RAPID7

RESEARCH

Industry Cyber Exposure Report (ICER): Nikkei 225

TABLE OF CONTENTS

Executive Summary 4

Key Takeaways 5

Email Security Among The Nikkei 225 6

Results 8

By Industry 9

CISO Takeaways 9

Web Service Security Among the Nikkei 225 10

HTTPS Support 11

HSTS Adoption 12

Summary 13

CISO Takeaways 13

Version Complexity Among the Nikkei 225 14

Version Dispersion Among Web Servers 16

Version Dispersion: Focus on Microsoft Exchange 17

CISO Takeaways 18

High-Risk Services Among the Nikkei 225 22

Findings: RDP, SMB, and Telnet 24

Windows Remote Desktop Protocol (RDP) 24

Windows Server Message Block (SMB) 25

Telnet 27

Exposure Overview 28

CISO Takeaways 29

TABLE OF CONTENTS

Vulnerability Disclosure Programs Among the Nikkei 225 30

Results: Prevalence of VDP Adoption 32

CISO Takeaways 33

Conclusion 35

CISO Actions at a Glance 36

Appendix: Prioritization in Times of Crisis 38

Executive Summary

As the world's knowledge workers were driven home amid a pandemic and cases of ransomware ran rampant across the internet, measuring the world's most critical businesses' internet exposure has become more important than ever. In this round of Internet Cyber-Exposure Reports (ICERs), researchers at Rapid7 evaluate 5 areas of cybersecurity that are both critical to secure to continue doing business on and across the internet, and are squarely in the power of CISOs, their IT security staffs, and their internal business partners to address.

These five facets of internet-facing cyber-exposure and risk include:

1. Authenticated email origination and handling (DMARC)
2. Encryption standards for public web applications (HTTPS & HSTS)
3. Version management for web servers and email servers (focusing on IIS, nginx, Apache, and Exchange)
4. Risky protocols unsuitable for the internet (RDP, SMB, and Telnet)
5. The proliferation of vulnerability disclosure programs (VDPs)

In this report, we examine the internet-facing cyber-exposure of the top companies listed on Japan's Nikkei 225[1]. Each section is accompanied by real-world, practical advice that practitioners can start implementing today. Note that this advice is not only for those CISOs who are privileged to hold positions in Nikkei 225 companies, but also for those security experts who find themselves in business and regulatory relationships with members of this prestigious collection of corporations.

Through the first half of 2021, Rapid7 will be releasing reports measuring these five critical areas of cybersecurity fundamentals across five of the most advanced economies of the world:

1. The United States Fortune 500[2]
2. The United Kingdom's FTSE 350[3]
3. Australia's ASX 200[4]
4. Germany's Deutsche Börse Prime Standard 314[5]
5. Japan's Nikkei 225 (this report)

[1] https://indexes.nikkei.co.jp/en/nkave/index?type=index
[2] https://www.rapid7.com/research/report/2021-industry-cyber-exposure-report/
[3] https://www.rapid7.com/research/reports/2021-industry-cyber-exposure-report-uk
[4] https://www.rapid7.com/research/report/2021-industry-cyber-exposure-report-anz/
[5] https://www.rapid7.com/research/report/icer-germany-2021/

Key Takeaways

The paper is divided into five detailed sections covering the areas mentioned above, and the overall takeaways of this research are as follows:

- **Nikkei 225 email security posture is lagging behind the US and UK.** At the beginning of 2021, email security among the Nikkei 225 isn't keeping pace with its peers in the US and UK. While DMARC adoption in the US and UK hovers around 50%, only about 13% of all the surveyed companies operating in Japan have any DMARC records configured, and of those, 25 out of 29 (about 86%) are set with a p=none (or passthrough) policy. In other words, only 4 (under 2%) of the Nikkei 225-listed companies are taking active measures to protect their brands, employees, and customers through DMARC p=quarantine or p=reject policies.

- **Exposed, dangerous services are less of a concern in Japan.** While dangerous protocol exposures of Windows Remote Desktop (RDP) file-sharing (SMB), and Telnet continue to be an issue across the surveyed companies, it does not appear to be nearly as much of a problem as we've seen among the U.S.-based Fortune 500: For RDP and SMB, over 90% of the Nikkei 225 had no exposure.

- **Telnet and HSTS remain concerning, however.** Telnet is a different story; about 27% of the Nikkei 225 has some legacy telnet exposed to the internet. Additionally, when we looked at secure HTTP (HTTPS) deployment, we found that while HTTPS is standard for 100% of the Nikkei 225 companies, very few listed companies (18%) have implemented HSTS directives to ensure that HTTPS infrastructure is actually being used all the time.

- **Version dispersion is on the right track in Japan.** Only 16 companies in the Nikkei are running their own Exchange servers (rather than managed cloud instances), and of these, about 75% are running at least 1 instance of the latest supported version. That said, we did count 93 distinct versions of Apache, 75 distinct versions of Nginx, and 17 distinct versions of IIS in the Nikkei 225.

- **The Japanese Technology sector stands alone in vulnerability disclosure.** Nearly all of the 16 VDPs we found across the 225 surveyed companies are either in the Technology sector proper, or in tech-heavy Consumer Goods companies. So, while this is pretty good for Japanese tech, it's not great for the rest of the Japanese businesses that have not normalized VDPs for their products and infrastructure.

With these key findings in mind, the remainder of this report explores each of the 5 areas of cybersecurity measurable in the Nikkei 225.

Before you dive in, we want to note that if your organization was and/or still is impacted by those events, you may be feeling like you are spending most of your time and energy dealing with emergencies rather than being able to focus on some of the more chronic issues outlined in this report. Since our goal is to help organizations become (and remain) safe and resilient, we have an appendix just for you. Consider jumping there first before tackling the sections below.

Email Security Among the Nikkei 225

We all know and love—or at least begrudgingly rely upon—email. It is a pillar of modern communications, but is unfortunately also highly susceptible to being leveraged as a mechanism for malicious actions, such as spoofing or phishing.

A core concern regarding email is the authenticity of the source, and in recent years, DMARC has arisen as the preeminent email validation system. DMARC builds upon the foundations of 2 older email authentication systems: Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM), which respectively check for mail-server authorization ("Is the sender authorized?") and email integrity based on key signatures ("Was the content altered?"). The various components of DMARC can serve to mitigate direct threats as well as potential reputational damage, such as spoofed emails intended to mislead partners, suppliers, or customers.

A properly implemented DMARC system can identify illegitimate emails and define how they should be handled. DMARC can be configured to handle emails of suspect provenance with different degrees of severity, depending on the aggressiveness of IT administrators. The DMARC policy options include:

- `None`, where suspect emails are reported to a designated email address that serves to monitor DMARC notifications.
- `Quarantine`, where suspect emails are punted to the spam folder and a report of its receipt is delivered to the monitoring email address.
- `Reject`, where in addition to notifying the monitoring email address, suspect emails are not delivered at all.

By virtue of its efficacy in mitigating malicious messaging via email, we consider DMARC a significant risk mitigator and highly recommend its implementation. Unfortunately, while the benefits of DMARC are profound, its implementation is not global. DMARC's implementations are tracked in public Domain Name System (DNS) records. To determine whether an organization utilizes DMARC only requires the examination of the organization's published DMARC record. We are able to discern the scale and types of DMARC implementations by comparing the primary, well-known domains of the Nikkei organizations against their corresponding DMARC records that appear alongside DNS.

Note that for the scope of this study, we focus primarily on the apex domains of organizations, and do not explore additional domains owned by particular organizations. We elected this approach because there can be significant variation in domain set ownership by organization. By focusing on apex domains, we are in effect treating it as a bellwether indicator of an organization's overall email security posture. After all, if an organization fails to implement DMARC on a primary domain, how confident should we be that the organization practices healthy email hygiene across far less-prominent domains?

These published DMARC records are intended to be highly accessible. They are the means through which email recipients determine how to validate emails using DMARC, what email address to notify when receiving emails that fail DMARC validation, and what DMARC policy to apply in handling invalid emails.

Results

We found that 29 (or approximately 13%) of the Nikkei 225 set of organizations had implementations of DMARC for their primary domains, all of which were validly formatted. Of the set of national indexes that we have examined so far in the ICER series, this is a remarkably low level of DMARC coverage in comparison.

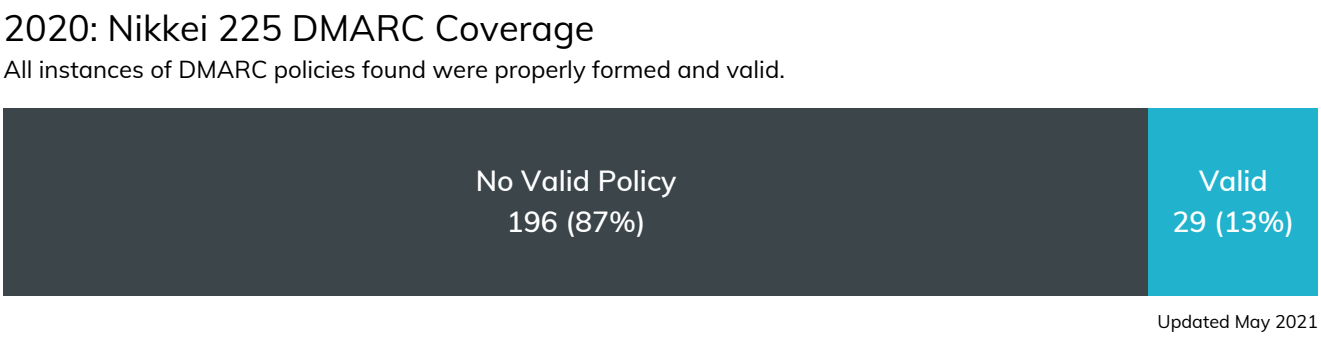

Figure 1: 2020 Nikkei DMARC Coverage

When we examine the DMARC policies in a bit more detail, we find that most valid DMARC policies are set to "none", or simply to monitor and inform, followed by "reject", which is the most aggressive approach. The least prominent policy implementation is "quarantine", a policy to isolate suspect emails. That being said, the numbers are all fairly small, so attempts to draw any sort of pattern would be meaningless.

