduction

Do you stay up at night wondering about the next zero-day on the horizon? Or, do you wake up screaming with the thought that someone in your organization might have created a new cloud asset outside of your security processes and not bothered with something simple, like disabling Remote Desktop Protocol (RDP)?

As a seasoned security professional, you already know that zero-days get the headlines, but the real problems always come from the dozens of decisions a day made inside your organization. Just one mistake or lapse in security protocols is all it takes to create a crack in defenses. It's the low-hanging fruit that attackers count on because it has become easy and inexpensive for attackers to find any vulnerabilities, exposures, or other unknown open doors and decide what path will likely offer the least resistance in a cyberattack.

Even lower-skilled attackers can put together scanning infrastructure to perform a rough scan of the internet to uncover assets ripe for compromise. Some may even take a shot at breaching that exposure, but far more enterprising attackers sell this scan data on the dark web to bidders who can then launch more sophisticated attacks.

For example, RDP instances (services to remotely log in to a device for work) sell for anywhere between $3–$10 to deploy ransomware on the unsuspecting target's network.[1] Luckily, attackers aren't the only ones able to scan the entire internet and discover exposures—Cortex® Xpanse™ can too. To help organizations fight fire with fire, the Xpanse research team studied the public-facing internet attack surface of some of the world's largest organizations.

From March to September, we monitored scans of 50 million IP addresses—over 1% of the entire internet—associated with 100+ global enterprises to understand how quickly adversaries can identify vulnerable systems for fast exploitation.

What follows are key findings on the state of the global attack surface based on observed scan data, not self-reported surveys, to help organizations fight back.

Here is a summary of our key findings:

- Cloud continues to be a security nightmare.
- Low-hanging fruit continues to hang.
- End-of-life software means end-of-life for your security.
- The unmanaged attack surface is growing.
- Issues are persistent, complex, and unique.

Key Findings

In the pantheon of security issues that exist, there is no shortage of issues and issue types, but that doesn't mean we can't group things together into meaningful categories. For example, if we look at the basics of cybersecurity, there are some topics most practitioners would agree are persistent more because of the frequency of occurrence or because they are non-starters when discussing security hygiene and posture.

In the case of frequency, our data continues to show that overall, cloud issues continue to be a problem, and so do exposed RDP servers. When looking at fundamental issues of poor security, we discovered a troubling amount of exposures in administrative login pages as well as in internet-facing end-of-life (EOL) software.

According to our data, risks and exposures are persistent because modern attack surfaces are inherently dynamic, constantly shifting, moving, and growing. All too often, this means without the right visibility and processes, more threats arise as you remediate current issues.

1. Brian Krebs, "Hacked Via RDP: Really Dumb Passwords," Krebs on Security, December 13, 2013, https://krebsonsecurity.com/2013/12/hacked-via-rdp-really-dumb-passwords/.

Takeaway #1: Cloud Continues to be a Security Nightmare

For the second year in a row, 80% of all observed issues were present on cloud infrastructures. This should come as no surprise given the aggressive move to the cloud that was accelerated by the pandemic came with risks.

The sheer volume of issues found on the cloud as opposed to on-premises indicate that deploying to the cloud is easy to do but difficult to secure. There are myriad reasons why security might be lacking in the cloud, from cloud assets being created outside of security controls, insecure defaults, or just the sheer amount of assets in the cloud can overwhelm under-resourced security teams. The cloud is the modern attack surface in a microcosm: moving and changing at such a rapid pace that traditional security practices often can't keep up.

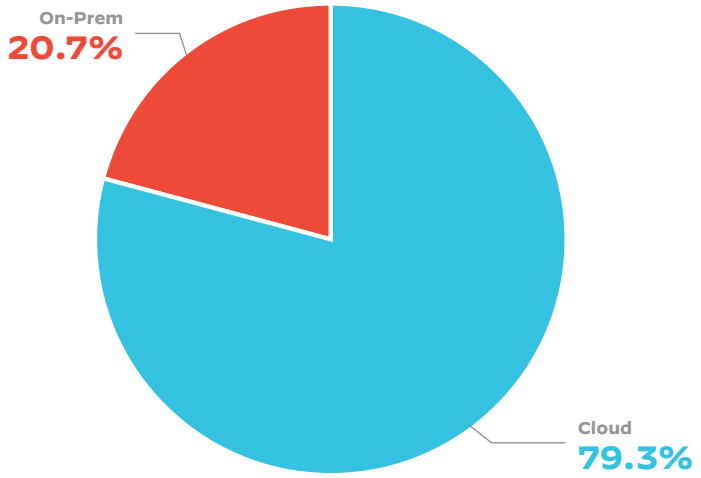

Figure 1: Nearly 80% of new issues discovered are in the cloud, highlighting the need for more comprehensive security

We took a deeper look into the cloud vs. on-premises split between the four recent issue types which have had observed exploits in the wild and were mentioned in several federal advisories (more details in Takeaway #4). We observed that insecure Apache issues are almost exclusively seen in the cloud (97%). This is similar to insecure Microsoft List Server and BIG-IP TMUI vulnerabilities as well. However, we see that this trend is inverted when looking at insecure Microsoft Exchange Servers, which might imply that these deployments are still primarily on-premises. Organizations without a clear view of where their deployments are will be limited in their ability to respond to new CVEs that affect their network.

