

Acronis

E-book by Alex Fields, ITProMentor.com
Commissioned by Acronis

Preventing cyberattacks in clients' Microsoft 365 environments:

A primer on purpose-built MSP cybersecurity services

Protecting Microsoft 365 environments matters


Most organizations today, both large and small, are moving or have already moved to cloud-based Software-as-a-Service (SaaS) platforms such as Microsoft 365. Like any other software, these cloud-based systems require a certain degree of technical skill to set up correctly and to secure against various threats and risks. Out-of-the-box productivity software is designed to be as easy to use as possible, with as few friction points as possible.

The threat landscape that has been emerging during the past few years has forced some cloud providers to re-evaluate their “default” configurations, however. Case in point, Microsoft has been taking steps to improve the default security posture of tenants in the Microsoft 365 service. For example, legacy authentication (a.k.a. ‘Basic auth’), which is susceptible to password spray and other common attacks, is [being deprecated](#) later this year.

Microsoft also has started enabling something called [‘Security Defaults’](#) for new tenants, which forces users to register for Multi-Factor Authentication

(which can [prevent 99% of identity-based attacks](#)). While older tenants also have access to this feature, it would need to be manually switched on.

Even though Security Defaults is a great improvement for the baseline configuration of new Microsoft 365 tenants, there are far more opportunities for managed service providers (MSPs) to manage risk and implement cybersecurity initiatives for their small and mid-sized business (SMB) clients in the cloud. In fact, most clients will probably want to tailor their security policies using the help of talented professionals. And advanced features like Conditional Access are actually quicker and much less expensive to configure in the cloud than most traditional security solutions that would have been deployed on-premises in the past.



Will our investments ever be enough?


One of my clients, feeling exasperated after the recent disclosure of a significant Exchange (on-premises) vulnerability, asked me whether our efforts in cybersecurity really amount to anything at the end of the day, when we have nation-state actors and other advanced adversaries out there continuously pushing the boundaries just beyond our reach. How are we supposed to keep up? “I feel like we are drowning, or at least always struggling to keep our heads above water!”

I get it: in recent months we can point to more than one major incident which was both high impact and widespread enough to touch even our SMB clients. So far, 2021 has been another stressful year in cybersecurity. But fighting the good fight still matters and there is plenty we can do to keep our heads above water. In this modern era, we must. The way we accomplish this is by adopting the principles of Zero Trust.




Three principles of Zero Trust

If you have seen references to this idea of Zero Trust in marketing materials for various cybersecurity products, you might still be confused about it. The first thing to note is that Zero Trust is not something you can buy in a box, nor is it a place that you can go. It is also more than just the latest catchphrase in the industry. It is best described as a fundamental stance or attitude toward security. Microsoft uses three basic principles to describe its approach to Zero Trust:




Verify explicitly:

This means validating as many attributes as possible when granting or denying access requests. The key Microsoft technology being leveraged here is Azure AD Conditional Access. Using Conditional Access, we can explicitly verify a user's identity, endpoint, location, and other risk indicators to determine if the request is legitimate or illegitimate.



Least privilege:

The principle of least privilege says that you should only grant enough access to perform the requisite job duties, and no more. This rule should be applied liberally and across the board. For example, you should limit administrative access within your cloud services, as well as for your on-premises infrastructure, and even your clients' endpoints such as Windows 10 workstations.



Assume breach:

At some point, your defenses will fail, and trust will be compromised. You must plan for this inevitability accordingly. That implies having people, tools, and processes in place that can be activated and leveraged in the event of a confirmed breach event. It also means actively looking or 'hunting' for indicators of compromise in your clients' environment.

Zero Trust means literally moving beyond the concept of trust

A question arises: “What about these [recent high-profile supply chain attacks](#) that fundamentally took advantage of trust? Are the attackers finding ways to circumvent the Zero Trust model?” It is a good question. When an attacker compromises your vendor (whom you trust) or takes control of your domain (which is the basis of trust in your own environment), do these principles fall apart? **The answer is No.** In fact, they are even more necessary.

That is the whole point. Zero Trust means literally moving **beyond** the concept of trust. That is why we talk about **explicit** verification, granting **just enough** access, and carrying an ever-present awareness of (and ability to respond to) possible breach events.

Therefore, the framework is sound. The real question is: “How far are you willing to go on this journey?”



Solorigate (Nobelium)

Let us apply the framework above to a recent real-world attack on trust. For those of you who do not already know about the Solorigate event, this was a major cybersecurity incident that was discovered in late 2020 and early 2021 that impacted many businesses and government agencies in the United States. The attack was carried out by an extremely sophisticated actor that Microsoft later named Nobelium; the group is most likely nation-state sponsored.¹

Solorigate was a supply chain attack carried out against SolarWinds: a software company with a product called Orion, which is a suite of IT management tools. This implies nearly automatic privileged access to corporate networks and infrastructure. The idea to infect a trusted software package used for daily IT work and, from there, spread to other targets within the organization.

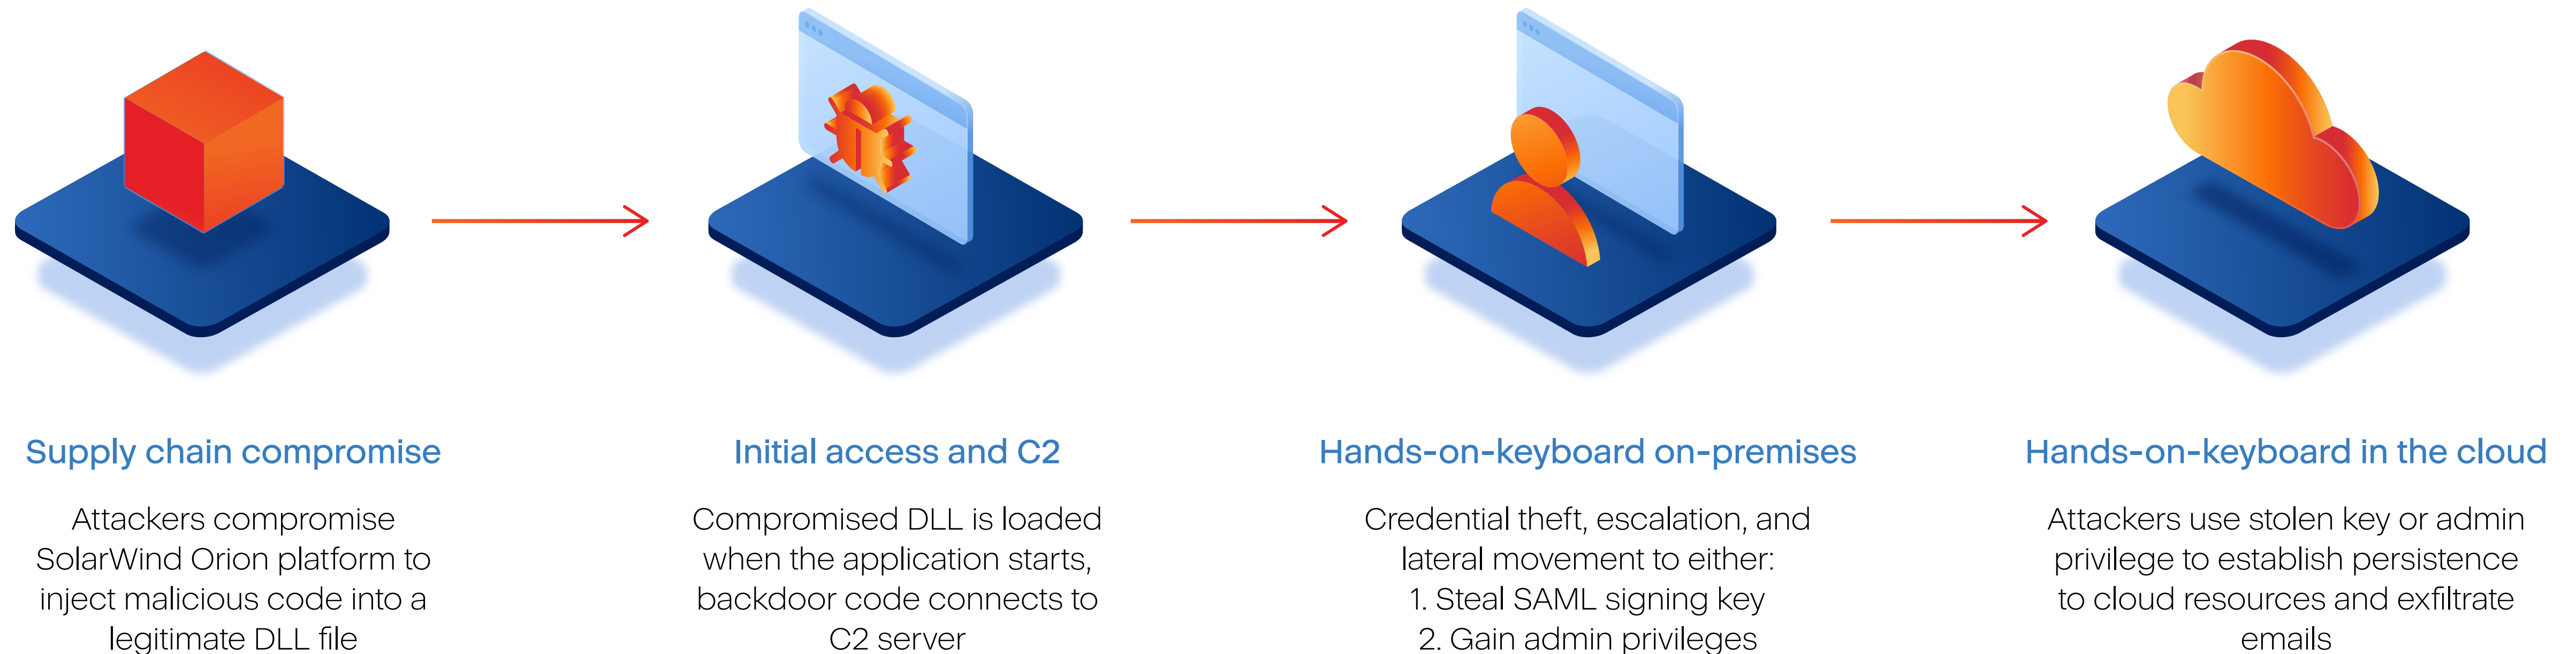
The idea is to infect a trusted software package used for daily IT work and, from there, spread to other targets within the organization.