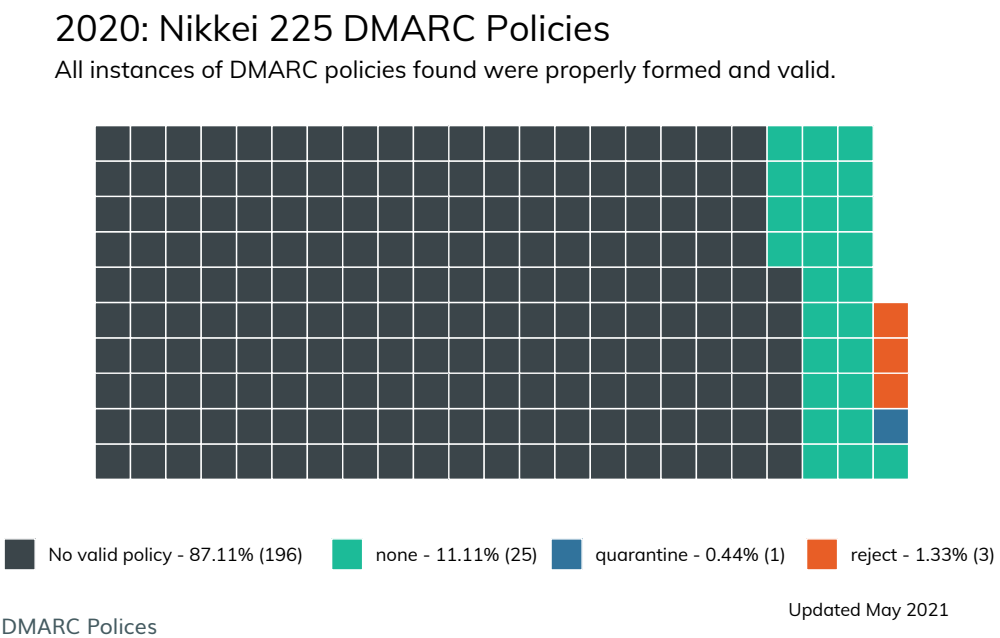

Figure 2: 2020 Nikkei DMARC Polices

By Industry

We can also separate the organizations by industry to get a better sense of DMARC variations across the sectors. The most prominently featured industries in the Nikkei 225 include technology and materials.

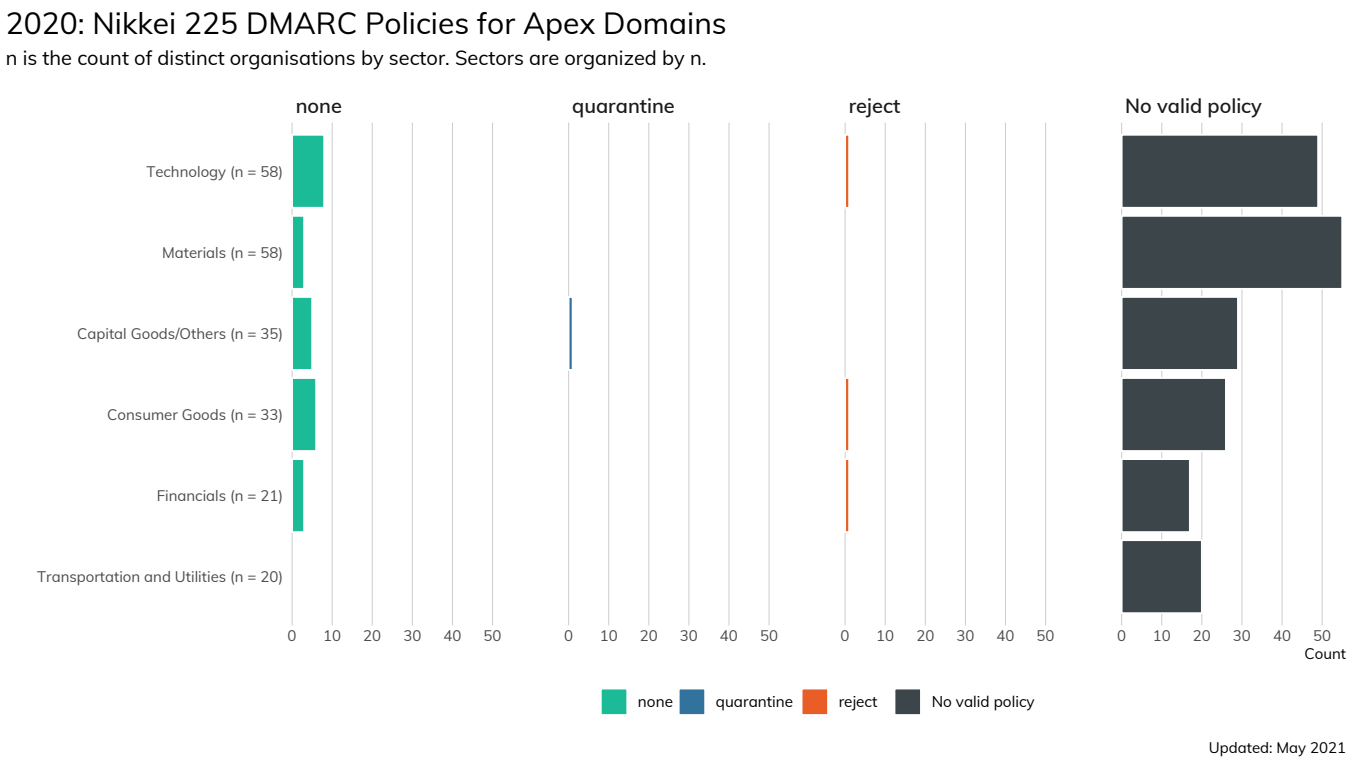

Figure 3: 2020 Nikkei DMARC Polices for Apex Domains

CISO Takeaways

If DMARC has not already been implemented in your organization, take proactive measures to get it set up.

Nowadays, DMARC can be thought of as a foundational fixture of email hygiene, and it broadly signals an organization's commitment to modern information security norms. Furthermore, lacking a DMARC implementation leaves an organization potentially blind to malicious email campaigns that are not captured through some form of DMARC monitoring that can be informative in terms of scale, source, and severity.

Once the decision has been made to implement DMARC, it's time to consider the policy implementation in a more nuanced manner. An aggressive reject policy is highly secure but might result in legitimate emails being blocked. A more forgiving quarantine policy could strike a balance between preventing aggravation and allowing for some form of recourse. At the very minimum, a DMARC implementation of some form should be in place to monitor for illegitimate or poorly configured email traffic.

Web Service Security Among the Nikkei 225

The vast majority of the interactions an average person has with technology is through some form of a web application, but what constitutes a "web app" can be considered quite nebulous, and the security controls for hardening these applications are equally broad. APIs, distributed authentication schemes, single-page applications, and static websites all might fall under the general category of "web application." There are very few security measures that should be applied to all web applications across the board without further subdividing what specific type of application we are referring to. However, there are a couple that we will examine here.

All web applications should require strong encryption, with a vanishingly small number of exceptions. While this is most critical for applications serving up critical or sensitive information, such as personally identifiable information (PII), it is important even if you serve only static informational content. There is a common misconception that the only risk of using an insecure connection is a loss of confidentiality—that the information a user is browsing could be observed by a malicious third party. While this certainly is a risk, it is often overlooked that a lack of encryption makes the connection vulnerable to modification (a loss of integrity). This means that malicious third parties could not only observe potentially confidential information, but that they could alter that information or inject their own content that could potentially compromise your users.

The risk of malicious content injection exists regardless of whether your web application serves sensitive information or just cute pictures of cats[6]. Due to this universal risk to a site's users and to the overarching brand reputation of the site owner, we will consider the support of strong encryption (in our case, TLS) and the enforcement of its usage via HTTP Strict Transport Security (HSTS). For the purposes of this section, we will look at the primary domain for each company, as it is the domain that is most responsible for a company's brand reputation.

HTTPS Support

HTTPS is the protocol that ensures web traffic is encrypted and secure. There are a few ways that HTTPS could be configured in an environment.

- Not available (HTTP only)
- Available and optional
- Required (HTTP "Strict Transport Security", or HSTS, configured)
- Required with HSTS preloading

Supporting HTTPS for your site is table stakes for having a web presence at all, with requiring encryption following very closely behind. HSTS preloading does carry some technological challenges, but they are challenges that a web security program should be working to proactively address.

With all this said, let's share some good news right off the bat: Among the sites we examined in the Nikkei 225, 100% of them supported HTTPS.

[6] https://www.sanrio.co.jp/

HSTS Adoption

The outlook for HSTS adoption was unfortunately a bit grim.

As you can see, only about 18% of the sites examined supported HSTS at all. This is substantially less than what we have observed in other reports. If the site already fully supports HTTPS (and these sites all do), it should be relatively simple to implement HSTS to guarantee your users visit the secure version of your site. Most of these sites do provide a redirect from the insecure version of their homepage—however, that will not mitigate a man-in-the-middle (MiTM) attack.

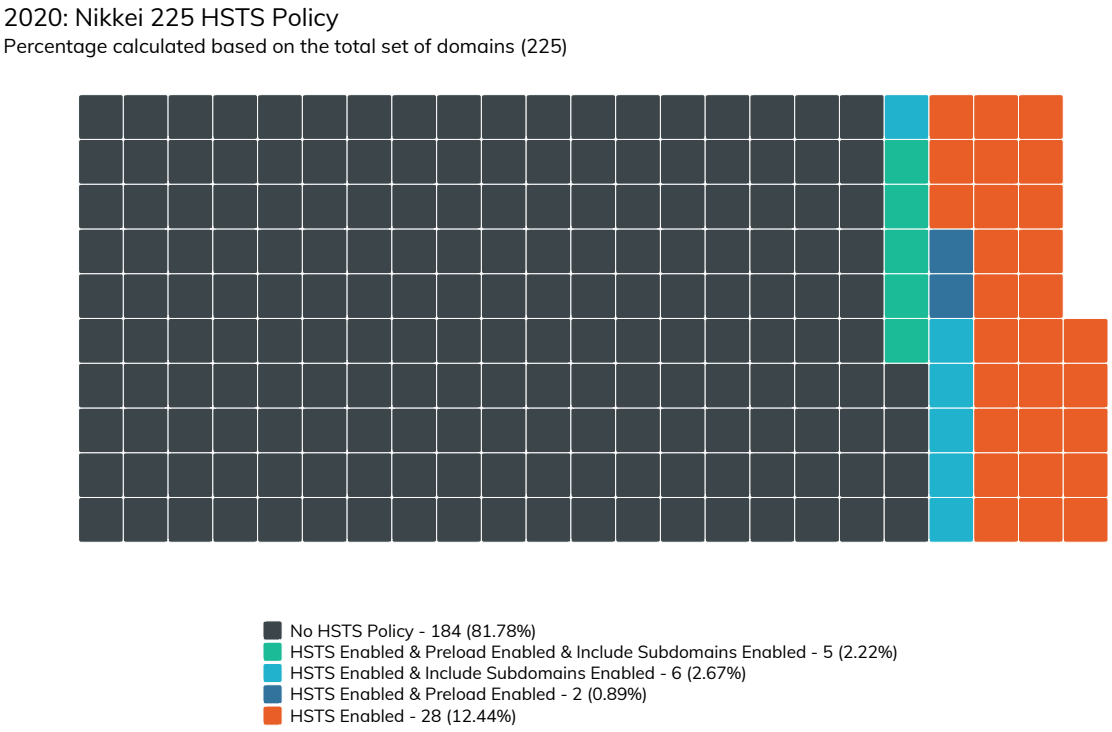

Figure 4: 2020 Nikkei 225 HSTS Policy

None of the observed domains have HSTS manually disabled. The percentage of domains with this configuration tends to be low, so this observation is likely due to the low total number of HSTS supporting domains in this list. 27% of sites that support HSTS also support the "includeSubDomains" directive, protecting the entire domain and all subdomains. This is a fantastic security feature, but it can be difficult to implement in certain situations.