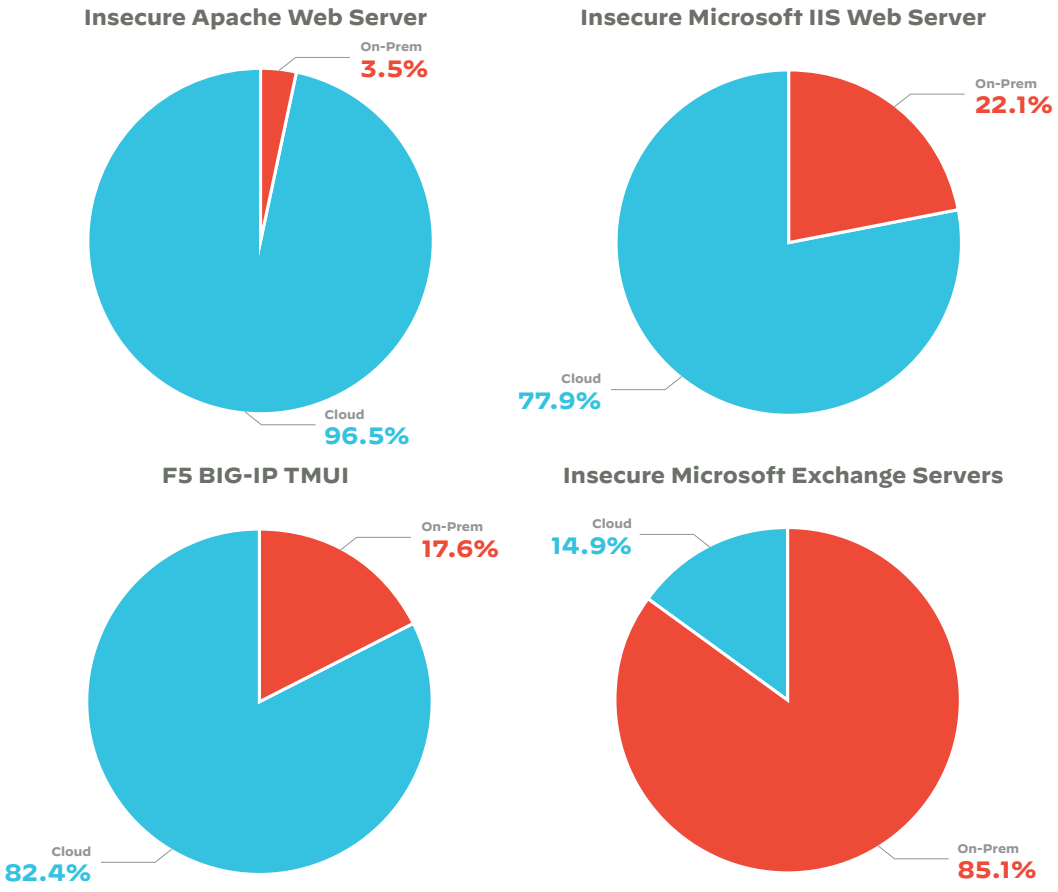

Figure 2: Whether issues are in the cloud can depend on the type of server being set up

Keeping in mind that the cloud is where 80% of issues occur lends valuable context to the rest of our findings.

Takeaway #2: Low-Hanging Fruit Continues to Hang

If an attacker is looking for an open door, they often don't need to look very far. On the internet, leaving an administrative login page exposed publicly is akin to having a neon sign inviting attackers to knock.

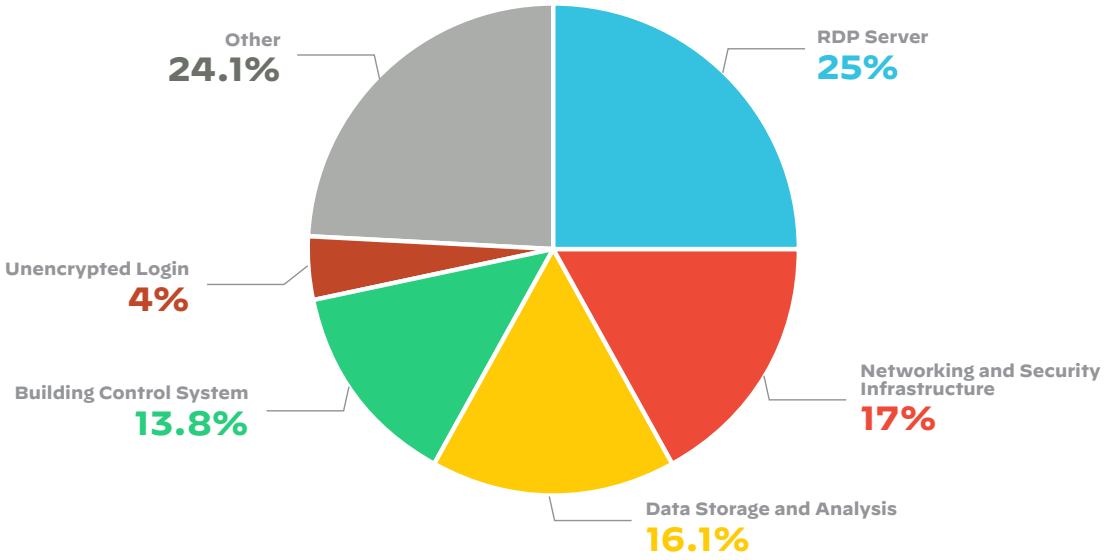

Figure 3: Distribution of risks across the global attack surface

Nearly one out of every four issues we found on the attack surface was related to an exposed RDP server, but even looking at the second most common issue, networking and security, the end result was often an exposed system administration login portal.

In the case of RDP, breaching that exposed login gives an attacker rights equivalent to logging in with legitimate user credentials. The networking and security-related issues revealed IT admin portals, which would give an attacker even more privileges and access to an organization's core networking infrastructure. Beyond that, Xpanse research uncovered over 700 unencrypted login pages for several IT services that were unencrypted and publicly exposed to the internet.

Unencrypted logins make it dramatically easy for attackers to steal credentials, and so organizations should identify these and shut them down. If those same credentials can be used on another exposed portal, that just makes the attacker's job all the easier.

However, any exposed portal leaves an organization open to simple brute-force attacks if multi-factor authentication isn't used, and in the case of unpatched systems, they can be accessed by exploiting known vulnerabilities and worse, these exposures are preventable by placing these login portals behind a VPN or firewall.

These exposures are not only prevalent but costly as well. RDP has become the ransomware deployment protocol, and the average cost of a successful ransomware attack in 2021 was $312,493, according to the latest 2022 Unit 42 Ransomware Threat Report. In light of the surge in ransomware attacks in recent years, C-suites around the world are deploying active ransomware prevention programs, and many boards of directors are asking for attack surface management plans to ensure no unknown assets become vectors for attack.

