

# Microsoft Security Intelligence Report

Volume 21 | January through June, 2016

*Protecting cloud infrastructure:  
Detecting and mitigating threats using  
Azure Security Center*





This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Copyright © 2016 Microsoft Corporation. All rights reserved.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

## Authors

|  |   |   |
|--|---|---|
| <b>Charlie Anthe</b><br><i>Cloud and Enterprise Security</i> | <b>Michael Johnson</b><br><i>Windows Defender Labs</i>          | <b>Siddharth Pavithran</b><br><i>Windows Defender Labs</i>                          |
| <b>Evan Argyle</b><br><i>Windows Defender Labs</i>           | <b>Jeff Jones</b><br><i>Corporate Communications</i>            | <b>Daryl Pecelj</b><br><i>Microsoft IT Information Security and Risk Management</i> |
| <b>Eric Douglas</b><br><i>Windows Defender Labs</i>          | <b>Tim Kerk</b><br><i>Windows Defender Labs</i>                 | <b>Ferdinand Plazo</b><br><i>Windows Defender Labs</i>                              |
| <b>Sarah Fender</b><br><i>Azure Security</i>                 | <b>Mathieu Letourneau</b><br><i>Windows Defender Labs</i>       | <b>Tim Rains</b><br><i>Commercial Communications</i>                                |
| <b>Elia Florio</b><br><i>Windows Defender Labs</i>           | <b>Marianne Mallen</b><br><i>Windows Defender Labs</i>          | <b>Paul Rebriy</b><br><i>Bing</i>   |
| <b>Chad Foster</b><br><i>Bing</i>                            | <b>Matt Miller</b><br><i>Microsoft Security Response Center</i> | <b>Karthik Selvaraj</b><br><i>Windows Defender Labs</i>                             |
| <b>Ram Gowrishankar</b><br><i>Windows Defender Labs</i>      | <b>Chad Mills</b><br><i>Safety Platform</i>                     | <b>Tom Shinder</b><br><i>Azure Security</i>   |
| <b>Volv Grebennikov</b><br><i>Bing</i>                       | <b>Nam Ng</b><br><i>Enterprise Cybersecurity Group</i>          | <b>Nitin Sood</b><br><i>Windows Defender Labs</i>                                   |
| <b>Paul Henry</b><br><i>Wadeware LLC</i>                     | <b>Hamish O'Dea</b><br><i>Windows Defender Labs</i>             | <b>Tomer Teller</b><br><i>Azure Security</i>  |
| <b>Aaron Hulett</b><br><i>Windows Defender Labs</i>          | <b>James Patrick Dee</b><br><i>Windows Defender Labs</i>        | <b>Vikram Thakur</b><br><i>Windows Defender Labs</i>                                |
| <b>Ivo Ivanov</b><br><i>Windows Defender Labs</i>            |   |   |

## Contributors

|  |   |  |
|--|---|--|
| <b>Eric Avena</b><br><i>Windows Defender Labs</i>              | <b>Satomi Hayakawa</b><br><i>CSS Japan Security Response Team</i> | <b>Heike Ritter</b><br><i>Windows and Devices Group</i>        |
| <b>Iaan D'Souza- Wiltshire</b><br><i>Windows Defender Labs</i> | <b>Sue Hotelling</b><br><i>Windows and Devices Group</i>          | <b>Norie Tamura</b><br><i>CSS Japan Security Response Team</i> |
| <b>Dustin Duran</b><br><i>Windows Defender Labs</i>            | <b>Yurika Kakiuchi</b><br><i>CSS Japan Security Response Team</i> | <b>Steve Wacker</b><br><i>Wadeware LLC</i>                     |
| <b>Tanmay Ganacharya</b><br><i>Windows Defender Labs</i>       | <b>Louie Mayor</b><br><i>Windows Defender Labs</i>                | <b>David Weston</b><br><i>Windows Defender Labs</i>            |
| <b>Chris Hallum</b><br><i>Windows and Devices Group</i>        | <b>Dolcita Montemayor</b><br><i>Windows Defender Labs</i>         |  |

# Table of contents

|  |          |
|--|----------|
| About this report .....  | iv       |
| How to use this report .....   | v        |
| <br>   |          |
| Featured intelligence .....  | 1        |
| <b>Protecting cloud infrastructure: Detecting and mitigating threats using Azure Security Center .....</b> | <b>3</b> |
| Threats against cloud deployments and infrastructure .....   | 3        |
| The cyber kill chain: On-premises and in the cloud.....  | 7        |
| Countering threats with Azure Security Center Advanced Threat Detection.....                               | 11       |
| Summary.....   | 20       |

# About this report

The *Microsoft Security Intelligence Report (SIR)* focuses on software vulnerabilities, software vulnerability exploits, malware, and unwanted software. Past reports and related resources are available for download at [www.microsoft.com/sir](http://www.microsoft.com/sir). We hope that readers find the data, insights, and guidance provided in this report useful in helping them protect their organizations, software, and users.

## Reporting period

This volume of the *Microsoft Security Intelligence Report* focuses on the first and second quarters of 2016, with trend data for the last several quarters presented on a quarterly basis. Because vulnerability disclosures can be highly inconsistent from quarter to quarter and often occur disproportionately at certain times of the year, statistics about vulnerability disclosures are presented on a half-yearly basis.

Throughout the report, half-yearly and quarterly time periods are referenced using the *nHyy* or *nQyy* formats, in which *yy* indicates the calendar year and *n* indicates the half or quarter. For example, 1H16 represents the first half of 2016 (January 1 through June 30), and 4Q15 represents the fourth quarter of 2015 (October 1 through December 31). To avoid confusion, please note the reporting period or periods being referenced when considering the statistics in this report.