17% of sites with HSTS also support the "preload" directive. This directive will cause crawlers to automatically add your site to a global list of known sites that support HSTS. If a supporting browser is directed to a site with HSTS preload enabled, it will guarantee that the first connection is always conducted over HTTPS, meaning it eliminates the one, single place where your site's users are vulnerable to MiTM attacks—the first connection to your site before an HSTS header has ever been encountered. This configuration option is a simple way to add an extra layer of protection for your users, and if you bother to enable HSTS, you should certainly add this option. While it's a somewhat newer directive with less browser support, there is no downside to including it (browsers that do not support HSTS will simply ignore it).

Summary

Securing and encrypting traffic to your user-facing domains is not only good practice, but it also protects your corporate brand. Securing HTTP with TLS has been a major point of focus for the web-security community for the past several years, and for good reason. All of the Nikkei 225 companies provided a secure version of their primary website, but they have a long way to go before they come up to snuff in terms of best practices.

The especially poor adoption of HSTS across the Nikkei 225 could be an indicator that their application security programmes are falling behind, especially since other, more sophisticated, mitigations can be significantly more complicated to implement. While the standards certainly move quickly, it's important to keep up to speed, especially when your brand reputation is on the line.

CISO Takeaways

If you haven't thought about your site's encryption for a while, now might be the time to revisit it. A company's brand reputation is on the line when consumer-facing web applications suffer from security failures, and it's important to consider this fact when making investment decisions in various security programmes.

If your company's website is not supporting HSTS, it might be worthwhile to find out why. Is it a technical, organizational, or budgetary constraint? Finding the cause could be a great springboard for re-evaluating your entire application security program.

Version Complexity Among the Nikkei 225

Complexity is the enemy when it comes to successful security outcomes in an organization. Diversity in systems, technologies, and business processes present real, daily challenges for even the most mature security teams, especially when it comes to patch and vulnerability management.

Patching even 1 major vulnerability can be a Herculean task in many places. Diversity compounds complexity within each technology component. That is to say, an organization may have many different web server technologies in use. Each technology, in turn, may have its own hodgepodge of versions, which directly (negatively) impacts configuration management and patch management.

To get a feel for how these well-resourced organizations are performing in this area, we looked at 3 separate factors:

1. The diversity of the portfolio of a selected technology—web servers—in use by each organization
2. How well maintained this portfolio is
3. How well organizations maintain critical services, such as email gateways

Our findings show that:

- Within a single technology stack (web servers), organizations in some key industries—Capital Goods/Others, Consumer Goods, Technology, and Transportation and Utilities—**expose 9 or more different versions of Apache and/or Nginx.** Capital Goods/Others, Consumer Goods, Financials, Materials, Technology, and Transportation and Utilities have 1 or more members exposing 3 or more different versions of IIS. **This increases their respective attack surfaces** and makes it difficult to deploy patches (when they bother to apply patches) due to testing and quality assurance complexity.

- Some organizations have **serious difficulty keeping critical IT infrastructure**—such as Microsoft Exchange—**current**. Impressively, around 75% (12 out of 16) of Nikkei 225 that still run self-hosted Microsoft Exchange are running at least 1 current/supported version of Exchange. On the flip side, about 56% (9 out of 16) are running at least one end-of-life version of Exchange 2010, putting them at **risk of future vulnerability exploitation**.[7]

We used Project Sonar[8] and Recog[9] to identify internet-facing technologies—e.g., web servers, file servers, DNS, SSH, etc.—that were in use for each organization in the Nikkei 225. We then mapped them to available Common Platform Enumeration[10] (CPE) strings. This methodology has some limitations in that the results are constrained by:

- The fingerprints available to Recog
- How promiscuous each fingerprintable service is (i.e., whether Recog can extract version information)

---

[7] This adds up to over 100%, because it's possible to run both a current and EOL'ed Exchange installation in the same organization.
[8] https://www.rapid7.com/research/project-sonar
[9] https://github.com/rapid7/recog
[10] Common Platform Enumeration definition and database: https://nvd.nist.gov/products/cpe

- The ports and protocols Project Sonar studies
- Our measurement of only IPv4-space
- Sonar honouring IPv4 opt-out requests

These constraints, if anything, generally result in underreporting of the magnitude of the findings.

Version Dispersion Among Web Servers

Back in 2018, when we began our first foray into analysing the cyber-exposure of the Nikkei 225, we created the term "version dispersion" to refer to the diversity of versions within a service component an individual organization was exposing to the internet. With the dramatic rise[11] in enterprise use of tooling such as Kubernetes[12], we expected to see a reduction in version dispersion of the 3 web servers—IIS, Apache, and Nginx—that we previously measured.

There are at least more than 93 distinct versions of Apache, 75 distinct versions of Nginx[13], and 17—yes, 17—versions of IIS[14] running across the entire set of Nikkei 225 members. Let's see how that stacks up by industry.

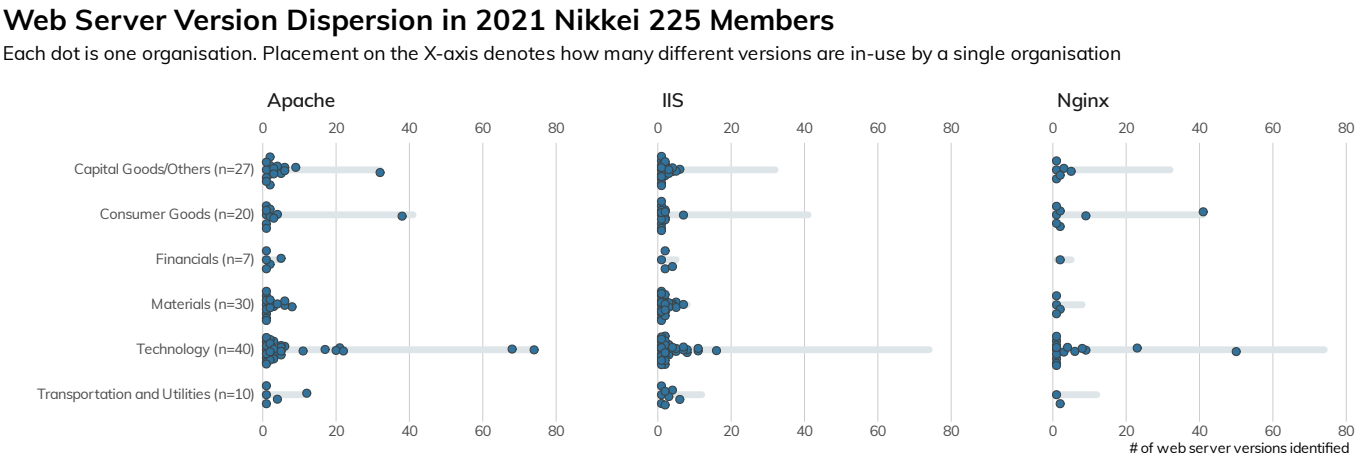

Figure 5: Web Server Version Dispersion in 2021 Nikkei 225 Members

A higher density of points toward "1" on the X-axis means that each of the organizations those points represent are running with a low version dispersion. This means they have better control over server/service deployments and configurations, have fewer versions to test patches against, and can make changes faster and with more confidence than others. It also likely means they have a more rigorous "you must be this tall to deploy a server on the internet" rules than organizations that are further to the right on the X-axis.

[11] A Cloud Native Computing Foundation 2019 survey notes **78% of respondents are using Kubernetes in production, a huge jump from 58% in 2018**

[12] Kubernetes main site: **https://kubernetes.io**

[13] Some organizations announce they use a particular server type but redact the discrete version number.

[14] We frequently see leaking of IIS build strings in announced Server header banners in IIS deployments.

Attackers and cyber-insurance assessors alike notice such things and may be more likely to target organizations that exhibit a more "wild, wild west" stature. There is a striking difference between web server version dispersion in the Nikkei 225 vs what we've reported in the FTSE 350 and Fortune 500 ICERs. One reason for this is that companies listed in the Nikkei seem to have a preference for "the cloud," possibly to ensure faster global connectivity to the information or services provided by the web services they expose. We do not measure "cloud" assets in the ICERs, so these positive results come with said additional caveat.

Version Dispersion: Focus on Microsoft Exchange

Some internet-facing services are more important than others. It's one thing to have a crusty old Apache HTTPD server attached to the internet, which may only have a denial-of-service weakness. It is quite another thing to run old versions of what most organizations would (or should) deem critical infrastructure, such as Microsoft Exchange servers or VPN/gateway/remote-access services.

