

# Cryptojacking: Understanding and defending against cloud compute resource abuse | Microsoft Security Blog

By Microsoft Threat Intelligence

Published: 2023-07-25 · Archived: 2026-04-05 16:02:23 UTC

In cloud environments, cryptojacking – a type of cyberattack that uses computing power to mine cryptocurrency – takes the form of cloud compute resource abuse, which involves a threat actor compromising legitimate tenants. Cloud compute resource abuse could result in financial loss to targeted organizations due to the compute fees that can be incurred from the abuse. In attacks observed by Microsoft, targeted organizations incurred more than \$300,000 in compute fees due to cryptojacking attacks.

While there are fundamental differences in how cloud providers handle authentication, permissions, and resource creation, a cloud cryptojacking attack could unfold in any environment where a threat actor can compromise an identity and create compute, and the attack lifecycle is largely the same. Microsoft security experts have surfaced tell-tale deployment patterns to help defenders determine, identify, and mitigate cloud cryptojacking attacks.

To perform cloud cryptojacking, threat actors must typically have access to compromised credentials obtained through various means, highlighting the need to implement common best practices like [credential hygiene](#) and cloud hardening. If the credentials do not have the threat actors' desired permissions, privilege escalation techniques are used to obtain additional permissions. In some cases, threat actors hijack existing subscriptions to further obfuscate their operations.

Once access to the tenant is gained, threat actors create large amounts of compute, preferring core types that allow them to mine more currency faster. Threat actors use these deployed resources to start mining cryptocurrency by installing cryptomining software in the newly created virtual machines (VMs) and joining them to mining pools.

In this blog post, we present insights from our research on how attackers launch cryptojacking attacks in cloud environments. These insights deepen our understanding of these threats, which in turn inform the protections that we continuously build into our cloud security solutions. We share patterns that administrators and defenders can look out for to identify if a cryptojacking attack is occurring within their cloud environment. We also provide information on how Microsoft Defender for Cloud, Microsoft Defender for Cloud Apps, and other solutions can detect cryptocurrency mining threats and related malicious activity.

While this blog covers mitigation and protections against cloud cryptojacking, in general, strengthening cloud security posture, protecting cloud workloads from threats, and better control of cloud app access can help organizations defend against a wide range of cloud-based threats and risks.

## Cryptocurrency mining in cloud environments

In incident response investigations and proactive research in the past year, we observed threat actors abusing administrative features to deploy and manage cryptocurrency mining resources in compromised tenants. Many of

these attacks take advantage of automation, which increases the potential threat to cloud environments.

Cryptocurrency mining using central processing unit (CPU) or graphics processing unit (GPU) compute in cloud environments is not financially viable if one is paying for the compute used. In order to profit, threat actors use malicious methods to avoid paying for the resources, such as abusing free trials or compromising legitimate tenants to conduct cryptojacking attacks.

Unlike free trial abuse, which the cloud provider may be able to detect, cryptojacking in compromised tenants is more challenging to identify since it involves the threat actor having access to a legitimate user account. This complex method impacts the user more directly, as it allows the threat actor to make more intrusive changes in the target environment:

- Utilize available compute quota from compromised tenants, and provision significantly more compute and other additional resources.
- Mask resource provisioning activity as legitimate when operating within a compromised tenant.
- Use access to the compromised tenant to do further lateral movement, achieve persistence, and conduct information theft.

Successful cloud cryptojacking attacks could result in significant unexpected charges to the compromised tenant and depletion of resources that the tenant might need for business continuity, potentially resulting in service interruption, highlighting the need to prevent, detect and mitigate cloud cryptojacking attacks.

## Attack lifecycle

Cryptojacking requires the threat actor to reach a certain level of access to the cloud environment, which we explain in more detail in the next sections. The diagram below shows the stages of a typical cloud cryptojacking attack.

