

Provision Lab Environment

This workshop will be using 2 different Azure Environments and access will need to be setup to both. For ease of identification we will call them Red Subscription and Blue Subscription. Different Scenarios below will use different Subscriptions.

Setup Red Subscription:

1. Collect Azure Pass Promo Code.
2. Make sure you have an email address that does not have access to a current subscription. If you do not have an email address available, please create a new email address with an online email service, such as www.outlook.com.
3. Visit microsoftazurepass.com to redeem code.
 - a. Review Instructions for Redeeming Code and Activating Subscription - <https://www.microsoftazurepass.com/Home/HowTo>
4. Ensure Azure Subscription is active by logging into <http://portal.azure.com> with same credentials used on #2.
5. Provision a single small, such as B2s size, Windows 2016 Virtual Machine with default options.

Setup Blue Subscription:

1. Collect Azure Login and Password.
2. Ensure Azure Subscription Access is working by logging into <http://portal.azure.com> with credentials used on #1.
3. If you can see a VM named 'aschack-vm1' then you are ready to go.
4. This is a shared environment so please do not change anything.

Azure Security Center Provisioning Scenario – Red Subscription

Azure Security Center provides unified security management and threat protection across your hybrid cloud workloads. While the Free tier offers limited security for your Azure resources only, the Standard tier extends these capabilities to on-premises and other clouds. Security Center Standard helps you find and fix security vulnerabilities, apply access and application controls to block malicious activity, detect threats using analytics and intelligence, and respond quickly when under attack. You can try Security Center Standard at no cost for the first 60 days.

In this scenario, you upgrade to the Standard tier for added security and install the Microsoft Monitoring Agent on your virtual machines to monitor for security vulnerabilities and threats.

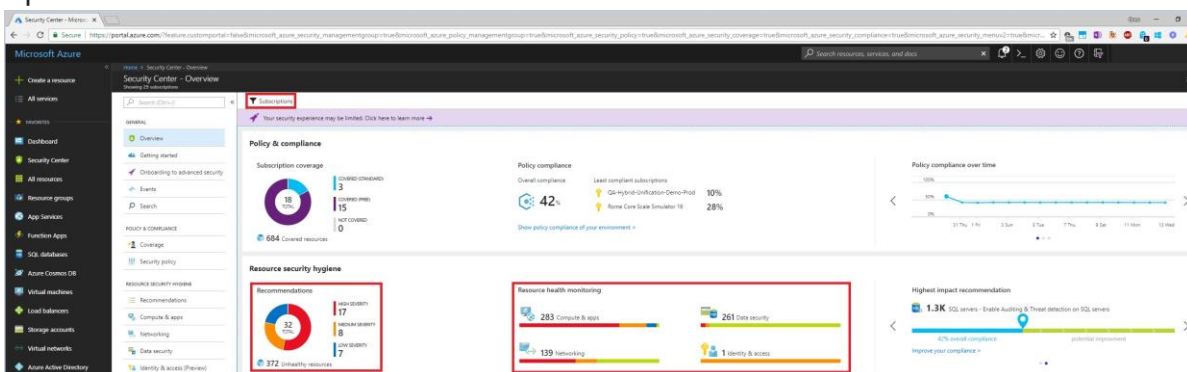
Prerequisites

To get started with Security Center, you must have a subscription to Microsoft Azure.

To upgrade a subscription to the Standard tier, you must be assigned the role of Subscription Owner, Subscription Contributor, or Security Admin.

Enable your Azure subscription

1. Sign into the [Azure portal](#).
2. On the **Microsoft Azure** menu, select **Security Center**. **Security Center - Overview** opens.



Security Center – Overview provides a unified view into the security posture of your hybrid cloud workloads, enabling you to discover and assess the security of your workloads and to identify and mitigate risk. Security Center automatically enables any of your Azure subscriptions not previously onboarded by you or another subscription user to the Free tier.

You can view and filter the list of subscriptions by clicking the **Subscriptions** menu item. Security Center will now begin assessing the security of these subscriptions to identify security vulnerabilities. To customize the types of assessments, you can modify the security policy. A security policy defines the desired configuration of your workloads and helps ensure compliance with company or regulatory security requirements.

Within minutes of launching Security Center the first time, you may see:

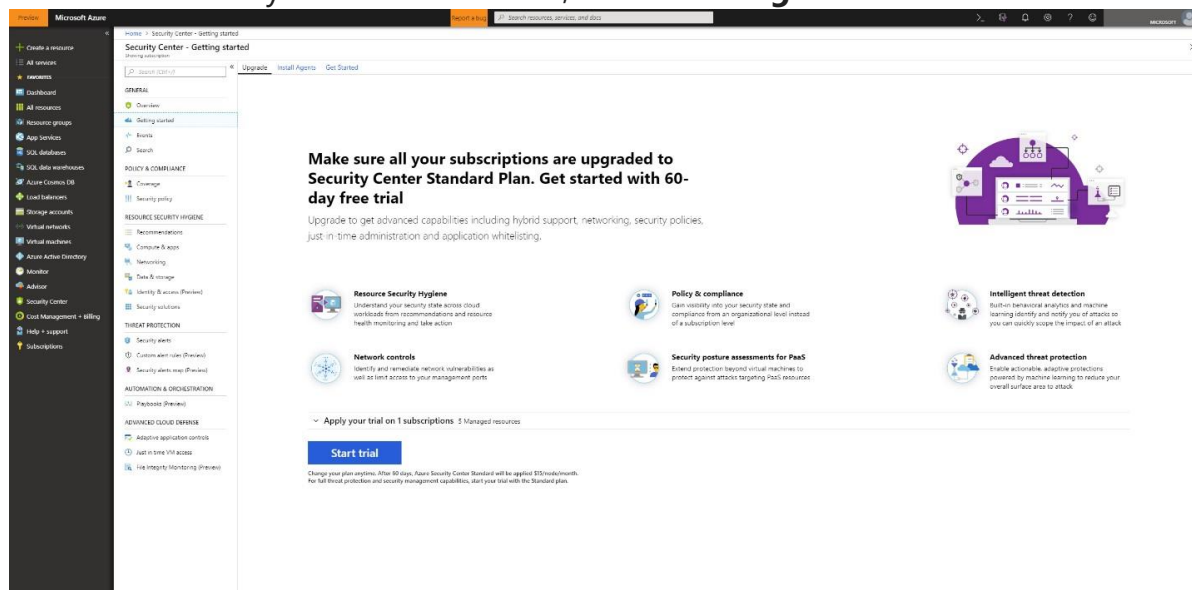
- **Recommendations** for ways to improve the security of your Azure subscriptions. Clicking the **Recommendations** tile will launch a prioritized list.
- An inventory of **Compute & apps, Networking, Data security, and Identity & access** resources that are now being assessed by Security Center along with the security posture of each.

To take full advantage of Security Center, you need to complete the steps below to upgrade to the Standard tier and install the Microsoft Monitoring Agent.

