

Industry Cyber-Exposure Report

ASX 200

TABLE OF CONTENTS

Executive Summary	5
Overview of Results	7
Phishing Defence Capabilities	7
Attack Surface by Sector	8
Critical Service Exposure	8
Vulnerable Service Exposure	9
Third-Party Risk Exposure	13
Evidence of System Compromise	17
Conclusions	21
Measuring Industry Exposure: Methodology Overview	23
Active Measurements with Project Sonar	24
Passive Measurements with Project Heisenberg	25
Measuring Web Server Third-Party Risk	25
Further Work	27
Improving Entity Internet Asset Attribution	27
Avoiding Opt-Out Opacity	27
Finding More DNS Records	28
Expanding Resource Safety Evaluation	28
Third-Party Dependency/Risk Analyses	28
Study Methodology	29
Why the ASX 200	29
Organisation Internet Asset and Metadata Attribution Methodology	30
About Rapid7	33

Executive Summary

In the face of growing cybersecurity threats, it is increasingly important to measure the cost and concentration of “exposure,” which we define here as weaknesses in the public-facing configuration of internet-connected services. Having an accurate view of the resilience of organisations and industries against cyber-attacks can facilitate more accurate cost models, help target efforts to reduce exposure to the industries that need it most, and enhance cooperative efforts between government and the private sector to better protect users and companies alike. Measurement of industry-level exposure can also inform industry-specific working groups that share cybersecurity information and threat intelligence, such as Information Sharing and Analysis Centres.

To understand current levels of exposure and resiliency in Australasia,¹ Rapid7 Labs measured the internet-facing security profiles of the S&P/ASX 200² (ASX200) during Q4 2018 for:

- Overall attack surface (the number of exposed servers/devices);
- Presence of dangerous or insecure services;
- Phishing defence posture;
- Weak public service and metadata configurations; and
- Joint third-party website dependency risks.

An important factor to consider in the context of discovered weaknesses is that members of the ASX 200 list are well-resourced organisations that typically attract top talent in all aspects of the business, including information technology (IT) and security. The discovery of such widespread weaknesses in the exposed services of these leading organisations makes it likely that there is even greater exposure and risk in smaller organisations with fewer human and financial resources available for securing their public internet resources.

Key findings include:

- ASX 200 organisations, on average, expose a public attack surface of 29 servers/devices, with many companies exposing 200–300+ systems/devices.
- Severely vulnerable services such as Telnet and Windows file-sharing are not prevalent for the most part, which is positive. However, most organisations in every sector have serious issues with patch/version management of business-critical internet-facing systems.
- Of the appraised ASX 200 organisations, 134 (67%) have weak or nonexistent anti-phishing defences (i.e., DMARC) in the public email configuration of their primary email domains.
- Every industry sector in the ASX 200 signals how many and which cloud service providers they use in their public domain name system (DNS) metadata, with 144 organisations using between two and five cloud service providers and some using 10 or more. This information can be used to craft highly effective, targeted attacks, among other actions.
- All industry sectors had at least one organisation with malware compromises, with the Consumer Discretionary and Information Technology sectors³ showing daily signs of ongoing compromise. These compromises ranged from company resources being co-opted into denial-of-service (DoS) amplification attacks to signs of EternalBlue-based campaigns similar to WannaCry and NotPetya.

The details behind these findings are presented in the remainder of the report.

¹ While the ASX primarily comprises Australian companies, some New Zealand companies also list on the ASX, as described here: https://www.asx.com.au/prices/new_zealand_based_companies.htm (last accessed Feb. 12, 2019).

² S&P/ASX 200 list and definition, <https://www.asx200list.com/> (last accessed Feb. 12, 2019).

³ ASX Industry Sectors overview, <https://www.asx.com.au/products/sector-indices.html> (Last accessed Feb. 11, 2019).

Overview of Results

Phishing Defence Capabilities

Phishing remains one of the most common cyber-attack vectors corporations face today; the Anti-Phishing Working Group (APWG) collected over a quarter-million phishing reports in the third quarter of 2018.⁴ As noted in the **Methodology** section (pg. 23) on domain name attribution, DNS records expose a means to identify how well an organisation has configured its email service for protection from spam and phishing through the analysis of Domain-based Message Authentication, Reporting and Conformance (DMARC) records.⁵ DMARC enables organisations to:

- Signal that they are using email authentication to prove emails are not forged;
- Provide an email address to gather feedback about messages using their domain, legitimate or not; and
- Apply a policy to messages that fail authentication (one of “none”, “quarantine”, or “reject”).


No DMARC records—or a DMARC record of “none”—means this first-line-of-defence protection from spam or phishing attacks is absent. However, a “none” record may be a signal that an organisation is on the path to email safety and is in the process of validating its DMARC configuration before enabling more active email defence measures.

Properly configured DMARC records with “quarantine” or “reject” have active email defence measures in place.

Figure 1 shows the percentage of DMARC adoption (by configuration category) of ASX 200 organisations within a given sector. Green indicates that organisations within that sector have either fully adopted and implemented DMARC or are on the path toward DMARC adoption. Unfortunately, the results indicate that the majority (67%) of the ASX 200 have not embraced modern email safety configurations, boosting their risk of phishing attacks. Several industry sectors have no organisations with DMARC configured to “quarantine” or “reject.”

Since there is no direct scanning involved, DNS records are not impacted by the Project Sonar opt-out blacklist (described in the **Further Work** section (pg. 27)). Therefore, we can paint a more complete picture of the email safety configurations of the entire ASX 200 than we can with active scanning for vulnerable services. The **Further Work** section outlines additional steps that can be used to increase the scope of the examination to paint a wider picture of email safety.

Figure 1: Email Safety Status of ASX 200 Primary Email Domains



⁴Phishing Activity Trends Report, http://docs.apwg.org/reports/apwg_trends_report_q3_2018.pdf (last accessed Feb. 12, 2019).

⁵DMARC, <https://dmarc.org> (last accessed Feb. 12, 2019).

Attack Surface by Sector

The **Methodology** section (pg. 23) details how Rapid7 uses Project Sonar to scan the internet for exposed systems and devices. On average, each ASX 200 organisation exposes roughly 29 services. This number is neither good nor bad, but each exposed node increases the attack surface of an organisation, potentially boosting the opportunity for attackers to gain a foothold. To put it another way, each exposed server or device must be properly configured, managed, patched, and defended to reduce the risk of a cyber-attack. There is no rule to indicate when the number of exposed services tips the balance of risk, since many factors can influence how well an organisation can protect internet-exposed resources. Still, the more exposed systems, the more opportunity attackers gain, regardless of defensive capabilities.

Taking a look at Figure 2, there are four significant outliers exposing well over 100 services: two in Financials, one in Telecommunication Services, and another in Information Technology.

Recommendation: Reduce Attack Surface

As organisations expand the number of assets they expose on the internet, they must also ensure they have commensurate vulnerability management, patching, and monitoring practices in place to facilitate a speedy response to discovered weaknesses or attempts by attackers to compromise services. Organisations should strive to only expose systems and devices on the internet if they support business processes and must further ensure they have robust asset identification and configuration management processes in place to help avoid these systems becoming enterprise entry points for attackers.

Critical Service Exposure

The type of service being exposed has a direct impact on the severity of exposure (i.e., some services are less “safe” than others). Figure 3 shows that organisations in the ASX 200 are not immune to attacks that target two critically vulnerable services: Telnet and Windows file-sharing.