To get a feel for how well these organizations maintain critical services, we'll take a peek at Microsoft Exchange hygiene. Unlike their Fortune 500 counterparts, only 6% of Nikkei 225 organizations still[15] have at least 1 internet-facing Exchange server handling business-critical email, and Exchange has had a fair number of weaknesses—of varying criticality—uncovered over the years:

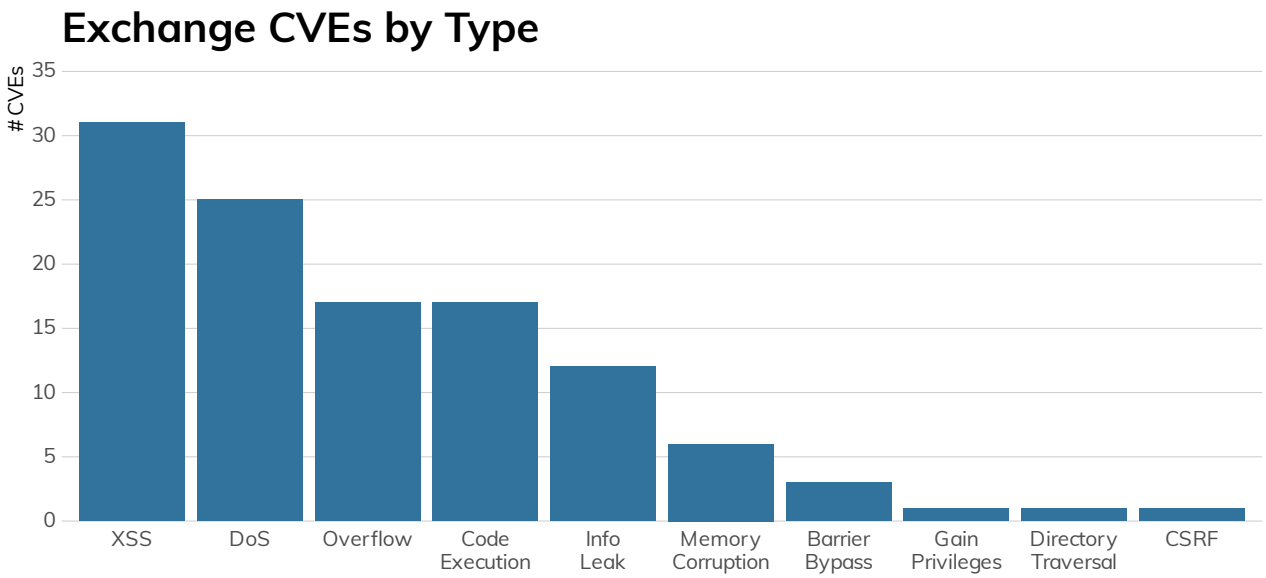

Figure 6: Exchange CVEs by Type

16 organizations (excluding 2 ISPs that allow general service hosting) have chosen to go it on their own when it comes to email hosting, so surely they know the dangers facing self-hosted Exchange and take care to ensure this vital service is at peak resiliency, at least when it comes to security patches. Right?

[15] Microsoft 365/Office 365 adoption continues to grow at a significant clip, with 70% of the Fortune 500 using one or more services, including hosted Exchange. Source: https://www.thexyz.com/blog/microsoft-office-365-usage-statistics/

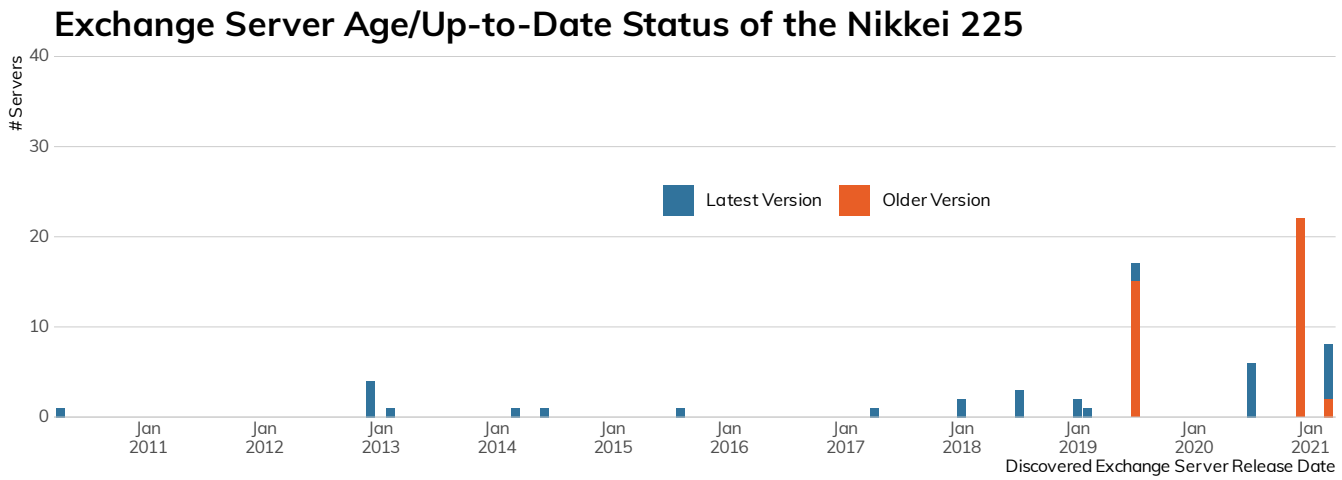

Figure 8: Exchange Server Age/Up-to-date Status of the Nikkei 225

The above figure paints a fairly disturbing picture of the state of Microsoft Exchange in the Nikkei 225 in both currency (i.e., age of some server versions) and whether the deployed version is supported[16] by standard Microsoft support contracts[17]. On the plus side, just over 50% of discovered, precise-version fingerprinted instances are 2020/2021 releases.

We were hoping to see a repeat of our ASX 200 ICER findings with zero presence of Exchange 2007 (which has been at end-of-life status for a while). Sadly, our hopes were crushed as we spotted a single instance at a major Technology company. Additionally, a few handfuls of the Nikkei 225 did not seem to get the memo[18] about Exchange 2010 reaching end-of-life status in October 2020.

[16] https://docs.microsoft.com/en-us/exchange/new-features/build-numbers-and-release-dates

[17] This does not take into account the fact that an organization may have a custom or extended support agreement with Microsoft, though that matters little when it comes to vulnerability exploitation.

[18] https://docs.microsoft.com/en-us/microsoft-365/enterprise/exchange-2010-end-of-support?view=o365-worldwide

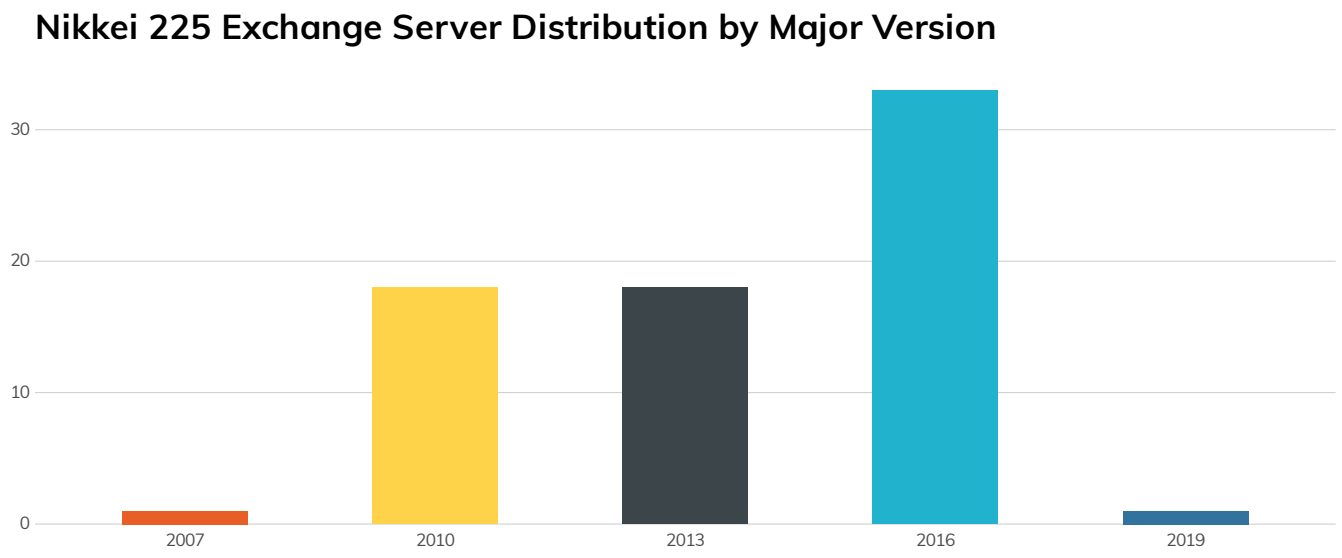

Figure 9: Nikkei 225 Exchange Server Distribution Major Version

If your organization is struggling to keep up with Exchange patching, you may have a bit of wiggle room when it comes to excuses since Microsoft does keep you busy, as seen in the volume of in-year updates for at least modern versions of Exchange:

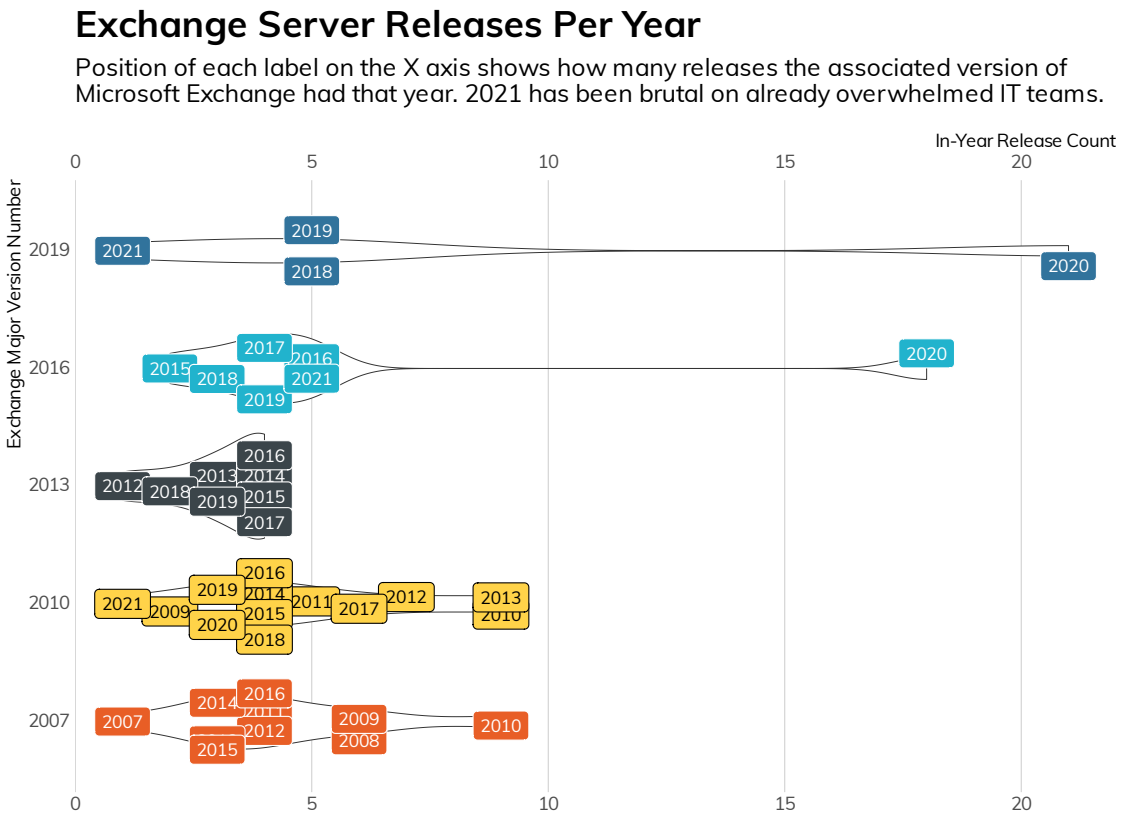

Figure 10: Exchange Server Releases Per Year

And, the outlook is still pretty grim across industries.[19] Figure 11 shows release and support status of Exchange deployments in each industry, and virtually all of them are having trouble keeping current.