SolarWinds Orion is deployed at many different private and public institutions, including several in the United States government (the primary target of this attack). It is equally important to understand what Microsoft believes was the ultimate goal of Nobelium: to exfiltrate email data from these government institutions (the email provider in many cases here is Microsoft's Office 365 Exchange Online service).

Now, if your intended destination is email accounts hosted at a third-party cloud provider like Microsoft, why would you start by targeting a supply chain to infiltrate internal networks? Why not go directly to the provider? First, realize that Microsoft's cloud service is a more abstract target; each individual organization has their own security boundary or 'tenant' within the Microsoft cloud. Organizations connecting to this public service have existing trust relationships set up to manage their own respective tenants.



At a high level, the attack chain was carried out in the following stages:

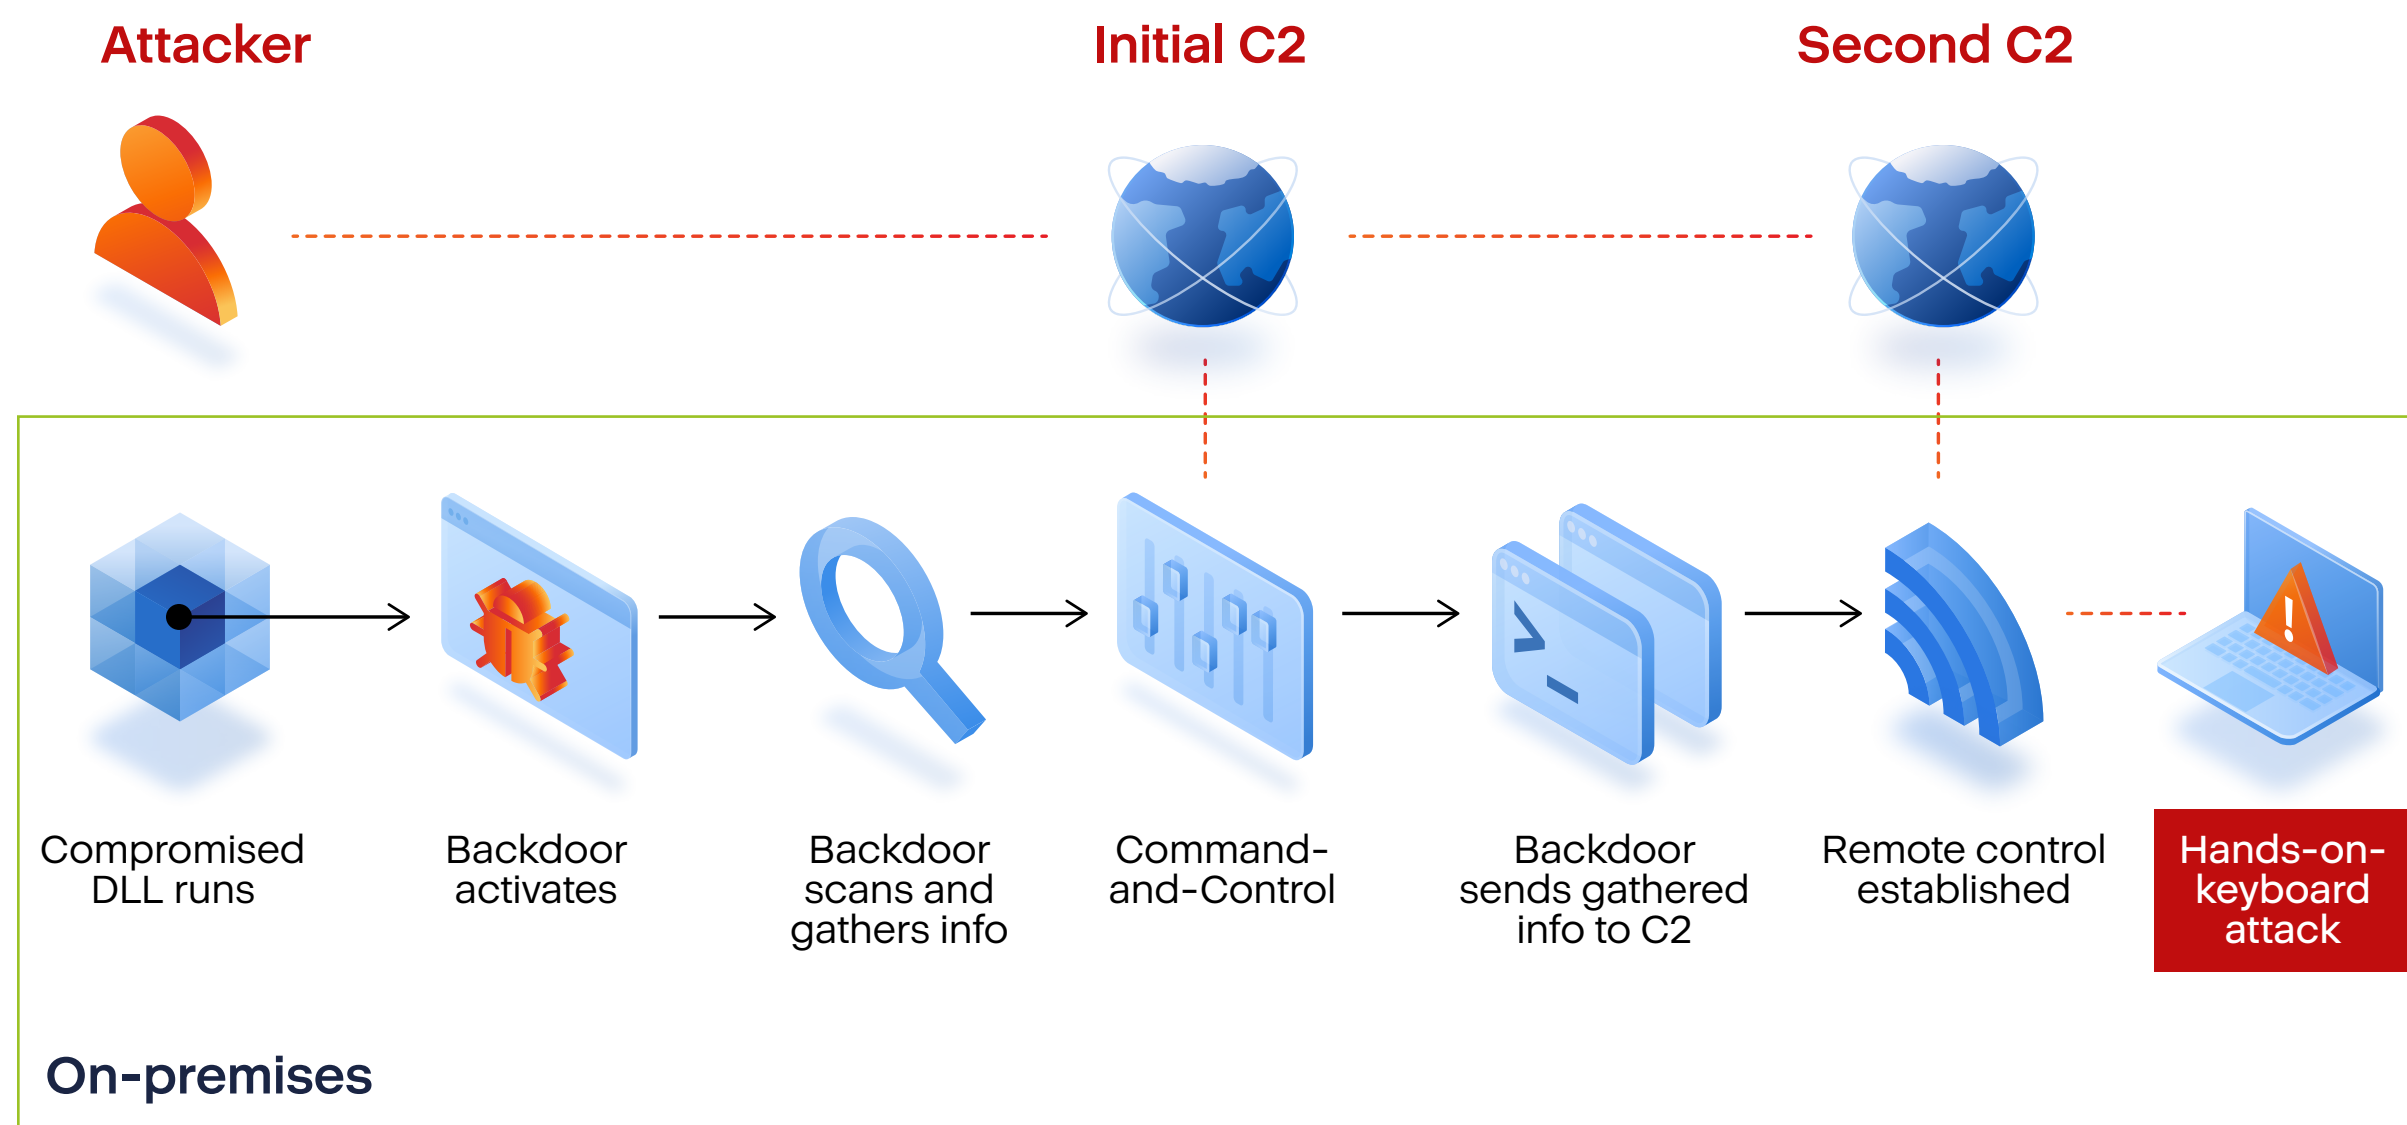


