Office 365
Advanced Threat Protection

Frequently Asked Questions

Updated 11/08/2016

Table of Contents

General

What is Office 365 Advanced Threat Protection?
Office 365 Advanced Threat Protection (ATP) helps to secure your mailboxes against advanced threats, providing time-of-click protection against unknown malware and zero-day attacks. Advanced Threat Protection delivers several capabilities including Safe Attachments, Safe Links and rich reporting to help combat sophisticated attacks.

Will ATP catch 100% of malicious attacks?
No. In fact, no advanced threat protection product can catch 100% of malicious attacks, despite claims to the contrary. The notion of 100% protection is a misperception that is driven by the marketing and sales messages of some vendors in this industry.

What is Microsoft's SLA on virus and spam detection?
Microsoft's SLA on known viruses is 100%, with a spam effectiveness SLA of greater than 90%, a false positive ratio SLA of 1:250,000, and a monthly uptime SLA of 99.999%. Additionally, we keep continuously updated lists of malicious URLs that is checked approximately every 20 minutes.

Do I need to assign licenses and configure policies in order for Advanced Threat Protection to work?
Enabling ATP requires the configuration of policies in order to activate and target specific users, groups or domains to be protected by the service. You can configure separate policies for ATP to check links, attachments, or both. Safe Links and Safe Attachments policies can each be applied to specific sets of users. Learn how to do this at Set up a Safe Attachments policy in ATP and Set up a Safe Links policy in ATP. You can also create individualized policies within the Safe Links and Safe Attachments settings so that subgroups of users can have custom protection settings.

Assigning licenses is not a technical requirement, but it is required to be compliant.

How long does it take for Advanced Threat Protection policies to be effective?
Once a change is made to an Advanced Threat Protection policy, it can take up to 30 minutes for that change to propagate.

Does Advanced Threat Protection only work for Exchange Online mailboxes?

Advanced Threat Protection works for Exchange Online cloud-hosted mailboxes; on-premises customers running Exchange Server 2010, Exchange Server 2013, and Exchange Server 2016; and on-premises customers running non-Microsoft mail servers.

Can a user be configured only for Safe Attachments or only for Safe Links?
Yes, there are separate policies for Safe Links and Safe Attachments. Each policy can be applied to a specific set of users, distribution groups, or domains. It is also possible to have unique policies within Safe Links and Safe Attachments so each group of users can have custom settings.

Does it protect only internal mailboxes?
Safe Attachments scans incoming mail from outside the organization for all customers, as well as internal emails between employees for hosted mailbox customers. Safe Links is only applied for inbound traffic from external senders to internal recipients.

Does Exchange Online Protection (EOP) anti-malware work with Advanced Threat Protection?
Yes, Advanced Threat Protection complements EOP anti-malware filtering. Only those attachments that successfully pass EOP anti-malware scanning are impacted by Safe Attachments or Safe Links policies. The EOP anti-malware filtering is also designed to learn from ATP. Our ATP customers are protected immediately when ATP identifies a new threat, but we are also working constantly to improve the protection across the entire service.

Can on-premises organizations use Advanced Threat Protection?
Yes, Exchange organizations with on-premises mail servers can use Advanced Threat Protection, so long as they use Exchange Online Protection to route incoming messages.

What additional bandwidth is required once ATP is enabled for a tenant of 80k users? Will I see latency in service now that all email is ATP-scanned?

- **Email delivery**. If the Safe Attachments policy that applies to a particular recipient has an action of Block, the email will not be delivered until the attachments can be detonated by the Safe Attachments technology in ATP. Safe Attachments will launch a unique hypervisor to open the attachment. This can result in a delivery delay of 2-30 minutes for each mail evaluated by Safe Attachments.
- **Dynamic Delivery**. Dynamic Delivery is a new Office 365 ATP capability that is scheduled to be released later this year. Dynamic Delivery will eliminate the latency described above by delivering the body of an email with a placeholder attachment, to be replaced by the actual attachment after it has undergone a Safe Attachments scan. This allows recipients to read and respond to the message immediately, while also notifying them that the original attachment is still being analyzed.
- **Web browsing**. If a link points to a website recognized as not malicious, Safe Links adds very little latency to loading the target page. If the link points to a website recognized as malicious, the user is routed to a warning page and has to go through it (if click-through is enabled) in order to continue on to the site.