## Conventions

This report uses the [Microsoft Malware Protection Center \(MMPC\)](#) naming standard for families and variants of malware. For information about this standard, see “Appendix A: Threat naming conventions” on page 135 of the full report. In this report, any threat or group of threats that share a common unique base name is considered a family for the sake of presentation. This consideration includes threats that may not otherwise be considered families according to common industry practices, such as generic and cloud-based detections. For the purposes of this report, a threat is defined as a malicious or unwanted software family or variant that is detected by the Microsoft Malware Protection Engine.

# How to use this report

The *Microsoft Security Intelligence Report* has been released twice a year since 2006. Each volume is based upon data collected from millions of computers all over the world, which not only provides valuable insights on the worldwide threat landscape, both at home and at work, but also provides detailed information about threat profiles faced by computer users in more than a hundred individual countries and regions.

To get the most out of each volume, Microsoft recommends the following:

## **Read**

Each volume of the report consists of several parts. The primary report typically consists of a worldwide threat assessment, one or more feature articles, guidance for mitigating risk, and some supplemental information. A summary of the key findings in the report can be downloaded and reviewed separately from the full report; it highlights a number of facts and subjects that are likely to be of particular interest to readers. The regional threat assessment, available for download and in interactive form at [www.microsoft.com/security/sir/threat](http://www.microsoft.com/security/sir/threat), provides individual summaries of threat statistics and security trends for more than 100 countries and regions worldwide.

Reading the volume in its entirety will provide readers with the most benefit and context, but the report is designed to provide value in small doses as well. Take a few minutes to review the summary information to find the information that will be of most interest to you and your organization. Consult the table of contents and the index to learn more about particular topics of interest.

## **Share**

Microsoft also encourages readers to share each released volume, or its download link, with co-workers, peers, and friends with similar interests. The *Microsoft Security Intelligence Report* is written to be useful and accessible to a wide range of audiences. Each volume contains thousands of hours of research disseminated in easy to understand language, with advanced technical jargon kept to a minimum. Each section and article is written and reviewed to provide the most value for the time it takes to read.

## **Assess your own risk**

Reading about the threats and risks that affect different types of environments presents a good opportunity to assess your own risks. Not every computer and entity faces the same risk from all threats. Assess your own risks and determine which topics and information can help you to best defend against the most significant risks.

The volume and scope of threats facing the typical organization make it important to prioritize. The greatest risk to any computer or organization is posed by currently and recently active threats. Pay attention to the threats that have most commonly affected your region or industry, focusing particularly on the most common successful attacks in the wild that cause the most problems. Give less consideration to very rare or theoretical-only attacks, unless your computers are at particular risk for such threats.

## **Educate**

Microsoft strives to make this report one of the most valuable sources of threat and mitigation information that you can read and share. We encourage you to use the *Microsoft Security Intelligence Report* as a guide to educate your employees, friends, and families about security-related topics.

Anyone, including a business, may link, point to, or re-use articles in the *Microsoft Security Intelligence Report* for informational purposes, provided the material is not used for publication or sale outside of your company and you comply with the following terms: You must not alter the materials in any way. You must provide a reference to the URL at which the materials were originally found. You must include the Microsoft copyright notice followed by "Used with permission from Microsoft Corporation." Please see [Use of Microsoft Copyrighted Content](#) for further information.

## **Ask questions**

Contact your local Microsoft representative with any questions you have about the topics and facts presented in this report. We hope that each volume provides a good educational summary and helps promote dialog between people trying to best secure their computing devices. Thank you for trusting Microsoft to be your partner in the fight against malware, hackers, and other security threats.

# Featured intelligence

Protecting cloud infrastructure: Detecting and mitigating threats using Azure Security Center... 3



# Protecting cloud infrastructure: Detecting and mitigating threats using Azure Security Center

Cloud computing introduces new challenges to security organizations of all sizes. Enterprise IT teams have established policies and procedures designed for enterprise infrastructure and applications, based on their decades of security experience dealing with on-premises threats. Many of these policies and procedures can be used effectively in public and hybrid cloud environments. However, security teams need to keep abreast of changes in the threat landscape brought on by the emergence of cloud computing.

## **Threats against cloud deployments and infrastructure**

New types of threats can be related to characteristics of the public cloud only, or to issues introduced by connectivity between on-premises environments and the public cloud. The following subsections provide descriptions of some new types of threats.

### **Disclosing secrets on public sites**

Public code repositories such as GitHub have become very popular with developers because they enable easy collaboration and source control and remove the responsibility for maintaining the repository infrastructure from developers. But public repositories can be a double-edged sword. Documented cases exist of developers accidentally publishing secret keys on GitHub and other public code repositories, which were discovered by attackers and used to compromise cloud services. Such incidents can sometimes give attackers access to a service's entire account/subscription database, or allow them to misuse its compute resources for malicious purposes.

## Pivot back attacks

A pivot back attack occurs when an attacker compromises a public cloud resource to obtain information that they then use to attack the resource provider's on-premises environment. Public facing endpoints in the cloud are often under constant brute force attack through protocols such as Remote Desktop Protocol (RDP) and Secure Shell (SSH). Although the overwhelming majority of these attacks fail, a very small percentage of them succeed. When they do, an attacker can sometimes find sensitive information in unexpected or obscure places.

Targeted attacks against on-premises and cloud infrastructures often focus on IT administrators.

For example, they could find such secrets in a Bash session history or a text file in the root directory of the virtual machine's desktop. Such information can be used to access resources such as databases, SharePoint sites, and cloud storage. If left unimpeded, an attacker could continue gathering information that could provide greater access to the enterprise infrastructure and data.

### Attacks against cloud administrators

Targeted attacks against on-premises and cloud infrastructures alike often focus on IT administrators. The intent is to take control of an email account that has a high probability of containing credentials that can be used to gain access to the public cloud administrator portal.