We do not have all the details on the initial supply chain compromise, but for the purposes of our analysis here today, it does not play a huge part. Everything that happens after that first stage takes place within your environment, and therefore you have opportunities to interrupt the attack only after the supply chain compromise has already taken place.

Note: For more information on the Solorigate/Nobelium attack, see Microsoft's Nobelium Resource Center and also CISA's Alert AA21-008A.

Initial Access and Command-and-Control

The latter three stages offer several good lessons for anyone trying to prevent attacks like these from impacting them in the future. Let us start by analyzing the first stage: Initial Access and Command-and-Control.



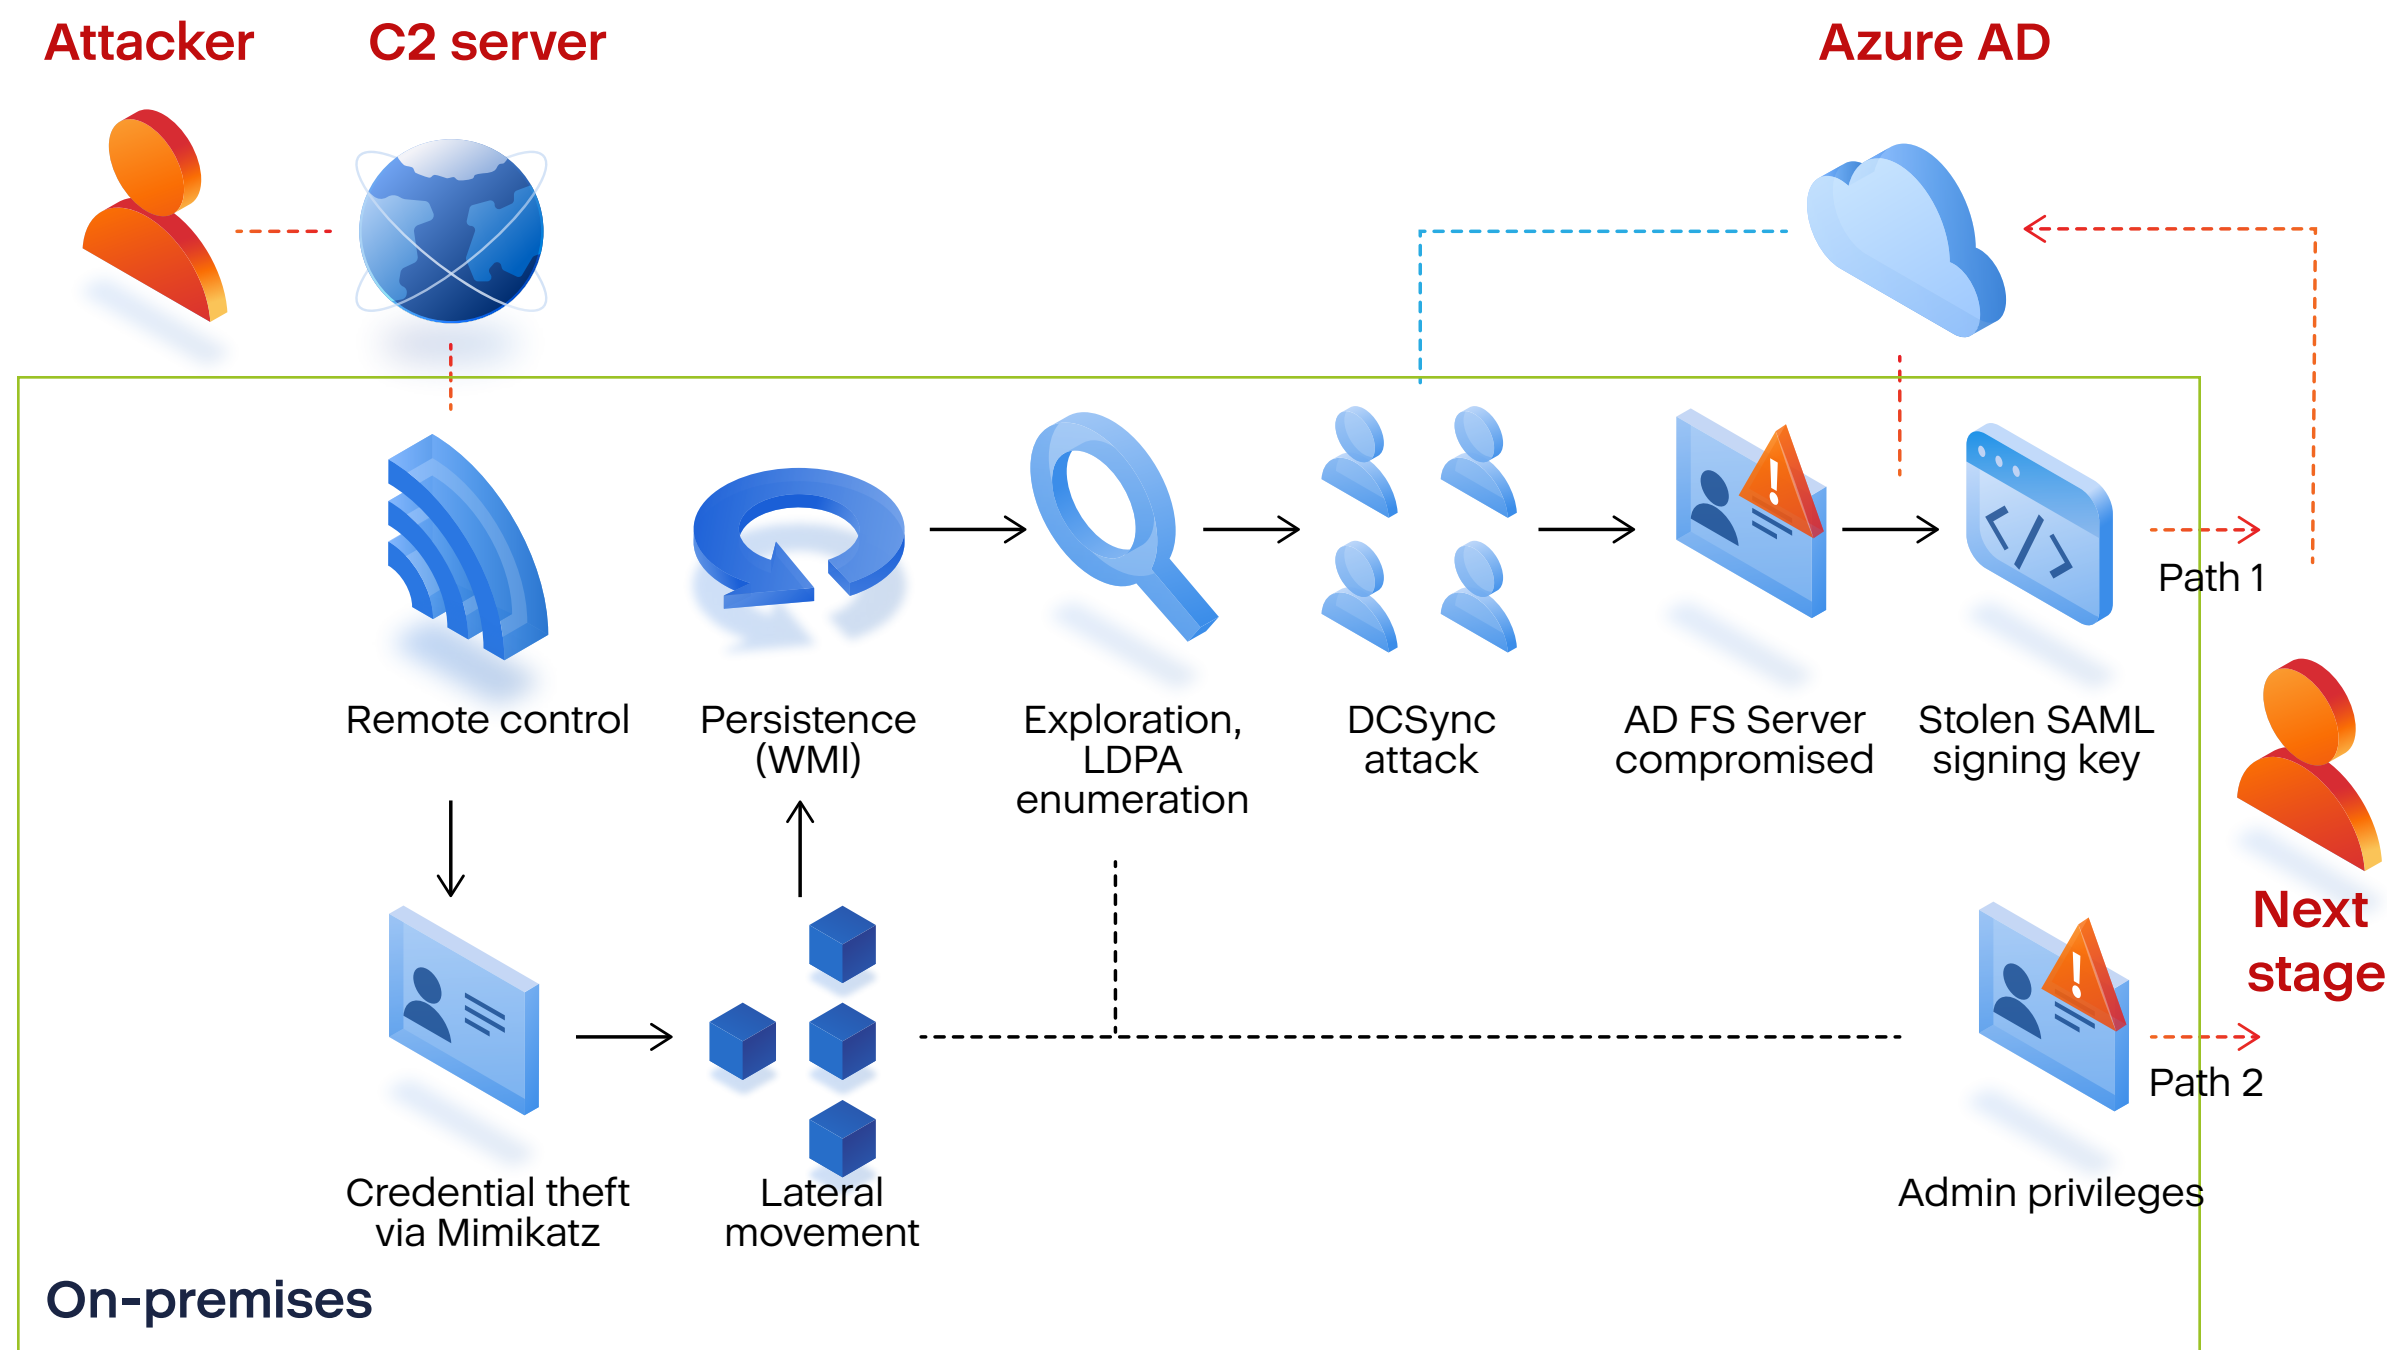
In this stage, the attackers chose carefully which organizations to target; not every backdoor that was activated was ultimately carried forward to the subsequent hands-on-keyboard attacks. Therefore, the initial command-and-control was established purely to gather information about the affected devices so that the attack group could determine whether they wanted to proceed with the attack in each instance.