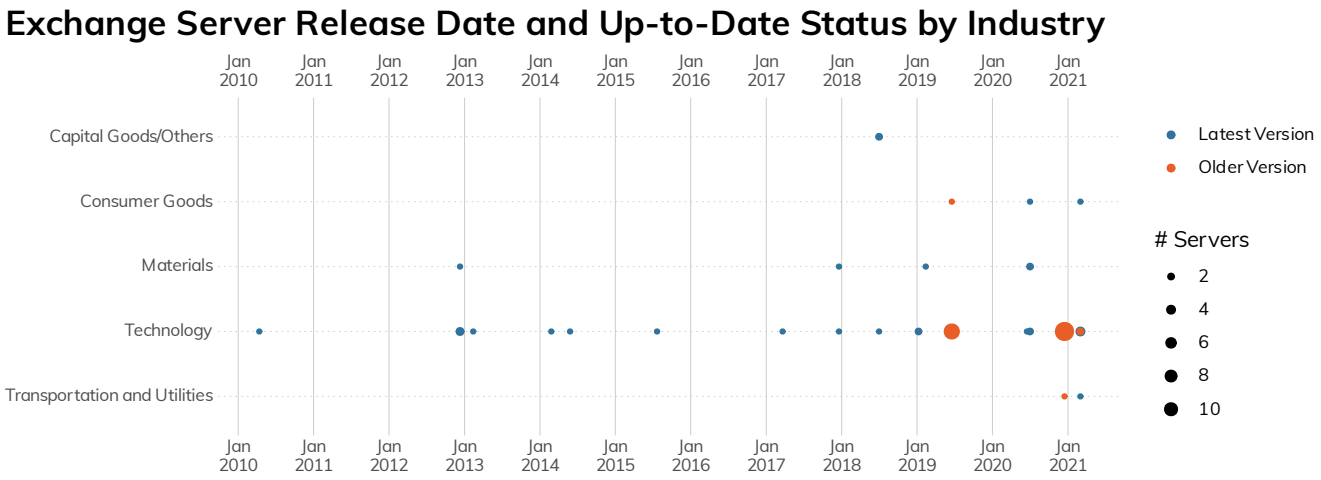

Figure 11: Exchange Server Release Date and Up-to-date Status by Industry

If keeping Exchange deployments updated, secure, and resilient is a challenge for you, take some comfort in the fact that even Microsoft has issues normalizing hosted Exchange (Microsoft 365) build levels, though this chart is far less shameful than the December 2020 snapshot used in the Fortune 500 ICER where their most current deployments were firmly "stuck in the middle" of the chart, with an almost equal number of dispersed versions lingering on the internet's edges.

[19] Yes, we took the obvious pun.

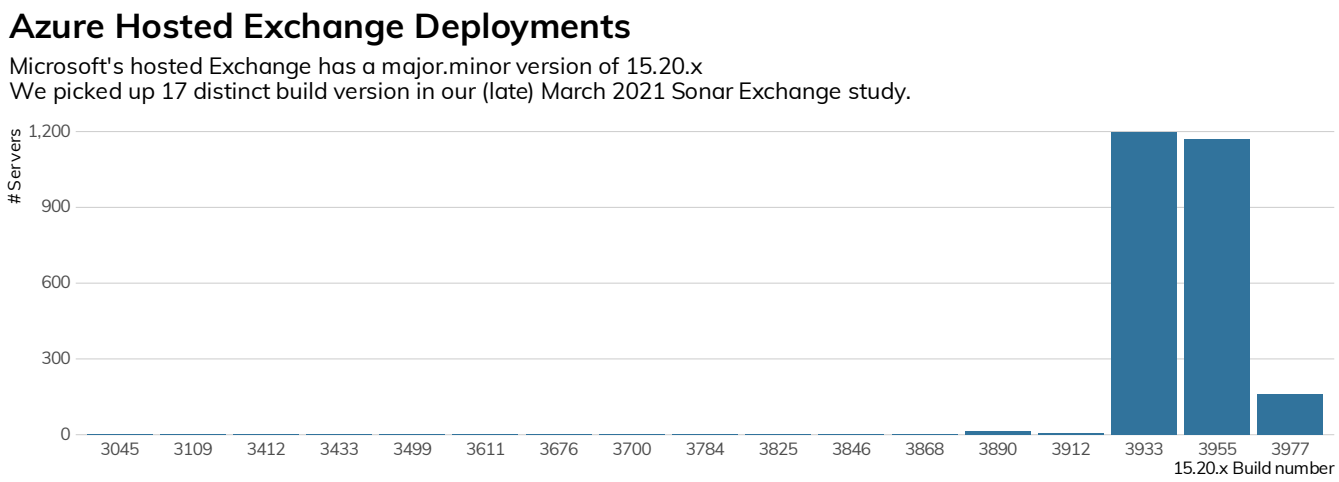

Figure 12: Azure Hosted Exchange Deployments

CISO Takeaways

For this chapter, we'll be talking to 2 different sets of CISOs: those who see their image reflected in the mirrors in each of the sections, and those who have organizations like this as business partners or suppliers.

If you're a security leader who is working to build resilience and safety into the DNA of your organization, issues such as technology sprawl, version management, and critical-service maintenance are non-negotiable must-haves. The good news is that these aren't just "security" issues. Organizations deploy services to meet a business need, and it is far easier to sustain service uptime and stability if there are fewer moving parts to maintain. To achieve buy-in with your peers, collect historical and current data regarding service degradation (and/or outages). Add to that data how long it takes IT, application, and operations teams to support each component of each business process. If you pair that up with information on the volume and severity of identified weaknesses (CVE-based or otherwise), you will find areas that have a solid business case to warrant partnering for improvement. As each area ameliorates, you'll have far more agency to affect change in other, lagging areas.

For those who shuddered at what this section revealed, make sure these are areas you look for when evaluating third parties on behalf of business-process stakeholders in your organization. It's fairly straightforward to both ensure you're asking about these potential areas of weakness and verifying[20] that the answers you receive are accurate. There's no guarantee that the internal exposure of organizations reflects what is seen externally. However, it is generally more likely that the internal picture is even worse than what is presented to the outside world. Holding your partners and suppliers to a higher level of safety and resilience will not only lessen the risk to your organization, but can also have a cascading positive effect as other organizations follow the standards you're setting.

[20] For free, even! https://opendata.rapid7.com/

High-Risk Services Among the Nikkei 225

There are certain services that are generally considered to be high-risk when found available on the public internet. For example, with very few exceptions[21], placing SMB file shares on the internet is considered a Bad Thing. Doing so may expose data, leak environmental information such as domain names, enable brute-force attacks against credentials, and provide a vector for exploiting vulnerabilities in the Windows Server Message Block (SMB) implementation, as was seen in the Conficker[22] and WannaCry[23] worms.

In our research across the public internet, we know that we're only seeing a surface level of information, and we often try to find ways to understand what it is telling us about the organisations that operate these services. We can look at configuration and protocol details and use them as proxy markers for the internal environment and security maturity of an organisation.

For example, if we discover an SMB service and can detect that it doesn't support SMBv2[24], which was introduced in Windows Vista[25] and Server 2008, we can make certain assumptions about the age of the operating system and/or requirements for legacy compatibility.

If an organisation permits Telnet[26] connections to routers from a different country, we can make assumptions about the age of the equipment as well as the security policies for secure protocols and network access control lists (ACLs).

In order to get a sense of how well the Nikkei 225 organisations were performing in this area, we surveyed SMB, Windows Remote Desktop Protocol (RDP), and Telnet on the default ports in their public IPv4 address space and reviewed service data where present.

Our findings show that:

- There is still significant usage of Telnet-based control for routers and switches.
- Of those hosts exposing SMB, all leaked the SMB hostname, DNS name, and fully qualified domain name (FQDN) configured on the host.
- 89 RDP services were found across 18 companies. These were heavily skewed toward the Consumer Goods industry vertical due to the outsized impact of one company.

We used Project Sonar and Recog to identify internet-facing SMB, Windows Remote Desktop Protocol (RDP)[27], and Telnet services on the default ports that were in use for each organisation in the Nikkei 225. In each case, we fully negotiated the protocol to verify that we were indeed communicating with the expected service. This methodology has some limitations in that the results are constrained by the fact that:

[21] https://docs.microsoft.com/en-us/sysinternals/
[22] https://en.wikipedia.org/wiki/Conficker
[23] https://en.wikipedia.org/wiki/WannaCry_ransomware_attack
[24] https://wiki.wireshark.org/SMB2
[25] Which now old enough to drive in most states (it was born in November 2006)
[26] https://en.wikipedia.org/wiki/Telnet
[27] https://en.wikipedia.org/wiki/Remote_Desktop_Protocol

This methodology has some limitations in that the results are constrained by the fact that:

- Services are only observed on the default ports. Telnet and, less commonly, RDP can be moved to non-default ports.
- Measurements are made only in IPv4-space.
- Certain IP ranges are not examined by Sonar, by request.
- Certain cloud and ISP-related ranges were excluded. The impact of this will vary greatly from company to company.
- Certain networks were excluded if they were believed to be assigned to customers or otherwise allocated to third parties.

All things being equal, these constraints generally result in underreporting of the findings.
We'd like to make a special call out here as it relates to SMB protocols on ISP networks. As is our standard, we've excluded some networks from our datasets which are being used by ISP subdivisions of Nikkei 225 companies. We've done this because the results aren't reflective of the corporate security practices of the organizations on whom we are focusing. That being said, we do want to call out that we observed nearly 900 more SMB endpoints on those ISP-assigned networks. Given the risks associated with SMB, we urge residential and commercial ISPs to start blocking SMB traffic entirely.

Findings: RDP, SMB, and Telnet

We should start this section by stating that **any non-zero number of these services made available to the general internet is considered to be unacceptable** in organisations with mature security programmes. Followers of the Rapid7 blog and past Rapid7 research reports will be quite familiar with this advice, but looking at the calendar here in 2021, we have to note that it's been a while since the last major worm outbreak on the internet. NotPetya (SMB) was 2018, WannaCry (also SMB) was 2017, and Mirai (Telnet) was way back in 2016. Despite all the vulnerability and exploit churn we saw in 2019 and 2020, we appear to be overdue for another self-replicating issue across open ports to insecure services. Closing off your exposure to these services will certainly save you weeks of cleanup later.