After logging into the administrator portal, an attacker can gather information and make changes to gain access to other cloud-based resources, execute ransomware, or even pivot back to the on-premises environment, as explained earlier.

## Man in the Cloud (MitC) attacks

Another new threat is posed by what the security company Imperva has dubbed "Man in the Cloud," or MitC attacks,<sup>1</sup> in which an attacker induces a prospective victim to install a piece of malware using a typical mechanism, such as an email with a link to a malicious website. After the malware is downloaded and installed,

---

<sup>1</sup> "Man in the Cloud (MITC) Attacks," Imperva Hacker Intelligence Initiative Report, [https://www.imperva.com/docs/HII\\_Man\\_In\\_The\\_Cloud\\_Attacks.pdf](https://www.imperva.com/docs/HII_Man_In_The_Cloud_Attacks.pdf).

it finds a cloud storage folder on the user's computer. It then switches out the user's cloud storage synchronization token with the attacker's token.

After the token switch, the attacker will receive copies of each file the user places in cloud storage, which effectively makes the attacker a "man in the middle" for cloud storage. One of the attacker's advantages in this threat scenario is that the malware is removed after the token is switched out, which makes it harder to detect the compromise.

### Side-channel attacks

In a side-channel attack, an attacker attempts to put a virtual machine on the same physical server as the intended victim. If such a successful co-location can be achieved, the attacker will be able to launch local attacks against the victim. These attacks might include local DDoS, network sniffing, and man-in-the-middle attacks, all of which can be used to extract information.

It should be noted that side-channel attacks are not trivial. Microsoft Azure employs a number of obfuscation methodologies to significantly decrease the chances of such an attack succeeding.

### Resource ransom

Ransomware is well-known in the desktop operating system space. This malware restricts access to components of an operating system or to files stored on disk, typically through encryption, and demands that the victim pay the attacker to get the keys required to restore access.

Attackers have made similar attempts to hold cloud resources hostage by breaking in to a prospective victim's public cloud account using any one of a number of methods, including some of the methods discussed in this section. When they have control of the account, the attackers attempt to encrypt or otherwise restrict access to as many cloud resources as possible. The attackers then require the victim to pay the ransom to release the restricted resources.

The challenge for the attacker is to inform the victim that the attack has taken place, and how to pay the ransom. Because servers usually don't have signed in users, attackers need to use methods other than those used for desktop ransomware. One way an attacker can inform cloud resource ransom victims is through the use of bot technology, which presents another, and perhaps unexpected, use case for the new and growing ecosystem of bot technologies.

## Cloud weaponization

In the cloud weaponization threat scenario, an attacker establishes a foothold within a cloud infrastructure by compromising and taking control of a few virtual machines. The attacker can then use these virtual machines to attack, compromise, and control thousands of virtual machines – some within the same public cloud service provider as the initial attack, and others inside other public cloud service providers.

Each of the compromised virtual machines has malware installed that establishes a backdoor connection to the attacker’s command and control servers, from which the attacker can issue commands to the thousands of compromised virtual machines to attack targets throughout the Internet.

Cloud weaponization can be implemented in a number of ways using a variety of attacks, including SSH, RDP, distributed denial-of-service (DDoS), unsolicited messaging (spamming), port scanning, and port sweeping.

Azure actively monitors for cloud weaponization. Figure 1 shows the distribution of the outbound attacks discovered (and in many cases mitigated) by Azure Security Center’s advanced detection mechanisms.

Figure 1. Outbound attacks from Azure virtual machines, September 2016

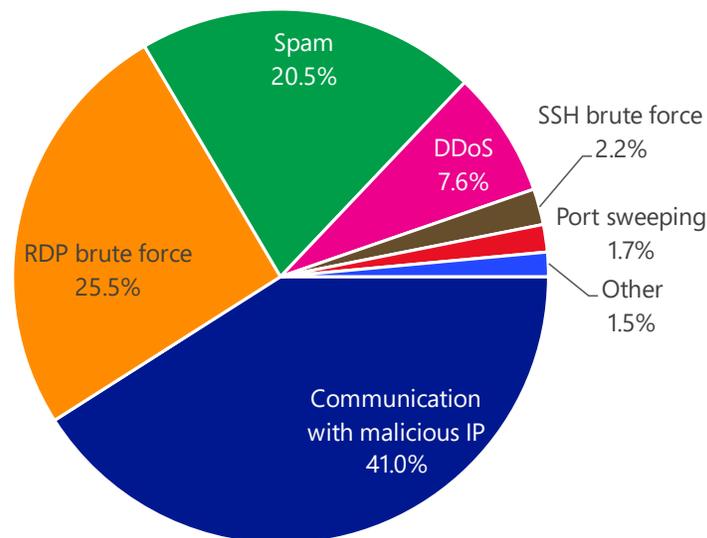


Figure 2 and Figure 3 show where incoming and outgoing attacks originate from.

Figure 2. Incoming attacks detected by Azure Security Center in September 2016, by country/region of origin

