



Microsoft Digital Defense Report

OCTOBER 2021

CHAPTER 1

Introduction

- 3 Introduction
- 5 Our 2021 focus areas

CHAPTER 2

The state of cybercrime

- 8 Introduction
- 8 The cybercrime economy and services
- 10 Ransomware and extortion
- 20 Phishing and other malicious email
- 34 Malware
- 38 Malicious domains
- 42 Adversarial machine learning

CHAPTER 3

Nation state threats

- 48 Introduction
- 49 Tracking nation state threats
- 52 What we’re seeing
- 57 Analysis of nation state activity this year
- 69 Private sector offensive actors
- 69 Comprehensive protections required

CHAPTER 4

Supply chain, IoT, and OT security

- 71 Introduction
- 72 Challenges in managing risk associated with the supplier ecosystem
- 73 How Microsoft thinks about supply chain
- 76 IoT and OT threat landscape
- 81 The 7 properties of highly secured devices
- 82 Applying a Zero Trust approach to IoT solutions
- 83 IoT at the intersection of cybersecurity and sustainability
- 84 IoT security policy considerations

CHAPTER 5

Hybrid workforce security

- 89 Introduction
- 91 A Zero Trust approach for securing hybrid work
- 93 Identities
- 98 Devices/Endpoints
- 99 Applications
- 100 Network
- 104 Infrastructure
- 105 Data
- 106 People

CHAPTER 6

Disinformation

- 110 Introduction
- 111 Disinformation as an emerging threat
- 113 Mitigation through media literacy
- 114 Disinformation as an enterprise disruptor
- 117 Campaign security and election integrity

CHAPTER 7

Actionable insights

- 122 Introduction
- 123 Summary of report learnings
- 128 Conclusion

Contributing teams at Microsoft

Introduction

TOM BURT, CORPORATE VICE PRESIDENT, CUSTOMER SECURITY & TRUST

Over the past year the world has borne witness to a burgeoning cybercrime economy and the rapid rise of cybercrime services. We have watched this global market grow in both complexity and fervency. We’ve seen the cyberattack landscape becoming increasingly sophisticated as cybercriminals continue—and even escalate—their activity in times of crisis. New levels of supply chain and ransomware attacks were a powerful reminder that we must all work together, and in new ways, to protect the cybersecurity of the planet.

We see transparency and information sharing as essential to the protection of the ecosystem. Knowledge brings power, and to that end, security professionals need diverse and timely insights into the threats they are defending against.

Microsoft serves billions of customers globally, allowing us to aggregate security data from a broad and diverse spectrum of companies, organizations, and consumers. Informed by over 24 trillion security signals per day, our unique position helps us generate a high-fidelity picture of the current state of cybersecurity, including indicators that help us predict what attackers will do next. Our goal in creating the Microsoft Digital Defense Report is to bring together integrated data and insights from more teams, across more areas of Microsoft than ever before. We will share what we’re seeing to help the global community strengthen the defense of the digital ecosystem, and we will include actionable

learnings that companies, governments, and consumers can use to further secure individuals and environments.

The Microsoft Digital Defense Report draws on insights, data, and signals from across Microsoft, including the cloud, endpoints, and the intelligent edge.¹ Thousands of Microsoft security experts across 77 countries interpret and contribute to the insights gained from our advanced engineering and threat signals. Our security experts include analysts, researchers, responders, engineers, and data scientists. We also share lessons learned from customers transitioning to a hybrid workforce and frontline stories from our incident responders. Of course, there is malign activity we do not see, some of which is reported on by others in the industry. While the defender community at Microsoft works hard to identify threats and keep our customers informed, the bad actors are skilled and relentless.


By continually sharing insights we and others in the industry derive from the work we do, we hope to empower everyone to defend the online ecosystem more effectively.

Microsoft has made significant and ongoing investments to increase and improve the knowledge we derive from our threat signals. These investments deliver the highly synthesized and integrated insights that we share here. Our goal in aggregating these learnings is to help organizations understand the ways in which cybercriminals are continually shifting their modes of attack—and determine the best ways to combat those attacks. We write and share this report in the spirit of empowering the global community to benefit from the insights, observations, and transparency generated by our unique mission and vantage point.

THE MICROSOFT
DIGITAL DEFENSE
REPORT DRAWS ON
INSIGHTS, DATA,
AND SIGNALS
FROM ACROSS
MICROSOFT,
INCLUDING
THE CLOUD,
ENDPOINTS, AND
THE INTELLIGENT
EDGE.

Microsoft security signals

Volume and diversity of signals processed by Microsoft



July 1, 2020 through June 30, 2021

Our 2021 focus areas

2021 brought powerful reminders that to protect the future we must understand the threats of the present. This requires that we continually share data and insights in new ways. Certain types of attacks have escalated as cybercriminals change tactics, leveraging current events to take advantage of vulnerable targets and advance their activity through new channels. Change brings opportunity—for both attackers and defenders—and this report will focus on the threats that are most novel and relevant to the community as we look to the months ahead.

Looking at the threat landscape, along with data and signals from cross-company teams, five top-level areas emerged as most critical to bring into the sharpest focus in this report: the state of cybercrime; nation state threats; supplier ecosystems, Internet of Things (IoT), and operational technology (OT) security; the hybrid workforce; and disinformation. To provide the greatest benefit, we also extract our recommendations and actionable learnings, and present them throughout the report and in our concluding chapter.

The state of cybercrime

In this chapter, we discuss new developments in the cybercrime economy and the growing market for cybercrime services. We provide updates and analysis of what we are seeing in ransomware and extortion, phishing and other malicious email, malware, and the use of domains by cybercriminals, presenting recommendations for mitigating risk in each area. Finally, we share what we’re seeing in adversarial machine learning and what we are doing to stay ahead of cybercriminals in this area.

Nation state threats

This chapter provides an update on what we’re seeing in nation state adversarial activity, including reports on seven activity groups we have not previously mentioned publicly. We provide an analysis of the evolving threats in this watershed year with an increased focus on on-premises servers and the exposure of widespread supply chain vulnerabilities. We conclude with a discussion about private sector offensive actors and our guidance for comprehensive protections.

Supply chain, IoT, and OT security

The highly publicized events of the last year have made clear that securing and managing risks associated with supplier ecosystems is critically important. This chapter covers some current challenges in doing so in the supplier ecosystem and presents how Microsoft thinks about end-to-end supply chain security in nine investment areas. Then we turn our discussion to what we’re seeing in the Internet of Things (IoT) and operational technology (OT) threat landscape, with guidance on the properties of highly secured devices. We include specialized use cases of IoT and present some new research informing IoT policy considerations.

Hybrid workforce security

This chapter is about our greatest asset, our people. As we have moved to a hybrid workforce over the past year, we've seen developments in the threat landscape which point to the importance of adopting a Zero Trust approach. We include threat signals and other data across the six pillars of Zero Trust—identities, endpoints, applications, network, infrastructure, and data—and provide guidance based on what we're seeing. We conclude with discussions about insider threats in hybrid work environments, and an empathy imperative for managing the new and significant challenges encountered by today's workforce.

Disinformation

This chapter addresses the unprecedented disinformation campaigns and related cyber operations by state and non-state actors, impacting public awareness and knowledge as well as enterprise operations. We look at some parallels in cybersecurity and discuss mitigation through media literacy. We include a discussion on disinformation as an enterprise disruptor, providing a four-point plan for enterprise executives. The chapter concludes with an in-depth exploration of political campaign security and election integrity, two areas that have been targeted by disinformation campaigns.

Actionable insights

We open this year's concluding chapter with a discussion of five paradigm shifts that will center the evolution of work around the inclusivity of people and data. The chapter concludes with a distilled look at the key learnings from all the previous chapters of this report: to minimize impact of attacks we must truly practice good cyber hygiene, implement architectures that support the principles of Zero Trust, and ensure cyber risk management is built into the business.



CHAPTER 2

The state of cybercrime

Introduction

The cybercrime economy and services

Ransomware and extortion

Phishing and other malicious email

Malware

Malicious domains

Adversarial machine learning

INTRODUCTION: The growing threat of cybercrime

AMY HOGAN-BURNEY, GENERAL MANAGER, DIGITAL CRIMES UNIT

Cybercrime, whether nation state sponsored or permitted, is a threat to national security. Cybercriminals are targeting and attacking all sectors of critical infrastructure, including healthcare and public health, information technology (IT), financial services, and energy sectors. Ransomware attacks are increasingly successful, crippling governments and businesses, and the profits from these attacks are soaring.

The cybercrime supply chain, often created by criminal syndicates, continues to mature allowing anyone to buy the services needed to conduct malicious activity for financial gain or other nefarious purpose. Sophisticated cybercriminals are also still working for governments conducting espionage and training in the new battlefield.

It is not hopeless, and there are two positive trends we have seen recently. First, more governments and companies are coming forward when they are victims. This transparency helps in several ways. It has made clear to governments around the world that cybercrime is a threat to security. Victim stories humanize and make clear the consequences of these attacks, drawing attention to the problem and allowing increased engagement from incident responders and law enforcement. Second, now that governments around the world recognize that cybercrime is a threat to national security, they have made combatting it a priority. Governments around the world are passing new laws regarding reporting, creating cross-government task forces, allocating resources, and seeking out private sector assistance.

The cybercrime economy and services

Through our investigations of online organized crime networks, frontline investigations of customer attacks, security and attack research, nation state threat tracking, and security tool development, we continue to see the cybercrime supply chain consolidate and mature. It used to be that cybercriminals had to develop all the technology for their attacks. Today, they rely on a mature supply chain, where specialists create cybercrime kits and services that other actors buy and incorporate into their campaigns. With the increased demand for these services, an economy of

specialized services has surfaced, and threat actors are increasing automation to drive down their costs and increase scale. For example, we are seeing an increasing offer of backconnect proxies (proxies that rotate between mobile, residential, and datacenter systems) in addition to Remote Desktop Protocol (RDP), Secure Shell (SSH), virtual private network (VPN), virtual private server (VPS), web shells, cPanels (webhosting management dashboard), and other anonymization systems.

Other examples include selling compromised credentials that may have been obtained from phishing, scraping botnet logs or other credential harvesting techniques, imposter domain names, phishing-as-a-service, customized lead generation (for example, victims by country, industry, or roles),

WITH NO
TECHNICAL
KNOWLEDGE OF
HOW TO CONDUCT
A CYBERCRIME
ATTACK, AN
AMATEUR
THREAT ACTOR
CAN PURCHASE
A RANGE OF
SERVICES TO
CONDUCT THEIR
ATTACKS WITH
ONE CLICK.

loads (malicious software used to update malware on an infected computer), denial of service (DoS), and more. As an illustration, in some marketplaces, compromised credentials are offered by different sellers for \$1.00 USD to \$50.00 USD, depending on a variety of variables including the perceived value of the enterprise target. The number of sites offering services has significantly increased in the past 12 months as well as volume of credentials and variety of phishing kits.

Among the services available to even amateur threat actors are the cryptocurrency escrow services (to ensure services are rendered as offered) that we often see in commodity ransomware campaigns where affiliate models have become firmly established. Nontechnical cybercriminals sign up with a ransomware affiliate where for 30% of the revenue, the affiliate network will supply the ransomware, recovery services, and payment services. The attacker then buys “loads” from a market and pushes the ransomware to the loads they purchased. They then sit back and collect their revenue.


At times there are geographic groups of actors who may offer certain services, but most of these cybercrime markets are global in nature. A buyer in Brazil can obtain phishing kits from a seller in Pakistan, domains from the United States, victim leads from Nigeria, and proxies from Romania.

These prices have remained fairly steady over the past several years, but like any other market they vary according to changes in supply, demand, and externalities such as politics.

KEY TAKEAWAYS:

- **Identity and password/phishing attacks** are cheap, and on the rise. Why would an attacker break in when they can log in?
- **Distributed denial of service (DDoS) attacks are cheap** for unprotected sites—about \$300 USD/month.
- **Ransomware kits** are one of the many types of attack kits designed to enable low-skill attackers to perform more sophisticated attacks.

Average prices of cybercrime services for sale



Organizations now face an industrialized attacker economy with skill specialization and trading of illicit commodities. As seen in this snapshot of average prices, many commodities that can be purchased in the dark markets are very inexpensive, making attacks cheaper and easier to conduct (which also drives up attack volume).

Not all attacks work. It’s critical that we keep improving our defenses to increase the failure rate of attacks and the associated cost to attackers.

Ransomware and extortion

Ransomware basics and taxonomy

Ransomware and extortion is a high-profit, low-cost business which has a debilitating impact on targeted organizations, national security, economic security, and public health and safety. What started as simple single-PC ransomware has grown to include a variety of extortion techniques enabled by human intelligence and is affecting the networks of all types of organizations across the globe.

This combination of real-time intelligence and broader criminal tactics, techniques, and procedures

(TTPs) has maximized the impact of these attacks and driven the profits from these attacks to levels that were hard to imagine a few years ago. To put it in perspective, the publicly reported profits from ransomware and extortion attacks gives these attackers a budget *that would likely rival the budgets of nation state attack organizations* (without even counting the profits from attacks that never made the headlines).

To counter ransomware, a global collaborative effort between the private sector, law enforcement, and government is necessary to reduce the profitability of this crime, make it more difficult to enter the ransomware market, and supply victims with effective tools for efficient prevention and remediation. Microsoft is a contributor to the Ransomware Task Force report, a comprehensive

framework designed for taking action in combatting ransomware.² Microsoft has also published a project plan with links to technical guidance to help organizations better prepare for and respond to these attacks and is contributing to a National Institute of Standards and Technology (NIST) publication containing a cybersecurity framework profile for ransomware risk management.³

A ransomware and extortion attack involves a threat actor deploying malware that encrypts and exfiltrates data and then holds that data for a ransom, often demanding payment in cryptocurrency. Rather than just encrypting a victim’s files and requesting a ransom in exchange for the decryption key, the attackers also exfiltrate sensitive data before deploying the ransomware. This practice prevents victims from disengaging from

negotiations and raises the victim’s reputational costs of not paying the ransom as the attackers likely will not only leave the victim’s data encrypted but also leak sensitive information.

A series of criminal activities occur long before the ransomware is ultimately deployed across computer systems in an organization. As a result, we created a taxonomy that focuses on the relationship between entities within the ransomware ecosystem because any entity may play a different role at any given time.

Ransomware attacks have evolved into human-operated ransomware, also known as “big game ransomware.”

Ransomware taxonomy

Primary role	Description
Develops	Writes the malware
Deploys	Sends phishing emails, deploys ransomware
Provides access	Malware that loads other malware, or a group that sells access as a service
Manages/operates	Leadership of a group (such as MAZE cartel membership) and/or function that provides coordination (such as managing or operating a central extortion leak site)
Publicly reported connection	A publicly reported connection exists

For example, as shown in the image at right, a threat actor may develop and deploy malware that gives one threat actor access to a certain category of victims, whereas a different threat actor may merely deploy malware.

Post-breach response


Just as the criminal enterprise that deploys ransomware typically involves several stakeholders each with a particular responsibility, the response to ransomware also involves several key stakeholders.

If a victim of a ransomware attack has cyber insurance, that carrier will employ certain service providers, including an incident response firm, a law firm, and an organization specializing in ransom negotiation. Even if a victim does not have a cyber insurance policy, these stakeholders are common to finding a resolution to the ransom.


Once a ransomware gang locks a victims’ network, exfiltrates data, and holds the network and data for ransom, an incident response team will investigate the root cause of the breach and drive remediation efforts depending on the victim’s level of preparedness prior to the attack. If the victim has sufficient backups of its data or data has not been stolen, often the incident response team will work to remove the threat actor from the victim’s system, restore business operations, and apply future mitigation measures. The incident response team

will often provide the victim a report which includes root cause, criminal actor movement inside the victim network, data exposure and exfiltration, and remediation recommendations.

Depending on the jurisdiction of the victim, the victim could be subject to data breach notification requirements. A law firm will often assess the exposure of the victim’s liability and assist the victim with meeting its regulatory obligations. Importantly, the law firm will interface with relevant law enforcement, where appropriate. Finally, if a victim is unable to return to business operations, an organization specializing in negotiating with ransomware criminal syndicates will work to obtain the decryption key on behalf of the victim.




Sample analysis of roles and relationships between entities within the ransomware ecosystem



Ransomware syndicates and affiliates are all working together toward these interconnected threats. Rather than one individual behind a ransomware attack, there are multiple groups of individuals, similar to a shared business model.

Stakeholders and roles involved in post-breach response




Criminal economics: A changing business model

The business model for ransomware has effectively evolved into an intelligence operation; criminal actors perform research on their target victim to identify an optimal ransom demand. Once a criminal actor infiltrates a network, they may exfiltrate and study financial documents and insurance policies. They may also understand the penalties associated with local breach laws. The actors will then extort money from their victims, to not only unlock their systems, but also to prevent disclosure of the victim’s exfiltrated data to the public. After they’ve collected and analyzed this intelligence, the criminal actor will identify an “appropriate” ransom amount.

The negotiation chat, at right, with a public school district to extort cash in exchange for a decryption key to unlock the Conti ransomware deployed on its network demonstrates the research performed by the criminal in advance of the negotiation. Here, the criminal actor explains “we examined all financial documents, bank statements for the last year, insurance. And came to the conclusion that you are exaggerating about poor financial condition [sic]. We also calculated your possible losses from lawsuits from both your staff and your students for the leakage of their personal data. These fines will exceed \$30 million. We are not talking about the loss of reputation, which in our opinion costs more.”

Ransomware negotiation chat



THE
RANSOMWARE
ENTERPRISE HAS
EVOLVED INTO
RANSOMWARE AS
A SERVICE DRIVEN
BY HUMAN
INTELLIGENCE
AND RESEARCH.

There are few barriers of entry into this criminal enterprise. A cybercriminal does not need specialized code development skills to profit from this crime. The ransomware enterprise has evolved into ransomware as a service driven by human intelligence and research. It is no longer solely the province of malware developers; rather, the business structure is modular. Malware developers are recruiting hackers with access to networks promising a “cut” of the profit. Criminals can purchase malware and access to specific networks and target specific industries. This is effectively a crime syndicate where each member is paid for a particular expertise.

In the example shown below, following the crypto flows, we can see where a criminal enterprise split its bitcoin “earnings” such that approximately 15% of the earnings flowed to the developer/manager and 75% of the earnings flowed to the attacker.

Regardless of where ransomware is deployed, typically the threat actors will demand payment via cryptocurrency through crypto wallets. Although the underlying blockchain technology facilitates transparent cryptocurrency flows, the owners of wallets remain pseudonymous. Nonetheless, they still need to find on- and off-ramps into the crypto ecosystem. At its core, the criminal actor needs to append the blockchain with a transaction and ultimately find a way to cash out. There are several stakeholders within the cryptocurrency ecosystem that facilitate ransom-related transactions and payments. These intermediaries often exist in jurisdictions with governments that are historically unwilling to cooperate with the United States and others. It’s these intermediaries that facilitate the flow of ill-gotten earnings from ransomware. The private sector through civil litigation, and the government through prosecution, regulatory enforcement, and international collaboration, can take coordinated action against intermediaries to disrupt the payment process.

SIDEBAR: TO PAY, OR NOT TO PAY?

In the aftermath of a ransomware attack, companies are often completely offline—their security systems tampered with, their backup systems deleted, their data encrypted, and their users unable to log in. When operations are offline and losses pile up, it is important to remember that paying the ransom demands does not guarantee the restoration of operations, nor does paying prevent future attacks.


In addition, we have in effect, the classic “Tragedy of the Commons”⁴— while it may make sense for individual victims to pay for their own individual benefit (restore critical business operations), the payment also contributes to the growth of this damaging model for everyone. Ransom payments can keep the cycle turning, as described below:

- The business model of extortionists gets reinforced, which also attracts more bad actors into the monetization strategy. Substantial revenue is supplied to the actors who then use part of it for research and development (R&D)

to improve their tooling and ability to buy breach access to potential victim organizations. Some ransomware teams have significant amounts of funds for R&D and for buying high-end 0-days. For example, some ransomware teams have budget to spend up to \$1 million USD, or more, per 0-day. While some high-end ransomware teams buy 0-days, others focus on traditional ways to gain remote access into victims’ networks.

- The ransomware tools become more automated and effective, allowing the bad actors to scale and accelerate their attacks—with more sophistication and less effort.

Transaction hashes and wallet addresses




Victim ransom wallet is shown on the far left. The funds split to two different wallets on the far right. Text intentionally blurred for publication

Important details to be aware of when making the decision whether to pay ransoms:

- On average, organizations that paid the ransom got back only 65% of their data, with 29% getting back no more than half their data.⁵
- Ransom decryptors are buggy and regularly fail to decrypt the largest, most critical data files (files 4 GB+ in size).
- Decrypting data files is a slow and labor-intensive process, most customers decrypt only their most critical of data files and restore the rest from backup.
- Restoring data does not undo any tampering performed by the attackers.
- Restoring data does not secure systems to prevent future attacks.
- Organizations must understand the legality of making payments in their country. Governments across the globe are instituting ransomware payment reporting requirements, may have penalties for payments that are made to sanctioned parties, and are considering laws that could make ransom payments illegal.

Paying a ransom fuels the ransomware syndicates



Paying a ransom gives the criminals more resources to expand their operations, helping them become more organized and specialized. With more funding available, the groups can improve their tools and code, enabling ransomware to spread through networks undetected by antivirus software.

EXAMPLE: CONTI RANSOMWARE

The Conti ransomware first appeared around July 2020 adopting the double extortion business model. In this double extortion model, a victim is first extorted for ransom and for the possible publishing of their stolen data. Conti is also a ransomware as a service (RaaS), which is a subscription-based service allowing affiliates of the service ready access to ransomware-building tools and builds. Affiliates of the service agree to ransom percentage payouts between the ransomware developer and threat actor who performed the exploitation. Conti usually gains access to the victim network via other threats like Trickbot, IcedID, or Zloader. Once inside the victim network, Conti has a configurable reconnaissance module where it can scan internal networks looking for network shares and other high-value targets. Once deployed, Conti begins to encrypt user-modifiable data and databases based on targeted file extension lists. Upon completion of the encryption, a ransom note is left in every file directory with instructions for the user on how to contact the ransomware actors:

Ransom note

```
All of your files are currently encrypted by CONTI strain.

As you know (if you don't - just "google it"), all of the data that has been encrypted by our software cannot be recovered by any means without contacting our team directly.
If you try to use any additional recovery software - the files might be damaged, so if you are willing to try - try it on the data of the lowest value.

To make sure that we REALLY CAN get your data back - we offer you to decrypt 2 random files completely free of charge.

You can contact our team directly for further instructions through our website :

TOR VERSION :
(you should download and install TOR browser first https://torproject.org)

http://conti[REDACTED].onion/


HTTPS VERSION :
https://contirecovery.top/

YOU SHOULD BE AWARE!
Just in case, if you try to ignore us. We've downloaded a pack of your internal data and are ready to publish it on out news website if you do not respond. So it will be better for both sides if you contact us as soon as possible.


---BEGIN ID---
[REDACTED]
---END ID---
```

The user must now upload the ransom note text file to the recovery site listed in the ransom note. The note serves as proof of encryption and victim identification for the ransomware actors.

Ransom recovery site




Conti News site



After the ransom note is uploaded and verified and a chat session is initiated.

Chat session following ransom note upload



The negotiation phase starts with the threat actor, as they prove they can decrypt any files provided by the victim. After a final ransom price is negotiated, the ransomware actor provides a Bitcoin wallet address for the victim to send payment. Conti ransomware actors maintain recovery and news sites on regular top-level domains (as on the open web) as well as on the dark web or Tor (also known as The Onion Router).


As part of the double extortion business model, the actors behind Conti maintain a news site, which serves as the publishing site to prove that if the ransom is not paid, the victim’s private information will be posted publicly and could be sold on the black market. The Conti News site currently lists hundreds of victims with various samples of their private data.

Conti victims are located mostly in the United States and Europe and include public schools, healthcare providers, manufacturing companies, US city governments, and even public utility providers.

What we’re seeing in ransomware data and signals


DEFENDER SIGNALS

Ransomware encounter rate (machine count): Enterprise customers




These charts show the overall increase in ransomware encounters, with notable surge to consumer and commercial encounters in late 2019,⁶ when RaaS started to grow, and in early 2020 at the onset of the COVID-19 pandemic.⁷

Ransomware encounter rate (machine count): All customers



Ransomware machine counts by country (July 2020-June 2021)




The stakes have changed. There is a massive growth trajectory for ransomware and extortion.

DART DATA

While the Colonial Pipeline ransom attack of May 2021 drew considerable public attention, our Detection and Response Team (DART) ransomware engagement data shows that the three most targeted sectors were consumer, financial, and manufacturing. Despite continued promises from ransomware actors not to attack hospitals or healthcare companies during a pandemic, healthcare remains in the top-five sectors victimized by human-operated ransomware.

DART ransomware engagements by industry (July 2020-June 2021)



SUMMARY OF RECOMMENDATIONS

The stakes have changed. There is a massive growth trajectory for ransomware and extortion. To help protect your organization from ransomware, we recommend that organizations:

Deploy ransomware protection

1



Prepare a recovery plan
Recover without paying

2



Limit the scope of damage
Protect privileged roles

3



Make it harder to get in
Incrementally remove risks

Prepare a recovery plan by making it harder to access and disrupt systems, which minimizes the monetary incentives for ransomware attackers and makes it easier to recover from an attack without paying the ransom.

Limit the scope of damage by forcing the attackers to work harder to gain access to multiple business-critical systems. Establish least-privilege access and adopt Zero Trust principles. These steps make it harder for an attacker who gets in to a network to travel across the network to find valuable data to lock up. Also, encrypt data at rest, and practice good backup-and-restore hygiene. This way, even if data is stolen it will be encrypted and not very useful to the attackers. In the unfortunate event that the attacker does encrypt your data, you will have a good backup to restore from and use to maintain business continuity.

Make it harder to get in by following basic cybersecurity hygiene steps that make it more difficult for attackers to gain access to the network. The most important of these steps is the use of multifactor authentication (MFA), which is important to raising friction for entry but will take time to complete as part of a larger security journey. Other steps, such as keeping up to date on patching and correct configuration, can be taken to identify and close off vulnerable entry points.

Use the phases as a starting plan for what to do first, next, and later to get the most impactful elements first. These recommendations have been prioritized using the Zero Trust principle of assume breach, which focuses on minimizing business risk by assuming the attackers can successfully gain access to your environment through one or more methods.

Microsoft supports the guidance presented in the Ransomware Playbook by the Cyber Readiness Institute.⁸

Learn more:

- [3 steps to prevent and recover from ransomware | Microsoft Security blog](#) (9/7/2021)
- [Rapidly protect against ransomware and extortion | Microsoft Docs](#) (8/24/2021)
- [Azure Sentinel Fusion Detection for Ransomware \(microsoft.com\)](#) (8/9/2021)
- [The growing threat of ransomware - Microsoft On the Issues](#) (7/20/2021)
- [Human-operated ransomware | Microsoft Docs](#) (5/27/2021)
- [Ransom mafia analysis of the world's first ransomware cartel pdf \(analyst1.com\)](#) (4/7/2021)
- [Ransomware Playbook - Cyber Readiness Institute](#)

Phishing and other malicious email

Threat to identity

In 2020, the FBI IC3 Report⁹ identified phishing as the top crime type for victim complaints. The number of reports doubled compared to the previous year. Phishing poses a significant threat to both businesses and individuals, and credential phishing was leveraged in many of the most damaging attacks last year.


From our investigations on online organized crime networks involved in business email compromise (BEC), we noted broad diversification of how credentials are obtained, verified, and later used that may explain the increased threat. Threat actors are increasing their automation and purchasing tools to increase the value of their criminal activities. Credentials belonging to unsuspecting victims could be obtained from phishing websites that impersonate a myriad of online services, automatically scraping and parsing logs belonging to infected devices that record the keys typed on keyboards to guessing where credentials from one breached online service were reused on another.

For each credential, there are services that enrich the information on the identity with additional details on the person's identity that includes name, company they work for, roles, seniority in company, and industry associated to the company. With this information, the identity could be used in BEC attacks, to send unsolicited messages (spam), to gather sensitive information, or to host phishing websites in related online accounts. Even when one attack occurs, accounts may be resold after automated systems verify that they remain compromised.


Identity is further threatened by impersonation as may be seen in BEC attacks where one party to a financial transaction is impersonated to divert payments to an unauthorized recipient. Our investigations identified that threat actors would monitor financially inclined messages to find an identity to impersonate and thereafter register homoglyph/imposter domains to resemble the email of the person being impersonated. In this case, the person whose credentials were stolen would cause another person to become a victim.



Emails determined as phish



Top 10 verticals affected by phishing (Defender detections, June 2021)



What we’re seeing

Types of malicious emails

Whether their goal is to phish credentials, redirect a wire transfer to their own bank account, or download malware onto a machine, attackers are most likely to utilize email as their initial entry vector for a campaign. While a lot of focus is given to credential phishing, malicious emails are used in multiple types of cyber incidents. Microsoft security researchers observe the following three most common types of malicious emails:

PHISHING

Phishing is the most common type of malicious email observed in our threat signals. These emails are designed to trick an individual into sharing sensitive information, such as usernames and passwords, with an attacker. To do this, attackers will craft emails using a variety of themes, such as productivity tools, password resets, or other notifications with a sense of urgency to lure a user to click on a link.

The phishing webpages used in these attacks may utilize malicious domains, such as those purchased and operated by the attacker, or compromised domains, where the attacker abuses a vulnerability in a legitimate website to host malicious content. The phishing sites frequently copy well-known, legitimate login pages, such as Office 365 or Google, to trick users into inputting their credentials. Once the user inputs their credentials, they will often


be redirected to a legitimate final site—such as the real Office 365 login page—leaving the user unaware that actors have obtained their credentials. Meanwhile, the entered credentials are stored or sent to the attacker for later abuse or sale.

Attackers also use consent phishing to send users links that, if clicked, will grant the attacker access and permissions to applications, such as via OAuth 2.0 authorization protocol. In these instances, users may unwittingly grant the attackers permissions to applications that enable them to access a wealth of sensitive information.

The number of phishing emails observed in Microsoft Exchange global mail flow increased during the period from June 2020 through June 2021. We saw a pronounced surge in November potentially related to holiday-themed phishing, and a subsequent decrease over the US winter holidays, potentially indicating that attackers send fewer messages when many people are not working.


While all industries receive phishing emails, some verticals within those industries tend to receive more phishing campaigns than others. The verticals most affected by phishing may change month to month depending on several factors, including attacker objectives, availability of leaked email addresses, or current events regarding specific sectors and industries.

Malware emails per week



There has been an overall downward trend in the number of emails containing malware.

Top 15 Defender detections (June 2021)



Spyware designed to steal credentials was the most common type of malware observed through email delivery and was detected three times as often as the next highest detection.

MALWARE DELIVERY

Malware delivery is another example of how threat actors utilize emails for their objectives. A variety of malware variants, such as Agent Tesla, IcedID, Trickbot, and Qakbot, use email as a primary method of distribution. These emails will use either links or attachments to deliver malware and many times use techniques that overlap with phishing emails. For example, both malware delivery email and phishing email may use links that direct to a CAPTCHA test to evade detection from security technologies.

Since malware does not rely on user interaction in the same way phishing does, attackers can design their delivery to be less noticeable to the user. For example, when using attachments as a delivery method, attackers may use a decoy document with macros that, when enabled by the recipient, download the malware in the background without the user's knowledge. In these cases, the user may think that the document is broken or isn't intended for them and may be completely unaware that malicious software is running on their machine.

One of the most common methods of malware delivery observed in the past year was through password-protected archive files. These emails contain archive files, such as ZIP attachments that are password protected, to prevent security technologies from detonating and analyzing them. However, the passwords for these files are often included in the body of the email to enable

the recipient to open the files and download the malware. By using these archive files to house the malicious document—frequently an Excel or Word document—the attackers can use a unique archive file for every recipient, making it more difficult for defenders to fully scope a campaign.

Interestingly, between July 2020 and June 2021 we observed an overall downward trend in the number of emails containing malware, indicating that attackers may be using other means of entry. In addition, a few notable malware takedowns—namely, Trickbot and Emotet—may have contributed to this overall decline. A large spike in October is associated with the distribution of those malware variants, and the rapid decrease following the spike aligns to when the Trickbot malware was taken down by Microsoft.

The malware that attackers distribute via email changes regularly for a variety of reasons, including malware takedowns and attacker objectives. As shown in the chart on Defender detections for June 2021, the most prolific malware observed by Microsoft was Agent Tesla, which is a credential-stealing spyware. The second most observed malware, Tisifi, which identifies social engineering lures, was seen only one third as much as Agent Tesla. EncDoc and CVE-2017-11882 as the third and fifth top detections indicate that attackers still favor malicious documents as a common method of delivering a variety of threats. The fourth top

detection, HTML/Phish, includes only phishing emails that use an HTML attachment. These types of phishing frequently take the form of fake voicemail phishing messages.

