Defending the New Perimeter

MODERN SECURITY FROM MICROSOFT

A Guide to the Microsoft Cybersecurity Stack for IT Decision Makers

Pete Zerger and Wes Kroesbergen

Contents

Copyright Notice

Version: 1.1

Publication Date: October 10th, 2016

Acknowledgements

We would like to take a moment to thank the following people who helped to make this project happen:

Our diligent content reviewers:

- Stuart Dankevy, Microsoft
- Aaron Hamilton, Microsoft
- Nick Romyn, Softchoice

Bettina Berntsen, Lumagate, who handled logistics and around distribution, allowing us to focus on writing.

Our patient families, for tolerating many hours of work on weekends and late evening phone calls.

About the Authors

Pete Zerger, CEO of Lumagate in North America, is an author, speaker, architect and 11-time Microsoft MVP, focusing on cloud architecture, automation and security, as well as DevOps. Pete has presented at popular technical conferences around the world including Microsoft Ignite, TechEd and System Center Universe. He has authored many technical whitepapers and is co-author of several books on management and cloud topics, including the popular “Inside the Microsoft Operations Management Suite”.

Wes Kroesbergen is a Global Blackbelt Specialist at Microsoft, focused on Enterprise Mobility and Security. Wes worked for a variety of Microsoft Partners prior to Microsoft, having previously architected and implemented SharePoint, Hyper-V, System Center, Azure, and EM+S solutions for enterprise customers across North America. He now focuses on enabling Microsoft customers to understand the incredible security capabilities offered within the Microsoft stack.

1 – Introduction

Virtually all organizations are turning to the cloud as an enabler of their digital transformation. This transformation will drive new business opportunities, and a modern work style. This transformation, enabled through mobile devices, cloud apps and services, big data and machine learning, is already delivering great leaps in productivity and new business opportunities. Organizations embracing their digital transformation are gaining competitive advantages and leapfrogging their competition. Users, and the business as a whole, demand productivity from anywhere on any device. This requires access to sensitive corporate data outside of the traditional workplace, modern line of business apps, and even access to legacy applications within corporate facilities, created long before the cloud was born.

One thing is clear: businesses are moving at rapid speed to reap the benefits of their digital transformation and are not tolerating delay of 'traditional' IT's pace of innovation. To retain their relevance, IT must enable and drive this digital transformation and prepare to support the business at this new pace. Though security remains a top concern for CIOs, CISOs and IT professionals, it is fast becoming a universal area of focus for boards of directors. Cloud and enterprise mobility impact security considerations and strategy across many areas, including identity management, data security, as well as application and device management. Just as the business must undergo this digital transformation to enable this new model of work, so too must the IT organization transform the way it views security around devices, data, applications and identity in the context of a cloud-first world.

This book is a guide to that transformation.

Democratization of Technology

Today we live in an increasingly connected world, with technology becoming accessible to more and more people. With 64% of American adults who now own a smartphone in 2015, up from 35% four years ago, a vast number of individuals now rely on Internet connectivity to services that enable their personal and professional lives. These services are often delivered on hyper-scale infrastructure providers such as Amazon or Microsoft. Indeed, the democratization of datacenter technology via these hyper-scale cloud providers has given rise to countless thousands of independent software vendors leveraging these new economics to provide unique, differentiated, and specialized applications, enabling reach to billions of users around the globe. Ultimately, these cloud services have also served to commoditize data storage and communications, enabling pay-as-you-go models for commercial activities and services, which often serve to subsidize free services for individuals. We live in a wonderful age where there are seemingly endless numbers of software applications tailored to suit almost every niche, whether personal or corporate. What this has led to is the consumerization of IT and mass scale adoption of applications that historically could only be deployed through a 'traditional' IT department.

The Importance of Trust

With the increasing reliance on cloud services and mobile devices to perform day-to-day personal or professional tasks, comes an explosion in data growth. Mobile devices contain high value data for malevolent attackers: personal health information, credit card information, photos, corporate communications, files, and resource access. Cloud providers host an increasing amount of corporate data and communications data gathered from consumer-facing services delivered on mobile devices. The efficiencies of cloud provider storage enable virtually unlimited storage of data by the application vendor, and this data in turn, is often mined by the application vendor via machine learning for analytics of past performance, and predictions or recommendations for future optimizations and revenue opportunities.

As a result, protection of data is of increasing concern. Transparency of cloud providers and application vendors becomes vital, as organizations have compliance and security standards they must meet. Since these compliance and security standards are now a shared responsibility between the organization and the cloud provider, trust has become more important than ever. Individuals and businesses alike, have an expectation and a right to privacy of their data in compliance with the laws of their land. Cloud providers and application vendors have a significant responsibility to fulfill in providing privacy, transparency and compliance for both the business and consumer customer. This includes providing controls for businesses to protect their users and data.

Operational Challenges

As access to a plethora of differentiated cloud services have become so readily accessible to consumers via their mobile devices, more and more consumers are carrying a desire and sense of entitlement to utilize these services in their professional life. If the organization does not provide corporate access to a particular SaaS service, individuals will often sign-up for a personal account with their corporate email address and store work-related information and data in the service. In the past, an organization could simply perform network inspection and lock down their PCs, and in general, were able to manage security of corporate data. Today however, users are typically accessing these services from non-corporate locations and devices.

Based on the fact that every SaaS provider has different security capabilities and compliance standards, it becomes a seemingly impossible challenge to manage. As a result, within the traditional perimeter model (think 'castle and moat' scenario where data and applications are housed within the datacenter protected by traditional technology and controls), protection of corporate data becomes exceedingly difficult, and for regulated organizations, compliance is an even greater challenge.

Far too many organizations continue to acquire tactical, 'point product' security solutions, with many organizations operating between 30 and 70 of these types of security solutions. The problem with a strategy of leveraging a variety of security solutions to address different tactical needs is that gaps start to emerge between solutions, causing the organization to end up with a patchwork security ecosystem. This often manifests itself with many operational and security cracks. For those that have invested in 'tight integration' across these solutions to try and address the gaps, they are saddled with an overly complex security infrastructure that is operationally intensive and challenging to manage from both a

technical, as well as a currency perspective. To truly secure a modern enterprise across cloud, mobile, and infrastructure, a security model must be developed, and a strategy formed accordingly.

Evolution of the Threat Landscape

The threat landscape, specifically the types and sources of threats, has changed significantly in the last five years. There are a number of readily identifiable causal factors, perhaps explained most concisely by the European Union Agency for Network and Information Security (ENISA), which cites culprits that include:

- **Changes in technology**. The introduction of new technology results in weaknesses that are related to low technological maturity, improper use, improper integration with existing systems, low user awareness, etc.
- **Changes in the capabilities of threat agents**. Available skills, available tools, available resources, information on exploits and motivation make up the profile of a threat agent. Threat agents include malicious actors ranging from hackers out for individual gain, to well-funded state sponsored intrusion specialists. Let us not forget about malicious employees with access to your trusted networks and an array of cloud apps and storage at their disposal, easing the process of siphoning sensitive data from trusted sources, unnoticed by traditional firewalls and network proxies.
- **Data growth**. 90% of all the worlds data has been produced in the last two years, and with the growth of myriad devices (the Internet of Things), this trend will only become more pronounced. Data is the asset primarily targeted by these threat agents, and the target just keeps getting bigger.

Together, these factors open the door to new attack methods and new threats, emerging from new types of assets, exposing new weaknesses and vulnerabilities. These threats emerge from the combination of new technologies, more skilled threat agents, and poor practices of the organizations and people that use them.

Digital Crime

The increasing connectedness of individuals and businesses, and the explosion of data associated with their profiles and transactions has not been without cost. Cyber-crime has transformed, transitioning from intermittent mischief by so-called 'script kiddies', to an increasing amount of sophisticated fraud and theft committed by organized crime rings across the world. Every vertical is now under attack. It is a matter of when, not if a business will be targeted by an advanced, persistent threat (APT).

Data is a wonderful tool for legitimate use, but in the hands of criminals it swiftly can be turned to cause real material damage to businesses and consumers. **In fact, a McKinsey study of 350 companies across 11 countries found that the average cost of a data breach costs around $3.5M, and this has in turn impacted liability costs for businesses to the tune of $500M**.

Beyond the fraud and theft just mentioned however, we see an increase in activity by nation-states and terror or activist groups. These attackers are very-well funded and very sophisticated, provided with all

the resources needed to cause damage and disruption. In fact, in their 2018 Security Outlook, the Canadian Security Intelligence Service (CSIS) now lists cyber-threat as a significant risk for the 2016-2018 period. CSIS points out that the damage done by the 'never let a good crisis go to waste' Snowden event was an initiator for the ongoing modernization of cyber-weaponry. As a result, the United States and other countries have identified cyber-insecurity as the paramount national security threat.

Anatomy of an Attack

Let us take a look at the anatomy of a typical breach of an organization.

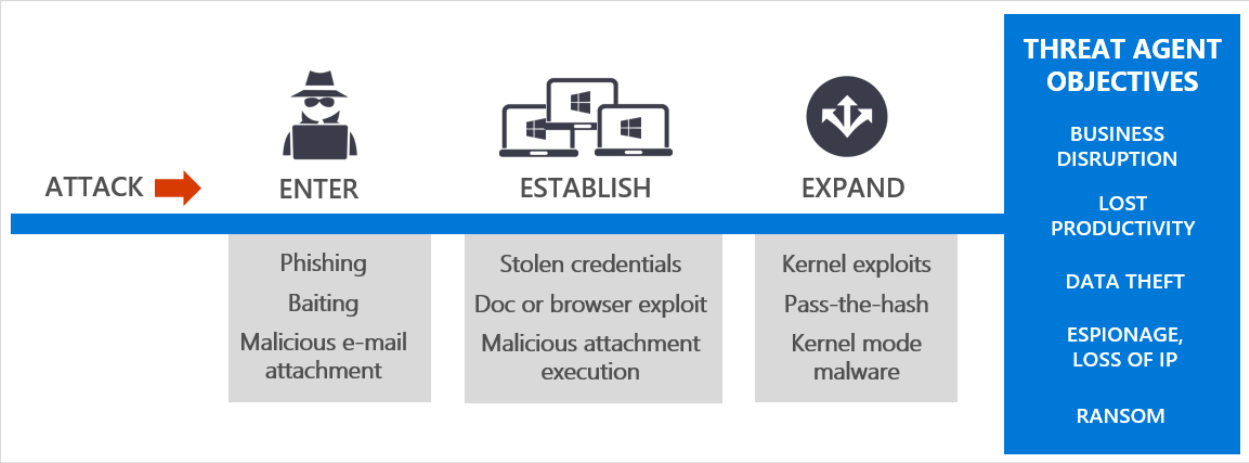

***Figure 1**. Anatomy of an Attack*

These attackers often penetrate their targets via the supply chain and/or compromised credentials. The attackers are often unbound by time-constraints, and can spend a significant amount of time probing a target for possible avenues of compromise. In many cases, the initial beachhead is established in preparation of the attack by tricking a user into clicking a link in an email, which initiates installation of a piece of malware for remote control by the attacker.

