



How I (could've) stolen your corporate secrets for \$100

Author:
Cameron Camp

Contributor:
Tony Anscombe

April 2023

CONTENTS

Executive summary	4
Getting started.	5
Details of routers acquired.	5
Customer data	8
Third-party data	8
Trusted parties	8
Specific applications	9
Extensive core routing information	9
Trusted operators	9
What's the risk?	9
The human element	11
Third-party trust	12
Trusting e-waste disposal companies	13
Who is really telling the truth?	13
Industries impacted	13
Notifications	14
What other devices might have secrets?	15
Best practices and guidance	15
How <i>should</i> you dispose of your routers?.	15
Should you engage a data disposal or e-waste service?	17
What if the device is "dead"?	17
What if you already have unsecured routers somewhere "out there"?	18
What if the previous owner's information is on a secondhand device?	18
Potential further research	18
Conclusion	19
References	19

LIST OF TABLES

Table 1. Routers obtained for this research 5

Table 2. Accessibility of uniquely identifiable network configurations 7

Table 3. Overview of the companies identified from the nine easily accessed
router configurations in this study 8

LIST OF FIGURES

Figure 1. Sample of routers purchased for this research 6

Figure 2. A Juniper SRX550 router 7

Figure 3. Extract from a Juniper SRX series configuration 10

Figure 4. Extract from another Juniper SRX series configuration 11

Figure 5. CompactFlash card storage slot on two Cisco devices (lower) and
the Fortinet Storage Module (FSM) slot on a FortiGate (upper) 17

EXECUTIVE SUMMARY

There are well documented processes for the decommissioning of hardware, yet ESET researchers found during a period of discovery and extensive analysis that core routers, the kind that are likely to be found in corporate networks, are often not wiped clean before they are decommissioned and offered for resale. This leaves critical and sensitive configuration data from the original owner or operator accessible to the purchaser and open to abuse.

Results reported here show that a majority of the secondary market core routers sampled contained recoverable configuration data from their previous deployments, replete with sensitive, and even confidential, data. This allowed ESET researchers to identify devices previously used in a data center/ cloud computing business (specifically, a router provisioning a university's virtualized assets), a nationwide US law firm, manufacturing and tech companies, a creative firm, and a major Silicon Valley-based software developer, among others. Obviously, these industry sectors were just the "luck of the draw"; such a profusion and array of victims suggests that many organizations – including some that really should know how best to handle such tasks – do not have reliable decommissioning processes in place, or perhaps place undue trust in their managed service providers or e-waste contractors.

The results of unreliable or incomplete decommissioning of such devices could be catastrophic for an impacted organization. We found, among the networks represented by the functioning routers we purchased, that 56.25% contained trivially accessible and sensitive corporate information. Among those identifiable networks:

- 22% contained customer data
- 33% exposed data allowing third-party connections to the network
- 44% had credentials for connecting to other networks as a trusted party
- 89% itemized connection details for specific applications
- 89% contained router-to-router authentication keys
- 100% contained one or more of IPsec or VPN credentials, or hashed root passwords
- 100% had sufficient data to reliably identify the former owner/operator

Much of this information would be very useful to anyone planning an attack.

Equally concerning was the difficulty the team experienced during the disclosure process when attempting to contact the companies concerned, to disclose that our researchers were in possession of a device with the company's sensitive network configuration data.

On many levels, this research is about human error compounding to create a potential breach, and the mitigation steps companies can take to reduce or avoid such pitfalls moving forward.

GETTING STARTED

Cybersecurity literature and popular media contain many studies and cautionary tales of secondhand computing equipment being found for sale with the previous owner's data – often of a very personal and sensitive nature – still intact. These studies have consistently shown that a majority of used hard drives [1] [2] [3] [4], USB flash drives [5], SD and other memory cards [6], and mobile devices such as phones and tablets [7] purchased secondhand have had no, or very incomplete, data wiping procedures applied to them before being offered for resale.

One might hope that these well-publicized studies – which have looked mainly at preowned *consumer* devices – would lead to a heightened sense of awareness in general, and especially among system admin and security staff disposing of used *corporate* equipment with storage functions, and whose configuration data can contain the keys to the (virtual) kingdom.

Such optimism, sadly, is often misplaced in cybersecurity circles...

Musings such as these, though, were not the inspiration for this research. About a year ago we set up a lab for testing real-world attacks against multiple attack surfaces like RDP, Microsoft Exchange, industrial control systems (ICS), and others. While sourcing equipment, we chose devices that would most closely mimic typical current production environments for a representative test, including core routers typical of organizations of various sizes. Generally, these would be a generation old, but still solid, well tested, and very widely deployed in current environments across the globe. Further, they are a class of device that is commonly available at bargain basement prices in the secondary market (US\$50–150). To test variants, we selected multiple brands and specific equipment types most suited for corporate use.

Much to our surprise, we found, as we staged the lab and these devices started to arrive, that many – notably core routers – still contained sensitive information. This raised a great deal of concern. For instance, the very first core router we received originally belonged to a reputable, major, Silicon Valley-based, household-name software vendor with a global footprint and contained such sensitive corporate data.

This surprising discovery temporarily diverted the path of our research – we branched out to investigate the prevalence of sensitive corporate network configuration data being, presumably unwittingly, available for sale on the secondary market. To determine if this initial finding was a one-off, we began procuring more device variations, as used in different market segments.

DETAILS OF ROUTERS ACQUIRED

Our budget allowed us to obtain 18 routers, covering a range of models from three of the big vendors in this space. Once received, all devices detailed in [Table 1](#), with some pictured in [Figure 1](#), were kept in a secure location and were never connected to the internet, or any network.

