Orange
Cyberdefense

Security Intelligence | Research Report

Monthly Report
July 23

www.orangecyberdefense.com

Contents

Introduction

Downfall

Downfall attacks, as presented at Blackhat by Google senior research scientist Daniel Moghimi, are a new type of side-channel attack that exploit a vulnerability in speculative execution in Intel processors. The vulnerability, CVE-2022-40982, affects Intel processors from the sixth generation Skylake series to the 11th generation Tiger Lake chips.

Downfall attacks work by exploiting the way that Intel processors use speculative execution to speed up the execution of code. When a processor is speculatively executing code, it may load data into its vector registers even if the code turns out not to be executed. This data can then be leaked to an attacker through a side channel.

In the case of Downfall attacks, the data that is leaked is the content of the vector registers. This data can be used to steal sensitive information from other users who share the same computer, such as passwords, encryption keys, and private data.

Inception

AMD CPU’s have also been found to be leaking data, the Inception attack is a new type of transient execution attack that can be used to leak sensitive data from all AMD Zen CPUs.

Transient execution attacks exploit a feature of modern CPUs called speculative execution. Speculative execution allows CPUs to guess what instructions will be executed next, and then execute those instructions even if they are not actually needed. This can improve performance, but it also opens up a security vulnerability.

As with Downfall above, the Inception attack can be used to leak sensitive data from anywhere in the memory of a computer powered by an AMD Zen processor. This includes data such as passwords, encryption keys, and other confidential information.

Electoral Commission Breached

The Electoral Commission has issued a public notification, confirming that due to what it termed a "complex cyberattack" the personal data of up to 40 million UK voters has been compromised by unknown "hostile actors”.

In the notification the election watchdog confirmed that it became aware of the attack in October 2022, but that the hostile actors had been able to access its systems from August 2021. The attack was reported to the Information Commissioners Office (ICO) and the National Cyber Security Centre (NCSC)

The attackers had access to email servers, control systems, and copies of the electoral registers. This means the unknown attackers would have been able to access the full names and addresses of all registered voters who registered between 2014 and 2022, as well as the names of overseas voters.

At a glance

The Police Service of Northern Ireland (PSNI) has disclosed details of a major data breach that exposed sensitive information about all serving police officers. A member of the public made a Freedom of Information (FoI) request about officer rank and staff grades, which led to the breach. Mistakenly a large Excel spreadsheet was shared that had the last names and initials of all current employees, as well as where they work and in what department.

World Watch Review

The Orange Cyberdefense CERT published a total of 14 new World Watch advisories during July 2023, along with adding updates to a further 24 previously published advisories.

