

2022

**Blumira's
State of
Detection &
Response**

Blumira

Table Of Contents

Introduction

Working Smarter, Not Harder

Our Approach To Findings

Time to Detect and Respond

The Impact on the Bottom Line

Our Findings At A Glance

Top 5 Findings Overall

Microsoft 365 Security

Top Microsoft-Related Findings

Creation of Office 365 security group

10 Windows user password reset attempts within 1 hour

Ps-Exec use on network

Modification of Microsoft 365 group

Clearing of Windows event logs

Business Email Compromise

Living Off The Land

The Problem With Living Off The Land

PowerShell: An Attacker's Favorite LotL Tool

5 Ways That Attackers Use PowerShell

PsExec: A Double-Edged Sword

Identity-Based Attacks

Common identity-based attacks

Hands In The Honeypot

How Blumira Can Help

Introduction

Threat actors can be evasive, clever, and complex — but fortunately for defenders, they are also predictable. There is only a limited number of methods to access an environment, and when an attacker finds a technique that works, they tend to reuse it.

That's not to say threat actors — especially those in state-sponsored, high-profile ransomware groups — aren't getting more sophisticated. Adversaries have the same access to endpoint detection software as customers do, and thoroughly test their attacks against them to hone their evasion techniques. More advanced attackers are always attempting to stay ahead of the curve by leveraging new exploit kits, vulnerabilities, or malware loaders.

Another concerning trend is the shortening of ransomware dwell time: the time it takes for an attack to complete, from initial access to exploitation. We're no longer seeing as many attacks in which adversaries lurk in an environment for weeks or months before exfiltrating data. Attacks happen quickly — and at inopportune moments, like holidays and weekends — and defenders, too, must work quickly to stop an attack in its early stages.

But there's good news. Although today's attacks may appear more sophisticated, the techniques, tactics and procedures (TTPs) used to launch those attacks remain the same. Adversaries often take the approach of working harder, not smarter; finding easy, low-cost and relatively simple methods to launch attacks.

By studying patterns in attacker behavior, we can better understand those methods — no matter how advanced — and detect them accordingly.

Behavior-based detection and signature-based detection are both valid approaches, but monitoring behavior can identify the paths that an adversary takes on the road to an attack — even if those behaviors seem legitimate. Focusing on attacker behavior and what initiated that behavior is a strong indicator of a potential threat or attack in progress.

As defenders, we're always interested to get inside of a threat actor's mind. Looking at their patterns in behavior is the closest way to achieve that.

Our Approach To Findings

Blumira's platform incorporates hundreds of different findings that detect suspicious behaviors that may indicate an attack in progress. This report is based on research from **33,911 key findings** from a sample including **230 organizations**, which took place over the course of 2021.

These 33,911 findings are filtered to exclude outliers and low-priority alerts that we considered less significant, including account lockouts and blocked websites. That's not to say that these alerts should be ignored, but we decided not to include them for the sake of accurate, relevant data.

To understand how we generate these findings, let's take a step back. Blumira's incident detection engineers (IDE) take an intentional approach to rule design to reduce alert fatigue.

1

First, our IDE team creates rules based on threat-based research, pulling data from various threat intel reports to determine how current threat actors operate.

2

Once the team emulates attacks in a lab environment, they identify and build detections based on the threat actors' behavior.

3

Then the detection is tested again across customer datasets to remove false positives, reducing noisy alerts to help customers focus on priority findings.

4

Blumira's platform stacks similar alert data to already-triggered findings until the case is closed, helping to prevent alert fatigue and providing all relevant evidence to assist with investigation.

But it's not enough to be able to detect and respond to an attack in progress. As attacks happen faster, security and IT teams must be able to **both detect and respond quickly** before real damage occurs.

Time To Detect and Respond

Time to detect and respond refers to the time it takes to identify a compromise and contain the threat (sometimes referred to as the “**breach lifecycle**”). It directly affects the bottom line of an organization, with the longer the breach takes to detect and contain, the higher the overall cost.

The Impact on the Bottom Line

In IBM/Ponemon’s 2021 Cost of a Data Breach report, they found that breaches that take longer than 200 days to resolve can result in **35% higher cost, from \$3.6 million to \$4.9 million on average.**

287 days

The total average breach lifecycle


(Source: IBM/Ponemon's 2021 Cost of a Data Breach)



Blumira's detection engine includes real-time, or instantaneous, individual findings that notify a customer almost immediately of a potential threat, such as detecting a virus on your network – the median time to detect for these types of findings is 50 seconds.