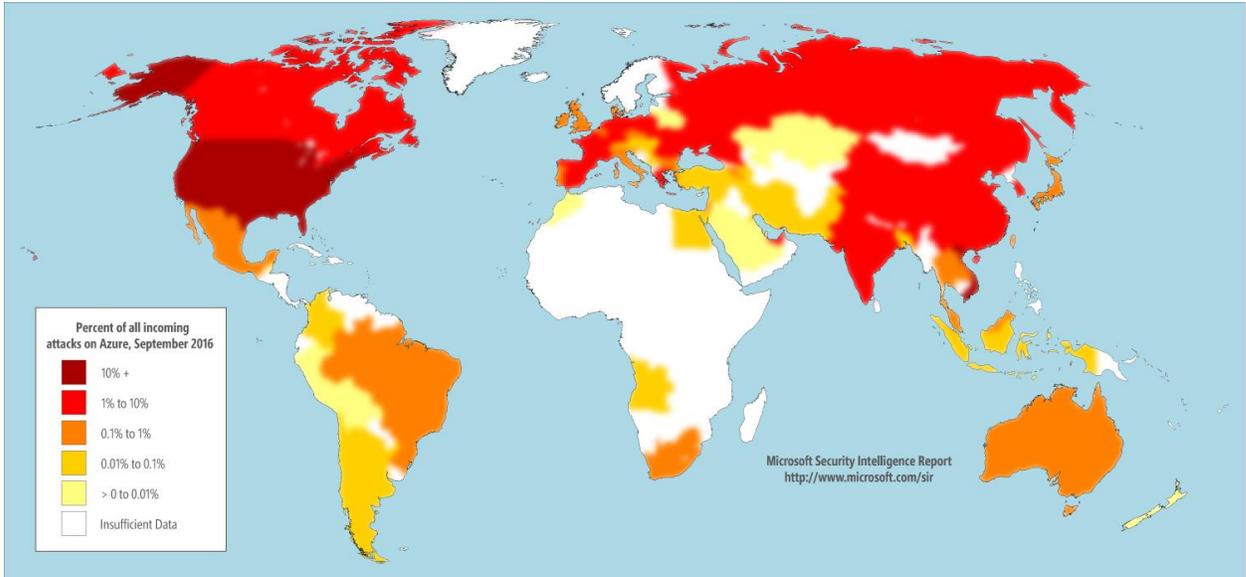
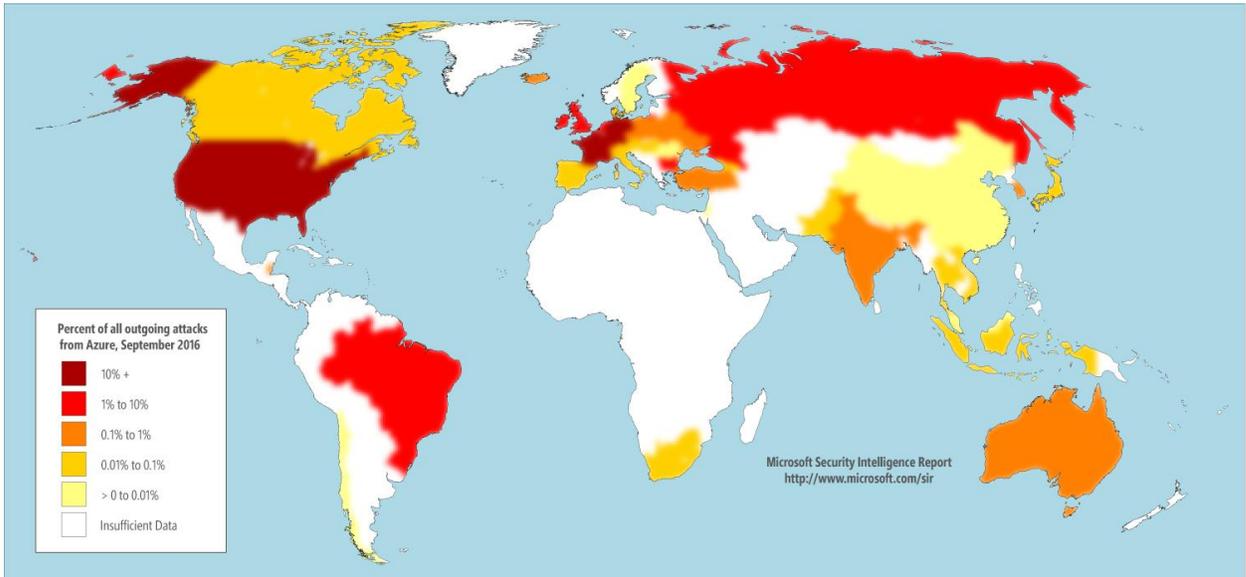


Figure 3. Outgoing communication to malicious IP addresses detected by Azure Security Center in September 2016, by address location



## The cyber kill chain: On-premises and in the cloud

The cyber kill chain is a model defined by analysts at Lockheed Martin to aid decision making with regard to detecting and responding to threats.<sup>2</sup> This

<sup>2</sup> Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph.D., "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," Lockheed Martin Corporation, 2011, [www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf](http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf).

model has become very popular with IT security groups within both small and large IT organizations. It includes the following phases:

- Reconnaissance. The attacker determines the best targets by probing a number of online and offline resources.
- Weaponization. Files (such as documents) can be changed in ways to make them useful “weapons” against a target system and can also be used to enable installation of malicious code.
- Delivery. Weaponized files are placed on the target.
- Exploitation. Weaponized files are “detonated” to take advantage of weaknesses in the target operating system or applications.
- Installation. A back door mechanism is installed on the compromised device so that that the attacker has persistent access.
- Command and control (C2). Malware on the compromised device communicates with a command-and-control system that provides the attacker with access to resources required to carry out actions.
- Actions on objectives. The attacker moves forward to carry out objectives, which may be predefined, or evolve based on discovery.

The cyber kill chain was defined at a time when cloud computing was still gaining traction and did not explicitly consider the some of the unique aspects of cloud computing. There are some differences in how to approach the various phases in the kill chain between on-premises and cloud computing scenarios.

Figure 4 reformulates the cyber kill chain phases to make it easier to understand some of the differences in the cyber kill chain between on-premises and cloud environments.