Windows Remote Desktop Protocol (RDP)

While some may think that RDP should be considered an exception to this rule, we'd argue that there are commonly available techniques and technologies such virtual private networks (VPNs), RDP Gateway servers, and firewall access control lists (ACLs) that remove the risk related to this technology and so, as a general rule, **RDP shouldn't be exposed to source addresses outside of the organisation.**

Since we're on the topic of RDP, let's start with the discussion with the findings there. On the default RDP port of 3389/tcp, we observed 89 services across 18 companies. One organisation in the Consumer Goods industry accounted for 45% of the observed RDP services.

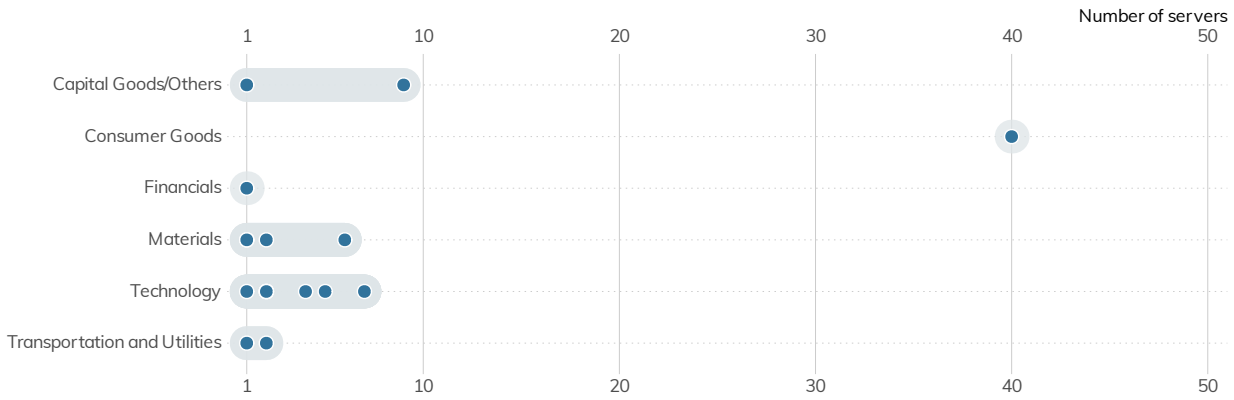

Figure 13: Port 3389 Distribution by Industry

The graphic above shows that, while the overall numbers are mostly attributable to just a few companies, we do see quite a few industries represented.

On a positive note, when we looked at the security requirements for RDP authentication, we found that 94% required Network-Level Authentication (NLA).[28] NLA, introduced in Windows Server 2008, enforces Transport Layer Security (TLS) protection of traffic in-flight, strengthens authentication options, and significantly reduces the risk and impacts related to brute force and certain denial-of-service attacks. NLA has been enabled by default since Windows 2012. The lack of NLA serves as a proxy indicator for older infrastructure, either on the server itself or a requirement for compatibility with older clients. The only other reason for not having NLA enabled is that it doesn't allow authentication with expired passwords. That is another reason to deploy RDP Gateway servers, VPNs, or other infrastructure to provide facilities for changing the password as well as enable security access to remote desktop services.

Windows Remote Desktop Protocol (RDP)

The SMB protocol is for file- and print-sharing as well as interprocess communication on Windows and compatible networks. **We say this in every report[29], but SMB should never be exposed to the internet**. The risks include data leakage from file shares, credential compromise via brute force attacks, and malware infection (think of the previously noted Conficker and WannaCry) via vulnerabilities in the host operating system or service. Given the plethora of options for securely sharing files, SMB shares aren't worth the risk.

[28] https://en.wikipedia.org/wiki/Network_Level_Authentication
[29] https://www.rapid7.com/research/report/nicer-2020/#smb-tcp-445

When we surveyed the Nikkei 225, we looked at 2 different SMB ports: 139/tcp and 445/tcp. Port 139/tcp is used for older variants of SMB, and its presence is generally a sign of very old software and legacy requirements. In our surveys, we found 25 servers across 6 companies exposed over port 139/tcp. They were all running an open source SMB server called Samba[30]. The oldest version of Samba we observed, 3.0.36, was released in late 2009 and contains quite a few critical vulnerabilities.

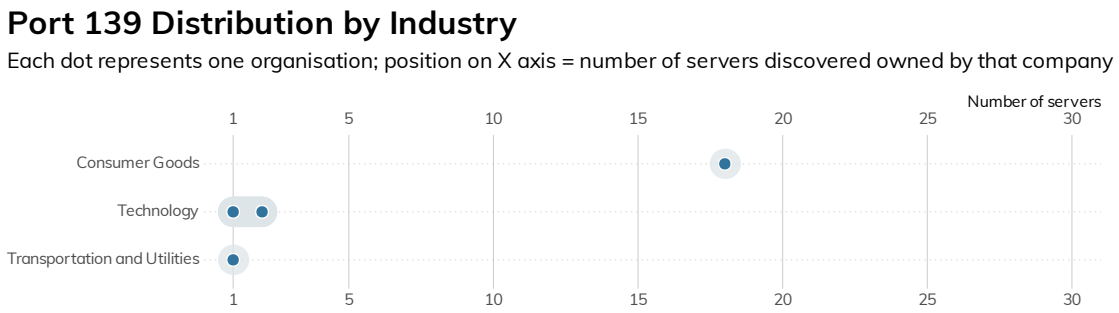

Figure 14: Port 139 Distribution by Industry

We also surveyed SMB on port 445/tcp. Introduced in Windows 2000, this transport for SMB removed some of the legacy protocol overhead. In our research, we observed 146 servers across 17 organisations on this port.

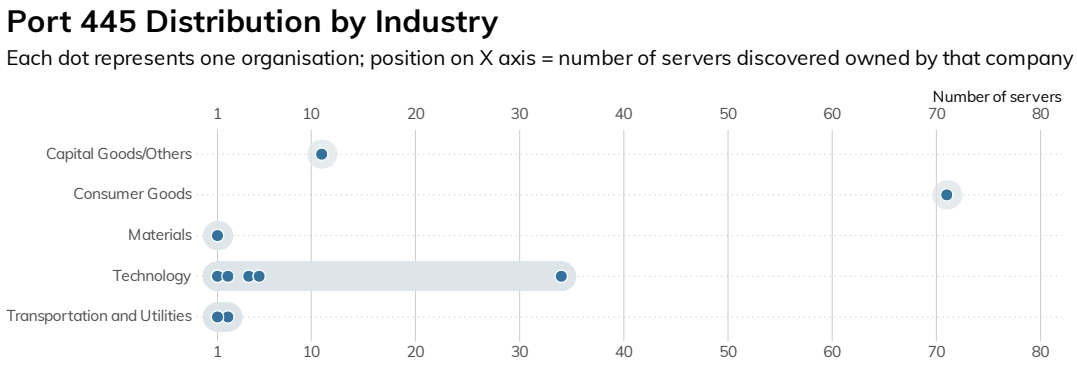

Figure 15: Port 3389 Distribution by Industry

[30] https://www.samba.org/

The mere presence of these SMB servers on the internet is cause for concern, but when we dug into the protocol configurations, the concern increased. All servers supported SMBv1, which means they are missing several critical security controls, and attackers can force clients to downgrade to SMBv1 from more secure versions of the protocol. All of the 146 servers we observed supported a newer version of SMB and so, absent a dependency by legacy systems, shouldn't need to have SMBv1 enabled. We strongly recommend Microsoft's guidance to disable SMBv1.[31]

SMBv3 was released with Windows Server 2012 and included many security and performance improvements[32], such as encryption of data on the wire and protocol downgrade protections. SMBv3 was supported on 68% of the observed servers.

These SMB services also leaked information about the organisation. All of the services provided a hostname,DNS name, and fully qualified domain name (FQDN) configured on the host. This information may indicate role (VCENTER01) or indicate internal organisational structure (db1.prod.us.corp.local).

Telnet

Telnet is a plaintext-based protocol used for providing remote console access to devices. It nearly always transmits credentials and data in cleartext and has no protections against man-in-the-middle (MiTM) injection of commands or data.

Originally specified in 1969, Telnet is well past its "Use By" date and has been superseded by other, more secure technologies such as SSH. Our survey found 243 hosts across 62 companies. The majority of these hosts were in the Technology sector.

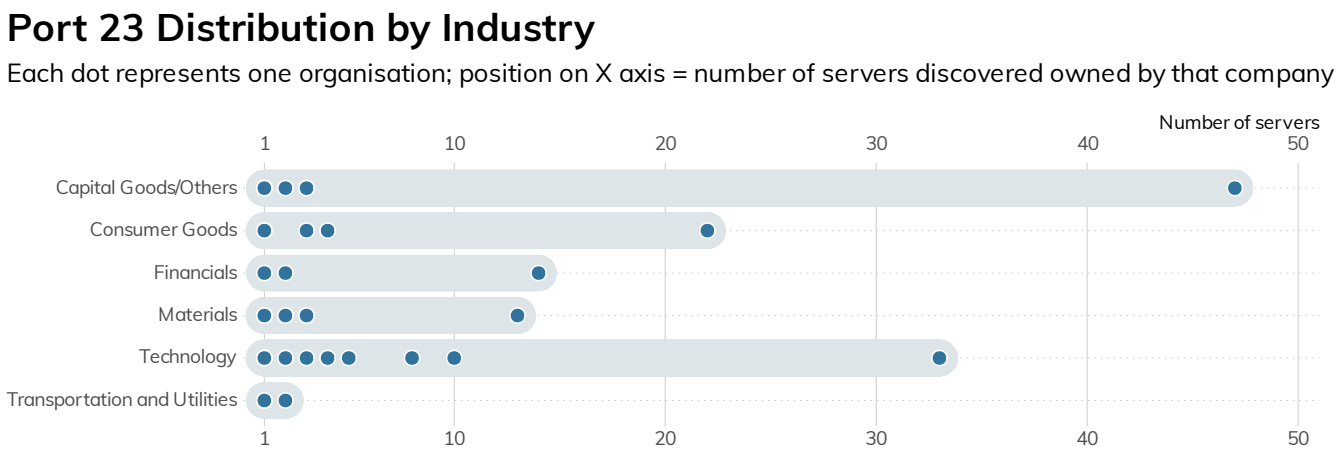

Figure 16: Port 23 Distribution by Industry

[31] https://techcommunity.microsoft.com/t5/storage-at-microsoft/stop-using-smb1/ba-p/425858
[32] https://docs.microsoft.com/en-us/windows-server/storage/file-server/file-server-smb-overview

Most of the equipment was found to be a router or switch. As a general rule, it's considered insecure to use Telnet as opposed to more secure protocols such as SSH.