Note: After a change is made to an ATP policy, it can take up to 30 minutes for that change to propagate to every server.

Why should I be interested in Office 365 Advanced Threat Protection if I'm already using Windows Defender Advanced Threat Protection?

Windows Defender Advanced Threat Protection (WDATP) is a new service that builds on the existing pre-breach security features and services Windows 10 offers today. Windows Defender ATP provides a new post-breach layer of protection to the Windows 10 security stack that enables customers to detect, investigate, and respond to advanced and targeted attacks on their networks. Office 365 Advanced Threat Protection further supplements these defenses by providing focused protection for customers' email and messaging environments. Together, Windows Defender ATP and Office 365 ATP provide a comprehensive and robust set of protection and threat analysis tools for our customers. For more information, see Announcing Windows Defender Advanced Threat Protection.

How do we respond to Proofpoint's collateral that accuses Microsoft of missing a lot of malware attacks and positions us as an incomplete solution?
The recent Proofpoint materials are comparing their full solution to Exchange Online Protection only, and not Office 365 ATP. This is not a valid comparison of our competing security offerings.

Can ATP protect me from crypto malware (i.e. ransomware)? How does it determine this?
Both Crypto Locker and Locky are detected by Office 365 ATP. Office 365 ATP uses an internal sandbox technology which detonates the attachments in VMs and detects any anomalies. It uses several internal tools to detect vulnerabilities triggered by the sample being examined as well as other behavioral analysis to identify malicious activity.

Ok, but I need a solution that protects me from ransomware on my desktops and tablets. Does Office 365 ATP offer this?
Windows Defender ATP, not Office 365 ATP, provides protection for Windows. However, Office 365 ATP and WDATP share signals which provide a holistic solution.

In June 2016, Office 365 was the target of a massive, zero-day Cerber ransomware attack. How did Microsoft respond to the attack when it affected customers of EOP/ATP security?
Microsoft has a responsibility to continue to evolve the service to be ahead of the malware attacks. As part of our ongoing effort to provide better malware protection, the July 2016 release of the Microsoft Malicious Software Removal Tool (MSRT) includes detection for Win32/Cerber, a prevalent ransomware family. These additions to MSRT complement our Cerber-specific family detection in Windows Defender, and our ransomware-dedicated cloud protection features.

Does ATP identify phishing? How does it determine this? If so, what is the false positive rate?
Here are some of the key features as part of EOP and ATP that help combat phishing email threats:

- EOP has strengthened its counterfeit detection by over 500 percent and helps to protect against insider spoofing, also known as whale attacks, that target high profile users in an organization.
- Advanced Threat Protection, such as "Time of Click" malicious URL and "Zero-day" unknown malware protection.
- Strengthened coverage against malicious URLs by EOP and ATP.
- Implementation of key sender authentication technologies, such as DKIM and DMARC, provided by EOP.
- Improved protection against bulk mail provided by EOP.

Pricing and Licensing

What are the license requirements for shared mailboxes?
Shared mailboxes need to be licensed for Advanced Threat Protection. See the following example:

Company has 5 users in Office 365. Each has a mailbox. There is 1 shared mailbox. They need a license for:
6 seats of Exchange Online Protection (5 users + 1 shared mailbox)
6 seats of Advanced Threat Protection (5 users + 1 shared mailbox)

If licensing for Shared Mailboxes is blocking the deal, please escalate to the business desk.