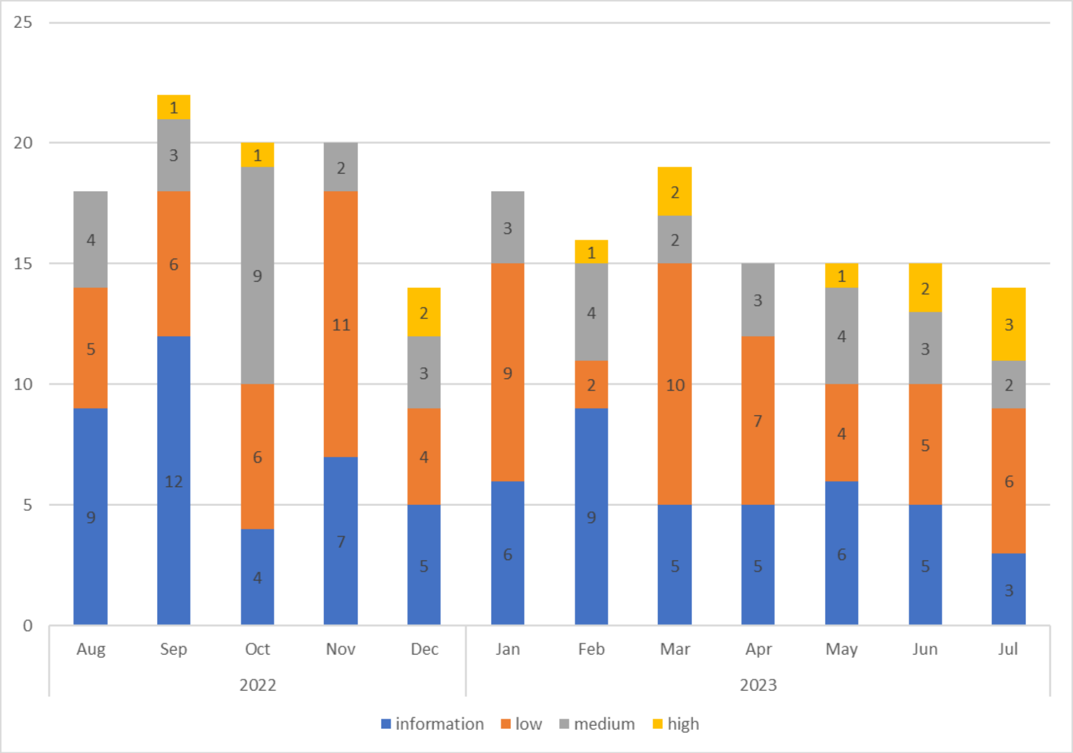

Breakdown of Published Advisories Previous 12 Months

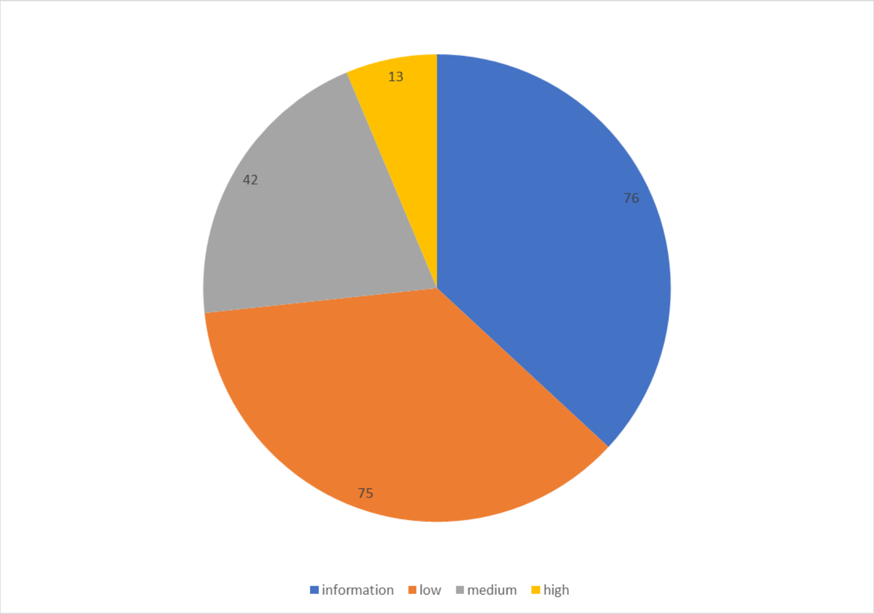

Breakdown of Advisory Criticality for Previous 12 Months

Advisory Summary

In July, as you can see in the first chart above, we had 3 advisories rated as high criticality with all other advisories given criticality ratings of low, medium or information when initially published. These ratings are based on our CERT's assessment of the risk and threat levels associated with the subject of the advisory at the time of publication, so even though an advisory may concern a vulnerability rated as critical by the vendor we may deem it to only initially be medium, if say there is no publicly available exploit. This is under constant monitoring however and subsequent updates will increase our criticality level as required if circumstances should change.

See below for a timeline of advisories rated Medium and higher:

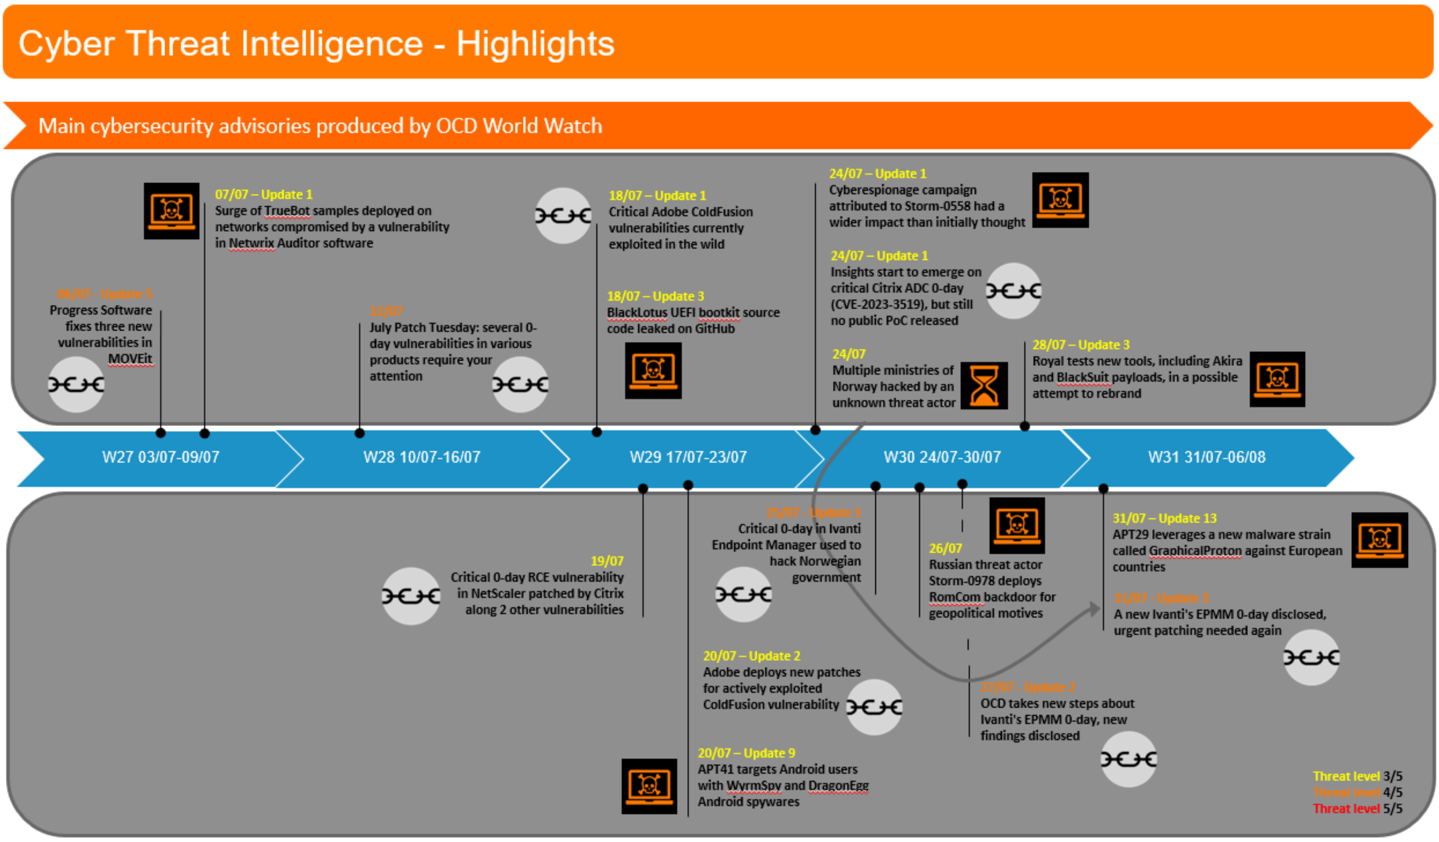

Cyber Threat Intelligence - Highlights

Main cybersecurity advisories produced by OCD World Watch

06/07 - Update 5
Progress Software fixes three new vulnerabilities in MOVEit

07/07 – Update 1
Surge of TrueBot samples deployed on networks compromised by a vulnerability in Netwrix Auditor software

12/07
July Patch Tuesday: several 0-day vulnerabilities in various products require your attention

18/07 – Update 1
Critical Adobe ColdFusion vulnerabilities currently exploited in the wild

18/07 – Update 3
BlackLotus UEFI bootkit source code leaked on GitHub

24/07 – Update 1
Cyberespionage campaign attributed to Storm-0558 had a wider impact than initially thought

24/07 – Update 1
Insights start to emerge on critical Citrix ADC 0-day (CVE-2023-3519), but still no public PoC released