[Table 1](#). Routers obtained for this research

Manufacturer	Device type	Number
Cisco Systems	ASA 5500 series	4
Fortinet	FortiGate series	3
Juniper Networks	SRX Series Services Gateway	11




Figure 1. Sample of routers purchased for this research

And there were also a couple of 2U devices like the one in Figure 2, which weren't in the rack for the previous photo.



Figure 2. A Juniper SRX550 router

Of these 18 core routers, and despite them all being advertised as fully functional, one was dead on arrival (when powered on, the fan spun up but there were no indicator lights and no accessible console). As we initially had decided not to make any heroic efforts, such as data recovery or forensic investigation, we have eliminated this device from our remaining results, although perhaps simply replacing its power supply might have rendered it bootable, and if so, it seems likely that would have resulted in the configuration data – which was present when the device “died” – becoming accessible.

Using only standard operating system commands, and utilities supplied by the device manufacturers, we discovered that two devices were a mirrored pair from a cluster configuration, failing over to each other. Of the remaining 16 routers, two were more heavily protected than the others, only disclosing the internal and external network IP space; thus, we were able to extract nine (56.25%) *complete*, unique sets of previous configuration data from those devices. The remaining five devices contained no trivially extractable configuration data, appearing to have been cleaned via their standard device-specific “wipe” procedures. Table 2 summarizes our results after removing the dead router and one of the mirrored routers; assuming that the functional devices with no readily recoverable configuration data were all from different networks, the 18 original devices represent 16 unique networks.

Table 2. Accessibility of uniquely identifiable network configurations

Network configuration status	Number	Percentage ¹
Complete configuration data available	9	56.25
Wiped properly	5	31.25
Hardened	2	12.50
Dead (no recoverable data)	1	N/A
Second device in mirror pair	1	N/A

No procedures or tools of a primarily forensic or data recovery nature were ever employed, nor were any techniques that required opening the routers’ cases.

¹ Percentages are based on the 16 functioning devices with unique network configurations.

In all nine of the networks from which we recovered configuration data, that data allowed us to ascertain with very high confidence the previous owners of those routers. [Table 3](#) outlines some characteristics of those companies where router configuration data was discovered and analyzed.

Table 3. Overview of the companies identified from the nine easily accessed router configurations in this study²

Vertical	Reach	Employees	Revenue (US\$, M)
Light manufacturing/supplier	Products/subassemblies integrated in larger companies' products	5–50	5–25
Legal	Nationwide (US) law firm	50–100	5–25
Creative	Services multiple tier one, household brand companies	100–500	25–100
Data center	Direct data services, as well as managed MSP services for region	100–500	25–100
MSP	Manages fintech companies	100–500	25–100
Open-source software	Has over 100 million users, worldwide	100–500	500–1,000
Events	Operates trade shows and equipment rentals	1,000–5,000	25–100
Multinational technology company	Global data company	10,000+	1,000+
Telecoms	This was CPE (Customer Premises Equipment) for a transportation company	10,000+	1,000+

CORPORATE SECRETS ABOUND

A range of sensitive data was found on many of the core routers we obtained for this research.

Even for those not familiar with these devices, it's no surprise that IP lists might be found in their configuration data – that's what routers *do*. While our data may not be representative of the information inadvertently available on such devices on the secondary market – or even on routers in general – it may come as a considerable surprise, and concern, what we found in our sample.

It is a treasure trove for a potential adversary – for both technical and social engineering attacks.

Customer data

In some cases, core routers point to internal and/or external information stores with specific information about their owners' customers, sometimes stored on premises, which can open customers up to potential security issues if an adversary is able to gain specific information about them. We found such data for two (22%) of the networks represented in our tests.

Third-party data

As we've seen in real-world cyberattacks, perhaps most (in)famously in the Target incident in 2013 [\[8\]](#) [\[9\]](#), a breach of one company's network can proliferate to their customers, partners, and other businesses with whom they may have connections. Sensitive data allowing such third-party connections to other company's networks was present for three (33%) of the configurations we recovered for these networks.

Trusted parties

Trusted parties (which could be impersonated as a secondary attack vector) would accept certificates and cryptographic tokens found on these devices, allowing a very convincing [adversary in the middle \(AitM\) attack](#) with trusted credentials, capable of syphoning off corporate secrets, with victims unaware

² Based on publicly available data, using multiple independent sources when possible.

for long periods. Four (44%) of the networks would allow AitM attacks, were the configuration data to fall into hostile hands.

Specific applications

Complete maps of major application platforms used by specific organizations, both locally hosted and in the cloud, were scattered liberally throughout the configurations of these devices. These applications range from corporate email to trusted client tunnels for customers, physical building security such as specific vendors and topologies for proximity access cards and specific surveillance camera networks, and vendors, sales and customer platforms, to mention a few. These include, in alphabetical order, Exchange, FTP, LDAP, Lync/Skype, PeopleSoft, Salesforce, SharePoint, Spiceworks, SQL, VMWare Horizon View, and VoIP. Three of the most popular of these applications have had, respectively, three CVEs in 2022 with a maximum severity of 4.9 (Medium severity), nine CVEs assigned so far in 2023 and 18 in 2022, with a maximum severity of 8.3 (High severity); and six CVEs assigned so far in 2023 and 28 in 2022, with maximum severity of 9.0 (Critical severity).

Additionally, ESET researchers were able to determine over which ports and from which hosts those applications communicate, which ones they trust, and which ones they don't. Due to the granularity of the applications and the specific versions used in some cases, known exploits could be deployed across the network topology that an attacker would already have mapped. In our sample, we found application data for eight (89%) of the networks.