The second command-and-control is established to provide remote access so that the attackers could escalate their activities to the next stage. Another interesting attribute about this phase of the attack is that the initial C2 servers generate unique domains for each affected device; this is another clue as to the maturity level of this attack, and of course, it is much harder to detect as well.

Hands-on-keyboard attacks on-premises

The next stage of the attack had one primary goal: gain access to the cloud (where the data is stored). This is possible through a couple of different paths:

- **Path 1:** Find and compromise the AD FS server (if it exists)
- **Path 2:** Find an administrator account that has privileged access in the cloud



In either case, many of the techniques that were used in this stage are familiar to us: credential theft via Mimikatz, lateral movement and exploration, establishing persistence, and so on. Since initial access happened via compromised IT management software, gaining domain admin was fairly trivial in this case. Where most attackers would consider this step a huge victory, it was more like a given for Nobelium. What they did next was a little more unique.

If the environment was set up for on-premises federation with an Active Directory Federation Services (AD FS) server, then this was the preferred path to move into the cloud. The attacker would gain control of the federation trust, and then they would steal the SAML signing key so that they could mint their own SAML tokens moving forward.

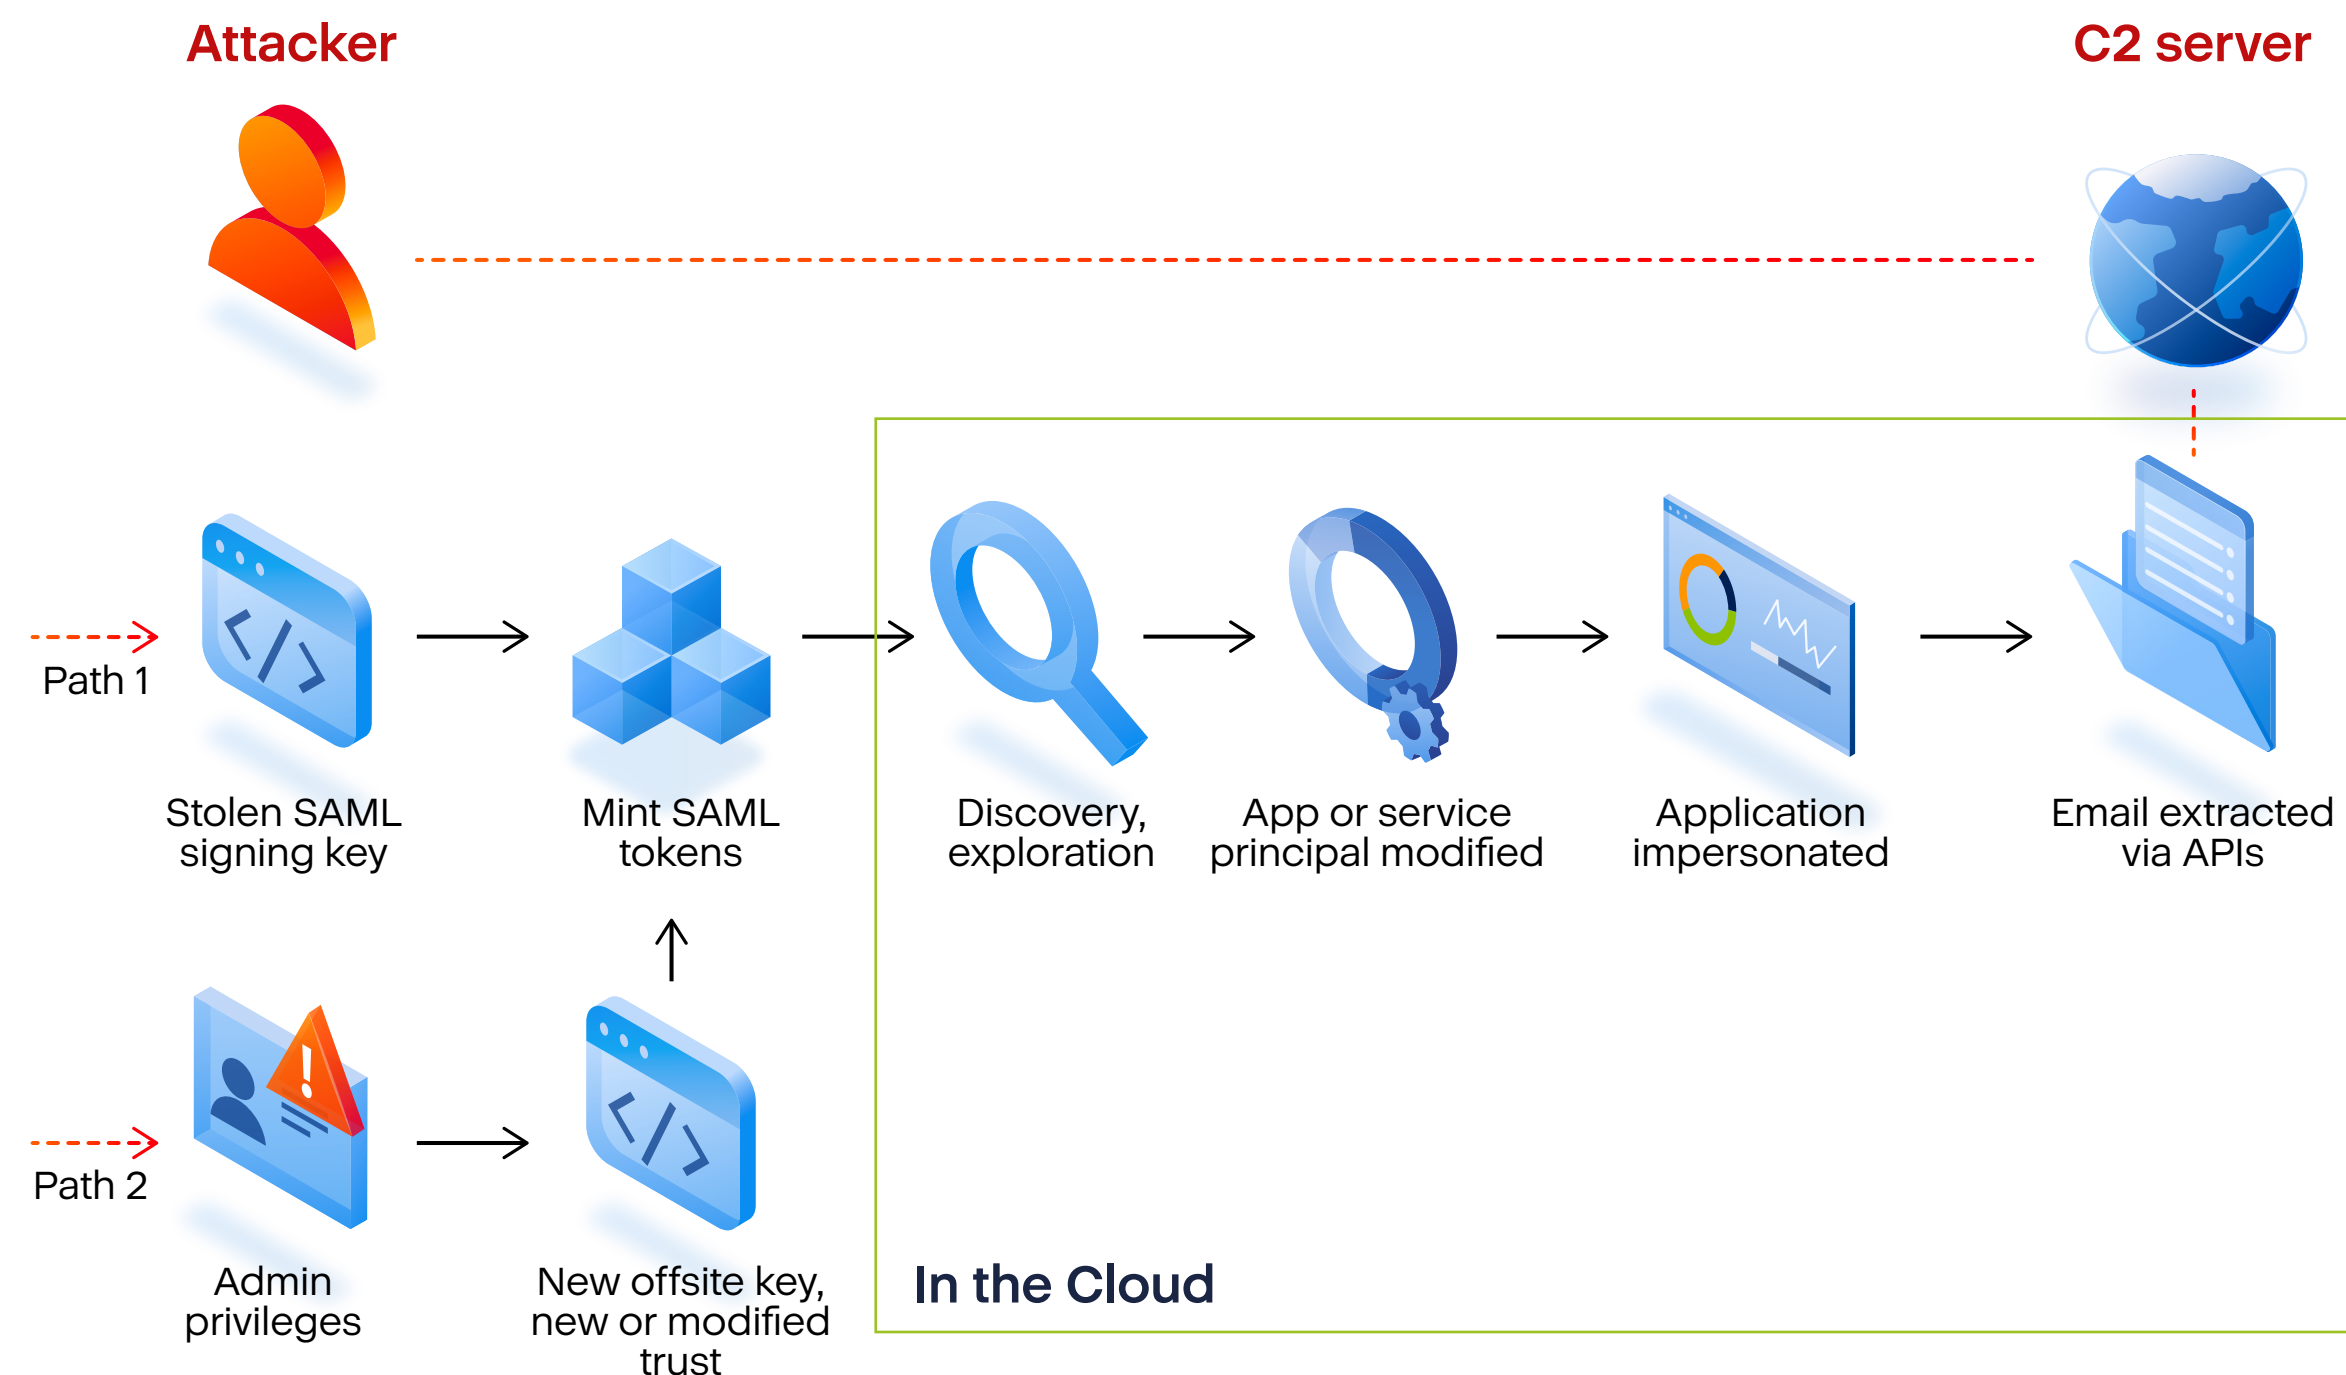
If the domain was not set up for federation (what Microsoft refers to as a “Managed” domain), then all the attackers had to do was gain access to an administrator account that held privileges both on-premises and in the cloud. Unfortunately, this is an incredibly common practice: syncing a domain admin account to the cloud and assigning that same account global administrator privileges.

Hands-on-keyboard attacks in the cloud

In this next stage, the idea was to establish persistence in the cloud and ultimately extract email data from Exchange Online. If the attacker already cloned the identity provider on-premises (AD FS) then they could issue their own SAML tokens. If they had started with only global administrator access, they may have had an extra step of creating a new trust with their own offsite key and then minting SAML tokens from there. Either way, at this point they no longer needed

access to the on-premises environment and, in some cases, they even took their time to go back and remove footprints left behind from the earlier stages.

At this point, they no longer needed access to the on-premises environment and, in some cases, they took time to remove footprints left behind from the earlier stages.



A less skilled attacker would have approached this phase differently. Most likely they would have accessed email directly either by impersonating accounts or signing into the mailboxes and then establishing forwarders. The problem with that more traditional approach is that it is very 'noisy.' It leaves a lot of audit data behind and it increases the risk that their activities will be discovered before they have accomplished their goals.

What is remarkable about this particular incident is not only that the attacker carried out a successful supply chain attack against the U.S. government and

other agencies downstream of SolarWinds, but that they were able to do so **without being detected.**

By looking for existing applications that were already connected to Microsoft 365, Nobelium was able to quietly add or modify permissions on those applications, and then impersonate the apps to run bulk extractions against the Exchange Online service. For example, let us say you have a third-party backup configured for your cloud-based email. That means you have an application object in Azure AD that already has access to read all emails. If you did not have such an application, then a different application object could be modified to add the permissions necessary to read emails.

Either way, the application was manipulated so that the attacker could impersonate the app at will. Thus, the requests for data being made by the application would appear to be perfectly legitimate. Essentially, the Nobelium actor was able to launder the exfiltration of massive amounts of email data. Again, they were quiet.

You might be thinking, “Well, this is highly interesting stuff, but so unlikely to touch my average SMB client!” Not true, my friend. Some SMBs were also impacted by this attack, in addition to the bigger fish. Furthermore, the reality

that we live in today is that all the less skilled actors out there are watching the pros and learning from them more quickly. There are already kits available on GitHub that make it easier for others to replicate some of these same techniques.

In short, your client doesn't have to be one of the larger targets to get hit by increasingly sophisticated attacks. This is the world we live in (sorry to be the messenger).

You do not have to be one of the larger targets to get hit by increasingly sophisticated attacks. This is the world we live in.



Lessons learned

What have we learned from the Solorigate incident?

It can be tempting to throw up our hands and conclude that there is simply nothing we can do to stay ahead in today's threat landscape. Perhaps we should all just buy cyber incident insurance and call it a day. But I disagree. Yes: it can be difficult to keep treading water in such a stormy sea, but why would you stop trying to save yourself or your clients from drowning? That would be suicide indeed. Sure, you might sink anyway, but do not give up so quickly!

I suggest that we double down on the Zero Trust approach. Let us explore some of our capabilities in Microsoft 365. As we look at each of the three principles again, we will build a short checklist describing a typical 'Zero Trust' journey in the Microsoft cloud. I will also map these items against a common cybersecurity framework called the [CIS Critical Security Controls](#).