Threshold-based findings are based on a certain event happening multiple times over a set period of time. For example, in a password spraying attack, an attacker will attempt to log in by trying a large number of usernames with a single password, which can help evade detection. In this case, notification will happen only after the behavior is observed over a certain period of time.

The Cost of a Data Breach



DOWNTIME

Downtime and disrupted business operations means a loss in revenue, especially for companies without a disaster recovery plan. **Lost revenue due to downtime accounts for 38% of overall costs, according to IBM.**

DAMAGE TO REPUTATION

A ransomware attack can make customers feel uneasy, leading to damaged reputation, and subsequently, customer churn. **86% of people are less likely to deal with companies that experienced a data breach, according to a Semafone study.**



CUSTOMER COMMUNICATION

Companies must follow up with their affected customers after a ransomware attack, and cover costs related to credit monitoring and identity protection services.

LEGAL COSTS

If customer data was breached as a result of the ransomware attack, then companies must incur legal costs related to third-party claims.



The Cost of a Data Breach (cont.)



REMEDIATION

Remediation costs include implementing forensics and investigative work, as well as containing the actual breach.

Remediation costs grew from an average of \$761,106 in 2020 to \$1.85 million in 2021, according to Sophos.

COMPLIANCE FEES

Paying a ransom could breach OFAC regulations and result in needing to pay compliance fees on top of that ransom.



WHAT'S THE TOTAL?



So it's clear that the time to detect and respond has a major impact on your business, and may be devastating for smaller organizations that have less resources to help them recover from lost revenue.

Small and medium-sized businesses (SMBs) that experienced a data breach in 2021 suffered costs of \$2.98 million, according to IBM.

Time to Detect

32 min

Blumira's average time to detect a finding

(Source: Blumira's 2021 dataset)



212 days

Average time to detect a breach

(Source: IBM/Ponemon's 2021 Cost of a Data Breach)

99.4% faster

Time to Respond

6 hours

Average time to respond, or how quickly a customer closed findings

(Source: Blumira's 2021 dataset)



75 days

(or 1,800 hours) Average time to respond to a threat

(Source: IBM/Ponemon's 2021 Cost of a Data Breach)

99.7% faster


Our Findings at a Glance

We've analyzed and compiled the top findings based on our data. Those top findings highlight a few trends, some of which we'll delve more deeply into later on in this report.

Here are some trends we've witnessed:

Microsoft 365 activity

Our findings revealed patterns of Microsoft-related activity, including activity associated with password spraying, lateral movement, and business email compromise.



Living off the Land

Adversaries are using built-in tools such as PsExec and PowerShell to evade detection from traditional security tools.

Identity-based attacks

Cloud environments are particularly vulnerable to identity-based attacks such as credential stuffing, phishing, password spraying and more.



Top 5 Findings Overall

Blumira

#5

50 GB+ Inbound Connection via Generic Network Protocol

MITRE ATT&CK technique: Data Exfiltration

What does it mean? This can indicate a business-related connection or data exfiltration. Depending on the protocol it may be important to consider the security of the connection if this is business related traffic. It is recommended to correlate with the source to determine if this is an expected connection as well.



#4

Admin-Level Account Added

MITRE ATT&CK technique: Persistence: Account Manipulation

What does it mean? It's uncommon for a threat actor to add an admin-level account, but it's important for IT and security teams to audited and validate each creation of an admin-level account when they occur to avoid scope creep or attackers gaining access.

#3

Service Execution with Lateral Movement Tools

MITRE ATT&CK technique: Execution: System Services

What Does It Mean? The Windows service control manager (services.exe) can enable threat actors to execute malicious commands or payloads via a temporary Windows service.



#2

Okta Log Failure

MITRE ATT&CK technique: n/a

What Does It Mean? Okta logs aren't flowing properly to your SIEM, meaning you may have a gap in detection coverage. It's important to be aware of IT operational failures for both compliance and security.



#1

Honeytrap HTTP Authentication Attempt

MITRE ATT&CK technique: Credential Access

What Does It Mean? Someone is actively attempting to access your honeypot and is unaware of its nature.