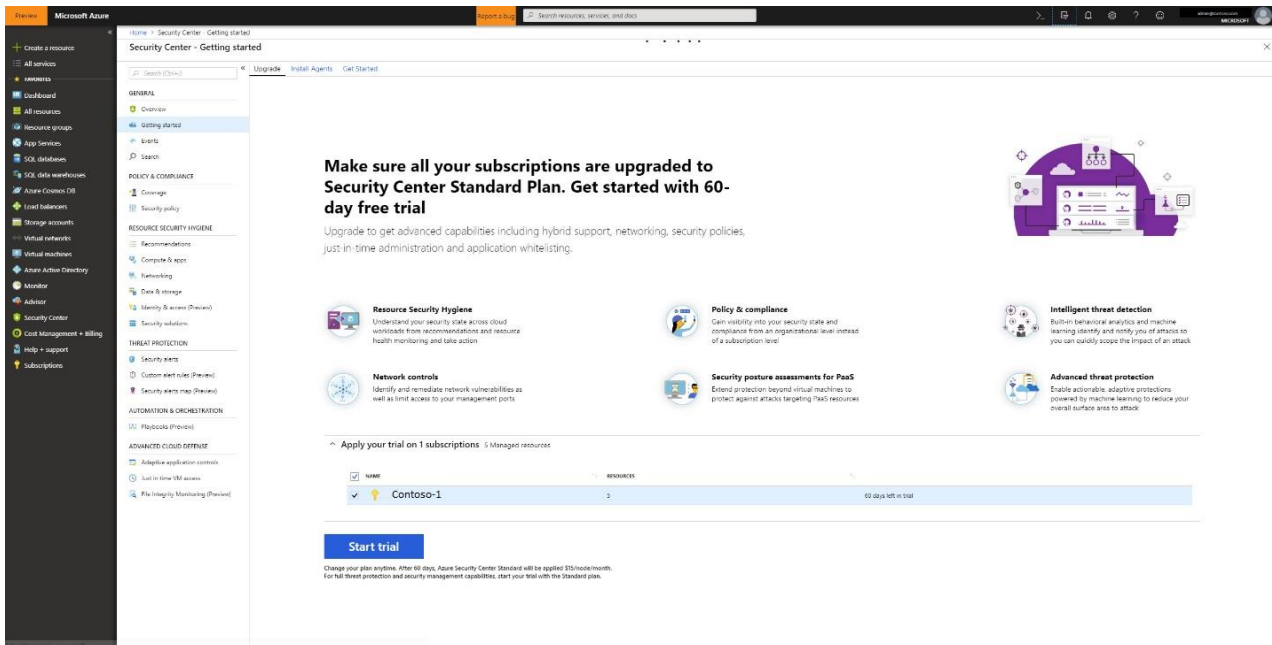
Upgrade to the Standard tier

For the purpose of this scenario you must upgrade to the Standard tier. NOTE: Your first 60 days of Security Center are free, and you can return to the Free tier any time.

1. Under the Security Center main menu, select **Getting started**.



2. Under **Upgrade**, Security Center lists subscriptions and workspaces eligible for onboarding.
 - You can click on the expandable **Apply your trial** to see a list of all subscriptions and workspaces with their trial eligibility status.
 - You can upgrade subscriptions and workspaces that are not eligible for trial. You can select eligible workspaces and subscriptions to start your trial.
3. Click **Start trial** to start your trial on the selected subscriptions.



Automate data collection

Security Center collects data from your Azure VMs and non-Azure computers to monitor for security vulnerabilities and threats. Data is collected using the Microsoft Monitoring Agent, which reads various security related configurations and event logs from the machine and copies the data to your workspace for analysis. By default, Security Center will create a new workspace for you.

When automatic provisioning is enabled, Security Center installs the Microsoft Monitoring Agent on all supported Azure VMs and any new ones that are created. Automatic provisioning is strongly recommended.

To enable automatic provisioning of the Microsoft Monitoring Agent:

1. Under the Security Center main menu, select **Security Policy**.
2. On the row of the subscription, select **Edit settings**>.
3. In the **Data Collection** tab, set **Auto provisioning** to **On**.
4. Select **Save**.

Home > Security Center > Security policy > Settings - Data Collection

Settings - Data Collection

Contoso IT - demo

Search (Ctrl+/) Save

Settings

- Data Collection
- Email notifications
- Pricing tier
- Edit security configurations

Security Center collects security data and events from your resources and services to help you prevent, detect, and respond to threats. [Learn more >](#)

Auto Provisioning

This enables the automatic installation of the Microsoft Monitoring Agent on all the VMs in your subscription. If enabled, any new or existing VM without an installed agent will be provisioned. [Learn more >](#)

Default workspace configuration

Data collected by Security Center is stored in Log Analytics workspace(s). You can elect to have data collected from Azure VMs stored in workspace(s) created by Security Center or in an existing workspace you created. [Learn more >](#)

Use workspace(s) created by Security Center (default)
Connect Azure VMs to report to workspaces created by Security Center

Use another workspace
Connect Azure VMs to report to selected user workspace

Any other solutions enabled on the selected workspace will be applied to Azure VMs that are connected to it. For paid solutions, this could result in additional charges. For data privacy considerations, please make sure your selected workspace is in your desired region.

With this new insight into your Azure VMs, Security Center can provide additional Recommendations related to system update status, OS security configurations, endpoint protection, as well as generate additional Security alerts.

Recommendations

Filter

Filtered by: State: Open, Resolved

DESCRIPTION	RESOURCE	STATE	SEVERITY	
Enable advanced security for subscriptions	1 subscriptions	Open	High	...
Endpoint Protection not installed on Azure VMs	6 virtual machines	Open	High	...
Endpoint Protection health failures	1 VMs & computers	Open	High	...
Add a Next Generation Firewall	2 endpoints	Open	High	...
Enable Network Security Groups on subnets	shsubnet	Open	High	...
Route traffic through NGFW only	vm3	Open	High	...
Enable Auditing & Threat detection on SQL servers	2 SQL Servers	Open	High	...
Remediate vulnerabilities (by Qualys)	2 virtual machines	Open	High	...
Apply system updates	2 VMs & computers	Open	High	...
Security configurations mismatch	60 VMs & computers	Open	Low	...

Permissions in Azure Security Center – Advanced

Azure Security Center uses [Role-Based Access Control \(RBAC\)](#), which provides [built-in roles](#) that can be assigned to users, groups, and services in Azure.

Security Center assesses the configuration of your resources to identify security issues and vulnerabilities. In Security Center, you only see information related to a resource when you are assigned the role of Owner, Contributor, or Reader for the subscription or resource group that a resource belongs to.

In addition to these roles, there are two specific Security Center roles:

- **Security Reader:** A user that belongs to this role has viewing rights to Security Center. The user can view recommendations, alerts, a security policy, and security states, but cannot make changes.

