How to Defend Your Attack Surface

Continuously Discover, Track, and Protect Your Attack Vectors

Bad actors are constantly looking for ways to attack organizations. They hunt for vulnerabilities on websites, exposed servers in the cloud, and other internet-connected systems and services that have been forgotten about or have little-to-no protection. Organizations need to understand their attack surface, all the ways their infrastructure is exposed and vulnerable to attack and prioritize activities that can help make that attack surface smaller.

The Potential Damage to Your Attack Surface Is Real

Security researchers found 1.2 billion records with individuals' personal data aggregated by People Data Labs on an exposed Elasticsearch server[1]

MoviePass exposed credit card information for thousands of customers on a server open to the internet that was unencrypted and not password protected[2]

Hackers compromised a reservation database for Marriott's Starwood division and accessed the data of 383 million guests[3]

A database managed by the Indian government was left open on the internet without a password, exposing the medical records of more than 12.5 million pregnant women[4]

A brute-force attack on an exposed RDP server from LabCorp resulted in 7,000 systems and 1,900 servers infected[5]

With so many attack vectors and limited resources to defend them, it's critical that organizations understand where the critical entry points are and how they can prioritize attack surface reduction in a smart, data-driven way. In this white paper, we'll explore what makes up an organization's attack surface, major risk drivers and critical exposures, and how to regain control of your attack surface. This data enables you to know your internet—what you own, where it resides, and what risks you have.

Do You Know Your Internet?

Organizations have a lot of public-facing systems and services, in addition to data about them. Their internet presence can include websites, networking equipment, mobile apps, LinkedIn data, or even Glassdoor reviews. Attackers scour the internet for this data to find an entry point and plan an attack. The classes of off systems and data fall into a few different buckets, some of which may not actually belong to the organization itself.

Corporate Systems and Services

This consists of the directly attackable parts of your network like websites, networking equipment, remote access protocols, and exposed user workstations. It can also include your larger cyber ecosystem, such as strategic supplier or remote subsidiary networks that might be targeted for an attack.

Brand Protection

This encompasses issues like malicious apps, domain squatting to phish your employees or customers, and counterfeit services.

Company Intel

A broad category of data, some of which is directly discoverable. Hackers discussing imminent attacks on dark web forums, your employees' personally identifiable information for sale, or even your executive's administrative assistant, who can then be targeted. All three of these are often lumped together, but each serves a distinct purpose and comes with a different level of importance. For example, learning that hackers are talking about targeting your organization on an underground forum is only useful if you can actually do something about it, like adding extra security staff to support an incident response or sending out a warning email to your employees to be extra diligent about incoming emails. In the same way, learning that your data has already been exposed can be useful, but everyone would prefer to prevent a breach in the first place instead of just having excellent detection capabilities.

Brand protection is also very difficult to do well. Many organizations register thousands of domains defensively, purposefully owning misspellings and typo domains to make it harder for attackers to scoop these up. But with so many combinations of domains, it's nearly impossible to cover them all and keep an up-to-date inventory. Phishing attacks are usually launched extremely rapidly, with the period

1. Corinne Reichert, "1.2 billion records exposed in unsecured database," CNET, November 22, 2019, https://www.cnet.com/news/1-2-billion-records-exposed-in-unsecured-database/.
2. Corinne Reichert, "MoviePass left customers' credit cards exposed online," CNET, August 22, 2019, https://www.cnet.com/news/moviepass-reportedly-left-customers-credit-cards-exposed-online/.
3. Alfred Ng, "Marriott says hackers stole more than 5 million passport numbers," CNET, January 4, 2019, https://www.cnet.com/news/marriott-says-hackers-stole-more-than-5-million-passport-numbers/.
4. Catalin Cimpanu, "Indian govt agency left details of millions of pregnant women exposed online," ZDNet, April 1, 2019, https://www.zdnet.com/article/indian-govt-agency-left-details-of-millions-of-pregnant-women-exposed-online.
5. Steve Ragan, "Samsam infected thousands of LabCorp systems via brute force RDP," CSO Online, July 19, 2018, https://www.csoonline.com/article/3291617/samsam-infected-thousands-of-laptop-systems-via-brute-force-rdp.html.

between domain registration to emails sent being as short as minutes, but the detection, alert, and action cycle can be much longer. Recent data suggest that a significant proportion of domains are used days, weeks, or months later, but dynamic learning and blocking of suspicious inbound emails is likely a more effective strategy.[6] Another common brand protection strategy is to look for logo images, but the false-positive rate can make this an unrealistic endeavor.