The bottom line is attacks happen quickly and are difficult to stop. According to the Verizon 2015 Data Breach Investigations Report, if an attacker sends an email to 100 people in your company:

- 23 people will open it…
- 11 people will open the attachment…
- and 6 will do so within the first hour

Once they choose to execute the attack, it typically only takes 24-48 hours before the attacker obtains complete ownership of the domain through lateral movement and further identity compromises. In what is perhaps the most disturbing piece of reality, it typically takes 4-6 months before the business detects that there is a compromise, and begins the difficult task of flushing the attackers and attempting to regain control of their environment.

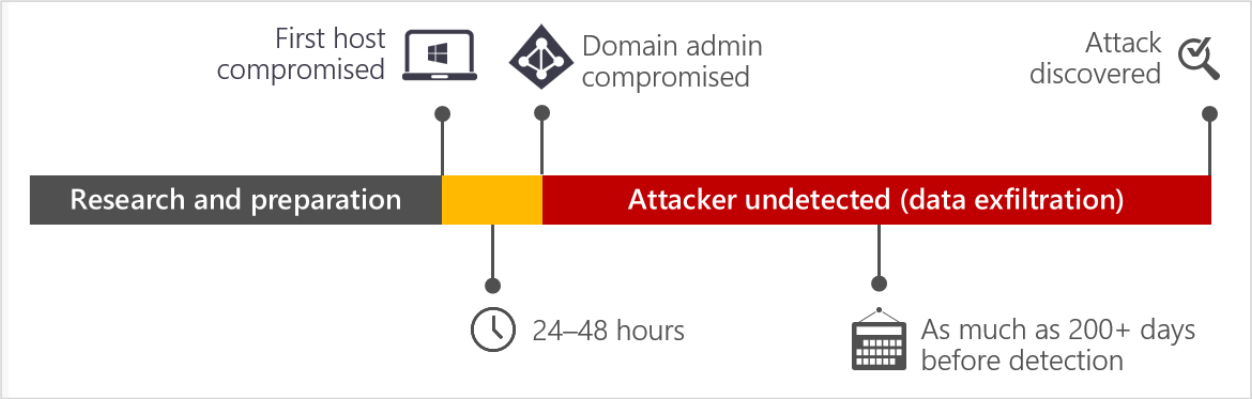

***Figure 2**. Attack Timeline*

While defenders typically think in lists (which assets need to be protected, what firewall solutions are in place, security logs being collected, etc.), attackers think in less linear (or graph) approaches. Where can they establish a beachhead, and what plethora of avenues might there be to the high value assets they wish to retrieve? Detection and response therefore, has become an area of increasing priority. Protection and defense in depth are still critical components to a security strategy, but at the end of the day, assumption of compromise (that security breaches will happen) is an important perspective to take when balancing protection, detection, and response.

2 – Identifying Threats at Scale

Microsoft has a unique position in cybersecurity in the industry, due in large part to the massive scale of information that Microsoft processes on a daily basis. Microsoft receives trillions upon trillions of pieces of data (signals) from billions of devices, its own solutions, and industry and research data from partners, enabling Microsoft to synthesize threat data far faster than your organization could on its own. All of this goes into the **Microsoft Intelligent Security Graph**.

With the incredible amounts of signal collected and experience gathered every day, human monitoring becomes untenable. It is impossible for humans alone to keep up with the data gathered, and so Microsoft has turned to machines to assist with analysis. Thanks to the power and economy of the cloud, Microsoft is able to leverage machines to detect the proverbial needle in a haystack, with algorithms searching for anomalous behavior and correlating various pieces of signal that would otherwise go undetected. As the algorithms detect threats, Microsoft's response mechanisms kick in, and are able to prevent exploitation not only for the customer immediately targeted, but also mitigate the same threat or pattern of attack for all of their other customers. Microsoft's unparalleled amounts of threat signal and reach across their customer base around the globe enable them to detect and respond to threats like no other enterprise vendor.

Machine intelligence only part of the equation. There is a human element as well. Microsoft invests $1B/year into security, employing data scientists and research analysts to deliver world-class protection. Microsoft also leverages their Cyber Defense Operations Center, which has direct access to thousands of security professionals, data analysts, engineers, developers, program managers, and operations specialists throughout Microsoft to ensure rapid response and resolution to security threats. Informed by decades of experience working with the industry to fight threats on a global scale, the center maintains critical connections with industry security partners, governments and enterprise customers, and engages Microsoft's Digital Crimes Unit (DCU) when law enforcement needs arise.

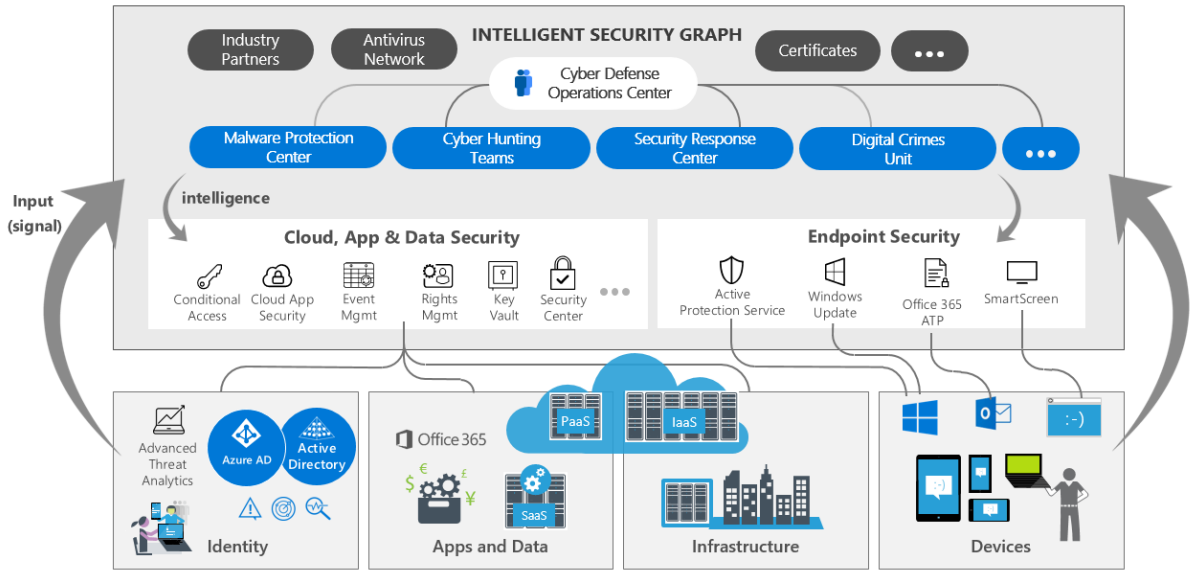

***Figure 3**. Microsoft Intelligent Security Graph*

The DCU is an international team of attorneys, investigators, analysts, data scientists, engineers, and business professionals working together to combat digital crime. The DCU works with law enforcement and governments across the globe not only to fight those who would breach systems to disrupt or steal personal information, but also to protect the most vulnerable among us while online. The DCU has partnered with the National Center for Missing & Exploited Children to fight against the online exploitation of children, and works with other organizations to dismantle criminal networks that target senior citizens with online tech scams and “click fraud”.

The DCU also works with law enforcement agencies to take control of malicious botnets. While Microsoft does not have authority to remove malware from infected devices, they are able to take over the Command & Control hosts, and are able to then map the IP addresses of those infected devices into the intelligent security graph. This enhances Microsoft’s ability to know if a user is signing on from an infected endpoint, and take appropriate action.

All of these elements come together to form the Microsoft Intelligent Security Graph, which today is a truly unique differentiator in the Microsoft Defense Stack when compared to all other security vendors.

3 – Leveraging Identity as a Gateway

The Cloud Security Alliance (CSA) and National Institute of Standards and Technology (NIST) espouse an approach of multi-factor and policy-based authentication that evaluates the full context of the authentication attempt (user, device, location, date/time, app and data). This strategy, which dramatically reduces the likelihood of breach based on compromised credentials alone, is a critical step in a secure approach to identity management, authentication and authorization.

The traditional model of the network perimeter, including firewalls and proxies and a perimeter network (aka DMZ) is dead. The perimeter, where access and authorization are enforced, can be the login screen on mobile device, or an app installed on that device. The app is the window to your corporate data (content), and the new perimeter is the content and context by which the user tries to access that data.

The CSA advises that "identity management in the context of cloud computing should not only manage the user identities. It should extend this to manage cloud application/services identities, access control policies for these cloud applications/services, as well as privileged identities for the applications/services, etc.".

With this in mind, organizations must rethink their approach to identity management, authentication and authorization in a world that did not exist when the concept of username and password entered on a PC behind a trusted network were conceived.

Facilitating Access to Applications

Traditionally, providing remote access to applications meant implementing a virtual private network (VPN) solution or perhaps publishing the application through a firewall. These are two solutions that share a few things in common, including:

- Many do not offer much insight into the circumstances around the login attempt, such as location, device health and variance in login hours
- They do not offer capabilities to provide additional challenges when one or more of the above risk factors are present
- They tend to be complex to configure and maintain
- They require expensive IT staff with specialized skillsets

This is where the Azure Active Directory (AD) App Proxy comes into play. The Azure AD App Proxy Connector establishes an HTTPS outbound connection to Azure, enabling a channel through which the App Proxy service is able to deliver connectivity to internal resources. Instead of opening up ports on the firewall, domain name entries for internal services can be pointed at the Azure datacenters, minimizing an organization's attack surface area. The Azure services become the attack service, leveraging Microsoft's significant security investments to mitigate attacks many organizations are not equipped to handle (e.g. largescale Denial of Service attacks). Just as importantly, you can leverage Azure AD's rich conditional access controls and multi-factor capabilities to govern and secure access, pre-authenticating

users via Kerberos to internal applications, and even legacy applications developed before the cloud was born!

Azure AD App Proxy is a much lighter-weight, and far more powerful mechanism to enable internal resource access to mobile users than traditional VPN. To achieve highly available deployment, the connector merely needs to be installed on additional servers, and the connectors can be collocated with other services to minimize server sprawl. The App Proxy service takes care of distributing incoming load across the individually established connections from the on-premises connectors.

The Azure AD App Proxy also provides a much friendlier user experience, enabling users to more securely and simply access corporate applications without the heavy overhead and confusion often associated with traditional VPN solutions.

The following diagram illustrates how simple the deployment of the Azure AD App Proxy is.

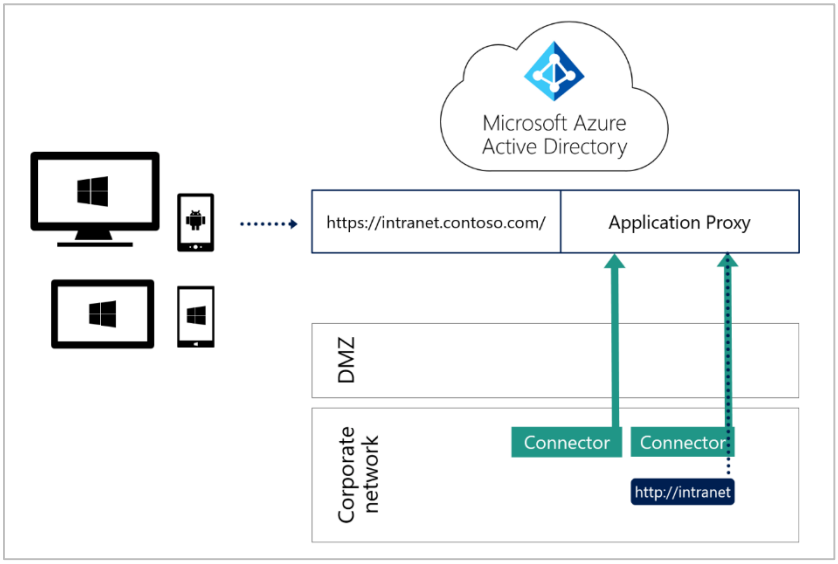

***Figure 4**. Azure AD Proxy Architecture*

Controlling Access from Anywhere

> *Since cloud applications are accessible widely through various devices, authenticating with simple userID / password should be deprecated as a solution. For cloud-based applications, authorization should not only be performed based on the content, but also by the context.*
>
> ~Cloud Security Alliance

Identity is at the same time both a *gateway* and a *perimeter*.