Also, if Telnet is unavoidable, firewall access control lists (ACLs) and other controls should be used to limit which internet IP addresses can access the devices. Since our survey process had to make connections from multiple IPs—in some cases in different countries—to validate a service, we can say that ACLs were likely not in place or were overly broad.

Exposure Overview

When we look across the surveyed protocols and industries, we can see that there are certain hotspots.

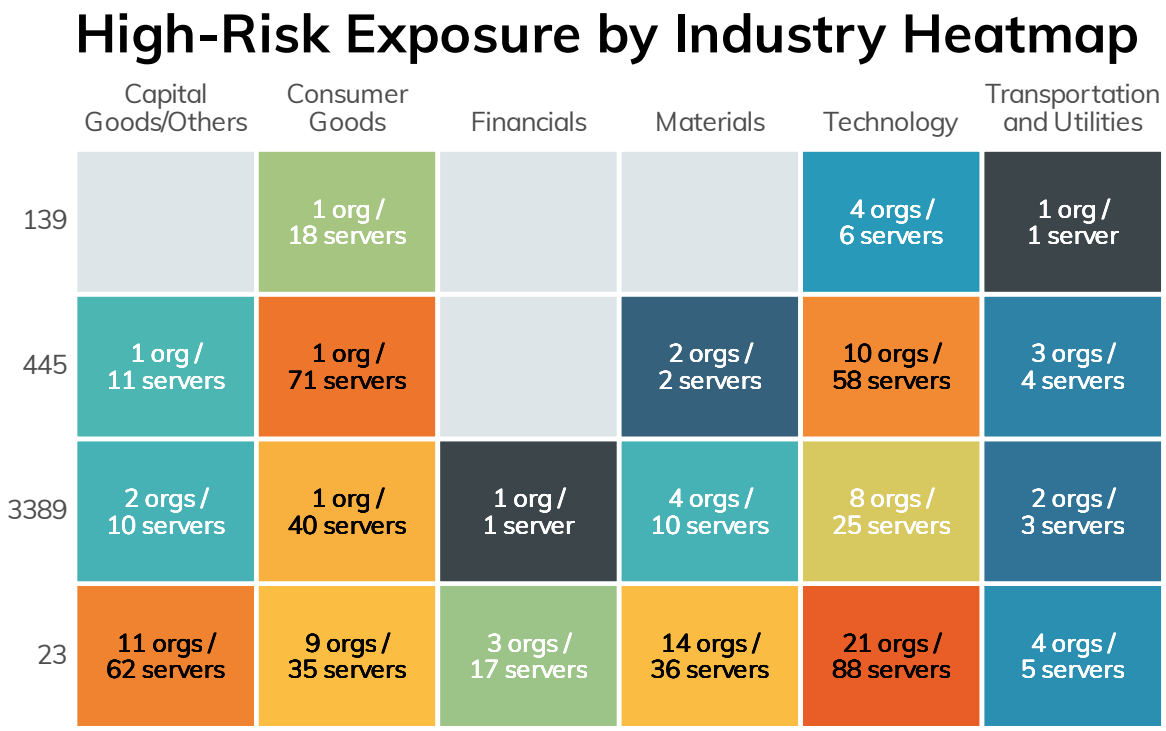

Figure 17: High-Risk Exposure by Industry Heatmap

There are a couple of points that should stand out in the graphic above. First, the SMB-related rows (135, 445) should be entirely empty, but they are not.

Though 1 organization in the Consumer Goods industry dominates the findings there are 16 others who are also exposing these services as well. Second, we see heavy and widespread Telnet usage across many industry verticals here. Finally, most of the RDP exposure is in the Consumer Goods industry and is due to the outsized impact of just 1 company (which we will not name for obvious reasons).

CISO Takeaways

The findings here indicate that even some of the most resourced companies are exposing services that have an outsized risk.

Our guidance for addressing the risks above isn't to implement some advanced security controls or software, but to instead return to the basics. You can find all of them in the early parts of the CIS Top 20[33] controls.

- Develop and maintain an inventory of internet-facing hosts that includes software versions, roles, and services that are expected to be exposed, as well as the reason why. Make sure that this inventory is validated by outside-in scans of all of your public-facing IP ranges.

- Implement security policies and supporting configuration standards that enforce the use of secure protocols and configuration settings. Using the example of Telnet, every device currently using Telnet should be able to support SSH—and if it doesn't, it is too old or insecure to be directly connected to the internet.

- Ensure that software and hardware are kept current. In many cases, such as with Microsoft Windows, newer software brings better security features and controls. Older software's lack of these features can force security trade-offs and require the implementation of compensating controls, which add complexity.

[33] https://www.cisecurity.org/controls/cis-controls-list/

Vulnerability Disclosure Programs Among the Nikkei 225

Every major corporation on Earth is a technology company[34]. It is unthinkable that a business that generates billions of Euros in revenue and employs thousands of workers worldwide would not have a significant technological investment in their products, processes, and logistics. We rely on fantastically advanced technology in every aspect of our modern lives. Of course, anyone who has spent any time analysing these technologies will notice we are routinely bedevilled with vulnerabilities, especially when it comes to internet-based technologies.

As it happens, we have a powerful and proven method to stem the tide of vulnerabilities in major technologies: coordinated vulnerability disclosure[35] (CVD), and a now-standard mechanism to participate in CVD, vulnerability-disclosure programmes[36] (VDPs).

The presence of a publicly accessible VDP is conspicuously lacking across most of the companies listed in the Nikkei 225, which, in turn, makes it difficult for those companies to ever learn about vulnerabilities in their products and technical infrastructure in a constructive way.

While VDPs are more common today among the U.S.-based Fortune 500 (about 20%), these programmes are largely absent in the Nikkei 225: Only 16 of the exchange-listed companies (or about 7%) have a discoverable VDP. Without vulnerability-disclosure programmes, these industries are telegraphing that they do not want to know about their own vulnerabilities, intentionally or not, to their shareholders' and customers' peril.
For this study, we searched for VDPs associated with the Nikkei 225-listed companies and the flagship brands of those companies, much in the same way we would if we were about to disclose a vulnerability about those companies' products or services. At the end of May 2021, we looked for the following, in this order:

- The presence of a VDP associated with all Nikkei 225-listed companies (or flagship brands of those companies) listed on either Bugcrowd's[37] or HackerOne's[38] crowdsourced bug bounty lists, or in the Disclose.io[39] program database.
- The presence of a standardized security.txt file on each company or flagship brand website to facilitate the sharing of discovered vulnerabilities with website maintainers.
- An obvious pointer to, or indication of, a VDP offered by the candidate companies by Googling the terms "vulnerability," "disclosure," and "security" along with the company name and flagship brand.

All this said, it is possible some of the surveyed companies that appear to not offer a VDP do, in fact, have a process for receiving vulnerability intelligence, but the lack of an easily discoverable VDP (in either the company's preferred language or in English) drastically undercuts the effectiveness of the VDP for both researchers and the companies.

[34] https://www.wsj.com/articles/every-company-is-now-a-tech-company-1543901207
[35] https://blog.rapid7.com/2018/10/31/prioritizing-the-fundamentals-of-coordinated-vulnerability-disclosure/
[36] https://blog.rapid7.com/2016/11/28/never-fear-vulnerability-disclosure-is-here/
[37] https://www.bugcrowd.com/bug-bounty-list/
[38] https://hackerone.com/directory/programs
[39] https://github.com/disclose/diodb/

Assessing the relative merits of individual VDPs is beyond the scope of this paper, but it should be noted that not all VDPs are created equal—some offer robust "safe harbor" protections for researchers and accidental discoverers when reporting and publishing vulnerabilities, while others seek to bind researchers in restrictive agreements about what can be assessed and how results are to be handled and communicated. For this paper, the mere existence of a VDP, no matter how liberal or restrictive, counts as a positive.

Results: Prevalence of VDP Adoption

In January 2019, Bugcrowd founder and noted Australian Casey Ellis remarked in a blog post that "only 9% of the Fortune 500 run vulnerability-disclosure programmes.[40]" This is just a touch higher than what we found in Japan in the first half of 2021.

We were able to discover a total of 16 vulnerability-disclosure programmes across the 225 ticker symbols investigated in May 2021, which accounts for about 7% of the Nikkei 225 listings.

With such a low showing, it's difficult to say that any particular industry or valuation quintile has normalized the practice of advertising a VDP. That said, the industry most represented in the positive findings is Technology (13), followed by Consumer Goods (3), and even these Consumer Goods corporations have a decidedly technical bent.