Extensive core routing information

From core network routes to BGP peering, OSPF, RIP, and others, we found complete layouts of various organizations' inner workings, which would provide extensive network topology information for subsequent exploitation, were the devices to fall into the hands of an adversary. Recovered configurations also contained nearby and international locations of many remote offices and operators, including their relationship to the corporate office – more data that would be highly valuable to potential adversaries. IPsec tunneling can be used to connect trusted routers to each other, which can be a component of WAN router peering arrangements and the like. In our sample, we found IPsec credentials for two of the nine networks (22%); for eight of the networks (89%), we found router-to-router authentication keys/hashes and the expected IP tables. We found it interesting that all networks in our sample had at least one of these types of sensitive data present.

Trusted operators

The devices were loaded with potentially crackable or directly reusable corporate credentials – including administrator logins, VPN details, and cryptographic keys – that would allow bad actors to seamlessly become trusted entities and thus to gain access across the network. Here, the IPsec info we mentioned above also applies (22%); for seven networks (78%), we found VPN credentials, and root hashes for eight (89%) of the networks.

WHAT'S THE RISK?

What could an adversary do with these types of corporate secrets? In short, plan and execute an attack with inside information normally visible only to highly credentialed personnel. Even within an organization, rarely do folks outside core network routing, systems, and security administration have access to this level of sensitive information. Core routing is normally within the realm of network administrators and their managers, and precious few others.

It's not just the network. Quite often lists of application types and network locations were available, and in some cases remote cloud applications hosted in specific remote data centers, complete with which ports or controlled access mechanisms were used to access them, and from which source networks. Additionally, there were firewall rules used to block or allow certain access from certain networks; often, specifics about times of day they could be accessed were available as well.

With this level of detail, impersonating network or internal hosts would be far simpler for an attacker, especially since the devices often contain VPN credentials or other easily cracked authentication tokens.

We also knew about their administrators. Some companies used a centralized admin identity to manage company-wide assets; others designated specific individuals by name with administrative access to the lowest levels of the device.

Moreover, we knew about each organization's approach to security in general. By noting how detailed or vague their security defenses were on these devices, we could make a reasonable approximation about the security levels in the rest of their environment. The irony is that the security expertise demonstrated in these devices' security configurations didn't seem to map directly to the size or sophistication of their organization: apparent security expertise was all over the map.

We would expect to see a large, multinational organization have a very structured, standards-driven, and complete set of security initiatives reflected in their devices' configurations, but that just wasn't always the case.

We also noted, significantly, that multiple devices were acquired following decommissioning from managed IT providers who operate networks for much larger organizations, so often the affected organizations would have no idea that they may now be vulnerable to attacks due to data leaks by some third party. This seemed like a massive security attack surface that was potentially wide open to a whole host of target organizations. Two such IT companies (an MSSP in one case) managed networks for hundreds of clients in a variety of sectors including education, finance, healthcare, manufacturing, and professional services, among others.

We sampled across multiple device manufacturers to emulate a typical environment and found unredacted configuration data distributed fairly uniformly across device types. Some devices had better default security settings that made some data harder to access, but all devices had settable options to guard against the proliferation of "residual data", even if they weren't implemented; settings that would have been free and fairly simple to implement had the previous owners or operators³ known – or cared – to enable them.

For example, here we provide a couple of heavily redacted samples from two routers we purchased.

Figure 3 shows a variety of administrative access and methodologies along with crackable passwords.

```
set system time-zone America/New_York
set system no-redirects
set system authentication-order radius
set system authentication-order password
set system root-authentication encrypted-password [REDACTED]
set system radius-server 10.[REDACTED] secret [REDACTED]
set system radius-server 10.2.[REDACTED] timeout 5
set system radius-server 10.[REDACTED] source-address 10.[REDACTED]
set system radius-server 10.1.[REDACTED] secret [REDACTED]
set system radius-server 10.[REDACTED] timeout 5
set system radius-server 10.[REDACTED] source-address 10.2.[REDACTED]
set system login user [REDACTED] admin uid [REDACTED]
set system login user [REDACTED] admin class super-user
set system login user [REDACTED] admin authentication encrypted-password [REDACTED]
set system login user [REDACTED] uid [REDACTED]
set system login user [REDACTED] class super-user
```

Figure 3. Extract from a Juniper SRX series configuration

³ For brevity's sake, we will use owner or owners from here on, but it should be read as "owner(s) or operator(s)".

