

Cloud Security – How We Are All Targets

You cannot afford to get it wrong. It is 15 years since Amazon / AWS launched their first cloud service and even after we have witnessed malicious actors launch the largest ransomware attacks, data breaches / theft and network infiltration, there is still a believe that “I won’t be the next target” or “who would target us?”.

Exposed Cloud Resources Cause Almost Instant Attacks

Across Amazon Web Services (AWS) and Microsoft Azure we found that common misconfigurations resulted in almost instant attacks to their systems.

Misconfigurations in this context refer to any configuration that results in a cloud system to be unintentionally accessible by 3rd parties. This can range from a cloud administrator / developer making a legitimate mistake to even just accepting default configurations by the cloud service provider.

Public cloud service provider IP address ranges are well-known and constantly scanned for vulnerable endpoints by individuals and bot networks.

100% of all misconfigurations we tested that caused cloud systems to be publicly available resulted in 3rd parties attempting to access those systems within 1.5 minutes of being deployed. These attempts were not targeted at us specifically but were non discriminating attacks.

Everybody is a Target

What are the top attacks we found?

- RDP brute-forcing of Windows Virtual Machines
- IIS / Apache exploits
 - where an unauthenticated webserver port was detected
- Database vulnerabilities
 - Where databases were detected to be accessible on the internet

These attacks can easily result in data loss / data exfiltration or potentially even privilege escalation and continued network traversal especially with unpatched or insecure systems.

This list is not complete; however, it highlights that many of the attack vectors are, simply put, basic techniques. They rely on the fact that often basic security hygiene is not a given, especially in cloud-based environments.

This and most data breach reports demonstrate that most cloud infrastructure does not have sufficient real-time security monitoring and shows the critical requirement for this capability.

The Cloud – Blessing and Curse

ARGOS researchers interviewed dozens of industry experts, analysts, and ARGOS customers to understand if our results match what they see. One question we asked everybody, as it was a question we are also being asked regularly:

What is the biggest security problem you see in the cloud?

The most common concern amongst security professionals was a fear of someone unintentionally exposing internal resources to the internet. The examples they cited were the "public AWS S3

Buckets", databases with public IPs or team members "just getting things done" and missing certain check boxes that result in a resource becoming publicly available.

Our analysis shows the cause of these common errors occur when employees copy and paste code or follow tutorials from pages like StackOverflow, community blogs and even vendor documentation. Unfortunately, most of these resources are geared towards usability or speed and often rely on these individuals to apply security practices later. This is not to say that employees make insecure configurations intentionally, but that deploying secure cloud environments is not straightforward or the priority at the time.

Point in time or one-off checks before deployment, i.e., code-reviews or static-code-analysis (SCA) as part of deployment (CI/CD) pipelines, are good, but not sufficient if environments are not continuously and automatically tested. Humans are great at working around roadblocks and will always find a way to "get things done".

Interviews ARGOS conducted with security professionals found that it was common for team members to manually deploy infrastructure instead of using the prescribed deployment pipelines, turned specific security features off or they even deployed using the prescribed methods and then changed configuration afterwards.

Our findings show a lack of visibility often means that organisations do not even know that this is happening until months later when it is too late. Cloud platforms are complex with often thousands or even hundreds of thousands of resources deployed. Resources that keep changing with every new deployment, but also due to the cloud providers themselves updating their services. Automated monitoring of configuration is a non-negotiable in order to avoid "not knowing" about security vulnerabilities.

100% Likelihood That You Will Be Targeted

Our research has shown that across Microsoft Azure and Amazon Web Services (AWS), there is a 100% likelihood that if you run any public facing cloud systems you will be tested almost immediately and any weakness in infrastructure or application configuration will be leveraged in very likely next steps, without exception. This does not imply that you will be a victim of a cyber-attack, but rather that parties will use you as a target. We will talk about remediating factors later in this paper.

Our tests included the following test cases:

- Microsoft Azure
 - Deployment of a Windows Server 2019 Virtual Machine with a public IP and a Network Security Group allowing RDP inbound from the internet.
 - Deployment of a CentOS Server Virtual Machine with a public IP and a Network Security Group allowing SSH inbound from the internet.
 - Deployment of a SQL Database (PaaS - public endpoint by default).
- AWS
 - Deployment of a Windows Server 2019 EC2 with a public IP and a Security Group allowing RDP inbound from the internet.
 - Deployment of a CentOS Server EC2 with a public IP and a Security Group allowing SSH inbound from the internet.
 - Deployment of an RDS SQL Database (public endpoint enabled).