Figure 4. The cyber kill chain on-premises and in the public cloud

| Phase                   | On-premises                                    | Public cloud                           |
|-------------------------|--|--|
| Active reconnaissance   | HUMINT, OSINT ( <i>users</i> )                 | Foot printing ( <i>services</i> )      |
| Delivery                | Browser, mail, USB ( <i>user interaction</i> ) | Hacking ( <i>no user interaction</i> ) |
| Exploitation            | Client-side vulnerabilities                    | Server-side vulnerabilities            |
| Persistence             | File system based                              | Memory based                           |
| Internal reconnaissance | Custom tools                                   | Built-in admin tools                   |
| Lateral movement        | Machine pivot                                  | Resource pivot                         |

## Active reconnaissance

During the active reconnaissance phase, the attacker learns about the intended victim to improve their chances of a successful attack. In the on-premises world, the attacker can take advantage of social networks to learn information about the target that can be used to induce the victim to download malware during the delivery phase.

The same ruse isn't as easy in the cloud. There's no social network for servers and services to help the attacker learn more about them. The attacker must go through a time and effort-intensive process of scanning the network, doing port scans to discover devices, and then testing active service ports. All this activity provides the defender an opportunity for discovering the attacker's activities.

Active reconnaissance isn't as easy in the cloud.

## Delivery

The attacker places malware on the target during the delivery phase. In the on-premises world, the attacker can create an email that has a malicious link to a website or include an attachment that leads to the installation of malicious code. Another option is to copy the malware onto a USB key and then place the USB key in a strategic location so that the intended target finds it. The victim then puts the USB key into their computer, which compromises it.

In the public cloud, the attacker needs to deliver the malicious payload to a server. Because it's unlikely to find a logged on user on a server to install malicious code, the attacker needs to find a way to gain direct access. One way to accomplish this is through a brute force attack. If such an attack is successful, the attacker will be able to place malware on the server.

The defender has an opportunity to detect the malware on the server before the attacker moves on to the exploitation phase.

## Exploitation

On-premises exploitation typically focuses on client-side vulnerabilities. In the public cloud the focus is on server-side vulnerabilities.

## Persistence

In most cases, attackers of client operating systems will use tools that persist on the compromised device by placing them on the local hard disk. Tools aren't placed in memory because client computers reboot relatively often for system updates, policy changes, or even simple password changes or bug checks.

In contrast, server uptime is much longer, which benefits attackers because they can load exploit code into memory and have the code persist for an extended period of time. Longer server uptime reduces the risk of detection because there is no persistent code on disk that's easy to detect.

Although traditional disk scanning techniques won't find evidence of in-memory malware, defenders can use crash dump analysis to discover and examine malware that exists only in memory.

## Internal reconnaissance

In many on-premises client attack scenarios, the attacker uses custom tools. Built-in toolsets are not as robust as those found on servers and therefore don't meet their needs.

Such custom toolsets aren't seen very often in the cloud. Attackers take advantage of built-in admin tools, which are typically more powerful than what's found on client operating systems. These built-in admin tools help attackers by reducing the risk of detection; they don't need to place custom attack tools on disk.

Because new attack tools aren't being installed on cloud-based virtual machines, they can't be detected with disk scanning techniques. Instead, defenders can use machine learning and behavioral analytics to differentiate between legitimate admin activity and malicious activity.

## Lateral movement

Lateral movement across on-premises networks uses a machine (or virtual machine) pivot. Attackers move from machine to machine by obtaining increasingly privileged credentials as they expand outward. Tools such as mimikatz are used by attackers to harvest such credentials.

The machine pivot isn't currently the norm in the cloud. There are a number of reasons for this, such as the fact that tenants maintain a number of resource

islands in the cloud. Also, in most cases there is limited trust between the cloud and on-premises deployments.

In the cloud, the primary pivot appears to cloud resources. For example, with resource pivoting, an attacker will compromise an IaaS virtual machine, find credentials for a storage service, where more credentials are discovered, some of which allow access to a SQL instance. The attacker hops services instead of virtual machines.

Without powerful detection, there can be no response.

This service hopping behavior enables the defender to focus on this type of activity and enables another avenue for detection.

### **Countering threats with Azure Security Center Advanced Threat Detection**

Azure Security Center helps protect, detect, and respond to security threats against Azure cloud-based resources. Security Center provides protection by analyzing the security status of Azure resources and then providing recommendations on how to increase the level of security.

Protection is just the first level. The ability to detect that an intrusion has taken place is critical. Without powerful detection, there can be no response. Azure Security Center uses advanced threat detection technologies and methodologies to detect threats that would have been very difficult to find prior to the advent of machine learning and big data.

Azure Security Center uses a number of methods that work together to provide advanced threat detection. These methods include:

- Atomic detections
- Threat intelligence feeds
- Behavioral analysis
- Anomaly detection
- Detection fusion

#### **Atomic detections**

Atomic detections are based on well-known malicious patterns that are consistent with indicators of compromise (IoC). These patterns are not subject to mutation, and therefore are considered unambiguous. They can be determined