Why is there an additional charge for Advanced Threat Protection?
Advanced security capabilities such as Advanced Threat Protection provide an additional level of security and control for those customers who require it. ATP is available as an add-on to all enterprise plans, in addition to being included in the E5 plan.

Am I able to try ATP without changing my MX record?
The supported method is for you to use Exchange Online Protection and point your MX record to Microsoft. If this is not possible, there is an alternative but not recommended (nor supported) method documented here: https://technet.microsoft.com/EN-US/library/jj937232(v=exchg.150).aspx.

Will my Exchange service change now that I've added the ATP feature?
Other than ATP adding 2 new layers of protection…no. Depending on the actions configured, email delivery delays may occur and/or all URLs will be rewritten, as mentioned previously.

Are any additional steps needed to verify ATP is working as expected?
The best way to verify that ATP is working is to test and then review ATP reporting trends to see whether the ATP actions have been taken. For Safe Links, reporting captures and displays the URLs that are recorded at the time of click. For Safe Attachments, you can run message trace to see whether a message with an attachment has been processed by ATP. The best perspective, however, will be provided by ATP's reporting features. Disposition reports show all of the actions that have been taken by ATP (i.e. Blocked, Monitored, etc.) and additional reporting breaks down those detections by file types. One thing to keep in mind is that with Safe Attachments, lower numbers of actions taken by ATP indicate higher performance by the overall EOP service. Safe Attachments policies are designed to capture zero-day (unknown) attacks. If the number of files you see being blocked in ATP is high, that means they've been able to slip through EOP. The real indicator of success is in the overall quantity of malware blocked by EOP and ATP, and the reduction in the number of escalations you have internally due to attacks getting through.

Please provide guidance on how to enable, manage, and monitor the service.
Learn how to do this at Set up a Safe Attachments policy in ATP and Set up a Safe Links policy in ATP.

How does the per user licensing model for Office 365 ATP work? Is it a tenant-wide solution or can customers choose to only protect a portion of the environment?
Customers can set up separate policies for ATP to check either links or attachments or both. Each policy can be applied to a specific set of users. Learn how to do this at Set up a Safe Attachments policy in ATP

and Set up a Safe Links policy in ATP. Customers can also create individualized policies within the Safe Links and Safe Attachments settings so that subgroups of users can have custom protection settings.

What happens (in tenant) after licenses are procured?
The ATP Advanced Threats tab should show up in the EAC within 30-60 minutes. When it does, this means that ATP is provisioned and ready to be configured to protect.

How can I get a trial set up for my customer?
You can set up a trial for your customers by sending your customer information to ATPTrials@microsoft.com and requesting a 30-day trial that includes 25 users. Alternatively, customers can try out Advanced Threat Protection and other advanced services when they obtain an Office 365 Enterprise E5 Trial subscription.

How do you enable ATP in my tenant?
Please refer to the TechNet article, Advanced threat protection for Safe Attachments and Safe Links.

Safe Attachments

What is Safe Attachments?
Safe Attachments helps to protect against zero day exploits in email attachments by blocking messages or attachments that could be malicious. Safe Attachments leverages sandboxing technology via a virtual environment to identify suspicious activity. Attachments that don't have a known malware signature are sandboxed and not released until a behavior analysis is performed and the attachments are determined to be safe. It is designed to detect malicious attachments even before antivirus signatures are available. If an email has multiple attachments ATP will treat them all as malicious if even one is identified as malicious.

What is the email delivery latency for attachments that are scanned by Safe Attachments?
This can result in a delivery delay of up to 30 minutes for each mail evaluated by Safe Attachments. The delay time varies depending on the file type with the average time being ~4 minutes.

Why is there a delay during scanning?
Typically, malicious payloads will activate in the sandboxing environment in a few seconds. As attackers become more sophisticated, they have built in delays of several minutes before activation in an effort to trick the sandboxing environment and evade detection. Office 365 ATP helps to safeguard organizations against delayed malware activation by conducting thorough scanning, which on average takes 4 minutes. Given enough time, all malware could be detected and blocked, but that would be very disruptive to mail flow because we would never be able to actually release the messages. Even with ATP protecting your mail flow it is important to ensure safe message handling practices are being used by your recipients.