Security Administrator: A user that belongs to this role has the same rights as the Security Reader and can also update the security policy and dismiss alerts and recommendations.

NOTE

The security roles, Security Reader and Security Administrator, have access only in Security Center. The security roles do not have access to other service areas of Azure such as Storage, Web & Mobile, or Internet of Things.

Roles and allowed actions

The following table displays roles and allowed actions in Security Center. An X indicates that the action is allowed for that role.

ROLE	EDIT SECURITY POLICY	APPLY SECURITY RECOMMENDATIONS FOR A RESOURCE	DISMISS ALERTS AND RECOMMENDATIONS	VIEW ALERTS AND RECOMMENDATIONS
Subscription Owner	X	X	X	X
Subscription Contributor	X	X	X	X
Resource Group Owner	--	X	--	X

Resource Group Contributor	--	X	--	X
Reader	--	--	--	X
Security Administrator	X	--	X	X
Security Reader	--	--	--	X

NOTE

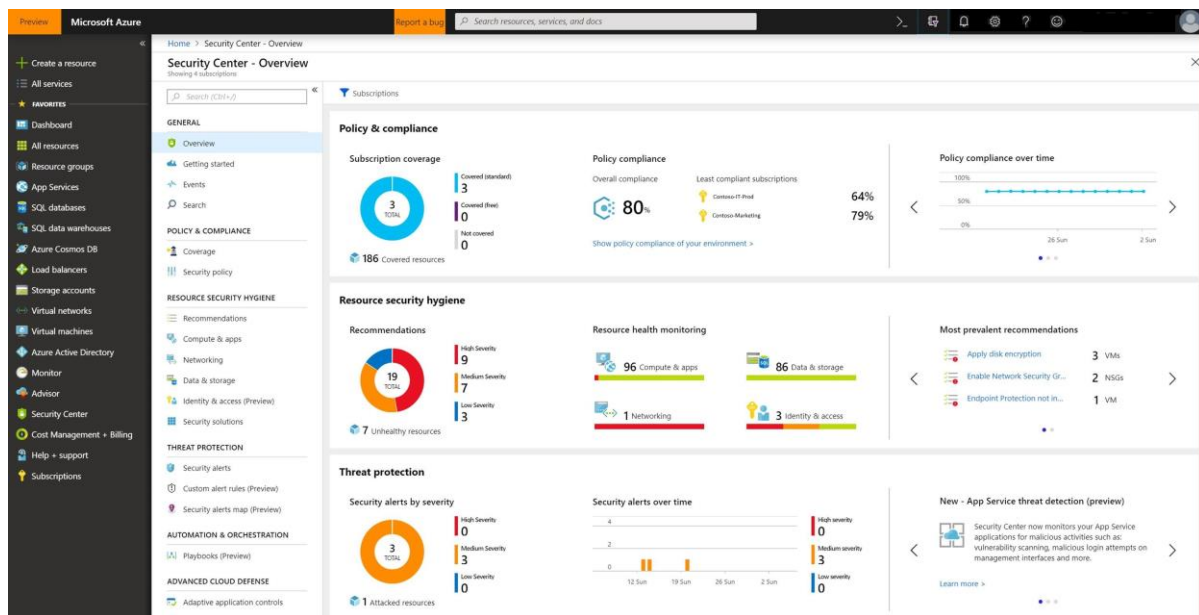
We recommend that you assign the least permissive role needed for users to complete their tasks. For example, assign the Reader role to users who only need to view information about the security health of a resource but not take action, such as applying recommendations or editing policies.

Onboard Windows computers to Azure Security Center – Review Only

After you onboard your Azure subscriptions, you can enable Security Center for resources running outside of Azure, for example on-premises or in other clouds, by provisioning the Microsoft Monitoring Agent.

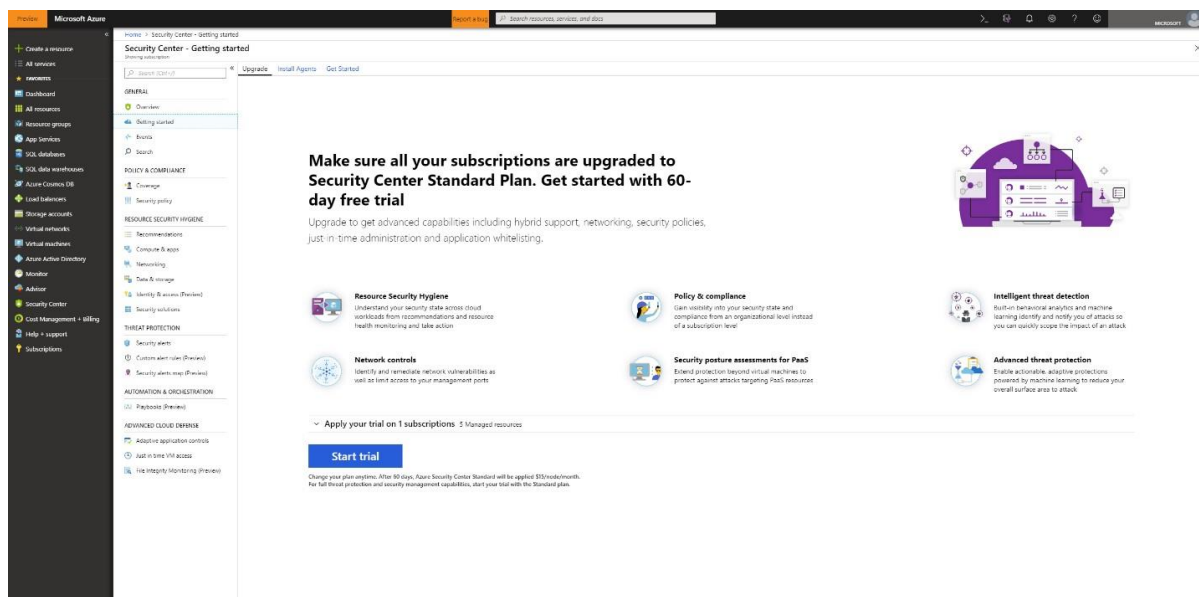
Add new Windows computer

1. Sign into the [Azure portal](#).
2. On the **Microsoft Azure** menu, select **Security Center**. **Security Center - Overview** opens.