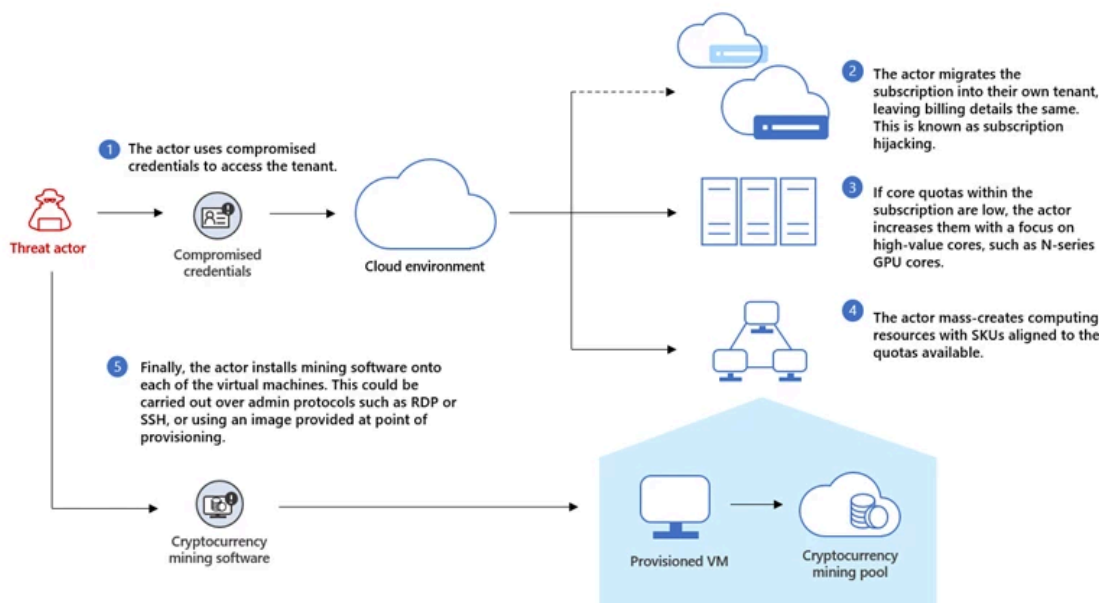


Figure 1. Diagram of cryptojacking attack on a compromised cloud tenant



After gaining access to the tenant and performing reconnaissance to determine available permissions, the attacker may proceed to hijack the subscription. Subscription hijacking has been covered previously in the blog entry [Hunt for compromised Azure subscriptions using Microsoft Defender for Cloud Apps](#).

Subscription hijacking is an evasion technique that allows the threat actor to hide some of their activities from the tenant administrator and security teams. Migrating a subscription directory requires the threat actor to have sufficient privileges in the target subscription. In cases observed by Microsoft, the destination tenant may be attacker-controlled or another affected tenant that the threat actor has access to.

Additionally, subscription hijacking is disruptive forensically. Microsoft Incident Response has observed instances where a threat actor compromised accounts in customer environments that were over-privileged. Abusing over-privileged accounts allowed the threat actor to migrate the subscription to a separate tenant (often attacker-controlled) to spin up additional resources. While activity logs at the subscription level remain with the subscription, anything recorded at the tenant role-based access control (RBAC) level is recorded in the new tenant, making forensic analysis, understanding the full timeline, or incident response by or for the customer, more challenging.

## Impact: Increasing core quotas

Once a threat actor has access to a tenant, they can either create compute using existing core quota, or they may choose to increase core quotas within the tenant. Increasing core quotas is potentially risky for the actor as quota increases undergo review. Some quotas can't be immediately adjusted and require a support ticket to increase.

Threat actors without permission to increase quotas use whatever is available. This often leads to them exhausting available core counts across multiple regions. Quota increases have occurred up to a month before resources are deployed by the threat actor.

GPU compute offerings are often targeted by threat actors. GPU compute provides access to high performance NVIDIA and AMD GPU cores, allowing cryptocurrency mining magnitudes more effective than any CPU compute offering. A complete overview of GPU compute types can be found in [GPU optimized virtual machine sizes](#).

The NVIDIA T4, V100, and A100 GPU compute options are most abused by threat actors. At time of writing, the NVIDIA A100 is the best mining card available that is not a dedicated application-specific integrated circuit (ASIC). When comparing NVIDIA GPU performance for cryptomining, the number of Compute Unified Device Architecture (CUDA) cores can be used as a rough representation of the card's performance. CUDA is designed specifically for high performance parallel computing, which allows more computations to take place at once. For NVIDIA GPUs, more CUDA cores generally means more mining potential. The table below shows the comparative hash rate for the top three most abused GPU compute cards within cloud environments based on mining [Ethereum Proof of Work \(ETHW\)](#).