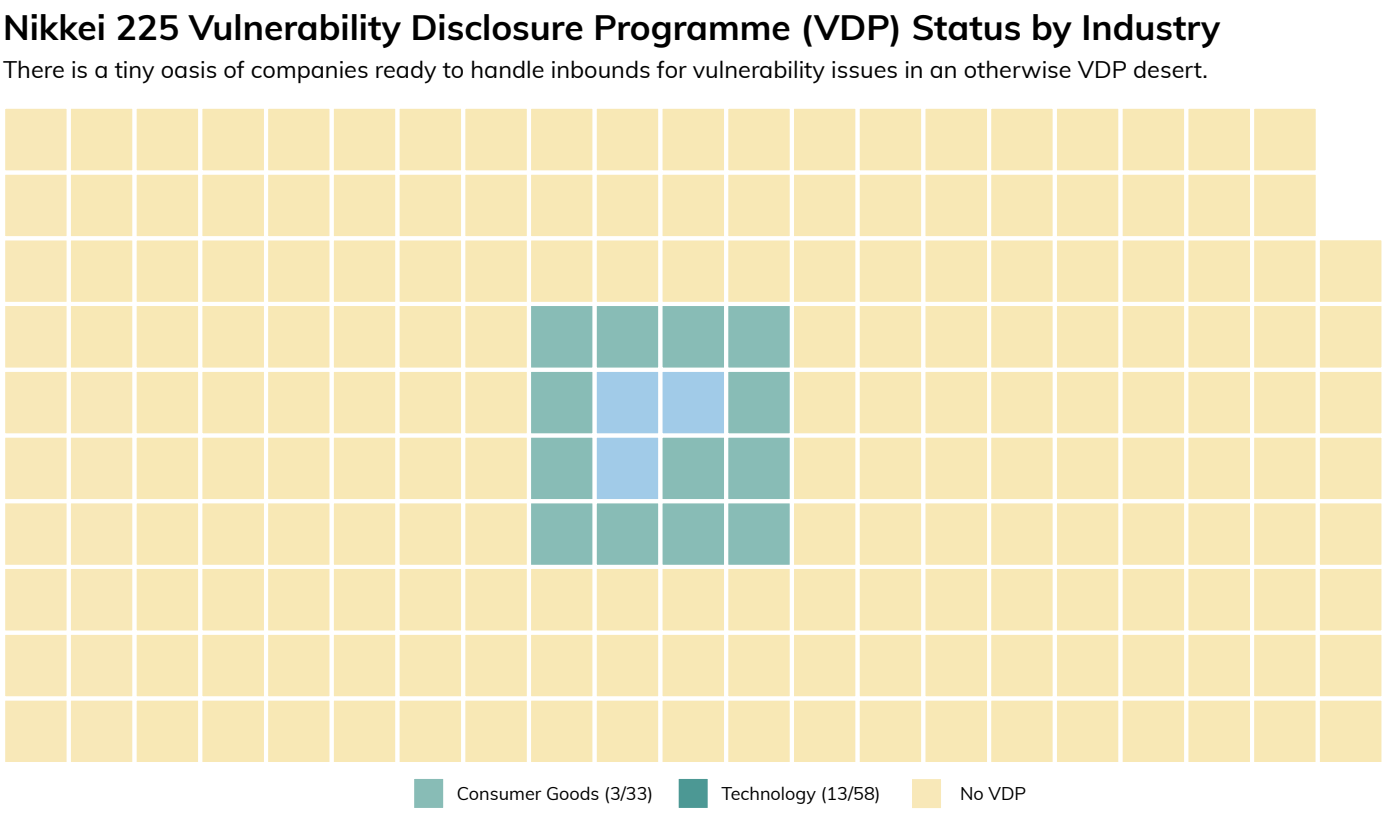

Figure 18: Nikkei Vulnerability Disclosure Program (VDP) Status by Industry

[40] https://www.bugcrowd.com/blog/3-reasons-why-every-company-should-have-a-vdp/

VDPs appear to be wholly absent in major segments of the Japanese economy; namely in Capital Goods/Others (which is largely heavy machinery and real estate), Financials, Materials, and Transportation & Utilities.

The key takeaway from this view of the Nikkei 225 is that, while all major companies have some technical component (and therefore have technical vulnerabilities), over 90% of these top companies in Japan lack a formal vulnerability-disclosure program.

Japan, however, is something of a unique case — since 2004, Japanese companies have been encouraged by the Japanese government to work with the Information-Technology Promotion Agency[41] (IPA) and JPCERT/CC[42]. The IPA was established directly by the Japanese government, while JPCERT/CC is a non-government organization that is largely funded by public money. These two organizations are dedicated to protecting national IT interests, with IPA handling website issues and JPCERT/CC picking up the remainder of software and firmware issues[43] for triage through resolution.

That said, these entities should be considered legacy infrastructure when it comes to VDP. There is simply too much technology floating around for one or two government-funded agencies to handle, on top of securing government assets themselves. The lack of VDPs across the upper echelons of the Japanese economy discourages the reasonable and responsible disclosure of newly discovered vulnerabilities in their products, services, and infrastructure. After all, VDPs aren't just for reporting software bugs in software applications, but are also useful for reporting the discovery of sensitive data found about customers or company internals left open on insecure cloud storage. It is, of course, possible to disclose vulnerabilities to companies in industries without a formal VDP, but the lack of VDPs introduces inefficiencies for the companies and legal risk to researchers.

Finally, a functioning VDP signals that a given company has made some investment in their overall information-security program, so it stands to reason that the lack of a VDP is signaling the opposite. Every company on this list has a website privacy policy, so every company should have some formal method for receiving and handling vulnerability reports.

CISO Takeaways

Hopefully, it is obvious by now that the authors of this paper are strong proponents of clearly defined, easily discoverable vulnerability-disclosure programmes. We believe that every company in the Nikkei 225 (and beyond) should adopt one.

Launching and running a successful VDP may be tricky—after all, the presence of a VDP implies a level of security maturity that may not yet exist at a given company, so CISOs at organisations without a VDP are strongly encouraged to familiarise themselves with the basics of vulnerability disclosure.

[41] https://www.ipa.go.jp/
[42] https://www.jpcert.or.jp/
[43] For more on the unique Japanese approach to vulnerability handling, see "Information Security Early Warning Partnership" at <https://www.ipa.go.jp/files/000059696.pdf>.

We believe there is a critical mass of CISO expertise in building and maintaining VDPs and that there is plenty of opportunity to learn from the experiences of others in the field. In our experience, not only do CISOs personally enjoy discussing their VDP experiences, but it can be hard to stop them when they get going.

ISO 29147[44] (Information technology—Security techniques—Vulnerability disclosure) and ISO 30111[45] (Information technology—Security techniques—Vulnerability handling processes) are excellent starting points for building, maintaining, and improving a vulnerability-disclosure programme. These ISOs were developed in partnership with internationally recognised experts in the field of vulnerability disclosure, and can help any CISO get a leg up.

Another, first-step approach to establishing a minimal VDP is a contact and policy document placed at <hxxps://your-company.com/.well-known/security.txt>. This is a relatively new standard for VDP communication that provides for basic contact information signalling, readable by both humans and machines.[46]

[44] https://www.iso.org/standard/72311.html
[45] https://www.iso.org/standard/69725.html
[46] Interested CISOs can read up on it at https://securitytxt.org/

Conclusion

The global COVID-19 pandemic forced many of these companies to abruptly shift to a large work-from-home workforce in short order, and each company is its own miracle of corporate survival in the face of such drastic and unprecedented changes to the workplace. In addition, Japanese companies are doing pretty well in stamping out dangerously exposed services and version dispersion when compared to other regions surveyed by Rapid7.

However, these companies are lagging their international counterparts in the 3 other areas we measured for this report: adoption of DMARC, HSTS, and VDPs. More progress must be made, and faster. Because of their outsized position in the Japanese business community, they also tend to have access to the best and brightest cybersecurity experts from around the world, and so it is incumbent upon them to behave more like model internet citizens. The researchers at Rapid7 who contributed to this report sincerely hope these companies—and the organizations that have business relationships with them—find this information and advice useful in our shared responsibility of advancing security for everyone.

CISO at a Glance

Throughout this report, we've kept our focus on what CISOs in the Nikkei 225 can do, today, to reduce their exposure to the most common issues we've discussed here. For the reader's convenience, those recommendations are summarised here.

Email Security: If you're on the Domain-based Message Authentication, Reporting & Conformance (DMARC) path, like 13% of the Nikkei 225, that's great! Now is the time to plan out how you'll move from a p=none to a p=quarantine policy, and ultimately a p=reject policy. This is not an easy journey, since you will certainly uncover pockets of shadow IT running their own email infrastructure, but the confidence of being able to authenticate mail from your major brand domains is a pretty great feeling, and a nice item to report to your board of directors.

Web Security: HTTP Strict Transport Security (HSTS) is rapidly becoming table stakes for running a reasonably secure website, and this is the kind of security feature that browser manufacturers like Google, Apple, Microsoft, and Mozilla are likely to enforce in future versions of Chrome, Safari, Edge, and Firefox. It's a relatively easy switch that CISOs can flick (compared to the universe of nice-to-haves in cybersecurity, anyway), so take some time to investigate whether your organisation is using HSTS and if not, why not?

Version Dispersion: For the mega-corporations that roam the fields of capitalism, mergers and acquisitions are a fairly common activity throughout the year. That means the Nikkei 225 CISO is never truly "done" with ensuring version consistency across the enterprise, even after investing in an excellent asset and vulnerability management toolchain. New networks and network services will join your ranks, and that means undertaking a fairly continuous modernization and normalization effort for those new assets. Taking on this continuous effort will pay off in easier, more straightforward planning for the next patch cycle, scheduled or surprise.

High-Risk Services: Telnet, SMB, and RDP have no business being exposed directly to the world at large, and are just waiting for the next self-replicating cyberattack to sweep across the internet. An up-to-date inventory of exposed services, sourced from internal and external scanning, is worth its virtual weight in Bitcoin, and will help you enforce a no-nonsense policy of network service exposure to the internet. As stated above, though, there are very few of these exposed services left in the Nikkei 225 as of 2021.

Vulnerability Disclosure Programs: As a CISO, you might have hired the best of the best software, QA, and platform engineers. But, without a good way to harness the smarts of the tens of thousands of talented hackers around the world, you may never learn about the most critical vulnerabilities in your products and services. A VDP is a bridge to that enormous community of well-meaning investigators who have goals aligned with your own: a safer and more secure internet. Getting that program spun up now will give you plenty of time to practice safer software production. As a bonus, most of the pioneering work is already done for you, in the form of ISO 29147 and ISO 30111.

Appendix: Prioritization in Times of Crisis

The disclosure of both the SolarWinds-related multiple-technology vulnerabilities (and associated campaigns), the release of the out-of-band Microsoft Exchange patches responding to active exploitation campaigns, and the Codecov compromise that will undoubtedly impact many, many software development CI/CD processes, have all strained virtually every single information-security team in every industry. We wanted to take a moment to help ensure you're on safer ground now, and also put each section into context, relative to some of the crises we've had to deal with this year.

The SolarWinds and Codecov situation brought third-party risk square into focus like it has never been before. If you had a solid list of partners/vendors and a well-oiled contact plan (which many organizations did), you may have weathered that portion of these extended incidents fairly well. If not, we hope you had the support required to put such things in place and have been able to use it in some subsequent serious vulnerability disclosures and exploit campaigns since.

When it comes to being able to get a feel for how well a partner/vendor values safety and resilience, you may want to heed the advice in the "CISO Takeaway" section. It's much easier to sleep at night knowing that the bulk of your third-party contacts prioritize email safety, avoid exposing dangerous services to the internet, and stay current with both patching and advanced encryption standards. You will also know how to contact them in the event you do discover a security issue with any of their products and services, since they'll have a vulnerability-disclosure program in place.

Similarly, the massive Exchange vulnerability and associated malicious campaigns further demonstrated how quickly 1 weakness in a component used by hundreds of thousands of organisations can come out of the blue to disrupt execution on even the most well-crafted enterprise information-security roadmap. Having current, accurate telemetry of what is deployed internally and externally, along with highly agile quality assurance and change management processes (as noted in the section on version complexity), can be the difference in having an unexpected patch (like Exchange) be a quick exercise with a slight bit of triage (to ensure attackers did not have time to target you) versus an "all hands on deck" massive incident.

We hope our quantification, context, and advice prove useful as you emerge from these 2 major incidents to take on the remaining challenges that await us all in 2021 and beyond.