Fraudulent mobile apps and social media attacks can cause harm, but they require much more effort from attackers and impact victims one by one (as opposed to a central data breach). Any exposed network services, on the other hand, are ticking time bombs just waiting to be found. By far, the cheapest and most successful attack against organizations is simply scouring a perimeter until an unsecured device is found.

Given the limited number of resources available to security staff, it's important to focus on the biggest risk factors when securing your organization. The data shows that for many organizations, lack of basic perimeter security and hygiene cause the largest number of data breaches and the biggest impact on the bottom line. Organizations should independently determine their biggest attack vector. Public data can be complicated to interpret. One of the most well-known studies is the Verizon Data Breach Investigations Report. This includes a wide range of cyber incidents, including hacking, phishing, denial of service, and other threat actions. Hacking (which excludes phishing and includes perimeter attacks) consistently leads all other threat actions as the major cause of data breaches. Servers lead the category of affected assets, most often targeted because they are publicly exposed, just waiting for a bad actor to find them.

The Primary Cause of Breaches

With all the advanced technologies that have been developed over the past decade, it's easy to lose sight of some of the basic tenets of security. Knowing your network and the devices on it is the number-one control listed by SANS/CIS, yet most organizations haven't conducted an IP list audit in years.

The Privacy Rights Clearinghouse maintains a public ledger of data breaches. The data convincingly show that "hacks" (which includes both perimeter attacks and phishing) are by far the most frequent and costly form of data breach.[7] A privately maintained database of incidents recorded 10 times more hacking incidents than phishing incidents.[8] And individual organizations that have been analyzed also show perimeter incidents occur at three to four times the frequency and impact compared with email phishing.[9]

The figure from the Cortex® Xpanse™ Internet Operations Management platform shows the combined perimeter exposures for a sample of a dozen Fortune 100 organizations. Cumulatively, they have over 700 certificate hygiene issues and the occasional Telnet or SNMP exposure on registered ranges alone. When looking across the cloud, the counts become even worse, with some organizations having dozens of RDP instances publicly exposed.

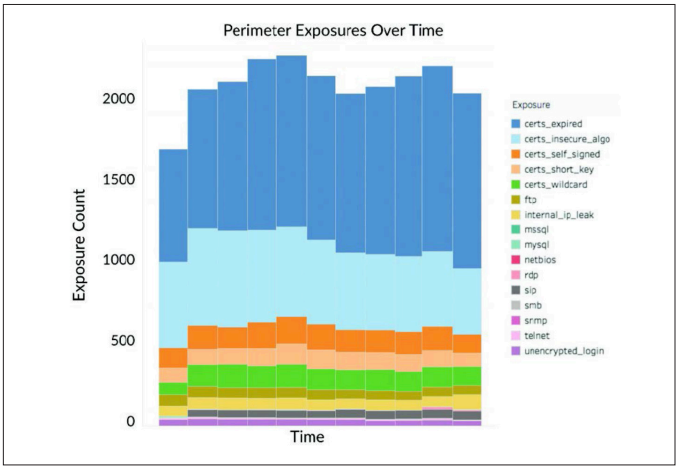

Figure 1: Perimeter exposures over time

6. Greg Aaron and Rod Rasmussen, "*Global Phishing Survey: Trends and Domain Name Use in 2016*," APWG Internet Policy Committee, June 26, 2017, https://docs.apwg.org/reports/APWG_Global_Phishing_Report_2015-2016.pdf.
7. Annette Hoffmann, Spencer Wheatley, and Didier Sornette, *Heavy-Tailed Data Breaches in the Nat-Cat Framework & the Challenge of Insuring Cyber Risks*, Research Gate, January 2019, https://www.researchgate.net/publication/330132704_Heavy-Tailed_Data_Breaches_in_the_Nat-Cat_Framework_the_Challenge_of_Insuring_Cyber_Risks.
8. Sasha Romanosky, *Examining the costs and causes of cyber incidents*, Journal of Cybersecurity 2.2 (2016): 121-135, August 2016, https://www.researchgate.net/publication/306927611_Examining_the_costs_and_causes_of_cyber_incidents.
9. Marshall A. Kuypers, Thomas Maillart, and Elisabeth Paté-Cornell, *An Empirical Analysis of Cyber Security Incidents at a Large Organization*, Semantic Scholar, 2016, https://www.semanticscholar.org/paper/An-Empirical-Analysis-of-Cyber-Security-Incidents-a-Kuypers-Maillart/f00c9721e0adbebe19630ef38497b696cd6c8567.