One service in particular—Server Message Block (SMB)—is one of the most dangerous services for a system to expose. SMB is an all-in-one file-sharing and remote administration protocol, usually associated with Windows, that has been an attractive target for attackers and researchers alike for decades. MS03-049 in 2003, MS08-067 in 2008, and MS17-010 (EternalBlue) in 2017 all arose from the complexity of this protocol and its central nature to Windows networking.⁶ Recently, vulnerabilities in the SMB service were at the heart of the WannaCry and NotPetya attacks, which crippled networks and caused significant outages to critical business processes that cost many companies millions of dollars in lost revenue.⁷

⁶Rapid7, National Exposure Index 2018, “Inappropriate Services,” pg. 14, Jun. 7, 2018, https://www.rapid7.com/globalassets/_pdfs/research/rapid7-national-exposure-index-2018.pdf.

⁷Bob Rudis, “No More Tears? WannaCry, One Year Later,” Rapid7, May 14, 2018, <https://blog.rapid7.com/2018/05/14/no-more-tears-wannacry>.

Figure 2: Distribution of Discovered Organization Asset Totals by Sector

Each dot represents one organization; position on axis = number assets discovered.



While ASX 200 organisations expose roughly 29 services on average, four organizations across the Financials, Telecommunication Services, and Information Technology sectors exposed more than 100.

Telnet exposure creates risks similar to SMB exposure. Telnet dates back to the early days of the internet, with the official “modern” standard dating back to 1983.⁸ Telnet is a cleartext protocol that is used to directly log in to servers and network equipment, usually to issue commands and run scripts directly at the operating system level of the device. Telnet

services have a history of vulnerabilities and exposures that put organisations at risk of credential theft, passive and active eavesdropping, and remote code execution. The cleartext nature of the protocol means that an attacker in the proper network position can read any usernames, passwords, or data being transmitted, and endpoints with weak, default, or eavesdropped passwords can be hijacked to run malicious code directly by the operating system.

The singular positive takeaway is that only five sectors have at least one member exposing either Telnet or SMB. While a total absence of Telnet and SMB on today’s internet would be ideal, the ASX 200 has far less SMB/Telnet exposure, both in absolute and relative terms, than the Fortune 500.⁹

Recommendation: Eliminate Public-Facing SMB and Telnet

Though the presence of these services—especially Windows SMB— is extremely small among ASX 200 organisations, there is no safe way to expose SMB services to the public internet. In light of this, Microsoft has made efforts to reduce SMB exposure for normal desktop and laptop clients. For example, current Microsoft documentation explicitly recommends blocking SMB on an internet perimeter firewall, and Windows 10 desktops automatically firewall access to port 445 by default.¹⁰ Even exposing one asset with SMB running could end up spreading (or re-spreading) WannaCry, NotPetya, or modern variants across an entire organisation.

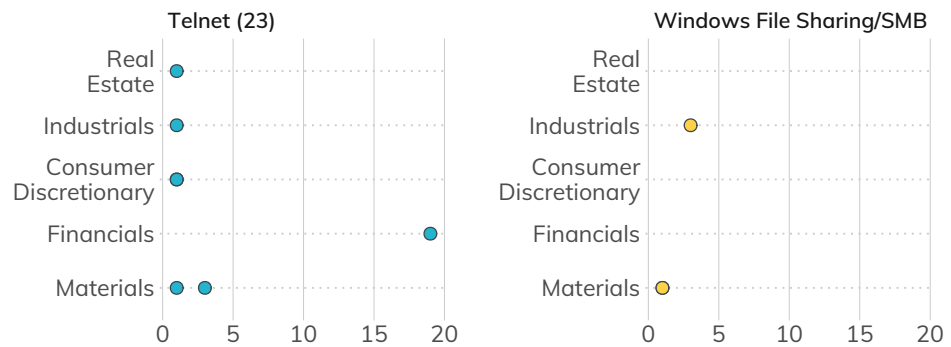
There is also no technical or practical justification for running a Telnet service today. It has been superseded by the Secure Shell (SSH) Transport Layer Protocol, which provides encryption-in-transport and encourages the use of digital certificates when authenticating connections.¹¹ If a device is truly incapable of running SSH rather than Telnet due to a lack of local computing resources, that device is simply too insecure by design to expose to the public internet, regardless of the reasoning for staying with a 40-year-old unencryptable protocol. Of note, nearly half (80) of the ASX 200 expose SSH services without exposing any Telnet services, so it seems there has been some acknowledgement of the strengths of this protocol.

Vulnerable Service Exposure

Keeping internet services configured, patched, and running supported versions of operating systems and internet-facing applications can go a long way toward thwarting attackers. Unfortunately, most organisations in the ASX 200 are running old and often unsupported versions of the three most prolific web servers: Microsoft’s Internet Information Services (IIS), Apache HTTPD, and nginx.

Figure 3: Distribution of Organisations Exposing Telnet/Windows File-Sharing (SMB) Services

Each dot represents one organisation; position on axis = number assets discovered.



⁸ J. Postel and J. Reynolds, Telnet Protocol Specification, Internet Engineering Task Force, May 1983, <https://tools.ietf.org/html/rfc854>.

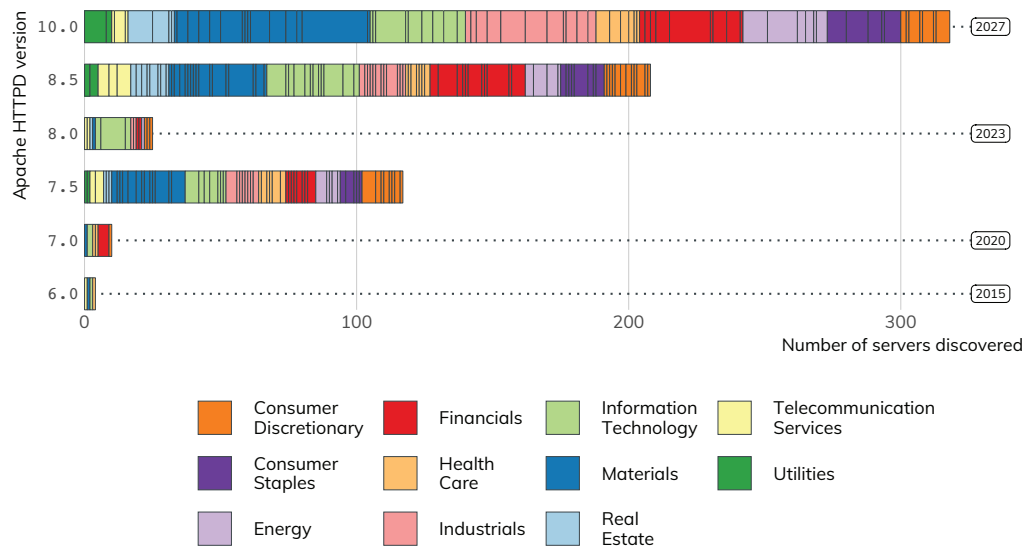
⁹ Rapid7, Industry Cyber-Exposure Report: Fortune 500, pgs. 13–14, Dec. 11, 2018, <https://www.rapid7.com/info/industry-cyber-exposure-report-fortune-500>.