Close to 3,000 database storage and analytics systems were also frequently left exposed to the public internet. These systems can contain critical customer data or intellectual property and were never meant to be accessed from the public internet, but these were still showing up likely on account of accidental misconfigurations.

Our research also uncovered over 2,500 critical building control systems (BCS) accessible from the public internet, which indicates that in a remote-first world, organizations should not only be concerned about their IT assets but also their operational technology (OT) assets. These assets are frequently operated and managed by facilities or office divisions, and are thus not necessarily tracked or monitored by IT security systems.

Takeaway #3: End-of-Life Software Means End-of-Life for Your Security

If exposed remote access or remote login protocols are like neon signs inviting attackers to knock, end-of-life (EOL) software is like leaving your valuables in a straw house in autumn. It may survive the rainy season if it was well-built, but winter is coming, and no one is around to patch all of the inevitable holes.

During the course of our observation, we saw that several organizations across industries were running EOL versions of software. Across the following applications below, we can see, on average, around 30% of organizations were running EOL software versions.

• Apache Web Server	• ~32% running EOL versions
• Microsoft Exchange Server	• 29% running EOL/unsupported versions

Additionally, organizations are still running unpatched versions of software with active observed exploits. Despite patches being available for between four and five months at the time of detection, Xpanse research discovered the following based on the application's self-reported version:

- 11,511 instances of Apache Web Servers running unsecured on the public internet were vulnerable to CVE-2021-41773, CVE-2021-42013.
- 2,700 instances of vulnerable to CVE-2021-26084 (Atlassian Confluence).
- Based on the self-reported application versions, 74% of instances of Zoho ManageEngine ServiceDesk Plus software (3,400 total) were vulnerable to two critical CVEs (CVE-2021-44077, CVE-2021-44526), one of which was actively being exploited in the wild.

Exploits like these can allow malicious actors to gain access to a victim's network, escalate privileges, move laterally, and execute remote code.

Business Impact Takeaway

Attackers don't need to preselect victims, because it is all too easy to find exploitable weaknesses, and nothing presents more weaknesses than end-of-life software. There is no reason why any asset running end-of-life software should ever be internet-facing. If an asset cannot be updated to secure versions of software, it should be isolated or decommissioned altogether.

The opportunist attacker can find potential victims by simply scanning the internet for assets or services exposed to accidents or misconfigurations. Organizations need to automatically discover end-of-life software, misconfigurations, and unknown assets to identify all non-zero-day vulnerabilities on their attack surface.

Takeaway #4: The Unmanaged Attack Surface Is Growing

The cloud is a magnet for security issues. Basic risks and exposures continue to plague attack surfaces. And, on top of all of that, unmanaged attack surfaces are growing. As noted before, risks and exposures are persistent because modern attack surfaces are inherently dynamic, constantly shifting, moving, and growing. Unfortunately, this means that as attack surfaces grow, so too does the number of unmanaged assets on those surfaces.

To take a deeper look at how the attack surface has evolved, we focused on the following four active vulnerabilities which have had observed exploits in the wild and were mentioned in several federal advisories:

- Insecure Apache Web Server
- F5 BIG-IP TMUI
- Insecure Microsoft Exchange Server
- Insecure Microsoft IIS Web Server

During the course of our analysis of these four issue types, we observed them on an unmanaged attack surface over the course of a month.

These four issue types we've used to perform deeper analysis only constitute a small sliver (<1%) of the overall issue types seen on an attack surface. A typical organization has 300-400 issue types of varying severity on their external attack surface. As a consequence, when you consider the entire range of issue types commonly seen on the attack surface, the potential exposures and risks are much more vast than what we are presenting here.

We can see that, across industries, even if organizations worked to remediate active issues in a month, newer issues kept cropping up throughout the month. One could make the claim that these organizations were never secure and remained vulnerable throughout the month as the unmanaged attack surface continues to grow and compounds security issues that an organization faces.