Azure VM versions	GPU	CUDA cores	ETHW*
NC T4 v3	NVIDIA T4	2,560	25.1MH/s

NCv3	NVIDIA V100	5,120	89.5MH/s
ND A100 v4	NVIDIA A100 (40GB)	6,192	175MH/s

\* Mining rates based on the Ethereum Proof of Work complexity in February 2023

As the table above shows, threat actors who can provision NVIDIA GPU cores can mine a meaningful amount of currency in a relatively short period of time. In attacks observed by Microsoft, cryptojacking activities were seen to incur compute fees more than \$300,000, illustrating how unprofitable mining is within cloud environments without committing resource theft.

## Impact: Deploying compute

There are several ways to deploy compute, and threat actors have adapted to abusing features to speed up deployment. As resource hijacking is an attack of scale, the threat actor needs a way to rapidly spin up and manage multiple devices. In observed cases, threat actors have employed [VM scale sets](#), [Azure Machine Learning compute instances](#), [Azure Batch](#), and [Azure Container Instances](#). Each of these systems allows compute to be deployed quickly and centrally managed.

Malicious provisioning behavior of compute using the above methods generally does not match existing compute provisioning patterns within the tenant. The graph below shows an attacker deploying NVIDIA compute cores within a target environment using VM scale sets. The Y axis shows the capacity of the VM whilst the X axis represents time, this activity spans a three-hour period. Each color represents a single region, with the attacker iterating the various regions to create compute.

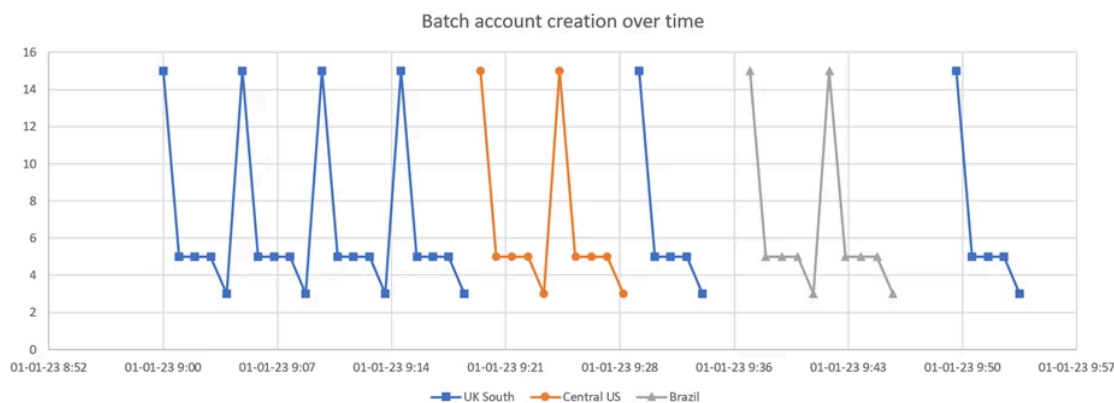


Figure 3. Attacker compute deployment pattern

In the graph above, the actor followed a predictable and anomalous deployment pattern across several hijacked subscriptions. Microsoft Threat Intelligence analysis shows that this deployment pattern is unique to a specific threat actor. While this specific pattern may change, the automated nature of malicious compute deployments means that an unusual pattern almost always emerges.

Some staggering of deployment is used, but the threat actor ultimately needs to provision compute very quickly to make the attack profitable. This time restriction means that patterns in provisioning generally emerge over relatively short periods of time. In the above case, the entire provisioning stage of the attack took place over a three-hour period.

In addition to the pattern of deployment, in this case, the following additional anomalies were also observed:

- The user accounts used to provision compute had never provisioned compute before.
- The compromised user provisioned GPU compute, when no GPU compute had been provisioned in this environment before.
- Compute was deployed to regions anomalous for the environment.

Other cases observed by Microsoft showed the following deployment anomalies:

- A user with a recent Azure AD anomaly creating large volumes of compute.
- A user suddenly causing multiple deployment failures spanning multiple core types due to a core quota unavailability.

Other than VM scale set deployment patterns, the same anomalous patterns can be identified within other automated deployment services such as [Azure ML compute instances](#), [Azure Batch](#), and [Azure Container Instances](#).

## Impact: Mining cryptocurrency

Once compute resources are deployed, the actor may need to install GPU drivers to take full advantage of the graphics card, especially on N-series VMs. Actors have been observed abusing Azure Virtual Machine extensions such as an NVIDIA GPU Driver Extension for [Windows](#) or [Linux](#), or an AMD GPU Driver Extension for [Windows](#), to facilitate driver installation. These extensions allow for the mass-deployment of drivers, reducing the threat actors' setup time before mining.