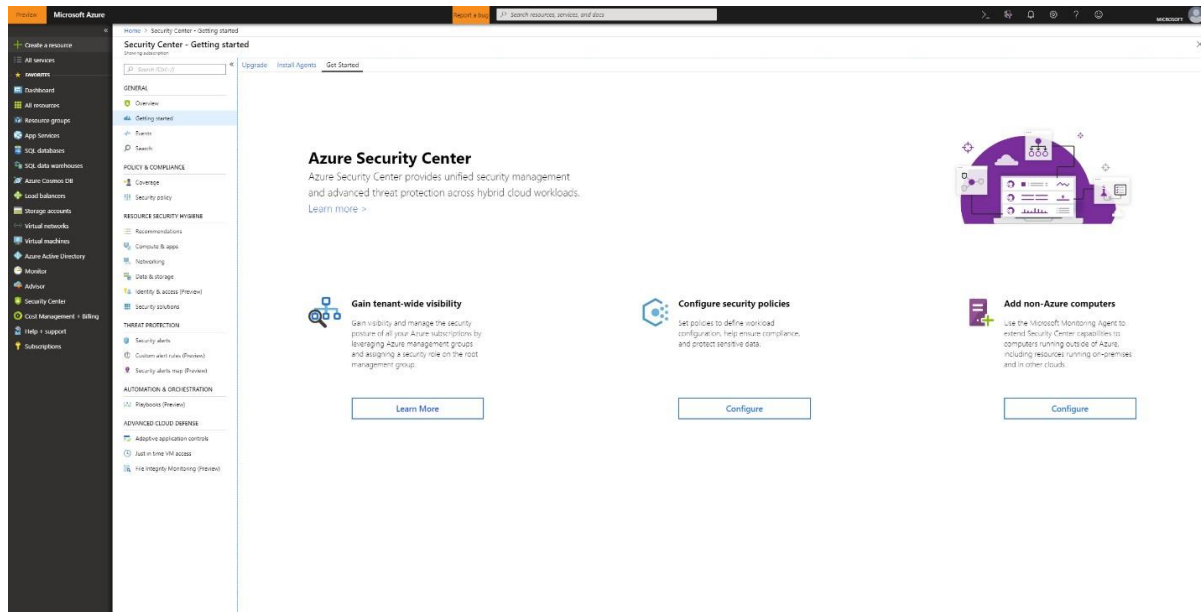


3. Under the Security Center main menu, select **Getting started**.

4. Select the **Get started** tab.



5. Click **Configure** under **Add new non-Azure computers**. A list of your Log Analytics workspaces is shown. The list includes, if applicable, the default workspace created for you by Security Center when automatic provisioning was enabled. Select this workspace or another workspace you want to use.



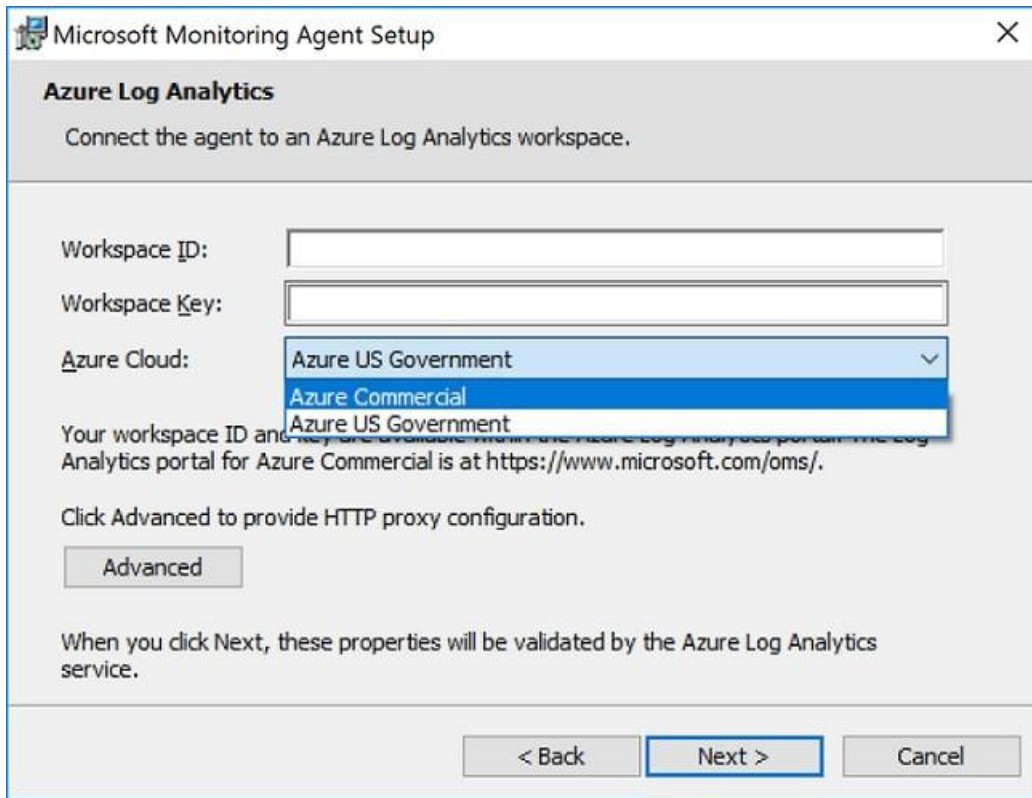
The **Direct Agent** blade opens with a link for downloading a Windows agent and keys for your workspace ID to use in configuring the agent.

6. On the right of **Workspace ID**, note the copy icon.
7. On the right of **Primary Key**, note the copy icon.

DO NOT Install the agent...

The following would be the steps to install the agent, please review.

1. Copy the file to the target computer and Run Setup.
2. On the **Welcome** page, select **Next**.
3. On the **License Terms** page, read the license and then select **I Agree**.
4. On the **Destination Folder** page, change or keep the default installation folder and then select **Next**.
5. On the **Agent Setup Options** page, choose to connect the agent to Azure Log Analytics and then select **Next**.
6. On the **Azure Log Analytics** page, paste the **Workspace ID** and **Workspace Key (Primary Key)** that you copied into Notepad in the previous procedure.
7. If the computer should report to a Log Analytics workspace in Azure Government cloud, select **Azure US Government** from the **Azure Cloud** dropdown list. If the computer needs to communicate through a proxy server to the Log Analytics service, select **Advanced** and provide the URL and port number of the proxy server.
8. Select **Next** once you have completed providing the necessary configuration settings.



9. On the **Ready to Install** page, review your choices and then select **Install**.

10. On the **Configuration completed successfully** page, select **Finish**

When complete, the **Microsoft Monitoring Agent** appears in **Control Panel**. You can review your configuration there and verify that the agent is connected.

For further information on installing and configuring the agent, see [Connect Windows computers](#).

Now you can monitor your Azure VMs and non-Azure computers in one place. Under **Compute**, you have an overview of all VMs and computers along with recommendations. Each column represents one set of recommendations. The color represents the VM's or computer's current security state for that recommendation.

Security Center also surfaces any detections for these computers in Security alerts.