What is Dynamic Delivery?
We recently announced a new capability in Office 365 Advanced Threat Protection called Dynamic Delivery. Dynamic Delivery eliminates latency by delivering the body of the email with a placeholder attachment while the actual attachment undergoes a Safe Attachment scan. Recipients can read and respond to the message, which includes notification that the original attachment is being analyzed. If the real attachment is cleared, it is automatically attached to the original message and replaces the

placeholder; if not, the message attachment is updated with an attachment which tells the recipient that their original message was identified as malware.

Can ATP identify malicious macros being executed? How does it determine this, since macros are commonly used by businesses?
Yes. A malicious macro would be detected with statistics and heuristics of our ATP detection system. In addition to ATP, a new policy-setting in Office 2016 allows administrators to block macros from untrusted sources.

Can ATP identify PowerShell being executed?
Yes. This would happen during detonation in the sandbox.

Can Advanced Threat Protection scan archives for malware? If so, what archive formats can it scan?
Advanced Threat Protection can drill into container files (e.g. archive, compression, special types) so long as they are not protected with a password. Advanced Threat Protection does not scan password protected archives or encrypted files.

How does Office 365 Advanced Threat Protection work with DRM/IRM encryption?
Office 365 ATP can drill into archive (compressed) files (such as .zip files). Advanced Threat Protection does not scan password protected archives or encrypted files. It does support container files (i.e. archive, compression, special types) as long as they are not protected with a password.

How does Advanced Threat Protection treat multiple versions of the same file? Does Advanced Threat Protection scan duplicates? For example, if 1000 users received the same file would all 1000 messages be detonated by Advanced Threat Protection?
Once the first file is scanned, the outcome will be applied to other recipients who have received the same file. For example, if File #1 was sent to Employee A and blocked, File #1 will be blocked for all other employees. File # 1 will also be blocked by reputation immediately for all other ATP tenants.

What file types will Safe Attachments detonate?
Safe attachments will analyze attachments that are common targets for malicious content, such as Office documents, PDFs, executable file types, and Flash files.

Is there a way to exclude a file from Safe Attachments inspection?
You can configure an Exchange Transport Rule to insert an X-Header, *X-MS-Exchange-Organization-SkipSafeAttachmentProcessing*, which ATP looks for in order to bypass ATP scanning.

How do I set up a Safe Attachments policy in ATP?
Please refer to the TechNet article, Set up a Safe Attachments policy in ATP.

Safe Links

What is Safe Links?
Safe Links provides real-time, time-of-click protection against phishing and malicious web sites by warning users when they click a link in email that has been determined to be unsafe. When a message is processed by EOP, all URLs are scanned and compared with our lists of known bad URLs. ATP Safe Links provides an opportunity to re-check the URL reputation lists when the recipient tries to follow the link.

Safe Links utilizes URL trace capabilities that allow you to track individual malicious links in messages that have been clicked to support faster remediation.

Does Safe Links add any latency to the browsing experience?
Safe Links adds very little latency to loading target web pages that have not been identified as malicious. Most people will not notice it at all. If the link points to a potentially malicious web site, then the user is routed to a warning page and must click through (if click through is enabled) to continue on to the site.

What logic does Safe Links use to detect malicious URLs? Does it have a special original malicious URL list or check the web site whenever a user clicks the URL?
We keep a continuously updated list of malicious URLs that is updated approximately every 20 minutes.

Does Safe Links also decrypt HTTPS and check what's inside of Host Header?
Safe Links does not decrypt https, it checks the destination URL to see if it is malicious or not.

How long is a Safe Links rewritten URL valid for? Will it ever expire?
The rewrite does not expire/time out.