Verify explicitly

1. Strong authentication (MFA)

Strong Multi-factor Authentication (MFA) is not a silver bullet, but it is still the single best thing you can do to protect your identities in the cloud (especially your admin accounts). Having another step to verify the identity of a user is crucial for preventing traditional identity-based attacks such as password spray (which is another technique that Nobelium used when they did not have the backdoor option), and is the best defense we have against phishing and credential theft. Yes, there are other email security protections available that can serve as a first line of defense, but in the context of a “Zero Trust” conversation, you have to assume that your clients’ end users’ credentials will be exposed or harvested at some point. And that is why MFA is no longer optional. As I like to tell my clients, if you do not implement MFA then you are essentially signing up to be compromised.

As I mentioned previously, we have multiple ways to turn this on and newer tenants may have this enabled by default.

Security defaults: This option is free and available with all subscriptions. Enable this feature from Azure AD > Properties (scroll to the bottom of the page to find a link to manage your security defaults).

Per-user MFA: Another free option. You can also enable or disable MFA on a per-user basis. See [this article](#) for more details. You will need to take additional steps to [block legacy authentication](#) if you use this method.

Conditional access: This is the preferred approach moving forward for most organizations. Here you have a chance to [customize your security policies](#) further (unlike the Security defaults, which cannot be customized). Conditional access policies require at least Azure AD Premium P1 (which is included with Microsoft 365 Business Premium or Enterprise plans).

If you plan to use Conditional access (recommended) then be sure to implement at least these policies to begin:

- [Block legacy authentication](#)
- [Require MFA for administrators](#)
- [Require MFA for Azure management](#)
- [Require MFA for all users](#)

These four policies when configured together mimic functionality enabled by [security defaults](#).

Implementing strong authentication satisfies several requirements from the CIS Critical Security Controls: 4.5 (Require MFA for admin accounts) and 16.3 (Require MFA for all users).

We also have [passwordless authentication](#), but we will discuss this option later.

In the case of Solorigate, MFA did not really concern them very much since they already had access to trusted networks and could mint their own SAML tokens at will. Still, it is worth mentioning that password spray did in fact play a role in Nobelium's overall strategy (remember password spray is 99.9% mitigated by disabling basic authentication and enabling MFA). It was just one of many options, however. They could also lean on the supply chain compromise and gain access through other means, as we saw in the attack chain.

2. Require compliant devices

[Conditional access policies](#) can also provide us with additional verification options. For example, [require devices to be marked as compliant](#). When implementing device-based policies, it is a good idea to also [block unsupported device platforms](#), so that you do not leave gaps in this type of verification. On its own, device compliance can be an effective defense against certain types of

advanced attacks (e.g. man-in-the-middle), but there are many other benefits hidden here in terms of your overall strategy.

Requiring device compliance means that every device will need to be enrolled for management before that device can be granted access to resources. This is significant for a number of reasons. First, a managed device is [71% less likely to become compromised](#) than an unmanaged device, according to Microsoft. That is because you are going to be able to implement device controls (e.g.: remove local admin, Attack Surface Reduction rules) that further diminish the chances an attacker can execute malware, steal credentials, move laterally, and so on.

Second, with device compliance in place, you start to gain inventory and control over your clients' the hardware and software assets. This helps you to satisfy CIS Critical Security Controls #1 (hardware) and #2 (software). Not to mention, once you control their inventory, you can put those additional measures in place that make it more difficult for would-be attackers to gain a foothold and move around in the environment (see CSC #5: Secure Configuration of Hardware and Software).

That brings us to our next Zero Trust principle.

Least privilege

There are many ways to apply the principle of least privilege, both on-premises and in the cloud. Let us start with the assumption that attackers gaining access to internal corporate resources is already a given (as it was for Nobelium).

Therefore, we want to put in place controls that create a natural limitation or 'blast radius,' which can help to mitigate the impact of the initial breach event, and to prevent the subsequent activities that lead to further compromise in the cloud.

3. Move to cloud-only accounts (at least admins)

We already touched on this point briefly as we went through the attack chain. Can you guess? Answer: you must take extra care with your clients' most privileged accounts. It has been a long-standing best practice to use separate, dedicated administrator accounts, and to minimize the number of accounts assigned superuser privileges such as domain admin (on-premises) and global admin (in the cloud). This is covered under CSC #4: Controlled Use of Administrative Privileges.

But taking it one more step, we should also separate the local and cloud administrative accounts. Many organizations sync all of their accounts and assign full global admin privileges in the cloud to regular domain administrators. Stop it. This is a very bad practice which was exploited in the

Solorigate attack (and in others). An attacker who compromises one domain should not have automatic access to another!

Many organizations sync all of their accounts and assign full global admin privileges in the cloud to regular domain administrators. Stop it.

Furthermore, for your own environment, you should consider whether it makes sense to keep legacy on-premises infrastructure at all. Especially for SMBs, moving to 100% cloud-only accounts is a real possibility. But even if you decide that it makes sense to keep on-premises infrastructure (perhaps to support a legacy line-of-business app), do you really need to connect that legacy infrastructure to the cloud?

Very often, the original reason we historically opted to implement technologies like Azure AD Connect or AD FS was to keep our passwords the same on-premises and in the cloud. But did you know about [passwordless authentication](#) in Azure AD? What is the point of keeping password sync in place if you hardly (if ever) have to touch the password in the cloud? ²

Many SMBs would benefit from simplifying their environment and removing all the extra overhead which attends directory sync: Azure AD Connect, AD FS, the last on-premises hybrid Exchange Server, and so on. More and more, these relics are becoming liabilities that frankly are not worth the hassle for most small and mid-sized organizations. Therefore, converting all of your client accounts to 'cloud-only' is a very good idea.

4. Enable admin consent request

Third-party apps are often granted permissions to interact with data in Microsoft 365, but we can choose to control how these permission grant requests happen. By default, all users have the ability to consent to permissions requests coming from outside applications. You can limit or remove standard users' ability to grant these requests from **Azure AD > Enterprise apps > Consent and permissions**.

Taking it further, we can enable [Admin consent requests](#) from **Azure AD > Enterprise apps > User settings**. This feature will allow you to delegate the responsibility of adding and managing applications and permissions to specific users within your client's organization.

This item helps us to meet sub-controls from CSC #2: Inventory and Control of Software Assets, as well as CSC #4: Controlled Use of Administrative Privileges.

5. Regularly audit privileged Azure AD roles + applications

As we have seen with Solorigate and other attacks, adversaries will often establish their own accounts or application permissions. Therefore, it is a good idea to periodically audit your environment to see if any accounts or permissions can be removed. You ought to do this for all of your accounts but pay special attention to Azure AD roles that have privileged access such as Global administrator, Global reader, User administrator, Exchange administrator, and so on. Enterprise apps should also be audited on a regular basis. To assist in this endeavor, check out the [Azure AD Incident Response PowerShell module](#). This tool includes some handy functions that allow you to output privileged accounts and application permissions, for example:

```
$TenantID = Get-AzureADIRTenantId -DomainName <InsertPrimaryDomain>  
Get-AzureADIRPrivilegedRoleAssignment -TenantId $TenantID -CsvOutput  
Get-AzureADIRPermission -TenantId $TenantID -CsvOutput
```

² **Note:** technically the password attribute in the cloud would still exist, but [Microsoft recommends](#) not rotating the password unless you believe the account has been compromised.

This output should be audited on a regular basis and compared to a known good sample or baseline from the tenant. These items correlate with sub-controls from CSC #16: Account Monitoring and Control, as well as #18: Application Software Security.

Is there more we can do in the category of least privilege? Of course. There is always more we can do, but these items here will be accessible to most managed environments and popular subscriptions such as Microsoft 365 Business Premium. That brings us to our last Zero Trust principle.



Assume breach

We cannot assume all of the previous defenses will leave us bulletproof. If only it were so! Yes, what we have done up to this point will turn the attack chain into a 'bumpier ride' for would-be attackers but it is still possible that someone will eventually get past your defenses. What then? You will need to go into Incident Response (IR) mode. But Incident Response does not start only once you discover a breach! It starts right now: you must prepare today, because preparedness is everything when it comes to IR.

6. Create and practice a formal incident response plan

Too many organizations overlook planning. But this is a very valuable service that MSPs could offer to smaller orgs with limited budgets and resources: Incident Response Planning. IR plans will outline who is responsible for which activities both during and after a breach or other cybersecurity event. It may also outline a few steps that the team will take under various circumstances. What happens when data is lost or stolen? What happens if the organization is hit with ransomware? What happens when you get locked out of your tenant? And so on. It is a good idea to document and even practice some of these scenarios; this will prepare you to be at your best when the real thing happens. You can find many policy templates to get you started online, such as [this one from SANS](#). As well, Microsoft publishes some Incident Response playbooks for several common scenarios like phishing and illicit application consent

grants; these resources are a good addition to any Microsoft 365 Incident Response Plan.

This item correlates with CSC #19: Incident Response and Management.

7. Audit log monitoring and alerts

To assist incident responders in their work, one of the best things you can do in advance is enable auditing. When an incident unfolds, you need accurate data that can be reviewed by cybersecurity professionals: this tells us what happened when, which accounts or apps are likely compromised, and more. There is nothing worse for an incident response team than learning that there is no data. It makes the job much more difficult because visibility issues must be fixed before we can really be helpful in recovery efforts.

One of the best things you can do is enable auditing. There is nothing worse for an incident response team than learning that there is no data. It makes the job much more difficult.