BUSINESS EMAIL COMPROMISE

While not the most prolific type of malicious email in terms of quantity, BEC has proven to be the most financially impactful type of cybercrime.¹⁰ BEC occurs when an attacker pretends to be a legitimate business account—utilizing either a compromised email address, a lookalike domain they have registered, or a free email service such as Hotmail or Gmail—and sends emails designed to trick recipients into taking some financial action, handing over sensitive information, or providing assets, such as gift cards, to the attacker.

The most common type of BEC observed by Microsoft in the past year was gift card scams. In these scams, attackers will usually create a multitude of free email accounts, changing the display name depending on the target, though attackers have also registered their own domains for these attacks or have created target-specific free email accounts. They will then pretend to be someone the recipient works with (usually their boss or an executive at their company) and ask them to purchase gift cards (often with company funds). Frequently, these emails suggest that the sender wants them for a family

member’s birthday gift or for employees as rewards. The recipient is typically asked to send the digital gift cards to the attacker once purchased, but we have also seen attackers asking the user to buy physical gift cards and send a photo of the code on the back of the card, enabling the attacker to resell them online or trade them for cryptocurrency.

A much more sophisticated and financially damaging type of BEC is wire transfer fraud. In this type of BEC, actors will insert themselves into expected financial transactions and ask the recipient to adjust the bank account information on an outgoing wire transfer. The actors will masquerade as the intended recipient of the funds, so this does not seem out of the ordinary to the victim. Once the victim wires the money to the new account, it is withdrawn by the actors and may be difficult to retrieve. Companies can help to avoid this type of scam by ensuring that financial policies require verification for changing accounts. Finance employees should verify via a means other than email—such as from a known, trusted phone number with the recipient—before making account number changes that originate from emails. Additionally, utilizing impersonation protections features in email security products can help prevent attackers from successfully conducting this type of scam.

Detecting web-based phishing

In the past year, web-based phishing attacks have continued to become more sophisticated. Phishing kits used by web-based phishing attacks typically use images, context-based content, and other advanced techniques to avoid detection. Our machine learning (ML) models and network heuristics must continuously evolve to maintain effective protection. The language used by attackers has also improved significantly; **past user guidance to look for poor spelling and grammar is now less effective, particularly against targeted, more advanced attacks. Modern kits are sufficiently sophisticated to masquerade as legitimate content in their use of spelling, grammar, and imagery.**

Phishers are increasingly leveraging legitimate infrastructure, but this pattern still accounts for a minority of detected phishing attacks. Microsoft SmartScreen detected more than a million unique domains used in web-based phishing attacks in the last year, of which compromised domains represented just over 5%. This 5% of domains typically host phishing attacks on legitimate websites without disrupting any legitimate traffic so that their attack remains hidden for as long as possible.

Domains created specifically for attacks tend to be active for shorter periods, and with fewer malicious URLs, than attacks that abuse legitimate infrastructure. Over the last year, Microsoft SmartScreen has seen an increase in attacks that begin and end within as little as an hour or two.

Malicious email techniques

Attackers have adapted over time to make their emails more likely to evade detections and protections by utilizing aspects of legitimate business emails. Defenders need to protect the company but also have a duty to maintain the flow of business—and attackers rely on this fact to get their foot in the door. In the last year, Microsoft security researchers have observed attackers using numerous techniques across multiple malicious email campaigns to make emails appear more legitimate to both end users and protection technologies.

Malicious email techniques




Some common techniques observed over the past year:

COMPROMISED SENDERS > COMPROMISED SERVICES

For years, attackers have used compromised senders to perpetuate phishing email chains, as they use the victim’s email account to send additional phishing emails. While this is still extremely prevalent, many companies have begun utilizing MFA, which reduces the effectiveness of this method. Attackers therefore are adjusting their methods to begin compromising entire email services, such as when NOBELIUM gained access to an email marketing solution that enabled the attacker to send as multiple, legitimate addresses.¹¹


Phishing email using compromised service on behalf of legitimate companies



ABUSE OF LEGITIMATE INFRASTRUCTURE

Defenders have often told their end users to verify aspects within an email were legitimate before interacting with that email, such as the sender and any links within the email. This advice is still valuable, but sometimes the links and senders can look legitimate but contain malicious content. Attackers are shifting to abusing legitimate infrastructure to mask the malicious content in their emails. For sender addresses, attackers may register trial tenants for services such as Office 365, which make their email appear much more legitimate. Additionally, attackers are using ways to mask the malicious domain in an email, either by using open redirects from legitimate domains or by abusing legitimate hosting platforms such as Google Drive or OneDrive. In these cases, it may be tricky for users to know when an email is legitimate or malicious.

Legitimate infrastructure abuse




Attackers abuse legitimate contact forms on websites to send emails, and legitimate Google sites to host malware.

<https://t.me/learningnets> ¹¹<https://www.microsoft.com/security/blog/2021/05/27/new-sophisticated-email-based-attack-from-nobelium/>

FAKE REPLIES

In addition to users being instructed to verify aspects of the email, such as the sender or links before interacting with them, users have also been instructed not to interact with emails they are not expecting to receive. This is still extremely valuable advice, but attackers are aware of this as well and have shifted their strategies to find ways to convince recipients that they are expecting the email. One way that they do this is by crafting fake reply emails. In these cases, the attacker will take the contents of a previous email from a compromised mailbox, or will craft an entirely new email, and include it in the body of the email in a way that appears that the new email is a reply. Users who have jobs that require them to email dozens of people per day may not remember each email they have sent. Seeing a fake reply may convince them that they are expecting the email and cause them to interact with malicious links or attachments. Utilizing email security features that can notify a user when an email is being sent from a user they have not interacted with before can help mitigate this technique. This technique is a favorite of malware variants such as Emotet and IcedID and is also frequently used in BEC emails.

Fake reply email




Gift card scam BEC email using fake reply to trick the user into thinking they were expecting this email.

DEFENSE EVASION

While attackers are focusing their techniques on convincing the recipient to interact with an email, they are also aware that all that effort will be worthless if the email is never delivered to the victim. Because of this, threat actors are developing new means of defense evasion in email. It used to be enough for an attacker to include a password-protected archive file to evade detection, but most security technologies can now input passwords included in the email to detonate them and identify malicious content. Attackers have shifted to including CAPTCHAs and legitimate login screens for services such as Microsoft or Google that prevent detection technologies from reaching the malicious content.

The digital journey of stolen credentials

Where do your stolen credentials go after they are entered into a fake web page?




Secret phishers: The hidden economy of sophisticated phish kits

It is a long-held perception among research and business communities that victim credentials are delivered to the individual (or group) operating phishing campaigns. Researchers within the security community¹² have begun to identify more sophisticated kits in which not only are victim credentials sent to the phishers running a phishing campaign, but they are also likely going back to the kit’s originating author or a sophisticated intermediary who has modified the kit with a hidden collection account before redistributing the kit. The Microsoft Digital Crimes Unit (DCU) has seen several variations of this technique.

This stratification of the cybercrime world is an increased threat to business infrastructure as phish kit authors are more technically skilled than the phishers and have a much broader reach by orders of magnitude. The anonymity of the dark web enables this technique. Phishing has generally been known to be a scheme in which a phisher (or group of phishers) buys a kit on a dark web market, obtains infrastructure components such as a server, a domain to host the imitation site, and an email account or other endpoint to receive victim information. Once the infrastructure is assembled, they essentially just “stick their poles in the water”—the entire process is designed to be as easy as possible by the authors and distributors of kits. Although it is important not to overgeneralize, these


plug-and-play phishers have largely been viewed as lower-level coders compared to the much more sophisticated phish kit author and mastermind of the overall operation.

<https://t.me/learningsnets> <https://t.me/learningsnets/blog/phishing-kit-exfiltration-methods>



Phish kits and credential harvesting

Phish kits: enabling credential harvesting at scale



Multi-factor authentication: using MFA protects against the risks associated with credential reuse

Phish kit creator writes code that allows phish kit to be configured by kit buyer to indicate collection account where phished credentials are sent. Also included in code is a hidden collection account that will also receive phished credentials.

Phish kits are sold on the dark web. Each kit buyer configures the kit to meet their phishing campaign needs, including their own collection account to receive phished credentials.

Who's phishing whom?

Kit creators have expertise and resources to carry out more sophisticated and targeted attacks at scale.

Each kit buyer deploys their own campaign. Phished credentials are delivered to both the kit buyer and the kit creator.

Lists of newly harvested credentials feed more targeted attacks at scale.

Even well-protected organizations can become victims of more costly attacks exploiting credential reuse if not using MFA.

What's in Microsoft's phishing defense toolbox? How we approach employee awareness

In 2020, the industry saw a surge of phishing campaigns that has remained steady throughout 2021. Internally at Microsoft, we saw an increase in overall number of phishing emails, a downward trend in emails containing malware, and a rise in voice phishing (or vishing).


Fortunately, we were prepared with an effective foundation of protective controls to reduce the number of successful phishing attempts, and acknowledging the evolving threat landscape, we had expanded our controls to cover other vectors that could be exploited (beyond email, such as Forms and Teams).

What's in our toolbox

There is not a silver bullet fix for phishing; it must be solved through a multipronged approach. We focus on four primary elements: Protective controls, User awareness, Reporting and insights, and Detect and respond.


With these methods, we have seen a 50% year-over-year reduction in susceptibility.

Trends at Microsoft, 2020-2021



No matter how many protective controls we have in place, mitigation remains a crucial element to combating the phish that make it through our defenses as threat actors increase their level of sophistication. Our security operations center is equipped with Microsoft Defender for Office 365’s tools and automation to quickly detect, investigate, and effectively remediate malicious emails. For us, the automated incident response features have been key in enabling our team to move quickly because **minutes matter**.

Training followed by simulations, reinforcement, and targeted simulations



Despite the increased sophistication of phish, susceptibility of employees has decreased, attributed to increased frequency of simulations and training.

Beyond detection and protection, it is imperative that we cultivate a security-conscious culture, equip our employees (our last line of defense) with skills to identify a phish, and provide a simple reporting mechanism that is a consistent experience across all platforms. Employee reporting is vital, but closing the feedback loop by validating the report is just as important.

Beyond detection and protection, it is imperative that we cultivate a security-conscious culture.

Our approach to employee phishing awareness includes annual foundational training, simulated exercises, and positive reinforcement. Simulations leveraging Microsoft Defender for Office 365’s Attack Simulator and Training are built on incidents insights to ensure we are exposing our employees to phish that are realistic to the level of sophistication we may see in our environment. Employees who repeatedly fall susceptible to simulations are phished on a more frequent basis to increase their opportunity to learn through experience and preventative guidance. We also run targeted campaigns focused on high-risk groups such as new employees, executives, and

their support staff. The findings from our simulated exercises are also leveraged to identify opportunities within the product to aid employees in identifying phish (such as safety tips).

Learn more:

[Automatically triage phish submissions in Microsoft Defender for Office 365 \(9/9/2021\)](#)

Employee phishing awareness

Simulations

Our approach to simulated exercises

- ✓ Combination of company-wide and targeted simulations.
- ✓ Construct mails that match the level of sophistication that is trending in the wild.
- ✓ Gain insights into more than just employee behavior.
- ✓ Use previous exercise results and real-world incident trends to determine our next move.

The data we collect and track

- ✓ Susceptibility (user action)
- ✓ User reporting of simulated mail
- ✓ Time to report simulated mail
- ✓ Channels used for reporting

Training

How we build skill and increase awareness

- ✓ Leverage large scale training programs for in-depth education.
- ✓ Use insights from simulations and incidents to focus content.
- ✓ Drive training reinforcement through focused modules to advance skill level.
- ✓ Keep the conversation going through awareness campaigns.


How we measure training effectiveness

- ✓ Module scoring
- ✓ Reduction of susceptibility
- ✓ Time to report
- ✓ Channels used for reporting

User Reporting

How we empower users to take action

- ✓ Ensure our community knows when and how to report.
- ✓ Make reporting as quick and easy as possible and expand across platforms.
- ✓ Reinforce reporting behavior by closing the feedback loop.
- ✓ Educate employees on how reporting helps our system get better!



Summary of recommendations about malicious emails

1. Use MFA to lessen the impact of credentials being phished by attackers.

2. Develop robust user education processes that use positive reinforcement to teach users how to identify potentially malicious emails. Create processes wherein users can report suspicious emails and can receive feedback on whether the email they submitted was indeed malicious. Focus extra training on groups that may be more heavily targeted, such as executives, executive assistants, and finance employees. Share real-world phishing examples that your company has received with your end users so that they understand the threat and know what to expect.

3. Surface external emails to recipients by appending a tag to the subject line of any email that originates from an email address outside of your organization.

4. Enable features that allow users to spot emails coming from senders they have not communicated with in the past.

5. Review mail flow rules to ensure that broad rules are not inadvertently allowing malicious emails to be delivered.
6. Create and enforce finance policies that require employees to verify any account information changes, including wire transfer information, with the account holder.

7. Ensure that all your email is effectively signed (DKIM) and verified on delivery (DMARC) so that your customers are protected from attackers trying to send messages as your domain/brand.

8. Enable advanced protections for your users,¹³ including:
 - Look-alike domains or impersonations of important users in the organization.
 - Deep analysis (detonation) of attachments and URLs across the
 - Collaboration suite to protect against 0-day attacks.
 - Post-delivery protection to remove mails that were delivered and later determined malicious.

9. Eliminate opportunities for attackers to bypass your security, such as allowing listing for senders, domains, or IP addresses.

Learn more:

[Trend-spotting email techniques: How modern phishing emails hide in plain sight | Microsoft Security Blog \(8/18/2021\)](#)

Malware

Trends we’re seeing

While phishing has grown over the last year, malware and the cybercrime infrastructure that supports attacks has also continued to evolve. There are key malware areas where Microsoft 365 Defender Threat Intelligence has observed changing trends in recent years, many of which require equal parts innovative defensive strategies and historically resilient mitigations such as multi-factor authentication and robust application security practices.

Individualized malware techniques and actions


Within the popular malware types and delivery methods analyzed over the past year, Microsoft observed many trends in individual tactics used during infection. Despite the wide range of outcomes such as ransom, data loss, credential theft, and espionage, most pieces of malware rely on similar strategies for establishing themselves in a network. Windows PowerShell launched by malicious processes with suspicious commands or encoded values was the most common behavior Microsoft observed from malware in recent months. The next most common were attempts by malware to rename payloads to mimic system processes or replace them entirely, and using malware to collect data such as credentials from browser caches.

Other noteworthy behaviors and protection opportunities for security operation centers are the use of specific reconnaissance commands, processes being added to startup folders, scheduled task or registry alterations, and malicious process execution by abuse of Office documents. These behaviors stand out due to widespread use among all malware regardless of sophistication, though Microsoft has also observed more specific tactics that are more difficult for enterprises to mitigate.

Fileless malware and evasive behavior

“Fileless” malware is malware that derives most of its components from system processes or legitimate tools already on a device, which can make it harder to remove and detect, since more than a single file needs to be removed. Persistence strategies can include registry, scheduled task, and startup folder persistence to remove the necessity for malware to remain a static item in the filesystem. Free or easily available remote access trojans (RATs), banking trojans, and offensive toolkits like Cobalt Strike are routinely utilizing process injection and in-memory execution. These are methods of abusing stolen administrative privileges to move malicious code into running benign processes rather than in static files, to circumvent easy removal. To combat these kinds of behaviors it is imperative that security teams within organizations review their incident response and malware removal processes to include sufficient forensics to ensure common malware persistence mechanisms have been fully remediated after cleanup by an antivirus solution.

Alert counts by activity (May-June 2021)



Legitimate service abuse in network communications

Another tactic used in many malware campaigns this past year utilized legitimate sites in almost every stage of malware: delivery, reconnaissance, command and control, exfiltration, malicious advertising, and cryptocurrency mining. Cloud services such as Google Drive, Microsoft OneDrive, Adobe Spark, Dropbox, and others are still very popular for use as initial delivery of malware, while content “pasting” sites such as Pastebin.com, Archive.org, and Stikked.ch are increasingly popular for covert use in multi-part and fileless malware. In the last case, the code used in the malware is pulled directly from the pasting site and executed immediately into memory, bypassing the need to download malware as a single file.


Larger trends in malware propagation and behavior

BOTNET RENOVATIONS

Botnet as a term has been evolving as well. Historically it was used to refer to a network of computers completing tasks for an operator. However, now most malware families could potentially be classified as having botnet components or behaviors.

As historically prevalent malware botnet infrastructures such as Trickbot and Emotet were disrupted, other malware families have replaced them. In their place, older botnets as well as a new

Emotet encounters



In January 2021, law enforcement performed a takedown, which led to the demise of the Emotet family of malware and a dramatic subsequent decrease in Emotet encounters.¹⁴

class of evasive malware began delivering more severe secondary components at faster speeds. In January 2021, law enforcement performed a takedown, which led to the demise of the Emotet family of malware and a dramatic subsequent decrease in Emotet encounters. Botnets such as Phorpiex¹⁵ gradually increased in number of infected base hosts and delivered numerous ransomware and secondary malware components to further monetize its behavior, including the Avaddon ransomware. Botnets such as Lemon Duck, Purple Fox, and Sysrv>Hello, surged this past year, incorporating new programming languages, new infrastructure,

and new infection methods as well. Lemon Duck, as with most emerging botnets, uses over 10 distinct methods of infection across Windows and Linux environments. Newer botnets are also quick to begin using new vulnerabilities to infect servers. Despite this, most methods still rely on unpatched edge applications, lateral movement via connected drives, and weak credentials on available services.

SEO AND MALICIOUS ADVERTISING

Search engine results and advertising are also an increasingly effective means of delivering malware to end users, both via abusing legitimate search engine


optimization strategies and by utilizing existing infections to install browser extensions to modify search results and to surface illicit material attacker content.

This was the case in 2020 with the Adrozek¹⁶ malware, a browser extension used where infected devices would use browser extensions to replace legitimate search results with links to malware impersonating Microsoft products and other legitimate software. The operators of Gootkit, a malware infection that can led to ransomware, used a slightly different technique to abuse search engines by purchasing advertising in 2020 to uplift the links to compromised sites hosting the malware. Other information-stealing malware, such as Jupyter or SolarMarker, used yet another method to appear in search results by using documents hosted on services such as AWS, Google, and Strikingly content delivery network to lead users searching for common terms via search results to PDF pages that would ultimately establish persistence on their device.

Information stealing, data exfiltration, and other areas of malware delivery can increasingly leverage browser modifications and search results to achieve their ends. This continues to solidify a class of malware leveraging the browser for delivery and exploit across both consumer and enterprise sectors.

¹⁴ <https://www.europol.europa.eu/newsroom/news/world%E2%80%99s-most-dangerous-malware-emotet-disrupted-through-global-action> ¹⁵ <https://www.microsoft.com/security/blog/2021/05/20/phorpiex-morphs-how-a-longstanding-botnet-has-evolved-to-keep-up-with-the-current-threat-environment/> ¹⁶ <https://www.microsoft.com/security/blog/2020/12/10/widespread-malware-campaign-seeks-to-silently-inject-ads-into-search-results-affects-multiple-browsers/>

Browser search results manipulation



Comparison of search results pages on an unaffected machine and one with Adrozek running.

MALWARE TOOLS

Malware has evolved to take advantage of tools that are available and, in some cases, are not inherently malicious. One prime example has been the use of Cobalt Strike, a commercial penetration testing tool. While Cobalt Strike is a penetration testing, it has been used more and more frequently in various attacks, ranging from nation state to human-operated ransomware, to

perform system and network discovery actions and move laterally through a network. Cobalt Strike is specifically designed to evade traditional detection methodologies and offers the operator a range of options for performing obfuscation of their attack commands. These obfuscation techniques themselves can however become a signal, and identifying Cobalt Strike has become more essential than ever as the cybercriminal economy leads to

malware that plants Cobalt Strike quickly, handing off to ransomware operators.

WEB SHELLS DEEP DIVE

Web shells remain popular with advanced persistent threat (APT) actors of all types, including NOBELIUM¹⁷ and HAFNIUM¹⁸ nation state activity groups. As DART and the Microsoft 365 Defender Research Team reported in both 2020¹⁹ and 2021,²⁰ web shell usage continues to climb among nation state groups and criminal organizations. Web shell is a piece of malicious code, often written in typical web development programming languages (such as ASP, PHP, or JSP), that attackers implant on web servers to provide remote access and code execution to server functions. Web shells allow adversaries to execute commands and steal data from a web server or use the server as a launch pad for further attacks against the affected organization.

The escalating prevalence of web shells may be attributed to how simple and effective they can be for attackers. Once installed on a server, web shells serve as one of the most effective means of persistence in an enterprise. We frequently see cases where web shells are used solely as a persistence mechanism. Web shells guarantee that a backdoor exists in a compromised network, because an attacker leaves a malicious implant after establishing an initial foothold on a server. If left undetected, web shells provide a way for attackers to continue

to gather data from and monetize the networks that they have access to. In addition, the volume of network traffic plus the usual noise of constant internet attacks means that targeted traffic aimed at a web server can blend right in, making detection of web shells a lot more difficult and requiring advanced behavior-based detections.

In February 2020, we reported a steady increase in the use of web shells in attacks worldwide. The latest Microsoft 365 Defender data shows that this trend not only continued, but it also accelerated; in every month from August 2020 to January 2021, we registered an average of 140,000 encounters of these threats on servers, which was almost double the 77,000 monthly average.

Throughout 2021 we saw an even bigger increase, with an average of 180,000 encounters per month. In March 2021, we saw a huge spike in web shell encounters, which we attributed to the HAFNIUM nation state activity group targeting Exchange servers with 0-day exploits.

In March and April of 2021, as exploit code became available for web-facing on-premises Exchange servers, we saw a large spike in web shell detection rates. This was due to multiple threat actors using a “compromise first, monetize later” approach that takes advantage of customer patching delays. Actors jump on opportunities as soon they arise.

¹⁷ <https://www.microsoft.com/security/blog/2021/02/11/web-shell-attacks-continue-to-rise/>

¹⁸ <https://www.microsoft.com/security/blog/2021/03/02/hafnium-targeting-exchange-servers/>


¹⁹ <https://www.microsoft.com/security/blog/2020/02/11/web-shell-investigating-web-shell-attacks/>

²⁰ <https://www.microsoft.com/security/blog/2021/02/11/web-shell-attacks-continue-to-rise/>

https://t.me/learningnets

To minimize risk, organizations should accelerate their deployment of security updates, especially for internet-facing systems. To minimize risk, organizations should accelerate their deployment of security updates, especially for internet-facing systems.

Web shell encounters, Defender signals (September 2020-June 2021)



Learn more:

- [HAFNIUM targeting Exchange Servers with 0-day exploits | Microsoft Security Blog \(3/2/2021\)](#)
- [Web shell attacks continue to rise - Microsoft Security \(2/11/2021\)](#)
- [Ghost in the shell: Investigating web shell attacks - Microsoft Security \(2/4/2021\)](#)

Summary of recommendations for malware prevention

1. Install security updates on all applications and operating systems promptly.
2. Enable real-time protection through an antimalware solution, such as Microsoft Defender.
3. Mitigate large attack vectors such as macro abuse, exposed edge services, insecure default configurations, legacy authentication, unsigned script types, and suspicious executions from certain file types delivered through email. Microsoft offers some of these larger mitigations through the use of attack surface reduction rules²¹ to prevent malware infection. Azure Active Directory users may also leverage security defaults²² to establish baseline authentication security for cloud environments.
4. Enable Endpoint Detection and Response functionality to analyze and respond to threats based on individual behaviors and techniques proactively.
5. Enable domain and IP-based protections on hosts as well as at network gateways, if possible, to ensure infrastructure-based coverage is complete.
6. Turn on potentially unwanted applications (PUAs) protection. Many antimalware solutions may label initial access threats such as adware, torrent downloaders, RATs, and Remote Management Services (RMS) as PUA. Occasionally, these types of software may be disabled by default to prevent impact to an environment.
7. Determine where highly privileged accounts are logging on and exposing credentials. Monitor and investigate logon events for logon type attributes. Highly privileged accounts should not be present on workstations.
8. Practice the principle of least privilege and maintain credential hygiene. Avoid the use of domain-wide, admin-level service accounts. Restricting local administrative privileges can help limit installation of RATs and other unwanted applications.
9. Educate users about malware threats, such as RATs, that can propagate through email as well as through web downloads and search engines.

Learn more:

- [Use attack surface reduction rules to prevent malware infection | Microsoft Docs \(6/23/2021\)](#)
- [Turn on network protection | Microsoft Docs \(6/14/2021\)](#)
- [Block potentially unwanted applications with Microsoft Defender Antivirus | Microsoft Docs \(6/02/2021\)](#)

Malicious domains

Any domain used in the pursuit of cybercrime can be considered malicious. Malicious domains can be legitimate sites which have been compromised to enable criminals to host malicious content on subdomains, or they can be entirely fraudulent infrastructure set up for the commission of a crime. Cybercriminals use malicious domains for three primary functions: information transmission, location obfuscation, and building resiliency against those seeking to interfere with their criminal activities.

Domains are used for data exfiltration, controlling ransomware communication, hosting phishing pages, and providing control to malware. They are also used as email domains to create visually identical imposter email aliases for deception. Fraudulent domains can use trademarks to deceive customers or provide a platform for fraud, such as fraudulent technical support sites.

Domain proliferation and threat mitigation

The number of domains available on the internet has mushroomed over the past several years. This includes country code top-level domains (cTLDs) such as .uk, .ca, and .cn; generic top-level domains (gTLDs) such as .com, .net, and org; and over 1,200 new gTLDs that were introduced into the Domain Name System (DNS) in 2013. Due to the sheer number of top-level domains, and the growing ecosystem of domain name registries, domain registrars, and domain registration service providers, mitigating cyber threats from malicious domains has been further complicated. Uniformity in cyber threat mitigation across the ecosystem is critical. There is movement in this direction as evidenced by recent language incorporated into the Internet Corporation for Assigned Names and Numbers (ICANN) agreements with registries, which includes terms of use that define and prohibit illegal activity, and a requirement for registries to develop their own anti-abuse policy and monitor and address abusive activity.

How domains are being used for malware

A malicious domain is often used as a destination to which malware victims are directed. In this way, the domain both initiates the establishment of a communications channel with the victim and reveals the infected victim’s location. Knowing a victim’s location is important as cybercriminals use a myriad of methods to disseminate their malware but are unable to anticipate where it will ultimately be successfully downloaded. Therefore, cybercriminals engineer their malware to “phone home” to a malicious domain. Newly infected computers immediately reach out to such domains, effectively “announcing” their location via IP addresses. There are two primary ways domains are used for these purposes in malware:

Domains aid in obfuscating and hiding the cybercriminals’ location and identity

Domains directly added to malware (or “hard-coded” domains) and incorporated into the communications infrastructure can effectively hide the cybercriminal’s true location. The cybercriminal sets up a domain as a proxy, or “stepping-stone,” which redirects the communications with a victim to another domain or IP. This process can consist of multiple “hops” spanning domains associated with top-level domains from around the world. It is common for cybercriminals to use fictitious names, emails, and addresses and pay for the domains with stolen credit cards or non-traceable digital currency.

Domains can be a mechanism to build resiliency into the infrastructure

Cybercriminals are now routinely adding domain generating algorithms (DGAs) to their malware, providing a fallback mechanism for when hard-coded malicious domains are seized by law enforcement, for example. To utilize DGAs, malicious software includes code to generate lists of domains built using random characters or strings that change based on the day, time, and year. For example, on Friday, June 4, 2022, the DGA might generate three domains, such as ahu3rrfsirraqrty.com, hyrssgu5oqr4cetc.com, and wkclsoqqpcaty.com, and the malware would attempt to contact those domains in that order. The use of DGAs increases the cost and complexity of disrupting cybercriminal communications infrastructure, requiring disruptors to monitor hundreds of thousands of potentially malicious domains, whereas the cybercriminal needs only one of them.

30,720

Potential new domains generated by DGAs in just 3-4 days

Disrupting malicious domain infrastructure

Disrupting domains on Microsoft-hosted services

Cybercriminals are now increasingly abusing Microsoft and third-party clouds as well as the services information workers and consumers use for day-to-day collaboration (such as email). Microsoft takes numerous steps to reduce cloud hosting abuse. We proactively detect abuse of the Microsoft cloud at the hosting source and neutralize it before attacks start or scale; we act on detections in our services (as in Office 365 email) and route this knowledge to internal services that can neutralize the threat; we act on customer and third-party reports; and we notify third-party industry partners of abuse on their cloud, detected by using our security services, so they can act to neutralize it at their hosting source. In the three-month period between May and July 2021, we disabled roughly 15,850 phishing sites hosted on Azure. We closely monitor abuse and evaluate new ways to detect and neutralize hosting of malicious sites

>168K

Phishing sites taken down
by Microsoft this year

Disrupting third-party-hosted domains through legal action

Given that cybercriminals are increasingly deploying private technical infrastructure, including malicious domains, to carry out a wide range of cybercrime, it is incumbent on organizations and individuals to establish the necessary legal and technical capabilities to disrupt this infrastructure through legal actions.

In recent years, the private sector has used a variety of legal theories to pursue disruption through civil actions in federal court. Criminal statutes directed at hacking and unlawful access frequently provide a civil cause of action to target malicious infrastructure. In particular, the Computer Fraud and Abuse Act (CFAA), the Wiretap Act, and the Stored Communications Act are frequently asserted legal theories. Very frequently trademark theories under the Lanham Act²³ are asserted as cybercriminals leverage trusted brands to deceive victims. In many cases, infringing domains provide the critical part of malicious technical infrastructure and often include wholesale counterfeit reproduction of legitimate content to confuse victims and advance criminal schemes.

Pursuing the appropriate remedy is an essential component of an effectively disruptive legal action. A well-crafted injunction that relies on the broad equitable authority of federal courts enables plaintiffs to obtain flexible court orders permitting them to exercise control over the cybercrime infrastructure. To obtain such relief, plaintiffs frequently invoke statutes that support seizure of physical devices, computers, and servers used for criminal purposes. Malicious domains engaged in criminal activity can also be subject to seizure under various federal statutes and equitable theories. In several civil cases, federal courts have granted this sort of injunction for violations of the CFAA, the Electronic Communications Privacy Act, the Lanham Act, and common law claims. These legal claims also support court orders that direct transfer or disablement of domains and/or IP addresses. In this scenario, courts grant the transfer or disablement of malicious infrastructure to a private plaintiff’s control, and away from the control of defendants, which effectively disrupts the technical capability of cybercriminals to launch attacks and inflict harm.

The next big threat: “Forever” (blockchain) domains

Blockchain domains are an emerging threat outside of regulation. Over the last two years, the adaptation of blockchain technology has skyrocketed across many business verticals. Real-life applications of blockchain technology range from supply chain management, identity management, real estate contracts, and domain infrastructure. In recent years, we have observed blockchain domains integrated into cybercriminal infrastructure and operations. We first observed this on a large scale when investigating the Necurs botnet that reigned terror worldwide for years with its ability to send malicious spam, often with ransomware payloads. Necurs contained a robust backup system, which incorporated a DGA. In one of its DGA versions, Necurs produced 2,048 new domains from 43 different TLDs approximately every 30 days, including the blockchain domain TLD “.bit.”

Blockchain domains are an emerging threat outside of regulation.

Unlike traditional domains that are purchased through internet registrars operating through the ICANN-regulated DNS system, blockchain domains are not governed by any centralized body, limiting the opportunity for abuse reporting and enforcement disruptions.

Traditionally, blockchain domain purchases are made through a crypto wallet with cryptocurrency from a blockchain DNS provider. Crypto wallets utilize asymmetric encryption, which involves both a private and public key for the blockchain transaction. After the transaction has been executed, the domain name, domain IP, and transaction hash are recorded into the blockchain. Moving forward, the only entity that can make changes to the IP recorded on the blockchain is the person with the wallet and private key who made the initial transaction to purchase the domain.

Blockchain domains work differently and pose challenges from both a utilization and disruption standpoint. Blockchain domains function either through software/browser plug-in or proxy resolution services. The challenge for cybercriminals with respect to blockchain domains is getting the most updated IP address from the blockchain to the computer trying to resolve the blockchain domain to an IP address. Because blockchain domains operate outside the normal DNS channels, malware authors must include additional resolution instructions for infected victims. These instructions are usually hard-coded into the malware and point the infected system to a blockchain proxy resolution service IP.