Identity is a gateway in that opens access to resources, provided the right conditions are met. When we consider identity with additional factors, one can determine if resource access should be allowed. While traditional role-based access control (RBAC) in Active Directory considers identity alone, there are additional factors that can help more accurately ensure identity, as well as risk associated with the collective conditions, including:

- **Device**. The state of the device from which they are authenticating?
  Is the device managed by IT, or is it a personal device that has been rooted (jailbroken)?
- **Location**. From where they are attempting to access resources, as well as where those resources are located. Is the location outside the normal set of locations from which this user has attempted to authenticate in the past?
  For example, what if a user based in North America is suddenly attempting to login from China?
- **Time**. Is this day and time within normal patterns represented in past authentication attempts from this user?
  While users may work from home during the evening hours, access attempts in the middle of the night may represent abnormal behavior that imposes an unacceptable risk.
- **Resources**. Which resources they are attempting to gain access to? Are these resources that should be accessed under these conditions?
  Some resources may be so sensitive you may not want users to access them except from your trusted network. For example, perhaps access to your HR application should not be allowed externally, in which case you may want to deny login based on the resource.

Incorporating these factors into resource access enables organizations to mitigate risks associated to sign-in, limiting the ability of an attacker to exploit a compromised identity or device without restoring the identity of device to a safe state. So in this sense, identity, together with these additional factors, becomes a perimeter, behind which resources are protected, allowing through only those requests that fall within on our tolerance for risk.

How does Microsoft solve for this scenario?

The Microsoft security stack for hybrid identity makes this possible through implementation of *policies*, enabling assurance of identity, device, and other conditions in our scenario, ensuring only authentication attempts within our defined risk tolerance threshold are allowed.

> *Enterprises should plan for using risk-based authentication for their cloud applications. This type of authentication is based on device identifier, geolocation, ISP, heuristic information, etc.*
>
> ~ Cloud Security Alliance

- **Ensuring identity**. Through Azure Active Directory, a user can be presented with additional challenges, such as providing a code delivered through an SMS message to their mobile device, to ensure their identify has not been compromised.

NOTE: In Summer 2016, NIST deprecated SMS message as a second factor of authentication in multifactor authentication (MFA) scenarios. A solution offering additional options for out-of-band verification (the additional challenge) is critical to meeting current best practices.

- **Ensuring device health**. Even a known user can pose serious risk from a compromised device. With Microsoft Intune, you can enforce device management policies to ensure devices meet your management standards, including policies requiring assurance the device has not been rooted or jailbroken, before allowing access to corporate resources.
- **Ensuring acceptable conditions for access**. You can use conditional access in Azure Active Directory Identity Protection to establish additional gates to resources based on your tolerance for risk and the sensitivity of the information the user is attempting to access.

To establish identity, with conditional access in Azure AD Identity Protection we can force the user to answer multiple challenges through multi-factor authentication (MFA). In these scenarios, you can leverage conditional access to restrict a user's ability to access resources based on several risk factors, including:

- **Leaked credentials**. If a user's credentials have appeared in a credential dump in one of the dark corners of the web.
- **Irregular sign-in activity**. Such as a user is signing in from new devices, outside of normal hours or days.
- **Sign-ins from possibly infected devices**. Such as if a device is unmanaged, or not up to date on security updates.
- **Sign-ins from unfamiliar locations**. If a user is attempting to login from a new location, this represents a potential risk.
- **Sign-ins from IP addresses with suspicious activity**. If the Microsoft DCU sinkholes have identified that an IP address is a source of suspicious activity, an organization can leverage this intelligence as another reason to restrict access.
- **Sign-ins where impossible travel is involved**. For example, if a user logs in from Tokyo at 3pm and then logs in from Miami 30 minutes later, we know that someone has invented a working version of that transporter from Star Trek, or that user's credentials have been compromised.

Microsoft also provides a multi-factor authentication service for Azure AD users called Azure MFA. Azure MFA supports a variety of easy authentication methods out of the box for mobile devices, such as:

- Phone call
- SMS text message
- Mobile app notification via the Microsoft Authenticator app, allowing users to choose the method they prefer (e.g. Apple's TouchID)
- Mobile app verification code via the Microsoft Authenticator app
- 3rd party OATH tokens

This makes avoiding MFA methods identified as risky, such as the NIST warning regarding SMS text messages as a second factor, an easy task.

Together, MFA in Microsoft Enterprise Mobility and Security (EM+S), Microsoft Intune and Conditional Access with user identity form an effective perimeter, ensuring security of your sensitive corporate data wherever it may be accessed.

Securing Your Trusted Network

Mobility scenarios, like those mentioned in the previous section, receive a lot of attention in our cloud-first world. However, many organizations have established a (seemingly) secure network perimeter behind firewalls, proxies and VPN appliances and moved on, assuming their on-premises environment was secure. While your trusted corporate network may seem like the simplest resource to secure, it may in fact be the most vulnerable.

As described earlier in the "Anatomy of an Attack" image, some of the most common points of entry to your trusted network are through browser exploits, malicious document delivery and phishing attacks. What these exploits all have in common are that they target the greatest vulnerability on your network – the end user. This is not intended to insult your valuable human resources, but pointing out the reality that trusting users can be fooled into clicking malicious URLs. This may result in the opening of infected e-mail attachments that install malware or ransomware on client computers, letting hackers and thieves through your secure network perimeter undetected.

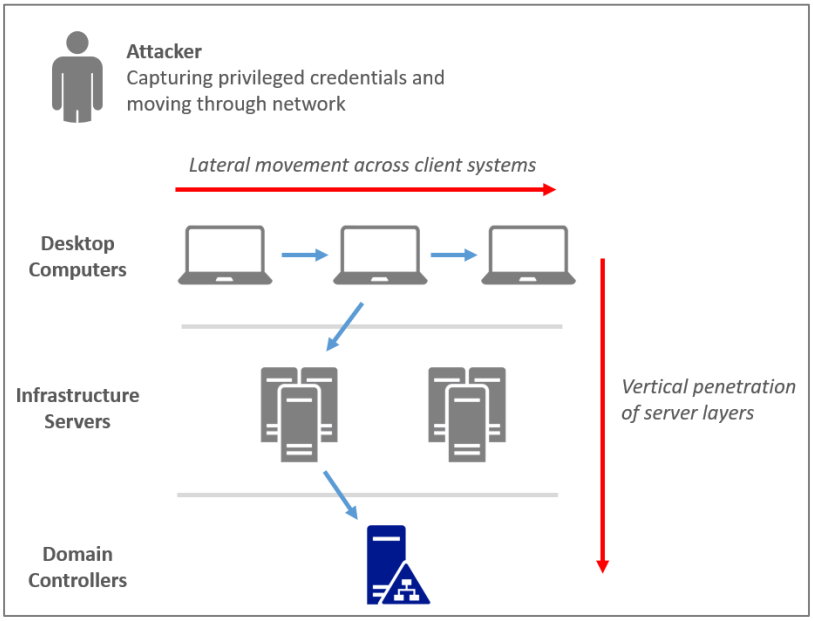

***Figure 5**. Lateral and vertical movement in a security breach*

This malware often lives undetected on your trusted network for an average of more than 140 days, listening to conversations, waiting to uncover network credentials, then stealing these secrets that enable lateral movement through your environment. This challenge is compounded by compromising more systems and uncovering more credentials, eventually enabling vertical movement from client to server. Sadly, there is little all the specialized security appliances on your network perimeter can do to stop it. Typical anti-virus software will block some known threats, but miss many others.

So, how do you defend against the weakest link in your trusted, on-premises network?

Microsoft Advanced Threat Analytics (ATA) is an on-premises solution that leverages behavioral analytics and machine learning to develop a rolling baseline of what constitutes "normal" behavior for users, devices and other resources in your environment. ATA takes information from multiple data-sources producing logs and events in your network to learn the behavior of users and other entities in the organization and builds a behavioral profile about them. ATA can receive events and logs from:

- SIEM solutions and appliances
- Windows Event Forwarding (WEF)

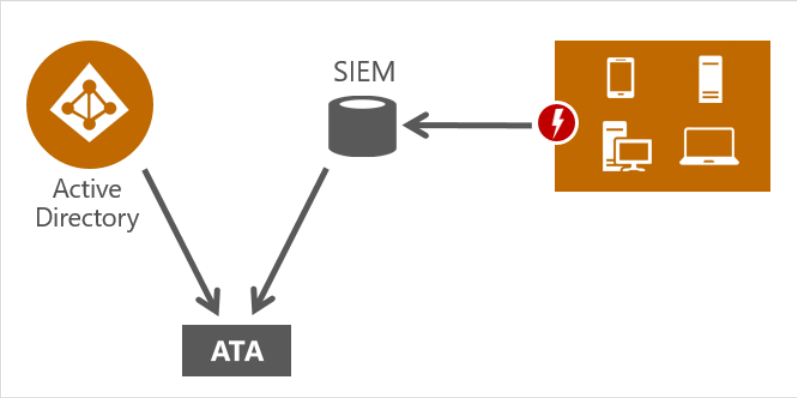

***Figure 6**. ATA Collection and Analysis Architecture*

In addition, ATA leverages deep packet inspection and a proprietary network parsing engine to capture and parse network traffic of multiple protocols (such as Kerberos, DNS, RPC, NTLM and others) for authentication, authorization and information gathering. This information is collected by ATA via:

- Port mirroring from Domain Controllers and DNS servers to the ATA Gateway
- Deploying an ATA Lightweight Gateway (LGW) directly on Domain Controllers

Microsoft created a new data source, entity-contextual (EC)-DPI technology, enabling ATA to analyze all levels of the network traffic to and from Domain Controllers. After analysis, ATA uses a unique deterministic detection engine to look for techniques commonly used by attackers, such as:

- **Reconnaissance**
  - Abnormal resource access
  - Account enumeration
  - Net Session enumeration
  - DNS enumeration
- **Compromised Credentials**
  - Abnormal working hours
  - Brute force using NTLM, Kerberos or LDAP
  - Sensitive accounts exposed in plain text authentication

    - Service accounts exposed in plain text authentication
    - Honey Token account suspicious activities
    - Unusual protocol implementation
    - Malicious Data Protection Private Information (DPAPI) Request
- **Lateral Movement**
    - Abnormal authentication requests
    - Abnormal resource access
    - Pass-the-Ticket
    - Pass-the-Hash
    - Overpass-the-Hash
- **Privilege Escalation**
    - MS14-068 exploit (Forged PAC)
    - MS11-013 exploit (Silver PAC)
- **Domain Dominance**
    - Skeleton key malware
    - Golden ticket
    - Remote execution
    - Malicious replication requests

These powerful detections enable new insights and faster discovery of compromises which ordinarily would go unnoticed by a breached organization.