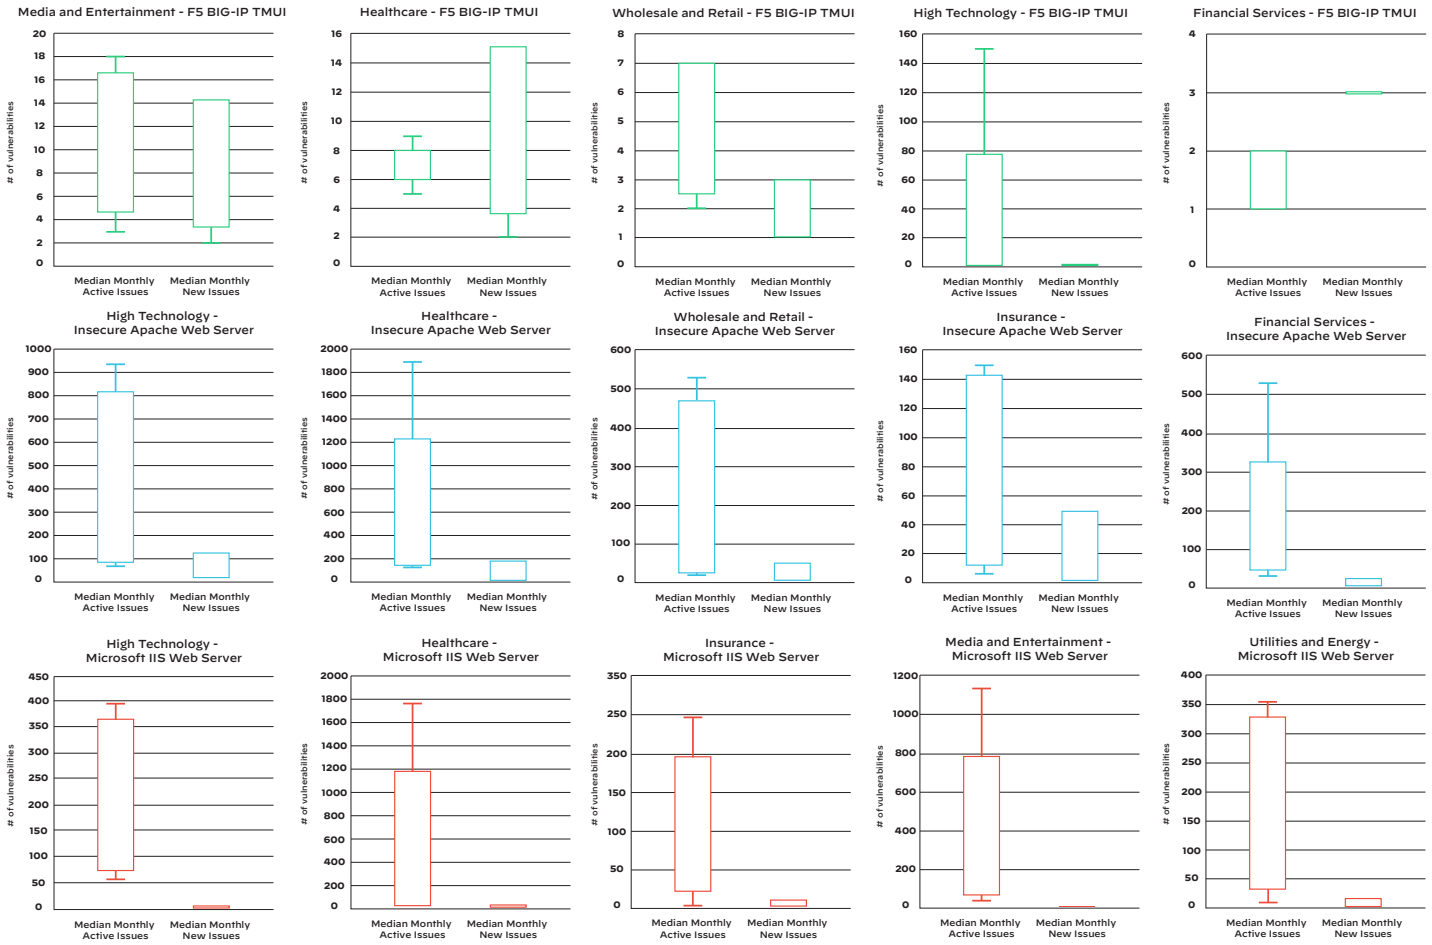

Figure 4: Median active and new issues per month per company in an industry

Business Impact Takeaway

The unmanaged attack surface goes from bad to worse. Organizations need to continuously monitor their exposures to the public internet and automate both discovery and remediation, or else they will always remain vulnerable to opportunistic attackers.

These four issue types chosen for this assessment are emblematic of the rate of change of unmanaged attack surfaces, but they are only the tip of the iceberg. Exposures on an unmanaged attack surface persistently grow, making organizations increasingly vulnerable over time.

How to Fix It

Organizations not only have hundreds of active issues on their attack surface but also continuously add new issues over the course of time. The modern unmanaged attack surface is fragmented and growing steadily. Organizations without a single source of truth for their asset inventory are only going to compound their problems every day as they cannot secure what they don't know exists.

Takeaway #5: Persistent, Complex, but Unique

Common security wisdom often results in something like victim-blaming, where organizations are asked to consider why you might be the target of an attacker. Adversaries may want your data if you're in financial services, or they might know you're more likely to give in to demands because to risk downtime would extoll a cost far too high, such as in healthcare.

Knowing why you might be a target should help in terms of isolating data and systems that don't need to be publicly accessible, but when it comes to internet-connected devices, the more important task is to identify exposures and vulnerabilities. So, from an attacker's perspective, organizations are often more similar than different.

Xpanse research showed similar types of issues across industries, but what issues were most common varied dramatically. So, organizations often had exposures falling into some combination of RDP server exposures, networking and security, or data storage and analytics, but the specific details regarding assets, types of exposure, and the reasons why they might be targeted by an attacker make situations unique.

Real-World Disruption

Looking at industry verticals in broad terms, there are two general reasons why attackers might target certain organizations: operational disruption or stealing high-value data. The first category includes industries like utilities and energy, healthcare, transportation and logistics, and (sometimes) wholesale and retail. The aim of attackers here is to disrupt or threaten to disrupt, the business operations of the organizations, either under political motives or because attackers hope the ransom demanded is a substantially smaller loss for the organization to incur than the disruption caused or threatened.

Utilities and Energy

Utilities and energy companies can almost be considered a class of their own because the issues seen are vastly different than most other industries, but also because historically industrial infrastructure has been a key target for attacks with political motives.

The most famous example of this is Stuxnet, a worm built by the US and Israeli governments for the purpose of disrupting the Iran, and reportedly, the North Korean nuclear programs.

The biggest issue type observed was related to exposed admin panels of IT infrastructure. Several critical OT systems like their building control systems which are not supposed to be on the public internet were also found during the course of our analysis.

The utilities and energy industry needs to focus on securing their key IT and OT infrastructures. Critical energy and utilities are key targets for nation-state attackers and, as a consequence, having a BCS system admin portal that is discoverable on the public internet makes it significantly easy for a bad actor to cause significant business disruptions by controlling building systems, including fire alarms, elevators, and fire suppression systems.

For example, the attack on Colonial Pipeline in April 2021 was due to a single compromised user account and VPN access that didn't have multi-factor authentication enabled, but it led to a ransomware attack and the entire pipeline being shut down for five days.