Over the years, there have been several projects on the internet to operate free unregulated DNS and support the resolution of blockchain domains. Most recently, the OpenNic project, which operates under the mission statement of “DNS neutrality and provide uncensored DNS access,” took on the task of resolving “.bit” crypto domains. Several years into the project, because of reported widespread abuse of “.bit” domains, the OpenNic project decided to stop resolving “.bit” domains.²⁴

Big threats using blockchain domains

The threat landscape of criminal infrastructure is constantly shifting to avoid detection and disruption. Within the past year, some of the bigger threat actors on the internet have started utilizing blockchain domains as part of their infrastructure. Trickbot, the notorious banking trojan which has evolved its business model into providing access to high-value targets in the ransomware space, started using “.bazar” domains provided by Emercoin blockchain DNS. The more recent threat, Bazarloader, which has connections to Trickbot, started deploying a unique version of its DGA that uses “.bazar” domains. This trend of threats leveraging blockchain domains as infrastructure with the means to create an undisputable criminal network should be taken seriously.


Investigating blockchain domains provides a unique challenge because there is no central WHOIS registration database tracking who registered the domain and when. Fortunately, some blockchain DNS providers like Emercoin have provided access to a block explorer tool,²⁵ which enables a search for domain names, transaction hashes, and other values that might be stored in the blockchain. The Emercoin blockchain is pseudo-anonymous but can reveal some interesting information about a blockchain domain, such as IP addresses and transaction dates.

Countering blockchain domains might not be as difficult as you think

The weakness in blockchain domains is the need for third-party proxy services or browser plug-ins to resolve blockchain domains to an IP. Disabling or blocking the blockchain proxy resolution services and disabling browser plug-ins will disable the ability for blockchain domain resolution. Many threat intelligence vendors provide malicious URL feeds, which sometimes include blockchain resolution proxies or the blockchain domain itself.

Blockchain domains have become a preferred choice for cybercrime infrastructure

Regular domains



Domain is registered to the owner and resolves to a specific server and IP address.

Blockchain domains



Domain is registered to the owner and the records are stored on a blockchain.



The crypto currency address is replaced with a human readable name. Many cryptocurrency addresses can be added to one domain.

Regulated



The content is stored on centralized servers owned by large companies.



The domain is regulated for safe and appropriate use by ICANN (Internet Corporation for Assigned Names and Numbers).

Not regulated



The domain is stored on a blockchain, just like cryptocurrency. No third party can move or seize the domain without the private key. The holder of the private key is the only person who can make changes to the blockchain record.



The domain is NOT regulated by ICANN.

Censorable



If the domain is used for malicious intent or other misconduct, it is censored, and can be taken down by the hosting company, by governments, or by law enforcement.

Not censorable



Since blockchain domains can be used to build uncensorable websites and simplify cryptocurrency payments while remaining fairly anonymous, they are becoming increasingly used for cybercrime.

Machine learning (ML) is an artificial intelligence (AI) technique that can be used in numerous applications, including cybersecurity. In responsible ML innovation, data scientists and developers build, train, and deploy ML models to understand, protect, and control data and processes to build trusted solutions.

However, adversaries can attack these ML-driven systems. The methods underpinning the production ML systems are systematically vulnerable to a new class of vulnerabilities across the ML supply chain collectively known as “adversarial ML.” Adversaries can exploit these vulnerabilities to manipulate AI systems and alter their behavior to serve a malicious end goal.

Microsoft worked with MITRE to create the Adversarial ML Threat Matrix because we believe the first step in empowering security teams to defend against attacks on ML systems is to have a framework that systematically organizes the techniques employed by malicious adversaries in subverting ML systems. We hope that the security community can use the tabulated tactics and techniques to bolster their monitoring strategies around their organizations' mission-critical ML systems.


We think that securing ML systems is an infosec problem. The goal of the Adversarial ML Threat Matrix is to position attacks on ML systems in a framework where security analysts can orient themselves in these new and upcoming threats. The matrix is structured like the ATT&CK framework, owing to its wide adoption among the security analyst community. This way, security analysts have a familiar framework to learn about threats to ML systems, which are inherently different from traditional attacks on corporate networks.

We seeded this framework with a curated set of vulnerabilities and adversary behaviors that Microsoft and MITRE vetted to be effective against production ML systems, enabling security analysts to focus on realistic threats. We also incorporated learnings from Microsoft's vast experience in this space into the framework. For instance, we found that model stealing is not the end goal of the attacker but in fact leads to more insidious model evasion. We also found that when attacking an ML system, attackers use a combination of traditional techniques like phishing and lateral movement alongside adversarial ML techniques.

Microsoft has incorporated AI- and ML-specific security practices into its Security Development Lifecycle (SDL) to protect Microsoft products and services against these attacks. In addition to threat detection and mitigation development work and automation, we have published guidance on steps customers can take to build defense in depth into their own AI and ML systems.

The centerpiece of the materials we've published is called Failure Modes in Machine Learning,²⁶ which lays out the terminology we developed jointly with the Berkman Klein Center for Internet and Society at Harvard University. It includes vocabulary that can be used to describe intentional failures caused by an adversary attempting to alter results or steal an algorithm, as well as vocabulary for unintentional failures such as a system that produces results that might be unsafe.

Attacks on ML models



Attack	Description	Example
Evasion attack	Attacker modifies the query in a way that causes a model misclassification.	Self-driving cars By manipulating a stop sign, or the environment observed by the car's image recognition system, the adversary causes a model misclassification. In this way, a self-driving car could be made to ignore the stop sign.
Poisoning attack	Attacker contaminates the training phase of ML systems to get intended result. The attacker wants to misclassify specific examples to cause specific actions to be taken or omitted.	Critical infrastructure systems By submitting antivirus software as malware, an adversary can force its misclassification as malicious thereby eliminating the use of the antivirus software on client systems. This could leave critical infrastructure systems open to attack.
Membership Inference	Attacker can infer if a given data record was part of the model's training dataset or not.	Healthcare information An adversary could look at a model trained on a body of data consisting of people with a specific surgical procedure. By knowing that a particular individual's data was in the training set, an adversary would then know the individual had the surgical procedure. This privacy violation could then be leveraged publicly.
Model Stealing	Attacker is able to recover the model through carefully crafted queries.	Finance algorithms Through model queries, an adversary could reconstruct the potential outputs of the ML model. This could be used to adversely affect proprietary algorithms designed for high frequency stock trading in a particular market.

Attacker evasions

An evasion attack is an exploratory attack against an ML model to cause an integrity violation. From a system’s security perspective, it is instructive to consider black-box evasion attacks, in which an attacker may have no specific knowledge of the inner workings of the ML model but instead effects change by submitting inputs and observing the corresponding system output. This threat model is common to many AI systems hosted as a cloud service or on a consumer device, for example, and is a concern for models in finance, healthcare, defense, fraud, and security.

Sophisticated black-box evasion attacks against ML models have been demonstrated repeatedly by white-hat researchers by using algorithms that iteratively determine what input will cause an integrity violation. Today, however, threat actors in the wild may also attempt evasion of ML systems in some domains, but usually do so through manual rather than algorithmic means and do not necessarily focus exclusively on ML as the evasion target. For example, content moderation filters are bypassed by mischievous or economically motivated users by obscuring payload content in creative ways. Security products that include antimalware or antiphishing models are evaded by adversaries using several obfuscation techniques. That these target systems rely on ML is not necessarily a consideration in these practical manual evasion attacks.

Whether or not an adversary is present in your business domain, the risk of adversaries evading an ML model exists in every domain. An ML model is an imperfect summary of a dataset, and as such, models have intrinsic failure modes even when trained on an ideal dataset. It is generally understood that the feasibility of evasion is a property of all ML models, rather than a failure mode to which only a few are susceptible to. These integrity violations may rarely be encountered during nominal use of the ML model but can be readily discovered by an adversary who is explicitly optimizing for the worst-case conditions necessary to cause them.

ML model/data poisoning

We are seeing a trend shift in adversarial ML security research. While in past years there was a focus on highly visible model evasion attacks where the fragility of some ML models could be easily demonstrated, the focus of security researchers is broadening to include less noticeable attacks. For example, in data poisoning attacks, the target is the training data that the ML models are built from. As new data is aggregated and incorporated into a dataset for training, it becomes increasingly important to validate that new training data has not been compromised. We have evidence of customer ML model compromise resulting from adversarial contamination of training data which, when left undetected, becomes an equally trusted part of existing training datasets. Without automation measuring for statistical drift in growing datasets, these types of attacks largely go undetected until an ML model has a critical failure.

As we’ve seen with security research into past vulnerabilities,²⁷ a pronounced uptick in research publications is soon followed by active exploitation. In anticipation of such a shift to focus on data poisoning attacks, Microsoft continues to focus on designing threat detections and mitigations to protect ML models and their datasets against these threats. Mitigations in this space can also be beneficial to detecting non-malicious training data drift, giving data scientists greater insight into the quality of their data over time and highlighting anomalies for investigation.

The risk of adversaries evading an ML model exists in every domain.

What we’re doing to stay ahead of the curve

Performing security assessments of production AI systems is not easy. Microsoft surveyed 28 organizations,²⁸ spanning Fortune 500 companies, governments, non-profits, and small and medium-sized businesses, to understand the current processes in place to secure AI systems. We found that 25 out of 28 businesses indicated they don’t have the right tools in place to secure their AI systems and that security professionals are looking for specific guidance in this space.

To address the growing needs of adversarial ML, Microsoft released Counterfit, an open-source tool to help assess risk by allowing users to attack their own AI/ML. This tool was developed out of our own need to assess Microsoft’s AI systems for vulnerabilities and proactively secure AI services, in accordance with our responsible AI principles²⁹ and Responsible AI Strategy in Engineering (RAISE) initiative. Counterfit started as a corpus of attack scripts written specifically to target individual AI models and then grew into a generic automation tool to attack multiple AI systems at scale. Today, we routinely use Counterfit as part of our AI red team operations.


Based on learnings from internal and external engagements, Counterfit is designed to be flexible in three key ways:

- 1. Environment agnostic:** It can help assess AI models hosted in any cloud environment, on-premises, or on the edge.
- 2. Model agnostic:** The tool abstracts the internal workings of AI models so that security professionals can focus on security assessment.
- 3. Strives to be data agnostic:** It works on AI models using text, images, or tabular input, and we continue to add data types.

Learn more:

- [Our approach to responsible AI at Microsoft](#)
- [GitHub – Azure/counterfit: a CLI that provides a generic automation layer for assessing the security of ML models](#)
- [AI security risk assessment using Counterfit | Microsoft Security Blog \(5/3/2021\)](#)
- [Adversarial Machine Learning – Industry Perspectives \(3/19/2021\)](#)

AI risk management



Security is one part of a larger emphasis in a burgeoning market called “AI risk management” and includes “model operations”—ensuring that your AI system is reliable, accurate and available. It also includes “responsible AI” with fairness, ethics, transparency and all the legal ramifications of having AI systems behave responsibly. This “security” element also deserves attention and is an important piece in rounding out a risk management posture.

Standards for addressing security of AI systems

The prevalent use of AI and ML across industry sectors, an emerging regulatory landscape, and widespread mistrust or misunderstanding in the use of these technologies has led to an increased need for standards to define good practice and provide guidance to improve trust and market adoption. The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) are developing AI standards, including defining key terminology and concepts for AI and ML, risk management, governance implications, data quality, and various topics related to trustworthiness. Also under development is a certifiable management system standard for AI, which will guide organizations to adopt a risk-based approach to responsibly use and develop AI systems as well as to demonstrate accountability and their duty of care toward stakeholders.

AI and ML are increasingly integrated into all types of systems, including critical and safety infrastructure, resulting in new security threats unique to the use of AI systems and highly undesirable consequences of attacks. Such consequences can include the detrimental performance of a chatbot, denial of essential services, theft of intellectual property, or even physical danger to humans. In some instances, security attacks on AI systems have already caused significant issues.

AI security cannot be considered in isolation of existing risk-based security, privacy, and governance foundations, which can address many of the threats that arise using AI systems. For example, using standards such as those for an information security management system (ISO/IEC 27001) and a privacy information management system (ISO/IEC 27701) can help an organization to implement processes and controls to address security and privacy risks associated with its objectives and activities, including security threats to AI systems. In addition to these existing practices, this evolved threat landscape will require new guidance, good practices, and organizational and technical measures to help organizations protect their AI systems. Effective security measures are a vital component of responsible development and deployment of AI systems. Microsoft is engaged in new standards work that has been initiated to provide guidance for addressing security threats and failures in AI and ML.

Learn more:

- [Responsible AI principles from Microsoft](#)
- [Threat Modeling AI/ML Systems and Dependencies – Security documentation | Microsoft Docs, \(11/11/2019\)](#)
- [AI/ML Pivots to the Security Development Lifecycle Bug Bar – Security documentation | Microsoft Docs \(11/11/2019\)](#)
- [Failure Modes in Machine Learning – Security documentation | Microsoft Docs \(11/11/2019\)](#)

CHAPTER 3

Nation state threats

Introduction

Tracking nation state threats

What we're seeing

Analysis of nation state activity this year

Private sector offensive actors

Comprehensive protections required

INTRODUCTION: Attackers increase use of deception to pursue national objectives

JOHN LAMBERT, DISTINGUISHED ENGINEER AND VICE PRESIDENT, MICROSOFT THREAT INTELLIGENCE CENTER

The last year has been marked by significant historic geopolitical events and unforeseen challenges that have changed the way organizations approach daily operations. During this time, nation state actors have largely maintained their operations at a consistent pace while creating new tactics and techniques to evade detection and increase the scale of their attacks.

Major cybersecurity events, like the SolarWinds attacks by NOBELIUM and on-premises Exchange Server attacks by HAFNIUM, and attacks by multiple other actors have focused our collective attention on securing our supply chains. Nation state actors and many cybercrime operations have focused efforts on exposing security vulnerabilities among their suppliers or discovering unpatched systems that organizations relied on for continuity of business during this unusual year. These recent events have demonstrated the increasing importance in maintaining current security updates in all deployed systems as the most effective way to protect against rapidly evolving threats.

The Microsoft Threat Intelligence Center (MSTIC) and the Digital Security Unit (DSU) have observed that most nation state actors continue to focus operations and attacks on government agencies, intergovernmental organizations (IGOs),

nongovernmental organizations (NGOs), and think tanks for traditional espionage or surveillance objectives. The victims of attacks often have information relevant to an adversary government's intelligence needs, which is why so many government agencies and think tanks are attacked. However, private industry's role in supporting remote workers, increased healthcare services, COVID-19 vaccine research, and COVID-19 vaccine distribution have also made them more common targets for these sophisticated actors seeking information for their government's national security or intelligence purposes. Our increased reliance on the global telecommunications backbone and virtual private network (VPN)/virtual private server (VPS) infrastructure for remote workers gave malicious actors new vectors to gain access to targeted private networks that were scrambling to support new ways of doing business.

These sophisticated attackers continue to focus on effective techniques that help them maintain stealth and access. We have seen continuing attacks on traditional security hygiene elements as well as focus on developing and refining new, breakthrough attacks targeting the supplier ecosystem in order to attack downstream customers. Well-established spear phishing and password spray campaigns by nation state actors continue to be successful against organizations that have not yet implemented multifactor authentication (MFA) or other protections against this common tactic. However, as more organizations invest in securing their accounts, the success rate of these techniques will decline, and detection of the attacks will increase. In response, nation state actors appear to be increasing the scale and volume of these attacks to evade detection and improve the likelihood of success across multiple targets. This volume-oriented approach to

NATION
STATE ACTORS
APPEAR TO BE
INCREASING
THE SCALE
AND VOLUME
OF ATTACKS
TO EVADE
DETECTION.

compromising credentials will continue to be a useful technique if poorly secured accounts are available as targets. Attacks against unpatched third-party software or on-premises infrastructure will likely become more pervasive and become more easily exploited by nation state and cybercrime actors. Postponing installation of security updates or incomplete knowledge of deployed systems and their patch state will leave organizations vulnerable to sudden large-scale attacks as they scramble to identify affected assets and catch up to a fully patched state. Running networks with unsupported software, or software that is no longer updated, increases risk exponentially. Organizations must maintain comprehensive asset inventory, patch state awareness, and thorough backup and containment plans to be resilient against sophisticated attacks. Adversaries will continue to evolve new techniques to target and compromise corporate resources requiring a comprehensive “assume breach” mentality that extends beyond basic hygiene needs and MFA and into a holistic set of Zero Trust security principles. Applying a Zero Trust security model³² will become increasingly critical in protecting corporate identities, devices, applications, data, networks, and infrastructure against sophisticated threats.

Looking forward, we know that adversarial governments will continue with their intelligence collection objectives, as well as explore the political boundaries of acceptable behavior in cyberspace. As a result, we expect nation state actors to continue

refining their techniques by leveraging new exploits against security weaknesses and unpatched systems of common supply chain vendors in order to gain access to and collect information from downstream customers. Spear phishing and password spray attacks show no sign of slowing as the common method for reconnaissance and infiltration, increasing the importance of implementing end-to-end MFA across accounts. The information Microsoft provides in this chapter captures much of the activity we have seen targeting our customers globally in the past year and captures the trends we anticipate nation state actors will continue to use in the next year. We recommend you use this information as a guide to understanding the tactics and techniques that these sophisticated actors may use to target your organizations so you can more effectively implement proactive defense.

Tracking nation state threats

Microsoft tracks nation state activities to protect our customers and our platforms and services. We use a variety of metrics and sophisticated data integration techniques to better understand targeting, motivations, and customer impact. MSTIC focuses on nation state actor activities because these tactics, techniques, and procedures (TTPs) often have significant impact on our customers and

are often unique, prompting deeper analysis and creation of customized detections. Understanding these TTPs also helps Microsoft better understand downstream actors such as cybercriminals and smaller nation states who often copy or reuse these methods. DSU focuses on the victims identified by MSTIC, connecting the victims of the attack to political objectives and stated intelligence goals of governments to help Microsoft provide fuller context to the world about why these nation state attacks occur.

We focus on nation state activities regardless of platform, targeted victim, or geographical region, and we maintain visibility and active threat hunting worldwide to write better detections for our customers. We also analyze why nation state actors are pursuing particular victims, sectors, or regions. Putting it all together, the information presented here provides a snapshot of our defensive efforts on behalf of our affected customers. It is important to note that even if a particular industry sector or geographic region is not represented in the following information, nation state activity spans nearly every industry sector and geographic region. In other words, protections against these tactics are critical for every organization and individual. Our intelligence is impacted by the degree to which our products and platforms are utilized in a particular geography or sector.

Nation state notifications

When a customer, whether it's an organization or individual account holder, is targeted or compromised by nation state activities that Microsoft tracks, we deliver a nation state notification (NSN) to the customer. Over the past three years, Microsoft has delivered over 20,500 NSNs. The charts and graphs in this chapter are derived from Microsoft's NSN process.

Countering nation state activity

Nation state actors are generally well-resourced and capable adversaries. As noted above, they are often pursuing intelligence collection against targets of interest to their governments. Our relentless pursuit of these adversaries and our continuous development of new capabilities to detect and deter malicious activity supports our commitment to customer protection. We are constantly improving our capability to understand nation state actors and their victims to help bring better context and understanding to our customers.

Our approach

Microsoft uses a five-pronged approach to disrupt nation state actors—providing direct customer notifications, leveraging technology to detect and defend, taking technical action against malicious operations, pursuing legal action, and participating in public policy discourse—and each one plays an important role in our commitment to protecting our customers and the ecosystem at large.

1. Empowering customers

Microsoft leverages its NSN process to inform customers of targeting or compromise from nation state actors we track, providing actionable information for customers to rapidly respond and protect themselves. Microsoft also provides alerts to industry sectors and customer segments to help raise awareness of malicious activity and guidance on how to respond.

2. Leveraging technology

Microsoft's cumulative knowledge of the global threat landscape enables our products and services to constantly create and update new security product detections, helping to protect and defend against nation state activities at scale. These collective defenses represent the most effective method to counter nation state threats, as they are informed by the extensive threat intelligence resources built into each product and enabled by world-class engineering.

3. Taking technical action against malicious operations

From time to time, Microsoft will have sufficient information to warrant a one-time deletion or shutdown of infrastructure or assets associated with a nation state attacker. By taking proactive action against malicious infrastructure, the actor loses visibility, capability, and access across a range of assets previously under their control, forcing them to rebuild.

4. Digital Crimes Unit

One of Microsoft's unique resources in the fight against nation state actors is the Digital Crimes Unit (DCU). Using litigation to seize domains and assets used by nation state actors against Microsoft customers, the DCU has been instrumental in shutting down those attack vectors. These cases have led to the takedown of hundreds of domains and the protection of thousands of customers, and Microsoft remains one of the only companies willing to pursue legal action against nation state actors in order to seize infrastructure and disrupt attacks. Lessons learned from the cases are shared with Microsoft engineering teams to help improve our operational and technical disruption capabilities.

5. Informing public discourse and policy

Microsoft uses its voice to raise awareness about nation state activities, highlighting the context and impacts of the incidents and sharing context about attacks and why they matter to the world. This helps drive a broad discussion about what can be done to combat malicious nation state activities across government entities, NGOs, enterprises, academia, and the public. Talking publicly about nation state attacks is an important part of deterrence.


Guide to the nation state actors discussed in this report

Throughout this chapter, we cite examples of nation state actors to provide a deeper view into attack targets, techniques, and analysis of motivations. Microsoft identifies nation state activities by chemical element names, just some of which are shown in the following table together with the countries of origin from which the actors operate. This small sample of the total nation state actors tracked by Microsoft represents those that were most active in the last year and made most effective use of the tactics detailed in this chapter.

Microsoft also tracks and investigates many malicious activities that are either new or unknown in origin to develop a full understanding of the tactics, techniques, and objectives.

When we take proactive action against malicious infrastructure, the actor loses visibility, capability, and access across a range of assets previously under their control, forcing them to rebuild.

Sample of nation state actors and their activities



What we're seeing

Nation state targets

In the 2020 Microsoft Digital Defense Report, we identified common aims (espionage, disruption/destruction) and common techniques (reconnaissance, credential harvesting, malware, and virtual private network (VPN) exploits) prevalent among major nation state cyber actors. These aims and techniques were as prevalent this year as the year before. Tried-and-true methods such as large-scale spear-phishing campaigns were still valuable tools in the kits of hackers. However, attackers worldwide, either affiliated directly with governments or with more loose connections, are continuing to perform research against targets in order to be more convincing in an attack, develop new techniques that have not been seen before, or even mimic criminal behavior in an attempt to obfuscate intent and objective. Microsoft responds by also working to improve our ability to keep up with the changes.

Espionage more prevalent than destructive attacks

The two main goals of nation state actors have not changed either. In the last year, espionage, and more specifically, intelligence collection, has been a far more common goal than destructive attacks.

Iran has been the only nation state actor willing to regularly engage in destructive attacks, mostly against Israel. These cyberattacks happened within a political environment in which both countries were trading blows just short of military strikes, including attacks on one another's cargo ships. With tensions already so high, the decision to use cyber for destructive attacks was less of a strategic leap for Iran than it would have been for North Korea, Russia, or China. While nations other than Iran mostly refrained from destructive attacks, they did continue to compromise victims that would be prime candidates for destructive attacks if tensions increased to the point where governments made strategic decisions to escalate cyber warfare.

The "Most targeted sectors" chart in this chapter section shows that nearly 80% of those targeted were either in government, NGOs, or think tanks. Think tanks often serve as policy incubators and implementers, with strong ties to current and former government officials and programs. Threat actors can and do exploit the connections between the more traditional NGO community and government organizations to position themselves to gain insights into national policy plans and intentions. As noted, it's the think tanks with ideas relevant to current or future government policy or political objectives that put these organizations into the line of sight for intelligence operations. When traditional NGOs have similar information, we also see them as an objective for nation state actors.

Nation state actors from North Korea added a third motive to their cyberattacks: monetary gain. North Korea targets companies in cryptocurrency trade or related research, likely seeking either to steal cryptocurrency or intellectual property. North Korea's economy is never strong, but the COVID-19 pandemic coming after years of UN sanctions has pushed it to its worst state in a generation, forcing North Korea to seek to find money by any means necessary. Although Iranian nation state actors frequently used ransomware attacks, Microsoft assesses that the ransomware was used more for covering the tracks of the attackers than for profit.


Targeting of IT companies is the big story of last year

A more revolutionary change, one common among all of the Big Four nation state cyber programs, has been the decision to target IT service providers in order to more successfully exploit victims downstream who receive services from those IT providers. The most glaring examples of the use of this kind of strategy from the last year are the Russian SolarWinds attacks and the Chinese exploitation of a vulnerability in on-premises Microsoft Exchange servers. These attacks are both covered in detail in the sections on Russia and China.

Although SolarWinds and the Exchange vulnerability were the two main cybersecurity stories of the year, Iran and North Korea also used similar tactics of targeting IT providers to find creative ways to exploit their real targets. For example, North Korean actor ZINC created online personas of apparent cybersecurity experts, including websites and social media pages, and used these personas to approach experts in cybersecurity vulnerabilities in an attempt to get the researchers they corresponded with to open content that would have downloaded exploits onto their machines. While this was not a direct attack on an IT company like SolarWinds, it did represent an attempt to indirectly find avenues to compromise North Korea's actual targets by going through the experts responsible for finding ways to protect them.


For more information on supply chain security, see the [Supply chain, IoT, and OT security](#) chapter of this report.

Most targeted countries (July 2020-June 2021)




Organizations in the United States remained the target of most of the observed activity this year. We also noted targeting increases consistent with increasing geopolitical tensions between nations. Russia-based NOBELIUM raised the number of Ukrainian customers impacted from six last fiscal year to more than 1,200 this year by heavily targeting Ukrainian government interests involved in rallying support against a build-up of Russian troops along Ukraine’s border. This year marked a near quadrupling in targeting of Israeli entities, a result exclusively of Iranian actors, who focused on Israel as tensions sharply escalated between the adversaries.

Most targeted sectors (July 2020-June 2021)



Every threat actor we tracked this year targeted entities within the government sector. NOBELIUM, NICKEL, THALLIUM, and PHOSPHORUS were the most active against this sector from the Big Four threat countries. Government sector targeting largely focused on ministries of foreign affairs and other global government entities involved in international affairs. (This chart excludes consumer accounts and depicts only enterprise targets’ corresponding sectors.)

Consumer versus enterprise targets (July 2020-June 2021)



It is likely that threat actors thought consumer email accounts could be an easier way to gain access to targeted networks during the global move to remote work. Separate from enterprise targets and the industries they represented, consumer accounts received the second highest number of notifications this year. THALLIUM and PHOSPHORUS invested heavily in spear-phishing campaigns targeting these accounts.

Critical infrastructure versus noncritical infrastructure targets


From July 2020 to June 2021, critical infrastructures were not the focal point according to the NSN information that was tracked. China-based

threat actors displayed the most interest and Russia-based threat actors accounted for the least in targeting entities in the critical infrastructure sector. Russian NOBELIUM's cyber operations are a perfect example


of displaying Russia's interest in conducting operation for access and intelligence collection versus targeting a critical infrastructure for potential disruption operations.

Targeting critical versus noncritical infrastructures (PPD-21)(July 2020–June 2021)


Russia's targeting of critical infrastructures




China's targeting of critical infrastructures




Iran's targeting of critical infrastructures



North Korea's targeting of critical infrastructures




Combination of China, Iran, North Korea, and Russia's targeting of critical infrastructures



Compromised versus targeted success rate

Success rates varied widely from threat group to threat group. Some groups, like North Korean THALLIUM, had a low success rate, because they used strategies like large-scale spear-phishing campaigns that rely more on using a wide net than a surgical strike. Password sprays are another example of a tactic with a low success rate, but where the attackers understand the success rate will be low. Other groups use very focused attacks that succeed much more often. HAFNIUM, for example, succeeded in 43% of its attacks. NICKEL succeeded at an astonishing rate of over 90%. The figures below represent an average of different tactics that are designed to succeed at different rates. The first quarter was extremely high, not necessarily because actors were more successful that quarter, but because Microsoft noted less activity stemming from low-success-rate attacks.

Compromise rate (July 2020-June 2021)




Activity origins

The following section depicts the frequency of attacks by country of origin, measured in the number of NSNs generated from attacks by each actor. Russia-based threat activity dominated this year, driven by NOBELIUM's large-scale targeting. North Korea-based actors also employed a strategy of ubiquitous targeting that earned North Korea the second highest percentage of notifications.

Country of activity origin

NOBELIUM, and its aggressive targeting of IT service providers and Western government institutions, catapulted Russia to the top spot for countries where attacks originated this year. That group was responsible for 92% of the notifications to customers about Russia-based threat activity. The outsized proportion of attacks coming from North Korea is a result of the strategy taken by threat actors THALLIUM and CERIUM. These groups rely on large quantities of attacks. While these attacks have a low percentage of success, because of the high number of attempts, the groups still manage to successfully infect some victims.


Attacks by country of origin (July 2020-June 2021)



Most active nation state activity groups

Like the “compromised versus targeted” section above, the data in this section is heavily affected by the tactics chosen by the attackers. If one group attempts a password spray attack on a hundred targets and successfully compromises one, while another group surgically focuses on only one victim that it compromises, they have both had the same number of successes. However, the first group will appear as the more “active” group, because it has attacked a hundred targets. The top three groups in this list all make use of high-fail-rate tactics. NOBELIUM, in addition to high-success-rate-focused attacks, also makes frequent use of low-success password sprays, while THALLIUM and PHOSPHORUS send spear-phishing emails to large groups. This chart, then, does not necessarily equate to the most dangerous groups, although it does say something about their relative levels of persistence and ubiquity.

Most active nation state activity groups (July 2020-June 2021)



Nation state attacker tools

The tools used by nation states to compromise victim networks are most frequently the same tools used by other malicious actors. To achieve their objectives, nation state actors may create or leverage bespoke malware, construct novel password spray infrastructure, or craft unique phishing or social engineering campaigns. However, actors like GADOLINIUM are also increasingly turning to use open-source tools³³ or common malware to impact a supply chain, attempt a man-in-the-middle attack, or launch a denial-of-service attack. These methods allow malicious actors to obfuscate their actions by hiding in plain sight.


Increased use of open-source tools provides some advantages for security professionals responsible for detecting and defending against these attacks. Increasingly, the same security and computer hygiene routines that protect from ordinary threats also help protect from nation states. Training employees to be skeptical can go a long way to stopping spear-phishing attacks and thwarting the common methods that Microsoft sees in the early stages of a compromise.

It is often the case that nation state actors develop and refine new attack techniques and that criminals adopt and further refine them over time. Microsoft expects tools designed to target and compromise IT supply chains to enter the mainstream and become more common, making concepts like Zero Trust architecture a priority from software development through deployment and updating. Well-funded nation state actors will continue to create unique tools to achieve their objectives, but just like any other streamlined organization, are just as likely to use common tools where they can to improve efficiency and effectiveness.

Learn more:

Protecting customers from a private-sector offensive actor using 0-day exploits and DevilsTongue malware | Microsoft Security Blog (7/15/2021)

Attack vectors used by nation state malicious actors



Nation states are advanced enough to do reconnaissance on their victims and select the attack method that best suits each goal or intended outcome.

Analysis of nation state activity this year

Evolving nation state cybersecurity threats have produced a watershed year with an increased focus on on-premises servers and the exposure of widespread supply chain vulnerabilities, most acutely

in software. The scope, sophistication, and success of HAFNIUM's 0-day exploits for on-premises Exchange servers in early 2021 and NOBELIUM's compromise of SolarWinds' network management software in late 2020 caught the world's attention, although the nation state threat extended beyond those immediate incidents. For one, Russia-based NOBELIUM and China-based HAFNIUM's targeting of on-premises resources and dumping credentials in those operations would have allowed them the opportunity to seize credentials to access and pivot to cloud-based resources. Russia-linked STRONTIUM

also developed target-specific capabilities for on-premises infrastructure of foreign and defense-related entities in Europe, marking a shift from the predominant cloud-first, cloud-only operations the group was known for in 2019 and 2020. Multiple Iranian actors also likely conducted supply chain operations, including one in early 2020 that likely sought intelligence from government agencies indirectly through IT and engineering services companies that support US defense and intelligence agencies.

Russia

Over the past year, Russia-based activity groups have solidified their position as acute threats to the global digital ecosystem by demonstrating adaptability, persistence, a willingness to exploit trusted technical relationships, and a facility with anonymization and open-source tools that make them increasingly difficult to detect and attribute. They have also shown a high tolerance for collateral damage, which leaves anyone with connections to targets of interest vulnerable to opportunistic targeting.