Microsoft Security

Blumira


Microsoft 365 Security

The most popular cloud collaboration tool is also highly targeted by attackers-- so how can small & mid-sized businesses protect themselves?

Over 260 million monthly users

As organizations move from on-prem Exchange to cloud-based M365, it's quickly become one of most used applications across the world.

(Source: Microsoft)




71%

Most exploited

Microsoft 365 is the top most commonly exploited application worldwide in Q3 2020


(Source: statista.com)



70% suffered an average of 7 account takeovers

In 2020, from a survey of over 1,000 IT security decision-makers using Microsoft 365

(Source: Vectra AI)



95% of breaches were financially motivated

Business email compromise (BEC) is the second most common social engineering attack targeting cloud-based email servers.

(Source: 2021 DBIR)

#5

Clearing of Windows Event Logs

MITRE ATT&CK technique: Defense Evasion

What Does It Mean? An insider or threat actor may be attempting to clear evidence to cover their tracks after malicious activity.



#4

Modification of Microsoft 365 Group

MITRE ATT&CK technique: Persistence: Account Manipulation

What Does It Mean? A threat actor using an admin account can modify a Microsoft 365 group to add users or grant additional permissions, resulting in data leakage and access by unauthorized users..



Ps-Exec use on network

MITRE ATT&CK technique: Lateral Movement: Remote Services

What Does It Mean? An attacker may be moving laterally within your environment and interacting with remote machines using compromised credentials.

#3



#2



10 Windows user password reset attempts within 1 hour

MITRE ATT&CK technique: Lateral Movement

What Does It Mean? An attacker may be moving laterally throughout your environment and attempting to reset passwords for other accounts.

#1

Creation of Microsoft 365 security group

MITRE ATT&CK technique: Persistence: Account Manipulation

What Does It Mean? Someone that creates a security group can grant members of that group access to certain things, such as a SharePoint site. This may lead to insider risk or elevation of privileges.



Business Email Compromise

Microsoft is the top three of brands impersonated in BEC attacks, according to Abnormal Security.

Over the last decade, as organizations moved from on-site email systems to cloud-based, scammers have adapted. Small and medium-sized businesses (SMBs) as the most vulnerable to these types of scams due to lack of resources and being priced out of most defensive security solutions.

\$2.1 billion

In losses from cloud-based BEC scams between 2014 and 2019, according to the FBI.



What is Business Email Compromise?

Business email compromise (BEC) is when a threat actor uses social engineering and impersonation to trick employees into sending payments or sensitive data to their accounts.

One example is the impersonation of an executive, sending an email to an employee asking for gift cards or wire payments. With a legitimate-looking domain name or compromised email account, a recipient may be fooled into fulfilling the request out of a sense of urgency. BEC is basically a form of phishing that involves sending money directly to fraudulent accounts.



Security Recommendations

If your organization uses Microsoft 365, it's likely that a lot of data flows in and out of it — making it a prime target for attackers.



Ensure you can detect suspicious activity such as creating inbox rules or external email forwarding rules.



Check for MFA misconfigurations or instances of MFA being disabled



Monitor your Microsoft 365 environment to be able detect threats early enough to stop an attack

Signs of Business Email Compromise

Activity from suspicious IP addresses



Disabling of MFA

Enabling external email forwarding



Mass downloading of files

Living off the Land

Blumira

What is Living off the Land (LotL)?

Living off the land techniques involve using tools that already exist within a system to conduct attacks. Many of these tools are used by sysadmins for legitimate work, making it difficult for defenders to distinguish between malicious behavior and an admin simply doing his or her job.



62%

of detections were malware-free in CrowdStrike's 2022 Global Threat report.

Why Do Attackers Use LotL?

- **Low cost.** These attacks take advantage of tools that already exist within an environment, so attackers don't need to buy or create malware or attack tools, saving money and time. You can't get better than free.
- **Easy and simple.** No need to build, test, and use tooling, which creates obstacles for adversaries wanting to launch attacks quickly.
- **Avoid detection.** A lack of malicious tools and files means a lack of signature (or known-bad behavior recognized by many security tools), making detection difficult.