¹⁰ Microsoft, Guidelines for blocking specific firewall ports to prevent SMB traffic from leaving the corporate environment, Aug. 31, 2016, <https://support.microsoft.com/en-us/help/3185535/guidelines-for-blocking-specific-firewall-ports-to-prevent-smb-traffic>.

¹¹ T. Ylonen and C. Lonvick, The Secure Shell (SSH) Transport Layer Protocol, The Internet Society, January 2006, <https://tools.ietf.org/html/rfc4253>.

Figure 4: Microsoft IIS Version Distribution

Each colour represents a different organisation; markers designate EOL year.



Microsoft IIS

Microsoft’s IIS was the most popular web server on the internet in January 2019, according to Netcraft.¹² This is no surprise, as IIS is usually found in enterprise settings due to software licensing agreements and large organisations’ continued reliance on Microsoft technologies. Figure 4 shows Project Sonar discovered 682 IIS servers with attributable version numbers in 143 organisations spanning all 11 ASX sectors.

Table 1 shows less than half of ASX organisations maintain only one version of IIS, with nearly 10% running over four separate versions. This version diversity and the discovery of end-of-life IIS versions increases defence and management complexity and further increases the likelihood of down-version IIS servers becoming an intrusion conduit for attackers.

Table 1: ASX 200 IIS Version Diversity

NUMBER OF IIS VERSIONS MAINTAINED	NUMBER OF ORGANISATIONS	%
1	67	46%
2	42	30%
3	24	16%
4	8	6%
5	2	2%

¹² Netcraft January 2019 Web Server Survey, <https://news.netcraft.com/archives/2019/01/24/january-2019-web-server-survey.html>.

Apache HTTPD

The web server version picture is even worse when we look at Apache. Figure 5 shows that Project Sonar discovered 418 Apache servers with 33 distinct version numbers in 80 organisations spanning 10 of the 11 ASX sectors (no versioned Apache servers were found in the Utilities sector). Table 2 shows a majority of organisations only expose a single version of Apache HTTPD, but 28% of them expose five or more distinct versions, which, again, increases the management complexity.

Figure 5: Apache HTTPD Version Distribution

Each colour represents a different organisation; markers designate release year.




Table 2: ASX 200 Apache Version Diversity

NUMBER OF APACHE VERSIONS MAINTAINED	NUMBER OF ORGANISATIONS	%
1	42	52%
2	7	8%
3	8	10%
4	2	2%
5	5	6%
6	6	8%
7	3	4%
8	2	2%
9	1	2%
10	2	2%
11	1	2%
13	1	2%

While the software version diversity alone is disconcerting, the fact that most discovered versions are over a year old is an indicator organisations aren't keeping Apache installations up-to-date. Unlike IIS, organisations must use a third-party vulnerability management tool to catalog and identify the version and patch levels of Apache servers. The Apache Foundation regularly releases new versions to add features, fix problems, and patch security issues. Also, because Apache HTTPD is open source, attackers have full access to the source code and can more easily discover flaws and develop exploits.

nginx

The nginx web server may come in third in the Netcraft survey, but it's No. 1 for the number of distinct versions Project Sonar discovered among the ASX 200 (35). Figure 6 shows that 291 nginx servers were found in 55 organisations spanning all 11 ASX sectors.

Table 3 is similar to the Apache findings, with many organisations running a single version but far too many running five or more.

Recommendation: Strive for Version Consistency


The need to patch and maintain servers may sound trivial and obvious, but scheduling change control and orchestrating outages of what may be core services can be a complex undertaking in large organisations. Though this task can be tedious, it's vital that organisations keep an up-to-date inventory of what they're exposing and work with their business and application teams to ensure they are using supported and patched versions of software.

Table 3: ASX 200 nginx Version Diversity

NUMBER OF NGINX VERSIONS MAINTAINED	NUMBER OF ORGANISATIONS	%
1	26	48%
2	7	12%
3	7	12%
5	2	4%
6	1	2%
7	2	4%
8	3	6%
9	4	8%
10	1	2%
11	2	4%
13	1	2%

Figure 6: nginx Version Distribution

Each colour represents a different organisation; markers designate release year.



Third-Party Risk Exposure

There is no question that the internet has become the backbone of international commerce in virtually every industry and locale. This interconnectedness means no organisation is an island and is a central reason why it is almost impossible to have a website, business process, or digital storefront without relying on some outside party. As a firm's digital footprint expands, the more the details of these third-party dependencies leak out through necessarily exposed metadata required to keep these services connected and operating smoothly.

An unfortunate result of this is that all organisations in the ASX 200 is vulnerable to targeted phishing attacks based on the third-party service metadata they expose in their DNS records. In addition, every ASX 200 organisation places itself and its website visitors at risk due to reliance on improperly configured third-party web services.

Figure 7: Third-Party JavaScript Execution by Sector and Service

Advertising

Consumer Discretionary		9 (0.3)			1 (0.03)
Consumer Staples		1 (0.07)		1 (0.07)	
Energy		6 (0.6)			
Financials		10 (0.4)		1 (0.04)	1 (0.04)
Telecommunication Services		3 (0.5)			
Real Estate		5 (0.24)			2 (0.1)
Materials	1 (0.03)	4 (0.1)		3 (0.08)	
Information Technology	1 (0.07)	16 (1.07)		4 (0.27)	1 (0.07)
Industrials		6 (0.29)		1 (0.05)	
Health Care		3 (0.21)			
	ads twitter	doubleclick	googleadservices	yahoo	

Analytics

Consumer Discretionary		10 (0.33)	5 (0.17)	14 (0.47)	4 (0.13)	1 (0.03)	4 (0.13)	2 (0.07)	
Consumer Staples		6 (0.43)		9 (0.64)	2 (0.14)		1 (0.07)		
Energy				7 (0.7)	2 (0.2)	1 (0.1)	2 (0.2)	2 (0.2)	
Financials	5 (0.2)	2 (0.08)	1 (0.04)	17 (0.68)	3 (0.12)	2 (0.08)	6 (0.24)	17 (0.68)	
Utilities				1 (0.2)		1 (0.2)	1 (0.2)		
Telecommunication Services		3 (0.5)	1 (0.17)	8 (1.33)				1 (0.17)	
Real Estate		3 (0.14)	4 (0.19)	11 (0.52)	4 (0.19)		3 (0.14)		
Materials		3 (0.08)	2 (0.05)	18 (0.46)	5 (0.13)		2 (0.05)		
Information Technology		4 (0.27)	1 (0.07)	15 (1)	8 (0.53)		2 (0.13)	1 (0.07)	
Industrials	1 (0.05)	5 (0.24)	1 (0.05)	15 (0.71)	4 (0.19)			2 (0.1)	
Health Care	1 (0.07)	6 (0.43)	5 (0.36)	9 (0.64)	4 (0.29)				
	en sight en	go mpulse	google	google analytics	google tag manager	hotjar	new relic	nr data	omtrdc