Activity Group Name	Other names	Country of origin	Industries targeted
STRONTIUM	APT28, Fancy Bear	Russia	Government, diplomatic and defense entities, think tanks, NGOs, higher education, defense contractors, IT software and services
NOBELIUM	UNC2452	Russia	Government, diplomatic and defense entities, IT software and services, telecommunication, think tanks, NGOs, defense contractors
BROMINE	Energetic Bear	Russia	Government, energy, civil aviation, defense industrial base

Russia



Abusing supply chain and other trusted technical relationships

Russia-based NOBELIUM proved how insidious and devastating software supply chain attacks can be with its compromise of the SolarWinds Orion software update. Although the group limited follow-on exploitation to roughly 100 organizations, its malicious backdoor malware was pushed to roughly 18,000 entities worldwide, leaving those impacted customers vulnerable to further attack.

NOBELIUM’s operational techniques were much more diverse than just the malicious backdoor and ranged from password spray and phishing to compromise of third-party providers to facilitate future attacks. The actor targeted cloud solution providers (CSPs) and leveraged the backdoor to steal a Mimecast private key. NOBELIUM went on to target downstream customers by masquerading as those CSPs and as the legitimate Mimecast app.³⁴ In May, the group compromised a US government agency’s account at a popular email marketing service, cloaking malicious components behind the service’s legitimate URL to send a phishing email to more than 150 diplomatic, international development, and nonprofit organizations mostly in the United States and across Europe.³⁵

Comparing the distribution of NOBELIUM victims identified in the first few months after discovery

of the SolarWinds compromise and the targeting picture of NOBELIUM’s activity through June 2021 highlights the tactical shifts and multi-vectored approach the threat group employs to gain access to desired systems. The first chart depicts victims that were subject to high-touch threat actor exploitation that in some cases leveraged the supply chain backdoor access. The second chart reflects the mass spear-phishing and password spray campaigns the actor used against targeted organizations in the first half of 2021. We can see NOBELIUM consistently targeted the government, NGO, IT services, and professional services sectors (included in “Other” in the latter chart), but the volume of compromise attempts fluctuated in line with the tactical changes.

Using a range of techniques to evade detection and attribution


Russian actors demonstrated varying degrees of adaptability and security consciousness that helped them evade attribution and network defenses. NOBELIUM showed a deep knowledge of common software tools, network security systems, and cloud technologies, as well as remediation methods incident response teams use, and they changed their operations accordingly to maintain persistence.³⁶ Responder surveillance was a tactic employed by another Russian threat group, YTTIRIUM, in the past.³⁷

In summer and fall 2020, STRONTIUM deployed an automated password spray/brute forcing tool that ran through more than a thousand anonymized Tor IPs,³⁸ making it large scale and hard to detect and attribute. The tool was deployed multiple times against more than 40 political organizations and advocacy groups based in the United States and the UK in the run-up to, and immediately after, the US presidential election.


Achieving higher rates of compromise and targeting more government organizations

Over the past year, Russia-based groups have improved their rates of successful compromise and increasingly set their sights on government targets, a confluence of trends that could portend more high-impact compromises in the year ahead. Year-on-year comparisons of NSN data depict a marked increase in successful compromises, from 21% successful between July 2019 and June 2020 to 32% since July 2020. The percentage of government organizations among Russian targets exploded from roughly 3% last period to 53% since July 2020. Russian threat actors will follow targets wherever they are, be it in the cloud or on-premises. This past year, STRONTIUM pivoted to more on-premises targeting, developing target-specific capabilities against on-premises infrastructure of foreign policy and defense-related entities in Europe. This strategy was a change from the predominant cloud-first,

NOBELIUM targets by industries/verticals (December 2020-Jan 2021)



NOBELIUM targets by industries/verticals (January-June 2021)



NOBELIUM: Variable target picture reflects diversity of tactics.

³⁴ <https://www.usnews.com/news/technology/articles/2020-12-24/solarwinds-releases-update-to-flagship-software-after-hack> ; <https://www.msn.com/en-us/news/technology/mimecast-reveals-source-code-theft-in-solarwinds-hack/ar-BB1eInqC>
³⁵ <https://blogs.microsoft.com/on-the-issues/2021/05/27/nobelium-cyberattack-nativezone-solarwinds/> ³⁶ <https://www.microsoft.com/security/blog/2021/03/04/goldmax-goldfinder-sibot-analyzing-nobelium-malware/> ³⁷ <https://www.youtube.com/watch?v=1K5tj1F0t88> ³⁸ <https://www.microsoft.com/security/blog/2020/09/10/strontium-detecting-new-patterns-credential-harvesting/>

cloud-only operations the group was known for throughout 2019 and into 2020.

For the first time since August 2018, government organizations were the most targeted sector for Russian threat actors Microsoft tracks, followed by think tanks. The government organizations were largely involved in foreign policy and national security or defense, although threat actor BROMINE concentrated its efforts against US state, county, and city governments, as well as aviation and

port authorities. The healthcare industry was the third most targeted sector this period, fueled by STRONTIUM’s credential harvesting attempts against organizations developing and testing COVID-19 vaccines and treatments⁴⁰ in the United States, Australia, Canada, Israel, India, and Japan through summer 2020.

Seeking intelligence on the United States and Europe
Russian threat actors attempted to access accounts

at organizations on almost every continent this period, but they predominantly focused on organizations based in the United States, followed by Ukraine, UK, and NATO allies and member states across Europe. On May 14, the Russian government officially named the United States and Czechia “unfriendly” countries, while Poland, Lithuania, Latvia, Estonia, UK, Canada, Ukraine, and Australia appeared on a preliminary list leaked in April.⁴¹ The top three countries most impacted by Russian cyber activity this past year—United States, Ukraine,

and UK—were on the “unfriendly countries” lists. Microsoft’s observations of Russian threat activity this past year suggest that intelligence collection was a primary motivation, as we saw data exfiltration but little evidence of disruptive or destructive activity from the groups we track. Gaining information on the policy plans and intentions of those perceived as adversaries would be standard intelligence requirements for the Russian government agencies to whom the US government attributes much of this activity.⁴²

What lessons might NOBELIUM have learned from the SolarWinds incident?

1. The US Government is still not sure where the red lines are for cyber operations.

As a sign of the ongoing debate within US and European policy communities about whether and how to respond to the SolarWinds breach, in March a former senior adviser to Britain’s Government Communications Headquarters cautioned the Biden administration not to react too harshly to Russia’s “surgical” espionage campaign.³⁹ Russian threat actors have exploited this policy ambiguity for years and could continue to do so for years to come.

2. The private sector is critical to the defense of US government networks.

Microsoft and FireEye were the public face of incident response during the SolarWinds attack. In the future, NOBELIUM and other groups could move early to handicap high-profile cybersecurity teams, anticipating that doing so will slow the time to identification and remediation of intrusions against high-value targets.

Information accessed	Operational aim
• Sanctions policy • Defense/intelligence policy • Russia policy • COVID-19 information	Espionage to gain policy insights
• Cyber incident response; threat hunting techniques • Assessments of Russian threat actors • Red Team tools • Detection signatures • Source code	Intelligence collection to improve countermeasures
• CSP accounts • Software certificates • Source code	Intelligence collection to support operational planning

Examples of the types of information NOBELIUM operators may have acquired, based on the victim accounts they accessed, and the operational aims that likely drove the intrusion.

³⁹ Top Biden cyber official: SolarWinds breach could turn from spying to destruction ‘in a moment’ (yahoo.com) ⁴⁰ <https://blogs.microsoft.com/on-the-issues/2020/11/13/health-care-cyberattacks-covid-19-paris-peace-forum/> ⁴¹ <https://tass.com/politics/1289825> ; <https://www.newsweek.com/russia-puts-us-top-unfriendly-countries-list-1586749> ⁴² <https://www.nsa.gov/News-Features/Feature-Stories/Article-View/Article/2573391/russian-foreign-intelligence-service-exploiting-five-publicly-known-cyber-attacks> ; https://media.defense.gov/2021/Apr/15/2002621240/-1/-1/0/CSA_SVR_TARGETS_US_ALLIES_UOO13234021.PDF/CSA_SVR_TARGETS_US_ALLIES_UOO13234021.PDF ; <https://home.treasury.gov/news/press-releases/jy0127>

<https://t.me/learningnets>

China

Over the past year, Microsoft observed Chinese nation state threat actors target the US political landscape for insight into policy shifts and target government entities that enact foreign policies in Europe and Latin American countries likely for intelligence collection. To accomplish their mission, several China-based threat actors exploited a range of previously unidentified vulnerabilities for different services and network components.

The following charts illustrate activity by China-based threat groups in July 2020–June 2021 based on NSNs issued to customers. These charts represent only a portion of the threat actors’ activities observed.

HAFNIUM and the Exchange vulnerabilities
In early March 2021, Microsoft blogged about HAFNIUM for the first time related to the detection of multiple 0-day exploits being used to attack on-premises versions of Microsoft Exchange Server.⁴³

HAFNIUM, a group assessed to be state sponsored and operating out of China, based on observed victimology, tactics, and procedures, primarily targets entities in the United States across a number of industry sectors, including infectious disease researchers, law firms, higher education institutions, defense contractors, policy think tanks, and NGOs. HAFNIUM has previously compromised victims by exploiting vulnerabilities in internet-facing servers and has used legitimate open-source frameworks, like Covenant, for command and control. Once


they’ve gained access to a victim network, HAFNIUM typically exfiltrates data to file sharing sites like MEGA. In campaigns unrelated to these vulnerabilities, Microsoft has observed HAFNIUM interacting with victim Office 365 tenants. While they are often unsuccessful in compromising customer accounts, this reconnaissance activity helps the adversary identify more details about their targets’ environments. HAFNIUM operates primarily from leased virtual private servers in the United States.

Activity Group Name	Other names	Country of origin	Industries targeted
MANGANESE	APT5, Keyhole Panda	China	Communications infrastructure, defense industrial base, software/technology
ZIRCONIUM	APT31	China	Government agencies and services, diplomatic organizations, economic organizations
HAFNIUM	---	China	Higher education, defense industrial base, think tanks, NGOs, law firms, medical research
NICKEL	APT15, Vixen Panda	China	Government agencies and services, diplomatic organizations
CHROMIUM	ControlX	China	Energy, communications infrastructure, education, government agencies and services
GADOLINIUM	APT40	China	Maritime, healthcare, higher education, regional government organizations

China



China: Top five targeted industries/sectors
(July 2020-June 2021)




The most prevalent targets of China-based threat activity were government entities worldwide. The targeting of three countries' government entities accounted for half of the NSNs issued and 23 countries accounted for the remaining half.

HAFNIUM used these vulnerabilities to access on-premises Exchange servers, which enabled access to email accounts and allowed installation of additional malware to facilitate long-term access to victim environments. MSTIC attributes this campaign with high confidence to HAFNIUM. The vulnerabilities exploited were CVE-2021-26855, CVE-2021-26857, CVE-2021-26858, and CVE-2021-27065.

The attacks included three steps. First, HAFNIUM would gain access to an Exchange Server either with stolen passwords or by exploiting the aforementioned vulnerabilities to gain initial


China: Target attempts vs successful compromise (July 2020-June 2021)



Chinese nation state threat actors were successful in compromising victims 44% of the time. However, because they are an advance persistent threat, if they are tasked to target an entity for intelligence collection, they will find another vulnerability to leverage to gain access.

access. Second, they deployed web shells on the compromised server. Web shells potentially allow attackers to steal data and perform additional malicious actions that lead to further compromise. Third, they used that remote access, which was typically run from the US-based private servers to exfiltrate data from an organization's network. Microsoft assesses HAFNIUM was associated with the initial activity with the 0-day exploits; however, after the vulnerability announcement, several nation state actors and criminal groups maneuvered quickly to take advantage of the vulnerabilities for their own gain.

HAFNIUM: Top targeted industries/verticals
(Prior to the increase in Exchange Server exploitation)



CVE	Description
CVE-2021-26855	Server-side request forgery (SSRF) vulnerability in Exchange, which allowed the attacker to send arbitrary HTTP requests and authenticate as the Exchange server.
CVE-2021-26857	An insecure deserialization vulnerability in the Unified Messaging service.
CVE-2021-26858	A post-authentication arbitrary file write vulnerability in Exchange.
CVE-2021-27065	A post-authentication arbitrary file write vulnerability in Exchange.

On July 19, 2021, the US government with its allies and partners took a stance against the Chinese government and issued a statement that China’s malicious cyber operation poses a major threat to the United States and its allies’ economic and national security.⁴⁴ Although slim on the technical details, the same statement attributed HAFNIUM with a high level of confidence to cyber actors affiliated with China’s civilian intelligence agency, the Ministry of State Security. These actors compromised tens of thousands of computers and networks worldwide in a massive cyber espionage campaign that mostly impacted private sector victims.

More 0-days and other exploitation of vulnerabilities

In July, SolarWinds released a security advisory for CVE-2021-35211, crediting Microsoft with the notification.⁴⁵ Microsoft detected the 0-day remote code execution exploit being used to exploit the SolarWinds Serv-U FTP software at entities in the US Defense Industrial Base Sector and software companies. This activity is attributed to a group operating out of China, based on observed victimology, tactics, and procedures.

In April 2021, FireEye released a blog and credited MSTIC for their contribution in identifying a Pulse Secure VPN 0-day exploit that was leveraged by Chinese nation state threat actors.⁴⁶ Microsoft associates some of the activity with MANGANESE and NICKEL. The Department of Homeland Security Cybersecurity and Infrastructure Security Agency (CISA) released an alert on the same 0-day activity indicating that it affected US government agencies, critical infrastructure entities, and other private sector organizations likely beginning in June 2020.⁴⁷ CISA stated that after the successful exploitation, the threat actor used their access to place web shells on the Pulse Connect Secure appliance for further access and persistence.

In addition to MANGANESE, MSTIC has observed ZIRCONIUM and two other threat actors who exploited small office or home office routers worldwide. These threat actors are likely compromising routers to use as infrastructure for their computer network operations. These compromised routers are likely in the same geographical area as their intended target to obscure scrutiny against the associated activity.

A worldwide intelligence collection operation

After the September 2020 Microsoft blog on multiple nation state threat actors targeting information on US elections, ZIRCONIUM did not stop their collection operations.⁴⁸ As the US presidential election day approached, ZIRCONIUM continued to employ web-bugged emails, targeting individuals with access to knowledge of potential shifts in US policy. On July 19, 2021, the United Kingdom’s National Cyber Security Centre released a statement that attributed APT31, which is roughly tracked as ZIRCONIUM by Microsoft, to the Ministry of State Security—China’s civilian intelligence agency.⁴⁹

Chinese nation state cyber operations did not overlook their neighbors. Since July 2020, activity tied to CHROMIUM targeted entities in India, Malaysia, Mongolia, Pakistan, and Thailand and the sensitive social, economic, and political issues surrounding Hong Kong and Taiwan. From Microsoft’s perspective, CHROMIUM activity was most active against universities in Hong Kong and Taiwan, followed by government entities and telecommunication providers in the other countries. In addition to targeting neighboring countries, there has been a steady drumbeat of intelligence

collection against Latin American countries and in Europe. Besides leveraging exploits for VPN devices in their cyber operations, NICKEL’s activity also targeted government foreign ministries throughout Central and South American countries and some European countries. As China’s influence continues to shift in the region and with countries that are partners in their Belt and Road Initiative, we assess that Chinese threat actors will continue to target entities to gain insight into investments, negotiations, and influence.

⁴⁴ <https://www.whitehouse.gov/briefing-room/statements-releases/2021/07/19/the-united-states-joined-by-allies-and-partners-attributes-malicious-cyber-activity-and-irresponsible-state-behavior-to-the-peoples-republic-of-china/> ; <https://www.ncsc.gov.uk/news/uk-allies-hold-chinese-state-responsible-for-pervasive-pattern-of-hacking> ⁴⁵ SolarWinds Trust Center Security Advisories | CVE-2021-35211 ⁴⁶ <https://www.fireeye.com/blog/threat-research/2021/04/suspected-apt-actors-leverage-bypass-techniques-pulse-secure-zero-day.html> ⁴⁷ <https://us-cert.cisa.gov/ncas/alerts/aa21-110a> ⁴⁸ <https://blogs.microsoft.com/on-the-issues/2020/09/10/cyberattacks-us-uk-allies-hold-chinese-state-responsible-for-pervasive-pattern-of-hacking>

Iran

Iran continued its streak of destructive cyberattacks against regional adversaries while taking a “wait-and-see” approach with the United States amid the prospects for sanctions relief from nuclear talks after US elections.

Focused on Israel with new attack tools amid broader escalation

As a covert war between Iran and Israel escalated, Iranian offensive cyber actors increased their attention to Israel and brought with them Iran’s newest tool of choice—ransomware. Iran also conducted ransomware attacks against at least one Gulf State adversary.⁵⁰ While it remains unclear whether Iranian actors are using ransomware for financial gain, in at least one case they used it as a cover for a destructive attack by deploying wiper malware on a company’s network while demanding a ransom.⁵¹


Microsoft detected an increased focus from a growing number of Iranian groups targeting Israeli entities since November, and with that focus came a string of ransomware attacks. An Iran-linked threat actor that we track as RUBIDIUM probably conducted the Pay2Key and N3tw0rm ransomware campaigns that almost exclusively targeted Israel in late 2020 and early 2021, respectively. One common element of RUBIDIUM’s ransomware campaigns was its targeting of Israeli logistics companies involved in maritime transportation. These targets indicate a link to Tehran’s broader objective of retaliating against Israeli pressure.⁵²

A wait-and-see approach toward the United States likely serves two purposes

Despite Tehran’s less aggressive approach toward the United States, relative to its regional adversaries, US entities remained Iranian threat actors’ top target, comprising nearly half of the NSNs we delivered to cloud-service customers. Iranian cyber operations toward US targets consisted of a two-pronged approach: acquiring strategic intelligence likely to gain insights into US policy views and planning and acquiring a foothold on networks likely to provide Tehran with contingency options in case the United States failed to provide sufficient sanctions relief.

Activity Group Name	Other names	Country of origin	Industries targeted
PHOSPHORUS	Charming Kitten	Iran	Diplomatic and nuclear policy communities, academics, and journalists
CURIUM	Houseblend Tortoise Shell	Iran	US military and defense contractors, IT services, and Middle Eastern governments
RUBIDIUM	Fox Kitten Parasite	Iran	Israeli logistics companies, IT services, and defense

Iran



⁵⁰ <https://labs.sentinelone.com/from-wiper-to-ransomware-the-evolution-of-agrius> ; <https://unit42.paloaltonetworks.com/thanos-ransomware/> ⁵¹ <https://labs.sentinelone.com/from-wiper-to-ransomware-the-evolution-of-agrius> ; <https://www.clearskysec.com/wp-content/uploads/2020/12/Pay2Kitten.pdf> ; <https://www.flashpoint-intel.com/blog/second-iranian-ransomware-operation-project-signal-emerges/> ⁵² <https://www.aljazeera.com/news/2021/4/25/top-iranian-commander-hints-at-future-operations-against-israel> ; <https://www.timesofisrael.com/eye-for-an-eye-iran-editorial-urges-retaliatory-attack-on-dimona-reactor/> ; <https://www.al-monitor.com/originals/2021/04/iranian-military-leader-threatens-israel-following-missile-strike-syria>


<https://t.me/learningnets>

In late 2020, PHOSPHORUS began targeting nuclear policy experts in signatory nations of the 2015 Joint Comprehensive Plan of Action, very likely for intelligence to gain an edge in anticipated talks on the accord following President Biden’s election. PHOSPHORUS conducted a credential phishing campaign by masquerading as fellow foreign and nuclear policy experts and sending links to nuclear-themed articles that directed victims to a credential harvesting site. They targeted fewer than 25 senior personnel at medical research organizations, as detailed by Proofpoint,⁵³ but the vast majority of the 100-plus targets we detected were nuclear policy or conflict resolutions experts—in line with the theme of the phishing emails—in the United States, United Kingdom, France, and Russia. PHOSPHORUS honed its targeting on this community as nuclear talks began in Vienna in April, including targeting diplomat participants.

Previously in 2020, PHOSPHORUS masqueraded as conference organizers to high-profile international conferences, as we detailed in this blog.⁵⁴ At Microsoft, we detected PHOSPHORUS sending spoofed email invitations with links to credential harvesting sites to more than 100 policy experts who were prospective attendees, several of whom they compromised. The group’s credential phishing campaign likely sought to acquire intelligence to better position itself in international engagement. Since April 2021, select Iranian actors also targeted US agriculture and media companies that are unlikely intelligence targets for Tehran.⁵⁵ These same operators employed ransomware on other companies, suggesting a potential aim to gain a foothold for contingency plans in case nuclear talks fail to meet Tehran’s expectations.


In some cases, the Iranian targeting of US entities that we have detected could be focused on intelligence collection, contingency planning, or both. Early this year, CURIUM conducted a spear-phishing campaign targeting companies that provide IT and engineering services for US defense and intelligence agencies, probably as a part of a supply chain operation to gain access to their customers.

Iran: Most targeted countries (July 2020–June 2021)



⁵³ <https://www.proofpoint.com/us/blog/threat-insight/badblood-ta453-targets-us-and-israeli-medical-research-personnel-credential> ⁵⁴ <https://blogs.microsoft.com/on-the-issues/2020/10/28/cyberattacks-phosphorus-t20-munich-security-conference/> ⁵⁵ DEV-0270 Compromise Agrinos on 30 May. <https://spectre.microsoft.com/#/entry/19f1e6f3fe899dc1f315a9c432c597a3d4518115604afd63c169948ba7bc95cf?nonce=72c1a4aa8705>; DEV-0270 compromised CURIUM on 15 May. <https://spectre.microsoft.com/#/entry/201936a4ffcde2e1eff5d21b43834c7e38631e47e1f66b9ab3bc1e1a135f074f?nonce=2b351540141b>

Flow of a typical PHOSPHORUS compromise from spear phish



Conferences, conventions, and trade shows are widely known throughout industry and the US government as a hotbed of intelligence collection activities, both by domestic competitive intelligence and foreign adversaries. Individuals have been known to collect information thrown out in the trash, record presentations, attempt to steal products, and solicit sensitive information from employees. Though these events were widely paused due to pandemic restrictions, major conventions are coming back to calendars.

North Korea

In the last year, North Korean threat actors have been extremely active relative to the country's size and resources, compared to the other major attacking states. For example, in the last three months of 2020, just over half the NSNs Microsoft issued were for North Korean state actors, in spite of North Korea being the smallest of the four most prolific nation state actors Microsoft tracks.

Feeding a vast appetite for intelligence

The vast majority of the North Korean targeting Microsoft noted was directed at consumer account targets. For the most part, these targets were probably selected based on the likelihood they could help North Korea obtain non-publicly available diplomatic or geopolitical intelligence. North Korean groups THALLIUM and ZINC continued to create much of the targeting Microsoft observed, but they were joined by other groups, such as OSMIUM and CERIUM. Together, these groups focused on diplomatic officials, academics, and think tank members from around the world. Most


of those targeted were in three countries: South Korea, the United States, and Japan. However, North Korean actors also targeted academics and think tank officials in Europe and even China and Russia, countries generally seen as friendly to North Korea.

The focus on diplomatic or geopolitical intelligence likely was driven by Pyongyang's anxiety for information in a volatile international situation. Diplomatic targeting was particularly heavy during and directly after the US election. North Korea's strong interest in intelligence collection probably had several key questions it sought to answer: Will

the international community continue to strictly enforce sanctions on North Korea? How does COVID-19 change international dynamics? What will the new US administration's policy be toward North Korea, and how will the tripartite US-South Korea-Japan partnership pursue that policy together?

Activity Group Name	Other names	Country of origin	Industries targeted
ZINC	Lazarus Labyrinth Chollima	North Korea	Utilities, private companies, think tanks, security researchers
THALLIUM	Kimsuky Velvet Chollima	North Korea	Think tanks, diplomatic officials, academics
CERIUM	Kimsuky	North Korea	Think tanks, diplomatic officials, academics, defense and aerospace
OSMIUM	Konni	North Korea	Diplomatic officials, think tanks


North Korea



Global pandemic creates a new type of cyberattack

COVID-19 also drove another North Korean focus in the last year: the targeting of pharmaceutical companies. As Microsoft reported in November 2020,⁵⁶ ZINC and CERIUUM targeted pharmaceutical companies and vaccine researchers in several countries, probably to speed up its own vaccine research or to gain intelligence on the state of research in the rest of the world.

North Korea: Top 5 targeted industries and sectors (July 2020-June 2021)




Many of the consumer accounts were likely personal accounts of academics, think tank members, and government officials.

The world's only known nation state Bitcoin thieves

Alone among nation state actors, North Korea continued in the last year to target financial companies with the intent of stealing cryptocurrency and intellectual property. North Korea's economy, already under strain from sanctions, was put under even greater stress when it closed its borders to trade after the outbreak of COVID-19. Cyber-enabled theft presented one opportunity to make up lost

North Korea: Failed attempts vs. successful compromise (July 2020-June 2021)



Relentless spear-phishing attempts by groups such as THALLIUM do not often succeed, but because they are so ubiquitous, even occasional success yields big results.

income. One such group Microsoft tracks, which we have not named, often targeted cryptocurrency or blockchain research companies with spear-phishing campaigns while posing as cryptocurrency or blockchain start-ups.

A sophisticated social engineering campaign targeting security researchers

Finally, North Korea also used social engineering in ways not seen from it before. As MSTIC reported in concert with Google in January,⁵⁷ ZINC targeted security researchers with a fairly sophisticated social engineering campaign. The campaign included spending months setting up fake profiles that looked like real security companies and researchers, with websites and social media platforms to support these personas. This targeting sought long-term effects beyond the immediate attack. It also showed that North Korea is more than capable of understanding the Western security landscape well enough to blend into it.

Vietnam

Vietnamese threat group BISMUTH utilized cryptocurrency miners to target private sector and government institutions in France and Vietnam. Because cryptocurrency miners tend to be seen as lower-priority threats by security systems, BISMUTH was able to take advantage of the smaller alert profile caused by their malware to slip into systems unnoticed.

As MSTIC reported in November 2020,⁵⁸ BISMUTH carefully planned attacks, conducting reconnaissance before creating uniquely crafted spear-phishing emails for each individual. Sometimes, BISMUTH actors, similar to PHOSPHORUS operators, would correspond with targets to build rapport before sending the email containing a malicious attachment. Once it compromised networks, BISMUTH sought to achieve continuous monitoring. Its targets included human and civil rights organizations.

Activity Group Name	Other names	Country of origin	Industries targeted
BISMUTH	APT32 OceanLotus	Vietnam	Human rights and civil organizations

Vietnam



Turkey

SILICON pursues intelligence collection for strategic Turkish interests from a variety of countries, primarily in the Middle East and the Balkans. Their reconnaissance indicates the group is most heavily focused on countries of strategic interest to Turkey including Armenia, Cyprus, Greece, Iraq, and Syria. They regularly target telecommunication and IT companies, likely to establish a foothold upstream of their desired target, and often seek access by scanning infrastructure for remote code vulnerabilities.

Activity Group Name	Other names	Country of origin	Industries targeted
SILICON	Sea Turtle UNC1326	Turkey	Telecommunications companies in the Middle East and the Balkans

Turkey



Private sector offensive actors

A growing industry of companies called private sector offensive actors (PSOAs) create and sell malicious cyber technologies that enable their customers to break into people’s computers, phones, and internet-connected devices. These private companies may not be nation state actors, but their business model presents a dangerous and rapidly growing challenge for organizations, companies, and individual consumers. These tools also threaten many global human rights efforts, as they have been observed targeting and surveilling dissidents, human rights defenders, journalists, civil society advocates, and other private citizens.

In December 2020, Microsoft’s efforts to protect our customers from the threats presented by this technology led us to file an amicus brief in support of WhatsApp’s case against Israel-based NSO Group Technologies (NSO Group) along with Cisco, GitHub, Google, LinkedIn, VMware, and Internet Association.⁵⁹ The brief encouraged the court to reject NSO Group’s position that it is not responsible for the use of its surveillance and espionage products by governments. Microsoft also worked with Citizen Lab, at the University of Toronto’s Munk School, to disable malware being used by an Israel-

based PSOA that Microsoft calls SOURGUM, and that Citizen Lab identified as Candiru.⁶⁰ SOURGUM created malware and 0-day exploits (fixed in CVE-2021-31979⁶¹ and CVE-2021-33771⁶²) as a part of a hacking-as-a-service package sold to government agencies and other malicious actors. The malware was used to target more than 100 victims around the world including politicians, human rights activists, journalists, academics, embassy workers, and political dissidents. To limit these attacks, Microsoft has created and built protections into our products against this unique malware, which we call DevilsTongue.⁶³ By examining how SOURGUM’s customers were delivering DevilsTongue to victim computers, we saw they were doing so through a chain of exploits that impacted popular browsers and our Windows operating system. We published details of the malware and 0-day exploits so that the world can better understand SOURGUM’s activity and address and mitigate the threat. Private companies should remain subject to liability when they use their cyber-surveillance tools to break the law, or knowingly permit their use for such purposes, regardless of who their customers are or what they are trying to achieve. Microsoft will continue to identify, track, and protect our customers and global digital ecosystem from the indiscriminate attacks caused by PSOA technology and pursue other methods to disrupt this growing threat to our customers.

Comprehensive protections required

Nation state actors have demonstrated that they will go to great lengths to accomplish a mission to collect information or intelligence. The skill and persistence of malicious nation state actors increase the difficulty of detecting and protecting against advanced threats. Their impact can be wide ranging and highly damaging. These adversaries are well-funded, employ techniques of tremendous breadth and sophistication, and are motivated by objectives of national significance—which may lead to their compromising networks for unexpected purposes. More than other adversaries, nation state attackers target individuals specifically for access to their connections, communications, and information. At the conclusion of an operation, they will assess what went well and what did not and refine tactics and techniques for more successful future missions.

Therefore, defense-in-depth strategies against nation state adversaries should include educating employees on how to avoid being targeted themselves. Applying Zero Trust principles across corporate resources helps more effectively adapt to the complexity of the modern environment, embrace the mobile workforce, and protect people, devices, applications, and data no matter where they are located or the scale of threats they face.

While nation state attacks are often sophisticated or can deploy 0-day vulnerabilities to gain access to networks, defense-in-depth strategies and proactive monitoring can greatly reduce the actor’s dwell time on a network, potentially enabling disruption of their activities before they reach their goals. Above and beyond enabling foundational basics like MFA, IT departments should prioritize steps to mitigate lateral movement by attackers; specifically, credential hygiene and network segmentation. To limit the damage of data exfiltration, information rights management can be applied to files. Building protective controls across your managed identities, devices, applications, data, infrastructure, and networks will raise the threshold for attackers, improving your organization’s ability to detect anomalous activity in the environment.

⁵⁹ Amicus Brief 12.20.2020 (microsoft.com) ⁶⁰ Hooking Candiru: Another Mercenary Spyware Vendor Comes into Focus - The Citizen Lab ⁶¹ CVE-2021-31979 - Security Update Guide - Microsoft - Windows Kernel Elevation of Privilege Vulnerability ⁶² CVE-2021-33771 - Security Update Guide - Microsoft - Windows Kernel Elevation of Privilege Vulnerability ⁶³ Fighting cyberweapons built by private businesses - Microsoft On the Issues

CHAPTER 4

Supply chain, IoT, and OT security

Introduction

Challenges in managing risk associated with the supplier ecosystem

How Microsoft thinks about supply chain

IoT and OT threat landscape

The 7 properties of highly secured devices

Applying a Zero Trust approach to IoT solutions

IoT at the intersection of cybersecurity and sustainability

IoT security policy considerations

INTRODUCTION: Innovation-driven opportunity in an exponentially larger attack landscape

MICHAL BRAVERMAN-BLUMENSTYK, CORPORATE VICE PRESIDENT, CHIEF TECHNOLOGY OFFICER, CLOUD AND AI SECURITY

In the past year, we have observed an abundance of incidents driving both physical and digital disruption of operations for many organizations. These incidents preyed at times on the physical realm, such as disruption of production lines and energy substations, and in other cases, they were conducted entirely in the digital realm, such as via a ransomware campaign.

Looking at the attack surfaces that were exploited provides an additional perspective: from legacy operational technology (OT) equipment to brand new Internet of Things (IoT) devices; from seemingly ordinary cloud-migration projects to 5G IoT-related endeavors; and from physical supply chain to digital supply chain. All of these are playing an increasing role as fertile attack surfaces. These are all topics we will elaborate on in this chapter.

Supply chain integrity

Supply chains, both physical and digital, have an explicit reliance on trust, and adversaries have taken notice. Over the last decade, successful organizations have been able to meet the demands of scale, efficiency, and speed by building expansive, and often complex, ecosystems to deliver value to stakeholders. Security adversaries today view these systems as

targets for exploitation, as witnessed recently in the highly visible and impactful SolarWinds and Kaseya attacks. While threats and attacks continue to intensify, supply chain complexity increases the costs of defending and the likelihood that an exposure can produce a significant return for an adversary.