The Problem With Living off the Land



Living off the land behaviors often take place over a period of days or weeks, and during this time, an attacker can go undetected by endpoint detection tools because the attacker is not using anything that is known to be malicious.

This means that endpoint detection and response (EDR) tools may have a hard time detecting attacker behavior until it is too late — for example, when an attacker introduces malware into the environment.

Even when an EDR tool does alert on questionable behavior, it's very easy for an admin to miss or dismiss an alert that looks like normal behavior without additional questionable behavior identified from other IT and security systems that provide context. A single agent alerting on a single machine often isn't enough visibility and context to stop savvy attackers.

Top LotL Techniques

Service Execution with Lateral Movement Tools

The Windows service control manager (services.exe) can enable threat actors to execute malicious commands or payloads via a temporary Windows service.

Psexec is a command-line tool in Windows that lets privileged users execute processes on remote systems and redirect console applications' output to the local system so that these applications appear to be running locally.

Attackers use it for the same reasons, providing a convenient way to move laterally and interact with remote machines using compromised credentials. Only authorized users should be utilizing Psexec on the network.

Psexec Use

Threat actors can use Psexec maliciously to move laterally throughout your network, to execute commands or payloads, or to conduct remote execution.

Potentially malicious PowerShell command

PowerShell is like the swiss-army knife of tools, enabling adversaries and admins alike to perform a variety of tasks.

.NET User: Recon commands

Microsoft's Net user command utility allows for queries about both local users and domain users. While useful for systems administrators, it is often used by malware, and hands-on threat actors as an unobtrusive way to begin discovery in an environment.

PowerShell: An Attacker's Favorite LotL Tool

PowerShell is one of the most powerful tools to control a Windows machine from within. Only necessary users should have the ability to use PowerShell. Each additional user opens up another security gap, enabling attackers to have an elevated foothold in your network as soon as they're able to access one of those users, hosts, or sessions.

208%

increase in PowerShell threats in Q4 of 2020

(Source: McAfee)

5 Ways That Attackers Use PowerShell

Execute local scripts

Inject malicious code into memory

Install PowerShell scripts as services

Encode payloads

Execute code without admin access

Security Recommendations

Detecting living off the land techniques requires an understanding of what legitimate behavior looks like in your environment.



Pare down access to PowerShell to only the necessary users can help more easily determine your organization's definition of normal PowerShell activity.



Once you establish a baseline, you can more easily identify spikes in activity and abnormalities that may indicate an attack in progress.



Combine EDR tools — that may miss LotL techniques — with a behavior-based detection approach.

Identity-Based Attacks

Blumira

Identity-Based Attacks

The pandemic forced many organizations to move to cloud services to support their remote employees. For organizations without a solid understanding of their exposed attack surface, moving to a cloud environment only highlighted that knowledge gap.

In identity-based attacks, threat actors take advantage of those knowledge gaps by exploiting, misusing, or stealing user identities.



80%

of breaches are
identity-driven

What Makes Cloud Vulnerable to Identity Attacks?

- **Lowered visibility** into employee actions
- **Cloud misconfigurations**, i.e. leaving an unencrypted data store exposed to the public internet without requiring authentication, or failing to apply the least privilege principle
- The sheer **volume of identities** in the cloud means that identity and access management policies are harder to manage

Identity-Related Findings

We found that identity-driven techniques were common; 3 out of Blumira's top 5 findings (60%) were identity-related:

Honeypot Authentication Attempt

#1

This indicates that a user is actively attempting to access a honeypot and is unaware of its nature.

#2

Okta Log Failure

An Okta log failure means that logs aren't flowing properly from an Okta instance, which greatly increases your risk of an identity-based attack.

Admin-Level Account Creation

Someone creating a new admin-level account should always be monitored, since admin-level access can be the keys to the kingdom for an adversary.

#4

Identity-related Findings

Other identity-related findings we observed included:

Pass-the-hash behavior

Attackers may attempt to capture a password hash, exploiting the authentication protocol.

10 Windows user password reset attempts within 1 hour

During a breach, an attacker will often attempt to move laterally through accounts with access to other resources such as shared drives, servers, etc. When a malicious user doesn't know what the current password policy is, they may attempt an invalid password reset for other accounts.



Common Identity-based Attacks:

- Password spraying
- Credential stuffing
- Man-in-the-middle attacks
- Phishing

Hands In The Honeytrap

Attempts to authenticate into a honeypot was Blumira's #1 finding of 2021.