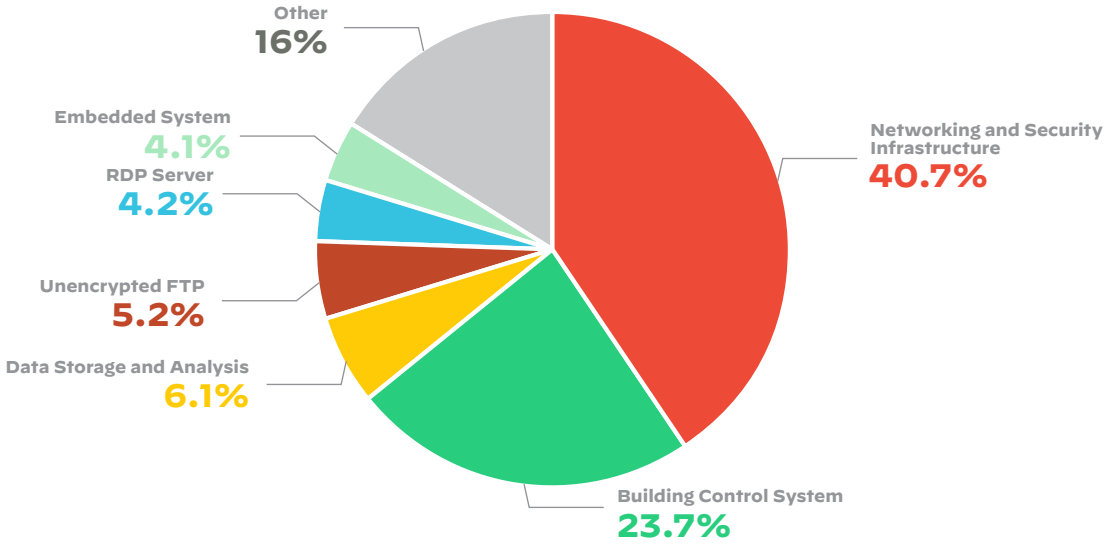

Figure 5: Distribution of risks across the utilities and energy industry attack surface

Healthcare

A HIPAA compliant RDP server allows healthcare professionals to work remotely and still have access to the same information they could view and update if they were working at a hospital. Remote desktop access allows healthcare professionals to work efficiently from home and while traveling.

According to our analysis, RDP risks are the single biggest avenues for compromise in the healthcare industry, and this leaves them more vulnerable than other industries to ransomware attacks. This is especially dangerous because disruption in healthcare facilities could lead to loss of life.

According to the 2021 ransomware report, RDP servers are now the preferred attack vector for ransomware gangs as it is easier to discover an accidentally exposed RDP server and brute force their way in.

The healthcare industry has been in the crosshairs of ransomware attacks for years now, and because the potential damage is so great, many hospitals have to pay the ransom rather than risk impacting operations, which makes them attractive ransomware targets.

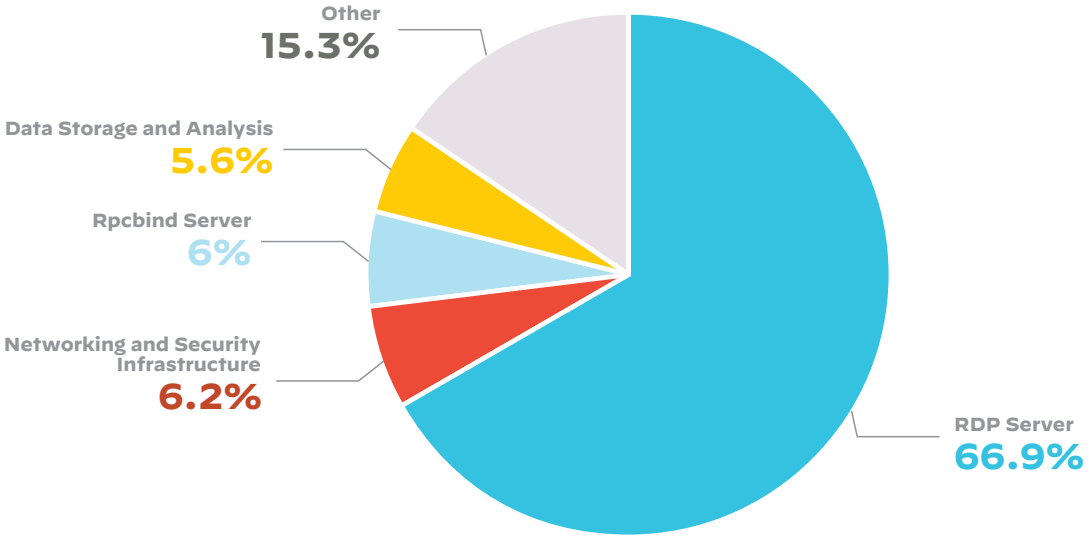

Figure 6: Distribution of risks across the healthcare industry attack surface

Transportation & Logistics

RDP risks are again the major issues in the transportation and logistics industry. Organizations should work with their supply chain logistics partners to identify and remediate their critical RDP exposures before there is a significant incident.

In 2017, Danish shipping company Maersk was hit by a ransomware attack that disrupted operations for two weeks and reportedly cost the company approximately $300 million.

Since the beginning of the pandemic, global supply chains have faced significant turbulence and as nations emerge from the pandemic, the demand for goods is extremely high. As a consequence, ransomware actors are now targeting organizations in the global transport and logistics industries.

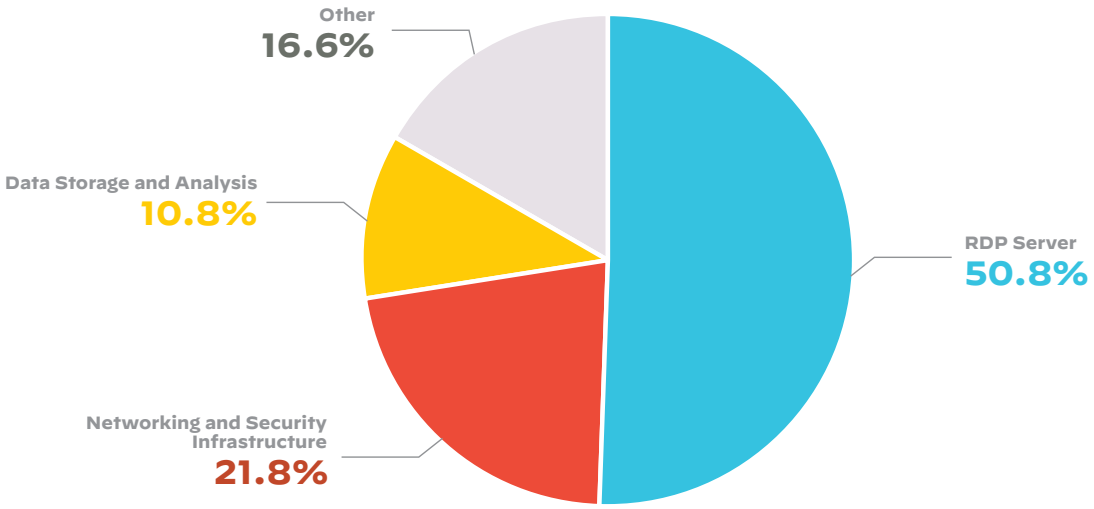

Figure 7: Distribution of risks across the transportation and logistics industry attack surface

Wholesale and Retail

Wholesale and retail lives somewhat on the border between targets attacked in order to create real-world disruption and those with high-value data that could be stolen. The disruption that could be caused by attacking retail is extremely high, but if it is point-of-sale systems being targeted, that puts credit card data of customers at risk.