Is our Safe Links technology a redirect or a fully-fledged proxy? How does it behave with a customer's own proxy device/service?
Safe Links is not a proxy service and does not follow the URL; it checks the health of the URL on an exact match at the time-of-click. We keep a continuously updated list of malicious URLs that is checked approximately every 20 minutes.

Can a URL be excluded from being rewritten?
Yes. This can be configured within a Safe Links policy using the "Do not rewrite the following URLs" option.

How does Safe Links identify a link within an email message?
For messages in HTML, Safe Links identifies any tag that uses the HREF attribute. For messages in plain text, Safe Links uses custom logic to identify any text resembling a URL.

How many links will Safe Links analyze (for example, when you click on a URL and you are redirected to another and another)?
It can handle the nested redirects if the first URL is a redirector that we trust (e.g., bit.ly). We are also working on enhancing Safe Links with URL detonation, which will sandbox-test URLs that point to a downloadable object like a PDF or application file.

Is there a way to exclude a URL from Safe Links inspection?
You can configure an Exchange Transport Rule to insert an X-Header, *X-MS-Exchange-Organization-SkipSafeLinksProcessing*, which ATP looks for in order to bypass ATP scanning.

Explain the URL routing for a customer utilizing EOP/ATP for Safe Links and a 3rd party URL category filtering product (for example: Websense, Barracuda).
Customers should have a good understanding of the routing and verify that other protections are not impacted by introducing this technology. The routing scenarios vary slightly if the customer is purely on-premises, online, or hybrid. For more information, see Transport routing in Exchange hybrid deployments.

How do I set up a Safe Links policy in ATP?
Please refer to the TechNet article, Set up a Safe Links policy in ATP.

New and future features

What is URL detonation?
URL detonation provides deeper protection against malicious URLs than can be provided by lists alone. In addition to checking the reputation of a link at time of click, Office 365 ATP's URL detonation feature will perform real-time behavioral malware analysis of suspicious attachments found at destination URLs within a secure sandbox environment. URL reputation checks are part of Advanced Threat Protection today; URL detonation will be in Preview by the last quarter of 2016.

How will ATP work across other Office 365 workloads such as SharePoint Online and Office applications?
ATP crawls the SharePoint library and detonates files after they've been uploaded to a SharePoint library. If a file is determined to be malicious, the icon representing the file will be marked to let the person who uploaded the file know that it's malicious. The same person will still be able to access the file. But to other users the file will not be visible.

Similar protection has been extended to Office applications Word, Excel, and PowerPoint. When a user clicks a link in a file created by one of these applications, Safe Links will provide time-of-click protection just as it does for links in email. And just like with email, URL trace capabilities will allow administrators to track individual malicious links that have been clicked to support faster remediation.

What is Office 365 Threat Intelligence?
The Threat Intelligence dashboard in the Office 365 Security & Compliance center will provide insight into the global threat landscape as well as the main risks to your organization. It will analyze data from Windows endpoints, Azure Advanced Threat Analytics, and Office 365 Advanced Threat Protection to alert organizations to threats being seen globally. Security administrators can use this data to track threat actors and trends as they unfold to better defend their organization's assets. The Threat Intelligence dashboard will go into preview the first quarter of 2017 and Generally Available (GA) the second quarter of 2017.

For more information about Office 365 Threat Intelligence, see the Threat Intelligence Internal FAQ.

What is Threat Explorer?
The Threat Intelligence dashboard helps admins understand the threat landscape in their organization and around the globe by displaying a variety of charts and widgets. Threat Explorer is the user interface where investigations are actually carried out. Using Threat Explorer, admins can get a detailed picture of the types of attacks they are facing, which of their users are being targeted, where malicious traffic is coming from, and remedial information to respond to those attacks. Threat Intelligence and Threat Explorer will greatly improve our reporting and analysis story related to ATP.