Figure 4 shows the relationship between the previous owner's company and many of their vendors and technologies, with specific IPs, to give a hacker a massive head start attacking from many potential angles. It also outlines real names of trusted personnel on the network, and some malware they were having problems with, so blacklisted a couple of entire IP blocks.

```

set security zones security-zone untrust address-book address BrowserCam 6 [REDACTED] 32
set security zones security-zone untrust address-book address Contegix1 199. [REDACTED]
set security zones security-zone untrust address-book address Contegix2 199. [REDACTED]
set security zones security-zone untrust address-book address Dev7FTP_U [REDACTED]
set security zones security-zone untrust address-book address Dev7FTP_U [REDACTED]
set security zones security-zone untrust address-book address Dev7Us [REDACTED]
set security zones security-zone untrust address-book address Dev7User_2 14 [REDACTED]
set security zones security-zone untrust address-book address Dev7User_3 24. [REDACTED]
set security zones security-zone untrust address-book address Han [REDACTED] /32
set security zones security-zone untrust address-book address Khar [REDACTED]
set security zones security-zone untrust address-book address nist1_ny.ustiming.org 64. [REDACTED]
set security zones security-zone untrust address-book address nist1_pa.ustiming.org 2 [REDACTED]
set security zones security-zone untrust address-book address rackspace_server 20 [REDACTED]
set security zones security-zone untrust address-book address Pmo_Ext [REDACTED]
set security zones security-zone untrust address-book address sitefi [REDACTED]
set security zones security-zone untrust address-book address Sony [REDACTED]
set security zones security-zone untrust address-book address SQL_Con [REDACTED]
set security zones security-zone untrust address-book address VirtualEdge [REDACTED]
set security zones security-zone untrust address-book address VirtualEdge [REDACTED]
set security zones security-zone untrust address-book address Amazon_Serve [REDACTED]
set security zones security-zone untrust address-book address Amazon_serv [REDACTED]
set security zones security-zone untrust address-book address Rackspace_Serv [REDACTED]
set security zones security-zone untrust address-book address TorpigMalware [REDACTED]
set security zones security-zone untrust address-book address ToupigMalware2 [REDACTED] 0/16
set security zones security-zone untrust address-book address MXLogics [REDACTED] 0/21
set security zones security-zone untrust address-book address MXLogic [REDACTED] .0/21
set security zones security-zone untrust address-book address 21 [REDACTED] .0/24
set security zones security-zone untrust address-book address Mark_LDAP [REDACTED]
set security zones security-zone untrust address-book address rbm_test 74 [REDACTED]
set security zones security-zone untrust address-book address Flex 50 [REDACTED]
set security zones security-zone untrust address-book address Lync- [REDACTED]
set security zones security-zone untrust address-book address Dev-8-Pub [REDACTED]
set security zones security-zone untrust address-book address ssh_dev [REDACTED]
set security zones security-zone untrust address-book address PMO_Ext [REDACTED]
set security zones security-zone untrust address-book address Gmail_Smtp d [REDACTED]
set security zones security-zone untrust address-book address Yahoo_Smtp dns-n [REDACTED]
set security zones security-zone untrust address-book address Yahoo_Pop dns-n [REDACTED]
set security zones security-zone untrust address-book address Gmail_Pop dns-n [REDACTED]
set security zones security-zone untrust address-book address NBPub [REDACTED]
set security zones security-zone untrust address-book address NBPubli [REDACTED]
set security zones security-zone untrust address-book address playbook [REDACTED]
set security zones security-zone untrust address-book address SPpub 2 [REDACTED]
set security zones security-zone untrust address-book address [REDACTED]

```

Figure 4. Extract from another Juniper SRX series configuration

THE HUMAN ELEMENT

Like many tales in cybersecurity, this is a story about human error and lack of resources.

It is tempting to assume that companies – especially large multinationals – have rooms full of security analysts working on security issues, but often that is simply not the case. Most network admins are hopelessly over-subscribed and often swamped by tech support tickets. And even if a company has an IT department, it may not have any security specialists at all [10]. (ISC)² reported in its 2022 Cybersecurity Workforce Study an estimated global cybersecurity workforce of 4.7 million people, but with a 3.4 million shortfall; an estimated 410,695 of those unfilled positions were in the US alone [11]. Further, ISACA's State of Cybersecurity 2022 report found that over 60% of surveyed information security professionals worked in organizations that had either, or both, unfilled cybersecurity positions or cybersecurity teams that were understaffed [12] [13].

This pushes the myriad tasks of rolling out and managing best practices on security gear to the back burner. Normally the emphasis is on getting the equipment up and running – then keeping it that

way – but hardly ever on training and continuing education to understand the latest security defenses available. Security is often viewed as a cost that's hard to justify. But then, so are breaches.

THIRD-PARTY TRUST

Third parties – from cloud providers to managed service providers (MSPs) to hosts of others – vie for our attention, but they also may represent vulnerable security services that are hard to understand, control, and remediate in the event of a security breach.

Speaking of trust, and its implied responsibility, who is most responsible for these corporate leaks? There actually is a chain of mistrust, starting with the party that originally owned these devices, and running through the third parties they trusted to manage, secure, deploy, and maintain them. Then there's the issue of what happens after these devices leave the premises, either expectedly or unexpectedly. Our research points to both, and lots of shades of responsibility.

When equipment and services move outside of your perimeter, there is a fresh chain of potential security issues, and little control determining which ones are trustworthy, and with what confidence level. It's also unclear if and when your primary provider swaps out their "trusted" third parties' tech on their integrated platforms. Some overzealous third parties tend to lead with marketing, not security, and then try to get customers to indemnify them against exposure to security breaches.

In our research we found some large gaps in the processes, which are basically divided into a couple of categories: MSPs and third-party equipment disposal companies. Both are partially culpable for allowing sensitive company information to leak, but not necessarily how you would expect.

In the current surge of corporate mergers and acquisitions, companies test a third party and then – if it's a good fit – purchase and integrate its technology into their own, a process that rarely goes smoothly in the initial stages.

Since the acquired third party does many things differently, there is a pruning process to make their technology work with the acquirer's stack. The acquirer, however, is then responsible for whatever tech is involved in the mashup.

In one case, a reputable data center had a business arrangement (we didn't deep dive into the details of what the arrangement was) where the data center asserted that a contracted third party did shoddy work for a college, which eventually leaked outside the organization on some resold routing equipment. Or so the data center said. There are several problems here:

1. Maybe the data center is ducking responsibility by shuffling the blame – the corporate equivalent of the common practice of blaming "the guy who left last month". It's uncanny how often the guy who just left has been found to have done bad things. Or at least that's a pretty convenient way of framing blame.
2. If the third party was, indeed, terrible, why was it engaged or acquired?
3. How are we to understand what other insecure or incompetent services may have been engaged by any given tech provider?