4 – Protecting the Modern Perimeter

As we have described previously in this book, firewalls, proxies and anti-virus alone cannot protect your environment against zero-day threats and bad user behavior, even when unintentional. Consequently, a defense-in-depth approach is needed, providing multiple safeguards should one defense be breached. Additionally, given the assumption that attackers can breach defenses, advanced detection mechanisms are needed in order to provide capability for an adequate, timely response.

Leaning on a provider with unparalleled reach across the technology stack inside and outside the perimeter enables organizations to protect, detect, and respond to new security threats in ways that were previously unavailable. With Microsoft's vast array of security signals stretching across the globe and $1B/year investment into security, they provide proactive discovery, evaluation and mitigation of malware and new attack vectors, enabling customers to benefit from the collective exposure and experience of Microsoft's many other customers.

This advanced crowdsourcing (of sorts) is possible through the infrastructure, tools and experience of Microsoft, enabling customers to adopt a comprehensive and mature security posture much more quickly than would otherwise be possible, and compensates for what would otherwise be an unaffordable security team for organization to assemble.

While these features may be impressive on their own, let us now take a look at how you can combine these features to secure your organization across a number of common scenarios, including:

- Protecting Against Malicious Email
- Protecting the Endpoint
- Protecting the Infrastructure
- Protecting & Governing Cloud Applications
- Discovering & Protecting Data
- Discovering & Responding to Compromise

Protecting Against Malicious Email

It is a terrifying reality that one wrong click in an email can lead to an enterprise-wide breach that goes undetected for more than 140 days. Email has become perhaps the most common point of access for bad actors to gain access to trusted computing environments, such as your corporate network. Therefore, it is essential that organizations implement protections that pragmatically address the reality that many users are not technically savvy, and so will be less able to recognize cleverly disguised malicious content in the form of attachments and URLs.

For e-mail based attacks, whether through malicious URLs or malware-infected attachments, **Office 365 Advanced Threat Prevention (ATP)** provides a layer of defense to protect your organization from an end user's ill-advised decision to click on an unfamiliar URL or open that e-mail attachment from an unknown source.

Office 365 ATP provides two critical defenses for malicious email:

- **Malicious Web Links:** As emails come in, Exchange Online Protection (EOP) scans the emails for URLs. These URLs are re-written so that when the user clicks the link, they are directed to the Office 365 ATP filter to access the original URL. Before the Office 365 ATP filter redirects the user, it checks the reputation of the URL, checking to see if Bing or Microsoft's other intelligence signals have detected that URL as being malicious. This protects the user at time-of-click, ensuring that the user is protected from clicking on a link that transitions from benign to malicious weeks after the email is delivered.
- **Malicious Attachments:** As emails come in, Office 365 ATP takes the attachment and launches it in a detonation chamber / sandbox. Office 365 ATP monitors the activity in the sandbox to detect if any unusual behaviors occur. If the document were to deliver an executable, modify a registry key, or execute a command it should not, Office 365 ATP will not deliver the attachment.

The following diagram describes at a high level how Office 365 ATP safeguards incoming email:

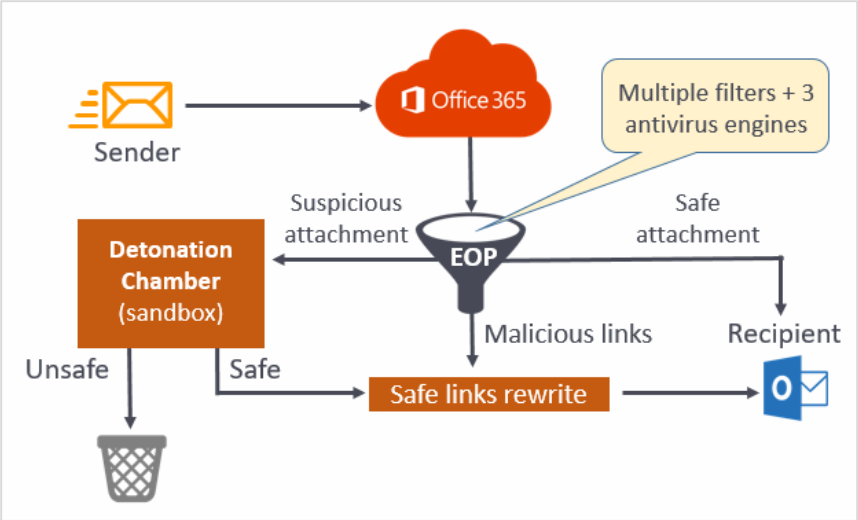

***Figure 7**. Protection from malicious email with Exchange Online Protection*

In addition to protecting against malicious links and attachments, Office 365 ATP provides advanced reporting, enabling the organization to see who clicked on a link discovered to be malicious. Paired with Windows Defender ATP (covered later in this chapter), the organization can gain deep insight into advanced persistent threats to the organization, and where to target remediation actions for endpoints or potentially compromised identities.

Protecting the Endpoint

At the machine level, protecting against credential theft from common attacks like 'pass-the-hash' mentioned earlier is a key point of prevention. While older versions of Windows desktop and server operating systems offer limited defense against such threats, Windows 10 and Windows Server 2016

offer a number of operating system features that defend against decisions of a trusting user consuming untrustworthy online content, including:

- **Device Guard** is set of features designed to work together to prevent and eliminate untrusted code from running on a Windows 10 system. Device Guard ensures that only trusted code runs from the boot loader onwards, and ensures the boot binaries and UEFI firmware are signed and have not been tampered with.
- **Credential Guard** is a specific feature that isolates and hardens key system and user secrets against compromise, helping to minimize the impact and breadth of pass-the-hash style attacks in the event that malicious code is already running locally or on the trusted network. It works by effectively isolating the process that manages credentials in an isolated, protected runspace. Remote Credential Guard extends this concept to protect user credentials in remote desktop connections.

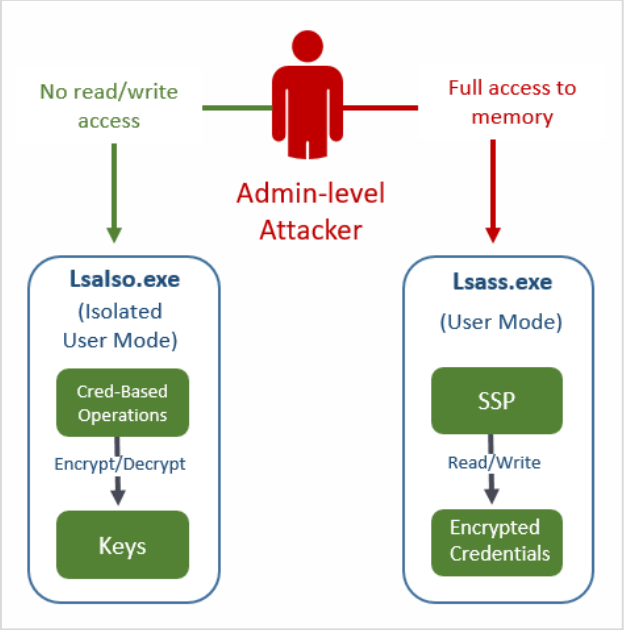

Figure 8. Process isolation with Credential Guard

- **BitLocker** is a drive encryption feature that integrates with the operating system and addresses the threats of data theft or exposure from lost, stolen, or inappropriately decommissioned computers. BitLocker helps render data inaccessible when BitLocker-protected computers are stolen, decommissioned or recycled.
- **AppLocker** enables organizations to deploy policies to control which apps and files users can run, including executable files, scripts, Windows Installer files, DLLs, packaged apps, and other installers.
- **Windows Defender Advanced Threat Protection (ATP)** is a cloud-based service that offers a *post-breach* layer of protection, enabling customers to detect, investigate, and respond to advanced threats on their networks. The service is composed of three components:

- The **client-end-point behavioral sensor**, built into Windows 10, that logs relevant security events and behaviors from the endpoint.
- Running on the Microsoft big data platform, the **cloud security analytics service** processes data from endpoints in combination with historical data and Microsoft's wide data repository to detect anomalous behaviors, adversary techniques and similarities to known attacks.
- **Microsoft and community intelligence**, composed of Microsoft's hunters and researchers who investigate the data, finding new behavioral patterns and correlating the data with existing knowledge from the security community.

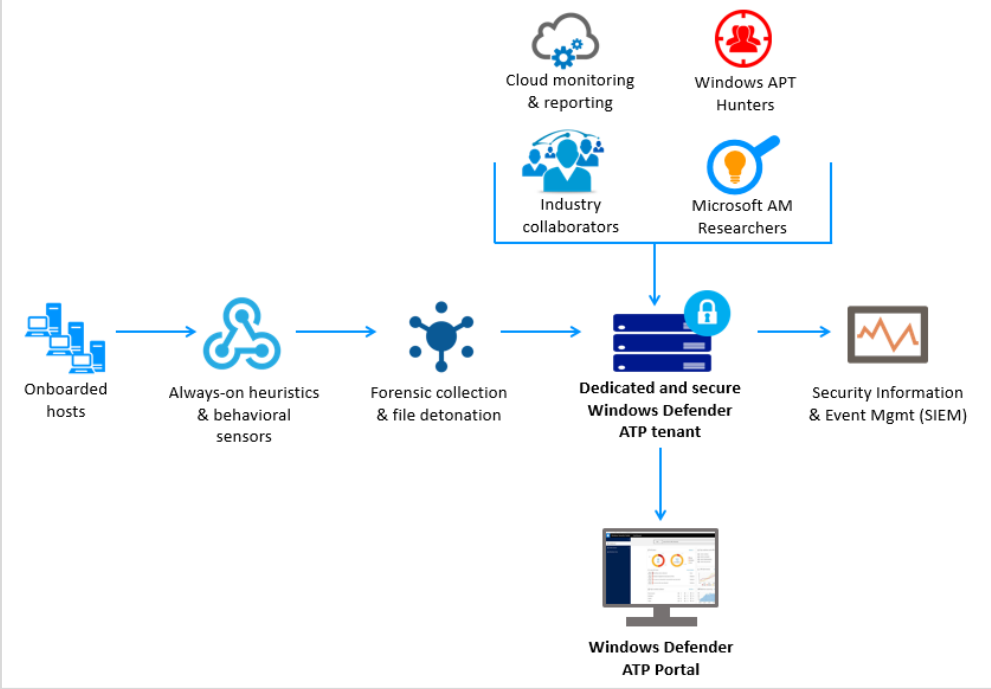

***Figure 9**. Windows Defender ATP Data Flow*

Windows Defender ATP provides the ability to drill down into anomalous behavior to read more about the specifics of the attack, how common the attack is, as well as if a known actor is responsible. Additionally, the service provides recommended mitigation or remediation steps.

In the end, a security strategy without a comprehensive post-breach component is not only incomplete, but represents an unacceptable risk to the business.