Attackers don't even need to resort to phishing or other attack techniques when so many misconfigured devices are sitting around. They can simply find open exposures to figure out how to access sensitive internet assets and data. Leaving these assets accessible from the public internet is like putting your servers out on the sidewalk.

Discover, Prioritization, and Risk Management

How can an organization begin to claw back its attack surface? How can an organization reduce the parts of its internet presence that actually cause breaches? It all starts with discovering what your attack surface really is.

Discovering and Mapping Your Attack Surface

The attack surface area of an organization has never been more distributed than it is today. Organizations have to identify, track, and manage more asset types across different locations than ever before. A discovery and mapping program should start with the basics:

- A system of record of every asset, system, and service you own that is on the public internet, including across all major cloud providers and commercial ISP space (not just known registered ranges)
- Comprehensive indexing, spanning all major port/protocol pairs (i.e., not limited to the old perspective of only tracking HTTP and HTTPS websites)
- Leverage multiple data sources for attribution (i.e., not just registration and DNS data)
- No reliance on agents (which can't find unknown internet assets)
- Continuous updating (i.e., not a two-week refresh rate)

Where Are My Internet Assets? The Importance of Global Coverage

In the past, the majority of an organization's attack surface was based on static ranges that were registered to that organization. Today, organizations need to search for their assets, systems, and services across the entire internet.

Core IP Space

Core ranges are table stakes. Organizations need to rapidly discover and monitor known ranges for inadvertent misconfigurations or device exposures. Any exposures on these ranges are highly attributable and are likely to be targeted quickly.

Subsidiary and Acquisition Networks

Attackers look for entry points anywhere they can, including nested subsidiaries and historical acquisitions. The Xpanse Internet Operations Management platform can easily identify internet assets that were orphaned during an M&A event and are unmonitored. Organizations should take care to search for abandoned servers, cloud instances, or remote protocols as an example of internet assets that may have been overlooked in the past.

Cloud Instances

Organizations are moving to the cloud, and it has never been easier for an employee to sign up a cloud instance outside of normal IT processes. Organizations should have focused discovery of internet assets pointed at all cloud environments, including AWS, Azure, Google, Oracle, Rackspace, and other cloud-hosting providers.

Remote Access Protocols

A mobile workforce has created new classes of risk that haven't existed before. Traveling employees may have misconfigured workstations that expose their laptops to the world via Remote Desktop Protocol. These exposures are highly ephemeral because they move as the employee travels from home to a coffee shop to a hotel.

Strategic Suppliers

Suppliers are more connected than ever. It's often impossible to do business without sharing sensitive data or permitting network access to critical business partners. Exposures on these fringe segments of your network can lead to data loss or network intrusions on your corporate enclave.

Overall, these different locations add up to the entire global internet. Organizations have networks that are so widely distributed that they need to monitor the entire internet to accurately track their internet-facing presence.

Attackers are constantly looking for Telnet, SSH, SMB, and RDP. If you leave any of these exposed, an attacker is almost certain to find them and find them quickly.

Bad Actors Constantly Hunt for Exposed Devices

Xpanse conducted an experiment to see what attack vectors bad actors were looking for across the internet. Honeypot listeners were deployed and using the Xpanse Internet Operations Management platform, key traffic patterns were captured and analyzed. Figure 2 shows the number of unique scanners and the number of unique scan attempts targeting our honeypots by port.

Key Takeaways

- Ports that are closer to the upper right corner were scanned more often by more people.
- TCP ports were scanned more than web ports.
- Attackers were constantly looking for Telnet, SSH, SMB, and RDP even more so than websites.
- Database servers and teleconferencing devices were popular targets as well.