24/07
Multiple ministries of Norway hacked by an unknown threat actor

28/07 – Update 3
Royal tests new tools, including Akira and BlackSuit payloads, in a possible attempt to rebrand

W27 03/07-09/07

W28 10/07-16/07

W29 17/07-23/07

W30 24/07-30/07

W31 31/07-06/08

19/07
Critical 0-day RCE vulnerability in NetScaler patched by Citrix along 2 other vulnerabilities

25/07 - Update 1
Critical 0-day in Ivanti Endpoint Manager used to hack Norwegian government

20/07 – Update 2
Adobe deploys new patches for actively exploited ColdFusion vulnerability

20/07 – Update 9
APT41 targets Android users with WyrmSpy and DragonEgg Android spywares

26/07
Russian threat actor Storm-0978 deploys RomCom backdoor for geopolitical motives

27/07 - Update 2
OCD takes new steps about Ivanti's EPMM 0-day, new findings disclosed

31/07 – Update 13
APT29 leverages a new malware strain called GraphicalProton against European countries

31/07 - Update 3
A new Ivanti's EPMM 0-day disclosed, urgent patching needed again

Threat level 3/5
Threat level 4/5
Threat level 5/5

Editor's Notes

Here the team will provide commentary on a news item, expansion on something specific such as a single incident, or could be as expansive as coverage on trends observed in the threat landscape.

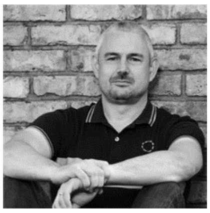

Carl

Everything* you need to know about ERP solutions and cybersecurity

*Ok, definitely not everything!!

Introduction

One of the biggest problems we face as researchers is identifying the relevant questions to ask. This makes curiosity a particularly important trait to have. In addition, you have to be aware of your own deficiencies and open to the fact that there are always going to be things you or your team do not know.

This became more than apparent to us recently when asked by a colleague what content we had around the subject of ERP security, to which we had to answer that we had nothing!! This post is now an attempt to rectify that, at least at a very high level, as after all we are not ERP experts.

What is an ERP solution?

An Enterprise Resource Planning solution, also known as an ERP solution, is a set of software programmes that is designed to manage and integrate a company's core business processes, such as financial management, supply chain management, human resources management, customer relationship management, and inventory management, into a single unified system.

ERP systems usually work from a centralised database, thus allowing all departments within an organisation to get access to the information contained inside the system. Because of this, they can simplify and automate their procedures leading to lower operating costs and boosting their overall efficiency.

Who are the major participants in the enterprise resource planning space?

On the market today, there is a wide selection of ERP providers, each of which offers a bespoke collection of features and options. Some of the top vendors/products in the Enterprise Resource Planning (ERP) space include the following:

- SAP is one of the most well-known and widely used enterprise resource planning (ERP) software suppliers in the world. They provide an extensive selection of ERP solutions that are geared towards a variety of business sectors and company sizes.

- Oracle is another industry-leading ERP vendor that provides solutions for a variety of businesses, including retail, healthcare, and manufacturing, amongst others.
- Microsoft Dynamics is a suite of enterprise resource planning (ERP) tools that can interface with other Microsoft products, such as Microsoft Office. It is geared towards both micro and medium-sized companies.
- Infor is a provider of ERP solutions that are tailored to various industries, like the healthcare industry, the manufacturing business, and the hotel industry.
- Epicor provide ERP systems for the manufacturing, wholesale distribution, retail, and service industries respectively.
- Sage is a provider of business management software and services, including ERP solutions for companies of all sizes, with a particular focus on those in the small and medium company sectors.

What different kinds of components does an ERP solution have?

An ERP solution is often made up of several different components that, when combined, form a comprehensive management system for a company's many business activities. These parts include the following:

- **Core ERP system**: This is the key component of the ERP solution that provides capability for managing core business processes, such as financial management, inventory management, supply chain management, human resources management, and customer relationship management.
- **Database**: An enterprise resource planning (ERP) system would normally make use of a centralised database to store the information it requires. This includes data on customers, suppliers, inventories, transactions, and any other business-related information that may be relevant.
- **Business intelligence and reporting**: ERP solutions frequently contain reporting and analytics capabilities that enable users to generate reports, analyse data, and get insights into business operations. These tools enable users to achieve a competitive advantage in their respective industries.
- **Integration modules**: Enterprise resource planning (ERP) solutions frequently feature integration modules that make it possible for the system to interface with other business applications including customer relationship management (CRM) software, e-commerce platforms, and supply chain management systems.
- **Tools for customisation**: Enterprise resource planning (ERP) solutions frequently contain tools for customization, which provide users the ability to customise the system so that it better meets their unique company requirements. This may involve user interface modifications, procedures, and specialised fields.
- **Access restrictions, data encryption, and auditing capabilities** are some examples of the security features that are included in ERP packages. These features are included to help safeguard important corporate data.

These components, when combined, work together to offer a comprehensive system for the management and automation of business processes, the enhancement of operational efficiency, and the provision of greater insight into business operations.

What kind of risk management model should be used for an ERP solution?

An enterprise resource planning (ERP) solution is an essential system that is used to handle and store sensitive corporate data, making it an attractive target for cyber criminals. As a result, a thorough threat model needs to be developed in order to detect potential dangers and openings in the ERP system's defences that could result in a security breach.

The following is a list of some of the most important aspects that make up a risk management model for an ERP solution:

- Identify prospective adversaries who could attempt to attack the ERP system, such as cybercriminals, hackers, company insiders, and other businesses in your industry.
- Identify possible vectors that attackers could use to exploit weaknesses in the ERP system. Some examples include vulnerability exploitation, phishing, and social engineering attacks.
- Identify potential vulnerabilities that may be exploited by attackers, such as unpatched software, weak passwords, unsecured network connections, or unsecured endpoints.
- Develop and put into place effective security controls to avoid and mitigate potential threats, such as encryption, access controls, monitoring, and frequent security audits. These methods are referred to as mitigation strategies.
- It is also necessary to determine the applicable compliance rules and regulations that relate to the ERP solution, such as GDPR, HIPAA, or PCI DSS, and to ensure that the system complies with these standards.
- Determine what the potential consequences of a successful attack on the ERP solution would be, such as monetary loss, damage to the company's reputation, or legal liability, and devise a strategy for how to react when security breaches occur.