While Windows 10 has a number of powerful features for endpoint detection, third-party mobile platforms are not left out. Microsoft Intune provides two very important capabilities for iOS and Android mobile devices:

- **Mobile Device Management (MDM).** MDM is the concept most people think about when they think of enterprise mobility. MDM provides the ability to configure various device capabilities, such as encryption / PIN requirements, delivery of wifi, VPN, or certificate profiles, and even delivery of applications and associated management profiles. Intune MDM is required to check the compliance of a device with the configurations pushed to it, and is the only MDM platform that is able to stamp mobile device objects in Azure AD with the compliance status. It is therefore a key component in the Conditional Access story. Intune can also act as an extension of an organization's System Center Configuration Manager infrastructure, providing a single pane of glass for the organization to manage both internal infrastructure and their mobile devices.
- **Mobile Application Management (MAM).** MAM is the concept most people think about when they think of 'containerization'. MAM provides the ability for an organization to wrap controls around mobile applications to block copy/paste/save-as actions to personal data repositories or personal applications. Intune MAM is an extremely important capability, as it is the only mechanism by which organizations are able to apply these controls to the Office mobile applications (e.g. PowerPoint, OneDrive, Word, Outlook). Microsoft has also gone one step further, and in addition to being able to manage the Office mobile (and other applications) on **managed devices**, Intune MAM can also apply these application-level controls to **unmanaged devices**. This allows organizations to enable productivity for their workforce who choose not to enroll their personal devices for full device management.

Protecting the Infrastructure

Protecting the infrastructure means protecting the virtualization hosts, as well as the VMs and applications that run on them from both external and internal threats. Privileged access to the server fabric, VMs, applications and other resources must be limited to the greatest degree possible. Privileged credentials and other secrets must be secured against "pass-the-hash" and other attacks that discover and then use administrative credentials to move laterally and vertically through your infrastructure.

Attackers will target virtual machines (VMs), and in a shared environment such as a private cloud, hybrid cloud or other hosted infrastructure, the exposure is greater than in a traditional virtualized infrastructure.

Infrastructure threats include:

- Any compromised or malicious fabric administrators can access guest VMs
- Health of hosts not considered before running VMs
- VM exposure to storage and network attacks
- VMs cannot take advantage of hardware-based security capabilities such as TPMs (see note)
- Excessive rights, granted permanently, to delegate administrators

Note: A trusted platform module, or TPM, is a dedicated microprocessor designed to secure hardware by integrating cryptographic keys into devices.

Without additional protections, a malicious datacenter administrator or other threat agent could simply copy the VM hard disks out of your environment and begin to siphon data without oversight. This is a problem inherent to virtualized infrastructure.

Securing On-Premises Fabric

In Windows Server 2016, Microsoft introduced a powerful new protection and VM theft called **Shielded VMs**. Quite simply, Server 2016 virtualization hosts can now provide virtualized TPM functionality, enabling BitLocker within the VMs. For these hosts to provide virtualized TPM functionality, they are integrated with a service in the datacenter which checks that the host is a legitimate, healthy host, and provides the keys necessary to boot the protected VM.

A secure fabric (aka "guarded fabric") consists of:

- **Host Guardian Service (HGS)**. A server, or more typically a small cluster of servers running a health **attestation service**, which evaluates health information provided by virtualization hosts (thereby protecting against rootkits or other malicious software on the host) and provides them with an attestation certificate. HGS also provides a **key protection service**, storing keys necessary to boot a "Shielded VM" once the host proves it is healthy. These keys can be stored in hardware security modules (HSM), physical devices dedicated to storing and protecting your encryption keys, providing complete protection against VM theft by insider threats.
- **Guarded Hosts**. A Hyper-V host that can run shielded VMs. Before a guarded host can boot a Shielded VM, it must prove to the HGS that it is healthy. Once the host obtains a certificate attesting to its health, it uses the health certificate to request the keys necessary to boot the Shielded VM(s).
- **Shielded VMs**. A Shielded VM is a generation 2 VM (supported on Windows Server 2012 and later) that has a 'virtual TPM', is encrypted using BitLocker and can only run on healthy and approved hosts in the fabric.

The diagram below outlines the flow for a Guarded Host to boot a Shielded VM.

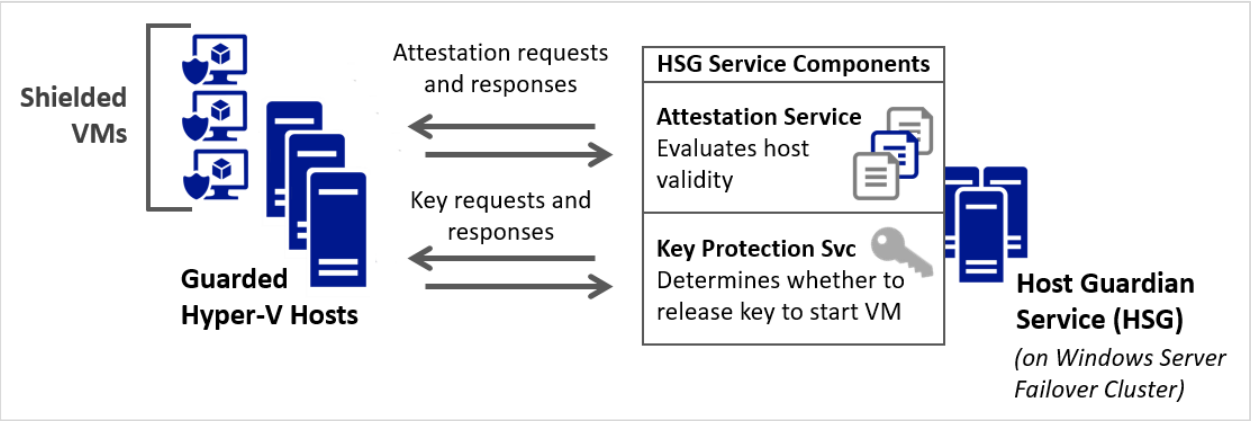

***Figure 10**. Secure Boot Process for Shielded VMs with HSG*

Operations Management Suite - Security & Audit

A new tool in an organization's security arsenal is the Operations Management Suite (OMS) Security and Audit solution. Organizations using System Center Operations Manager (SCOM) can plug into OMS for additional insights, or install the OMS agent directly on their servers. OMS Security & Audit provides insights into the following key areas:

- **Security Domains**: in this area you will be able to further explore security records over time, access malware assessment, update assessment, network security, identity and access information, computers with security events and quickly have access to the Azure Security Center dashboard.
- **Notable Issues**: this option will allow you to quickly identify the number of active issues and the severity of these issues.
- **Detections (Preview)**: enables you to identify attack patterns by visualizing security alerts as they take place against your resources.
- **Threat Intelligence**: enables you to leverage Microsoft's significant security intelligence resources to identify attack patterns against your organization by visualizing the total number of servers with outbound malicious IP traffic, the malicious threat type and a map that shows where these IPs are coming from.
- **Common Security Queries**: this option provides you a list of the most common security queries that you can use to monitor your environment. When you click on one of those queries, it opens a **Search** blade with the results for that query.

Securing Cloud Fabric

As organizations begin to build or migrate applications to the cloud, the traditional approach of leveraging internal public key infrastructure (PKI) services for certificate and cryptographic key storage begins to fall apart. Connectivity to internal certification authorities (CAs) and HSMs are not necessarily available to applications hosted in a public cloud provider. Additionally, as applications are increasingly built to serve audiences at scale, the limitations of internal certificate and key providers begins to impede the efficiencies and scale promised by public cloud services. Fortunately, Microsoft has developed a powerful tool to address this.

Azure Key Vault

Azure Key Vault helps safeguard cryptographic keys and secrets used by cloud applications and services. Key Vault is an HSM-backed service that provides encrypted storage of keys and secrets (such as authentication keys, storage account keys, data encryption keys, .PFX files, SSL/TLS certificates from third-party CAs, and passwords). These HSMs are fully FIPS140-2 compliant, and the service is designed so that Microsoft does not see or extract your keys, paired with robust monitoring and auditing. The keys are processed in the same datacenters as the Azure applications, reducing latency for application usage. These Azure Key Vaults are also regional, which means that secrets honor data residency requirements.

Azure Key Vault also provides full support for separation of duties. This allows separation of key vault management from management of the infrastructure or applications that might use the secrets. Access can be granted to Service Principals, enabling deployment of applications by DevOps personnel who do

not have knowledge of the secret. All access to the secrets are audited, enabling full compliance with any internal or regulatory requirements. With full support of template-able deployments of key vaults and resources which use them, organizations are easily able to follow the principles of least privilege access.

> *Keys used in existing encryption technology such as DRM and disk encryption products should be managed by central, internal to the enterprise, key storage technology. Hardware Security Modules (HSM) should be used to store keys as well as process cryptographic operations such as encryption/decryption, signing and verifying.*
>
> ~Cloud Security Alliance

Azure Key Vault supports bring-your-own-key (BYOK) scenarios. This allows organizations to leverage HSMs within their four walls to generate secrets, and securely import them into the Azure Key Vault service. This enables organizations to meet any sort of regulatory or compliance requirements which might require the secrets are only generated on-premises and always protected by HSMs. In fact, Azure Key Vault BYOK is the method in which Microsoft enables organizations to use BYOK encryption of Office 365 storage at-rest, and also the method used to enable BYOK for Azure Rights Management Services for data in-flight.

Azure VM Disk Encryption

While we've discussed leveraging BitLocker and Host Guardian Service to encrypt and protect VMs on-premises, organizations are also deploying VMs in Microsoft's datacenters where their Host Guardian Service is not necessarily available. Disk encryption is now available for Azure VMs, supporting both Windows (via BitLocker) and Linux (via DM-CRYPT). Azure Disk Encryption can be applied to both operating system and data disks, and maintains the cryptographic keys inside the aforementioned Azure Key Vault.

Azure Security Center

While many organizations are accustomed to having on-premises security systems to monitor system configuration, securing resources running on someone else's infrastructure is another ball game. Monitoring and configuration approaches need to be re-thought, and to make this easy Microsoft has introduced a security center for Azure resources. Azure Security Center helps you prevent, detect, and respond to threats with increased visibility into and control over the security of your Azure resources. Azure Security Center monitors the resources in your environment, and provides two key capabilities: Security Recommendations and Security Alerts.

Azure Security Center - Recommendations

The following list describes some key recommendations which may be triggered by one of your cloud resources:

- **Remediate OS Vulnerabilities:** Recommends that you align your OS configurations with the recommended configuration rules, e.g. do not allow passwords to be saved.
- **Apply Missing System Updates:** Recommends that you deploy missing system security and critical updates to your VM.
- **Apply Disk Encryption:** Recommends that you encrypt your VM's disks using Azure Disk Encryption (Windows and Linux VMs).
- **Enable Network Security Groups:** Recommends you apply Network Security Groups (essentially firewall rules) to networks which may have been provisioned without a security ruleset.
- **Enable SQL Server or Database Auditing:** Recommends you apply auditing capabilities to your Azure SQL server or database if they are not enabled.
- **Enable Transparent Database Encryption:** Recommends you enable SQL Transparent Database Encryption (TDE) to any databases that may have been provisioned without TDE enabled.

Azure Security Center provides many other recommendations, enabling organizations to get a handle on resources deployed in the cloud. Many of these recommendations provide simple ways to remediate discovered vulnerabilities right from the recommendation itself. Azure Security Center Recommendations is an excellent enabler for operational agility in the cloud without compromise of security control.