Content Delivery Network

Consumer Discretionary	6 (0.2)		3 (0.1)		3 (0.1)	1 (0.03)	
Consumer Staples	2 (0.14)			5 (0.36)	1 (0.07)		
Energy	6 (0.6)		1 (0.1)		1 (0.1)	2 (0.2)	1 (0.1)
Financials	41 (1.64)	2 (0.08)	1 (0.04)	7 (0.28)	2 (0.08)	1 (0.04)	
Utilities				9 (1.8)	1 (0.2)		
Telecommunication Services	2 (0.33)					1 (0.17)	
Real Estate			1 (0.05)			1 (0.05)	2 (0.1)
Materials			2 (0.05)		5 (0.13)	1 (0.03)	
Information Technology				5 (0.33)		3 (0.2)	
Industrials	9 (0.43)		4 (0.19)	3 (0.14)	3 (0.14)	5 (0.24)	1 (0.05)
Health Care	3 (0.21)				1 (0.07)	1 (0.07)	
	adobedtm	en25	jquery	tiqcdn	typekit	youtube	yting

Social

Consumer Discretionary	13 (0.43)		3 (0.1)		1 (0.03)
Consumer Staples	4 (0.29)		2 (0.14)		
Energy	2 (0.2)		3 (0.3)		
Financials	13 (0.52)		8 (0.32)		3 (0.12)
Utilities	2 (0.4)				
Telecommunication Services	8 (1.33)		1 (0.17)		
Real Estate	8 (0.38)		2 (0.1)		
Materials	7 (0.18)		2 (0.05)		4 (0.1)
Information Technology	10 (0.67)		7 (0.47)		3 (0.2)
Industrials	8 (0.38)		2 (0.1)		4 (0.19)
Health Care	8 (0.57)		1 (0.07)		
	facebook		linkedin		twitter

When an organisation uses third-party resources to supplement its online assets, it takes on risks associated with those third-party resources. Vulnerable third-party resources can be used as a conduit to attack the first-party organisation. For example, in September 2018, security researchers noted that many sites are vulnerable to web-based credit card-skimming attacks due to their reliance on third-party content delivery networks (CDNs).¹³ In another example, the Syrian Electronic Army used a compromised CDN in 2015 to take over a major news media outlet’s web presence and use it to send custom push notifications to readers.¹⁴

For the purposes of this report, “third-party risk” exposure is defined as being present either when:

- A measured organisation is seen to be relying on resources from a third-party site when building its own websites and applications; or
- A measured organisation exposes which third-party services it actively uses by leaving potentially sensitive artifacts in its published metadata.

¹³ Kevin Beaumont, “Magecart – new tactics leading to massive unreported fraud,” DoublePulsar, Sep. 19, 2018, <https://doublepulsar.com/magecart-new-tactics-leading-to-massive-unreported-fraud-5211c9883dea>.

¹⁴ Thu Pham, “Malicious Hackers Take Over Media Sites via Content Delivery Network Providers,” Duo Security, May 19, 2015, <https://duo.com/blog/malicious-hackers-take-over-media-sites-via-content-delivery-providers>.

The **Methodology** section (pg. 23) outlines how attributes of third-party risk are collected and analysed. Figure 7 shows the breakdown by general categories of resource: advertising-oriented, site analytics, sourced from a CDN, or incorporating code from social networking sites such as Facebook and Twitter.

In Figure 7, heatmap cells are coloured by an index score in which the number of discovered third-party resources in that sector is divided by the total number of industries in that sector. This normalises the amount so that over-represented sectors (such as Financials) aren't being penalised for having more representation in the ASX 200 list. The number on top represents the number of resources, and the parenthetical value is the index score.

Some of these third-party services are likely resilient to cyber-attacks and do not meaningfully contribute to the first-party organisation's degree of exposure. For example, it is unlikely that Google would be sufficiently breached as to be an unwitting conduit for malicious activity to client organisations. However, widespread common exposure for third-party services such as DoubleClick—which has regular occurrences of malicious ads in its network—does increase the shared risk across the sectors.

Figure 8 focuses on the latter component of third-party exposure: detecting the use of vendor applications/cloud services.¹⁵

In addition to providing the connection address for names such as <www.rapid7.com>, DNS records can identify secure email configurations, as detailed in the Phishing Defense Capabilities section. DNS records can also reveal which third-party providers an organisation uses for everything from application development to cloud hosting environments to file-sharing and more.

One way these services are exposed is through the use of verification information stored in free-form TXT records. To illustrate, Table 4 shows a sample of DNS TXT records for rapid7.com:

Table 4: Rapid7 DNS TXT Records Sample

DNS RECORD KEY	DNS TXT RECORD VALUE
rapid7.com.	smartsheet-site-validation.rapid7.com=wFJFw8OnJ0WwBCBDF7NuqH
rapid7.com.	MS=ms93061892
rapid7.com.	atlassian-domain-verification=+Mx+hFjC77g1TvA7K9Tp/5x7LvbyawRYOeZpkXhE/Xys/xciI66aaIgyQQAD88E7
rapid7.com.	citrix-verification-code=3d0b3642-a1b3-4cf3-8616-c9fb8cd0c2da
<ul style="list-style-type: none"> • "smartsheet-site-validation" signals that Rapid7 uses SmartSheet, a cloud spreadsheet service. • "atlassian-domain-verification" signals that Rapid7 uses cloud-based services by Atlassian, a provider of popular software development tools and platforms. • "citrix-verification-code" signals that Rapid7 uses services offered by Citrix. 	

¹⁵ The colour-coding normalisation on Figure 8 is identical to that of Figure 7's.

Rapid7 researchers used Project Sonar DNS collection data to examine the TXT records of the ASX 200 organisations in this study. Only well-known domain names were used (expanding on this effort to use additional domains is covered in the **Further Work** section (pg. 27)), and Figure 8 only focuses on the most prevalent or well-known third-party services.

It may come as no surprise that every industry sector uses Microsoft Office 365, and it is highly unlikely that Microsoft is going to fall prey to a non-nation-state attack that would enable Office 365 to be a malicious gateway into organisations. There is a high prevalence of Google Apps across the ASX 200 as well, plus a fairly well-populated use of all the other resources of interest. If organisations begin to stray from established and resilient service providers, they boost their risk of successful phishing and other types of attacks by observant, capable attackers who simply need to make a handful of DNS queries to create a list of targets.

Recommendation: Reduce Third-Party Risk Exposure

These findings may not seem like major risks when reviewed individually. In truth, many of these “validation” records are only required once and can be removed after the initial validation has occurred. These records prove that one is the true owner of a given domain, since in theory, only the true owner can add, modify, or delete DNS entries. If one were to look at those records in aggregate, it might be possible to find a common, shared third-party service in use by a large number of organisations or a boutique service provider used by only a handful of organisations. These may be high-value targets for malicious actors that seek to compromise multiple organisations, making resiliency of these third-party services all the more important.