Ransomware gangs are more likely to target wholesale and retail organizations as they need to function at all times and, therefore, will quickly pay the ransom to continue operations. Wholesale and retail organizations should be extremely concerned about this. We discovered that RDP servers account for more than 60% of issues found on the wholesale and retail industry attack surface. It is the single biggest issue type in this industry.

One of the biggest cyberattacks in history befell Target in 2013 when the credit card information of 40 million customers was stolen, and other information on 70 million customers was exposed. Target reportedly spent more than $200 million in legal fees and ended up paying states $18.5 million in damages.

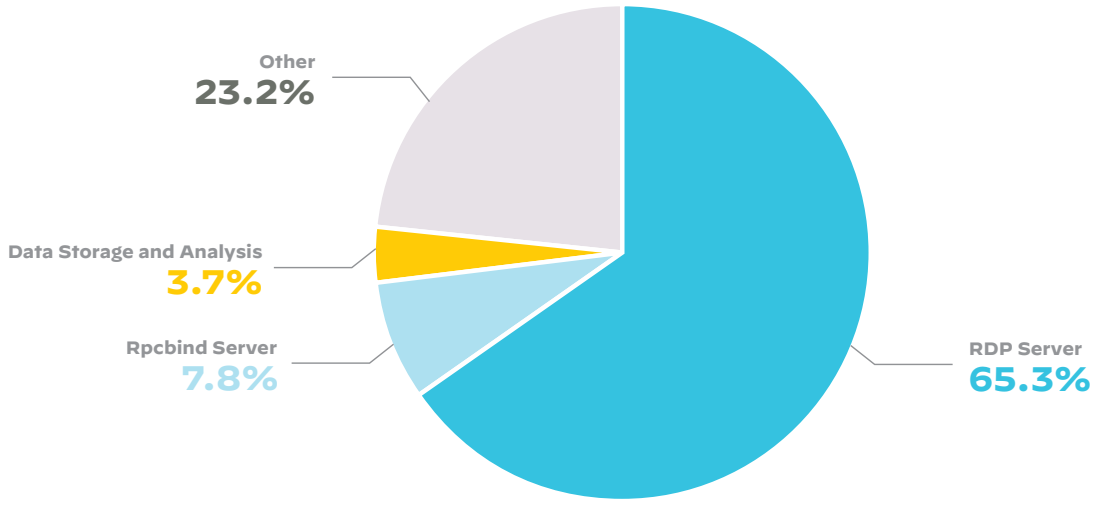

Figure 8: Distribution of risks across the wholesale and retail industry attack surface

Valuable Data

When real-world disruption isn't enough, attackers tend to target high-value data. This could mean financial data that could then be resold to identity thieves or it could mean data that is valuable to an organization, like intellectual property.

Financial Services

IT regulations and compliance directives are perhaps the strongest in the financial services industry. One result of this is that financial industries tend to be early adopters of cybersecurity practices and products, but malicious actors still target these organizations due to the critical data they hold.

More than 80% of all issues observed in the financial industry's attack surface were either related to exposed RDP servers or worse, were related to accidentally exposed database storage and analytics systems. Database systems should never be available over the public internet and should always be behind a firewall or a VPN.

Data storage exposures in this industry are worrying as they could contain key customer or transaction data. Additionally, with everyone working from home, these exposures become even riskier as employees are accessing critical information through potentially vulnerable, non-consumer-grade access points.

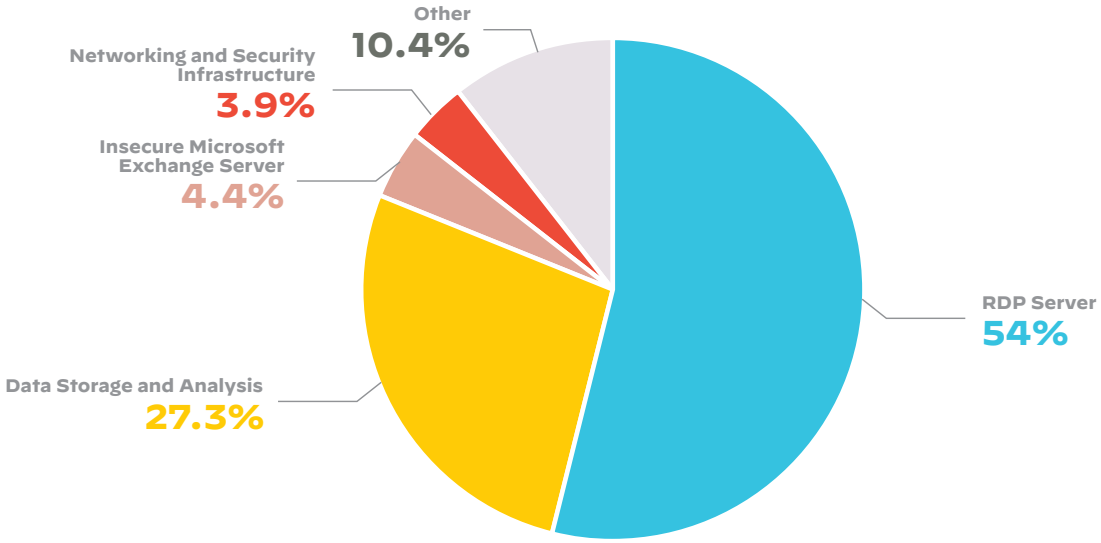

Figure 9: Distribution of risks across the financial services industry attack surface

Professional & Legal Services

The combination of the most prevalent issues we've observed in the professional and legal services industry makes for an incredibly dangerous scenario. Organizations that have data storage exposures with unencrypted logins are opening themselves up to a "when" and not an "if" scenario as it relates to a serious data breach or ransomware.