NAME	MONITORED	SYSTEM UPDATES	ENDPOINT PROTECTION	VULNERABILITIES	DISK ENCRYPTION
mgmtvm2	▲	●	●	▲	●
dr01knWinSrv	▲	●	●	▲	●
dr01knWinSrv-test	▲	●	●	▲	●
ContosoWebVM1	▲	●	●	▲	●
ContosoAzASRV1	▲	●	●	▲	●
ContosoAzASRV2	▲	●	●	▲	●
ContosoClient1	▲	●	●	●	●
ContosoClient2	▲	●	●	▲	●
aks-nodepool1-85388480-0	▲	●	●	▲	●
ContosoAzLnx1	▲	●	●	●	●
vm0	●	●	●	▲	●
infoweb02.contoso.com	●	●	●	●	●
OpsInsights02.contoso.com	●	●	●	●	●

There are two types of icons represented on the **Compute** blade:



Non-Azure computer



Azure VM

Installing the agent on a Linux machine is very similar until the Installation, which is explained here.

1. On your Linux computer, open the file that was previously saved. Select the entire content, copy, open a terminal console, and paste the command.
2. Once the installation is finished, you can validate that the *omsagent* is installed by running the *pgrep* command. The command will return the *omsagent* PID (Process ID) as shown below:

```
FileEditViewSearchTerminalHelp
root@kronos:~# pgrep omsagent
7899
root@kronos:~#
```

The logs for the Security Center Agent for Linux can be found at: [/var/opt/microsoft/omsagent//log/](#)

```
<> [Search] [Menu] [Close]
Recent
Home
Desktop
Documents
Downloads
Music
Pictures
Videos
Trash
Computer
Floppy Disk
volume_1 on 192.16...
Connect to Server

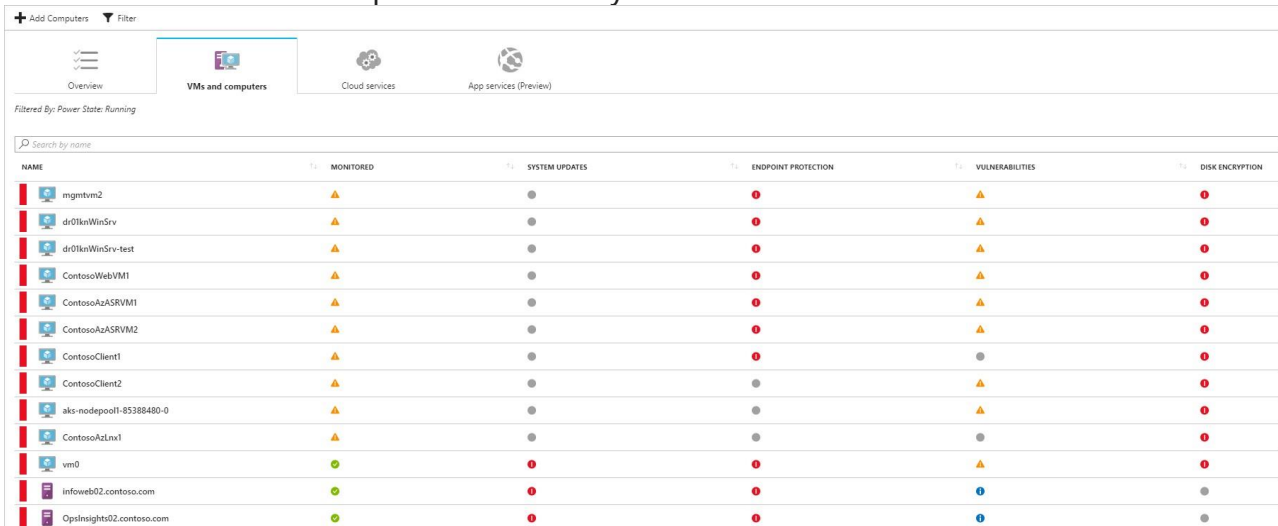
omsagent.log  urp.log.pos

Open [Icon]
/var/opt/microsoft/omsagent/[redacted].log

2017-12-12 13:04:47 -0600 [info]: reading config file path="/etc/opt/microsoft/omsagent-[redacted]/conf/omsagent.conf"
2017-12-12 13:04:47 -0600 [info]: starting fluentd-0.12.24 without supervision
2017-12-12 13:04:47 -0600 [info]: gem 'fluentd' version '0.12.24'
2017-12-12 13:04:47 -0600 [info]: adding filter pattern="oms.health.**" type="filter"
2017-12-12 13:04:47 -0600 [info]: adding match pattern="oms.health.** oms.heartbeat" type="out_oms"
2017-12-12 13:04:47 -0600 [info]: adding filter pattern="oms.operation.auditd_plu" type="filter"
2017-12-12 13:04:47 -0600 [info]: adding filter pattern="oms.operation.**" type="filter"
2017-12-12 13:04:47 -0600 [info]: adding filter pattern="oms.syslog.**" type="filter"
2017-12-12 13:04:47 -0600 [info]: adding match pattern="oms.blob.**" type="out_oms"
2017-12-12 13:04:47 -0600 [info]: adding match pattern="oms.** docker.**" type="out_oms"
2017-12-12 13:04:47 -0600 [info]: adding match pattern="diag.oms diag.oms.**" type="out_oms"
2017-12-12 13:04:47 -0600 [info]: adding source type="heartbeat_request"
2017-12-12 13:04:47 -0600 [info]: adding source type="monitor_agent"
2017-12-12 13:04:47 -0600 [info]: adding source type="oms_heartbeat"
2017-12-12 13:04:47 -0600 [info]: adding source type="dsc_monitor"
2017-12-12 13:04:47 -0600 [info]: adding source type="tail"
2017-12-12 13:04:47 -0600 [info]: adding source type="syslog"
2017-12-12 13:04:47 -0600 [info]: adding source type="exec"
```

After some time, it may take up to 30 minutes, the new Linux computer will appear in Security Center.

Now you can monitor your Azure VMs and non-Azure computers in one place. Under **Compute**, you have an overview of all VMs and computers along with recommendations. Each column represents one set of recommendations. The color represents the VM's or computer's current security state for that recommendation. Security Center also surfaces any detections for these computers in Security alerts.



NAME	MONITORED	SYSTEM UPDATES	ENDPOINT PROTECTION	VULNERABILITIES	DISK ENCRYPTION
mgmtvm2	▲	●	●	▲	●
dr01knWinSrv	▲	●	●	▲	●
dr01knWinSrv-test	▲	●	●	▲	●
ContosoWebVM1	▲	●	●	▲	●
ContosoAzASRV1	▲	●	●	▲	●
ContosoAzASRV2	▲	●	●	▲	●
ContosoClient1	▲	●	●	●	●
ContosoClient2	▲	●	●	▲	●
aks-nodepool1-85388480-0	▲	●	●	▲	●
ContosoAzLnx1	▲	●	●	●	●
vm0	●	●	●	▲	●
infoweb02.contoso.com	●	●	●	●	●
OpsInsights02.contoso.com	●	●	●	●	●

There are two types of icons represented on the **Compute** blade:



Non-Azure computer



Azure VM

Connect security solutions to Security Center – Review Only

In addition to collecting security data from your computers, you can integrate security data from a variety of other security solutions, including any that support Common Event Format (CEF). CEF is an industry standard format on top of Syslog messages, used by many security vendors to allow event integration among different platforms.