IoT Security

IoT security is a geometrically expanding frontier with innovation-driven opportunity and an exponentially larger attack landscape. The adoption of IoT and the huge acceleration in remote services, both at home and in the workplace since the onset of the COVID-19 pandemic, increases the likelihood of risk materializing. This is a trend that will continue as technologies like 5G and innovative IoT applications become more ubiquitous.

OT Security

OT devices, such as industrial control, hospital monitoring, or water management systems, represent the public infrastructure that many societies have come to depend on for decades. Many are lagging in adopting and leveraging modern security standards. As evidenced by recent attacks on water, transportation, and energy utilities, disruption in these areas have profound and broad impact

The chapter includes discussions about how organizations can understand and improve their IoT, OT, and supply chain security posture. We share in these discussions our data-driven perspectives related to the IoT and OT threat landscape, research findings by the Azure Defender for IoT team, global initiatives such as the Global Cyber Alliance, and more.

THE ADOPTION OF IOT AND THE HUGE ACCELERATION IN REMOTE SERVICES, BOTH AT HOME AND IN THE WORKPLACE, INCREASES THE LIKELIHOOD OF RISK MATERIALIZING.

Challenges in managing risk associated with the supplier ecosystem


As outsourcing for applications, infrastructure, devices, and human capital expands, the adoption of tools to monitor multiple tiers of suppliers for quality, security, integrity, and resilience risks is also growing, and the plethora of frameworks and approaches that organizations are leveraging today continues to grow in tandem. Additional complexity emerges when frameworks are applied inconsistently within an organization and across suppliers, or if multiple frameworks are in play.

When it comes to risk assessment and management, siloes can create additional problems. Different teams have different priorities, which can lead to completely different risk appetites, priorities, practices, and cultures. This inconsistency can be inefficient and create a duplication of effort, gaps in risk analysis, and an inability to effectively share risk information across the organization.

An always-on, automated, integrated approach is needed, but current processes aren’t well-suited to evolve:

- Supplier assessment and review processes often include just a questionnaire.
- Once a supplier is onboarded, there is only a point-in-time annual review cycle.
- Often, different teams within the same company have different processes and functions and no clear way to share information across teams. This can make it difficult to create a holistic and automated view of organizational risk.

Siloed environments pose challenges to risk assessment and management



What we’re hearing

Recently, the Microsoft M365 Security, Compliance, and Management team hosted an event to discuss the challenges and strategic needs of CxOs (CISOs, CIOs, VPs/Heads of IT and Governance, and others) and their organizations to gain a deeper understanding of their security and risk management experiences. These are some of the takeaways:

1. The selection and management of suppliers is shaped by a host of factors leading to lack of clarity and low trust, as many struggle to know their environment.

Organizations must balance protecting themselves from human liability, issues inherent with hybrid work, shadow IT (unknown or unmanaged apps, services, and infrastructure developed and managed outside standard policies), diversity of their digital estate, and evolving threats and vulnerabilities. Adding complexity, there are usually several parties weighing in on vendor/supplier selection. Security is just one of the factors considered and often not a top priority, even though it is an area that needs immediate attention. Organizations end up with long lists of suppliers that must be managed (some of them unknown), often with a limited look into their security practices and posture. Add to this that working within the confines of a supplier contract often limits the assessments. As a result, it is difficult to have desired visibility and trust in suppliers.

2. Proactive management of a supplier ecosystem is ideal, but difficult. Designated critical suppliers require focus due to the potential risk they pose.

Organizations are often forced into a split approach to supplier management—proactive, when possible, but typically more reactive. Reactive approaches are a result of limited resources and projects scoped without security as a consideration. To ensure their protection is as comprehensive as possible, organizations keep critical suppliers (those critical to an organization’s mission) on a shorter leash, with less flexibility and more oversight.

3. Leaders are devoting more resources to supplier security, but recent breaches demonstrated that traditional models couldn't guarantee safety.

IT teams are appreciative of the influx of budget and resources following the SolarWinds and Colonial Pipeline breaches. That appreciation comes with an equal amount of trepidation. Those breaches demonstrated that the very strategies they are implementing likely would not protect them from a similar attack, and continuous assessments and push for remediation across tiers of suppliers are necessary.

4. Greater visibility and unique solutions are top requests to managing suppliers.

Within the four key pillars of digital estates—identities, applications, infrastructure, and devices—suppliers' personnel are a top concern. Across that risk hierarchy, organizations are looking for suppliers to provide:

- Greater visibility into security and who ultimately has access to the organization's data.
- Customized solutions demonstrating a working knowledge of the industry and the company-specific needs.

Zero Trust security model for supplier ecosystem risk

For supplier risk management, having customized solutions and greater visibility into who ultimately has access to an organization's data across domains are top priorities. While there are many places to begin your Zero Trust journey, from a supplier ecosystem and risk management standpoint, instituting multifactor authentication (MFA) should be a priority.

For more information on Zero Trust strategy, see the [Hybrid workforce security](#) chapter of this report.

How Microsoft thinks about supply chain

The end-to-end supply chain and supplier ecosystem is complex and opaque, extending from development to build, chips to firmware, drivers, operating system, third-party applications, manufacturing/factory, and all the way to secure updates. Governments and critical infrastructure providers are looking for a new level of assurance for supply chain security and continuity. It's important that a repeatable process will continue to scale as organizations continue to innovate. Supply chain security rigor is foundational to how an organization should work and is expected by partners and customers who interact with an organization's products and services.


Nine secure supply chain focus areas

Outlined below is a framework to efficiently evaluate your supply chain ecosystem with considerations for how you might approach protecting it. We group our investments into nine secure supply chain workstreams to methodically evaluate and mitigate risk of exposure in each area.

1. First-party engineering systems for hardware and software

Massive software development companies aren't the only ones doing engineering work. Even small IT shops that may make minimal development investments and organizations building on top of existing infrastructure have teams writing code.

Nine areas of investment for a secure end-to-end supply chain



To produce secure software or hardware, organizations must ensure that the first-party engineering system is secured from the various threat vectors that could be exploited by attackers.

Developer environment: The developer environment encompasses the tools and platforms that developers author code on, such as operating systems, code editors, research tools (for example, browsers and associated websites), local build tools, and other general code authoring tools. The primary target is developer identity, and solutions must be put in place to mitigate this risk.

Source code: The risk of malicious code can arise when developers take source code and binaries from a variety of sources, including from internal sources, open-source software (OSS), or from another organization. Each source needs to have a level of trust ensured and known to have appropriately secure supply chain checks. In most reported OSS-as-malware cases, the malware is designed to steal developer credentials and build environment variables, exfiltrating them to a remote attacker-controlled server.

Build: Without appropriate protections, the build pipeline presents product compromise risks that can be highly effective and difficult to trace. While source control will usually have some form of change management, without appropriate controls in place in a build system, it can be difficult to identify what might have been injected during the build process.

Release: Ensuring you are releasing what you intended is the end goal of a secure software supply chain. Enumerating inputs and verifying the final product takes a comprehensive understanding of complex systems and is an essential final step in ensuring an organization knows what it is releasing.

2. Firmware and driver security

Firmware and drivers are the foundation of most hardware devices. If hacked and embedded with malware, they pose a huge risk to the hardware device and the organization that depends on it, potentially creating unauthorized access, making it inoperable or even unbootable. Organizations need to ensure that all firmware and drivers installed on servers or end-user equipment follow the required security requirements and have the necessary documentation to prove their compliance.

3. Physical security

Organizations must define, implement, and manage appropriate security controls to ensure the security

4. Manufacturing security

Manufacturing standards must be defined and enforced to enable detection, protection, and recovery from cyberattacks. Organizations must also ensure that samples or prototypes are securely handled and stored, and that appropriate monitoring is in place to track and maintain the chain of custody for all proprietary items or finished goods.

5. Logistics security

Logistics functions must be safeguarded to prevent tampering, loss, or theft of the products during transportation and storage. Hardware and devices teams should have appropriate operational controls and frameworks in place with the suppliers to ensure receiving, shipping, storage, and other logistics management nodes are secure and compliant with the company’s standards.


6. Supplier security

When engaging with suppliers, organizations must ensure that the suppliers comply with well-defined supplier security and privacy assurance requirements throughout the duration of their partnership.

LEVERAGING MACHINE LEARNING IN CONTINUOUS SECURITY MONITORING OF SUPPLIERS

Microsoft leverages machine learning (ML) to scan active supplier contracts. This model is trained to recognize commonly negotiated security clauses and determine whether or not they meet the intent of the original requirement. The outputs of this continuous scanning are leveraged to advise the third parties operating within our environments and ensure their expectations and accountabilities are clearly defined.

Continuous security monitoring using ML



Learn more:

[Transforming how Microsoft connects with its 58,000 suppliers – Inside Track Blog](#)

7. Security validations and assurances

The active nature of the adversaries requires continuous evaluation of the security posture of the end-to-end supply chain to identify and prioritize security investments. A mix of internal audits, lessons learned from recent events, and penetration testing provides assurance that the right controls are in place to detect, prevent, and/or mitigate these attacks. The findings from these engagements help identify the next set of requirements to be addressed.

8. Trust chain governance and resilience

Customers trust their supplier organizations protect them as they use their products and services. First-party trust chains provide identity, integrity, and non-repudiation for the organization’s platform, products, and services. The main elements of trust chains are public key infrastructure, cryptographic algorithms, hardware, and the supporting teams and facilities. Effective governance of third-party trust chains is critical.

9. Monitoring and detections

Each of the supply chain focus areas above require monitoring, detection, and follow-up actions when questionable activity is identified. Increasingly monitoring, detection, and initial response are automated in today’s hyper-scaled cloud operations. These actions cover common cases but can also be adaptable to deal with new or unimagined scenarios.

Learn more:

[Firmware security - Azure Security | Microsoft Docs \(6/24/2021\)](#)

[Microsoft Open Source Software Security](#)

[Datacenter security overview - Microsoft Service Assurance | Microsoft Docs \(8/23/2021\)](#)

[Physical security of Azure datacenters - Microsoft Azure | Microsoft Docs \(7/10/2020\)](#)

[Supply Chain Security - Microsoft Research](#)

[Microsoft Security Development Lifecycle](#)

US Executive Order and supply chain security

Issued on May 12, 2021, the Executive Order (EO) on Improving the Nation’s Cybersecurity (EO 14028) outlines significant new steps for US federal agencies and their technology providers to strengthen IT modernization, improve incident response, and enhance software supply chain security.⁶⁶ Section 4 focuses on software supply chain security, enumerating areas of requirements to be developed for both software providers and federal agency users of software. For software providers, the EO calls for requirements to enhance ability to resist tampering or attack and foster greater transparency into components, including through secure software development practices and environments, use of tools or processes for software verification and vulnerability checks, providing of Software Bill of Materials (SBOM) information, participation in a vulnerability disclosure program, and other practices. For federal agency users of

Learn more:

[Microsoft and NIST collaborate on EO to drive Zero Trust adoption | Microsoft Security Blog \(8/17/2021\)](#)

[CYBER EO | Microsoft Federal](#)

[Microsoft - Executive Order - NIST workshop position paper 4- Testing software source code Microsoft Corporation.pdf](#)

[Microsoft - Executive Order - NIST workshop position paper 5- Software integrity chains Microsoft Corporation.pdf](#)

[Microsoft’s approach to coordinated vulnerability disclosure](#)

[NTIA_RFC_SBOM_Minimum_Elements_MSFT_Response_061721.docx.pdf](#)

software with privileged access or other attributes that make it especially critical, the EO calls for security measures, published by the National Institute of Standards and Technology (NIST) in July, to manage operational risk. Microsoft has long invested in developing best practices for secure software development, software testing, and vulnerability disclosure and management programs, and we’ve contributed to efforts to define industrywide practices and consensus standards, including through SAFECode,⁶⁸ ISO,⁶⁹ and NIST.⁷⁰ Along with GitHub, Microsoft has also contributed to efforts to develop best practices and specifications to define use cases for and support the delivery of SBOMs, which identify what software is composed of and allow software providers to associate information with components. Microsoft and GitHub support delivering SBOMs to enable vulnerability and integrity checks, and we’re committed to leveraging SBOMs as part of a broader evidence store that would verify end-to-end supply chain integrity.

Know your suppliers and understand their supply chains.

⁶⁶ Executive Order on Improving the Nation’s Cybersecurity | The White House ⁶⁷ Security Measures for EO-Critical Software Use | NIST ⁶⁸ Fundamental Practices for Secure Software Development, Third Edition - SAFECode ⁶⁹ ISO - ISO/IEC 27034-1:2011 - Information technology — Security of information technology — Part 1: Overview and concepts ⁷⁰ Secure Software Development Framework | CSRC (nist.gov)

IoT and OT threat landscape

Successful solutions today often depend on the convergence of many components, including hardware, software, and cloud services, which often come together in an IoT solution. IoT is more than connected devices—it's about the information those devices collect and the powerful, immediate insights that can be garnered from that information. Accordingly, IoT and other embedded and OTs have become critical business, operational, and security topics. More than ever, IoT and OT security is finding its way into corporate boardrooms and state and federal legislature discussions as a high-priority issue, in part due to the increasing frequency and severity of attacks in the past year. This proliferation of attacks has also driven increased awareness of the extent to which cyberattacks in the digital realm can impact the physical realm: the Colonial Pipeline cyberattack directly led to shutdown of the largest conduit for gasoline in the United States. The compromise of the Oldsmar water

plant⁷¹ led to a hazardous situation, in which cyber actors obtained unauthorized access and used the SCADA system's software to increase the concentration of sodium hydroxide—a caustic chemical—in the water. The hack of a security camera provider⁷² exposed sensitive footage from hospitals, police departments, and a plethora of other companies.

All these developments underscore the need for organizations to secure their IoT and OT footprints. Organizations are interconnected more than ever, resulting in increased exposure of legacy OT devices and environments, including those that have existed in relative isolation. On the other hand, the newest IoT devices (such as smart TVs and smart sensors) reside in both OT and IT environments. Putting all of this together, with the added context of privacy concerns and regulatory compliance, stresses the need for a holistic approach that enables seamless security and governance across all OT and IoT devices.

Learn more:

[SCADA Hacking: The Most Important SCADA/ICS Attacks in History \(hackers-arise.com\)](#) (4/12/2021)

Evolving cyberthreats

Enterprises are having to contend with evolving cyberthreats and novel malware. These issues include supply-chain attacks such as HAVEX⁷³ and SolarWinds,⁷⁴ 0-day industrial control systems (ICS) malware such as Triton⁷⁵ and Industroyer,⁷⁶ fileless malware,⁷⁷ and living-off-the-land tactics using standard administrator tools,⁷⁸ which are harder to spot because they blend in with legitimate day-to-day activities. These attacks have also increased in frequency and severity in the past year.

From a technical perspective, the Triton attack on the safety controllers in a Middle East petrochemical facility was intended to cause major structural damage to the facility and possible loss of life. The attackers got their initial foothold in the IT network and subsequently used living-off-the-land tactics to gain remote access to the OT network, where they deployed their OT purpose-built malware.




⁷¹ [Lessons Learned from Oldsmar Water Plant Hack – Security Today](#) ⁷² [Hackers breach Verkada's giant trove of security-camera data collection | Fortune](#) ⁷³ <https://en.wikipedia.org/wiki/Havex>

⁷⁴ <https://msrc-blog.microsoft.com/2020/12/13/customer-guidance-on-recent-nation-state-cyber-attacks/> ⁷⁵ <https://www.fireeye.com/blog/threat-research/2017/12/attackers-deploy-new-ics-attack-framework-triton.html>


⁷⁶ <https://www.zdnet.com/article/industroyer-an-in-depth-look-at-the-culprit-behind-ukraines-power-grid-blackout/> ⁷⁷ <https://docs.microsoft.com/en-us/windows/security/threat-protection/intelligence/fileless-threats>

How an attacker can get into an enterprise through IoT




What we’re seeing: IoT-related malware in the wild


Distribution of IoT command and control services by country (July 2020–June 2021)



Top IoT malware detected in the wild (July 2020-June 2021)



Distribution of IoT malware CPU architecture (July 2020–June 2011)



Findings: Industrywide IoT and OT vulnerabilities

The Microsoft Defender for IoT team conducts research on various types of equipment ranging from legacy industrial control system controllers to cutting-edge IoT sensors. Upon discovery of a vulnerability, the findings are shared with relevant vendors through a responsible disclosure process led by the Microsoft Security Response Center and the US Department of Homeland Security (DHS), enabling these vendors to investigate and patch the vulnerability.

In April 2021, we uncovered a series of critical memory allocation vulnerabilities in IoT and OT devices that adversaries could exploit to bypass security controls to execute malicious code or cause a system crash. The group of vulnerabilities was dubbed BadAlloc. **These remote code execution vulnerabilities affected a wide range of industries and verticals, from consumer and medical IoT, to industrial IoT, OT, and industrial control systems. They covered more than 25 common vulnerabilities and exposures (CVEs).** In the context of IT environments, an exploitation of such a vulnerability can result in a loss of confidentiality. In the context of OT environments, it can be used to trigger a disruption of operations.

The vulnerabilities stem from using vulnerable memory functions, such as malloc, calloc, realloc,

memalign, valloc, pvalloc, and others. Our research showed that memory allocation implementations written throughout the years as part of IoT devices and embedded software have not incorporated proper input validations. Without these input validations, an attacker could exploit the memory allocation function to perform a heap overflow, resulting in execution of malicious code on a target device.

The memory allocation vulnerabilities can be invoked by calling the memory allocation function, such as malloc(VALUE), with the VALUE parameter derived dynamically from external input and being large enough to trigger an integer overflow or wraparound. The concept is as follows: When sending this value, the returned outcome is a freshly allocated memory buffer. While the size of the allocated memory remains small due to the wraparound, the payload associated with the memory allocation exceeds the actual allocated buffer, resulting in a heap overflow. This heap overflow enables an attacker to execute malicious code on the target device.

BadAlloc is a case that illustrates the extensive impact these vulnerabilities can have because the risk exists in IoT and OT devices across all major industries. This example highlights one of the biggest challenges in mitigating IT, IoT, and OT risks: they share attack surfaces, and attackers look at the entire ecosystem.

The following is an example of BadAlloc:

```
void * pvPortMalloc( size_t xWantedSize )
{
    BlockLink_t * pxBlock, * pxPreviousBlock, * pxNewBlockLink;
    static BaseType_t xHeapHasBeenInitialised = pdFALSE;
    void * pvReturn = NULL;

    vTaskSuspendAll();
    {
        /* If this is the first call to malloc then the heap will require
         * initialisation to setup the list of free blocks. */
        if( xHeapHasBeenInitialised == pdFALSE )
        {
            prvHeapInit();
            xHeapHasBeenInitialised = pdTRUE;
        }


        /* The wanted size is increased so it can contain a BlockLink_t
         * structure in addition to the requested amount of bytes. */
        if( xWantedSize > 0 )
        {
            xWantedSize += heapSTRUCT_SIZE;

            /* Ensure that blocks are always aligned to the required number of bytes. */
            if( ( xWantedSize & portBYTE_ALIGNMENT_MASK ) != 0 )
            {
                /* Byte alignment required. */
                xWantedSize += ( portBYTE_ALIGNMENT - ( xWantedSize & portBYTE_ALIGNMENT_MASK ) );
            }
        }
    }
```

The protection of IoT and OT devices from exposure to IT risks becomes more important as they converge. Too often these risks are addressed in isolation with a rigidly siloed approach. To be successful in countering attacks, risks should be dealt with holistically while also accommodating domain expertise in each area. It is also critical

to ensure that modern digital environments are not held back by threats from legacy technology connected to OT systems. Mitigation requires an integrated approach that spans the entire enterprise. Organizations should look for opportunities to harden, patch, or segment systems to reduce the attack surface.

BadAlloc impact matrix



BadAlloc is an example of an IoT/OT family of vulnerabilities that poses a risk across all industries. The extent and nature of the risk depends on the specific context of usage of the device. BadAlloc shouldn't be treated as only an OT or IT issue; instead, organizations should take a holistic approach to mitigating the risk.

Mitigating IoT and OT vulnerabilities such as BadAlloc

We recommend the following mitigation strategies for organizations with IoT and OT devices:

Patch, patch, patch. Follow vendor instructions for applying patches to affected products.

If you can't patch, monitor. Since most legacy IoT and OT devices don't support agents, use an IoT/OT-aware network detection and response (NDR) solution⁷⁹ and a SIEM/SOAR solution⁸⁰ to auto-discover and continuously monitor devices for anomalous or unauthorized behaviors, such as communication with unfamiliar local or remote hosts. These elements are essential in implementing a Zero Trust strategy for IoT/OT.

Reduce the attack surface. Eliminate unnecessary internet connections to OT control systems and implement virtual private network (VPN) access with MFA when remote access is required. The US DHS warns that VPN devices may also have vulnerabilities and should be updated to the most current version available.

Segment. Network segmentation is important for Zero Trust because it limits the attacker's ability to move laterally and compromise assets after initial intrusion. IoT devices and OT networks should be isolated from corporate IT networks by using firewalls.

Learn more:

[Eliminating IoT vulnerabilities using CIS Benchmarks and Azure Defender for IoT – Microsoft Tech Community](#) (8/8/2021)

The 7 properties of highly secured devices

We suggest ensuring the hardware and operating system of both your and your suppliers’ devices are designed and implemented securely, have high barriers to compromise, and incorporate mechanisms and processes that continually monitor, alert, and restore security when necessary.

Through extensive research and testing, Microsoft identified the seven properties that are present in all standalone, internet-connected devices considered to be highly secured. In many cases, these highly secured devices apply additional security measures, but in all cases each of the seven properties is present. Collectively, these seven properties provide a baseline foundation of security throughout device silicon, software architecture and OS, cloud communications, and cloud services. The complexity of maintaining all seven properties could be a barrier for some organizations, despite the exceptional cost that often results from a fallout of incomplete device security.

 Hardware root of trust	Device identity and integrity are protected by hardware. Physical countermeasures resist side-channel attacks. Does the device have a unique, unforgeable identity that is inseparable from the hardware? Is the integrity of the device software secured by hardware?
 Defense in depth	Multiple mitigations applied against threats. Countermeasures mitigate the consequences of a successful attack on any one vector. Does the device remain secured even if one security mechanism is breached?
 Small trusted computing base	Private keys stored in a hardware-protected vault, inaccessible to software. Division of software into self-protecting layers. Is the device’s security enforcement code protected from bugs in other software on the device?
 Dynamic compartments	Hardware-enforced barriers between software components prevent a breach in one from propagating to others. Is a failure in one component of the device contained to that component? Can new compartments be added in field to address new security threats?
 Password-less authentication	Signed token, signed by an unforgeable cryptographic key, proves the device identity and authenticity. Does the device authenticate itself with certificates or other tokens signed by the hardware root of trust?
 Error reporting	A software error, such as a buffer overrun induced by an attacker probing security, is reported to cloud-based failure analysis system. Does the device report errors for analysis to enable verification of the correctness of in-field device execution and identification of new threats?
 Renewable security	Update brings the device forward to a secure state and revokes compromised assets for known vulnerabilities or security breaches. Is the device’s software updated automatically? Can the device’s security TCB software be updated rapidly without repackaging other device code?

Applying a Zero Trust approach to IoT solutions

Securing IoT solutions with a Zero Trust security model⁸¹ starts with non-IoT specific requirements—specifically ensuring you have implemented the basics to securing identities and their devices and limiting their access. These requirements include explicitly verifying users, having visibility into the devices they’re bringing on to the network, and being able to make dynamic access decisions by using real-time risk detections. Meeting these requirements helps to limit the potential blast radius of users gaining unauthorized access to IoT services and data in the cloud or on-premises. Otherwise, you could face both mass information disclosure (such as leaked production data of a factory) and potential elevation of privilege for command and control of cyber-physical systems (such as stopping a factory production line).

After basic security requirements are met, you can shift your focus to the specific Zero Trust requirements for IoT solutions:

Strong identity to authenticate devices

Register devices, issue renewable credentials, employ passwordless authentication, and use a hardware root of trust to ensure you can trust its identity before making decisions.

Least privilege access to mitigate blast radius

Implement device and workload access control to limit any potential blast radius from authenticated identities that may have been compromised or running unapproved workloads.

Device health to gate access or flag devices for remediation

Check security configuration, assess for vulnerabilities and insecure passwords, and monitor for active threats and anomalous behavioral alerts to build ongoing risk profiles.

Continual updates to keep devices healthy

Utilize a centralized configuration and compliance management solution and a robust update mechanism to ensure devices are up to date and in a healthy state.

Security monitoring and response to detect and respond to emerging threats

Employ proactive monitoring to rapidly identify unauthorized or compromised devices.

Learn more:

[Azure Defender for IoT | Microsoft Azure](#)

[Azure Sentinel – Cloud-native SIEM Solution | Microsoft Azure](#)

<https://aka.ms/7properties>

[Nineteen cybersecurity best practices used to implement the seven properties of highly secured devices in Azure Sphere \(microsoft.com\) \(July 2020\)](#)

[Zero Trust Cybersecurity for the Internet of Things \(4/30/2021\)](#)

IoT at the intersection of cybersecurity and sustainability

Organizations generally deploy IoT to improve the bottom line: better quality, higher productivity/less downtime, production optimization, and reducing operating costs and/or increasing output nonlinearly. Improving the business in these ways also reduces waste: less resource use for the same or greater output, higher uptime, reduction in scrap, and so on. While IoT investment specifically for sustainability is less common, sustainability as an initiative is no longer regarded as a zero-sum effort in conflict with business value.

As organizations face increasing pressure to improve their environmental footprint—from shareholder calls and new government regulations—

sustainability will become a primary driver for operational deployment of IoT. To make significant improvements to their environmental footprint, organizations must assess and monitor their behavior and then use automated or remote-control methods to optimize it.

The challenge is figuring out how to measure, monitor, and automate these systems securely. These systems might contain sensitive data, connect to business systems across your organization, and increasingly impact the physical operation of your enterprise. Attacks like the aforementioned Triton⁸² or Crash Override⁸³ that specifically target OT systems demonstrate that these systems are attractive targets for nation states and cybercriminals, and they have the potential to both disrupt business operations and potentially create environmental damage.

It is essential to assess the security of OT systems with the same rigor and comprehensiveness as IT systems. As we have observed, attackers will choose the “soft targets” as a point of ingress. Spear phishing or similar attacks allow access to IT systems

that can then provide a pathway for attackers to reach OT systems, and the reverse is also possible. In one example, attackers used an aquarium system to access a casino’s high-roller databases,⁸⁴ demonstrating that any device with connectivity can present a motivated attacker with an opening.

While many organizations are evolving their IT security approach (moving away from a perimeter-based security model to a Zero Trust model), IoT is often overlooked and lagging. For example, organizations know to encrypt sensitive data from applications, but many have not considered that their control systems rely on the Modbus protocol, which by design lacks any authentication and sends data in the open. While PCs are routinely required to have updated certificates, IoT devices are often deployed with factory default passwords.

Compromises in these OT systems may disrupt operations, but attackers are also focusing on how IoT and OT interact. Industrial control systems are often updated or retrofitted with remote capabilities, introducing new attack vectors that allow virtual attacks to cause harm in physical scenarios. Earlier this year, a water treatment plant in Florida fell victim to an attacker that remotely accessed critical systems and attempted to alter the amount of chemicals in the water supply.⁸⁵

Moreover, it is critical to understand the security posture of supply systems that are not on the organization’s IT/OT network, but nonetheless affect operations. Just as an organization looks to improve efficiency and sustainability, its suppliers do too. These supply systems may be connected outside of the network (such as cellular) to measure and monitor the device operation, reduce truck rolls, and deliver more uptime. Compromises in externally managed infrastructure components can directly impact downstream businesses. For example, turning off the chillers in a building could halt operations and spoil inventory, air quality sensors may not alert workers to unsafe conditions, and so on.

While IoT can and will enable better environmental practices, it is essential that all connected systems—which may be in place for a decade or longer—are designed, evaluated, and operated securely. As the world adapts to the new priorities that emerged during the pandemic, companies are looking to address the growing sustainability challenges ahead. Secured IoT will play a critical role in enabling businesses to both sustainably use and protect vital resources and utilities today and into the future.

It is essential to assess the security of OT systems with the same rigor and comprehensiveness as IT systems.

⁸² Hackers use Triton malware to shut down plant, industrial systems | ZDNet ⁸³ Crash Override Malware Took Down Ukraine’s Power Grid Last December | WIRED ⁸⁴ A smart fish tank left a casino vulnerable to hackers (cnn.com)

⁸⁵ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁸⁶ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁸⁷ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁸⁸ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁸⁹ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁰ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹¹ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹² Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹³ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁴ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁵ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁶ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁷ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁸ Smart fish tank left a casino vulnerable to hackers (cnn.com) ⁹⁹ Smart fish tank left a casino vulnerable to hackers (cnn.com) ¹⁰⁰ Smart fish tank left a casino vulnerable to hackers (cnn.com)

IoT adoption for sustainability

% PHASE OF AI STRATEGY	Those adopting IoT for Sustainability	Those adopting IoT, but not for Sustainability	Those adopting IoT for sustainability are more likely to be in the implementation phase of their AI strategy than those adopting IoT for other reasons. ⁸⁶
Implementing	38%	28%	
Developed strategy but not implemented	29%	25%	
Developing strategy	26%	26%	

IoT security policy considerations

Policymakers around the world are acknowledging the profound implications of IoT security for privacy, safety, critical infrastructure protection, and digital transformation in general. Approaches to IoT security policy range from voluntary programs to mandatory security requirements. The range of IoT device types, growing number of devices, and the volume of interactions between devices, the physical world, and the internet make developing effective

and appropriately tailored cybersecurity policy a complex task. To tackle the IoT threat landscape, the global community of IoT manufacturers and cybersecurity experts has developed sets of best practice standards for IoT device cybersecurity. These standards have demonstrated effectiveness against common attacks, and industry and policymakers alike can leverage them for immediate improvements to the global state of IoT device security for consumer, enterprise, and government users.

Minimum security baselines

Standards for minimum IoT security baselines provide a promising start to improving global

cybersecurity health. These standards are intended to provide guidance for manufacturers, developers, and users to identify and adopt best practices for device security. Prominent examples of international standards include the European Telecommunications Standards Institute (ETSI) standard for consumer IoT security released in June 2020, [ETSI EN 303 645](#). The ETSI standard is now a public set of resources that governments and companies around the world can use to enhance the security of IoT devices, and it includes both governance and technical recommendations. It was created through a rigorous and collaborative multistakeholder process involving experts from industry, government, and academia. Similarly, after an iterative public consultation process, in May 2020, NIST released NISTIR 8259A, which details a baseline set of device capabilities necessary for common cybersecurity controls. The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) are also developing a minimum-security baseline.

Policy can help manufacturers adopt international standards in a consistent way to improve security across a range of consumer products and promote an advanced state of security in critical applications. In the United States, policy initiatives based on standards include the IoT Cybersecurity Improvement Act of 2020,⁸⁷ which references NISTIR 8259A for managing risks associated with US federal

agency use of IoT devices, and EO 14028, which proposes a consumer IoT device labeling program that may also reference NISTIR 8259A as well as compatibility with ETSI or ISO standards. Other examples include the proposed mandatory security requirements for consumer smart devices in the UK,⁸⁸ and the voluntary device labeling schemes in Singapore and Finland, all based on ETSI EN 303 645. In leveraging international standards and widely used best practices, policymakers can also help to ensure that mandatory requirements are consistent and mutually recognized across regions to avoid fragmentation that would work against IoT innovation, interoperability, and security.

Global Cyber Alliance project: How policy and standards improve IoT security

Standards, laws, and proposed requirements for minimum IoT device security baselines share many commonly recommended or required controls. The first three provisions of the ETSI standard for consumer IoT security, ETSI EN 303 645, are the most highly recommended by ETSI and make appearances in several national-level policies. These include:

- 1. No default passwords
- 2. Implement a vulnerability disclosure policy
- 3. Keep software updated

⁸⁶ Microsoft IOT Signals report, to be published November 2021 ⁸⁷ Text of H.R. 1668 (116th): IoT Cybersecurity Improvement Act of 2020 (Passed Congress) ⁸⁸ [UK: Drafting new cyber security laws to protect smart devices amid pandemic sales surge – GOV.UK \(www.gov.uk\)](#)

<https://t.me/learningnets>

State laws in California and Oregon both prohibit default passwords for IoT devices. Similarly, recommendations or requirements for updated software and use of secure communications protocols like HTTPS also frequently appear in standards and policies around the globe.

Microsoft supported a research study conducted by the Global Cyber Alliance (GCA)⁸⁹ to demonstrate the effectiveness of commonly recommended controls in preventing attacks. The analysis can be

used to help policymakers and industry leaders understand the benefit of best practice approaches to IoT device security, and as proof points for manufacturers to adopt standards and comply with policy.