A potential breach in this industry could put intellectual property, critical customer information, and other highly sensitive information at risk.

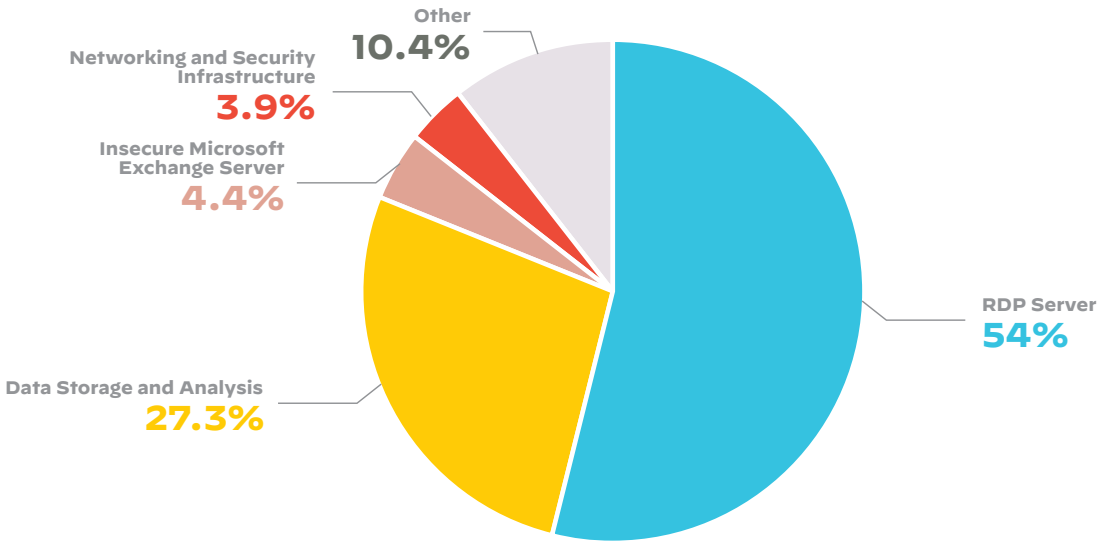

Figure 10: Distribution of risks across the professional and legal Services industry attack surface

High Technology

In high-tech companies, exposed IT admin portals are the biggest issues found on attack surfaces, and as we noted above, these exposures are extremely appealing to attackers because a successful exploit here grants attackers access to the entire network and privileges to cause untold damage. High-tech companies need to be able to quickly discover they're exposed IT admin login pages and immediately put them behind a firewall or VPN.

FTP is not an industry-standard protocol anymore and is in violation of numerous regulatory compliance standards. It relies on clear-text usernames and passwords for authentication and does not use encryption, but they seem to be very prevalent in this industry from our observations. Data sent via FTP is vulnerable to sniffing, spoofing, and brute-force attacks, among other basic attack methods.

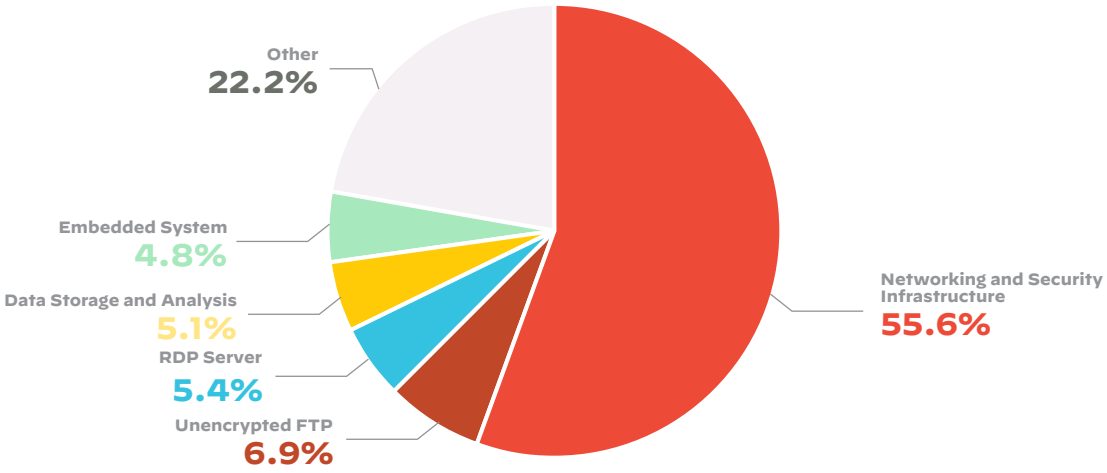

Figure 11: Distribution of risks across the high technology industry attack surface

Media & Entertainment

Data storage and analysis infrastructure potentially containing key IP and critical customer data were the biggest risks observed on the attack surface of the media and entertainment industry.

One of the most notable attacks in this space was on Sony Pictures in 2014, when attackers stole terabytes worth of data and Sony had to shut down its network for days. Following the breach, five movies, four of which were unreleased at the time, were leaked by the attackers.

RDP services were the second biggest issue type, and potential attacks could cripple the industry in its ability to service its customers.