These test cases are the most common scenarios that Enterprises all around the world are leveraging to run applications in the cloud.

Cloud Infrastructure Paints a Bullseye on One's Back

In our tests it did not matter which cloud region we deployed a resource to, it also did not make a difference at what time of day we did it or particularly even which service we deployed. As soon as we created a cloud resource that used a public endpoint the resource was exposed to hundreds of access attempts each second.




Figure 1 Map of remote locations accessing our test instances

Note: We see many access attempts here from countries not necessarily associated with “malicious attacks”. There are many companies continuously scanning the internet for research or commercial purposes.



RDP Brute Forcing

Here is an output of usernames and the number of times they were used in order to log on to an Azure Windows Virtual Machine.

Account	Count
\ADMINISTRATOR	650345
\administrator	387911
\Administrator	76308
DOMAINCONTROLLE\Administrator	56669
IZZET\Administrator	14235
\Admin	6426
DOMAINCONTROLLE\Admin	6422
DOMAINCONTROLLE\Administrateur	5466
\Administrateur	5466
\admin	4800
\hp	3376
DOMAINCONTROLLE\administrateur	3372
DOMAINCONTROLLE\admin	3370
DOMAINCONTROLLE\domaincontrolle	3369
DOMAINCONTROLLE\administrador	3367
DOMAINCONTROLLE\admindomaincontrolle	3365
DOMAINCONTROLLE\domaincontrolleadmin	3365
DOMAINCONTROLLE\administrator	3357
\ADMIN	1352
\sysadmin	1194
DOMAINCONTROLLE\sysadmin	1192
\Administrador	1084
\User	1026
\de	1002
DOMAINCONTROLLE\DOMAINCONTROLLE	930
\usuario	914
\accounts	891
\for	882
domainname\administrator	866
domainname\Administrator	748
DOMAINCONTROLLE\Administrador	630
\behsazi	616
domaincontrolle\domaincontrolle	552
\Cuentas	546
\ROOT	456
\Invitado	444
DOMAINCONTROLLE\sadmin	399
\sadmin	399
\ASUS	397
\TEST	383
domain\Administrator	380
\AIO	378
\AZUREUSER	340
\USER	333
\COMPUTER	310
\MANISHA COMPUTER	308
\MANISHA	294
localhost\admin	280
localhost\Lab1	280
localhost\administrator	280
\NAGIOS	266
\MYSQL	217
\xxsoft	204

Figure 2 Windows usernames

Overall, over a time frame of 7 days a single Windows Virtual Machine that exposed its Management Port 3389 (RDP = Remote Desktop Protocol) via a public IP address, had **1,460,273** logon attempts. That is approximately **160 attempts per second**. This number was only slightly different on an AWS EC2 Windows Virtual Machine.

An example of these events can be seen here using AWS CloudWatch when searching the Windows Event Log for [EventId 4625](#).

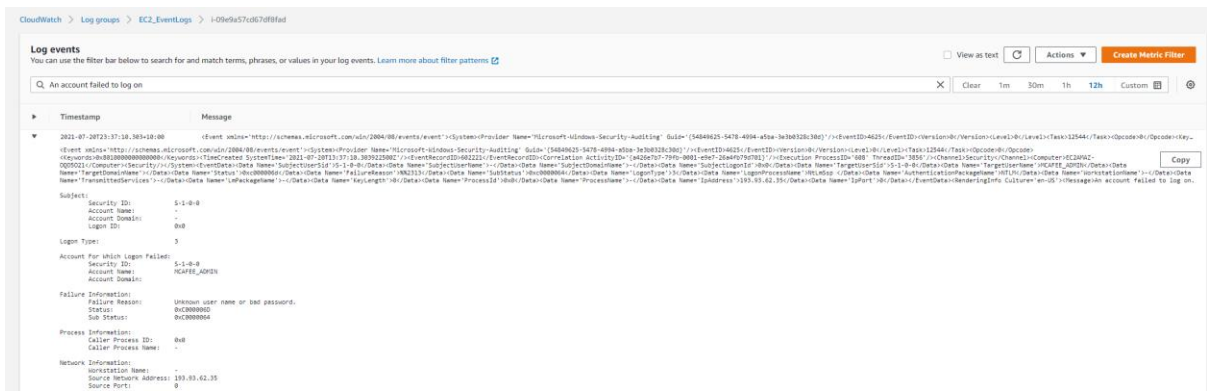


Figure 3 AWS CloudWatch Windows Event Log

In addition to being potentially disastrous if someone managed to enter the environment these logon attempts have other side-effects as can be seen here by looking at this Virtual Machine’s CPU metrics.