Reporting and analysis

Does ATP perform and capture static analysis, network analysis, and dropped files?
The ATP sandbox solution performs a combination of static and runtime analysis on files and URLs to detect security vulnerabilities. It is not a simple behavioral monitoring, but also involves tools running

along with the app associated with the attachment (i.e. Word, Excel). We are building solutions to take advantage of signals across not only cloud services, but also desktop endpoints to detect suspicious behavior across the entire organization.

Does ATP permit syslog export?
No.

Are files quarantined in EOP or ATP?
In EOP, spam is quarantined, but malware is not. In ATP, a copy of the original message along with any attachments tagged as unknown malware can be forwarded to a specific email address for further investigation, but they are not quarantined.

Can an email that was identified as malware be secured by a Security professional for additional analysis? Can this occur without releasing it to the intended recipient?
Yes. In ATP, a copy of the original message along with the attachment seen as having unknown malware can be redirected to a specific email address.

How can a customer get a report on Advanced Threat Protection detections?
Safe Attachments has two traffic reports which show aggregated data for a tenant by disposition (blocked, replaced etc.) and by file type. The report also shows detailed data (i.e. date, sender, recipient, ID, subject). Message trace and URL trace provide query options to see the details of the Advanced Threat Protection flow and actions on URLs and attachments. A few examples of behavior that would indicate the presence of malware would be content that:

- Writes to the system folder
- Injects a thread into an existing process
- Re-launches/deletes itself

The full list is dynamic and continually evolving as new methods are implemented by the adversaries. One thing to keep in mind is that with Safe Attachments, the lower the number in the report, the better the overall EOP service is working. Safe Attachments policies are designed to capture zero-day (unknown) attacks. If the number of files you see being blocked in ATP is high, that means they've been able to slip through EOP. The real indicator is the overall malware blocked by EOP and ATP, and the reduction in the number of escalations you have internally due to attacks getting through.

How can we get more detailed Advanced Threat Protection reports?
We are in the process of collecting feedback on which reports are most useful for our customers. We are committed to continually improving Advanced Threat Protection and publishing regular updates. To see a recent example, visit the Office blog where we updated the Advanced Threat Protection service to include Dynamic Delivery of attachments.

What is the latency for "Message Trace" results to show up in the reports?
Message Tracing shows every message in near real-time detail, including ATP verdict, evidence, and disposition. This information can also address helping organizations respond to mass attack situations and identify if there is a false positive since it includes the file name as well.

My customer is getting a lot of false positives. How can they report these to Microsoft and get feedback?

We recommend the customer take the following steps:

1. Check the following configuration settings:

    - **Safe Attachment settings**: Timeouts with Safe Attachments may be creating false positive. Check ATP Safe Attachment configuration and verify that the "Apply the above selection if malware scanning for attachments times out or error occurs" is not checked. You may want to keep the checkbox set to OFF to prevent losing any good attachments in case of error/timeout. The timeout for Safe Attachments scanning is 30 minutes. With the check box on, if a verdict is not reached within 30 minutes, the selected Safe Attachments unknown malware response action will be applied. If the checkbox is OFF, the message hitting the timeout/error will be delivered and assumed clean.

    - **Safe Attachment redirection**: In case of false positives, we recommend enabling Safe Attachments attachment redirection. This will help administrators recover any false positives due to error/timeout during scanning of the message by sending a copy of the original email and attachment to email address configured as part of the Safe Attachments policy.

    - **Safe Links policy setting**: A policy setting for Safe Links allows admins to configure the "Do not rewrite the following URLs" option which can help in addressing false positives by excluding certain URLs from being rewritten. Customers should also leverage the malware (anti-spam, antimalware) False Positive/False Negative (FP/FN) submission processes to submit email attachments.

2. If the customer is still seeing false positives, they can submit a report using the following link: https://www.microsoft.com/en-us/security/portal/submission/submit.aspx. To view the status of the report, click **Track your submission**. Note, response times may vary so we recommend opening a support case as well. For large customers or time sensitive issues, we also recommend opening a case with Premier Support.