Azure Security Center - Alerts

The following are some examples of security issues that Azure Security Center can detect in an organization's Azure environment:

- **Outbound communication to a malicious IP address:** outbound traffic to a known botnet or darknet likely indicates that your resource has been compromised and an attacker is attempting to execute commands on that system or exfiltrate data. Azure Security Center compares network traffic to Microsoft's global threat database, and alerts you if it detects communication to a malicious IP address.
- **Suspicious process execution:** Attackers employ several techniques to execute malicious software without detection. For example, an attacker might give malware the same names as legitimate system files but place these files in an alternate location, use a name that is very similar to a benign file, or mask the file's true extension. Security Center models processes behaviors and monitors process executions to detect outliers such as these.
- **Hidden malware and exploitation attempts:** Sophisticated malware is able to evade traditional antimalware products by either never writing to disk or encrypting software components stored on disk. However, such malware can be detected using memory analysis, as the malware must leave traces in memory in order to function. When software crashes, a crash dump captures a portion of memory at the time of the crash. By analyzing the memory in the crash dump, Azure Security Center can detect techniques used to exploit vulnerabilities in software, access confidential data, and surreptitiously persist within a compromised machine without impacting the performance of your machine.
- **Lateral movement and internal reconnaissance**: To persist in a compromised network and locate/harvest valuable data, attackers often attempt to move laterally from the compromised machine to others within the same network. Azure Security Center monitors process and login

activities in order to discover attempts to expand an attacker's foothold within the network, such as remote command execution, network probing, and account enumeration.

- **Malicious PowerShell Scripts:** PowerShell is being used by attackers to execute malicious code on target virtual machines for a variety of purposes. Azure Security Center inspects PowerShell activity for evidence of suspicious activity.
- **Outgoing attacks:** Attackers often target cloud resources with the goal of using those resources to mount additional attacks. Compromised VMs, for example, might be used to launch brute force attacks against other VMs, send spam, or scan open ports and other devices on the internet. By applying machine learning to network traffic, Azure Security Center can detect when outbound network communications exceed the norm. In the case of spam, Azure Security Center also correlates unusual email traffic with intelligence from Office 365 to determine whether the mail is likely nefarious or the result of a legitimate email campaign.
- **Inbound RDP/SSH brute force attacks:** Your deployments may have busy virtual machines with a lot of logins each day and other VMs that have very few or any logins. Azure Security Center can determine baseline login activity for these VMs and use machine learning to define what is outside of normal login activity. If the number of logins, the time of day of the logins, the location from which the logins are requested, or other login-related characteristics are significantly different from the baseline, then an alert may be generated. Again, machine learning determines what is significant.

Securing Credentials In-Flight

Protecting credentials from being passed across the network and captured in pass-the-hash and similar attacks is ensured with:

- **Credential Guard**. Credential Guard prevents pass-the-hash and pass-the-ticket attacks by protecting stored credentials and credential artifacts using virtualization-based security (VBS)
- **Remote Credential Guard**. Remote Credential Guard works in conjunction with Credential Guard for RDP sessions providing single-sign on (SSO) for RDP sessions while eliminating the need for credentials to be passed to the RDP host

Least Privilege Operation

Implementing the practice of "least privilege", granting only the permissions necessary to a delegate administrator, has been around for a long time. Microsoft implements this concept at a new level of granularity in Windows with:

- **Just Enough Administration (JEA)**. JEA limits administrative privileges to the bare-minimum required set of actions (limited in space).
- **Just in Time Administration (JIT)**. JIT provides privileged access upon request through a workflow that is audited and limited in time.

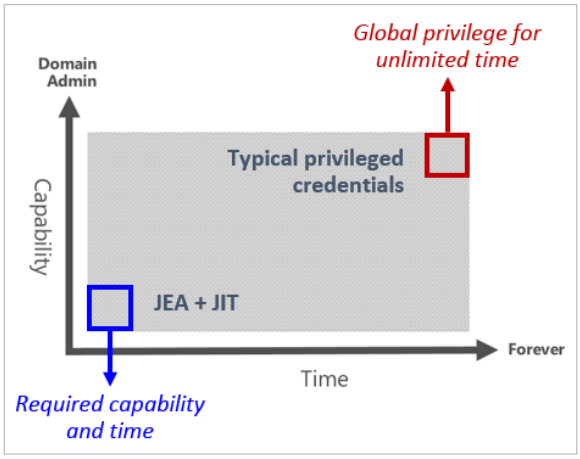

Figure 11. Least Privilege implementation with JEA + JIT

This also prevents the delegated administrator from granting privilege to other users outside approved corporate security policies.

Additionally, in Windows 2016, antivirus is installed and enabled by default. Windows Defender, Microsoft's in-box anti-malware solution that is also server workload aware, installed by default in Windows 2016, even on GUI-less VMs running Windows Server Core.

Protecting & Governing Cloud Applications

The path to the cloud for many organizations began with a credit card. Whether it was a Box cloud storage subscription to facilitate file sharing or an Amazon subscription to eliminate the weeks of waiting for development and test environments, the cloud gained momentum because it was a business enabler. The process of circumventing IT to facilitate greater agility for the business has become known as *shadow IT*. While the people responsible for this issue may seem like the enemy of IT, they are generally just well-meaning people trying to do their jobs. As the number of cloud applications in use multiplied over time, often with little or no oversight or governance from the IT or information security organizations, and just as frequently outside of planned IT budgets, a new problem was born. This problem is known as *cloud sprawl*. Today, it is common to hear of 100s or 1000s of 'unsanctioned' cloud applications in use within enterprise organizations.

Shadow IT places an organization at significant risk. These applications can become unmonitored points of data egress, with users saving corporate content to a third-party SaaS repository, where there is no organizational management of external sharing, and perhaps breaking regulatory and compliance rules on where particular types of data may reside. These applications can also have a negative financial impact to an organization, where users may be claiming the cost of using these applications through uncorrelated expense reports or unexpected spikes in usage of the non-managed service.

Discovery

The first step to securing and governing cloud applications is to secure the front door to applications in use via the organization's identity system and controls, as outlined in Chapter 2. In order to secure the front door however, an organization must know which applications are in use which must be secured.

This is where the Microsoft Cloud App Security (MCAS) comes into play. MCAS ingests logs from your network perimeter (firewalls and/or proxies) so you can easily identify which cloud applications are being used within your environment, who is using them, and how much of your network resources they are consuming. More importantly, MCAS has a portfolio of metadata on 13,000+ SaaS applications, containing information on the business (e.g. datacenter location, ownership, associated domains), compliance with industry certifications and standards (e.g. ISO 27001, HIPAA, FINRA, SOC 3), and security controls (e.g. data-at-rest encryption, admin audit trail, patched for Heartbleed vulnerability) that constitute a pre-defined risk-score assigned to each application. MCAS allows the organization to define importance of each of these metadata attributes, and dynamically computes a risk score tailored to your organization. This provides an immediate, intuitive understanding of risk exposure to discovered cloud applications, and assists the organization in deciding to block the application or to sanction and integrate via managed identity to secure the front door.

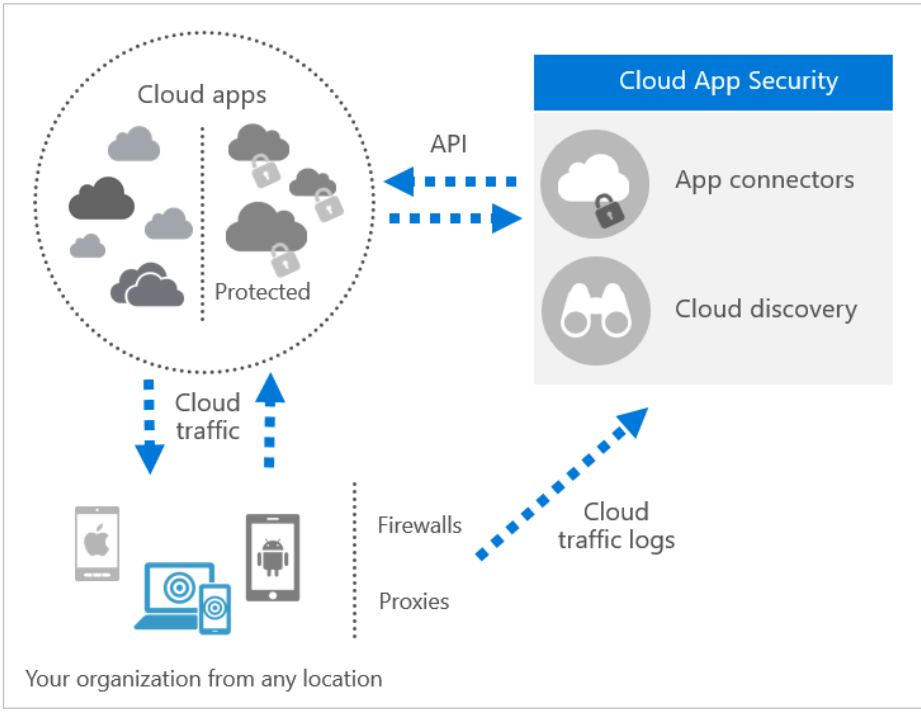

***Figure 12**. High level data flow and architecture of Microsoft Cloud App Security*

Investigation

Once an application has been sanctioned for usage, the organization can begin to get deeper insights on activities and files contained within that application. MCAS can plug into a variety of SaaS applications through their APIs, enabling insight into activities and files in that application, regardless of whether the

user is using the app from inside or outside the perimeter. Depending on the SaaS application being integrated (different app vendors support different capabilities in their APIs), MCAS can provide visibility into the following:

- **Account information:** Visibility into users, accounts, profile information, status (suspended, active, disabled) groups, and privileges.
- **Audit trail:** Visibility into user activities, admin activities, log on activity.
- **Data scan:** Scanning of unstructured data using two processes: periodically (every 12 hours) and a real-time scan (triggered each time a change is detected).
- **App permissions:** Visibility into issued tokens and their permissions.

These capabilities enable MCAS to automatically detect anomalous behavior within the application, or detect files that have been uploaded with particular content or metadata properties. It also allows the organization to perform investigations, searching for particular activities, files, or policy violations.

Control

While visibility into sanctioned cloud applications is a wonderful capability, being able to take action on the results is where MCAS completes the story. Depending on the SaaS application being integrated (different app vendors support different capabilities in their APIs), MCAS can execute the following types of actions:

- **Account governance:** Ability to suspend users, revoke passwords, etc.
- **Data Governance:** Ability to quarantine files, including files in trash, and overwrite files.
- **App permission governance:** Ability to remove tokens.

Completing the story with these capabilities, MCAS is able to take remediation actions when it detects particular activities or types of content stored in a cloud application. For example, a user might upload a file containing credit card information into Box from a managed BYOD device. While the front-door access to the application was secured through Azure AD's Conditional Access capability, this file may now violate an organizational policy on where credit card data may be stored. MCAS will detect the file having been uploaded, scan it, and then automatically quarantine the file (or perhaps remove sharing) and notify the user and their manager.

Protecting & Governing Data

No discussion of enterprise security would be complete without a look at data protection and governance. For purposes of this discussion, data comes in two forms:

- **Structured**. Structured data refers to kinds of data with a high level of organization, such as information stored in a relational database, as in Microsoft SQL Server.
- **Unstructured**. Unstructured data refers to data that is not contained in a database or some other type of data structure. Examples include email messages, Word documents, PowerPoint presentations and instant messages.