Are there any threats that are unique or specific to ERP solutions?

There are, without a doubt, several dangers that are exceptional to or peculiar to ERP systems. These include the following:

- **Inaccurate data**: For proper operation, ERP solutions require data that is both accurate and complete. It is possible for there to be errors, inefficiencies, and security concerns if there is missing data, incomplete data, or erroneous data.
- **Complexity**: ERP solutions can be notoriously difficult to manage and secure due to their high degree of complexity and the level to which they can be customised to specific business needs. It can be difficult to identify and mitigate risks due to these complexities, which may be made worse by the introduction of security vulnerabilities caused by customisations.
- **Third-party integrations**: ERP solutions frequently integrate with third-party systems and applications, which can present a security risk if the integrating systems are not well secured.
- **Insider threat**: ERP systems have the potential to be subject to insider attacks, which can include personnel who abuse their access, purposefully create security flaws, or even just inadvertently make mistakes.
- **Attacks from the internet**: If exposed and not properly secured, ERP solutions are a desirable target for attacks from the internet, leading malware infection, Cyber Extortion and the theft of sensitive or valuable corporate data.

It is essential for businesses to implement a comprehensive security strategy that includes regular security assessments, vulnerability testing, employee training, and stringent access controls in order to reduce the impact of these unique threats and risks and protect their ERP solution from security breaches.

What does security look like for ERP solutions, and what components are involved in providing that protection?

The security of ERP solutions is comprised of several different components, all of which collaborate to safeguard the system from various dangers and weaknesses. These parts include the following:

- **Access controls**: Access controls are used to ensure that the ERP system is only accessible to people who have been specifically authorised to do so. This includes features such as multi-factor authentication, strong passwords, and role-based access controls, and others.
- **Encryption**: Encryption is utilised for the purpose of protecting sensitive data both while it is in transit and while it is at rest. This includes precautions such as encrypting network communication with SSL/TLS and encrypting sensitive data stored in the database.
- **Network security**: The ERP system should be protected from network-based attacks and denial-of-service attacks by using network security measures. Firewalls, intrusion detection and prevention systems, and

routine network scans are some examples of the security measures that fall under this category.

- **Patch management**: Conducting routine patch management is essential for mitigating vulnerabilities and ensuring that the ERP system is always current with the most recent security patches and updates.
- **Employee training**: It is vital to train employees in order to guarantee that users understand the dangers and best practises for using the ERP system. The risks and best practises they need to be aware of include maintaining good password hygiene, data classification, and recognising social engineering attempts.
- **Auditing and monitoring**: Both auditing and monitoring are essential for identifying security breaches in a timely manner so that appropriate action can be taken. The ERP system should be subjected to logging, auditing, and monitoring in real time as part of these precautions.
- **Disaster recovery & business continuity**: It is essential to have an incident response plan in place including planning for disaster recovery and maintaining business continuity. This ensures that the ERP system can be swiftly recovered following a security incident or outage and minimises the impact that this has on business operations.

It should also be noted that there are third party vendors, such as Onapsis & Safe O'Clock, providing security solutions specifically for some of the ERP platforms, although primarily aimed at SAP & Oracle. These solutions, as well as being able to help with some of the above best practice configuration items, also offer some of the following capabilities:

- Assess – Scan for vulnerabilities, unnecessary or dangerous services and misconfigurations.
- Detect/Defend & Respond – Ongoing threat monitoring generating alerts for unauthorized changes, anomalous user behaviour or cyberattacks. Integration with Incident Management & SIEM solutions.
- Compliance – Automated checks to ensure compliance with a variety of security and regulatory frameworks and controls.

In summary

Like most cybersecurity, there's no magic here. Integrated cloud / on-premise deployments increase the complexity and indirect risk, but for the most part ERP security is achieved by consistently getting the security of the composite elements right. The biggest challenge for security managers is likely to be identifying what and where these components are, an exercise that will become increasingly more difficult as the complexity of the system increases.

Wicus

Bypassing Authentication Storm

A look back at SolarWinds

In December 2021 we published a blog post on the Golden SAML (Security Assertion Markup Language) attack against Microsoft Active Director Federated Service (ADFS)[1]. That blog post was inspired by the larger event that impacted many companies in late 2020 and pivoted mostly around the cyberattack of SolarWinds. Our blog post looked at how attackers targeted ADFS to steal trusted signing certificates. These certificates enabled the attackers to forge digital signatures that enabled them to gain authenticated access to services that had a trust relationship with the compromised ADFS, such as cloud services. With a legitimate username the attacker can now forge a valid authentication response that any third-party service (also known as a resource partner organization) in an existing trust relationship with ADFS will accept [2]. A password or MFA is not required, and the attackers bypassed the authentication mechanism.

The latest STORM

More recently, an announcement by Microsoft that an attacker, tracked as STORM-0558, managed to gain unauthorized access to Exchange Online data hosted in Azure by using Outlook Web Access (OWA) has caused quite a stir [3]. Microsoft indicated that the attackers targeted a subset of accounts belonging to specific organizations. Microsoft also indicated that the attack path was closed when they revoked several certificates associated with what is called Microsoft Account (MSA) consumer signing keys and fixed a validation flaw on their end. At the time of writing Microsoft admitted that they could not explain how these attackers obtained a copy of a private key of an MSA certificate used in the attack and were still investigating the matter. Microsoft indicated that this inactive MSA key enabled attackers to fool the Relying Party (RP) process that checks authentication token signatures, as the forged authentication token was signed by the trusted certificate.

A company called Wiz.io then published a blog in which they shared their views of some of the technical aspects of the attack, including the types of accounts that could be impacted by this type of attack [4]. At the end of the blog post Wiz indicated that a Microsoft team reviewed their blog to ensure it is “technically correct”. It is not clear if “technically correct” means that Wiz’s hypothesis about the attack vector is accurate, or if the technical details of elements of the MSA

[1] https://www.orangecyberdefense.com/global/blog/cloud/exploring-the-golden-saml-attack-against-adfs

[2] https://learn.microsoft.com/en-us/windows-server/identity/ad-fs/technical-reference/understanding-key-ad-fs-concepts#ad-fs-terminology-used-in-this-guide

[3] https://www.microsoft.com/en-us/security/blog/2023/07/14/analysis-of-storm-0558-techniques-for-unauthorized-email-access/

[4] https://www.wiz.io/blog/storm-0558-compromised-microsoft-key-enables-authentication-of-countless-micr

account as per technical documentation is correct. Irrespective, it seems clear that a team from Microsoft reviewed the blog post and did not object to the implication of the content. However, this is still not official Microsoft communication regarding technical aspects of the respective compromise, and the Wiz analysis is still only plausible speculation.