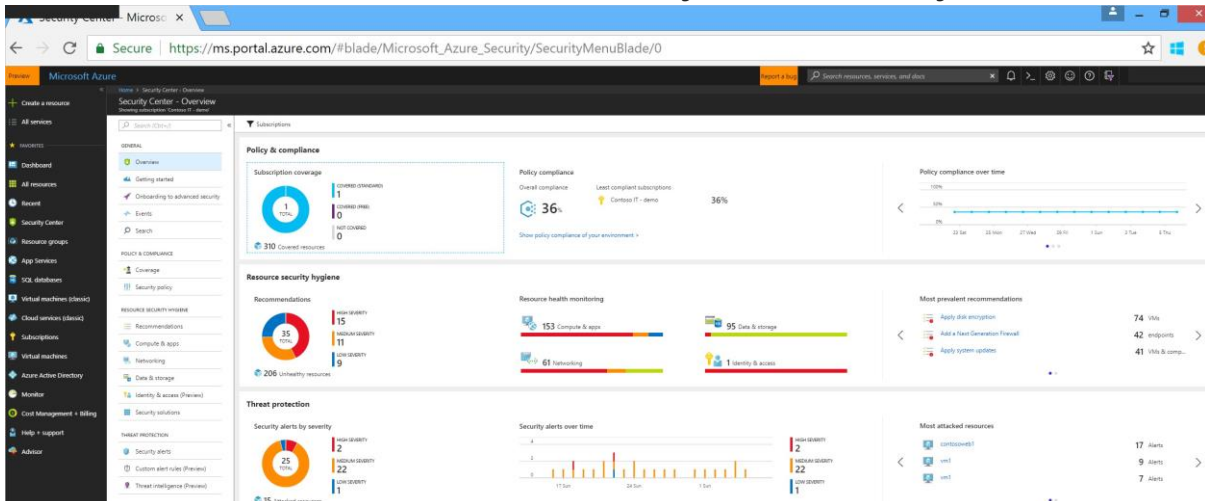
Prerequisites

You would need a [Linux machine](#), with Syslog service that is already connected to your Security Center, but this is a View Only scenario.

Connect solution using CEF

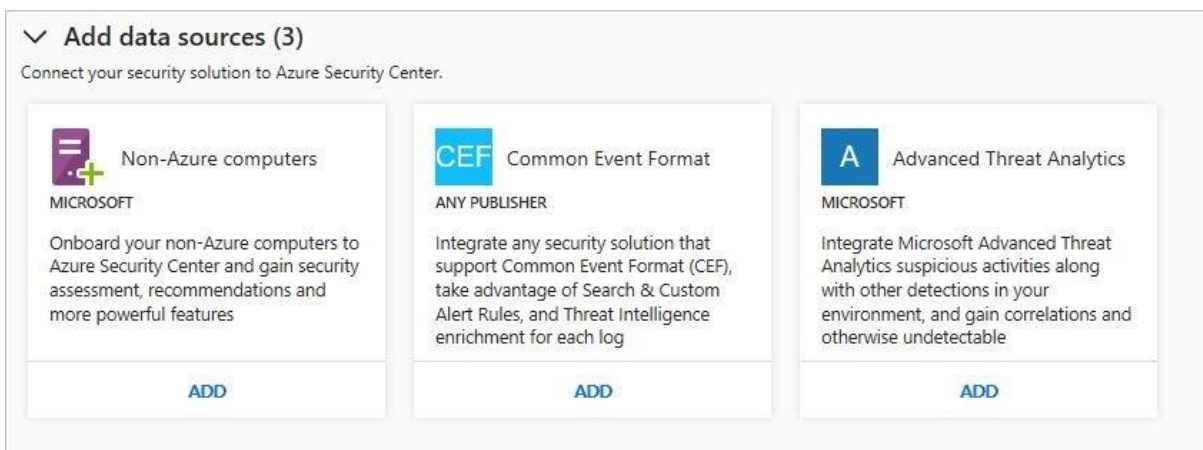
1. Sign into the [Azure portal](#).

2. On the **Microsoft Azure** menu, select **Security Center**. **Security Center - Overview** opens.



3. Under the Security Center main menu, select **Security Solutions**.

4. In the Security Solutions page, under **Add data sources (3)**, click **Add** under **Common Event Format**.



5. In the Common Event Format Logs page, expand the second step, **Configure Syslog forwarding to send the required logs to the agent on UDP port 25226**, and follow the instructions below in your Linux computer:

2. Configure Syslog forwarding to send the required logs to the agent on UDP port 25226

On the agent's computers, the events need to be sent from the syslog daemon to local UDP port 25226. The agent is listening for incoming events on this port. The following is an example configuration for sending all events from the local system to the agent (you can modify the configuration to fit your local settings):



If the agent computer has an rsyslog daemon:

In directory `/etc/rsyslog.d/`

Create new file

`security-config-omsagent.conf`

with the following content:

```
#OMS_facility = local4
local4.debug @127.0.0.1:25226
```

If the agent computer has a syslog-ng daemon:

In directory `/etc/syslog-ng/`

Create new file

`security-config-omsagent.conf`

with the following content:

```
filter f_local4_oms { facility(local4); };
destination security_oms { tcp("127.0.0.1" port(25226)); };
log { source(src); filter(f_local4_oms); destination(security_oms); };
```

6. Expand the third step, **Place the agent configuration file on the agent computer**, and note the instructions below for a Linux computer:

3. Place the agent configuration file on the agent computer

3.1. Place the following configuration file on the agent computer: `security_events.conf` Fluentd configuration file to enable collection and parsing of the events.

Destination path on agent:

`/etc/opt/microsoft/omsagent/{workspaceId}/conf/omsagent.d/`

3.2. Restart the syslog daemon:

rsyslog:

```
sudo service rsyslog restart
```

syslog-ng:

```
/etc/init.d/syslog-ng restart
```



7. Expand the fourth step, **Restart the syslog daemon and the agent**, and note the instructions below for your Linux computer:

4. Restart the syslog daemon and the agent

Restart the agent:

```
sudo /opt/microsoft/omsagent/bin/service_control restart [{workspaceId}]
```

Confirm that there are no errors in the agent log:

```
tail /var/opt/microsoft/omsagent/log/omsagent.log
```

The events will appear in Log Analytics under the CommonSecurityLog type.

[Click here to learn more](#)



To Validate the connection

You cannot do this in this scenario.

1. In the left pane, of the Security Center dashboard, click **Search**.
2. Select the workspace that the Syslog (Linux Machine) is connected to.

3. Type `CommonSecurityLog` and click the **Search** button.

The following example shows the result of these steps:

The screenshot displays a search interface with two main panels. The left panel shows a list of IP addresses and their counts, followed by a summary of device vendors and products. The right panel shows the search criteria and a detailed log entry.