Important considerations in data protection and governance include data classification and rights management, encryption at-rest and in-flight, as well as management and storage of encryption keys and other secrets related to securing data.

Protecting Structured Data

Today it seems that not a month goes by without a report of some new massive data breach, with millions and millions of customer records being leaked to the internet. These large volumes of structured data being leaked are exceptionally damaging to an organization's reputation, and typically are the result of database infrastructure compromise. Sometimes these types of records are compromised through misconfigured applications or security controls, enabling attackers to run queries in a manner that should be unavailable, resulting in the database returning the attacker whatever data they are looking for. Other times these structured datasets are the result of insider threat or infrastructure compromise, where the attacker has been able to get hold of the dataset directly and siphon the file(s) out of the organization through surreptitious means.

Securing Structured Data In-Flight & In Use

SQL Server 2016 (both SQL in VMs and Azure SQL) introduces some new capabilities to prevent unintentional leakage of data by misconfigured applications or security controls. Key highlights are listed below:

- **Always Encrypted:** This is a client-side encryption capability, enabling the application to encrypt data so the SQL server (or service if using Azure SQL) can never see the data. This is particularly useful for protecting content such as SIN/SSN, Credit Card, and private health identifiers.
- **Row-Level Security:** This allows the organization to create policies which only return data rows appropriate for the user executing the query. For example, this allows a hospital to only return health information of patients directly related to a nurse, or a bank teller to only see rows returned which are relevant to their role.
- **Dynamic Data Masking:** This allows the organization to create policies to mask data in a particular field. For example, an agent at a call center may identify callers by the last few digits of their social security number or credit card number, but those pieces of information should not be fully exposed to the agent. Dynamic Data Masking can be configured on the SQL server to return the application query for the credit card numbers as XXXX-XXXX-XXXX-1234.

These capabilities help prevent and mitigate accidental exposure of data while it is in-flight or in-use by a front-end application.

Securing Structured Data At-Rest

In order to protect structured data at-rest, Microsoft first introduced SQL Transparent Data Encryption in SQL Server 2008. This technology protects data by performing I/O encryption for SQL database and log files. Traditionally a certificate that SQL Server manages (and is stored locally within the SQL master database) would protect this data encryption key (DEK). In June 2016, Microsoft made a significant enhancement to this capability by making generally available a **SQL Server Connector** for **Azure Key Vault**. This allows organizations to separate SQL and Security Administrator roles, enabling a SQL Administrator to leverage a key managed by the security operators in Azure Key Vault, with a full audit trail should the SQL administrator turn rogue. This connector can also be used for encrypting specific database columns and backups, and is backward compatible all the way back to SQL 2008.

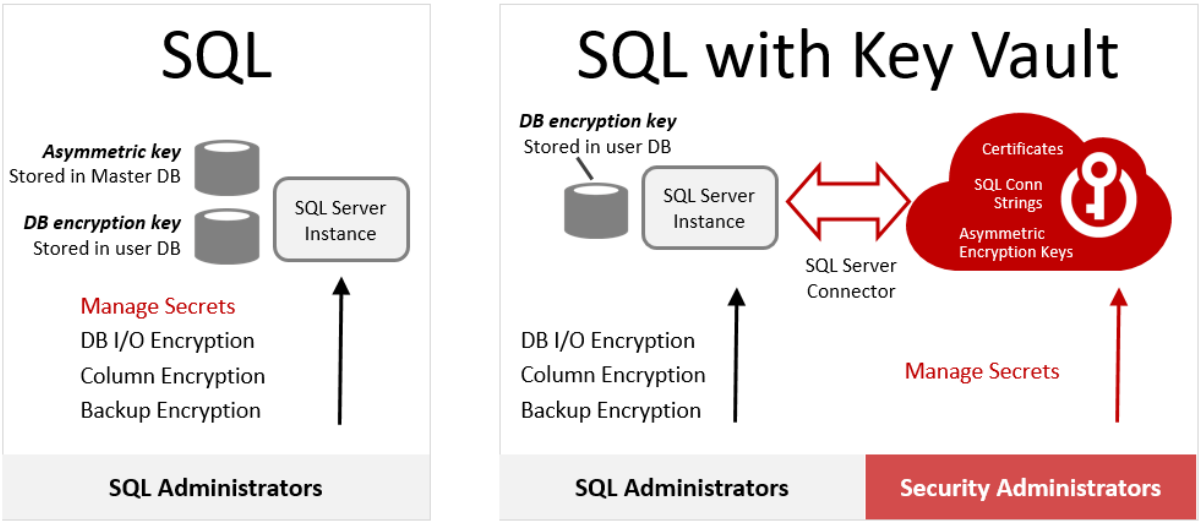

Figure 13. Key storage with SQL Server Connector and Azure Key Vault

Detecting Structured Data Threats

Running SQL in the cloud brings some additional benefits. For databases running on the Azure SQL service, the new **SQL Threat Detection** service monitors database activity and access, building profiles to identify anomalous behavior or access. If suspicious activity is detected, security personnel can get immediate notification about the activities as they occur. Each notification provides details of the suspicious activity and recommendations on remediating the threat.

SQL Threat Detection for Azure SQL Database can detect threats such as the following:

- **Potential Vulnerabilities:** SQL Threat Detection will detect common misconfigurations in application connectivity to the SQL data, and provide recommendations to the administrators to harden the environment.

- **SQL Injection Attacks:** One of the most common approaches to data extraction is to insert a SQL query into an unprotected web form, causing the form to return data that was unintended. SQL Threat Detection can identify if an attacker is attempting to leverage this mechanism to extract data.
- **Anomalous Database Access:** If a compromised database administrator account starts to execute queries from an abnormal location, SQL Threat Detection can detect and alert on the potential insider threat or identity compromise, enabling the security personnel to update firewall rules or disable the account.

SQL Threat Detection for Azure SQL Database is a powerful new tool in detecting potential data leakage threats.

Protecting Unstructured Data

While structured datasets tend to be the most frequent types of information breaches, leaks of unstructured data can be equally damaging. One only has to look at Wikileaks for evidence of the incredibly damaging power of leaks, like the Panama Papers. Unstructured data tends to be a much more difficult problem to solve, as this type of sensitive information could live anywhere, in any sort of file format.

Over a decade ago, Microsoft introduced a powerful technology called Rights Management Services (RMS) for encrypting information and providing usage policy enforcement (e.g. block copy/paste/print from users authorized to decrypt the content, or expiry at a date or interval). While very powerful, it was challenging to deploy and operate. Exposing the service, and particularly sharing protected data with partners involved a lot of IT overhead. With the advent of Azure AD, Microsoft was able to take advantage of the cloud directory and deliver RMS as a service called Azure RMS. Leveraging Azure AD reduced IT overhead for external partner collaboration to effectively zero, enabling very easy adoption of the service.

This service was then enhanced to enable two additional powerful capabilities: tracking and revocation. This allows users who have shared protected content to understand who is attempting to access the information, and from where. Should a document be discovered to have an unexpected amount of access attempts, the user can revoke access, effectively killing that document so it can never be accessed, even if the right account attempts to access it.

While protection, tracking and revocation are powerful capabilities for an organization to use, without understanding the corporate data which needs to be protected, it is effectively a useless capability. One of the most difficult challenges for organizations looking to protect information is to understand the data they have today. In order to understand existing data, it must be reviewed and classified. Traditionally, the approach to do this has been to go through a business exercise defining where corporate data is located, define what sort of information needs to be protected, and then configure the infrastructure to inspect for data at rest and in-flight and block or protect the content accordingly. This tends to be a prohibitively long, expensive initiative, causing many organizations to forgo protection of information at the item-level.

Classifying Unstructured Data

With the acquisition and integration of a company called Secure Islands, Microsoft completed the information lifecycle story, adding classification and labeling capability to the service now known as Azure Information Protection. Now at time of content creation or edit, the document's sensitivity can be detected and stamped with metadata, and optionally protected before it every traverses the network. As documents are classified, visual cues such as headers, footers, or background watermarks can be applied which describe the content or even the user who classified the content. Additionally, encryption and usage rights can be applied based on the classification, and what protection the organization has deemed appropriate to apply to that particular sensitivity classification. It is key to note that Azure Information Protection has virtually eliminated the historical challenges organizations faced in relying on the users to appropriately classify and protect data.

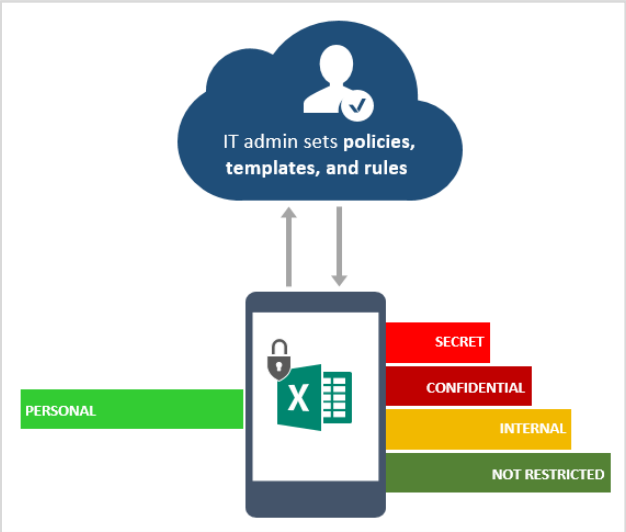

***Figure 14**. Data Classification with Azure Information Protection*

Classifying a document's sensitivity is accomplished in the following ways, alleviating complete dependency on a user to take the initiative to classify content:

- **Default Classification.** This is a default label (such as Internal) that the organization can apply to any unclassified content. This enables the organization to begin getting a handle of what data is considered corporate data, regardless of where it resides. The organization can enforce the application of sensitivity classification as content is being authored.
- **Manual Classification.** In the Office applications, a ribbon is surfaced with the available sensitivity classifications defined by the organization. The user can manually select one of the classifications, which triggers the stamping of metadata, and optionally visual cues or protection.
- **Recommended Classification.** The organization can define detection rules which look for particular patterns (such as credit card numbers or personally identifiable information, or PII data) or keywords (such as internal project codenames). Not every type of classification should be automatically implemented for the user however. For example, automatically classifying

financial data may impact some compliance or regulatory requirements. In the case of an organization defining a particular classification as recommended, when the user authors content which triggers a detection rule, a yellow bar appears in Office with a recommendation the user change the classification and a one-click button to do so.

- **Automatic Classification.** The organization can define detection rules which look for particular patterns (such as credit card numbers or PII data) or keywords (such as internal project codenames). In the case of an organization defining a particular classification as automatic, when the user authors content which triggers a detection rule, the document is automatically classified, and optionally stamped with visual cues or protection.
- **Override Classification.** In some cases, a document's sensitivity might change. For example, a product might be released to market, causing any marketing documents to no longer be considered secret information. Or perhaps a user has classified a document with a higher sensitivity than it actually is. The organization can allow the users to re-classify the content, but require a business justification. This enables the organization to audit and understand any reclassification applied to corporate data.