The MSA certificates revoked by Microsoft were associated with public services on Azure and enabled user access to applications hosted on Azure. The MSA certificates are used to ensure the integrity of Azure Active Directory (AAD) issued tokens. These applications need to follow prescribed practices detailed by Microsoft to ensure correct functioning behavior, including various validation steps such as validating the token authenticity using a certificate. As stated earlier, there was a flaw that allowed an unintentional side effect, namely granting access to data of unrelated parties. The flaw fixed involved absent validation of the Issuer specified in the certificate information (Issuer claim). The specified Issuer was not compared with the logically associated Issuer for the respective Azure tenant. In other words, each Azure tenant has their own unique certificates issued by Azure and there is an identification value that must be checked as part of the authentication process to ensure that the certificate is used in the correct context and is associated with the tenant listed in the certificate's issuer claim.

The MSA certificates were trusted in certain contexts and, combined with the validation flaw, enabled the attackers to gain implicit access due to the absence of additional verification steps. Another limiting factor for the attackers, according to the Wiz blog, was the attackers had to tailor their attack to focus on multi-tenant or mixed audience applications. Single tenant applications could not be abused in the same way. This limitation was hardly relevant, however, as the accounts the attackers were targeting fitted this profile.

The cloud may be safer, but don't Jump the gun

Another cyber-attack that made headlines around the same time was the breach of JumpCloud, a cloud-based identity provider [5]. JumpCloud detected the compromise and opted to force an admin API key rotation for all clients [6]. The admin API keys could potentially allow an attacker to further compromise JumpCloud's clients environments by abusing the privileged API access. Mandiant published a blog describing how the attackers compromised one of the JumpCloud victims by using the privileged API key access [7]. The attackers then used the JumpCloud agents deployed on the endpoints to push malicious scripts to further their actions on objectives.

Old man shouts at three clouds

[5] https://jumpcloud.com/blog/security-update-incident-details
[6] https://jumpcloud.com/support/mandatory-jumpcloud-api-key-rotation
[7] https://www.mandiant.com/resources/blog/north-korea-supply-chain

All three incidents described above involved stolen authentication key material that enabled attackers to gain access to respective infrastructure or data [8]. In the ADFS Golden SAML attack the attackers had to extract key material per compromised organization, thus multiple victims had to be probed for the specific services and then the attackers could pivot from there. In the JumpCloud instance the attacker had to compromise JumpCloud itself that held all the privileged admin API keys and then had to pivot to the respective organizations based on their API keys. In the MSA STORM-0558 incident, the attacker had to acquire one or more of the eight MSA private keys and then target user accounts that might be associated with multi-tenant or mixed user applications. From the disclosed information at least Exchange Online data was accessed without authorization.

In the case of the MSA incident, it is unclear if the attackers simply got lucky regarding the flaw relating to the missing check on the certificate issuer claim value. The attackers might not have known about the validation flaw and probed at the Azure applications using the stolen certificate and accidentally achieved access. Alternatively, the attackers might have known about the absence of the checks and knew that if they got hold of a MSA private key they could get in. According to Wiz, Azure multi-tenant application must be configured in a way that allows the MSA certificates to be usable. This constitutes a misconfiguration according to Wiz, but Microsoft's documentation is not clear on this even though it explicitly mentions issuer validation [9]. As with many vulnerabilities, several stars had to align for this exploit to be feasible.

Both JumpCloud and Microsoft shared high level details about the respective cyberattacks. In the JumpCloud incident we know the starting point was a spear phishing attack, but beyond that and the Mandiant blog not much else is publicly known. In the case of the MSA incident, Microsoft say they are still investigating how a MSA private key got leaked. Besides the Wiz.io blog, we do not know if any other techniques were used by the attackers [10].

Using a single certificate as part of an authentication process can be a costly mistake, especially in the ADFS Golden SAML and MSA attacks. Grouping the Golden SAML and the MSA attacks under the same umbrella seems like a stretch, but these two attacks are more alike in nature than the JumpCloud incident. The Golden SAML attack can be used to forge authentication tokens for multiple users associated with that authentication domain. The impact of the stolen MSA private keys is even larger as it is not limited to just one organization, while the Golden SAML attack is limited to a specific Active Directory domain.
All three incidents resulted in the potential to pivot into other environments to further the goals of the attackers. The extent to which attackers could move around seems to be varied, with the MSA incident scoped to specific Azure based applications. In the Golden SAML attack and JumpCloud incidents attackers

[8] Another commonalty not touched on is that all three incidents are attributed to well-resourced groups with known government ties.
[9] https://learn.microsoft.com/en-us/azure/active-directory/develop/access-tokens#validating-tokens
[10] As on July 26, 2023

could get free reign to a variety of services and even devices (as in the JumpCloud incident).

An alternative user authentication?

The Fast Identity Online (FIDO) Alliance is an open organization that seeks to reduce password dependency and has been publicly operating since 2013. The FIDO standard has given us a phishing resistant authentication process by building on the advances made in web browser security as well as the adoption of the cryptographic network protocols like Transport Layer Security (TLS).

The most recent FIDO standard, FIDO2, is based on public key cryptography, which is what ADFS and AAD MSA authentication rely on as well. This makes one wonder whether FIDO2 may also be susceptible to this kind of private key compromise. In the broader sense, the short answer is ‘no’.

Comparing FIDO2 with ADFS and AAD MSA is not an apples-to-apples comparison either. ADFS’ intended purpose is to provide a federation of trust for parties that wish to authenticate several applications using Active Directory credentials. ADFS cannot vouch for the actual credentials as it only indicates by means of trusted signature if the provided credentials, which another party verifies, are valid. The AAD with MSA authentication is a better comparison but falls short on the non-repudiation principle because a single public/private key pair is used across all accounts.