At a minimum, you must [enable the unified audit log search](#) so that events are centrally recorded and searchable. Once this is done, you can also turn on the built-in [Alert policies](#) that Microsoft provides via the Security & Compliance center; suspicious events within the tenant will then be forwarded to a monitored mailbox.

If you want to take this journey further, consider a Security Information and Event Management (SIEM) tool such as [Azure Sentinel](#), so that you can aggregate other log sources and customize your retention of the log data. This will help not only during incident response, but also in detecting potential security threats by analyzing or ‘hunting’ across the log data for indicators of compromise (IOCs). Sentinel is great for MSPs, too, because it can integrate with Azure Lighthouse to provide multi-tenant management for your customers. See this article for more details: [Build a scalable security practice with Azure Lighthouse and Azure Sentinel](#).

However, it is worth mentioning again that Microsoft also publishes the [Azure AD Incident Response PowerShell module](#): a free option that allows incident responders to gather data from Azure Active Directory regardless of whether the client has a SIEM in place or not. The amount of historical data available to this tool is dependent on what subscription level you have for Microsoft 365. It

is recommended to have at least Azure AD Premium P1 (available in Microsoft 365 Business Premium and up), which gives you 30 days of audit data to query. Obviously a SIEM would be more ideal, but just know that you always have other tools to fall back on where others do not exist.

These items correlate with CSC #6: Maintenance, Monitoring, and Analysis of Audit Logs.