Infrastructure-Enforced Data Loss Prevention

Once a document has been stamped with metadata describing its sensitivity, it becomes much easier for the organization to take action on the data. A couple of common scenarios are described below:

- **Detect Sensitive Information in Email**. When a document is sent through Exchange (on-premises or in Office 365), Exchange Data Loss Prevention (DLP) rules can inspect the document's metadata or content to ensure that it is not tagged as highly sensitive information, and can take appropriate action. For example, for an organization just beginning to classify their data, perhaps an automatic reply to a user sending information tagged as Confidential to educate him/her on sending this type of information, or perhaps a rule to automatically encrypt any out-bound email tagged as Secret data. The possibilities for DLP are endless once information has been classified and becomes available as a trigger.
- **Detect Sensitive Information in Cloud Applications**. When a document is uploaded into SharePoint, OneDrive, or perhaps even third-party applications, MCAS can inspect the document's properties for sensitivity, and take remediation actions accordingly. For example, perhaps the user has uploaded a document to the corporate Dropbox account which is tagged with a default 'Internal' label, and has created a share link to give to an external user. MCAS could be configured to automatically remove the sharing link for that document and notify the user that what they have done violates corporate policy.

Operating System-Enforced Data Loss Prevention

Thus far we have been talking primarily about classifying, labelling, and protecting content at the user-level. In Windows 10 however, we have an additional tool available to us: Windows Information Protection (formerly Enterprise Data Protection). Windows Information Protection (WIP) is a capability of Windows 10, where the operating system tags and encrypts data as corporate based on where it comes from. The organization can define a whitelist of corporate locations (such as cloud service domain names, internal domain names, or subnets and IP ranges), and a list of applications that are allowed to access data from those locations.

Windows Information Protection provides the following key benefits:

- Obvious separation between personal and corporate data, without requiring employees to switch environments or apps. For example, imagine an HR person wants to copy a job description from an allowed app to the internal career website, but makes a mistake and pastes into the wrong web browser tab (perhaps their personal email account). With WIP, the paste action fails and a notification pops up, saying that the app couldn't paste because of a policy restriction. The HR person then pastes to the career website without issue.
- Additional data protection for existing line-of-business apps without a need to update the apps. After adding an app to your allowed apps list, the app is trusted with enterprise data. All apps not on this list are blocked from accessing your enterprise data, depending on your WIP management-mode. You don't have to modify line-of-business apps that never touch personal data to list them as allowed apps; just include them in the allowed apps list.
- Ability to wipe corporate data from devices while leaving personal data alone. This is a benefit when an employee leaves your company, or in the case of a stolen device. After determining that the data access needs to be removed, you can use Microsoft Intune to unenroll the device, and the next time it connects to the network the user's encryption key for the device is revoked, rendering the enterprise data unreadable.
- Use of audit reports for tracking issues and remedial actions. This allows attempts by insider threats to siphon data to be detected, and appropriate action taken.
- Integration with your existing management system (Microsoft Intune, System Center Configuration Manager, or your current mobile device management (MDM) system) to configure, deploy, and manage WIP for your company.

Leveraging the Windows 10 operating system's native functionality to prevent data loss is a powerful weapon in any organization's information protection arsenal.

Discovering & Responding to Compromise

A core tenant of the Microsoft defense stack is "assume compromise". People are not perfect and no security solution is infallible, so a layered defense that can detect a security breach (and scope of the breach) based on anomalous behavior and can recommend appropriate response is critical to data loss prevention. The Microsoft defense stack includes post-breach capability across several components, providing post-breach capabilities in a variety of scenarios.

Post-Breach Detection

Detecting both weak spots and actual breaches in the context of your computing environments, as opposed to a single device, is absolutely critical to providing context and visibility into the scope of items that need attention. It is one thing to see an alert on an infected computer in your trusted network. It is quite another to see lateral movement of a malicious entity in your environment through a suspicious pattern of behavior with a common set of compromised credentials.

While post-breach detection may feel "too little too late", it is actually a critical layer of defense, particularly as your efforts to mature your security posture in a race against an ever evolving threat

landscape. In this case, detecting and squashing lateral movement at the client tier can prevent the next step in the intrusion process…listening for and capturing privileged credentials that enable vertical movement into server and application tiers containing sensitive business and customer data. This is exactly what tools like Microsoft ATA's machine learning-in-a-box protects against, continuously evaluating normal user and device behavior, then alerting on anomalous events whether they affect one, ten or even a hundred devices or corporate data repositories.

Targeted Tooling for Actionable Insights

Moving up a level from these machine-level defenses, there are additional solutions and services enabling collation of machine-level events to paint a picture of malicious activity across the enterprise.

Too often, security solutions deliver information in a form that raises more questions than answers, requiring time-consuming research by the customer. Capturing all of the events, potential exposures and actual breaches is of little value if the information is not presented in a usable format with appropriate context. In short, the keys to effective detection and response from a cloud provider include:

- Making data easily searchable data, enabling tagging and saving searches where appropriate.
- Enable alert notification
- Present clear explanations of issues, with recommended remediation steps, rather than complex technical jargon that requires extensive research by the customer.

Incident prevention and response is made easier through greater correlation of events into fewer interfaces.

> *Cloud customers must understand the cloud service provider's (CSP's) support for incident analysis, particularly the nature (content and format) of data the CSP will supply for analysis purposes.*
>
> ~Cloud Security Alliance

The Microsoft security and identity stack does exactly this through targeted, function-specific security components. These security components deliver visibility and context for a variety of functions, including everything from auditing usage to proactive incident prevention, remediation and response. Key interfaces delivering comprehensive detection and response capabilities include:

- **Advanced Threat Analytics**. Presents alerts of suspicious activities, known security issues and malicious attacks detected in near real-time, providing clear, functional, actionable information on a simple attack timeline.
- **Azure Security Center**. Provides a central view of the security state of all of your cloud (Azure) resources. At a glance, verify that the appropriate security controls are in place and configured correctly. And quickly identify any resources that require attention.
- **Azure AD Identity Protection**. The Identity Protection dashboard enables investigation in response to risk events triggered by authentication attempts that violate your risk-based conditional access policies.

- **Azure RMS Tracking**. The RMS document tracking site enables easy visualization of document use, along with the ability to stop sharing or revoke access when undesirable sharing patterns and activities emerge.
- **Azure SQL Threat Detection.** Detect suspicious database activities (such as a compromised account), and detection for attacks such as SQL injection attacks from an application front-end using the database.
- **Microsoft Cloud App Security.** Detect anomalous activities occurring within third-party SaaS applications, and automatically execute remediation activities.
- **Office 365 Security & Compliance Center**. Manage compliance for all of your organization's data across Office 365. You can manage eDiscovery searches and holds, manage access for mobile devices, as well as view and manage security alerts.

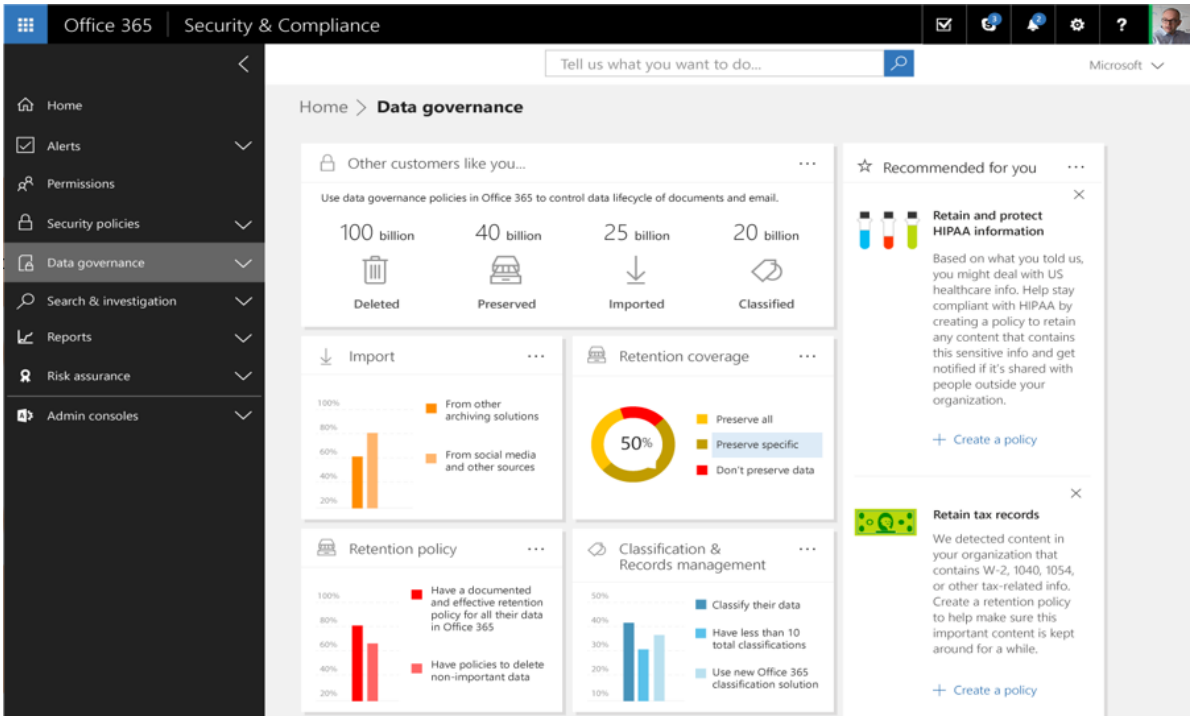

***Figure 15**. Office 365 Security and Compliance Center*

- **Office 365 Advanced Threat Protection**. Discover and mitigate malicious links in emails when the user clicks on the link, mitigating one of the most common threat vectors. Also detect attachments that run malicious code before they are ever delivered to the user. Track who in your organization has received and/or clicked on malicious attachments or links.
- **Operations Management Suite**. Discover malicious behavior occurring in your existing infrastructure, such as legitimate executables communicating with IP addresses from known botnets, or changing trends in authentication traffic or administrative group memberships.
- **Windows Defender Advanced Threat Protection.** Discover and mitigate malicious behavior occurring on the Windows endpoint, including usage of legitimate executables communicating with IP addresses belonging to known botnets.

5 – What's Next

We hope you have found the last couple of hours with us a good investment! Now that you have a better understanding of the challenges and key considerations in defending your network, as well as how the Microsoft Cybersecurity Stack can help protect your organization against modern threats, you may have an entirely new set of questions.

Next Steps

Ready to take the next steps with Microsoft security and identity? Contact the authors directly with any questions or for guidance on next steps for your organization! Click HERE or e-mail Pete and Wes at the addresses below.

Pete Zerger: Pete.Zerger@Lumagate.com
Wes Kroesbergen: Wes.Kroesbergen@Microsoft.com

Download the latest version of this book, get announcement of new versions and others in the series only at http://modernsecurity.info

Free Azure and EMS Trial

You can test drive all the components in the Microsoft cybersecurity stack for 30 days. Links to trial sign-up pages are listed below.

Sign up for a free 30-day trial of Microsoft Azure at
https://azure.microsoft.com/en-us/offers/ms-azr-0044p/

Sign up for a free 30-day trial of Microsoft Enterprise Mobility and Security at
https://portal.office.com/Signup/Signup.aspx?OfferId=2E63A04D-BE0B-4A0F-A8CF-407C1C299221&dl=EMS&box-skip-hip=1&ali=1#0