FIDO2 describes 16 Security Goals (SG) and 29 Security Measures (SM), each with their own unique features that contribute to the strength of the standard [11] [12]. Each SG has a SM mapping that helps us further understand the design strengths of FIDO2 [13]. To address the question of whether FIDO2 is susceptible to private key theft and token forgery we can look at SG-5 - Verifier Leak Resilience and SG-6 - Authenticator Leak Resilience in the SG – SM mapping table.
FIDO2 makes private key theft less feasible and more difficult since each user has their own private key for each identity provider login [14] [15]. FIDO2 requires a type of secure enclave, which is a hardware component that complies with several requirements for hardware attestation [16]. Strong cryptographic capabilities are also mandated that increase the difficulty level for attackers to clone authenticators or exploit potentially weak cryptographic algorithms [17]. A possible, but more expensive option is for attackers to physically gain access to the FIDO2 hardware device by physical theft, for example.

[11] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#h2_fido-security-goals
[12] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#h2_fido-security-measures
[13] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#h3_relation-between-measures-and-goals
[14] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#dfn-sm-2
[15] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#dfn-sm-6
[16] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#dfn-sm-9
[17] https://fidoalliance.org/specs/fido-v2.0-id-20180227/fido-security-ref-v2.0-id-20180227.html#dfn-sm-16

FIDO2 and ADFS

From documentation provided by Microsoft and Yubico, a FIDO2 compliant hardware authenticator vendor, AAD, must be used directly as ADFS is incapable of handling FIDO based authentication [18] [19] [20]. Microsoft has also indicated that their Azure Multi-Factor Authentication Server (MFA Server) will be deprecated at the end of September 2024 [21] and that any new deployments of MFA Server will not be possible. Microsoft Entra ID (the new name for Azure Active Directory) will be the new solution to handle MFA or passwordless authentication in the future [22].

Windows Hello for Business with FIDO2 support is another approach that can accommodate cloud and hybrid deployments where passwordless and FIDO-based authentication are used[23]. ADFS can still be present in environments with Windows Hello and FIDO, but the interactions will require additional authentication considerations for Single Sign On (SSO).

The world on its shoulders

The World Wide Web (WWW) has changed significantly since the mid 1990's. A combination of several key concepts such as the Domain Name Service (DNS), public key cryptography, secure Hyper Text Transport Protocol (HTTPS) with Transport Layer Security (TLS), and universally accepted web standards is today manifest in the form of a web browser. This has made it possible to navigate the web with a degree of certainty that would never have been possible before. Without this assurance the Internet as we know it today would be a minefield or toxic wasteland.

On the Internet it is important to authenticate not just the website but also the user's identity. But how does the website know which user is trying to access it? It is possible to use public key cryptography, like how websites are validated, by issuing a certificate for each registered user. However, this is difficult to execute in practice using traditional approaches, due to the additional enrollment requirements imposed on end users. A much simpler approach has thus been widely adopted, namely the ubiquitous combination of username and password. Supplying a username and password during a login process enables the website to authenticate the user. The authenticated user is issued a session token, also referred to as a session cookie, which is handled by the web browser seamlessly. We know that the username and password are very portable and can be stolen or leaked.

Multi-Factor Authentication (MFA) was introduced to make it more difficult for attackers to abuse the stolen credentials, as out-of-channel authentication components such as push notifications to mobile authenticators are much more difficult to compromise at scale. Attackers thus adapted their techniques to performing Person-In-The-Middle (PITM) attacks that focused on stealing the session cookie directly. With this session cookie the attacker can now act as that authenticated user, even though they do not have the username, password, or additional authentication factors.

Mutual authentication using public key cryptography, when combined with cryptographic concepts to protect against replay attacks, makes it more difficult for attackers to perform a PITM style attack. Passwords and MFA with session cookies do not provide those types of protections.

The FIDO2 standard was created to guard against all known attack scenarios that could result in credential theft. It even has protection against PITM style attacks where the base URL of the host requesting authentication is included in the verification process. In the case of a classic PITM based attack, the authentication process cannot be concluded because the party that sits between the victim and spoofed website must have a distinct domain name and certificate identical what was used during the initial FIDO registration process. Does FIDO2 eliminate the associated problem with portable session cookies? The answer is 'yes', as a session must still be maintained to identify the user with the website for each interaction. Now attackers must find ways to inject themselves into the browser through other means, such as malicious browser extensions or vulnerabilities, where the latter is much more difficult to achieve than the former. Another approach would be to gain physical access to a host and then try to dump the browser cache containing the session cookies. It is unclear what risk the use of classical web proxies with valid TLS certificates may pose with approaches such as FIDO, but such proxies have caused an outcry in the past due to the privacy risk and inherit security concerns introduced by their use [24].

Think of the traditional session cookie as the titan Atlas from Greek mythology, who carries the earth or sky, depending on which version of the story you prefer, on his shoulders [25]. Like the titan, the session cookie is responsible for a user's identity and carries the weight of that responsibility on its shoulders. One day someone will come to take away that burden, like Heracles who built the Pillars of Hercules to do so.

The bark is worse than the byte

A PITM Attack against FIDO2 is conceptually feasible but appears hard to execute in practical terms. There is an edge case where an attacker, in control of the DNS lookup mechanism, poisons the DNS cache of a victim that resolves to a spoofed

[18] https://support.yubico.com/hc/en-us/articles/8315620915996-Phishing-Resistant-Authentication-for-hybrid-environments-with-AD-and-Azure-AD-using-FIDO2
[19] https://learn.microsoft.com/en-us/azure/active-directory/authentication/howto-authentication-passwordless-security-key-on-premises
[20] https://learn.microsoft.com/en-us/azure/active-directory/authentication/how-to-migrate-mfa-server-to-azure-mfa
[21] https://learn.microsoft.com/en-us/azure/active-directory/authentication/howto-authentication-passwordless-deployment#technical-considerations
[22] https://www.microsoft.com/en-us/security/business/identity-access/azure-active-directory
[23] https://learn.microsoft.com/en-us/azure/active-directory/authentication/howto-authentication-passwordless-security-key-on-premises
[24] https://www.eff.org/deeplinks/2015/02/dear-software-vendors-please-stop-trying-intercept-your-customers-encrypted
[25] https://en.wikipedia.org/wiki/Atlas_(mythology)

web site. The attacker would need a legitimate certificate for their spoofed web site, which is non-trivial. Having achieved this means that the base URL will match what the FIDO validation processes the browser follows will attest to. This scenario is possible if the attacker has control of a router and DNS service, but modern browsers have already implemented DNS over HTTPS (DOH) that makes it exceedingly difficult to interfere with. DOH is not always enforced especially if Wi-Fi captive portals are in play. A now defunct web standard called HTTP Public Key Pinning would also have been a valid mitigating strategy against spoofed digital certificates [26].