Using home-grown honeypot technology, GCA created an enticingly accessible target for would-be attackers to learn as much as possible about the existence, source, and prevalence of attacks. GCA operated a large honey farm of several hundred


emulated IoT devices distributed worldwide to provide real, at-scale, long-term attack data to understand trends and changes in IoT attack methodologies.

The research tested three controls commonly referenced in IoT security standards and policy:

- Secured access control (“no default passwords”)
- Device capability to update software and recommendation to keep software updated
- Data in transit is protected


SUMMARY OF GCA CONCLUSIONS
GCA's analysis of real attack data shows that default passwords factory-set by device manufacturers and never changed by users, along with weak passwords set by users, together represent the most exploited security vulnerability for IoT devices. Policy and regulatory frameworks can help drive adoption and harmonize implementation of the requirements in IoT device security standards such as NISTIR 8259, ETSI EN 303 645, and ISO/IEC 27402 to promote secured access control best practices and address this risk.

GCA honeyfarm attacker traffic by protocol



Roughly equal attacker traffic scanning for open ports through which to launch attacks.

GCA honeyfarm attacks with activity (commands and downloads) by protocol



Telnet is the clear protocol of choice for exploitation in attacks launched on IoT devices.

The research also showed widespread prevalence of attempts to exploit security vulnerabilities in the software stacks of IoT devices. This finding offers support for the effectiveness of keeping software updated and using secure communications protocols to prevent attempts to compromise security vulnerabilities.

The prevalence of attack attempts observed on security vulnerabilities in IoT device software also supports the notion that commonly recommended nontechnical controls would reduce the likelihood of attack success. In particular, requirements for manufacturers to implement a vulnerability disclosure and management policy and to disclose the security support status of their products would provide valuable information to consumers.


Learn more:

[Global Cyber Alliance: IoT Policy and Attack Report](#)

Microsoft data and threat signals support these findings

The GCA project findings align well to what Microsoft sees across its IoT sensor network. Generic HTTP scanning and scraping forms the bulk of requests we receive, followed by Secure Shell (SSH) and Telnet, respectively. Both protocols are frequently seen across IoT/OT devices, with Telnet in particular being a favorite of botnets like MIRAI and others based on it. The two protocols are also commonly associated with brute force password attacks, and when taken together become the most prevalent type of attack we see against IoT/OT devices.

Distribution of attacks against popular IoT/OT protocols (July 2020-June 2021)




When we look at attacks against these two ports, we see things are anything but static. Different actors come and go, in many cases repurposing bots with leaked source code such as MIRAI. Often these attacks replicate existing ways of working, but we have started to see more advanced use cases where bots are leveraging new exploits.⁹⁰ It is worth noting that many of these attacks continue to use poor passwords as a basis for lateral movement between infected hosts.


Microsoft’s sensor network gives us raw data on these types of attacks and the passwords in use. We looked at over 280,000 attacks and analyzed the password data we collected. Perhaps unsurprisingly we saw that 96% of attacks used a password with fewer than 10 characters, 92% had fewer than 8, and slightly more shockingly, 72% of all attacks required only trying a password of 6 characters or less. Within these password attempts, only 2% included a special character and 72% didn’t even contain a number.

The attacks we are seeing in the wild line up with the GCA data and other reports. Default passwords, which are generally a single short English word, make it trivial for an attacker to access an organization’s infrastructure. If secure identity management with an organization’s IoT devices is not an option, longer passwords—especially those with special characters—are strongly advised.

Attacks against Telnet and SSH ports



Passwords seen in 45 days of sensor signals



>20 Million

NUMBER OF TIMES
WE OBSERVED
THE PASSWORD
“ADMIN” USED IN
IOT DEVICES OVER
A 45 DAY PERIOD.

CHAPTER 5

Hybrid workforce security

Introduction

A Zero Trust approach for securing hybrid work

Identities

Devices/Endpoints

Applications

Network

Infrastructure

Data

People

INTRODUCTION: The basics matter

BRET ARSENAULT, CHIEF INFORMATION SECURITY OFFICER

This past year continued to challenge us in profound ways. While most industries made the shift to remote work due to the pandemic, it created new attack surfaces for cybercriminals to take advantage of, such as home devices being used for business purposes. In the first half of 2021, there were three significant assaults: NOBELIUM (the SolarWinds supply chain attack), HAFNIUM (an on-premises Exchange server attack), and Colonial Pipeline (a ransomware attack).

Many lessons can be learned. First, a continuing threat vector is email compromise. In fact, phishing is responsible for almost 70% of data breaches.⁹¹ Second, cybercriminals are using malware that is posed as a legitimate software update to target unsuspecting employees. Third, ransomware attackers have raised the stakes to focus not only on double or triple extortion tactics in terms of a payout but are also offering ransomware as a service (RaaS), which uses a partner network to carry out an attack, making it tough to determine who the real bad actor is. Finally, adversaries are targeting on-premises systems, reinforcing the need for organizations to move infrastructure to the cloud where security is more difficult to penetrate.⁹²

While these incidents taught us tough lessons, a key takeaway is that *the basics matter*. A primary way criminals get in is through an unlocked door.

If compromised organizations had applied basic security hygiene like patching, applying updates, or turning on multifactor authentication (MFA), they may have been spared or less impacted. In fact, it is shocking that less than 20% of our customers are using strong authentication such as MFA⁹³ (which is free to our customers and can be turned on by default). Organizations that do not apply or maintain these basic hygiene practices will face much greater exposure to attacks.

Along with the security basics, Microsoft relies on an approach called Zero Trust,⁹⁴ which assumes the network has been breached. Zero Trust means we don't assume any identity or device on our network is secure—we continually verify it. Zero Trust helps us strike a balance in making sure employees can be productive, secure, and healthy beyond the corporate network from home, the office, or anywhere in-between.

Recommendations for getting started with Zero Trust:

- **Identities are validated and secured with MFA everywhere.** Using MFA helps eliminate the need for passwords. The added use of biometrics (such as retina eye scans or fingerprints) also ensures strong authentication of a user's identity.
- **Devices are managed and validated as healthy.** As a condition of access to any company resource, all device types and operating systems should be required to meet a minimum healthy device state before being validated.
- **Monitoring and threat signals are pervasive.** Make sure to collect all available logs, data, and signals to understand the current security state, which will help you identify gaps in coverage, discover anomalies, and drive a better employee experience.

⁹¹Email scams in particular are surging, according to the cyber defense firm Barracuda. Phishing was responsible for almost 70% of data breaches. <https://blog.barracuda.com/2020/03/26/threat-spotlight-coronavirus-related-phishing/>

⁹²Through 2024, workloads that leverage the programmability of cloud infrastructure to improve security protection will demonstrate improved compliance and at least 60% fewer security incidents than those in traditional data centers. (How to

Make Your Own Data Center, Neil MacDonald, Tom Croll (April 29, 2021)) ⁹³Based on Azure Active Directory protection telemetry as of August 2021. ⁹⁴[Zero Trust Security Model and Framework | Microsoft Security](#)

ORGANIZATIONS
THAT DO NOT
APPLY OR
MAINTAIN
THESE BASIC
HYGIENE PRACTICES
WILL FACE MUCH
GREATER EXPOSURE
TO ATTACKS.

Our goal with this report is to share insights from our own internal teams who are doing this work each day, to help educate and empower others to improve their cyber defense techniques to protect their employees, their companies, and our online ecosystem.

And remember, you can't secure future work if you don't secure your past work.

Moving toward a hybrid workforce at Microsoft


No one knows how or when the COVID-19 pandemic will end, but at some point, COVID-19 will no longer place a significant burden on our communities and will present itself more like an endemic virus such as influenza. Local

health situations and local government guidance determines how many employees we allow at a worksite and what services are made available. Once a location fully reopens, we return to somewhat normal operations with services that were provided pre-pandemic. As Microsoft offices globally invite more employees back to the worksite, we are gradually seeing an increase in the number of employees badging-in (scanning their badges for building entry) each week. We expect this number will continue to rise as more locations fully reopen. However, we do not expect these numbers to reach pre-pandemic levels of attendance. Our flexible work approach means that we expect to see most employees spending at least some time each week working remotely, even in locations where COVID-19 is no longer a significant burden.


Learn more:

- [Securing a new world of hybrid work: What to know and what to do - Microsoft Security \(5/12/2021\)](#)
- [Work Trend Index: Microsoft's latest research on the ways we work.](#)
- [The Next Great Disruption Is Hybrid Work—Are We Ready? \(microsoft.com\)](#)
- [Securing Microsoft's network with an internet-first, Zero Trust model \(4/16/2021\)](#)

Global pre-COVID onsite work and the rapid move to remote work, followed by gradual return



Global weekly unique badge scans (January – August 2021)



A Zero Trust approach for securing hybrid work

The increasing prevalence of cloud-based services, mobile computing, Internet of Things (IoT), and “bring your own device” (BYOD) in hybrid work environments has changed the technology landscape for today’s enterprise. Security architectures that rely on network firewalls and virtual private networks (VPNs) to isolate and restrict access to corporate technology resources and services are no longer sufficient for a workforce that regularly requires access to applications and resources that exist beyond traditional corporate network boundaries. The shift to the internet as the network of choice and the continuously evolving threats led Microsoft to adopt a Zero Trust security model. Zero Trust has become a priority of enterprise security leaders around the world.

We are facing a moment of reckoning as the world witnesses a rise in increasingly sophisticated and expansive cybersecurity attacks. This reality—coupled with work entering its next great disruption, the move to hybrid environments—has ushered in

an urgent opportunity for all companies around the world to adopt a Zero Trust approach and assume all activity, even by trusted users, could be an attempted breach. Signals across the industry highlight that every company needs to create a culture of security and modernize their approach to ensure they are protected.

Zero Trust principles

Zero Trust eliminates the inherent trust that is assumed inside the traditional corporate network. An effective Zero Trust architecture is designed to reduce risk at every opportunity across the digital estate. In practice, this means that every transaction between systems must be validated and proven trustworthy before the transaction can occur. For an effective Zero Trust strategy, we recommend these guiding principles:

Verify explicitly

Always authenticate and authorize based on all available data points, including user identity, location, device health, service or workload, data classification, and anomalies.

Use least privilege access

Limit user access with just-in-time and just-enough-access (JIT/JEA), risk-based adaptive policies and data protection to help secure both data and productivity.


Assume breach

Minimize blast radius and segment access. Verify end-to-end encryption and use analytics to get visibility, manage insider risk, drive threat detection, and improve defenses.

An integrated security philosophy and end-to-end strategy

Zero Trust controls and technologies are deployed across six foundational technology pillars. Each pillar is a source of signal, a control plane for enforcement, and a critical resource to be defended. In a Zero Trust architecture, they are interconnected by automated enforcement of security policy, correlation of signal and security automation, and orchestration.

Zero Trust across the digital estate



1. Identities

Identities can represent people, services, or IoT devices. When an identity attempts to access a resource, verify that identity with strong authentication, and ensure access is compliant and typical for that identity. Follow least privilege access principles.

2. Endpoints

Once an identity has been granted access to a resource, data can flow to a variety of different endpoints—from IoT devices to smartphones, BYOD to partner-managed devices, and on-premises workloads to cloud-hosted servers. This diversity creates a massive attack surface area. Monitor and enforce device health and compliance for secure access.

3. Applications

Applications and application programming interfaces (APIs) provide the interface by which data is consumed. They may be legacy on-premises, workloads moved to the cloud, or modern software as a service (SaaS) applications. Apply controls and technologies to discover shadow or unsanctioned IT, ensure appropriate in-app permissions, gate access based on real-time analytics, monitor for abnormal behavior, control user actions, and validate secure configuration options.

4. Network

All data is ultimately accessed over network infrastructure. Networking controls can provide critical controls to enhance visibility and help prevent attackers from moving laterally across the network. Segment networks (and do deeper in-network micro-segmentation) and deploy real-time threat protection, end-to-end encryption, monitoring, and analytics.

5. Infrastructure

Infrastructure—whether on-premises servers, cloud-based virtual machines (VMs), containers, or micro-services—represents a critical threat vector. Assess for version, configuration, and JIT access to harden defense. Use logging and monitoring to detect attacks and anomalies, and automatically block and flag risky behavior and take protective actions.

6. Data

Ultimately, security teams are protecting data. Data should remain protected throughout its lifecycle even if it leaves the devices, apps, infrastructure, and networks the organization controls. Use data classification and labeling as context to encrypt, minimize access to, control the flow of, and mask or delete sensitive information at the end of its useful or legally mandated life.

Learn more:


[The critical role of Zero Trust in securing our world | Microsoft Security Blog](#) (6/30/2021)

Zero Trust adoption

The Zero Trust Adoption Report⁹⁵ illuminates the path of Zero Trust adoption across diverse markets and industries. We hope that the learning gained by this research can help accelerate your own Zero Trust strategy adoption, shed light on the collective progress of your peers, and provide insights on the future state of this rapidly evolving space.


No single security risk area stands out as a primary starting point for Zero Trust strategy, as fewer than 15% start with the same security risk area. Organizations are beginning in different places, likely based on their needs and available internal resources. Most organizations approach Zero Trust as an end-to-end strategy to be completed over time. Eventually, they seek to adopt this strategy across all security risk areas to ensure even more protection against threats.

Hybrid workplace intent



The shift to a hybrid workplace is driving broader adoption of Zero Trust strategy. 81% of enterprise organizations have begun the move toward a hybrid workplace, with 31% already fully adopted.

Current zero trust implementation – security risk areas



There is no one-size-fits-all approach to Zero Trust implementation, giving permission to start anywhere.

Learn more:

- [Zero Trust Adoption Report \(7/27/2021\)](#)
- [Resources for accelerating your Zero Trust journey - Microsoft Security \(5/24/2021\)](#)
- [Zero Trust Security Model and Framework | Microsoft Security](#)

Identity is more important than ever.

Identities

In Azure Active Directory we observe 50 million password attacks daily, yet only 20% of users and 30% of global admins are using strong authentications such as MFA.⁹⁶ Password-based attacks remain the main source of Identity compromise. However, other types of attacks are emerging, including consent phishing and attacks on nonhuman identities.


Password-based attacks

Azure AD is the front door to all Microsoft cloud services. Our sign-in service sees 90 billion authentication requests per day, which gives us great visibility on identity attacks happening across all our clouds. Password-based attacks on user identities are still the most prevalent vector of identity compromise. While password spray and credential stuffing used to be the largest vectors of identity compromise, in the last year, we observed a significant change of tactics by the bad-actor ecosystem. In the last few years, but most notably during the pandemic, Microsoft and our customers have invested significantly in reducing the attack surface for bad actors to compromise Azure AD-backed accounts. The enhancement of security posture (security defaults, MFA adoption, password protection, legacy authentication block),

as well as our investments in detection (malicious IP address, password spray detection), have caused a reduction in password brute-force attacks that use legacy authentication protocols, such as IMAP or SMTP. The volume of attacks that we attribute to breach replay and password spray have decreased significantly, while at the same time we are seeing a large increase in attacks attributed to phishing or other more sophisticated techniques, such as credential harvesting with malware. The data shows that bad actors have adapted their tactics to keep compromising accounts.

Although blocking legacy authentication and enabling MFA are still the most important defenses for any organization, phishing protection is becoming more relevant than ever. Even with MFA enabled, users can still have their credentials phished by real-time man-in-the-middle phishing tools that replicate the sign-in page and replay the MFA prompt to collect the one-time password sent to the user. To protect from this type of attack, customers can adopt Fido2-based credentials (based on a public key cryptographic key pair), such as security keys or Windows Hello for Business.⁹⁷

Monthly compromised users by attack category (June 2020 – June 2021)



Rise in phishing emails using OAuth request URLs



Geo-distribution of IP addresses issuing password brute-force attacks in July 2021



Emerging trends in attacks

OAuth consent phishing


In typical phishing, an attacker looking to steal credentials will craft a convincing email, host a fake landing page, and expect the user to fall for the lure. On a successful phishing attempt, the user credentials are passed on to the attacker. Consent phishing is a bit different. This method attempts to trick users into granting permissions to a malicious attacker-owned application and uses the obtained access tokens to retrieve the users’ account data. This is a very sophisticated attack as

the access tokens do not require knowledge of a user’s password, and the user’s password is never shared with the attacker. Most importantly, as this is not a credential based attack, strong authentication requirements such as MFA do not prevent attacks that use this technique.

In the last six months, the monthly average of phishing emails using OAuth URLs has almost doubled. These phishing emails target a variety of legitimate cloud services such as Microsoft, Google, and Facebook. We are seeing an upward trend in the number of unique emails with OAuth phishing links.


Microsoft recommends limiting user consent to allow user consent only to apps from verified publishers.

Microsoft has deployed several defenses to counter this trend, including specific machine-learning-based detections to identify, isolate, and disable malicious applications. Strategies that can help prevent such attacks include granular user consent policies, blocking consent-phishing emails, anomaly detection, and user awareness training.



As a result of these countermeasures, there was an 89% increase in disabled apps from January to June 2021 compared to July to December 2020. Upon detailed investigation, we saw that the major malicious vectors have shifted from a multi-tenant phishing attack to an abusive type of attack.

Hackers don't break in,
they log in.



For more information on phishing trends and mitigation, see the [State of cybercrime](#) chapter of this report.

Learn more:

[Microsoft delivers comprehensive solution to battle rise in consent phishing emails | Microsoft Security Blog \(7/14/2021\)](#)

Attacks on nonhuman accounts

We saw a significant uptick in the volume of attacks against applications and service principal identities.⁹⁸ Unlike users, these identities are often more vulnerable because organizations do not consistently apply many of the typical safeguards such as strong authentication, rigorous lifecycle management, and security monitoring to this category of accounts. As a result, we have seen attackers exploiting application identities that already have a privileged role or are scoped to a wide set of permissions. Unlike the typical “Initial Access -> Privilege Escalation” attack chain, this attack vector gains initial access through a compromised user or a leaked application credential.

We have seen attackers who gain initial access through a compromised user use their elevated privilege to conduct reconnaissance to identify applications with existing permissions, add a new set of application credentials if needed, and even assign privileged roles to the application. As a result, the application behaves normally in the environment while still operating as a tool for the attacker. Customers should conduct a review of application roles and permissions to conform to the principle of least privilege and monitor their Azure AD Audit logs for unfamiliar activity on applications and service principals.

The challenge of credentials in code is not a new one. While most major vendors including Azure DevOps and GitHub offer safeguards to flag such credentials for removal, the practice persists. Those credentials provide attackers a direct path into an organization’s environment while flying under the radar in the Azure AD Audit logs. [Microsoft recommends that customers audit their engineering artifacts, including code repositories for credentials.](#)


Adoption of security posture

Strong authentication adoption

For identities, verifying explicitly means ensuring the identities are using strong authentication when accessing resources. Microsoft research shows that requiring strong authentication can protect against 99.9% of the identity attacks because the majority of the attacks are related to passwords. While augmenting passwords can help defend against those attacks, [eliminating passwords altogether with passwordless authentication methods can provide the most usable and secure authentication experience.](#)


As companies have been adapting their security posture for remote work and a hybrid workforce, we have seen over 220% increase in strong authentication usage in the last 18 months.

Strong authentication usage in Azure AD



Microsoft enables security defaults for all new customers to ensure that everyone has at least a basic level of protection, including controls like MFA for administrators.⁹⁹

Tenants with security defaults enabled




Growth in Zero Trust access policy usage

Azure AD Conditional Access is a security policy engine for verifying explicitly and granting least privilege access. In the last year, the number of conditional access policies deployed has more than doubled, as organizations revamped their security postures to account for a remote workforce.

When a user accesses an application, administrators can use conditional access to configure which additional requirements are enforced to grant access. As organizations embrace Zero Trust security principles, we have seen greater adoption in managed device and managed app requirements in addition to MFA.


Most frequently required access controls in Azure AD (July 2020 – June 2021)



Legacy protocols: A preferred entry point for adversaries

One of the leading sources of compromise in organizations is authentication from legacy protocols, such as IMAP, SMTP, POP, and MAPI. These protocols do not support MFA, so they are preferred entry points for adversaries. In fact, 99% of password spray and 97% of credential stuffing attacks use legacy authentication, according to authentication data from Azure AD. Data from Azure AD reveals that the compromise rate for authentications using IMAP clients is 22 times higher than the compromise rate for authentications from a browser.


Account compromise rate by authentication protocol (July 2020 – June 2021)



Fortunately, end users and admins have begun making progress toward using modern authentication in their organizations.

The hybrid world is largely “perimeterless,” so wrapping protections around identity and devices is critical. As part of Zero Trust, we also think the future is passwordless,¹⁰⁰ and we will start to see that transition this year.

Percentage of legacy authentication among monthly active users of Azure Active Directory




Over the past year, we have seen the percent of users with legacy authentication clients decrease over 30% to nearly 10% of users in Azure AD.

Devices/Endpoints

What we’re seeing

Mostly driven by necessity as the world shifted to a remote or hybrid work model, users are working from anywhere, from any device, more than any time in history, and attackers are quickly adjusting their tactics to take advantage of this change. Enterprises are left weighing the benefits of enabling BYOD (allowing their end users to access corporate resources that traditionally required VPN or on-premises access) against the increased risk of the same users unintentionally installing ransomware or other malware while performing non-work-related functions on their personal devices. **By enabling BYOD using a Zero Trust model, enterprises can reduce provisioning costs and avoid additional hardware purchases for work-from-home use, but they need to be able to protect their corporate assets on these devices, while still allowing the users to perform non-work functions on these same devices.**

Mobile device management (Intune data)



Sharp rise in mobile device management growth as workers moved from office to BYOD and home PCs

Recommendations for mitigating BYOD risk

To mitigate the increased risks of any BYOD model, it’s important to ensure constant verification of specified security standards as well as validation of the identity of the device and user to gate control of critical company resources. For example, you can block access to a personal device that has been jailbroken (modified to remove restrictions imposed by the manufacturer or operator) to ensure that enterprise applications are not exposed to known vulnerabilities.

For more information on securing devices, see the [Supply chain, IoT, and OT security](#) chapter of this report.

Learn more:

- [Protect Data and Devices with Intune | Microsoft Docs](#)
- [Device compliance policies in Microsoft Intune - Azure | Microsoft Docs \(4/29/2021\)](#)
- [What is Conditional Access in Azure Active Directory? | Microsoft Docs \(1/27/2021\)](#)

Applications

Moving from legacy to Zero Trust-ready applications

As we move from a security model centered on the corporate network to one based on identities, thousands of apps and services with internally facing posture remain heavily reliant on network firewalls and VPNs to isolate and restrict access. They were built around legacy authentication mechanisms, keeping them grounded within corporate networks. These traditional architectures built for legacy apps were designed for lateral connectivity rather than micro segmentation. They violate the fundamental principle of least privilege access and are more vulnerable to lateral movement across the network by an adversary, which in turn could expose

confidential data. For more on recent attack trends in this area, please see the [Legacy protocols](#) and [OAuth consent phishing](#) sections of this report.

Using modern apps and data solutions

In a cloud-centric architecture, we treat our applications and data differently. With more user-design models becoming available, engineers no longer function as the only developers in an organization. Users are taking advantage of platforms and tools that offer no-code or low-code development methods to create business solutions. Organizations should invest in creating the right guardrails for these new paradigms. Tracking cloud resources and applying correct policies and templates help to ensure that modern solutions immediately use the correct controls.

Learn more:

[Build 2021: Build Zero Trust-ready apps with the Microsoft identity platform - Microsoft Tech Community \(5/26/2021\)](#)

[Integrated Threat Protection | Microsoft Security](#)

[Safeguard your multicloud resources | Microsoft Cloud Security](#)

How do you modernize applications? Successfully deploy one of the three solutions listed here:

Move to a serverless, cloud native solution using SaaS- or PaaS-based services.

Segment the application extending the front end to a cloud access point, leveraging modern identities, and tiered access to internal components obfuscated from the user.

Add an internet hybrid endpoint such as a reverse proxy or other secure access service edge (SASE) solution.

MODERNIZED
APPS AND
SERVICES
REQUIRE
USERS TO BE
AUTHENTICATED
PRIOR TO
HAVING
ACCESS.

Network

A Zero Trust approach encourages organizations to assume they are always under attack and that a security incident can happen at any time. To be prepared with a setup that minimizes the blast radius of such an incident, networks should be segmented when the layout is designed. Implementing these software-defined perimeters with increasingly granular controls will increase the “cost” to attackers to propagate through the network and thereby dramatically reduce the lateral movement of threats.

What we’re seeing in Azure Firewall signals

As migration to cloud accelerates, we are seeing that most new deployments are hybrid with connectivity back to on-premises and configured for internet breakout (micro perimeter) from Azure. All data is ultimately accessed over network infrastructure.

2 Trillion

Flows blocked in the past year, including malicious flows detected by threat intelligence engines, and unwanted traffic blocked by firewall rules

Customers are using network segmentation for organizing workloads and deploying firewalls for securing traffic across subnets, VNETs, to on-premises, and to the internet. Customers are using firewalls to control and enhance visibility into all the network flows.

Web application firewall

Web application firewalls (WAFs) have evolved from an initial focus on injection-type attacks, such as SQL injection and cross-site scripting, to include attacks from malicious bots and API abuses. When combined with distributed denial of service (DDoS) protection, WAF forms an integral part of defense-in-depth strategy for protecting web and API assets.

WAF over the past year has upwards of 25 billion rules triggered on a per-week basis. Approximately 4% to 5% of incoming traffic on average is deemed malicious and is blocked either at the network edge or at the regional data center depending on where WAF is configured.

25 Billion

WAF rules triggered per week over the past year

~4-5%


Incoming traffic deemed malicious

Most common network attack types


Azure Firewall blocks millions of attempted exploits daily. Our signals show that attackers most commonly used malware, phishing, web applications, and mobile malware in their attempts at network attacks during July 2021. Also in July, there was a significant uptick in the use of coinminers, a type of malware that uses the network to mine cryptocurrency.

The protocols leveraged most often in attacks were HTTP, TCP, and DNS, as they are open to the internet.

Top 10 network threats (July 2021)



Top 5 most common protocols used in network exploit attempts (July 2021)



DDoS attacks

In our 2020 DDoS retrospective,¹⁰¹ we highlighted shifts in the very active cyberthreat landscape. With the COVID-19 pandemic and an increase in internet traffic, DDoS attacks on internet-facing endpoints ramped up significantly. We continue to see similar trends in the first half of 2021 as well. The trend is characterized by large TCP attacks, mainly SYN, SYN ACK, and ACK floods and user datagram protocol (UDP) reflection attacks. Online gaming and the gaming vertical continues to be a very attractive target of DDoS attacks. The majority of attacks on the gaming vertical have been mutations of MIRAI botnets and low-volume UDP protocol attacks.

There is also an uptick in DDoS attacks against IPv6 this year. This is an indication that as more IPv6 is adopted in the enterprise networks, the risk of DDoS attacks on IPv6 will continue to increase.

Daily attacks and volume

Compared to the latter part of 2020, the average daily number of attack mitigations in the first half of 2021 increased by 25% while the average attack bandwidth per-public IP increased by 30%. Azure DDoS Protection mitigated 1,200 to 1,400 unique DDoS attacks every day in the first half of 2021. As of July 2021, the average attack size in 2021 (325 Gbps) was 25% larger than in 2020 (250 Gbps).

Attack duration

Microsoft’s DDoS Protection team continues to see that most attacks are of short duration, with 75.35% being 30 minutes or less and 87.60% being one hour or less. This trend is similar to what we observed in 2020.

Attack vectors, applications, and regions involved


More than 35% of the attack volume targeted HTTPS and 10% targeted HTTP. UDP attacks represented 43% of the overall attack vectors, with amplification attacks accounting for 11% of attacks. Besides DNS and NTP reflection attacks, there has been a surge in remote desktop protocol (RDP) and datagram transport layer security (DTLS) reflection attacks.

DDoS attack mitigations




Number of unique DDoS attacks mitigated daily by Microsoft

Attack Duration (July 2020 – June 2021)



Over 96% of the attacks are of short duration (<4 hours)

DDoS attack type (July 2020 – June 2021)



The average packet size of TCP flood has been much smaller compared to UDP attack packets. UDP Fragment attacks averaged 1145 bytes per packet while TCP invalid syn averaged 512 bytes. TCP attack vectors such as SYN ACK, TCP Zero Seq, FIN-ACK and RST Floods averaged below 100 bytes per packet.

Europe, Asia, and the United States remain the most attacked regions because of the concentration of financial services and gaming industries in these regions. The UAE has been increasingly hit by DDoS attacks on government, oil and gas, telecom, and healthcare industries/sectors.


Top source regions from where the DDoS attacks originated were from Russia, Romania, Turkey, Indonesia, Vietnam, and the Philippines, owing to the prevalence of DDoS attack-for-hire services in those regions.

Learn more:

[DDoS Protection and Mitigation Services | Microsoft Azure](#)

[Azure DDoS Protection Standard documentation | Microsoft Docs](#)

Top 10 DDoS attack destination regions



Cybercriminal DDoS services

Between November 2020 and May 2021, the average price of a DDoS attack increased over 500%. The reasons for the price increase are multifaceted:

The rising cost of goods sold for DDoS services has made them more expensive and resulted in a decrease in supply relative to demand.

- The price for loads that are turned into DDoS bots increased as demand for loads for all monetization strategies increased beyond supply. (Loads are freshly infected devices which are sold to third parties, who install their malware on the device.) Attackers have found more effective strategies to monetize loads that are delivering more value than they would get using them for DDoS attacks. These include

- buying loads to conduct ransomware, spyware, adware, and banking theft monetization strategies.
- Some antivirus products use outbound DDoS attacks to detect malware. Therefore, the lifecycle of the bots used in DDoS attacks is shorter than it used to be, requiring the DDoS service to continuously buy new loads to maintain their capabilities.
 - Improved DDoS mitigation strategies make it more complex to perform quality attacks.

It has become more common for actors that want to conduct DDoS attacks to rely on these DDoS services to conduct them. As a result, many attackers no longer have the skills or infrastructure to conduct their own DDoS attacks. In other words, costs are rising due to lack of alternatives.

Starting in late June, the advertised cost of DDoS services decreased back toward the equilibrium that was observed until November 2020. This shift was likely caused by two factors. First, there was a recent increase in the supply of fresh loads available for purchase from the DDoS services. Second, there was a decrease in ads for premium DDoS services. The margins for DDoS services were being squeezed by increased costs and lack of demand for high-priced DDoS services, leading some of the premium DDoS services to move from a “publicly available service” in the cybercriminal forums to more private services that don’t advertise.


Learn more:

[Sharing how Microsoft now secures its network with a Zero Trust model - Inside Track Blog \(4/22/2021\)](#)

[Zero Trust networking: Sharing lessons for leaders \(microsoft.com\) \(1/5/2021\)](#)

[Lessons learned in engineering Zero Trust networking \(microsoft.com\) \(1/5/2021\)](#)

Average price of DDoS attack services



Infrastructure

Infrastructure represents a critical threat vector. IT and applications infrastructure, whether on-premises, in the cloud, or multi-cloud, is defined as all the hardware (physical, virtual, containerized), software (open source, first- and third-party, platform as a service (PaaS), SaaS, functions, and APIs), micro-services, networking infrastructure, and facilities that are required to develop, test, deliver, monitor, control, or support IT services and applications. It is an area where Microsoft has invested tremendous resources to develop a comprehensive set of capabilities to secure future cloud and on-premises infrastructure.

Collaboration with MITRE on an ATT&CK-style matrix

The flexibility and scalability of containers encourage many developers to move their workloads to Kubernetes, an open-source system for automating deployment, scaling, and managing containerized applications. While Kubernetes has many advantages, it also brings new security challenges

that should be considered. Therefore, it is crucial to understand the various security risks that exist in containerized environments, and specifically in Kubernetes.

The MITRE ATT&CK framework is a knowledge base of known tactics and techniques that are involved in cyberattacks. Started with coverage for Windows and Linux, the matrices of MITRE ATT&CK cover the various stages that are involved in cyberattacks (tactics) and elaborate the known methods in each one of them (techniques). Those matrices help organizations understand the attack surface in their environments and make sure they have adequate detections and mitigations to the various risks.

When the Azure Security Center team began mapping the security landscape of Kubernetes, they noticed that although the attack techniques are different than those targeting Linux or Windows, the tactics are similar. For example, a translation from OS to container clusters would be “initial access to the computer,” which becomes “initial access to the cluster.” So, the team created the first Kubernetes attack matrix: an ATT&CK-like matrix comprising the major techniques that are relevant to container orchestration security.