Figure 8: Third-Party App/Cloud Usage Exposure via DNS Metadata

	Adobe	Atlassian	Cisco	Citrix	DocuSign	Dropbox	Facebook	GlobalSign	Google App Domains	Microsoft Office 365
Consumer Discretionary	3 (0.1)	7 (0.23)	3 (0.1)		2 (0.07)	2 (0.07)	1 (0.03)	3 (0.1)	30 (1)	31 (1.03)
Consumer Staples		1 (0.07)	1 (0.07)		1 (0.07)		1 (0.07)	1 (0.07)	3 (0.21)	10 (0.71)
Energy	2 (0.2)	3 (0.3)	2 (0.2)		1 (0.1)		1 (0.1)	3 (0.3)	6 (0.6)	8 (0.8)
Financials	6 (0.24)	7 (0.28)	3 (0.12)		1 (0.04)		3 (0.12)		13 (0.52)	18 (0.72)
Health Care	2 (0.14)	2 (0.14)			3 (0.21)		2 (0.14)	4 (0.29)	8 (0.57)	12 (0.86)
Industrials	15 (0.71)	5 (0.24)	2 (0.1)				4 (0.19)	3 (0.14)	8 (0.38)	16 (0.76)
Information Technology	9 (0.6)	6 (0.4)		2 (0.13)	1 (0.07)		7 (0.47)	4 (0.27)	25 (1.67)	16 (1.07)
Materials	10 (0.26)	2 (0.05)	2 (0.05)		1 (0.03)		1 (0.03)		7 (0.18)	25 (0.64)
Real Estate	10 (0.48)	2 (0.1)	1 (0.05)	1 (0.05)	3 (0.14)	1 (0.05)	1 (0.05)		5 (0.24)	15 (0.71)
Telecommunication Services	1 (0.17)	1 (0.17)			1 (0.17)				4 (0.67)	6 (1)
Utilities	2 (0.4)	1 (0.2)			2 (0.4)	1 (0.2)	1 (0.2)		1 (0.2)	6 (1.2)

Evidence of System Compromise

Much of the data for this report was gathered through active scanning and DNS queries, but Rapid7 also maintains a global passive sensor network, Project Heisenberg. The Heisenberg sensor network is described in detail in the **Methodology** section (pg. 23), but in short, it is a collection of unadvertised services such as HTTP/HTTPS, Telnet, SMB, and so on, where no legitimate internet traffic should be expected. Figure 9 shows the unique, daily connections with the Heisenberg sensor network for all organisations in a given sector. This chart should ideally be blank, but it shows lapses in control across every sector in this data set. Some sectors, such as Consumer Discretionary, appear to have slightly higher systemic rates of control failures, but this view does not tell the whole story, since many modern networks sit behind a single internet address through which hundreds to thousands of employees, contractors, and devices communicate. This chart is handy to show presence, but we need another view to show volume.

Figure 9: Daily Time Series of Unique Connections to Project Heisenberg by Industry Sector




In contrast to the unique connection footprint view, Figure 10 shows the total daily connections to Project Heisenberg across organisations in the measured industry sectors. Note that the Y-axis is not uniform across the panels. This freescale lets us “zoom in” on each industry and more easily distinguish potential patterns and problems. We see that just because an industry has a small number of unique nodes connecting to Heisenberg sensors does not mean they are inactive. Larger volumes in this view could indicate a mass malware infection internal to an organisation (i.e., dozens, hundreds, or thousands of infected systems reaching out to the internet) or may be indicative of a few systems being co-opted into DoS campaigns.

Figure 10: Daily Time Series of Total Connections to Project Heisenberg by Industry Sector

NOTE: Free Y Scales



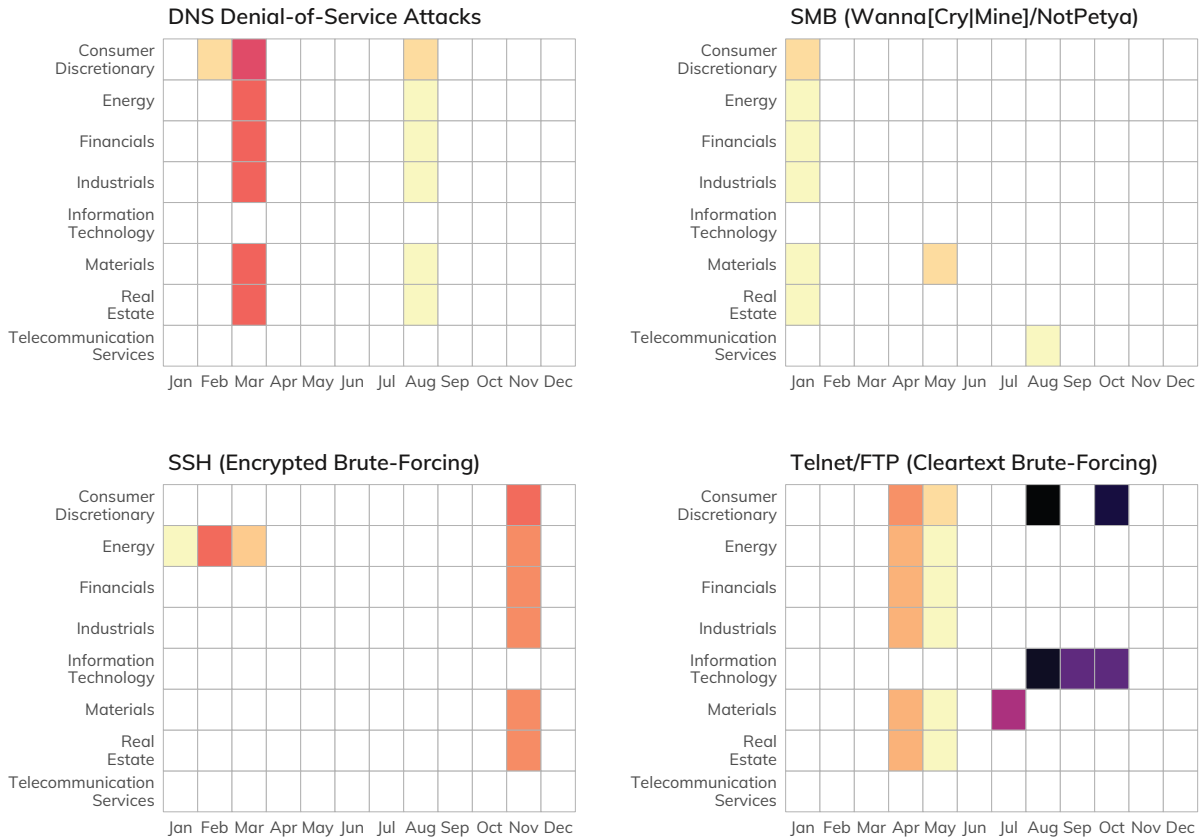

Figure 11: 2018 Heisenberg Connection Ratios by Industry Sector



To further compare industries, we can combine the data from the previous two charts to complete the macro-exposure picture. Figure 11 shows the distribution of connection ratios (total connections in a day / unique sources in a day) by industry sector. Sectors with a greater number of points have organisations with more frequent gaps in configuration control or malware containment. Sectors with points further out on the axis also have gaps in monitoring as well as containment. Consumer Discretionary and Information Technology organisations in this data sample appear to be lagging behind their peers in other sectors.

Figure 12: Signs of Serious Malicious Activity per Sector

Total attacks by identified service (log10 scale)



Some connections are more serious than others, and four of the top connection types to Heisenberg from organisations in this study are especially bad. As Figure 12 shows, throughout the first half of 2018, Heisenberg recorded daily connections that indicate multiple organisations were impacted by the following:

- Malware associated with SMB (i.e., WannaCry, WannaMine, and NotPetya);
- DNS DoS attacks;
- Telnet/FTP cleartext credential brute-forcing; and
- SSH encrypted credential brute-forcing.