In this case the data center actually has a good reputation, suggesting this sort of thing can happen to reputable companies even if (which we couldn't verify) they did their due diligence. The lesson here might be that even if you're doing your best work, relying on third parties to perform as expected is a process that is *far* from perfect.

Either way, this kind of thing could definitely have repercussions for the parent organization, which is probably a good lesson to learn from someone *else's* experience.

TRUSTING E-WASTE DISPOSAL COMPANIES

Many companies, instead, engage a third-party e-waste disposal service, charged with representing and verifying secure destruction and disposal of digital equipment and the data contained therein. In one of the cases we researched, this is exactly what the original owner of a router we had purchased had done. That clearly didn't go as planned.

The medium sized manufacturing business that used this type of third-party disposal service was shocked by the data we had (when we finally were able to get them on the phone). This data revealed company specifics like where their data centers are (complete with IPs) and what kinds of processes happened at those locations. From this information an adversary could get a critical view into proprietary processes that could be invaluable to the company – their “secret sauce” – which could be quite damaging. In an era where potential competitors digitally steal technical research, product designs, and other intellectual property to shortcut engineering R&D processes, this could have had a real financial impact.

The manufacturing company, in this case, was relieved to know that at least part of its sensitive data was held securely by a cybersecurity company that then notified them, rather than that their data was still in the open marketplace. A representative of the company then had plenty to say about the third-party disposal company they contracted with, much of which is not suitable for print, but which strongly suggests that this manufacturing company was not pleased with the “secure service”.

WHO IS REALLY TELLING THE TRUTH?

Of course, this could also just be a scapegoat for other mishandled internal processes at the manufacturing company – it's hard to know with our limited information. But it still begs the question whether you really can trust a third-party disposal process, and if it can be verified that your data was truly securely disposed of in the end. It may also depend on how much your data may be worth to other nefarious parties or for resale to the highest bidder.

In general, throughout this investigation, it was hard to know who was telling the truth. One of the companies that we identified as a previous owner of an investigated device, and that we successfully contacted, claimed that the device was stolen and asked us to return it to them. We requested proof, such as a copy of a police report or direct contact from the relevant law enforcement agency, eventually returning the device anyway. We still haven't been contacted by law enforcement, so we really cannot confirm this claim.

It's clear that trust could have been violated at a number of steps along the circuitous path before the devices ended up in our research lab. What is not in question is that they arrived in our research lab after being legally purchased. None of the original owners we managed to contact were thrilled to hear that these devices were out of their control and still contained sensitive company information. Some were further surprised to learn that their former device was still in existence, *having paid to have it shredded*. They were pretty happy that the devices had not fallen into the wrong hands, and that we engaged in a coordinated disclosure process and that we would not identify them beyond broad sector and regional categorizations. And it wasn't just tech companies we were talking to; the previous device owners that we identified came from across the business landscape.

INDUSTRIES IMPACTED

Surprisingly, the issues were spread across multiple verticals, affecting everything from light manufacturing to financial, legal, and creative service companies; from SMBs to large multinational technology companies. The takeaway? Basically, anyone who runs a moderately sophisticated network,

or hires someone to assist them in operating one, could be affected. We broke down the specifics of the affected organizations seen in our research in [Table 3](#).

While their responses, and indeed abilities to respond, differ greatly across sectors and organizations, we found that technology companies generally did a better job. They usually had more clearly defined processes for both disposing of equipment and responding to incidents (see the *Notification* section, below). Other companies apparently have no prepared response other than fright, prayer, or ignoring the messenger. None of these should be considered an acceptable security strategy.

NOTIFICATIONS

Once we identified organizations involved, we also initiated a separate process to attempt to notify the former owners of the routers we had purchased. While some were responsive, others were incredibly difficult – or impossible – to reach.

You might expect – as we did – that affected organizations would be anxious to hear about the issues we had uncovered. That would be wrong. While it varied by organization, it was perplexingly difficult to get responses from most of them. We discussed contacting the Cybersecurity and Infrastructure Security Agency (CISA) – just so some of these organizations would respond. But that seemed like overkill for what should have been simple data leakage notifications.

After trying for many months, we still have been unable to get a response from three organizations. One is a creative/PR agency handling major accounts for large healthcare vendors, and whose networks contain those clients' information. You might imagine PR agencies would be easy to get ahold of, given their field and their understanding of potential brand damage, but it was simply not the case here. On the other hand, they perhaps best understand the reputational damage they might suffer should information about such a leak become public. After two months of trying to contact them, we were left with radio silence.

We had another idea of contacting companies through LinkedIn. One fintech executive we tapped told us that such an approach might be more effective, raising our alert above the noise threshold. Using LinkedIn worked in a couple of cases, but was still ineffective with the creative agency, even after contacting two senior VP staff there. The same with the nationwide law firm, famous for its high-volume television commercials. We got the distinct impression that neither really wanted to hear from us. Here are some potential reasons why:

1. Notification fatigue – Staging and deploying a network security suite is an exercise in tuning. If you set the gear to default and just turn it on, you may get 30,000 notifications a week. Some operators create an email rule that pours this torrent into a folder, which they can then ignore until they can tune the system. All too often, though, this flood of security notifications causes the security team to get sloppy and ignore important ones. This is a common reason why we see teams deploy the latest technology and still have security issues.
2. They don't want to hear it – IT staff is famously overtaxed by support requests during normal workflows, leaving precious few cycles to track down potentially difficult problems that take many cycles to triage, and may come from a multiplicity of potentially noisy sources.
3. They don't have a point person – At many organizations there's not a well-defined single point of contact for reporting security issues, so it becomes a "best effort" to determine who should be responsible, which can lead ultimately to little or no response.
4. They don't have a plan – It's hard to tell, from an outsider's perspective wanting to notify an organization, what plan they would be expected to follow to report. We got the distinct impression from multiple organizations that this was because they didn't *have* a plan, or only a very loosely defined one. There were exceptions: for example, the Silicon Valley-based software developer had multiple security and system

admin staff in a virtual “war room”, analyzing our initial, low-fidelity report within a couple of hours of receiving it (*and* it was late Friday afternoon on the US West Coast).