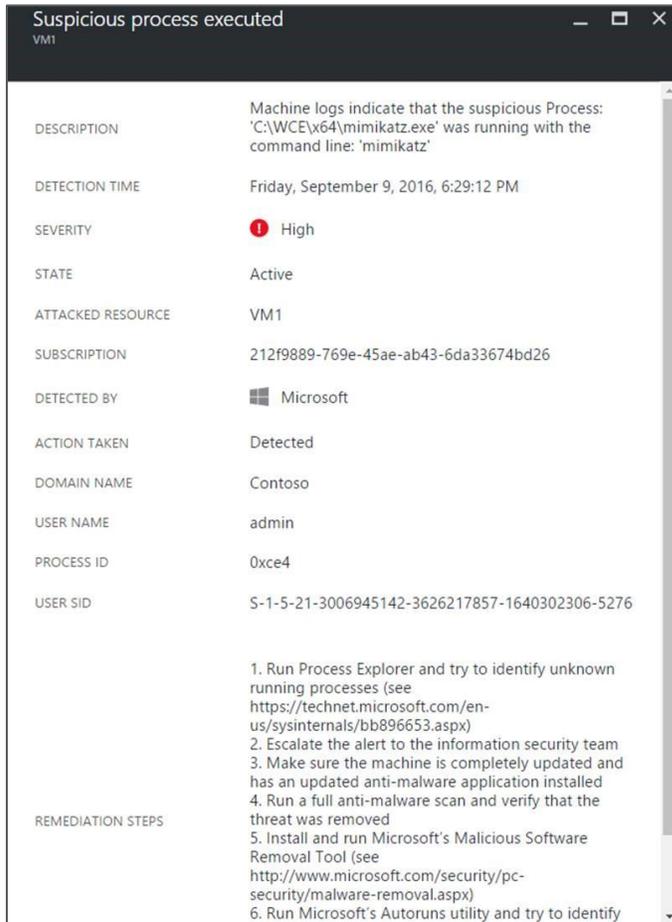
by a single entity, such as a single packet, single behavior, or single event (recorded in a log entry). This determination is similar to how an intrusion detection system (IDS) works, but instead of using on-the-wire packet analysis, atomic threat detection typically uses log entries.

Atomic threat detection has a high return on investment because development overhead is relatively low. In addition, there is also a very low false positive rate. Most commodity malware can be found with atomic detections.

A disadvantage of atomic detection is that it isn't the best method for detecting more sophisticated attacks. Atomic detections are very threat specific, and so it is relatively easy for skilled attackers to evade them. However, this isn't a problem because Azure Security Center uses a multi-tier detection strategy that provides the ability to detect attacks at multiple levels.

Suspicious processes provide an example of an attack type that lends itself to atomic detection. In the example seen in Figure 5, you can see that Azure Security Center has detected that the `mimikatz.exe` process is running. Mimikatz is a tool that has been used by malicious attackers to steal credentials from a compromised machine.

Figure 5. Azure Security Center detects and alerts on the mimikatz malware



## Threat intelligence feeds

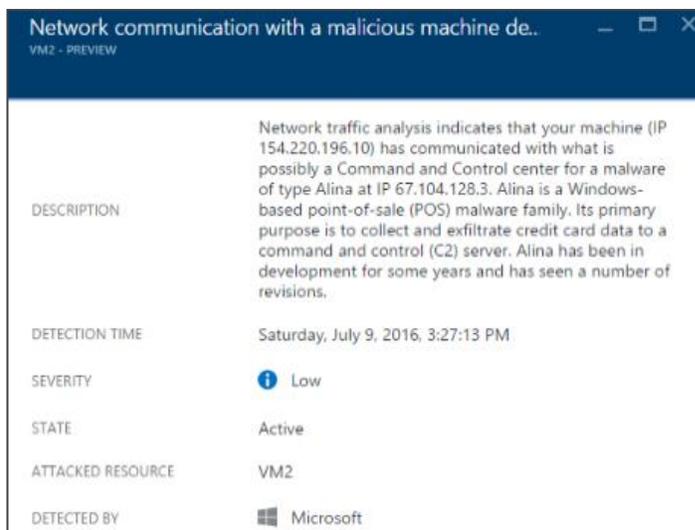
Azure Security Center uses a number of threat intelligence (TI) feeds, such as those from the Microsoft Digital Crime Unit, to help detect potential threats against Azure resources. Azure Security Center uses these feeds primarily for bot detection.

Several actions are possible if a virtual machine hosted on Azure appears in one of these feeds. For example, observing network traffic can confirm that the potentially compromised virtual machine is in contact with a command-and-control server. If this network communication is successfully verified, it's possible to take over the compromised VM's DNS, which can provide additional insight into the botnet infrastructure and IP addresses used by the command-and-control servers.

Similar to atomic detections, TI feeds provide a high ROI for threat detection because of the simple logic required to attain high fidelity alerts. And as more TI providers are added, more detections are possible.

Azure Security Center alerts users when their virtual machines are discovered to be communicating with command-and-control servers. These connections are consistent with bot links from computers infected with malware similar to Alina or Conficker. Figure 6 shows an alert generated by Azure Security Center based on such a detection.

Figure 6. Azure Security Center showing communication with a C&C server



## Behavioral analysis

Atomic and threat intelligence-based detections are essentially pattern matching. To detect more complex threats, more advanced methods of threat detection are required.

One such method is behavioral analysis. In contrast to pattern matching, behavioral analysis moves beyond pattern matching and signatures and focuses on the malicious *behavior*. This focus enables defenders to counter attackers who generate an almost infinite number of variants for a particular malware. Malware designers can change the hash, switch a bit or byte, change a packet or pattern; each of these changes would require a new signature.