Conclusions

The methodology outlined in this report describes several ways, based on openly available internet connections, to measure the exposure of specific organisations and industry sectors to certain cybersecurity risks. To reiterate, while far from a complete picture of the organisations' overall cybersecurity posture, the results of this research indicate significant levels of exposure among ASX 200 companies:

- Companies in the ASX 200 expose an average of 29 servers/devices, with four companies exposing 200–300+ systems/devices. More systems equals more exposure to opportunistic and targeted attacks.
- An overwhelming majority (67%) of ASX 200 companies do not use enhanced email safety configurations, creating a greater risk of phishing attacks. To reduce exposure, organisations should evaluate and strengthen their DMARC configuration settings.
- Critically weak services such as Telnet and Windows file-sharing had only a light presence in the ASX 200 corpus, but they were present. Each instance creates a greater risk of susceptibility to exploitation of SMB vulnerabilities. To reduce exposure, organisations should close port 445 whenever possible and migrate from Telnet to SSH.
- ASX organisations in every sector had serious issues with patch/version management of business-critical internet-facing systems. It is vital that organisations make configuration and patch management of internet-facing systems a top priority.
- Dozens of exposed third-party services are shared among the ASX 200 organisations, creating a greater risk that a vulnerability in a shared third party could lead to the compromise of multiple organisations. Organisations should ensure their third-party service providers are taking appropriate steps to strengthen their own security, as well as use tools such as subresource integrity signatures when sourcing these services to help reduce the likelihood of shared compromise.

Because the ASX 200 organisations typically have more resources and access to technical expertise, the findings suggest that the severity of exposure may be greater for organisations outside the ASX 200. The digital ecosystem could benefit from an ongoing conversation with key stakeholders on the reasons for this continued exposure, along with steps to mitigate the cybersecurity risks it poses.

Measuring Industry Exposure: Methodology Overview

This report documents findings regarding organisations' exposure to certain cybersecurity risks using data made available through interactions with public-facing systems over the internet. That data was then used to quantify the exposure of members of the AUS-based ASX 200, with results aggregated by industry sector. Measuring exposure at this level can help target cyber-risk reduction efforts, improve cybersecurity information-sharing within industry sectors, and build awareness of practices organisations can undertake to avoid future exposure.

Since 2016, Rapid7 has annually measured and reported on the exposure of specific countries to certain cybersecurity risks.¹⁶ With this information, we engage country-level Computer Emergency Response Teams (CERTs) to analyse the exposure in more detail and support action to reduce their overt exposure of critical services. To generate these reports, Rapid7 uses our internet-wide scanning platform, Project Sonar,¹⁷ and our passive sensor network, Project Heisenberg,¹⁸ to determine whether online assets are advertising vulnerable internet services or making suspicious outbound connections. We then aggregate the results at the nation-state level.

Aggregating the exposure data at the nation-state level is relatively straightforward. We use high-quality, regularly updated databases that match country location to internet addresses, with over 98% accuracy.¹⁹ However, it takes additional effort to measure exposure at a deeper level. More robust exposure measurement of specific organisations is possible by analysing the dedicated internet address space that those organisations own and use as part of their business operations. After matching organisations to internet addresses, exposure to certain cybersecurity risks can be quantified through publicly available data obtained with active scans and passive sensors. This section details the steps involved in the following:

- Attributing internet addresses and primary internet domain names to ASX 200 organisations;
- Using Project Sonar's active scan data to identify exposure to vulnerable services and systems within the internet address space attributed to these organisations;
- Enhancing this exposure measurement by identifying the frequency and nature of interactions from this attributed internet address space with Rapid7's Project Heisenberg global passive sensor network;
- Supplementing these direct exposure measurements with inferred exposure. To do this, we analysed metadata from organisations' attributed internet address space, such as email "safety" configurations stored in internet DNS records, and detectable operating system and application version information; and
- Further validating and enhancing the study's data through our research partnership with SecurityTrails, which catalogues, vets, and incorporates data from a wide array of internet scanning and intelligence services into its data warehouse.²⁰

¹⁶ Rapid7, National Exposure Index, June 7, 2018, <https://www.rapid7.com/info/national-exposure-index> (last accessed Feb. 12, 2019).

¹⁷ Rapid7, Project Sonar, <https://www.rapid7.com/research/project-sonar> (last accessed Feb. 12, 2019).

¹⁸ Rapid7, Project Heisenberg, <https://www.rapid7.com/research/project-heisenberg> (last accessed Feb. 12, 2019).

¹⁹ MaxMind, <https://www.maxmind.com> (last accessed Feb. 12, 2019).

²⁰ SecurityTrails, <https://securitytrails.com/#asx200> (last accessed Feb. 11, 2019).

The measurements can be broken down into three primary areas, each of which is covered in the following sections:

- **Inferential measurements** using public DNS records, the most significant of which is the measurement of an organisation’s defences against phishing attacks;
- **Active measurements** using Rapid7’s Project Sonar, which includes measuring both the presence of public-facing systems and services as well as the content those systems and services expose; and
- **Passive measurements** using Rapid7’s Project Heisenberg, which records when systems from an organisation’s network contact this honeypot collection and what actions they were trying to perform during these connections.

Active Measurements with Project Sonar

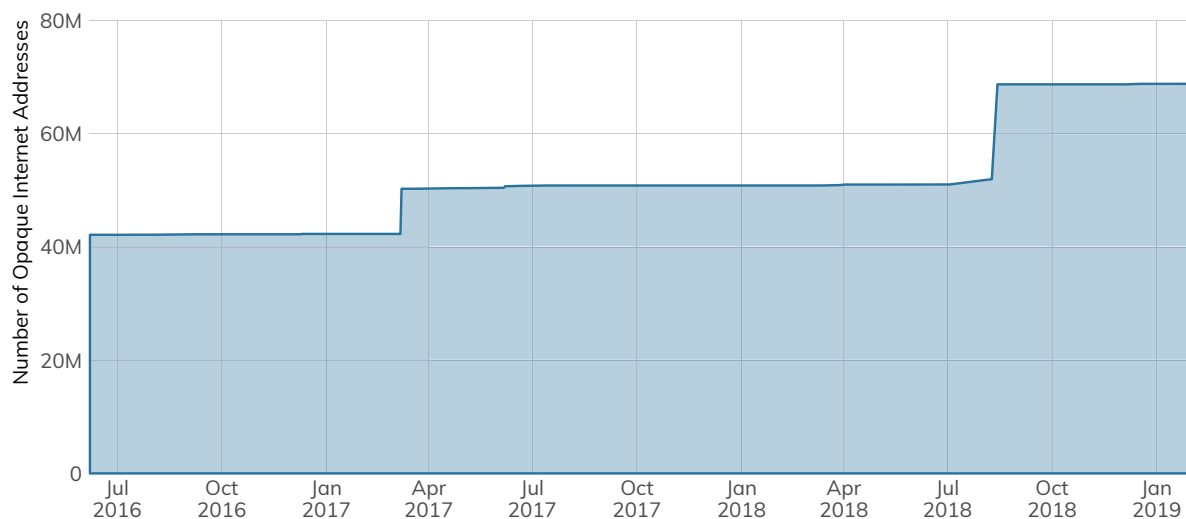
Project Sonar scans the internet across a wide array of services. A “service” could mean a web server, mail server, file server, database server, network equipment, or even cameras, along with many other types of servers that listen for requests over the internet. When a service on a given internet address responds positively to a probe, the positive result is recorded along with the response data. Depending on the service being scanned, this response data can include detailed version and configuration information of the scanned service.