5. They thought we were selling something – Fear can be a sales tactic, and nothing strikes fear like a notification of a potential breach, or security issue. Unscrupulous sales teams use similar tactics often, so thresholds for ignoring such contacts are artificially inflated just due to the noisiness of it all.
6. They thought we were scammers – Claiming something bad has happened is often the first step in a con, so it's likely easy to dismiss.

In the end, it's unclear why many organizations responded so poorly, if at all. While we focus on security as practitioners – [WeLiveSecurity](#) after all – it's pretty clear this isn't top of mind for many organizations who simply want to use technology, but not necessarily focus on security. Some organizations engage managed service providers (MSPs) and hope for the best. That's a reasonable thing to do, but it requires a level of trust that your MSP is doing the right thing as well. We found that wasn't necessarily the case.

If anything, this whole process highlights the issues with third-party trust from a whole bevy of parties ranging from MSPs to “secure equipment disposal” companies, and others.

WHAT OTHER DEVICES MIGHT HAVE SECRETS?

It's not just routers: as mentioned earlier, all kinds of hard drives and removable media in the secondary market have already been investigated and found to be positively oozing the previous owners' most sensitive data, and there promises to be a proliferation of stored data on IoT devices throughout the corporate environment. If miscreants manage to exploit one of a family of IoT devices, it seems likely that they would be able to gather corporate secrets on the secondary market for a whole class of devices, and then sell that data to the highest bidder or do the exploiting themselves.

With less of a set perimeter these days, the network is more like a data swarm all around the corporate environment, so data will be widely available wherever an individual happens to be at the moment. Protecting what remains of that perimeter is difficult, but policing potential IoT breaches that allow corporate security failures will be exponentially more difficult in practice.

Whether it would have been more secure to keep data within your corporate walls in this case is unclear, but cloud providers are notoriously opaque related to security issues, except to possibly say “we're being very secure”, which is not an answer at all. In the end, you should expect that cloud providers care far less about your data than you do ... and will try far less hard to get it back should the inevitable happen.

BEST PRACTICES AND GUIDANCE

Everyone has a router, sometimes in production long after security updates are no longer available. But one day its time will come. When it does, what should you do to protect yourself while disposing of it and upgrading? Sometimes the original network administrator that set the unit up isn't even around anymore to assist. Hopefully the following observations and recommendations are helpful.

How should you dispose of your routers?

Each manufacturer has specific guidelines they suggest for restoring the devices to factory default. We researched three of the big vendors in the space; there are many others, each with their own processes for securely wiping configuration data, so consult the manufacturer's website for specific information on your router. For example, instructions to reset a Cisco ASA 5500 series device to its factory default configuration are available in [14], a Fortinet FortiGate series device in [15], and a Juniper SRX550 Services Gateway in [16].

The irony is that these devices are typically fairly simple to wipe, often with just a command or two. Some units, however, store historic configurations that may still be accessible, so you should carefully verify that there really is none of your information left on any of these devices.

Many of the devices we obtained for this research were still supported by the manufacturer, despite not being bleeding edge or current models. In some instances, it was challenging to obtain end-of-life information about the no longer supported devices. This was partially due to equipment manufacturers removing information about discontinued or end-of-life devices from the public internet, and where that was not the case putting device information behind subscription paywalls. And, even when information was readily available, the lack of common terminology and phrasing among the different equipment manufacturers made these instructions more difficult to follow.

For end users who are operating core networking gear (routers and switches), and other devices whose configuration may expose sensitive company information, we recommend doing the following:

1. Because the information for your device can go away at any time, we recommend that you save copies of all relevant information, such as manuals and knowledge base articles, while the device is still covered by its maintenance contract. Do this regardless of whether that information is available in a public forum or behind a subscription paywall. This includes:
 - copies of device firmware (all versions used in your organization),
 - copies of any software used with the device (all versions used in your organization),
 - any documents (bulletins, documentation, FAQs, knowledge base articles, etc.) you have previously referred to when configuring, troubleshooting, or updating the device,
 - any documents explaining how to decommission and securely wipe the device, and
 - information from all support tickets opened on the device, including symptoms, troubleshooting, and steps to resolve.
2. This information should be stored in a secure location on the company's network. This could be an intranet, wiki, network share, or whatever method works best for organizing the information.
3. For devices with support subscriptions, create 30, 60, and 90-day alerts in your calendar for when the subscriptions expire. Activities to occur on alerts:
 - 90 days out: Determine whether support subscription is to be renewed and begin the purchase process for it.
 - 60 days out: Begin collecting and organizing information mentioned in steps 1 and 2, above.
 - 30 days out: Begin testing procedures for securely wiping a device, verifying that company sensitive information is no longer present, and re-provisioning/staging it for use. In the event of any issues arising with any of these procedures, resolve them by opening a ticket with the manufacturer while the device is still under support contract.