The following anomalies have been observed when actors use these extensions:

- Sudden or unusual high-volume provisioning of GPU drivers using a GPU Driver Extension.
- A user account suddenly deploying GPU extensions, especially where that user account has no history of deploying VM extensions.

With compute prepared, the threat actor can begin mining cryptocurrency by deploying mining software to the newly created VMs. The installed mining software joins the VM to a mining pool, which allows the threat actor to pool their stolen processing power from multiple compromised tenants.

Data from Microsoft Defender for Cloud shows some of the most recent pools in use by threat actors using already-compromised Azure tenants. Below is the list of the top 10 mining domains observed being used:

1. nanopool[.]org
2. nicehash[.]com
3. supportxmr[.]com
4. hashvault[.]pro
5. zpool[.]ca
6. herominers[.]com
7. f2pool[.]com
8. minexmr[.]com

9. monerocean[.]stream

10. miner[.]rocks

Seeing connections to any mining pool from a VM within an environment is a strong indication of compromise. Microsoft Defender for Cloud has multiple detections for this behavior.

## Recommendations to identify and mitigate cryptojacking attacks

Security teams should monitor and regularly review alerts specific to these scenarios. In environments where the creation of compute or increases in quota are uncommon, additional alerts should be built to monitor associated operations within your SIEM tool like Microsoft Sentinel. These are highly environmentally specific.

While every situation is unique to the customer and their environment, Microsoft Incident Response has identified several recommendations that are broadly applicable to help identify and mitigate cryptojacking attacks, alongside specific product detections. These recommendations are based on observations from responding to multiple resource abuse engagements.

- **Separation of privileged roles:** Keep administrator and normal user accounts separate. Non-administrator users who require privileged roles in the environment for specific functions should utilize [Privileged Identity Management](#) to access the roles on an as-needed basis in a way that can be audited and tracked, or also have separate accounts created. In most resource abuse cases Microsoft Incident Response has investigated, the initially compromised user is over privileged in some way. Thus, it is good practice to limit the number of accounts that have the [virtual machine contributor](#) role. In addition, accounts with this role should be protected by MFA and [Conditional Access](#) where possible. Also, since a global admin must enable the [elevate access](#) option to have permissions over all Azure resources, it should be considered a very sensitive activity that should be monitored and reviewed.
- **Multifactor authentication:** Tenant administrators should ensure that [MFA](#) is in use comprehensively across all accounts. This is especially important if the account has virtual machine contributor privileges. Users should also be discouraged from reusing passwords across services. Microsoft Defender for Cloud provides a range of recommendations to secure cloud environments. A full list can be found in [Security recommendations – a reference guide](#).
- **Risk-based sign-in behaviors and conditional access policies:** In cases investigated, attackers who have signed in using compromised credentials have triggered high Azure Active Directory (Azure AD) risk scores. Monitoring risky user alerts and tuning detections that take advantage of this security information help prevent these attacks. In addition to analyzing Azure AD risk scores, correlating risky Azure AD behavior with follow-on activity can help produce additional true positive detections. Risk-based conditional access policies can be designed to require multifactor reauthentication, enforce device compliance, force the user to update their password, or outright block the authentication. In many cases, policies such as these can be disruptive enough to provide security teams with enough time and signal to respond or alert the legitimate user to an issue before the resource abuse begins. Standard login anomaly detections were also found applicable in cases investigated by Microsoft Incident Response, with threat actors commonly using proxy services, signing in from anomalous locations, and accessing accounts using anomalous user agents. One group of activity tracked by Microsoft Threat Intelligence used Python requests and the default user agent (*python-requests/2.26.0*) for all operations.

Microsoft 365 Defender uses detections such as *Access elevation by risky user* and *Risky user performed suspicious Azure activities*, which correlate users marked as risky by Azure AD with anomalous actions to raise the severity of alerts in Microsoft 365 Defender.

Lastly, authentication to a tenant from an IP that is outside of that tenant should be anomalous. Defenders can identify which IP addresses are allocated within a tenant using the [az vm list-ip-addresses](#) command.