Search Results (Left Panel):

10.51.211.245	4K
10.51.211.234	4K
131.107.71.82	1K
131.107.71.98	1K
131.107.38.1	1K
[+] More	
DEVICEVENDOR (2) ×	
Cisco	19K
Barracuda	10
DEVICEPRODUCT (2) ×	
ASA	19K
WAF	10

Search Criteria (Right Panel):

Type=`CommonSecurityLog`

19K Results [List](#) [Table](#)

11/13/2016 12:25:03.590 AM | CommonSecurityLog

- ... `TimeGenerated` : 11/13/2016 12:25:03.590 AM
- ... `ReceiptTime` : 1478402264139
- ... `DeviceVendor` : Barracuda
- ... `DeviceProduct` : WAF
- ... `DeviceEventClassID` : 1000
- ... `Activity` : GET
- ... `LogSeverity` : 5
- ... `Computer` : barracuda
- ... `DestinationIP` : 10.0.0.10
- ... `ReceivedBytes` : 123
- ... `RequestURL` : /
- ... `RequestMethod` : GET

Azure Security Center Prevent Scenario – Red Subscription and Blue Subscription

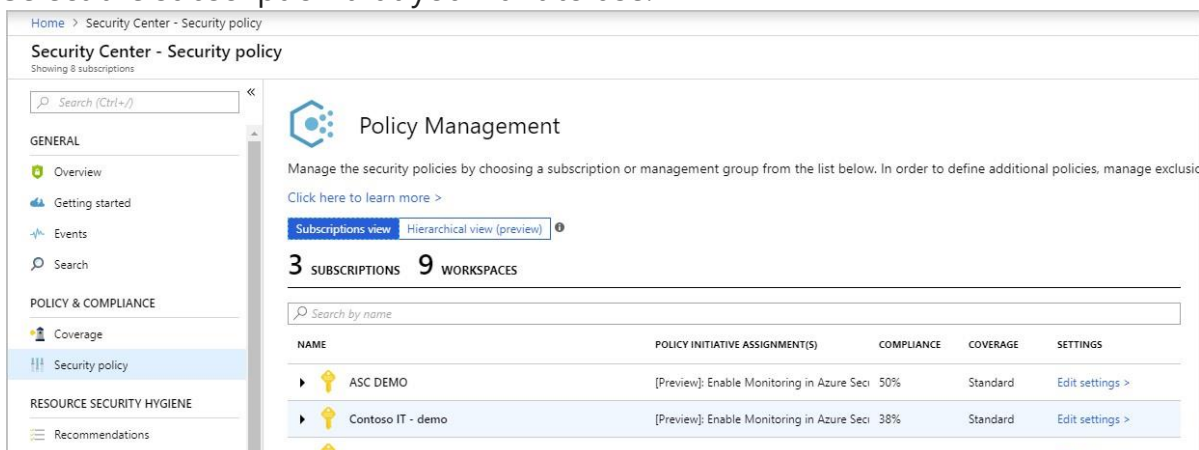
Security Center helps ensure compliance with company or regulatory security requirements by using security policies to define the desired configuration of your workloads. Once you define policies for your Azure subscriptions, and adapt them to the type of workload or the sensitivity of your data, Security Center can provide security recommendations for your compute, application, networking, data & storage, and identity & access resources. In this scenario, you will learn how to:

- Configure security policy – Red Subscription
- Assess the security of your resources – Blue Subscription

Configure security policy – Red Subscription

Security Center automatically creates a default security policy for each of your Azure subscriptions. Security policies are comprised of recommendations that you can turn on or turn off according to the security requirements of that subscription. To make changes to the default security policy, you need to be an owner, contributor, or security administrator of the subscription.

1. At the Security Center main menu, select **Security policy**.
2. Select the subscription that you want to use.



The screenshot displays the 'Security Center - Security policy' interface. On the left is a navigation pane with sections: GENERAL (Overview, Getting started, Events, Search), POLICY & COMPLIANCE (Coverage, Security policy), and RESOURCE SECURITY HYGIENE (Recommendations). The main area is titled 'Policy Management' and includes a search bar, view toggles for 'Subscriptions view' and 'Hierarchical view (preview)', and counts for '3 SUBSCRIPTIONS' and '9 WORKSPACES'. Below is a table of subscriptions:

NAME	POLICY INITIATIVE ASSIGNMENT(S)	COMPLIANCE	COVERAGE	SETTINGS
ASC DEMO	[Preview]: Enable Monitoring in Azure Seci	50%	Standard	Edit settings >
Contoso IT - demo	[Preview]: Enable Monitoring in Azure Seci	38%	Standard	Edit settings >

3. Under **Compute and apps**, **Network**, and **Data**, set each security configuration you want to monitor to **On**. Security Center will continuously assess the configuration of your environment and when vulnerability exists, Security Center will generate a security recommendation. Select **Off** if the security configuration is not recommended or not

relevant. For example, in a dev/test environment you might not require the same level of security as a production environment. After selecting the policies that are applicable to your environment, click **Save**.

Home > Security Center - Security policy > Security policy

Security policy

Visual Studio Enterprise

Save

Assess the following policies:

Compute And Apps (6 of 6 policies enabled)

System updates ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Security configurations ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Endpoint protection ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Disk encryption ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Vulnerability Assessment ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Adaptive Application Controls ⓘ	<input type="checkbox"/> On <input type="checkbox"/> Off UPGRADE

Network (4 of 4 policies enabled)

Network security groups ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Web application firewall ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
Next generation firewall ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
JIT Network Access ⓘ	<input type="checkbox"/> On <input type="checkbox"/> Off UPGRADE

Data (3 of 3 policies enabled)

Storage Encryption ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
SQL auditing ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off
SQL Encryption ⓘ	<input checked="" type="checkbox"/> On <input type="checkbox"/> Off

Wait until Security Center processes these policies and generates recommendations. Some configurations, such as system updates and OS configurations can take up to 12 hours, while network security groups and encryption configurations can be assessed almost instantly. Once you see recommendations in the Security Center dashboard, you can proceed to the next step.

Setting security policies in Security Center or in Azure Policy - Advanced

Azure Security Center policies integrate with Azure Policies, so you can set them either in Security Center on a specific subscription, or in [Azure Policy](#), which enables you to set policies across Management groups and across multiple subscriptions..

What are security policies?

A security policy defines the desired configuration of your workloads and helps ensure compliance with company or regulatory security requirements. In Azure Security Center, you can define policies for your Azure subscriptions and tailor them to your type of workload or the sensitivity of your data. For example, applications that use regulated data, such as personally identifiable information, might require a higher level of security than other workloads. To set a policy across subscriptions or on Management groups, set them in [Azure Policy](#).

NOTE

If you previously configured security policies on a subscription that is part of a management group, or has multiple policy assignments, those policies appear greyed out in Security Center so that you can manage the policy at the management group level via the Azure Policy page.