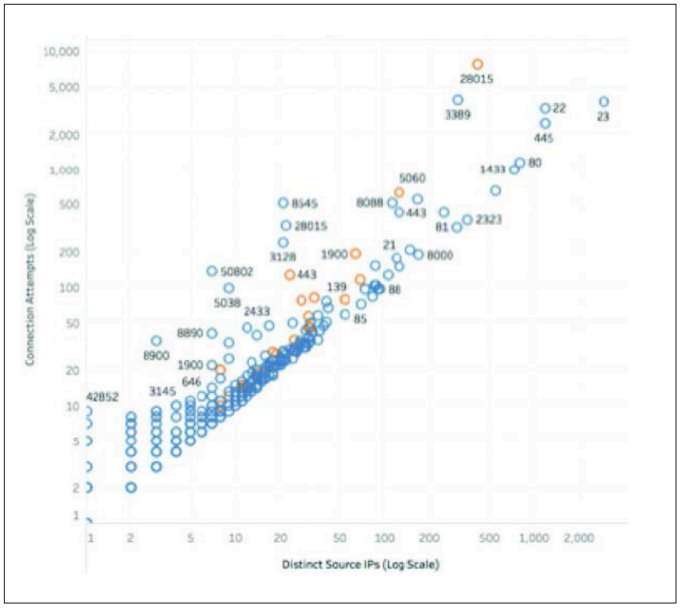

Figure 2: Unique number of scans targeting our honeypot ports

Table 1: Exposures Attackers Will Find, and Find Quickly

Port	Protocol	Common Device
80	HTTP	Web Server
443	HTTPS	Web Server
23	Telnet	Network Infrastructure
3389	RDP	Workstation, Servers
1433	54	Database Server

Defending Your Attack Surface

Discovering and mapping your attack surface is just the first step. Ongoing monitoring is essential to remaining secure and building a proactive defense posture. Critical capabilities include:

- A continuously updated system of record
- Continuous monitoring
- Processes that account for a dynamic infrastructure, including new cloud providers and network ranges
- Conducting periodic audits

You Need Independent, Global Data

Security teams have a tendency to use the same data to manage and audit a system. Audits will look at vulnerability scan outputs to confirm that all critical vulnerabilities have been patched. But what about the systems that aren't being scanned because they reside outside of the known network? Or what about the alerts that have been muted because they have waivers in place to allow an insecure machine to be public-facing for a critical business project? Audits that don't use independent data won't catch these and will fail.

A better strategy for auditing your system is to obtain independent, global data. Independent data avoids your bias, exceptions, and errors. Global data gives you the best chance of finding unknowns, and it anchors the analysis in ground truth instead of focusing on where you happen to have the best visibility.

When you have a comprehensive global program for discovering, monitoring and managing your attack surface, you can avoid some of the most common risks facing organizations today. These risks include:

Remote Ransomware: The RDP Problem

RDP is a top threat vector for ransomware attacks. A workstation with RDP exposed on the public internet is the equivalent of leaving a laptop open to its login screen sitting on the street, where anyone can try a username and password. Most organizations think that they're blocking RDP across their networks and devices, but the Xpanse Internet Operations Management platform regularly finds RDP instances for organizations on the public internet, including a majority of the Fortune 100.

The most common attack against RDP starts out with a brute-force password-guessing attempt. If the password isn't complex enough or if there aren't lockout attempts, then attackers will eventually compromise a device. Once this happens, ransomware is typically installed, which can spread throughout the organization, causing significant business interruption incidents. Data is encrypted or destroyed, leaving organizations with a crippled network caused by an unknown exposure that occurred in IP space that they were not monitoring.

These exposures are especially difficult to track because they often occur outside of places regularly monitored by the organization's IT and security staff. Without the complete, current, and accurate indexing of the entire internet provided by Xpanse, organizations don't have a way of tracking these findings themselves.

By indexing the global internet multiple times per day, Xpanse helps customers detect exposures like RDP before they are targeted, not weeks after the exposures have occurred and been found and exploited by attackers.

Cutting Down Your Cloud Risk

A number of the data breach stories of 2019 weren't typical hacks but rather were caused by an employee who may have initialized a database server in the cloud without telling the security team. Developers often use production data instead of fake data for testing, and if these databases become inadvertently exposed, it's just a matter of time before they are found. IT and security teams aren't informed and can't see them because they exist outside of a standard policy.

These exposures occur easily and can expose sensitive data to malicious actors with virtually no effort required. Organizations need to remain diligent by looking for common database misconfigurations (SQL, Elasticsearch, MongoDB, Memcached) and not just on their core IP space, but across cloud environments, too.

The Time to Protect Your Attack Surface Is Now

Fifteen years ago, if you exposed a device accidentally, it might go unnoticed for months or even years. Things are different today. Attackers can find every device on the internet in 45 minutes. Any misconfiguration or accidental exposure is likely to be discovered very quickly. Internet-scale attacks like WannaCry and NotPetya demonstrate how a new wave of attacks doesn't target specific companies but rather seeks out and attacks vulnerabilities across the entire globe.

The attack surface has grown to include the cloud and even commercial ISP space, creating new challenges for organizations trying to reduce entry points into their networks. Xpanse specializes in identifying high-risk internet asset types that occur outside of known IP space, especially in cloud environments and commercial ISP space. With the Xpanse Internet Operations Management platform, you can discover your attack surface and take steps to minimize it, resulting in a more secure organization.

Do you know your internet? Xpanse does. Visit start.paloaltonetworks.com/demo-request.html to set up a demo.

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