Figure 4 Virtual Machine CPU metrics

This otherwise idle Virtual Machine must react to all the logon attempts as can be seen in the CPU spikes. This can have a negative impact to applications otherwise running on that Virtual Machine or even have cost implications if cloud resources are configured to automatically scale up (deploy larger resources) or out (deploy more resources).

There are scenarios where an Operating System can also stop responding to legitimate logon attempts.

After 7 days the number of requests slowed down significantly with a peak after approximately 3 days.

reasons however, not many organisations make use of these services, especially early in their cloud journey, arguably the most critical time, and unintentionally and/or unknowingly expose themselves to many attacks.

Other Management Ports

The second most seen vector in our tests were other management ports attempted like the Windows “Server Management Block” (SMB) service most commonly on port 445.

In just 7 days we saw almost **200,000 attempts** to connect to the Windows SMB port.

SMB is often used by Ransomware crews. Remember [EternalBlue](#) and [EternalRomance](#)?

Microsoft also recommends to disable SMB where not required to ensure that even if the port was exposed that there was no service to connect to: <https://docs.microsoft.com/en-us/windows-server/storage/file-server/troubleshoot/detect-enable-and-disable-smbv1-v2-v3>

In addition to SMB, which is very Windows specific, another almost ubiquitous deployment scenario is (S)FTP ((Secure) File Transfer Protocol). Many organisations have processes that rely on the delivery of data via FTP. This sees many environments expose the FTP ports to the internet that only require a username and password to authenticate, even on Linux machines.

Why is it so easy for malicious actors?

On-premises in one’s datacentre, with a smaller perceived attack surface, one probably figured that if you did not tell anyone about a service then people would not be easily able to discover it. (Side note: This might not be absolutely true anymore with much more advanced internet-scale scanning happening.)

In the cloud however this is very different. In the cloud everybody can easily find anybody as IP address ranges for all cloud providers are freely available to read by everyone on the internet.

Microsoft Azure: <https://www.microsoft.com/en-au/download/details.aspx?id=56519>

Amazon Web Services: <https://docs.aws.amazon.com/general/latest/gr/aws-ip-ranges.html>

Google Cloud Platform: <https://support.google.com/a/answer/10026322?hl=en>

Those IP ranges are constantly scanned by both “good” and “malicious” actors for anything that is accessible from the internet. Once something is found it is immediately tested for more available entry points into the environment.

This means that typically no individual / organisation is a target but that everybody is. Of course, there are targeted attacks, but the takeaway here is that most attacks are like spam, hoping to find someone who responds.

What can you do?

It sounds simple, but basic cloud security hygiene needs to become a focus.

This does not mean that cloud engineers / developers must be utterly restricted in all their capabilities when building cloud environments and similarly cloud must not become the scapegoat.

Here are the first steps to a more secure cloud environment:

1. You cannot protect / secure what you do not know exists. A real-time inventory of assets must be created.
2. Identification of those cloud assets that are exposed to the internet for prioritised remediation.
3. Education of teams in basic cloud security practices. Allow teams to use the cloud to its fullest but put guidelines in place, then monitor and educate for adherence to those guidelines.
4. Continuous, automated monitoring of cloud assets for insecure configuration is vital.

ARGOS Cloud Security

ARGOS finds insecure configuration **WHEN**, not **IF**, they happen, which means exposed cloud resources will only be exposed for minutes instead of going undetected for days, weeks or even months.

In less than 20 minutes ARGOS is implemented, is ready to reduce any cloud environment's exposure and remediate cloud security issues in almost no time, even for teams just starting on their cloud journey.

Check ARGOS out at <https://argos-security.io/> and sign up for a free trial to welcome ARGOS into your team as your new **full time Cloud Security Expert**.