What's a honeypot?

A honeypot lures attackers with a network device that appears to contain valuable data. Once an attacker tries to log in, scan the device, or attempts to access a file on the device – the honeypot will notify your team.

Types of Honeypots

- **Honeynet** – A collection of honeypots and other deception techniques.
- **Honeytoken** – A piece of data that is used to lure in an attacker, such as API keys, database entries, executable files, and keys to cloud resources (e.g. AWS key).
- **Honeycred** – A username or ID that is used to identify specific types of attacks on systems.
- **Honeyport** – A job that listens on specific TCP Ports. When a connection is established, it can either simply log or add a local firewall rule to block the host from further connections.

Security Recommendations

As identity-based threats become more common – especially for cloud services – aim to get more visibility into your environment:



Enable multi-factor authentication to reduce the risk of unauthorized access due to credential compromise.



Limit domain access to small groups to limit exposure and lower your chances of a malicious actor gaining access to domain accounts.



Use honeypots to stay one step ahead of attackers and to be aware of potential intruders

How Blumira Can Help

Blumira

Easy, Effective Security

Blumira makes detection and response easy and effective for small and medium-sized businesses (SMBs) so they can respond to threats faster to prevent ransomware and data breaches.

✓ Fastest time to security

IT admins can deploy Blumira's platform in minutes to hours for broad on-premises & cloud coverage.

✓ Reduce noise, focus on critical threats

Prioritize your team's time with Blumira's behavior-based detections, fine-tuned to reduce noisy false-positive alerts.

✓ Faster, effective response

3-step rapid response: Block threats with automated response, follow our playbooks or contact Blumira's SecOps team for support

✓ Predictable, per-user pricing

Blumira is 40% more affordable than the industry average and doesn't charge based on data volume

Detect Real Threats

Blumira's behavior-based detections notify you of:

- **Cloud threats** - Misconfigurations, compromised cloud instances, C2 server communication
- **Identity-based attacks** - Anomalous logins, fraudulent MFA attempts, brute-force attacks
- **Email & document risks** - External doc sharing, email forwarding, new inbox rules
- **Endpoint security threats** - Malware/ransomware, blocklisted apps, compromised processes



Simple Security For IT Admins

*"We chose Blumira for its **simplicity** - I needed a solution that would simplify, consolidate and show me what I really need to see."*

- Jim Paolicelli, IT Director, Atlantic Constructors



See all available integrations: blumira.com/docs

Sign Up Free!
blumira.com/free

Blumira's Free Edition

Sign up for free, no credit card or sales conversation required.

With Free edition, you'll get:

- Coverage for unlimited users and data for Microsoft 365
- Easy cloud SIEM setup in minutes with Cloud Connectors
- Detections automatically activated, fine-tuned to reduce noise
- Summary dashboard of key findings & basic reports
- Playbooks to guide you through response steps
- 7 days of log data retention (upgrade for 30 days or one year)




Detect M365 Threats:

Get notified whenever Blumira detects:

- Attacker activity
- Ransomware & malware
- Unusual behavior
- User & access security

With over 38 detection rules and 21 security reports available for free, **Blumira does the heavy lifting for you** to keep your organization protected against the latest threats.



Get Security Support, Broader Coverage & Advanced Visibility

Upgrade to any unlimited data, paid edition to unlock 24/7 SecOps support for critical priority issues, in addition to:

- **Microsoft 365** - 30 days of log data retention, advanced reporting & security, manager & responder dashboards
- **Cloud** - Coverage for Microsoft 365, Duo Security, SentinelOne and one year of data retention
- **Advanced** - Unlimited integrations for all on-prem & cloud services, automated response, honeypots and one year of data retention

Blumira provides predictable, per-user subscription-based pricing, with no hidden fees or licensing required -- and no data caps.

See which plan is right for you: blumira.com/pricing

Easy for MSPs & SMBs

"I just finished setting up Blumira, and one word: WOW! I like the simplicity of your product and the easy-to-follow instructions for setting up logging. I am sold on Blumira's ease of use and capabilities."

-- Amitaf DaSilva, Principal, CompuNET Consulting (MSP)

Sign Up Free!
blumira.com/free