Learn more:

[The evolution of a matrix: How ATT&CK for Containers was built | Microsoft Security Blog \(7/21/2021\)](#)

[Threat matrix for Kubernetes \(microsoft.com\) \(5/10/2021\)](#)

[Secure containerized environments with updated threat matrix for Kubernetes | Microsoft Security Blog \(3/23/2021\)](#)

[Crypto-Mining Attacks Targeting Kubernetes Clusters via Kubeflow Instances \(6/9/2021\)](#)

[Update: Help Shape ATT&CK for Containers \(2/18/2021\)](#)

Threat matrix for cloud storage

As the move to the cloud enables a more secure hybrid workforce, organizations are also increasing their dependency on cloud storage services. They require effective threat protection, mitigation strategies, and tools in place to manage access to their cloud storage. For example, Azure Defender treats data-centric services, such as cloud storage accounts and big data analytics platforms, as part of the security perimeter and provides prioritization and mitigation of threats for data storage. As Microsoft cloud security researchers examined the attack surface of storage services, they noted potential risks to be aware of when deploying, configuring, or monitoring a storage workload. We’ve produced a threat matrix for storage¹⁰² to help organizations identify gaps in their defenses. We expect the matrix to dynamically evolve as more threats are discovered and exploited, and techniques can also be deprecated as cloud infrastructures constantly progress toward securing their services.

Learn more:

[Safeguard your multicloud resources | Microsoft Cloud Security](#)

[Threat matrix for storage services - Microsoft Security \(4/8/2021\)](#)

Infrastructure represents a critical threat vector.

Data

Digital transformation and tech intensity¹⁰³ have led to exponential data growth, and organizations are dealing with never-before-seen volume, velocity, and variety of data. This growth, combined with increasingly remote workforces and cloud migrations, has created a complex data sprawl often pitting security and regulatory requirements against the access business needs to deliver value from that data.

Data governance is an integral part of data security

Organizations that effectively manage the lifecycle and flow of their sensitive data as part of their business operations make it that much easier for

data security and compliance teams to reduce exposure and manage risk.

As valuable as data is to organizations, data value can drop faster over time than the risk associated with it.¹⁰⁴ Accumulating sensitive data through management neglect is not only wasteful from a storage perspective but it can also be a recipe for accumulating risk. In a rapid data growth environment, it is not difficult to envision that data governance will increasingly have an impact on data security.

There is a lot to be learned and reused from recent privacy initiatives. Successful organizations with mature privacy processes have inherently had to combine data governance and security to minimize

their privacy data footprint and surface of attack to achieve sustainable compliance. These security and compliance benefits can also be achieved by using similar approaches with other sensitive or regulated datasets.

Reducing the risk associated with data is not only about pruning old data and securing it where it resides now. It is also about reevaluating how the organization conducts business with sensitive data going forward to ensure proper storage, access, flow, and lifecycle so that this sensitive data does not persist or propagate uncontrollably.


Over time, this means reengineering high-risk business processes and adhering to data governance

and security principles as organization evolve their applications and infrastructures or take on new business ventures.

Learn more

- [Azure Purview for Unified Data Governance | Microsoft Azure](#)
- [Information Protection and Governance | Microsoft 365](#)
- [Microsoft Information Protection in Microsoft 365 - Microsoft 365 Compliance | Microsoft Docs \(8/26/2021\)](#)
- [Microsoft Information Protection and Microsoft Azure Purview: Better Together - Microsoft Tech Community \(12/7/2020\)](#)


Information rights management use (July 2020 – June 2021)



Dip in the chart correlates to decreased usage during major holidays

¹⁰³ [How technology intensity accelerates business value – Microsoft Industry Blogs](#) ¹⁰⁴ [Implement Your Data and Analytics](#)
Microsoft Industry Blogs. Published 6 July 2020 – ID G00729295 – By Guido De Simoni, Saul Judah

Cumulative impact of unified data governance and security on sensitive data risk



Data governance can maximize the business value of data while helping minimize the security and compliance risk of that data.

People

Helping people adapt to new ways of working is key to any successful transformation. While organizations are empowering people to work securely when, where, and how they want, we have found the most successful are the ones who are also empathetic to the end-user experience.

The previous sections of this chapter have presented a view of the threat landscape across the six foundational pillars of Zero Trust, with key steps we recommend taking to protect them. However, it is imperative to remember that every step we take to implement Zero Trust strategies impacts the people within the organization—and that successful implementation depends as much on them as it does on the systems and tools we put in place. How companies engage with their workforces around remote work and security matters. We conclude the chapter with a discussion about people, the human element of any enterprise, and ultimately our greatest asset.

Some guidance about insider risk in the hybrid workplace

As described in this report’s sections on Zero Trust, it is possible to give employees seamless information access while mitigating the risk of inadvertent leaks. Regarding more malicious insider threats, using a framework of common factors and patterns typically seen helps to enable proactive detection. Microsoft’s Insider Risk program has adapted numerous

preventative and detective controls that decrease risk through all stages of the insider threat attack path. In general, we should assume that the threat of loss to the organization or its stakeholders increases as an attack progresses down the path. Preventative controls such as awareness trainings that instruct the organization where to report insider threat concerns, or controls designed to limit risky behavior (such as limiting the ability for those leaving the organization

to share files) are both initiatives that Microsoft has undertaken to mitigate risk. As organizations work toward increasingly advanced capabilities in the preventative control space, appropriate stakeholders from across the organization should always be consulted. For example, before implementing any preventative controls, verify with the business that protections still enable employees and contractors to perform their legitimate business activities.

Insider threat attack path¹⁰⁵



Insider threat attacks share several common factors and patterns, which can be described in a critical-path approach. Identifying indicators across phases of the critical path can help to enable more proactive detection.

The Insider Risk program leverages a variety of artificial intelligence (AI) and machine learning (ML) and signature- or rule-based detections, including those from our Insider Risk Management solution, to accomplish detection goals. While AI and ML are incredible tools that help to decrease some of the noise associated with traditional rule-based alerts, it is critical that the organization has the appropriate people with the “tribal knowledge” of business-acceptable behavior in the correct roles to create trustworthy AI and ML models. This human input takes standard supervised or unsupervised ML a step further in predicting which alerts are the most actionable for incident response teams. Lack of this knowledge can create a pool of noisy alerts with limited investigative value.

Key considerations for insider threat detections

AI and ML are not always the answer	Rules and signature-based detections are better suited for some use cases
Data is key	An effective AI or ML solution requires understanding and acquiring the data to develop, train, and enrich the solution, as well as track performance metrics
You still need people	Human review of alerts and identification of tuning opportunities is critical – continue to “shift left”
Cybersecurity talent pool is evolving	While security experts are still important, data scientists also bring valuable skills to the table

AI and ML are not always the answer

The empathy imperative

Flexible work is here to stay and with that comes several challenges and stressors. Teams have become more siloed this year, and digital exhaustion is a real and unsustainable threat. One in five global survey respondents say their employer doesn't care about their work-life balance.¹⁰⁶ Fifty-four percent feel overworked. Thirty-nine percent feel exhausted. And trillions of productivity signals from Microsoft 365 quantify the precise digital exhaustion workers are feeling.

A positive corporate culture mitigates risk

A recent study out of CyLab, Carnegie Mellon University's Security and Privacy Institute, found that negative deterrence actions like employee constraints, monitoring, and punishment don't work to reduce insider risk.¹⁰⁷ What does work: putting employee engagement, connection, and well-being front and center.

To support the well-being of your people, it's important to create channels and mechanisms to listen to their concerns, providing an opportunity to give and receive feedback and embrace collaboration. Taking a holistic, purpose-built approach that can pull signals together into a cohesive view across an organization provides a better understanding of the relevant trends in the organization and better risk reduction. For this reason, organizations are turning to ML to uncover hidden signs of workplace risk such as inappropriate communications, threatening behavior, or actions that would negatively impact employees and the business. By identifying patterns and violations, technology can flag risk while intervention is still possible, while continuing our commitment to end-user privacy.

Learn more:

- [To Thrive in Hybrid Work, Support Flexibility in Work Styles \(microsoft.com\) \(9/9/2021\)](#)
- [Insider risk: Protect company data with insider goodwill — Quartz \(qz.com\) \(June 2021\)](#)
- [Data security: Eliminating insider risk in the hybrid workplace — Quartz \(qz.com\)](#)
- [Learn about insider risk management - Microsoft 365 Compliance | Microsoft Docs \(3/17/2021\)](#)
- [Reducing Code of Conduct and Regulatory Compliance Violation Risks - Microsoft Tech Community \(5/12/2021\)](#)
- [Fostering safe communication at work - The Washington Post \(11/17/2020\)](#)
- [Microsoft Customer Story-Avanade uses a light touch—and Microsoft Insider Risk Management—to lessen insider risk \(3/24/2021\)](#)

Assume positive intent; mistakes happen.



CHAPTER 6

Disinformation

Introduction

Disinformation as an emerging threat

Mitigation through media literacy

Disinformation as an enterprise disruptor

Campaign security and election integrity

INTRODUCTION:

Critical attention required on the increasing sophistication and scope of disinformation

ERIC HORVITZ, CHIEF SCIENTIFIC OFFICER

Disinformation refers to the deliberate use of false information with the intention of influencing public opinion. Efforts to fabricate falsehoods for the purposes of manipulating the masses have a long history. However, new forms of disinformation have come to the fore over the last decade, enabled by advances in computing methods and infrastructure that have transformed the power, scope, and efficiency of disinformation campaigns.

Widely used consumer platforms and services, such as social media, creator platforms, search engines, and messaging services, now provide state and non-state actors with powerful channels for distributing disinformation. Beyond channels, these services provide malevolent actors with ready-made tools to experiment, monitor, iterate, and optimize the impact of disinformation campaigns.

Commercial online platforms have been harnessed by these actors as engines of disinformation to power messaging programs aimed at political influence, polarization, and chaos. Disinformation strategies are growing in sophistication, including the concerted use of multiple services¹⁰⁸ to reinforce messages across platforms.

On a second front, advances in machine learning (ML) and graphics have led to widely available tools for fabricating high-fidelity audiovisual content, referred to as synthetic media and deepfakes. For decades, photos and comments by political leaders have been

manipulated or taken out of context in disinformation efforts, often with dramatic effects. However, technologies for generating deepfakes are providing malevolent actors with powerful, general palettes for fabricating behaviors and events. These methods are injecting new powers of persuasion into disinformation campaigns.

In a third area of concern, artificial intelligence (AI) methods can be used by state and non-state actors to formulate and drive powerful psychological operations that leverage insights and data about human cognition. ML and reasoning can be used to profile

individuals and groups and to generate personalized programs of disinformation aimed at influencing beliefs, opinions, and actions.

The repurposing of consumer computing infrastructure, use of tools for generating synthetic media, and harnessing AI to guide psychological operations are troubling separately and in synergy. Together, they are supercharging disinformation, with grave implications for the health and vibrancy of democracies that depend critically on educated and aware citizenries.

THESE
METHODS
ARE INJECTING
NEW POWERS OF
PERSUASION INTO
DISINFORMATION
CAMPAIGNS.

What might we do in the face of these developments?

We need to critically attend to the increasing sophistication and scope of disinformation—and to engage on multiple fronts. First and foremost, we need to invest deeply in modern media literacy, to educate people about how to understand, expect, and recognize disinformation and misinformation. Work on media literacy extends beyond education and includes efforts to provide new kinds of tools that can help people to critique the source and veracity of news and information. Second, we need to support high-quality journalism, including trusted news organizations. It is essential to have committed reporters on the ground to see, hear, and report with clarity on events and incidents. In addition, we need to assure the health and vibrancy of local journalism.

On the technical front, there is promise in applying AI pattern recognition technologies to detect patterns of communications and content that reveal an intent to deceive. Such work includes efforts to identify audiovisual and text-based media as fabricated. On another front, efforts in networking technologies

can be aimed at identifying primary locations and organizational sources of disinformation. Finally, there are promising developments with technologies that employ a set of methods, including cryptography, security, and database technologies in production tools and pipelines that serve to certify the origin and history of edits to online media content, referred to as the *provenance*¹⁰⁹ of the content. Exciting progress with media provenance and authenticity is being nourished by strong cross-organization collaborations.

We are facing unprecedented disinformation campaigns and related cyber operations by state and non-state actors. These campaigns target public awareness and knowledge with disinformation, while others target enterprise operations and confidence. It is important to stay aware of developments and to come together to address the challenges with awareness, technologies, and policies. Addressing the new and evolving challenges will take ongoing investments, innovation, and activity on multiple fronts. This important chapter reviews some of these challenges and provides insights on directions forward.


Disinformation as an emerging threat

Disinformation has been a steadily evolving method of information warfare, most recognized in the United States after the 2016 election, but well-known globally since the Cold War. This approach is high-stakes and effective for creating social discord, increasing polarization, and in some cases, influencing the outcome of elections. Nation state actors with geopolitical aspirations, proponents of radical ideologies, violent extremists, and economically motivated enterprises can manipulate online narratives with easy and unprecedented reach and scale, creating significant societal impact.

The general motive to spread disinformation is to damage the reputation of an entity, mislead consumers about the information, or influence the outcome of a proscribed event. It is one of the greatest threats to democracy, open debate, and free and modern society.

Commodity cloud computing, open-source research, AI tools and algorithms, and the speed and scale of social media have created a perfect storm for the rise of disinformation and malicious synthetic media popularly dubbed deepfakes.¹¹⁰ AI techniques to create hyper-realistic digital falsification, also known as deepfakes, include manipulated audio, video, images, and text, which will seriously challenge our ability to discern truth from falsehood. In this report, we use the term “deepfake” as AI-generated manipulated media used for malicious purposes.

Mapping the problem



¹⁰⁹ A promising step forward on disinformation – Microsoft On the Issues ¹¹⁰ Despite its popular usage, this term technically refers to a specific type of ML (a process by which exposure to large and diverse amounts of data allows a computer to improve its own performance) that uses layered neural networks (computer processors connected in a way that mimics how information travels in the human brain) to enhance the accuracy of the ML algorithms. Deepfakes use two algorithms: the first algorithm creates a video, and the second one tries to determine if the video is real or not. If the second algorithm fails, the video is fake, the first algorithm tries again, until the resulting image looks sufficiently believable.

In the attention economy, websites and platforms earn revenue from the time users spend on them. This advertisement-based business model incentivizes recommendation engines and curated timelines with clickbait headlines and scandalous news, opinions, and falsehoods. The timeline curation algorithms give rise to creating echo chambers, filter bubbles, and unintentional tribalism. **Because of echo chambers and filter bubbles, users predominantly see the content that matches their beliefs, biases, and desires, reinforcing their confirmation bias and filtering out opposing viewpoints.** Nefarious actors misuse the attention economy and take advantage of the advertisement business model to manipulate social media narratives to create divisions and sow discord.

“Filter bubbles” and “echo chamber”



In the attention economy, users are increasingly exposed only to information that matches their preferences (blue areas in the graphic). Filter bubbles represent algorithms that choose content based on the user’s previous search histories and other online activity.

Parallels in cybersecurity

Cyberattacks compromise the confidentiality, integrity, and availability of digital systems. The difference between a disinformation attack and a cyberattack is the target; disinformation is also an attack and compromise of our cognitive being. While cyberattacks are exploits of computer infrastructure to create disruption, disinformation exploits human infrastructure (our inherent cognitive biases, logical fallacies, and psychological vulnerabilities), and the attack compromises logical, analytical, and critical thinking.

We can therefore think of the threat posed by disinformation and computational propaganda as *cognitive hacking*. A cognitive hack attempts to change the target audience’s thoughts and actions using disinformation to manipulate the way they perceive reality. Nefarious actors accomplish cognitive hacks utilizing various techniques, including manipulating, mis-contextualizing, or misappropriating information. Ultimately, these hacks can create social discord, exacerbate

polarization, influence election outcomes, disrupt democratic principles, enable financial fraud, and threaten life and property.

Both disinformation and cyberattacks are used by an adversary to achieve disruption. A well-coordinated disinformation campaign can fill the broadcast and social channels with false information and noise, creating narratives that play with emotions and drown out the true narrative. This maneuver is similar to a distributed denial of service (DDoS) attack that floods the target services and networks with superfluous requests to connect and overload the system to prevent legitimate requests from being fulfilled. Disinformation can also be used for social engineering threats on a mass scale. Like phishing attacks to compromise IT systems for data extraction, disinformation campaigns play on emotions, giving cybercriminals another feasible method for scams. Deepfake videos and audio can trick employees into releasing or sharing login credentials, which can then be used to gain access to an enterprise’s network.

Deepfakes

Deepfakes are photos, videos, or audio files manipulated by AI in hard-to-detect ways. The weaponization of deepfakes can have a massive impact on an economy and national security. By eroding public trust in the media, deepfakes undermine the credibility of journalism. As this credibility is eroded, deepfakes also give rise to the “Liar’s Dividend” phenomenon. In an environment where it is unclear what is real and what is fake, it becomes easier to discount any inconvenient or unpopular truth as fake. As an example, deepfakes could enable a public figure to claim that their real actions are just a fake.

The proliferation of deepfake technology can directly harm individuals. Deepfake pornography has been used to objectify and victimize people without consent, especially women or those who identify as women. This revenge pornography is now a significant problem: over 96% of deepfake videos

are nonconsensual pornography.¹¹¹ A person’s reputation can be damaged beyond repair by a single deepfake video posted to social media. Criminals can and will take advantage of AI and deepfakes to increase the effectiveness of scams and other nefarious activities. We have seen recent examples of business¹¹² and romance¹¹³ fraud using audio deepfakes. Even after effectively debunking the deepfake, the damage remains. Deepfakes also present direct threats to society. Nefarious nation state actors work around the clock to spread disinformation campaigns, which increasingly include deepfakes, to create ideological conflicts in democracies around the world. They have effectively used disinformation and empowered domestic actors to disrupt elections and undermine efforts to contain the COVID-19 pandemic. Technology innovations like replacing video codecs with deep neural networks are making it easier to create deepfakes in live video calls.

Microsoft Research, in coordination with Microsoft’s Responsible AI team and the Microsoft AI, Ethics and Effects in Engineering and Research (AETHER) Committee, continues to develop technology for training and testing deepfake detection technologies.

Mitigation through media literacy

Pew Center Research indicates that people are increasingly concerned about their ability to discern false information and sift through the volume of information they encounter in their daily lives.¹¹⁴ Media literacy is one of the most effective tools available to improve individual resilience to disinformation. From helping people understand minor editorial corrections, to preventing fake news from gaining traction, and reducing the likelihood that foreign influences undermine an election, media literacy inoculates individuals by helping them think critically about information they encounter.

Even a short intervention with media literacy education has been shown to make a significant difference in understanding disinformation, identifying the motivations and context, and reducing belief in inauthentic content.¹¹⁵ News media organizations, technology companies,


universities, and social media companies have all started to implement media literacy efforts, which include resources such as labeling, contextualizing, providing further reading or resources, and even gamified tools like online quizzes.

Improving media literacy is vital to addressing disinformation. Interventions should be contextualized and tailored to the audience, starting from a young age with curriculum for responsible digital citizenship incorporated into standard civics education. Some US states, most notably Florida and Ohio, have already started piloting and developing media literacy curriculum in classrooms.¹¹⁶ In Finland, when the country was targeted with foreign influence operations and disinformation campaigns by Russia

in 2014, the government introduced multi-platform information literacy and strong critical thinking as a core component of a national curriculum.¹¹⁷

Media literacy curriculum should help individuals understand information they encounter, evaluate the plausibility, verify the source and search for other reputable sources on the topic, and interpret context. In today’s online world with essentially limitless information available, these skills may not be intuitive, but are certainly ones that can be learned and sharpened over time. Further, media literacy education does not need to be confined to the classroom: these skills are universally relevant, especially for digital nonnatives, or anyone who regularly consumes news and information online.

Approaches to countering disinformation



¹¹¹ *The State of Deepfakes: 2019 Landscape, Threats, and Impact* | Sensity AI ¹¹² *Thieves are now using AI deepfakes to trick companies into sending them money – The Verge* ¹¹³ *Romance Scammer Used Deepfakes to Impersonate a Navy Admiral and Bilk Widow Out of Nearly \$300,000 (msn.com)* ¹¹⁴ *Many Americans Believe Fake News Is Sowing Confusion* | Pew Research Center (journalism.org) ¹¹⁵ *A digital media literacy intervention increases discernment between mainstream and false news* | *South China Morning Post* (India. (2020)) ¹¹⁶ *Media Literacy Around the States* | *Media Literacy Now* ¹¹⁷ <https://www.theguardian.com/world/2020/jan/28/fact-from-fiction-finlands-new-lessons-in-combating-fake-news>

<https://t.me/learningnets>

Microsoft has and continues to develop media literacy resources to help consumers discern information. Ahead of the US 2020 elections, we released two online quizzes, “Spot the Deepfake”¹¹⁸ and “Know My News,”¹¹⁹ aimed at raising awareness about synthetic media technology and understanding news sources. It also released a “VaxFacts Quiz”¹²⁰ to heighten awareness around COVID-19 misinformation as vaccines were becoming more widely available. Microsoft has also developed a hybrid threat training curriculum that highlights how **threat actors are increasingly using cybersecurity and disinformation attacks in tandem to accomplish their goals**. These training modules have been customized for the political community, including political campaigns, parties, and government entities, as well as for journalists and human rights organizations.

Learn more:

An update on our effort to help preserve and protect journalism – Microsoft On the Issues (6/16/2021)

Disinformation as an enterprise disruptor

Disinformation traverses a social diffusion chain comprised of both intentional and unsuspecting agents. Intentional actors include hackers, disruptors, and other agents that generate or propagate disinformation. These agents could include nation states targeting a society, industry, or social arena. They could also include agents engaged in corporate counter-espionage activities or anti-competitive disruption. Unintentional agents might include personnel, vendors and suppliers, other partners and stakeholders, customers, and even competitors that continue the propagation of disinformation, often not realizing what they are doing. They concurrently become victims and unsuspecting agents of disinformation. Through these agents and other means of information dissemination, a parallel and concurrent wave of disinformation traverses or envelops the enterprise. This wave can be simple and anticipated, like the rise of a tide, or it can have the effect of an unpredictable tsunami about to weigh down on an unsuspecting coastal village. In this case, the coastal village would be the enterprise information

community. Like its disruptive effect on social discourse, such disinformation carries with it the potential to disrupt corporate decision making, create commercial confusion and discord, and sow doubts in the minds of employees, customers, and markets. The following three areas cover some of the more common considerations that enterprises could prioritize to ensure preparedness to mitigate the disruptive effect of disinformation.

1. Parallel to its effect through social media, disinformation has made its way into enterprise workflows that are dependent on data collection, aggregation, and distribution practices. These workflows and data practices may be automated or manual in nature, and even more advanced practices, such as those leveraging AI capabilities, are prone to breach. Sophisticated AI algorithms can be fooled, or their underlying models overwhelmed by excess disinformation and other types of information attacks, leading to erroneous or anomalous insights and outcomes. In general, automated practices face the threat of breach and corruption in the absence of adequate detection, timely notification, isolation and eviction, and defense in depth. In contrast, manual practices, such as copy/paste, swivel chair (entering data into one system and then entering the same data

into another system), screen-scraping, and aggregation of public information through human evaluation and automated means, are also prone to disruption, corruption, or nuanced modification. To address this growing threat, enterprise decision makers should evaluate which of the enterprise’s critical information gathering and distribution processes could benefit from more resilient practices. Which inflection points, or points where information is interchanged or interpreted, need more rigor, controls, or other forms of checks and balances? Based on this analysis, controls and even out-of-band crosschecks and validations should be introduced to mitigate the threat of polluted or mischaracterized information that could potentially corrupt the data stream and subsequent insights that are derived from the data.

2. Signals and data within the enterprise could be compromised through security vulnerabilities or attacks and infused with disinformation. Instrumentation and threat signals for performance of enterprise resilience are often seen as having secondary importance and are therefore not given the level of scrutiny that core or primary services or applications receive. Data and signal collection are sometimes relegated to supporting applications and systems that are bolted on as

¹¹⁸ <https://www.spotdeepfakes.org/en-US> (in partnership with the University of Washington Center for Informed Public) ¹¹⁹ <https://www.knowmynews.com/en-us> (in partnership with NewsGuard)

¹²⁰ <https://www.knowmynews.com/en-us> (in partnership with NewsGuard)

adjunct environments. Attackers and disruptors know this and strike where they believe weaknesses could exist. Their primary areas of reconnaissance include operational security and control systems alongside discovery of code and configuration vulnerabilities, especially within secondary or supporting systems. Successful attacks could disrupt threat signals, or infuse misleading data, if they are able to bypass intrusion detection. Enterprise decision makers should reevaluate the scope and security of their protected assets and associated instrumentation and signals and ensure that intrusion detection perimeters around core or critical assets have been set to encompass these areas as well. This review includes evaluation of controls that measure the effectiveness of network segmentation, validate adherence to data and signal classification schema, inbound and outbound traffic features and patterns, classification of data and dependencies, and instrumentation that ensures data integrity as it traverses across systems, at rest and in backups. Instrumentation and signals that evaluate and indicate the performance of critical systems should be cataloged and put through a security and operational review for vulnerabilities and possible causes of compromise and injection of disinformation.

3. Situational intelligence could be supplanted with disinformation or nuanced in ways to generate bias or create doubt in the minds of decision makers or front-line personnel. Situational intelligence includes threat intelligence, crisis intelligence, disaster data, and other types of information that helps to increase an enterprise’s understanding of an operating or competitive environment. One of the key strategies of disruptors is to supplant situational intelligence with disinformation. This strategy is often deployed in environments that are considered out-of-band or beyond the control and reach of the enterprise, including social and other media platforms, public sites and portals, and other types of campaigns that could target the enterprise or its complex and myriad supply chains. Disruptors sometimes generate so much out-of-band disinformation that it overwhelms facts and otherwise accurate data. Their disruption tactics go beyond the enterprise and its direct information, extending often to suppliers and third-party affiliates. Their goal is to introduce doubt in the minds of decision makers and/or front-line personnel. Disinformation that could disrupt suppliers, logistics, deliveries, orders, and fulfillment poses threats to the supply chain and confidence in product and resource availability.

The impact of diminished confidence could drive higher pricing and impact production costs. Doubt in the enterprise’s ability to maintain a grip on prices, supplies, and suppliers could exacerbate disruption and force accelerated timelines or induce imprecise or incorrect decisions. Alternatively, it could introduce latency or doubt in decisions, such as recovery, failover, or fallback. Disruptors could also try to diminish front-line confidence in leadership by corroding confidence or by presenting false or inaccurate information and leading or trailing indicators. Enterprises that are dependent on external sources of intelligence for critical functions need to be vigilant about information they ingest or consume. Information gatherers and decision makers need to validate sources and have means of cross-checking intelligence. Rating systems and schema should be used, and the veracity and provenance of the intelligence needs to be confirmed independently.

Recommendations: Four-point plan for enterprise executives

As enterprises develop and mature their capabilities to respond to information disruption and disinformation, there are new imperatives that are emerging in the areas of security controls, threat classification, and analysis. To support and enhance the enterprise's ability to improve its response, prioritize mitigations to address the highest risks, and to make investments that propel or accelerate its corporate objectives, the following practices could be leveraged as a starting point to develop sustainable mitigations and improvements to counter the effects of information disruption and disinformation.

1. Catalog enterprise exposures to disruption, manipulation, and disinformation.

Enterprises should begin with the arduous task of cataloging or keeping track of disinformation attacks on their information systems and data. This catalog would be an addition to the catalog of threats and attacks that are part of enterprise security controls.

- Include a classification of information manipulation and disinformation with clear indicators of targeted outcomes and impact of the attacks.
- Note the sources of the attacks or information leading to their origins and motivation, if known.

- Categorize and record the content that was manipulated or introduced so that patterns can be determined or detected.
- Determine the means of propagation and diffusion through the enterprise.
- Identify characteristics of actors and agents to the best ability that is available at the time.
- Conduct candid and objective identification of corporate functions, processes, and systems that consumed or were impacted by the disinformation.

The net impact of the disruption or information infiltration must be evaluated and eventually included in considerations of corporate liability.

2. Assess the impact of manipulation or disinformation.

An emerging area for enterprise management and human resources is to conduct impact analyses of disruption and disinformation. Managers and human resource departments should study the behavior of enterprise employees, partners, and customers as they reacted to disruption or interacted with disinformation. Security and data science teams should run A/B tests on AI and trained algorithms to ensure they have not been corrupted or their models tampered with. Finance teams and economists could look at the costs of disruption, including opportunity cost, and the liabilities associated with impact of market confidence, revenue, growth, and operational cost increases associated with disinformation.

3. Quantify the consequences of disruption.

By examining dependency maps of data flow, analysis, and the points at which humans and systems were impacted, enterprise leadership can estimate or calculate the consequences and collateral damage caused by disinformation and disruption. The goal of such an examination is to quantify the blast radius of any potential disruption and assess its severity on enterprise operations, functions, and reputation. The likelihood and impact of such disruption should be used to inform investment in mitigations and controls and help prioritize safeguards that need to be implemented to give the enterprise the most appropriate defense against the impacts of disruption and disinformation.

4. Assess privacy implications of disruption.

Disinformation and information disruption campaigns sometimes attack or threaten customer and other protected data, which can have significant privacy implications for enterprises. Primary sources or stores of such data are difficult targets as they are kept deep inside the defensive network of controls, intrusion detection, and layers of authentication and authorization. Modern enterprise controls include Zero Trust approaches and tightly controlled or restricted access to data that is subject to privacy regulations, controls, and guidelines. Sovereign rules and obligations might also apply. It is under such constraints and restrictions that enterprises are expected to provide data resiliency, including

hot and cold backups, or active-active models of redundancy or concurrency. Information disruptors and attackers aggressively search for backup facilities, dormant data stores, or unattended data and data services that are slated to be deprecated. When they find a weakness in security controls, not only do they threaten breach of data but their actions can also trigger potentially impactful privacy incidents. Enterprise privacy and security teams must take a closer look at controls, including security measures, encryption at rest and in transit, of primary data sources and any backups or alternate data stores. Gaps in these environments and controls must be reported to enterprise leadership and prioritized to be addressed in a timely manner.

Campaign security and election integrity

Threats to democratic processes from cyber-enabled interference are a critical concern. Microsoft has shut down dozens of websites widely attributed to nation state actors that have been used to target elected officials and candidates, democracy-promotion organizations, activists, and the press. We’ve also seen attempts by nation states to target and exploit key building blocks of democratic systems, including voting systems and political campaigns. We also endured the calculated manipulation of social media platforms in efforts to sow misinformation and disinformation.

The Defending Democracy Program at Microsoft was created to advance technology, increase cyber resilience, and engage with government, campaigns, and democratic stakeholders to address threats to democratic processes from cyber-enabled interference. In partnership with the rest of the Microsoft Security community, the Defending Democracy team has helped protect two major US elections and dozens of national elections around the globe. This unique and important issue space raises an interesting set of cybersecurity insights, challenges, and solutions that we must address to ensure that our democratic institutions remain secure.

Unique cybersecurity challenges in political campaign security

The structure and lifecycle of political campaigns present unique cybersecurity challenges. Campaign organizations comprised of a mixture of staff, volunteers, and consultants are often created and expanded rapidly once an individual declares their candidacy. As a result, team members are often asked to use their own personal devices, including cell phones or laptops, throughout the campaign. Further, with constrained budgets and limited IT expertise, decisions about which email and file sharing provider to use are more likely tied to personal preference than extensive security and risk evaluations. With people regularly joining and leaving a campaign and utilizing a variety of personal devices to conduct campaign business, it becomes difficult to manage and enforce strong cybersecurity practices.

These realities expose political campaigns to significant cybersecurity threats, compounded further by the asymmetric advantage that malicious actors maintain in terms of sophistication, resources, and capacity. Political campaigns are high-value, high-visibility targets, but may have as few as one person dedicated to cybersecurity for the organization.

It is critically important that we in the security community continue to raise awareness of cybersecurity issues and best practices within

political campaigns. This can be as simple as establishing strong communication channels between campaign staff and the private sector to ensure they have a point of contact if issues arise. It is also imperative for candidates to establish a culture of cybersecurity within their campaigns, not just among the IT staff but inclusive of everyone who interacts with the campaign. Relatively easy actions such as turning on multifactor authentication (MFA), utilizing a password manager for strong unique passwords, and using secure communication channels for campaign communications make a tremendous difference in improving the resilience of the entire organization.

Microsoft has developed a service to address the challenges associated with campaign staff using personal devices and accounts for campaign business. Microsoft AccountGuard¹²¹ is a security service that unifies threat detection and notification across all accounts, and is available at no cost to campaigns utilizing Microsoft 365 products. When a campaign enrolls in AccountGuard, they can also extend coverage to anyone who interacts with the campaign, including staff, volunteers, interns, consultants, and more. The service then provides notification and remediation guidance in the event of a nation state threat or compromise on campaign-related accounts. The service currently protects over 40,000 accounts for political customers, including political campaigns, political parties, technology vendors, and elections departments.