- **Limit unused quota and monitor for unexpected quota increases:** Looking for multiple unexpected quota increases occurring in a short period of time, quota increases across multiple regions, or quota increases within regions that the environment does not normally use might allow for early detection of a resource abuse attack. Quota increases are one of the first signals Microsoft Incident Response looks for when investigating suspected resource abuse attack. Quota increase detections can potentially be refined by looking for increases to commonly abused core types, especially if their usage is otherwise rare in an environment.
- **Monitor for external Azure IP addresses authenticated with your tenant:** Threat actors performing these attacks also use Azure compute resources to conduct their operations. Monitoring for successful sign in activity from Azure IP addresses that are not owned by your tenant is often a strong indicator of suspicious activity. Seeing multiple authentication attempts from Azure IP addresses using the same browser user agent is another strong indicator of potential password guessing.

## Detection details

### Microsoft 365 Defender

Microsoft 365 Defender is becoming Microsoft Defender XDR. [Learn more.](#)

Microsoft 365 Defender uses its cross-workloads detection capabilities to provide enhanced protection against cryptocurrency mining attacks. Microsoft 365 Defender customers who have enabled their [Azure connector](#) in Microsoft Defender for Cloud Applications can benefit from the following alerts:

- Access elevation by risky user
- Suspicious Azure activities related to possible cryptocurrency mining
- Mass provisioning of GPU virtual machines for possible cryptocurrency mining
- Suspicious creation of multiple Azure ML clusters and workspaces
- Suspicious role assignment in Azure subscription
- VM quota modified after risky user signed in

### Microsoft Defender for Cloud Applications

The following Microsoft Defender for Cloud Application alerts indicate threat activity related to the attack discussed in this post:

- Multiple delete VM activities
- Multiple VM creation activities

### Microsoft Defender for Cloud

Microsoft Defender for Cloud detects threat components associated with the activities outlined in this article with the following alerts:

- Azure Resource Manager operation from suspicious proxy IP address
- Crypto-mining activity
- Digital currency mining activity (Preview)
- Fileless attack toolkit detected
- Possible Cryptocoinminer download detected
- Process associated with digital currency mining detected
- Potential crypto coin miner started
- Suspicious Azure role assignment detected (Preview)
- Suspicious creation of compute resources detected (Preview)
- Suspicious installation of a GPU extension was detected in your virtual machine (Preview)
- Suspicious invocation of a high-risk 'Execution' operation by a service principal detected (Preview)
- Suspicious invocation of a high-risk 'Execution' operation detected (Preview)
- Suspicious invocation of a high-risk 'Impact' operation by a service principal detected (Preview)
- Suspicious invocation of a high-risk 'Impact' operation detected (Preview)
- Suspicious subscription transfer to external tenant was detected (Preview)

## Microsoft Defender for Endpoint

The following Microsoft Defender for Endpoint alert can indicate associated threat activity:

- Possible cryptocurrency miner

## Hunting queries

### Microsoft Sentinel

Microsoft Sentinel customers can use the TI Mapping analytics (a series of analytics all prefixed with 'TI map') to automatically match the malicious domain indicators mentioned in this blog post with data in their workspace. If the TI Map analytics are not currently deployed, customers can install the Threat Intelligence solution from the Microsoft Sentinel Content Hub to have the analytics rule deployed in their Sentinel workspace. More details on the Content Hub can be found here:

In addition, Microsoft Sentinel customers can leverage the following content to hunt for and detect related activity in their environments:

- [Crypto currency miners](#)
- [Suspicious cryptocurrency mining related threat activity detected](#)
- [Detecting Anomaly Sign-In Event](#)
- [Administrator Authenticating to Another Azure AD Tenant](#)
- [Creation of an anomalous number of resources](#)

## Appendix

Top 10 mining domains used by threat actors:

1. nanopool[.]org
2. nicehash[.]com
3. supportxmr[.]com
4. hashvault[.]pro
5. zpool[.]ca
6. herominers[.]com
7. f2pool[.]com
8. minexmr[.]com
9. monerocean[.]stream
10. miner[.]rocks

## Further reading

For the latest security research from the Microsoft Threat Intelligence community, check out the Microsoft Threat Intelligence Blog: <https://aka.ms/threatintelblog>.

To get notified about new publications and to join discussions on social media, follow us on Twitter at <https://twitter.com/MsftSecIntel>.

---

Source: <https://www.microsoft.com/en-us/security/blog/2023/07/25/cryptojacking-understanding-and-defending-against-cloud-compute-resource-abuse/>