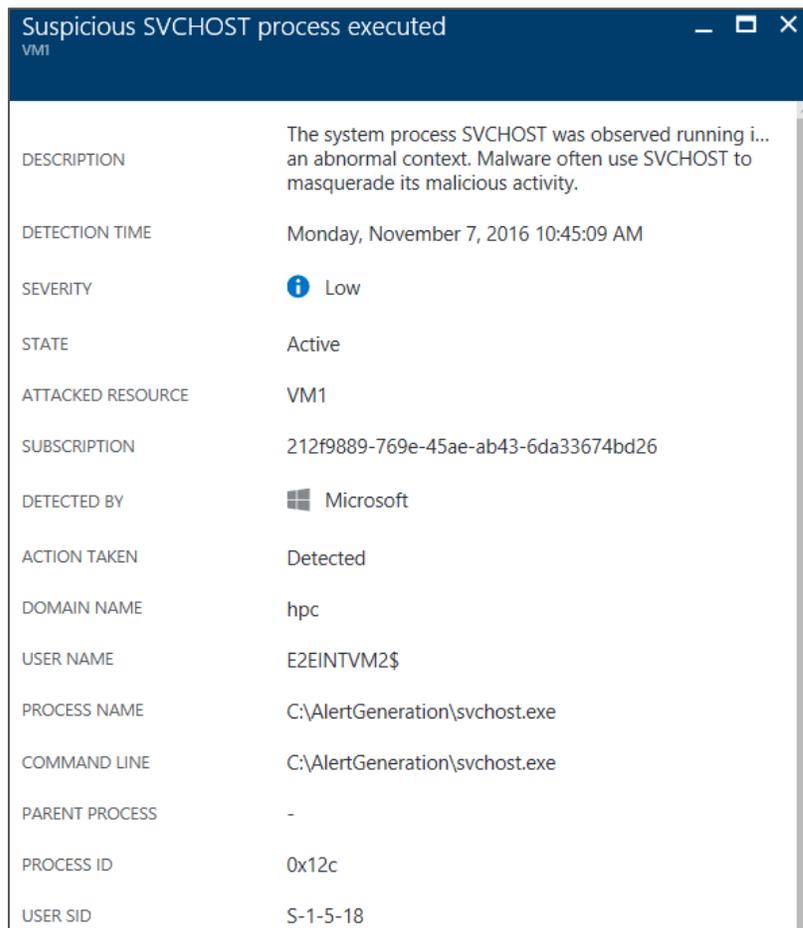
Behavioral analysis drills down to what the malware is *doing* on the system. There's no need to pattern match each malware variant if the *behavior* of the malware can be identified. In the final analysis, it's the behavior that is of most

interest. Each malicious behavior can represent literally thousands of individual signature variants. Behavioral analysis is variant resistant.

However, there are gray areas in some of the behaviors. Such borderline cases could lead to false positives. When the system detects these ambiguous behaviors, Azure Security Center looks for other behaviors or detections to confirm the initial suspicion. More information about the confirmation process is provided in the discussion on detection fusion later in this section.

An example of behavioral analysis: system processes that run in an abnormal context. Azure Security Center can detect these situations and fire an alert. In the example seen in, you can see that Azure Security Center has detected that SVCHOST was running in an abnormal context. SVCHOST is a container process for many other system processes and malware often tries to take advantage of this to hide its activity.

Figure 7. Azure Security Center detects processes running in an abnormal context



The screenshot shows a window titled "Suspicious SVCHOST process executed" with a sub-header "VM1". The window contains a table of alert details:

|                   |   |
|-------------------|---|
| DESCRIPTION       | The system process SVCHOST was observed running i...<br>an abnormal context. Malware often use SVCHOST to<br>masquerade its malicious activity. |
| DETECTION TIME    | Monday, November 7, 2016 10:45:09 AM  |
| SEVERITY          |  Low   |
| STATE             | Active  |
| ATTACKED RESOURCE | VM1   |
| SUBSCRIPTION      | 212f9889-769e-45ae-ab43-6da33674bd26  |
| DETECTED BY       |  Microsoft   |
| ACTION TAKEN      | Detected  |
| DOMAIN NAME       | hpc   |
| USER NAME         | E2EINTVM2\$   |
| PROCESS NAME      | C:\AlertGeneration\svchost.exe  |
| COMMAND LINE      | C:\AlertGeneration\svchost.exe  |
| PARENT PROCESS    | -   |
| PROCESS ID        | 0x12c   |
| USER SID          | S-1-5-18  |

## Anomaly detection

Behavioral analysis discovers known threats by detecting known behaviors. This process is immensely useful, and significantly extends detection capabilities beyond simple signature-based detections.

The logical next step is to detect unknown threats, which is where anomaly detection comes into play.

With anomaly detection, the system builds a baseline. The baseline is defined by the history of a certain element of the virtual machine. If a statistically significant deviation from that baseline is detected, an alert might be generated.

It's important to note that while the system *might* generate an alert, it's possible that no alert will be generated. The reason for this possibility is that not all deviations from baseline are detrimental. Similar to other detections, when the system detects ambiguous activity, supporting evidence and correlation with other detections are sought to confirm.

Anomaly detection makes it possible to move past what is already known, and discover possible new exploits.

One of the major advantages of anomaly detection is that it enables detections to move past what is already known, and discover possible new exploits. Anomalies can lead researchers to dig deeper, and come up with new analytics that define new "known" threats.

An example of an anomaly-based detection is a brute force attack. A brute force attack is characterized by repeated attempts to log onto a virtual machine with guessed user names and passwords. Azure Security Center can detect that the number of failed log on attempts has reached a statistically significant level and generate an alert, as seen in Figure 8.

Figure 8. Azure Security Center detects a failed brute force attack

| Failed RDP Brute Force Attack |  |
|-------------------------------|--|
| VM1                           |  |
| DESCRIPTION                   | Several Remote Desktop login attempts were detected from FreeRDP (96.81.218.10), none of them succeeded. Event logs analysis shows that in the last 48 minutes there were 93 failed attempts. 32 of the failed login attempts aimed at non-existent users. 1 of the failed login attempts aimed at existing users. |
| DETECTION TIME                | Sunday, November 6, 2016 9:45:07 AM  |
| SEVERITY                      | Low  |
| STATE                         | Active   |
| ATTACKED RESOURCE             | VM1  |
| SUBSCRIPTION                  | 212f9889-769e-45ae-ab43-6da33674bd26   |
| DETECTED BY                   | Microsoft  |
| ACTION TAKEN                  | Detected   |
| SUCCESSFUL LOGINS             | 0  |

### Detection fusion

As stated several times in this section, instances exist when a specific detection is non-specific, which requires supporting evidence to reduce the probability of a false positive. A very effective method for reducing ambiguity (and the false positive rate) is to correlate individual alerts generated throughout the cyber kill chain.