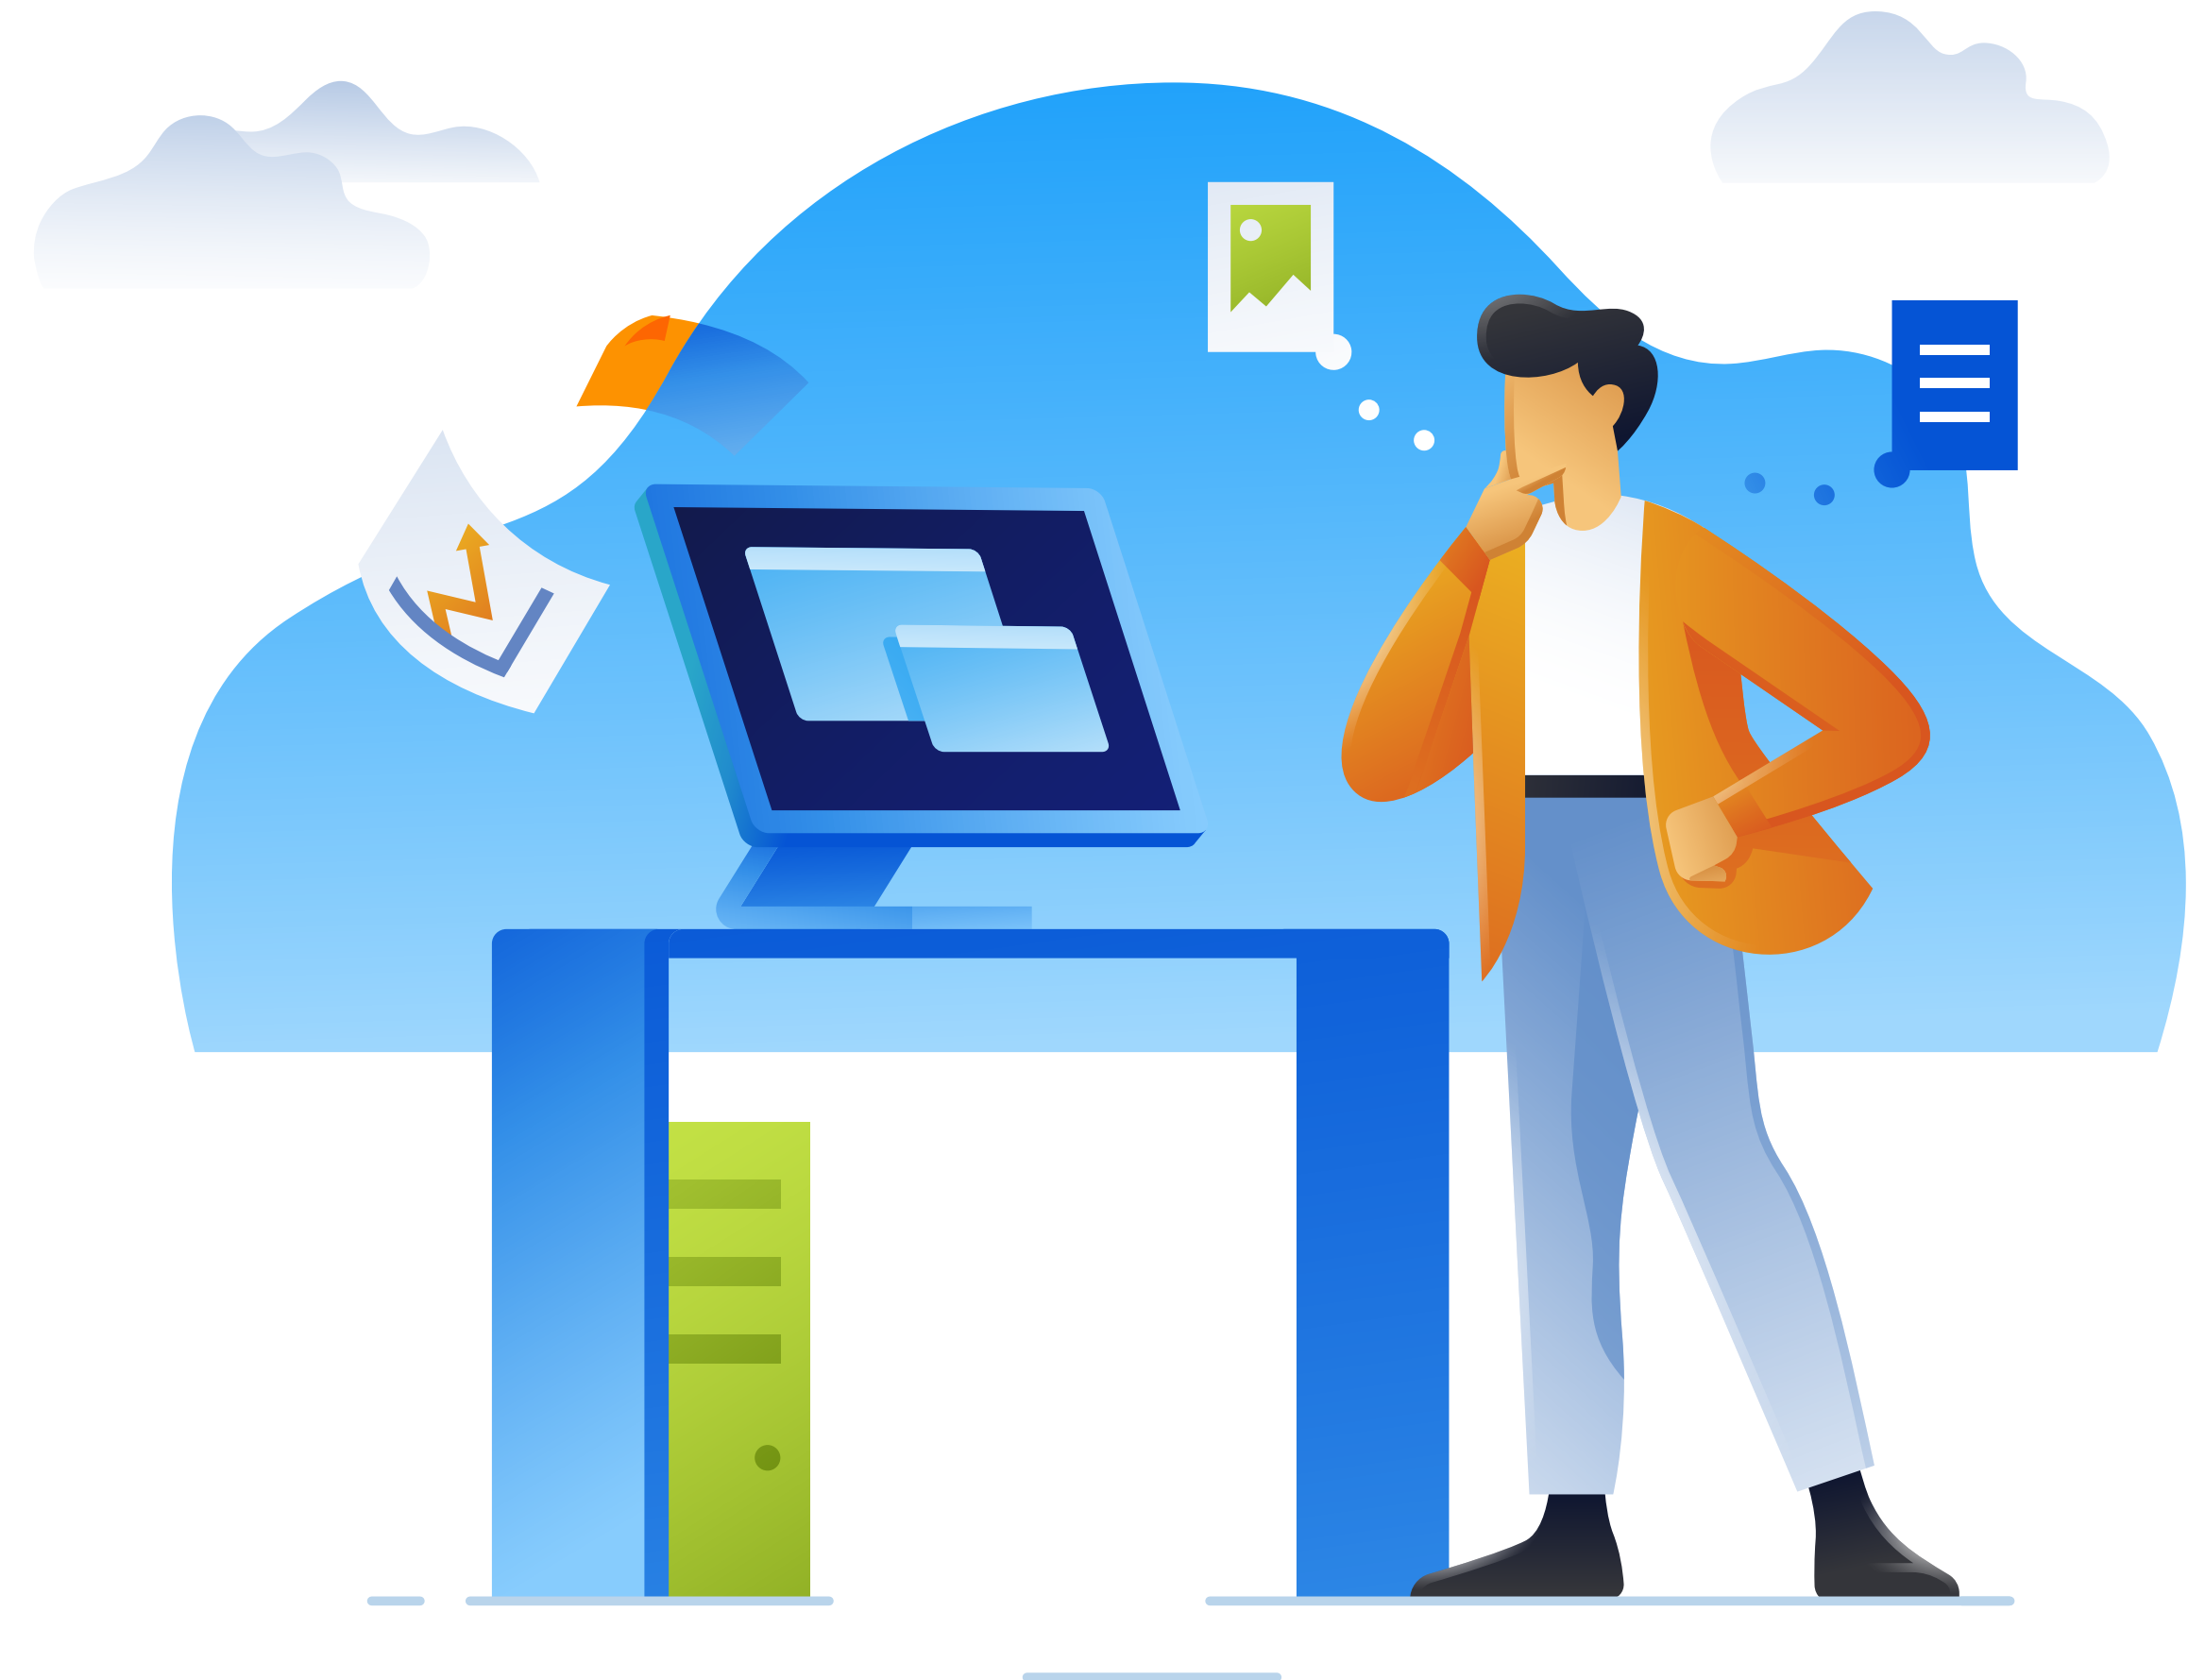


8. Consider adding a backup solution to Microsoft 365

What about backup? Does Microsoft back up your data in the cloud? Actually [yes, they do](#). However, Microsoft's backup is not your backup. Microsoft does not provide clients with any type of self-service access to the backups they keep. Some have pointed to other native features like [Retention policies](#) and [Litigation Hold](#), but these protections are intended to be used for compliance and legal scenarios, not to be used as a backup. And many clients still would not be comfortable keeping all of their eggs in the Microsoft basket, so to speak.

Whether or not you have protection such as retention policies in place, if you get into a bad situation where you have lost data due to ransomware or otherwise, Microsoft Support will only be able to offer you limited "best effort" recovery. If a restore operation does not work, then you are out of luck. Therefore, most organizations will be interested in third-party solutions such as [Acronis Cyber Protect](#) that protect the data across all systems and applications, including Windows, Microsoft 365, Active Directory, and SharePoint while also providing built-in anti-ransomware, anti-malware, backup encryption capabilities, MFA, technologies for instant restore, and ability to track the protection status of critical data across the organization.

This item correlates with CSC #10: Data Recovery Capability.



Conclusion

Today's attacks against Microsoft 365 are going beyond the traditional password spray and phishing attempts that we have seen in years past. We are now witnessing more sophisticated actors such as Nobelium teach other attackers how to better maneuver in these new cloud-first / mobile-first environments.

It is worth acknowledging that there many other things that we can and should do to prevent or at least minimize the chances of a trust compromise in the first place. It's essential to implement a rigorous vulnerability assessment and patch management process to keep your Windows and Microsoft 365 environments up-to-date and free of backdoors that can be exploited. Anti-malware solutions that focus on emerging threats such as zero-days or fileless attacks are also crucial for protecting against modern cyberthreats. Implementing email security can also help reduce the risk and protect your Microsoft 365 email boxes against phishing, impersonation attempts, malware, advanced persistent threats (APTs) and other attacks that try to steal sensitive data or gain unauthorized access or financial profit.

Managing multiple point solutions and training your team on them could prove especially painful for SMBs and MSPs, however. This is why many service providers look for integrated solutions to build their services on, such as

[Acronis Cyber Protect Cloud](#), as they consolidate multiple functionalities in a single platform, which greatly reduces the management burden and total cost of ownership.

While these classic protection layers are still a good and necessary thing for any organization, the writing is on the wall: those who fail to embrace a more thorough Zero Trust approach will be left behind and subjected to an increasing number of threats. You simply cannot guarantee, for example, that 100% of phishing or credential harvesting will be prevented at all times, and yes, sometimes threats may come through otherwise trusted sources. There is no one silver bullet here; instead, it is necessary to implement proven security controls at multiple layers to ensure the best all-around coverage that you can.