For device manufacturers, we recommend the following:

1. Provide to the public, at no charge, full and complete instructions for securely wiping devices, even for devices that are no longer sold or supported, and even if the information was previously behind a paywall.
2. For secure-wipe instructions, use similar language and terms not just across product lines and models, but also across the industry so that the instructions are clear and unambiguous: remember that the person performing the wipe may not be as familiar with your vendor-specific terminology or equipment, and may make assumptions about how to securely wipe a device based on familiarity with a different vendor's tools and commands. Ensuring that your instructions are as similar and consistent as possible with those of other manufacturers will reduce the risk of sensitive data exposure due to wiping processes being applied incorrectly or incompletely.
3. In future designs, consider *only* using easily removable media to store configuration information. Depending upon the type and complexity of a device's configuration, this could be a CompactFlash

(CF) card, SDXC card, USB flash drive, SATA DOM, or even a 2.5" or M.2 SSD. The use of proprietary form factors is not recommended, as they decrease a legitimate operator's ability to confirm that the storage has been successfully wiped by removing it from the device and connecting it to a computer for external verification and remediation, if required.

Note that some Cisco devices have such an option, using a CF card, already; see the two lower devices in Figure 5. In the same image, the (upper) FortiGate unit has a Fortinet Storage Module option. Other manufacturers have multiple options for device and configuration storage.




Figure 5. CompactFlash card storage slot on two Cisco devices (lower) and the Fortinet Storage Module (FSM) slot on a FortiGate (upper)

Should you engage a data disposal or e-waste service?

Some of the original owners of the devices in our research had trusted third-party data disposal companies to securely wipe their devices. All we can say is that these devices still made it into our research white paper. Of course, for the minority of devices we purchased that had been wiped, we cannot tell if any of them had been handled by such third-party services or by their previous owners.

If you decide to use a third party, make sure you get some kind of certification that the work has been completed, with an option to audit their process on request. It's also good to know how they plan to dispose of your devices once wiped. Regardless, it's pretty straightforward to do a factory reset yourself.

What if the device is "dead"?

Recall that one of the 18 routers we initially purchased would not boot? It seems quite likely that was the reason it was disposed of and ended up in the secondary market.

However, were any steps taken to clean this device of its configuration data? Our results suggest that 56.25% of the time, the answer would be "no effective steps were taken"; we also took no further steps to check that specific device. If you have such a "dead but configured" device, how do you wipe your configuration data from a router that has no functioning console access?

One answer is to ensure that the whole device goes into the e-waste stream as an item that must be physically shredded. Other more complicated options depend on the device, but if you are sure that the only place that sensitive data is recorded is some form of removable storage medium – an internal hard drive, or internal or external removable storage media such as a CompactFlash card or similar – then physically separating the storage media from the router, and taking appropriate data wiping and disposal steps for that media should be sufficient.

What if you already have unsecured routers somewhere “out there”?

That’s a hard one since a targeted attack won’t be very noisy. If an attack methodology has highly specific information about your network, it would often suggest an insider threat, but this may be another attack vector to consider in your red/blue team exercises.

If you have a policy to rotate your cryptographic keys periodically, that can also help, but we didn’t assess whether this was happening frequently on devices we tested. In something like Active Directory there are ways to enforce password policy, but we didn’t see any examples of centrally managed keys. Many of the keys were used by things like IPsec for peering with other data centers or locations, which probably rarely – if ever – change. This means if someone has the configuration for a disposed router, they may be able to gain trusted access to other devices on your network with key reuse.

If you’ve already implemented Zero Trust in your organization, you’ll do better than other organizations. There are two organizations we researched on these devices that actually had reasonable controls in place and highly segmented networks, which is very good. Although the devices, complete with information, still left their organization, so it would still classify as a failure.

What if the previous owner’s information is on a secondhand device?

Should your organization purchase a pre-owned device and you discover the previous owner’s information is still on it:

1. Start taking notes. These could be paper, electronic, or both. Note the time and date the device was purchased, from whom and the order number, when it arrived, when you first accessed it, and what you discovered. These do not have to be down to the second, but with as many of the details as you can recall. This will be helpful if an investigation needs to be performed for insurance or legal reasons.
2. Disconnect the device from any network(s) it is attached to, even if they are only internal networks with no outside connection to the internet.
3. If the device is in a “public” area within your facilities, move it to a more secure area. This could be a locked office or a file cabinet.
4. Contact a CISA [regional office](#) if you are in the United States, or a similar agency for your country.

POTENTIAL FURTHER RESEARCH

As already described, we did not employ any procedures or tools of a primarily forensic or data recovery nature, nor did we open the cases of the devices purchased for this research. However, many of these devices contain internal hard drives, CompactFlash, or other common removable media devices, in internal slots. Simply opening the cases of apparently wiped devices, removing such storage media, and analyzing the devices with suitable forensic or data recovery tools might reveal how well – or how poorly – the standard wiping procedures are implemented. Devices from all three brands involved in the current research are based on common Unix-like OSes (Linux: Cisco IOS XE and FortiOS; Linux plus FreeBSD: JunOS; and QNX: Cisco IOS XR), so are likely to feature well-known file systems.

CONCLUSION

Our research clearly shows there are gaping security holes in modern IT environments that can produce data leaks that have the capability of allowing more serious data breaches, corporate theft, and other security risks with significant potential impacts to the organization, its customers, and the industry. Perhaps more concerning still is that these leaks *in the corporate sector* apparently are occurring at close to the rate of analogous leaks in the secondary consumer market for hard drives and other storage media. And that is despite the research documenting the latter failure to properly wipe such storage devices being covered repeatedly in the popular media.

Companies should take a prominent and direct role in protecting their data and devices within their responsibility, both when on premises, and when obsolescence or other imperatives dictate their disposal and movement outside the organization's boundaries.