The correlations provide the context needed to confirm that the findings of each of the individual alerts represents an actual security event. We call this combination or correlation of multiple alerts along the kill chain a security *incident*.

Security incidents are explicitly identified in the alerts section of Azure Security Center. In addition, incidents help to reduce investigation time by providing insight into what steps the attacker took, and what specific resources were affected.

Incidents tie together alerts during attack progression. A simplified characterization of the cyber kill chain places incidents into one of three phases:

- Target and attack
- Install and exploit
- Post breach

### **Target and attack**

The target and attack phase represents the reconnaissance and deliver phases of the cyber kill chain.

For example, a brute force attack against a virtual machine fits into this phase; alerts related to brute force attacks are placed here.

When an attacker launches a brute force attack against a virtual machine, Azure Security Center will use anomaly detection to determine whether the number of logon attempts exceeds what is expected. If so, Azure Security Center surfaces a failed brute force attempt alert to the user.

### **Install and exploit**

The system analyzes the results of the initial attack during the install and exploit phase.

Some of the things considered during this phase include:

- Evidence of existing malware signatures (using Microsoft antimalware or partner solutions)
- In-memory malware (using crash dump analysis)
- Suspicious process execution (using behavioral analysis)
- Lateral movement
- Internal reconnaissance

Suppose an attacker were able to gain access to a virtual machine using a brute force attack (which would have taken place during the target and attack phase). Malware is installed on the machine (during the install and exploit phase) and the malware ends up causing a process to crash.

Azure Security Center will collect a copy of the crash dump and scan it for evidence of in-memory malware. If an exploit (such as malicious shellcode) is

found, an alert will be generated and assigned to the target and attack phase of the kill chain.

## Post breach

Attackers execute their plans during the post breach phase, which includes all the activities attackers carry out using automated or manual processes on compromised virtual machines.

For example, a virtual machine is compromised by a brute force attack during the target and attack phase and an alert is generated. The attacker installs malware, and another alert is generated during the install and exploit phase. Finally, the malware generates large amounts of SMTP traffic. This SMTP traffic is correlated with the Office 365 SPAM database to determine whether this traffic is legitimate or SPAM. If the assessment is SPAM, an alert is generated by Azure Security Center.

In addition to these alerts, an *incident* (defined as a collection of alerts) is generated by Azure Security Center indicating a very high probability that a successful compromise has taken place because of the correlation and verification of and by multiple alerts.

Figure 9 shows a number of security incidents as reported by Azure Security Center.

Figure 9. A list of security incidents in Azure Security Center

|   | DESCRIPTION ^              | COUNT ^ | DETECTED BY ^ | DATE ^   | STATE ^ | SEVERITY ^ |
|---|----------------------------|---------|---------------|----------|---------|------------|
| 🔍 | Security incident detected | 1       | Microsoft     | 09/19/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 08/30/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 08/30/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 08/17/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 08/17/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 08/17/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 07/20/16 | Active  | 🔴 High     |
| 🔍 | Security incident detected | 1       | Microsoft     | 07/20/16 | Active  | 🔴 High     |

Azure Security Center also provides the ability to drill down on a security incident and provide detailed information on the individual alerts that were correlated to create the incident, as shown in Figure 9.

Figure 10. Azure Security Center provides additional information about security events

**Security incident detected**  
Incident Detected - Preview

**DESCRIPTION** The incident which started on 2016-09-18T13:58:47.6860842Z and most recently detected on 2016-09-19T22:58:47.6860842Z indicate that an attacker has attacked other resources from your virtual machine VM1

**DETECTION TIME** Monday, September 19, 2016 5:58:47 PM

**SEVERITY** ! High

**STATE** Active

**ATTACKED RESOURCE** VM1

**SUBSCRIPTION** 212f9889-769e-45ae-ab43-6da33674bd26

**DETECTED BY** Microsoft

**ACTION TAKEN** Detected

**REMIEDIATION STEPS**

1. Escalate the alert to the information security team.
2. Review the remediation steps of each one of the alerts

Alerts included in this incident

| DESCRIPTION   | COUNT | DETECTION TIME    | ATTACKED RESOURCE | SEVERITY  |
|---|-------|-------------------|-------------------|---|
| SQL injection blocked                               | 1     | 09/18/16 08:58 AM | VM1               | Low   |
| Failed RDP Brute Force Attack                       | 1     | 09/18/16 09:58 AM | VM1               | Low   |
| Successful RDP brute force attack                   | 1     | 09/19/16 09:58 AM | VM1               | <span style="color: red; font-weight: bold;">!</span> High      |
| Suspicious SVCHOST process executed                 | 1     | 09/19/16 10:58 AM | VM1               | Low   |
| Multiple Domain Accounts Queried                    | 1     | 09/19/16 11:58 AM | VM1               | Low   |
| Network communication with a malicious machine d... | 1     | 09/19/16 12:58 PM | VM1               | <span style="color: orange; font-weight: bold;">!</span> Medium |

## Summary

The cloud introduces a number of new attack vectors that were previously unavailable to intruders in the on-premises world. These new attack types require that we evolve our methods of attack detection. Detecting cloud-based attacks requires us to address and act on the differences between the on-premises and cloud kill chains. There are also security benefits from running your workloads in the cloud, as you'll benefit from Microsoft's comprehensive threat intelligence and security expertise. Azure Security Center takes advantage of a multi-layered approach to threat detection, which ranges from rudimentary signature based systems all the way up to machine learning driven approaches and detection fusion. See [azure.microsoft.com](http://azure.microsoft.com) for more details and to take advantage of a 90 day free trial of Azure Security Center.





One Microsoft Way  
Redmond, WA 98052-6399  
[microsoft.com/security](https://microsoft.com/security)