Rapid7 adheres to the legal restrictions associated with internet scanning. As a result, the probes performed by Project Sonar never involve the use of credentials, exploits for known vulnerabilities, or payloads that may cause harm to the service being probed, no matter how obvious or well-known those passwords or exploits may be. While this places some limits on what we can scan and the types of service metadata we can retrieve, we can still capture a wide array of useful information.

A further, self-imposed restriction comes as a result of Rapid7’s “opt-out” process. Organisations may request that Rapid7 exempt specific internet address ranges from Project Sonar scans. Rapid7 obliges these requests and places the address range onto a blacklist that is restricted from the scanning process (Figure 13).

Unlike the 2018 Industry Cyber-Exposure Report (ICER) for the Fortune 500 corpus, there were no networks present on the opt-out blacklist from any organisation in the Q4 2018 ASX 200 list.²¹

Figure 13: Rapid7 Project Sonar Blacklist Growth



²¹ Rapid7, Industry Cyber-Exposure Report: Fortune 500, pgs. 10–11, Dec. 11, 2018, <https://www.rapid7.com/info/industry-cyber-exposure-report-fortune-500>.

Passive Measurements with Project Heisenberg

Rapid7's Project Heisenberg is, at heart, nearly 200 unadvertised systems hosting a variety of fake services, such as HTTP, SMB, and SSH. These honeypots are closely monitored for unsolicited connections but do nothing to attract or entice those connections. Other than internet-wide scanning research, there are no legitimate reasons for an organisation to connect with the Heisenberg sensor network, so any recorded activity in Heisenberg is a high-quality indicator that an organisation does not have control of its outbound connections—which further suggests either malicious activity or misconfigured service traffic coming from the organisation. In essence, if there is any contact with Heisenberg, there is some type of exposure occurring in that organisation.

Measuring Web Server Third-Party Risk

To get an idea of third-party risk when exposing web servers/services to the internet, we can examine the resources each web page loads when the page is loaded into a web browser. Project Sonar can perform this task at scale by controlling a virtual web browser, visiting the pages of the well-known domains of the organisations in the study, and capturing all the activity as each site loads resources.

These websites load a great quantity of third-party resources, so the complete list would be difficult to visualise and comprehend. The resultant list was pared down to only the most prevalent third-party resources used across the target study list.

Further Work

The processes and procedures used for the exposure analyses in this report are the initial steps at communicating the overall “cyber-health” of industries based on a subset of possible internet telemetry measurements. Possible measurement deficiencies have been identified and will be addressed in this section.

Improving Entity Internet Asset Attribution

The most common Internet Protocol (IP) address space (version 4, IPv4) is fully exhausted, meaning there are no “spare” blocks of IP addresses to assign an entity. However, organisations that currently do own IPv4 address space are not utilising said space to capacity. The scarcity of this finite resource has resulted in the creation of a marketplace in which IPv4 space can be bought and sold.²² While some long-standing organisations have sold portions of their IPv4 address space blocks to other parties, some retain ownership and manage the leasing of this space to others on their own. This practice results in attribution errors, which are especially vexing when corporate address space is leased in a non-attributable way to third-party hosting providers and/or cloud providers.

For this report, Rapid7 researchers initially used a manual processes for both preliminary attribution attempts, as well as identification of attribution anomalies, by comparing address space utilisation and service composition with that of known hosting and cloud service providers. As noted in the **Methodology** section (pg. 23), this approach was ultimately tabled in favour of working with directly attributable resources from organisations’ DNS records and inferring organisation-owned IPv4 space from these records. Further work will be performed to shore up IP space attribution and automate this classification, which will enable filtering out hosting and cloud service provider blocks at scale.

Note especially that this DNS record approach differs from that of Rapid7’s inaugural Fortune 500 ICER in 2018.

Avoiding Opt-Out Opacity

Research work like this report depends on continuous, light-touch scanning like the kind provided by Rapid7’s Project Sonar, so if enough organisations decide to opt out of these scans, the internet research community will undoubtedly suffer. There are two future paths that can reduce the impact of the Project Sonar “opt-out” list opacity issue. As a responsible internet citizen, Rapid7 keeps the opt-out list process in place, but it may be possible to augment current processes and have the opt-out be an annual process whereby organisations may re-acknowledge their desire to have their IPv4 space remain on the opt-out list. This would provide an opportunity to restate the advantages of allowing Project Sonar scans, reduce the size of the opt-out list, and preserve the statistical integrity of the surveys.

The second path is to just expand the sample size to cover more industry participants regardless of where their headquarters are. To this end, there are other notable organisation lists—e.g., Inc. 5000, S&P 500, FTSE 100—that can be mined to significantly expand the sample sizes in each industry and reduce the size of the opaque IPv4 address space to (perhaps) less than 1%. The previously noted attribution accuracy and expansion enhancements are key components to ensuring the validity and efficacy of this expansion process.

²² IPv4 Brokers, ARIN IPv4 Market Prices & Transfer Statistics, <https://ipv4brokers.net/arin-ipv4-prices-transfer-statistics/> (last accessed Feb. 12, 2018).

Finding More DNS Records

The initial exposure and email safety analyses utilised the well-known DNS domains of the organisations on the ASX 200 list. Most of those corporations have subsidiaries and brands, each with their own set of DNS domains. Further work will be performed to develop a machine learning-based web-crawling and data-mining process to identify these additional domains and incorporate the data associated with them into the analysis framework used in this report.

Expanding Resource Safety Evaluation

The further work to discover additional domain names will have a direct impact on the email safety analyses used for this report. Furthermore, this report only looked at one aspect of email safety (DMARC). There are additional resource records that describe other email safety configurations, such as Sender Policy Framework (SPF), which further helps reduce spam and prevents attackers from misusing email domains. This will be included in future analyses.

Other types of DNS records (i.e., non-email-related ones) also communicate other types of both exposure and safety. That information will also be explored for inclusion in future analyses.

Third-Party Dependency/Risk Analyses

Finally, by analysing the overall configuration of an organisation's DNS records, discovering how an organisation's IPv4 networks are routed on the internet, enumerating which third-party resources an organisation relies upon in its web and web application services, and conducting other indirect, public measurements, it is possible to report on both the potential fragility of an organisation's overall internet presence and provide exposure views of third-party dependencies across all organisations.

Study Methodology

Why the ASX 200

Aggregating exposure for specific Australasian industry sectors poses a unique problem. First, IP address space is fairly expansive (Australia, in particular, regularly comes in around 10th or 11th place each time autonomous system and BGP records are updated). IPv4 alone supports over 4.2 billion addresses (a portion of which are not assignable), without taking into consideration the exponentially more massive IPv6 space. These addresses are assigned to various governments, companies, and service providers around the world. Second, with the onset of dynamic infrastructure (“the cloud”), it is increasingly common for companies to lease IP address space from other companies to host their services. This makes traditional methods of attributing IP addresses to particular organisations (such as by using the WHOIS lookup tool) incomplete, since the owner of the IP address may not be the owner of the service evaluated for exposure.²³

Instead of attributing IP addresses to companies and filtering by AUS industries, we focus on the Q4 2018 ASX 200 as a representative sample, from which we attribute and filter global IP address space and services hosted on dynamic infrastructure.