If the attacker can overcome the challenges of poisoning the DNS and obtain a digital certificate that will please the browser, then it might be possible to steal the mythical Atlas session cookie. Microsoft Hello for Business offers viable mitigation by storing the session cookie in the Trusted Platform Module (TPM) of the device running Windows. The session cookie is also renewed frequently and seamlessly, making this attack path more costly and less viable.

Does this cloud have a silver lining?

There is no doubt that in many cases the adoption of ‘cloud’ based systems for traditional use cases will improve the technical security posture of the business. There is a saying that “Nobody is qualified to configure and manage Microsoft Exchange except Microsoft themselves.” This is true in many similar cases, and many businesses will benefit from outsourcing technology systems and platforms to specialist operations in the cloud.

However, this reduction in technical risk comes at the cost of an increase of less obvious, non-technical risks, as follows:

1. The threat will adapt

   Since cybercrime and other threats are driven by powerful systemic factors, we see historically, and we can predict that attack vectors will adapt to changes in the technology landscape. Crime will go where the ‘money’ is and the hacking ecosystem will evolve its capabilities to be effective in an emerging ‘cloud centric’ world. We may not know how this will happen yet, but the attacks described in this post expose some probable future trajectories. Another example we have seen is the deployment of malicious insiders.
2. Homogeneity & Contagion
   a. Over the last two decades security and resilience has really suffered because of the ubiquity of the homogenous Microsoft desktop and server platforms, and the opportunities to specialize and scale that this presents attackers, i.e., an attack that works against one Microsoft system will work against all Microsoft systems. This dynamic persists

[26] https://en.wikipedia.org/wiki/HTTP_Public_Key_Pinning

and even accelerates as homogenous SaaS (Software as a Service) and PaaS (Platform as a Service) are adopted.

b. Furthermore, homogeneity also exacerbates contagion, i.e., the impact of vulnerability, attack or compromise can spread rapidly across interdependent environments, as we saw with WannaCry, notPetya and SolarWinds. The more we adopt homogenous cloud systems, the more we expose ourselves to this kind of contagion risk.

c. The more standardized platforms are and centralized, the less opportunity there is for alternative approaches to survive. As a result, we collectively actually lose our access to alternative technologies and approaches, which further reduces our resilience, for example by removing the option of 'falling back' in the case of a compromise or other failure. E.g. When M365 mail goes 'down', are we be able to find an on-prem mail server to replace it, and someone who knows how to use it?

3. Attack Surface Management

   Attackers have always understood that most compromises happen because they are able to find a system that is vulnerable, rather than because they find a vulnerability in a system. As cloud adoption grows businesses will face the growing challenges of understanding and managing their evolving attack surface. Rather than track and reduce internet-exposed IP addresses and Ports, they now must learn to manage ephemeral systems, complex user and role permissions, diverse storage locations, API keys, compliance, and geopolitical risks and the like. This is already leading to frequent non-technical compromises where data or capabilities are simply exposed onto the Internet for anyone to access.

4. Two pillars

   As SaaS, PaaS and other cloud-based systems become standardized, we are seeing an inevitable migration toward 'web applications' in which the code and data reside on a 3rd party cloud system and the rendering and UX are performed in a browser. That means that increasingly all the responsibility for security now rests on these two pillars – the cloud service provider and the browser vendor. These players have proven to be very capable in the past – and there is a lot of technical benefit to this approach – but from a system perspective we should be aware that the security and resiliency of cyberspace increasingly rests on just these two pillars.

5. Geopolitical Threats

   Although we think of the cloud as something 'ephemeral', it is in fact comprised of actual computers located in actual datacenters and managed by actual people. These factors are all linked to specific geopolitical realities and therefore under the influence of the political forces and powers that govern those places. This concept was illustrated when Russia took control of the Internet in occupied Ukraine simply by redirecting the traffic from those physical locations into systems located in Russia, to facilitate censorship and surveillance. As hacking and cybercrime become increasingly influenced by politics and power, the geopolitical context of a cloud-based system becomes increasingly important. This is especially true when one considers that political realities are increasingly volatile and can easily change within the lifetime of a technology platform. In other words, what's politically acceptable today may not be politically acceptable tomorrow. In adopting cloud-based systems and platforms, businesses must recognize that they are making themselves vulnerable to threats that may emerge when the political realities in the physical and political 'homes' of these platforms change.

6. Switching costs

   Subscription-based businesses models (as we predominantly observe in the cloud) are highly incentivized to make 'switching' difficult for the customer. This is not always apparent but is deeply baked into the business prerogatives of these offerings. Once a business has chosen to adopt cloud platforms and systems it may be very difficult to switch to alternative options. This represents an obvious risk but is also a risk to resilience.

7. Responsibility, accountability, and transparency

   Businesses should recognize that, while cloud providers may assume responsibility for certain elements of cybersecurity, the accountability for security failures will almost always still rest wholly with the businesses. In cloud offerings, where so many of the technologies, people and processes are obfuscated from the end user, this can make it very difficult for the client to understand and manage what their real risk is.

This can be illustrated by an article posted by Tenable CEO Amit Yoran in which Amit alluded to inherit risks of running on the cloud where the cloud vendor has critical vulnerabilities that could be exploited to gain access to tenants' data. Amit's post mentioned that a researcher at Tenable found a vulnerability in Microsoft Azure that allowed the researcher to gain access to a financial institution's cloud infrastructure. Tenable raised the issue with Microsoft in March 2023, but Microsoft has pushed the complete fix out to end of September 2023 [27]. From a risk point of view what does this mean? Who is responsible for any breach related to this? All we know is that Microsoft is aware of the weakness and hopefully they are keeping an eye on it.

Conclusion

What we know about the Golden SAML attack again ADFS, stolen JumpCloud admin API keys, and stolen MSA private keys, are that clever and well-resourced

[27] https://www.linkedin.com/pulse/microsoftthe-truth-even-worse-than-you-think-amit-yoran/

attackers have demonstrated that compromising authentication and signing keys are viable avenues of attack.