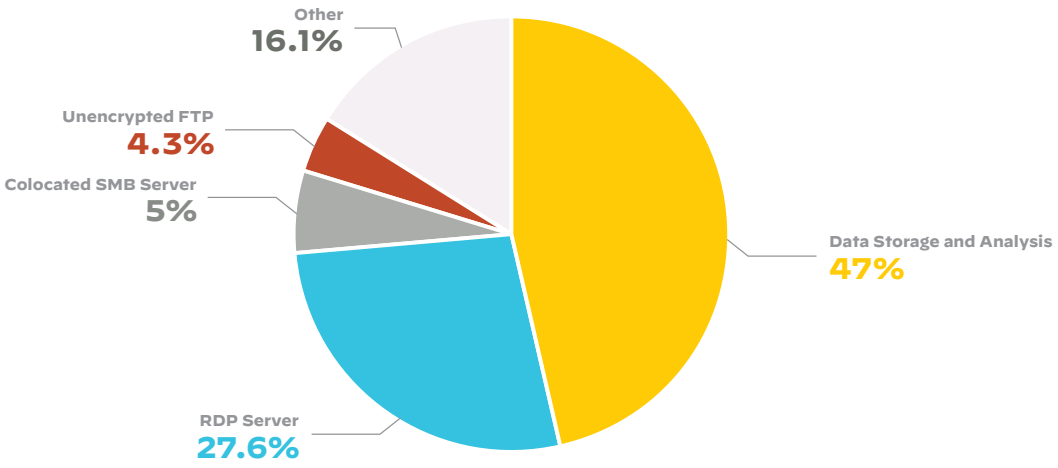

Figure 12: Distribution of risks across the media and entertainment industry attack surface

Insurance

Data storage, RDP servers, and networking infrastructure vulnerabilities show up in almost equal measure in the insurance industry.

In May 2021, insurance company CNA Financial reportedly paid $40 million in ransom after being unable to restore systems for two weeks after an attack. Additionally, the personal information of more than 75,000 customers was compromised in the breach.

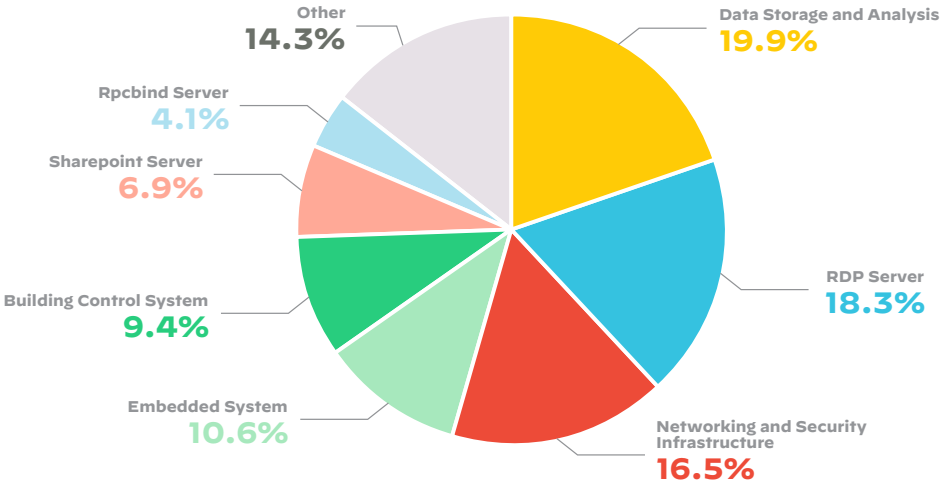

Figure 13: Distribution of risks across the insurance industry attack surface

Business Impact Takeaway

While attack surfaces may seem unique, to an attacker there are certain exposure types that are usually more readily available, so the initial attack vector may not vary too much. As we can see in the data, there are multiple ways remote access protocols or remote logins can and are exposed, leading to easy entry points for attackers.

Whether the ultimate aim is real-world disruption or high-value data, industries have to know what exposures are putting them at risk because no one wants to deal with a ransomware attack, stolen data, or both.

Knowing the most common risks can help vulnerability management to focus efforts, but even beyond the major issues for each industry, there are many other exposures lurking in the shadows. The best bet is to have a comprehensive and continuously updated inventory of all assets and potential exposures in order to minimize all risks.

Conclusion and Recommendations

Security is hard. Sometimes it's as simple as that. Security teams do the best they can with the resources and the data they have, but visibility is often the deciding factor as to whether an asset is secure or not. The number one tool in the SecOps arsenal should be an attack surface management platform that can provide a comprehensive and continuously updated inventory of all internet-connected assets and potential exposures.

If you don't know where exposures live, it's impossible to ensure issues are remediated. For many organizations, the cloud and RDP are going to be persistent issues to target, but the constellation of exposures and vulnerabilities on your attack surface will only continue to grow as attack surfaces get more complex.

Unfortunately for defenders, attackers just need one crack to find their way in. Attackers thrive on the complexity and ever-changing nature of attack surfaces because they can scan the entire internet looking for those weak points. The best option for security teams is to ensure they have the same view of their own attack surface. With an attacker's point of view, identifying and prioritizing issues for remediation gets far easier.

This also means focusing on metrics like mean time to detect (MTTD) and mean time to respond (MTTR) is inherently flawed. In the case of a breach, MTTD and MTTR are acceptable, but security should be focused on doing all they can to prevent breaches before they happen. That means putting more stake in mean time to inventory (MTTI), because it is impossible to secure unknown assets and unknown exposures.

Modern attack surfaces are dynamic. Without clear visibility that is constantly updated, it is all too easy to have persistent exposures and unmanaged assets. Security practitioners can only be as good as the data they have, so having a strong foundation of continuous discovery and monitoring ensures you can keep up with modern, dynamic attack surfaces in order to find, prioritize, and mitigate exposures as they arise.

Want to see your attacker's view into your attack surface/network? Reach out to a Palo Alto Networks representative for a custom outside-in view into your organization.

Methodology

Cortex 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.

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 non-technical audiences.
- Deploy proactive security measures.

In this analysis, Xpanse looked at 2021 data (beginning of March to end of September) from 100+ tenants spanning multiple industries. Data for a given industry was only displayed in this report if the industry included data from seven or more organizations in that industry category. Key Critical Vulnerabilities and Exposure (CVE) observations are based on observed data from January to February 2022.

It is important to note that the majority of metrics in this report are medians. The reason for this is that, when it comes to global internet asset data, there are frequent outliers that can heavily skew the mean/average for a given issue type or industry and using medians helps us avoid seeing a biased view.

About Cortex Xpanse

Cortex® Xpanse™ is an automated attack surface management (ASM) platform that provides a complete and accurate inventory of an organization's global internet-facing assets and misconfigurations to continuously discover, evaluate, and mitigate security issues on an external attack surface. Xpanse customers include leading Fortune 500 companies as well as US government organizations, including all five branches of the military.

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