While the checklist presented herein is not exhaustive, I find that most small and mid-sized organizations have plenty of work to do within the items that we have shared here today.



Verify explicit

- Strong authentication (MFA)
- Require compliant devices



Least privilege

- Move to cloud-only accounts
- Enable admin consent requires
- Regularly audit roles and apps



Assume breach

- Incident response planning
- Audit log monitoring and alerting
- Consider a third-party backup

About the author



Alex Fields

Having grown up in a family-owned small business, Alex started adopting Microsoft IT solutions at a young age. In his spare time after school, he gained experience with Windows Server, Exchange Server, and the migration to Small Business Server.

After completing college, he returned to IT with a local managed service provider, where he gained experience on more modern Microsoft cloud solutions over the better part of a decade. He has been awarded the

Microsoft MVP for his technical writing and other contributions to the community.

Today, Alex works independently, helping MSPs and other IT consultants to adopt modern technologies and build managed services practices in the cloud. He currently lives in Minneapolis, Minnesota (USA). He loves to write about technology topics and shares his ideas, experiences, best practices, and lessons learned with the community.

Additional insights from Acronis

[Acronis Blog](#): Provides the latest updates and insights from the world's cyber protection leader.

[Acronis YouTube Channel](#): Delivers frequent videos of use cases, demos, cyberthreat analysis, and company news.

[Acronis Resource Center](#): The go-to hub for cyber protection white papers, e-books, in-depth articles, tutorials, infographics, etc.

[Acronis Events](#): Ongoing series of events, webinars, interviews, etc., including details on joining.



About Acronis

Acronis unifies data protection and cybersecurity to deliver integrated, automated [cyber protection](#) that solves the safety, accessibility, privacy, authenticity, and security ([SAPAS](#)) challenges of the modern digital world. With [flexible deployment models](#) that fit the demands of service providers and IT professionals, Acronis provides superior cyber protection for data, applications, and systems with innovative [next-generation antivirus](#), [backup](#), [disaster recovery](#), and [endpoint protection management](#) solutions. With award-winning [AI-based antimalware](#) and [blockchain-based data authentication](#) technologies, Acronis protects any environment – from [cloud to hybrid to on-premises](#) – at a low and predictable cost.

[Founded in Singapore in 2003](#) and incorporated in Switzerland in 2008, Acronis now has more than 1,500 employees in 33 locations in 18 countries. Its solutions are trusted by more than 5.5 million home users and 500,000 companies, including 100% of the Fortune 1000, and top-tier professional sports teams. Acronis products are available through 50,000 partners and service providers in over 150 countries in more than 40 languages.



Acronis

Learn more at www.acronis.com

Copyright © 2002-2021 Acronis International GmbH. All rights reserved. Acronis and the Acronis logo are trademarks of Acronis International GmbH in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners. Technical changes and Differences from the illustrations are reserved; errors are excepted. 2021-10