In the cloud age, attackers continue to try bypass, weaken, or abuse Identity Providers to gain access by pickpocketing the keys and walking in through any door. Properly handling sensitive key material will become more important than ever as Identity and Access Management are fundamental in the cloud computing era.

Account activity and account creation monitoring are therefore important means of detecting anomalous and malicious activity. Microsoft did provide free access to their Azure Purview Premium log auditing service for Azure clients to help identify suspicious account activity, but only after the incident caused public outcry [28]. It is not clear how long this courtesy will last.

Ric

Dead Man's PLC

We've been working on some interesting research for a few months, and although it is under review at a couple of places for publication, we recently released it on arxiv. That research is Dead Man's PLC[29], a novel and pragmatic cyber extortion (Cy-X) technique against operational technology (OT) devices; in particular, programmable logic controllers (PLCs) and their accompanying engineering workstations.

Historically, traditional, encryption-based ransomware has not been used against OT devices such as PLCs for a couple of reasons. Firstly, the adversary would require specific vendor/device exploits to attain root level access on each device they want to target, which means attacks across multiple organizations that utilize different vendor ecosystems are hard to scale. Secondly, typical engineering response and recovery practices involve replacing faulty devices with new ones and flashing the configuration back to them, which would render encrypting individual devices ineffective. However, you don't need to encrypt PLCs to perform Cy-X against OT, because in OT we have something we can target that isn't possible in IT Cy-X attacks – the physical world.

Dead Man's PLC starts at the engineering workstation, the asset where engineers will create configurations and load them onto PLCs across the OT environment. 34.7% of attacks in OT environments are facilitated by engineering workstations[30] and we see no shortage of attacks reaching it as it is one of the last bastions of IT equipment that adversaries may see.

[28] https://www.microsoft.com/en-us/security/blog/2023/07/19/expanding-cloud-logging-to-give-customers-deeper-security-visibility/
[29] https://arxiv.org/abs/2307.09549
[30] https://www.nozominetworks.com/downloads/US/SANS-Survey-2022-OT-ICS-Cybersecurity-Nozomi-Networks.pdf

When the adversary is on the engineering workstation, they can view existing ‘live’ PLC code in their project files, edit them, and download new configurations to the PLCs. Dead Man’s PLC takes advantage of this capability, as well as existing OT functionality and seldom-used security controls, to hold the victim’s entire operational process, and by proxy the physical world, to ransom.

Dead Man’s PLC works by adding to the legitimate, operational PLC code to create a covert monitoring network, whereby all the PLCs remain functional but are constantly polling one another (as depicted in the image below). If the polling network detects any attempt from the victim to respond to the attack, or the victim does not pay their ransom in time, polling will cease, and Dead Man’s PLC will trigger akin to a Dead Man’s switch and detonate. Detonation involves the deactivating the legitimate PLC code, responsible for the control and automation of the operational process, and activation of malicious code that causes physical damage to operational devices. This leaves the victim with no option but to pay their ransom; their only other alternative recovery method is to gracelessly shut down and replace every affected PLC in their operational process, which will cost them in damaged goods, lost production time, and the cost of new materials.

In the past it was believed that PLC ransomware presented an unlikely risk, due to the requirements placed on adversaries from a technical perspective. The inability to easily recycle an attack across multiple environments also acted as a deterrent, due to the time and effort required to attack each victim. However, with Dead Man’s PLC, we show that an effective and pragmatic technique towards holding the entire operational process to ransom is possible. This quick, reliable, powerful, and recyclable approach and is also vendor/device agnostic. Furthermore, Dead Man’s PLC acts as a starting point for operators to rethink the risk ransomware could pose to operational processes, and that adversaries can now move beyond recoverable encryption-based attacks, with little added technical knowledge and effort.

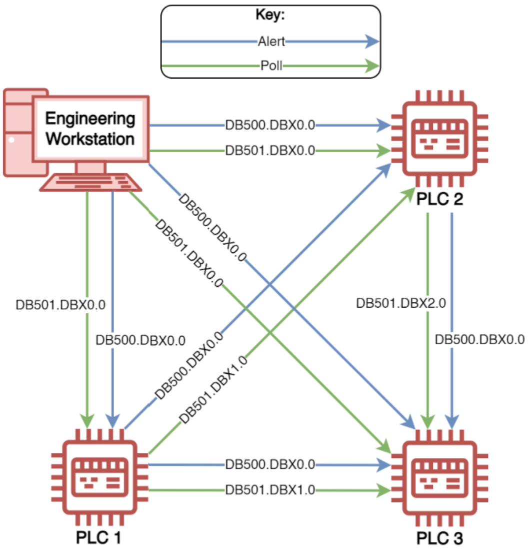

Good News Cyber

16shop Phishing-as-a-Service Platform Dismantled

In a great example of international cooperation, INTERPOL was able to take down the notorious "16shop" phishing-as-a-service (PaaS) platform, which led to the arrest of the platform's operator and two helpers.

The "16shop" platform sold "phishing kits" to hackers giving them the means to use email scams to take advantage of unsuspecting Internet users. By clicking on malicious PDF files or links, victims were tricked into giving out sensitive information, like credit card numbers. The stolen information was then used to scam people out of their hard-earned money.

The fact that law enforcement agencies and private sector partners are working together to fight cybercrime shows how powerful it is to share information.

Lolek Bulletproof Hosting Provider Taken Down

Authorities from the U.S. and Poland have taken down the popular bulletproof hosting provider <Lolek>Hosted. The providers website now shows a banner from the FBI & IRS stating “This domain has been seized by the Federal Bureau of Investigation and Internal Revenue Service - Criminal Investigation as part of a coordinated law enforcement action taken against <Lolek>Hosted.

As well as the several U.S. agencies involved in taking down the provider, there was also “substantial assistance” by two Polish authorities: the Regional Prosecutor's Office in Katowice and the Central Bureau for Combating Cybercrime in Krakow.

Bulletproof hosting services tend to disregard the content their customers post and promise to keep their identities secret. Criminals often rent IP addresses, servers, and domains from these companies to spread malware, build botnet armies, and carry out other activities related to fraud and cyberattacks.