The Q4 2018 ASX 200 list was chosen for many reasons. First, it is a diverse (see Table 5) list curated by experts²⁴ who use well-established criteria for selecting firms for inclusion. When revenues are combined, the composite list equates to approximately two-thirds of the AUS GDP, with aggregate employment reaching nearly 2 million individuals globally. Furthermore, over 90% of these organisations are incorporated in Australia, enabling the creation of a AUS-centric view of exposure and the development of potential economic impact models.

Table 5: ASX 200 Sector Counts

INDUSTRY SECTOR	NUMBER OF ORGANISATIONS
Consumer Discretionary	30
Consumer Staples	14
Energy	10
Financials	25
Health Care	14
Industrials	21
Information Technology	15
Materials	39
Real Estate	21
Telecommunication Services	6
Utilities	5

²³ ICANN WHOIS, <https://whois.icann.org/en> (last accessed Feb. 12, 2019).

²⁴ Australian Securities Exchange, <https://www.asx.com.au/> (last accessed Feb. 12, 2019).

Finally, ASX 200 member organisations attract and employ top talent at every level. This includes internal and external network and systems management personnel, as well as highly skilled and experienced application development and operations staff. Many of these organisations have representatives on committees who provide leadership and governance of groups that develop IT and internet standards—a large number of these organisations have been incorporated for over 20 years and were early adopters of internet technologies. In other words, if there are exposure issues in this group of organisations, it may signal that exposure conditions are even more substantial in companies that do not have similar stature.

Organisation Internet Asset and Metadata Attribution Methodology

The Internet Assigned Numbers Authority (IANA) coordinates the governance of key elements that enable smooth operation of the internet. Two key governance elements relevant to the process of attribution include internet address space (or “IP” addresses) and domain names (the system that helps turns web addresses such as <http://www.example.com/> into internet addresses so systems can connect to internet resources).

Attributing Internet Address Space to an Organisation

IANA delegates the management of internet address space to a small number of global and regional internet registries. These registries further delegate blocks of internet addresses and coordinate the metadata associated with these assignments to national and “local” registries that ultimately coordinate with internet service providers (ISPs), which assign internet addresses to users and organisations.

The metadata associated with these internet address assignments, such as the organisation names, location information, points of contact, and potentially the parent internet service provider, is stored in a distributed set of databases called the WHOIS service. The WHOIS service is a public resource that allows a user to retrieve information about an IP number, including the organisation that owns the internet address and the organisation’s point of contact. Each registry maintains its own WHOIS database.²⁵ Individuals can use WHOIS to make interactive queries to these systems, and bulk copies of WHOIS database information are made available to organisations that will use the data for technical research purposes.

When an organisation wishes to manage its own internet-connected resources, it makes a request to a local ISP or local registry and is assigned one or more contiguous sets of addresses to use. This attribution metadata is stored in the appropriate WHOIS service. To illustrate what this looks like, Table 6 shows the internet address block assignments for Rapid7:

Table 6: Rapid7 WHOIS Record Summary

INTERNET ADDRESS ASSIGNMENT	WHOIS ATTRIBUTION
71.6.233.0/24	Rapid7 Labs. Traffic originating from this network is expected and part of Rapid7 Labs Project Sonar sonar.labs.rapid7.com (C07045996)
208.118.237.0/24	Rapid7 LLC (C02934565)

Unlike the Fortune 500 ICER, using the IANA registry methodology to locate ASX 200 company-owned space proved significantly more difficult. First, less than half (35%) of the organisations had identifiable entries in any IPv4 registry. That does not mean the other 65% do not have blocks assigned to them, but it does mean there is an increase in the error rates when attempting to attribute those blocks.

²⁵ APNIC WHOIS Database Index, https://www.apnic.net/about-apnic/whois_search/about/what-is-in-whois/which-whois/ (last accessed Feb. 12, 2019)

Rather than penalise/overweight organisations that do have attributed IPv4 space, this report used DNS records as the initial ground truth of what an organisation has made available on the internet. We used Project Sonar FDNS study data, which we compared to and enhanced with that data from SecurityTrails to ensure we had the widest available coverage.²⁶ Additional enhancement was performed by expanding the definition of an organisation's IPv4 space when DNS address ("A") records clearly indicated they were using a full or partial CIDR block of associated addresses. While not nearly as comprehensive as using full blocks from ARIN records, this data set is fair to all ASX 200 organisations and has clearly identified numerous resources across every organisation with distributions on par to that of the Fortune 500 ICER.

This approach also avoided the serious issue of over-counting IPv4 space that may have been leased to cloud providers and further enabled the inclusion of additional telecommunications providers, since it was trivial to separate customer/carrier space from their own enterprise space from DNS records.

Attributing DNS Records to an Organisation

A similar WHOIS registration and database service exists for DNS assignments, except this is a far more distributed service that places direct control of all the underlying records for a domain into the hands of an organisation. Once assigned a domain name (e.g., "rapid7.com"), an organisation sets up its own DNS server (or uses one from a DNS service provider or cloud provider), then publishes and maintains records that map DNS names to a wide array of record types and values. Organisations can add, change, or delete records at will.

DNS "A" (address) records map names to internet addresses (e.g., <www.rapid7.com> currently maps to 13.33.37.212), but it is also possible to associate other types of information with an internet name.

DNS "TXT" (text) records facilitate storing arbitrary text strings with internet names. A number of formal standards exist that provide rules for crafting specially formatted text records to convey additional metadata about that internet name resource or the domain name's proper owner.

DMARC²⁷ and the SPF²⁸ are two key TXT records for inferring the "safety" of an organisation's email configuration. These standards enable an organisation to communicate which systems are authorised to send mail on its behalf and what should be done with forged emails sent by attackers or spammers. Missing, improperly configured, or overly permissive configurations of these records put organisations at risk for both increased spam and phishing attacks. Since phishing attacks have been the primary means of attackers gaining a foothold within an organisation for the past few years, lack of care and attention to appropriate DMARC and SPF configuration significantly increases the likelihood of successful attacks against that organisation. Anyone can query the DNS for these and other records. As part of our research efforts into ecosystem-wide cybersecurity, Rapid7 performs millions of DNS lookups every month and stores the time-stamped record results in a large, historical database, which makes it possible to perform large-scale queries and track changes over time.

The Q4 2018 ASX 200 list includes the primary, well-known domain names of the members of the list. For example, "amp.com.au" is the well-known domain for AMP (a large financial services organisation), and while "ampaws.com.au" is also owned by AMP, this domain isn't expected to handle email, nor does it have a valid email configuration. These sites were systematically scanned by Project Sonar, and the associated DNS names for the attributed organisations were used to determine the presence of DMARC and SPF.

²⁶ SecurityTrails, <https://securitytrails.com/#asx200> (last accessed Feb. 12, 2019).

²⁷ The DMARC Standard, <https://dmarc.org/> (last accessed Feb. 12, 2019).

²⁸ The SPF Standard, April 26, 2014, <http://www.openspf.org/>.

About Rapid7

Rapid7 (Nasdaq: RPD) is advancing security with visibility, analytics, and automation delivered through our Insight cloud. Our solutions simplify the complex, allowing security teams to work more effectively with IT and development to reduce vulnerabilities, monitor for malicious behavior, investigate and shut down attacks, and automate routine tasks. Over 7,200 customers rely on Rapid7 technology, services, and research to improve security outcomes and securely advance their organisations. For more information, visit our website, check out our blog, or follow us on Twitter.

QUESTIONS?

Email us at research@rapid7.com