Heading into the US 2020 Presidential election cycle, the AccountGuard program also offered enhanced identity protection features and resiliency reviews to customers involved in the election. Organizations that took advantage of these resources saw an average increase of 18% to their Microsoft Secure Score. In addition to political customers, the service is available to other high-risk groups including human rights organizations, journalists and media organizations, and healthcare organizations.

AccountGuard distribution



AccountGuard is currently available in 32 democracies around the world.

Learn more:
[Expanding AccountGuard protections for high-risk customers in 31 democracies – Microsoft On the Issues](#) (3/9/2021)

Threats to election infrastructure

As the Microsoft security community partners across the public and private sectors to secure election infrastructure, the differences between how cloud and on-premises systems are used in elections becomes clear, as well as the various types of election technologies that must be secured. We form our threat models around three main categories of systems, each with unique security considerations and risk profiles.

Globally, voting systems are statutorily required to be always disconnected from the internet and are effectively “air gapped” systems. Though this poses certain challenges for maintenance and updating devices, it acts as an extremely effective mitigation

against remote threat actors. Though some form of connected or online voting systems are starting to rise in popularity, they remain in the minority of systems.

Securing election systems

UPDATES AND PATCHING

The primary recommendations for elections officials and private vendors who support elections IT is to have a comprehensive strategy to patch systems regularly and keep software up to date, and to use MFA. This is particularly relevant to any internet-connected systems, such as election support technologies and back-office IT systems. During the lead-up to the 2020 US election cycle, the ZeroLogon vulnerability (CVE-2020-1472¹²²) was announced, less than three months before the

election. This event served as a good reminder that vulnerabilities can be announced at inopportune moments, and the security of an environment relies on the ability to apply patches rapidly. Customers who were using cloud infrastructure were secured against this vulnerability significantly faster—often automatically and immediately—than those running on-premises servers that needed to be patched.

SHARED SERVICES

Unlike a private sector company that is typically in charge of their entire IT estate, local governments tend to share IT systems with their state, regional, and national-level counterparts. Although this can be cost effective and beneficial from a management perspective, it spreads out the risk of cyber intrusion from an unrelated government organization into an elections department. For example, if a research lab at a state-run university were compromised, because of shared IT systems attackers could move laterally to attack a county election’s voter information portal. Furthermore, responsibility for maintenance and security is often split between private service providers and local governments. Sharing services highlights the need for network segmentation wherever possible to limit the impact of a cyber incident that either starts within or could migrate to an elections organization.

Voting systems	• Vote casting systems (electronic ballot marking devices, “voting machines”) • Vote tabulation devices (paper ballot scanners)
Connected election support systems	• Voter registration databases • Election-night results reporting sites • ePollbooks (voter check-in tablets) • Information portals • Not regulated, connected to the internet, shared ownership ← Highest risk
Election back-office IT	• Voter information portals • Day-to-day workstations and server infrastructure • Email and files

There are many connected components of an election beyond just casting a ballot at the polls—from data integrity concerns, to disinformation, to protecting election officials’ online identities.

While many of these things may be invisible to the average voter, each offers a potential pathway for an adversary to attack the integrity of an election.

SMALL TEAMS AND SMALL BUDGETS

In some countries, including the United States, elections are not managed federally; each municipality is responsible for running and managing its own elections infrastructure. The asymmetric threat of advanced adversaries targeting local-level elections offices that are time and resource constrained is of real concern. As discussed earlier in this report, advanced persistent threat actors have IT knowledge and monetary resources that exceed the annual budgets of local municipalities. Therefore, steps must be taken by both the private and public sector to concentrate cybersecurity talent and resources in a way that helps to offset and balance that asymmetry.

To help elections organizations address these concerns, the Defending Democracy team focuses heavily on securing and supporting customers using cloud services to support their elections, in both the private and public sectors. Our goal is to use the cybersecurity expertise at Microsoft to offset the asymmetric threats faced by local elections or small elections vendors.

Learn more:

[Keeping your vote safe and secure: A story from inside the 2020 election – On the Issues \(microsoft.com\)](#)

[Protecting political campaigns from hacking – Microsoft On the Issues \(5/06/2019\)](#)

Election integrity

To address growing concerns in voter trust among election systems, and to help drive forward principles of software independence, election auditability, and security and cryptography, Microsoft has developed and launched a project called ElectionGuard. Microsoft ElectionGuard is a free open-source software development kit designed to make voting systems more secure, auditable, verifiable, and efficient. This is done by implementing principles of end-to-end verifiability and answers the question: *How can I trust the accuracy of an election outcome if I worry that the software, hardware, transmission infrastructure, or personnel responsible for conducting the election could be untrustworthy?*

Unlike banking software or other high-security industries, secret ballot elections require that an individual’s data (votes) must be a secret to all parties, and there can be no direct tie between a person’s identity and their vote, other than the acknowledgment that a voter has cast a ballot.

ElectionGuard implements these end-to-end verifiable features:

- Enables every individual to verify that their vote is included in the final election tally, with a unique verification code.
- Ensures that no one specific individual’s vote can be revealed.

- Cryptography guarantees that no votes can be changed after being cast without being detectable.
- Ensures that no one specific individual’s vote can be revealed.
- Provides assurances that the system is recording votes properly and not “changing votes.”
- Allows interested third-party watchdog organizations to verify that the tally was summed properly, and that the system operated without errors.


Learn more:

[LinkedIn post: What is ElectionGuard? \(Microsoft on the Issues\) \(3/27/2020\)](#)

MOVING OUT OF THEORY AND INTO REAL ELECTIONS

ElectionGuard had its first real-world pilot election in Fulton, Wisconsin in February 2020 when it was added to VotingWorks’ touchscreen ballot-marking devices for a municipal election in Rock County. Hundreds of citizens voted for the first time on devices running ElectionGuard. After their ballot was cast, each voter received a verification code that they could take home with them to check online that their unique vote had been included in the final election tally.

ElectionGuard is free, open-source, royalty-free software, and all source code is publicly viewable on GitHub.¹²³ Microsoft believes that election security



should not exist in a vacuum, and independent security researchers should be able to validate the integrity of the software. To that end, we have implemented a bug bounty program where Microsoft will award the security community for finding vulnerabilities in the ElectionGuard software. Since launching in October 2019, the program has awarded tens of thousands of dollars in bounties to researchers across multiple continents. In this way, security issues were discovered and responsibly reported by the community, all patches were issued within 90 days of reporting, and potentially vulnerable code was never used by voters in production election systems. Public trust in elections is, in part, dependent on the independent auditability of the systems and processes that support our elections. We continue to welcome the trained eyes of security researchers across the world to suggest improvements to the service and to independently test ElectionGuard's security and cryptography.

Learn more:

Hart and Microsoft announce partnership to make elections more secure, verifiable – Microsoft On the Issues (6/3/2021)

Microsoft wants to make voting more trustworthy. Its partnership with Hart InterCivic will help – CNN (6/3/2021)

Microsoft hopes this technology can help fix America's elections (cnn.com) (2/22/2020)

Trainings as mitigation

To date, Microsoft has trained more than 1,500 participants from the political, civil society, media, and human rights sectors on topics spanning cybersecurity and disinformation threats. These trainings have been tailored to audiences around the world, including political campaign staff, political parties, and other government entities in democratic countries. Ahead of the US 2020 elections, Microsoft partnered with NYU's Brennan Center for Justice and the CISA to train more than 400 US election officials on topics including ransomware, pandemic preparations, threat intelligence, and an election day simulation exercise. Further, Microsoft has collaborated with PolitiFact at the Poynter institute to develop [a hybrid threat training curriculum exploring the intersection of cybersecurity and disinformation, as these threat vectors increasingly overlap](#), especially in regard to attacks against political and media organizations.

Learn more:

Defending Democracy Program – On the Issues (microsoft.com)

Countermeasures needed

To ensure access to credible information and preserve freedom of expression, we need a multistakeholder and a multimodal approach. The main objective of any countermeasure is to mitigate the negative societal impact of disinformation.

Specifically, when it comes to malicious synthetic media, the approach must be twofold: (1) to reduce the exposure to malicious deepfakes, and (2) to minimize the damage it can inflict.

Media literacy efforts can be enhanced to cultivate a discerning public. Deepfakes will have a limited adverse impact if media consumers use logical thinking and common sense to differentiate between fiction and reality. Deepfakes have such a potentially devastating impact because many individuals assume that a video, a photograph, or an audio is real if it aligns with their preconceptions and biases. Alternatively, often people believe it's fabricated if it contradicts their beliefs.

We need meaningful regulations and appropriate laws to govern disinformation and deepfakes so that, at the very least, the perpetrators are held accountable. Without legislation and legal remedies, people are vulnerable to disinformation campaigns and deepfake revenge pornography, fraud, and other harms. Of course, legislation must take a balanced approach toward freedom of expression and speech.

The technical countermeasures are not simple. The synthetic media's technological development continues to outpace what is possible by algorithms and other technologies. Technical solutions for deepfakes fall into two categories: (1) detection and authentication and (2) provenance. Detection of deepfakes and disinformation is difficult. Human-AI collaboration can help, but context, cultural differences, and intent make it hard to decipher disinformation objectively. In the long term, provenance solutions will help. Media authentication, authoritative content, and standards akin to SSL/HTTPS for web traffic must be developed for media. Technology that can help "sign" the media to find the media publisher, machine-readable fingerprinting, and watermarking technology to tamperproof media would reduce media manipulation to effectively combat the spread of deepfakes.

We're not saying ElectionGuard makes it impossible to hack voting machines; we're saying ElectionGuard makes it pointless to hack voting machines.

CHAPTER 7

Actionable insights

Introduction

Summary of report learnings

Conclusion

INTRODUCTION:

Five cybersecurity paradigm shifts

ANN JOHNSON, CVP, SECURITY, COMPLIANCE & IDENTITY

In working with organizations from around the world, we recognize the need to enable people to work productively and securely and from a variety of non-traditional locations and a variety of devices. Through these interactions, we’ve learned a lot about the role that cybersecurity plays in helping organizations maintain business continuity as we adapt to a hybrid work world. As a result, we anticipate five cybersecurity paradigm shifts that will support the evolution of work in a way that centers on the inclusivity of people and data.

1. The rise of digital empathy

When we consider building security into the productivity experience, there can be a tendency to focus on security from purely a technology perspective. However, during times of constant disruption and change, when people are susceptible to increased stress reactions, the importance of digital empathy comes into greater focus. Digital empathy involves thinking about the ways in which people behave and engage with technology. In this way, sociology and humanities considerations are essential to the evolution of technology and cybersecurity.

Empathy isn’t just for in-person interactions. By applying empathy to digital solutions, we can make these solutions more inclusive. In cybersecurity, that means building tools that accommodate more diversity with respect to

people and their ever-changing circumstances, their diverse perspectives, and varied abilities. For example, rather than requiring people to do unnatural things to engage with security, which can also increase their likelihood of getting distracted when busy, factoring in digital empathy leads to inclusion of security professionals with a broader range of abilities, skill sets, and perspectives—for greater diversity and effectiveness of cybersecurity solutions. It also means developing technology that can forgive mistakes.

Digital empathy will be critical to how we move forward as an industry. Whether it’s an organization—or an individual—our ability to be empathetic will help us to understand and adapt during times of constant change.

2. The Zero Trust journey is becoming increasingly important

Zero Trust is an “assume breach” security posture¹²⁴ that treats each step across the network and each request for access to resources as a unique risk to be evaluated and verified. This model starts with strong identity authentication everywhere. Multifactor authentication (MFA)—which we know prevents 99% of credential theft¹²⁵—and other intelligent authentication methods¹²⁶ make accessing apps easier and more secure than traditional passwords.

As we look past the pandemic to a time when workforces and budgets rebound, Zero Trust will become the biggest area of investment for cybersecurity. This means that right now, every one of us is on a Zero Trust journey—whether we know it or not.

WE KNOW A
COMPREHENSIVE
APPROACH TO
OPERATIONAL
RESILIENCE MUST
INCLUDE CYBER
RESILIENCE.

3. Diversity of data matters

Microsoft tracks more than 24 trillion daily signals from a diverse set of products, services, and feeds around the globe. We were able to identify new COVID-19-themed threats—sometimes in a fraction of a second—before they reached customers. This is just one example of how the power and scale of the cloud has a clear advantage when it comes to combating threats.

As one example, in the last year, the diversity of data also allowed us to understand COVID-19-themed attacks in a broader context. Microsoft cyber defenders determined that adversaries were primarily adding new pandemic-themed lures to familiar malware.¹²⁷

4. The resiliency of a business is tied to its cyber resilience

Cyberattacks are increasing in frequency and sophistication and are deliberately targeting core business systems to maximize the impact of the attack or likelihood of a ransomware payout. With this context, we know a comprehensive approach to operational resilience must include cyber resilience.¹²⁸ At Microsoft, we benefit from a strategy that focuses on four basic threat scenarios: Planful events such as weather incidents, unplanned

events such as earthquakes, legal events such as cyberattacks, and pandemics like COVID-19. Cloud technology helps organizations develop a comprehensive cyber-resilience strategy and makes preparing for a wide range of contingencies less complicated due to its scalability.

5. A greater focus on integrated security

The first half of 2021 brought into stark reality the agility and callousness of our adversaries. To uncover shifting attacker techniques and stop them before they do real damage, organizations must be able to see across their apps, endpoints, network, and users. Facing a new economic reality, organizations will also be driven to reduce costs by adopting more of the security capabilities built into their cloud and productivity platforms of choice. To maximize the effectiveness of security organizations, tools must be fully integrated to improve efficacy and provide end-to-end visibility.

While digital acceleration will continue to influence the paradigm shifts that shape our industry, one thing remains the same; security technology is fundamentally about improving productivity and collaboration through secure and inclusive user experiences.

Summary of report learnings

What became clear as we compiled this report is how much technology is now baked into everything we do. We can't afford to treat technology and cyber risk as something separate and contained that IT and security teams are left to manage on their own. The examples in this report show that criminals will seek to exploit whatever technology we develop and introduce; the challenge is in understanding what form that will take. Because we can't always predict how technology will be exploited, we need to assume that anything we create or use will be a potential target and prepare ourselves to be as resilient as possible.

The key actionable learning from all the elements of this report is that to minimize impact of attacks we must truly practice good cyber hygiene, implement architectures that support the principles of Zero Trust, and ensure cyber risk management is integrated into every aspect of the business.

The following section summarizes some of the key learnings reinforced by the findings and insights in the report.



The cybersecurity bell curve:

Basic security hygiene still protects against 98% of attacks



Enable multifactor authentication

Make it harder for bad actors to utilize stolen or phished credentials by enabling multifactor authentication. Always authenticate and authorize based on all available data points, including user identity, location, device health, service or workload, data classification, and anomalies.

Apply least privilege access

Prevent attackers from spreading across the network by applying least privilege access principles, which limits user access with just-in-time and just-enough-access (JIT/JEA), risk-based adaptive polices, and data protection to help secure both data and productivity.

Keep up to date

Mitigate the risk of software vulnerabilities by ensuring your organization’s devices, infrastructure, and applications are kept up to date and correctly configured. Endpoint management solutions allow policies to be pushed to machines for correct configuration and ensure systems are running the latest versions.

Utilize antimalware

Stop malware attacks from executing by installing and enabling antimalware solutions on endpoints and devices. Utilize cloud-connected antimalware services for the most current and accurate detection capabilities.

Protect data

Know where your sensitive data is stored and who has access. Implement information protection best practices such as applying sensitivity labels and data loss prevention policies. If a breach does occur, it’s critical that security teams know where the most sensitive data is stored and accessed.

Cyber hygiene

Taking basic security precautions can help your organization prepare for and mitigate the overwhelming majority of modern cyber threats and helps to prepare for the evolution of threats as technology advances. The “cybersecurity bell curve” shows the activities that will have the biggest impact on reducing threats. Some of those actions, the impact they have, and recommendations for implementing them are described here.

Enable multifactor authentication

This continues to be the top recommendation as it was last year. MFA can stop credential-based attacks dead in their tracks. Without access to the additional factor, the attacker can’t access the account or protected resource. The introduction of passwordless technology and architectures makes this easier for employees and customers to use and also provides more security than traditional text (SMS) or voice approaches. MFA should be enabled on all accounts that support it, in a way that it is easy for all users to use it. It’s also important to ensure that people understand that they should not approve an MFA request unless they were trying to log in or access a system—many people automatically click to approve any pop-up that they receive. Digital empathy can be useful in understanding this behavior and helping to nudge people toward less risky decisions.

Apply least privilege access and secure the most sensitive and privileged credentials

When attackers breach an organization, they look for privileged credentials to provide them with access to sensitive information and systems. In addition to using MFA to protect login to an identity and ensuring that they have least privilege to access systems, the credentials that support that identity and provide access must be secured. Among other things, this will help to minimize the impact and breadth of pass-the-hash-style attacks in the event that malicious code is already running on a local machine or network. This includes securing hardware such as with a trusted platform module or hardware security module or using cloud authentication services that provide credential protection.

Separate accounts should be used for privileged access versus general internet and email access. Dedicated hardened workstations should be used for privileged accounts and to perform privileged tasks to prevent the chances of infection through general internet activity and email. Using JIT/JEA systems ensure they will only get exactly the rights needed to perform a task and only for as long as needed to undertake it. This should be combined with risk-based adaptive policies that deny access to resources when there is any doubt over the hygiene of the account or device.

Secure and manage devices (keep up to date)

An essential part of good cyber hygiene is ensuring that devices are kept up to date and configured correctly. Use endpoint management software to enforce policies that ensure the correct configuration settings are deployed and that systems are running the latest software. Wherever possible, take an evergreen approach to ensure all devices are constantly running the latest version of software. This includes ensuring a means of updating every piece of software or application so that there are no dependencies that prevent you from implementing the latest updates and patches. For those devices missing critical patches, restrict them from accessing sensitive resources.

The same approach should be taken for cloud services, using cloud security posture management to ensure that systems are configured correctly. Keeping software and systems up to date can be easier in the cloud where update domains enable migration to updated infrastructure for testing with the option to roll back easily if issues occur.

For systems, such as OT, where updating software is not as easy, a strong inventory of systems is needed to understand which equipment exists, and how vulnerable it may be to certain attacks. Incorporate add-on modules or replace hardware to ensure it achieves all seven properties of secured devices.

In systems where this is not possible, the environment should then be protected from other systems and monitored to detect any unexpected traffic or attempts to compromise the systems.

Use antimalware and workload protection tools

Antimalware and detection and response technologies should be deployed across the ecosystem to prevent attacks and provide warning of any anomalies or threats that may be attempting to breach the environment. This includes OT and IoT environments. For cloud systems, workload protection should be deployed across all systems from virtual machines and containers to machine learning (ML) algorithms, databases, and applications.

Protect data

Good cyber hygiene as outlined in the previous four steps can protect data, but it is also important that organizations know which sensitive data they have, and ensure that they have appropriate steps in place to protect it. In fact, there is often a regulatory requirement to do so. The General Data Protection Regulation (GDPR), for example, requires a risk-based approach to protecting the data of residents of the European Economic Area (EEA). To take a risk-based approach it is important to know your data—to understand what is sensitive and what may be subject to regulatory requirements. While there have been standards for data governance and data

protection for over 30 years, many organizations have struggled to implement them. As we move into a world where we increasingly collaborate and share data, it's important to ensure we understand what data we have, classify it accurately, and apply sensitivity labels as appropriate. This practice will enable us to use information protection and data loss prevention technologies to protect data with more confidence.

In the event of a breach, these practices can also help security teams to know where the most sensitive data is and whether it was exposed to attackers.

Adopt Zero Trust principles

This report has highlighted the sophistication and complexity of many of the attacks we are seeing and why it is increasingly difficult to prepare to counter these attacks. Zero Trust is important to reduce the exposure of sensitive data by limiting the inherent trust within an organization that an attacker would exploit—especially when people are connecting from everywhere and will not necessarily be coming from a “trusted” location. This is why adopting a Zero Trust approach is now a top priority for most organizations. In a world where it's harder to predict or prevent the attacker, it's important to assume they will get in and limit their exposure.

Isolate legacy systems

Not every system is capable of running the tools to enable Zero Trust. For example, many operational technology (OT) systems have long technology lifecycles and may run operating systems and software that can't be updated.

Network segmentation should be used to restrict access to these systems. This can help to ensure that operational technology is not exposed to the risks from hybrid working, and that IoT devices and sensors have access and connections only to the smart ecosystems they support.

This means that rather than trying to run the modern and legacy alongside each other, we can isolate these systems from exposure to the risks that come from a modern connected infrastructure and avoid having that legacy technology hold back modern architectures. It also allows monitoring of the environment that hosts operational technology and Internet of Things (IoT) devices to be highly focused on detecting and responding to unusual activity in an environment where it may not be possible to install software on the system.

Integrate cybersecurity into business decision making

Now that technology is an essential element of business operations—a process that has only been accelerated over the last 18 months—cybersecurity must be a factor in overall business decision making and not just something that is left to the technology department.


Treat cyber as a business risk

Cybersecurity should no longer be viewed as a specialized risk that falls only within the purview of the IT department. Technology expertise sits in the IT department, just as expertise in financial risk management generally resides in the finance department, but ultimate responsibility and

accountability for the risks lie within the business function. As the chapters in this report illustrate, addressing the threats that we face require a mix of technology, policy, and people expertise—as does all business decision making.

Every leader in the organization should consider how they enable employees and customers to have a better digital experience, while also considering what's needed to mitigate the associated risks. This includes consulting the cybersecurity team on how to manage the risks that arise as they undergo digital transformation. As this report shows, there are inherent risks in all the new technologies and business practices we adopt, and those risks must factor into any decisions about technology, policy, or business practices.

Cybersecurity's role in digital transformation



Security decision makers should truly embrace a risk management mindset as they consider how the steps they take can protect the organization, while also helping achieve operational goals.

Resilience includes cybersecurity

In the connected world we live in today we need to consider resilience as a key success factor in everything we do. Digital transformation is bringing increasing complexity to our security solutions including greater collaboration with third parties and the expectation that systems will be available 24x7. The platforms we are building to support businesses need to be fully resilient against attacks. Cybersecurity and resilience should be considered together. Operational resilience planning should include understanding the cybersecurity threats to the system and making appropriate investment to ensure continued success.

It is crucial to implement strong backup and recovery solutions, but it is equally important that organizations plan for how operational decisions will be made in the event of a cyberattack and practice their crisis management and response as well as their technical response to incidents.

Build a third-party risk program

Partners, suppliers, and contractors interact with data and applications connected to our corporate environment all the time. Attackers are increasingly targeting third-party providers to gain access to their systems and networks with a view to gaining access to their customers.

Ensure that the organization has a strong supply chain assurance process, built on an understanding of your suppliers' exposure to cyberattacks, how they configure their systems to be secure, and what steps they take to protect any information you share with them. Ensure that you are managing your third-party risk through robust service-level agreements, attestations, and shared assessments like SSAE 18 SOC 1 and SOC 2, PCI-DSS, GDPR, and ISO 20001.

Third-party access to systems should also follow the Zero Trust principles you apply to your own organization to limit exposure to attacks originating from a compromise of their systems.

Use digital empathy in implementing security controls

As we connect more and more systems together, security can become more complex, but we need to ensure that we value diversity of skills, areas of expertise, work and learning style, and background, among other things. Do not expect or require everyone to be technology experts, in the traditional sense of the term, in order to engage with the security of these systems.

When you implement security controls, apply digital empathy to ensure that the controls you are providing consider the environment in which those using the system are working and allow them to easily engage with the environment. For example:

Invest in user training that educates and informs.

Implement security training that helps employees understand the risks they face and the best way they can help to protect the organization. This training should be ongoing and designed in a way that increases engagement. User training is not just a compliance activity but an essential part of the early detection and response to an attack. Ensure that the training you provide explains risks in the context of the employees' work, and provide the context and

tools they need to understand appropriate behavior, recognize attacks, and report unusual activity. A culture of enablement, trust, and engagement will significantly improve reporting and provide earlier warning of attacks.

Build security into productivity.

When you implement security controls, think about the impact on the experience of those using it, whether employees or customers. What is their background, expertise, and cultural experience? Any controls you put in place should consider the experience of those who do not have a background in technology. Is it intuitive, can it be understood, and does it fit into their workflow as naturally as possible? Too much friction without an understanding of why a control is in place can lead people to circumvent technology or ignore processes to get things done.

If, in addition to training people, we ensure that security fits into their working practices rather than those of a technology or cybersecurity professional, we increase the chance of their understanding the risks and taking the appropriate actions. Where possible, cybersecurity should be invisible to the user except where it can help nudge them to take appropriate actions to manage risk.

Conclusion

As technology becomes more integral to our society, attackers are increasingly seeking to exploit this cultural shift. From cybercriminals to nation state groups, these are sophisticated and well-researched organizations with the resources, investments, and research to deploy complex and well-informed attacks against an organization. They are professional enterprises with their own sophisticated supply chains and their own well-researched, well-engineered lures that seek to exploit the way your organization works.

As we increasingly do more of our work online, so do criminals and nation state attackers. You must take this realization into consideration as you plan digital activities. For any new venture, consider the threat alongside the opportunity, and think about how you can manage risk for your entire organization.

This kind of thinking will require fundamental changes in the way we operate. We must consider risk as a whole and across the organization, rather than within siloes or individual viewpoints. We must look at where we need to change the way we work, and where we need to do the things we are already doing—but better.

A number of key themes arise throughout the different sections of this report that we encourage you to consider as you think about improving your security posture:

- **Do the basics well.**

A running theme throughout many of the chapters is that, although attackers are becoming more sophisticated, good cyber hygiene and implementation of basic security measures is often the best way to disrupt, prevent, and detect their attacks.

- **Take a holistic view.**

Too often the way we organize security and risk is driven by our own organizational structure and siloes. Attackers will look for vulnerabilities across these siloes, so we need to consider risk and the best approach to mitigating risk at an organizational level. This may require some standardization or translation of approaches across the different teams in an organization. It also underlines the importance of standards as we seek to harmonize between companies, which is increasingly important to managing supply chain risk.

- **Any element can be used as an attack vector.**

Attackers will look for the weakest link across an organization’s ecosystem, so we must manage it holistically. The weakest link may be a connected freezer or building management system that is used to gain access to the corporate network, or it may be a user or device that is compromised via a phishing email in an attempt to gain access to the operational technology running a factory or production plant. We need to consider and manage the organization’s entire attack surface.

- **Think about people.**

People engage with technology and can be used as a way of gaining access to the digital environment. Think about how to engage with them in a way that will help them to understand the risks they face. Understanding, engaging, and educating people will allow them to become a key line of defense against modern threats, whether that is misinformation seeking to influence peoples’ decisions and undermine democracy or phishing emails seeking to gain access to and compromise an organization’s digital assets.

- **Zero Trust is an architectural principle.**

The threats we have seen underline the importance of Zero Trust in designing and managing the risk in an organization. The last year has emphasized why there should be no such thing as a trusted application, trusted user, or trusted device with unrestricted access. The risk and context of every connection needs to be considered before allowing access to resources. Zero Trust is not a technology but an approach to managing risk. When implemented properly, it can enable us to unlock the potential of modern technology while limiting our exposure in a hyperconnected world.

Contributing teams at Microsoft

Contributing teams at Microsoft

The insights in this report, as well as the actionable learnings above, have been provided by a diverse group of security-focused individuals, working across dozens of different teams at Microsoft. Collectively, their goal is to protect Microsoft, Microsoft customers, and the world at large from the threat of cyberattacks. We are proud to share these insights in a spirit of transparency, with a common goal of making the digital world a safer place for everyone.

Azure Networking, Core

A cloud networking team focusing on the Microsoft WAN, data center networks, and the software defined networking infrastructure of Azure. This includes the DDoS platform, the network edge platform, and network security products such as Azure Firewall and Azure WAF.

Cloud Security Research team

A team working to secure the Microsoft cloud and build security products, with a mission to protect and empower customers to securely transform their organizations. The team's focus is on research and feature productization for Azure Defender, Security Center, and Azure Sentinel.

Customer Security and Trust (CST)

A cross-disciplinary team driving continuous improvement of customer security in our products and online services. Working with engineering and security teams across the company, the mission of CST is to ensure compliance and enhance security and transparency to protect our customers and promote global trust in Microsoft. They formulate and advocate cybersecurity policy globally; advancing Digital Peace through multistakeholder collaboration, focusing on Digital Safety to protect customers from harmful online content, and collaborating with public and private organizations to disrupt cyberattacks and support deterrence efforts.



Contributing teams

Cyber Defense Operations Center (CDOC)

Microsoft's cybersecurity and defense facility is a fusion center that brings together security professionals from across the company to protect our corporate infrastructure and the cloud infrastructure to which customers have access. Incident responders sit alongside data scientists and security engineers from across Microsoft's services, products, and devices groups to help protect, detect, and respond to threats 24x7.

Defending Democracy Team

A Microsoft team who works with stakeholders including governments, nongovernmental organizations, academics, and industry all in democratic countries globally to protect campaigns from hacking, increase political advertising transparency online, explore technological solutions to preserve and protect electoral processes, and defend against disinformation campaigns.

Detection and Response Team (DART)

A Microsoft team whose mission is to respond to security incidents and help Microsoft customers become cyber-resilient. DART leverages Microsoft's strategic partnerships with security organizations around the world and with internal Microsoft product groups to provide the most complete and thorough investigations possible. DART's expertise has been leveraged by government and commercial entities around the world to help secure their most sensitive, critical environments.

Digital Security & Resilience (DSR)

A Microsoft organization developed with a mission to enable Microsoft to build the most trusted devices and services, while keeping our company safe and our data protected. Across the company, DSR is continually evolving the security strategy and taking actions to protect Microsoft assets and the data of our customers.

Digital Security Unit (DSU)

A team of cybersecurity attorneys and strategic cyber intelligence analysts who provide legal, operational, geopolitical, and technical subject matter expertise to protect Microsoft and our customers. DSU's analysis and proposed solutions to complex digital security problems help to build trust in Microsoft's enterprise security capabilities and defenses against advanced cyber adversaries worldwide.



Contributing teams**Enterprise Risk Management (ERM)**

A team focused on key risks to Microsoft's business objectives, the ERM team works across business units to enable prioritized risk discussions with Microsoft's senior leadership and, ultimately, the Board of Directors. The team manages the company's NIST Cybersecurity Framework internal assessment and the enterprise risk framework, which connects to multiple operational risk teams, and coordinates with the company's internal audit function.

GitHub Security Lab

An open-source software-focused security research team. Its mission is to help secure the world's code and build bridges between the security research and software development communities through contributions including security research, tooling, and meetups.

Global Cybersecurity Policy

A team that works with governments, NGOs, and industry partners to promote cybersecurity public policy that empowers customers to strengthen their security and resiliency as they adopt and use Microsoft technology.

Microsoft AI, Ethics and Effects in Engineering and Research (AETHER) Committee


An advisory board at Microsoft that helps to ensure that new technology is developed and fielded in a responsible manner.

Microsoft Customer and Partner Solutions

Microsoft's unified commercial go-to-market organization responsible for field roles such as security and technical sales specialists and advisors.

Microsoft Defender for IoT

A team composed of domain-expert reverse engineers and data scientists. The team continuously performs reverse-engineering and analysis of large amounts of data related to IoT threats and threat actors to gain better visibility into the IoT landscape and uncover related trends and campaigns.



Contributing teams

Microsoft Defender Team

Microsoft’s largest global organization of product-focused security researchers, applied scientists, and threat intelligence analysts. The Defender Team delivers innovative detection and response capabilities in M365 security solutions and Microsoft Threat Experts.

Microsoft Digital Crimes Unit (DCU)

A team of attorneys, investigators, data scientists, engineers, analysts, and business professionals that fight cybercrime globally through the innovative application of technology, forensics, civil actions, criminal referrals, and public/private partnerships while protecting the security and privacy of our customers.

Microsoft Security Response Center (MSRC)

Part of the defender community on the front line of security response evolution. For over 20 years, MSRC has been engaged with security researchers working to protect customers and the broader ecosystem. An integral part of Microsoft’s Cyber Defense Operations Center (CDOC), MSRC brings together security response experts from across the company to help protect, detect, and respond to threats in real time.

Microsoft Threat Intelligence Center (MSTIC)

Microsoft’s centralized team focused on identifying, tracking, and collecting intelligence against the most sophisticated and advanced adversaries impacting Microsoft customers, including nation state threats, malware, phishing, and more. The threat intelligence analysts and engineering teams in MSTIC work closely with Microsoft security product teams to both develop and refine high-quality detections and defenses across our security product portfolio.

Security, Compliance, and Identity Business Development Team

A team that supports Microsoft security product teams in providing market insights into the latest cybersecurity trends to inform product development decisions. The team works on building partnerships with independent software vendors working with Microsoft’s security ecosystem.