Further, it is not enough to assume proper disposal by third parties; this must be verified to determine no future impact to the organization, and third-party equipment disposal companies need to up their game. Clearly, there is room for ongoing research in this space to determine whether such reforms actually happen, and whether more network configuration data might be being leaked if forensic tools are applied to the storage media of apparently wiped network devices.

REFERENCES

- [1] S. L. Garfinkel and A. Shelat, "Remembrance of data passed: A study of disk sanitization practices," *Security & Privacy, IEEE*, vol. 1, p. 11, 19 02 2003. [Also online: https://www.researchgate.net/publication/3437324_Remembrance_of_data_passed_A_study_of_disk_sanitization_practices, accessed 26 03 2023].
- [2] A. Jones, C. Valli, I. Sutherland and P. Thomas, "The 2006 Analysis of Information Remaining on Disks Offered for," *Journal of Digital Forensics, Security and Law*, vol. 1, no. 3, p. 14, 2006. [Also online: <https://commons.erau.edu/cgi/viewcontent.cgi?article=1008&context=jdfsl>, accessed 26 03 2023].
- [3] M. K. Pratt, "Have you resold your data to crooks?," 16 02 2007. [Online]. Available: <https://www.computerworld.com/article/2543619/have-you-resold-your-data-to-crooks-.html>. [Accessed 26 03 2023].
- [4] A. Jones, C. Valli, G. S. Dardick, I. Sutherland, G. Dabibi and G. Davies, "The 2009 Analysis of Information Remaining on Disks Offered for Sale on the Second Hand Market," *Journal of Digital Forensics, Security and Law*, vol. 5, no. 4, p. 22, 2010. [Also online: <https://commons.erau.edu/cgi/viewcontent.cgi?article=1083&context=jdfsl>, accessed 26 03 2023].
- [5] P. Bischoff, "Two-thirds of secondhand USB drives still contain previous owners' data: study," 19 03 2019. [Online]. Available: <https://www.comparitech.com/blog/information-security/secondhand-usb-drive-memory-stick-study/>. [Accessed 26 03 2023].
- [6] A. Jones, O. Angelopoulou and L. Noriega, "Survey of Data Remaining on Second Hand Memory Cards in the UK," *Computers & Security*, vol. 84, pp. 239-243, 07 2019. [Also online: https://uhra.herts.ac.uk/bitstream/handle/2299/22333/UK_Micro_SD_paper_2018_Submitted.pdf, accessed 26 03 2023].
- [7] O. Angelopoulou, A. Jones and G. Horsman, "A Study of the Data Remaining on Second-Hand Mobile Devices in the UK," *Journal of Digital Forensics, Security and Law*, vol. 17, p. 14, 27 10 2022. [Also online: <https://commons.erau.edu/jdfsl/vol17/iss2/5/>, accessed 26 03 2023].
- [8] L. Myers, "Target targeted: Five years on from a breach that shook the cybersecurity industry," 18 12 2018. [Online]. Available: <https://www.welivesecurity.com/2018/12/18/target-targeted-five-years-breach-shook-cybersecurity/>. [Accessed 26 03 2023].

- [9] C. Camp and T. Anscombe, "Supply-chain attacks: When trust goes wrong, try hope?," 07 04 2021. [Online]. Available: <https://www.welivesecurity.com/2021/04/07/supply-chain-attacks-when-trust-goes-wrong-try-hope/>. [Accessed 26 03 2023].
- [10] S. Cobb, "Desperately seeking cybersecurity skills," 26 01 2017. [Online]. Available: <https://www.welivesecurity.com/2017/01/26/desperately-seeking-cybersecurity-skills/>. [Accessed 26 03 2023].
- [11] (ISC)², "(ISC)² 2022 Cybersecurity Workforce Study," 18 10 2022. [Online]. Available: <https://www.isc2.org/-/media/ISC2/Research/2022-WorkForce-Study/ISC2-Cybersecurity-Workforce-Study.ashx>. [Accessed 26 03 2023].
- [12] ISACA, "State of Cybersecurity 2022 Part 1 Infographic," 16 03 2022. [Online]. Available: https://www.isaca.org/-/media/files/isacadp/project/isaca/resources/infographics/state-of-cybersecurity-2022-part-1-infographic_0322.pdf. [Accessed 26 03 2023].
- [13] ISACA, "State of Cybersecurity 2022," 03 2022. [Online]. Available: <https://www.isaca.org/go/state-of-cybersecurity-2022>. [Accessed 26 03 2023].
- [14] Cisco Systems, Inc., "Cisco Security Appliance Command Line Configuration Guide, Version 7.2," Cisco Systems, Inc., 13 10 2013. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/security/asa/asa72/configuration/guide/conf_gd/start.html#wpi055130. [Accessed 26 03 2023].
- [15] Fortinet, Inc., "Technical Tip: How to reset a FortiGate with the default factory settings/without losing management access," Fortinet, Inc., 07 09 2015. [Online]. Available: <https://community.fortinet.com/t5/FortiGate/Technical-Tip-How-to-reset-a-FortiGate-with-the-default-factory/ta-p/196896>. [Accessed 26 03 2023].
- [16] Juniper Networks, Inc., "Resetting the SRX550 Services Gateway," Juniper Networks, Inc., 13 09 2017. [Online]. Available: https://www.juniper.net/documentation/en_US/release-independent/junos/topics/task/operational/services-gateway-srx550-resetting.html. [Accessed 26 03 2023].

ABOUT ESET

For more than 30 years, ESET has been developing industry-leading IT security software and services to deliver comprehensive, multilayered protection against cybersecurity threats for businesses and consumers worldwide. ESET has long pioneered machine learning and cloud technologies that prevent, detect and respond to malware. ESET is a privately owned company that promotes scientific research and development worldwide.