We are always working to reduce false positives and appreciated your feedback as it helps to improve the customer experience. We are committed to continually improving our service and publish regular updates. To see a recent example, visit the Office blog where we updated the Advanced Threat Protection service to include Dynamic Delivery of attachments.

I have mail that I suspect contains a malicious URL. Can I submit this to Microsoft?
Yes, you can send email by using the Subject, "[Potential Malicious URL Submission]", to safelinksfeedback@microsoft.com. Please modify the URL so that the link cannot be accidentally clicked (i.e. change the '.' within the link to '_').

Compete Scenarios

How does Office 365 ATP compare to your competitor's solution in providing protection against unknown malware via email which typically goes undetected by traditional signature-based technologies?
The EOP/ATP filtering pipeline offers world class malware protection. Verdicts and reputation are shared between EOP and ATP to provide one integrated solution against malware. Both technologies protect

you from spam, commodity malware, targeted malware, advanced persistent threats, phishing, and spoofing by using a combination of reputation blocking, filtering, heuristics clustering, behavioral analysis, sandboxing, etc. Furthermore, Microsoft is uniquely positioned to respond to the evolving threat space. We have the team that built the hypervisor where we do the detection, the team that built the sandbox on top of the hypervisor, the Windows engineering group that built the operating systems that are being attacked, and a collection of cyber-attack hunters and the people that protect Microsoft from malware, all of which provide full end-to-end coverage. And coming soon, ATP will be extended across Office 365 to protect documents and files in SharePoint Online, and Office applications including Word, Excel, and PowerPoint.

Ease of use and administration. How much effort is involved?
ATP can take less than a minute or two to set up. Setup simply requires the quick creation of Safe Attachments and Safe Links policies.

Does Office 365 ATP provide malware sandboxing capabilities and threat insight to reveal who was targeted and a reporting summary to show incident history?
Yes, Safe Attachments scanning occurs in a sandbox environment where we spin up multiple hypervisor environments, each running various versions of the OS, Office, and common 3rd party applications. Evidence gathered as a result of detonation is presented within message trace. ATP Safe Attachments provides various traffic reports which show gathered data for a tenant by disposition (was it blocked, replaced, etc.?) and by file type. Reports also show detailed data (i.e. date, sender, recipient, ID, subject). Message trace and URL tracing also provide query options to see the details of the ATP flow/evidence and actions on URLs and attachments.

Does Office 365 ATP provide visibility and reporting for user actions like clicks on URLs? (Ex. What was the malicious URL that was blocked, who clicked on it and when did it occur?).
Yes, message trace and URL trace provide insight as to the URL that was clicked, the user that clicked it, and the time it was clicked.

Tell me what threats exist in my environment. Tell me which users are compromised by advanced threats. Protect me from commodity threats.
EOP and ATP protect you from commodity threats, targeted threats, and advanced persistent threats. We are actively working on improving our reporting. By the end of 2016 we will be rolling out a new UI called Threat Explorer which will allow you to drill down and navigate through your threat data to pinpoint hot spots to get a deep understanding of attacks, and let you quickly identify top threats, top targeted users. Threat Explorer will also support full search capabilities. We will also be releasing a Threat Intelligence dashboard in 2017 that will give you insight into the threat landscape both within your organization and around the world.

Data about threats is interesting but I want to see and understand the big picture.
See above – regarding Threat Intelligence dashboard.

Will your solution work by exposing APIs that allow organizations to pull alerts and threat intelligence data into SIEMs and log processing tools?
Common SIEM vendors include Splunk, AlienVault, and LogRhythm. While it is too soon to say which solutions we will integrate with or how, we will be exposing APIs that will allow organizations to pull our alerts and threat intelligence data into their SIEMs and log processing tools.

Do you provide a UI that is specific to the investigation of threats?
Both Threat Explorer and the Threat Intelligence dashboard will provide UIs with charts, graphs, data points, etc. that you can use to investigate threats and trends.