How security policies work

Security Center automatically creates a default security policy for each of your Azure subscriptions. You can edit the policies in Security Center or use Azure Policy to do the following things:

- Create new policy definitions.
- Assign policies across management groups and subscriptions, which can represent an entire organization or a business unit within the organization. Monitor policy compliance.
-

For more information about Azure Policy, see [Create and manage policies to enforce compliance](#).

An Azure policy consists of the following components:

- A **policy** is a rule
- An **initiative** is a collection of policies
- An **assignment** is an application of an initiative or a policy to a specific scope (management group, subscription, or resource group)

A resource is evaluated against the policies that are assigned to it and receives a compliance ratio according to the number of policies the resource is compliant to.

Who can edit security policies?

Security Center uses Role-Based Access Control (RBAC), which provides built-in roles that can be assigned to users, groups, and services in Azure. When users open Security Center, they see only information that's related to resources they have access to. Which means that users are assigned the role of owner, contributor, or reader to the subscription or resource group that a resource belongs to. In addition to these roles, there are two specific Security Center roles:

- Security reader: Have view rights to Security Center, which includes recommendations, alerts, policy, and health, but they can't make changes.
- Security admin: Have the same view rights as security reader, and they can also update the security policy and dismiss recommendations and alerts.

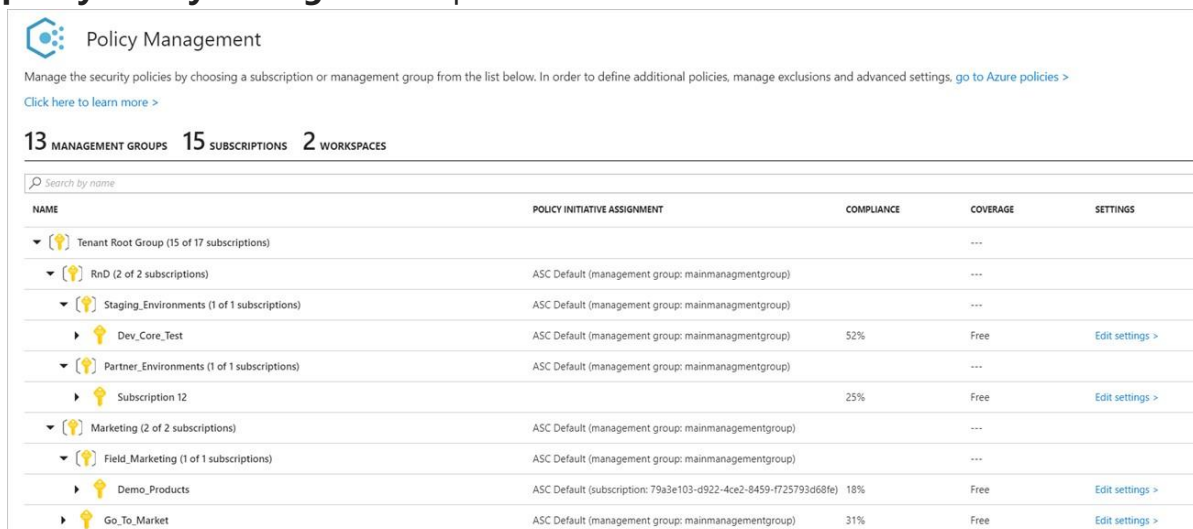
Edit security policies

You can edit the default security policy for each of your Azure subscriptions and management groups in Security Center. To modify a security policy, you must be an owner, contributor, or security administrator of the subscription or the containing management group. To view your security policies in Security Center:

NOTE

Any policies set on a subscription that is part of a management group, or has multiple policy assignments, will appear greyed out in Security Center. You can edit these policies in [Azure Policy](#).

1. On the **Security Center** dashboard, under **POLICY & COMPLIANCE**, select **Security policy**. **Policy Management** opens.



The screenshot shows the 'Policy Management' page in Security Center. It includes a search bar, a table with columns for NAME, POLICY INITIATIVE ASSIGNMENT, COMPLIANCE, COVERAGE, and SETTINGS, and a list of management groups and subscriptions with their respective policy assignments and compliance percentages.

NAME	POLICY INITIATIVE ASSIGNMENT	COMPLIANCE	COVERAGE	SETTINGS
▼ Tenant Root Group (15 of 17 subscriptions)			---	
▼ RnD (2 of 2 subscriptions)	ASC Default (management group: mainmanagementgroup)		---	
▼ Staging_Environments (1 of 1 subscriptions)	ASC Default (management group: mainmanagementgroup)		---	
▶ Dev_Core_Test	ASC Default (management group: mainmanagementgroup)	52%	Free	Edit settings >
▼ Partner_Environments (1 of 1 subscriptions)	ASC Default (management group: mainmanagementgroup)		---	
▶ Subscription 12		25%	Free	Edit settings >
▼ Marketing (2 of 2 subscriptions)	ASC Default (management group: mainmanagementgroup)		---	
▼ Field_Marketing (1 of 1 subscriptions)	ASC Default (management group: mainmanagementgroup)		---	
▶ Demo_Products	ASC Default (subscription: 79a3e103-d922-4ce2-8459-f725793d68fe)	18%	Free	Edit settings >
▶ Go_To_Market	ASC Default (management group: mainmanagementgroup)	31%	Free	Edit settings >

Policy Management displays the number of management groups, subscriptions, and workspaces as well as your management group structure.

NOTE

The Security Center dashboard may show a higher number of subscriptions under **Subscription coverage** than the number of subscriptions shown under **Policy Management**. Subscription coverage shows the number of Standard, Free, and "not covered" subscriptions. The "not covered" subscriptions do not have Security Center enabled and are not displayed under **Policy Management**.

The columns in the table display:

- Policy Initiative Assignment – Security Center built-in policies and initiatives that are assigned to a subscription or management group.
- Compliance – Overall compliance score for a management group, subscription, or workspace. The score is the weighted average of the assignments. The weighted average factors in the number of policies in a single assignment and the number of resources the assignment applies to.

For example, if your subscription has two VMs and an initiative with five policies assigned to it, then you have 10 assessments in your subscription. If one of the VMs doesn't comply to two of the policies, then the overall compliance score of your subscription's assignment is 80%.

- Coverage – Identifies the pricing tier, Free or Standard, that the management group, subscription, or workspace is running on. See [Pricing](#) to learn more about Security Center's pricing tiers.

Settings – Subscriptions have the link **Edit settings**. Selecting **Edit settings** lets you update your subscription settings such as data collection, pricing tier, and email notifications.

2. Select the subscription or management group that you want to enable a security policy for. **Security policy** opens.
3. Under **Security policy**, select the controls that you want Security Center to monitor for and provide recommendations on by selecting **On**. Select **Off** if you don't want Security Center to monitor that control.
4. Select **Save**.

Management groups

If your organization has many subscriptions, you may need a way to efficiently manage access, policies, and compliance for those subscriptions. Azure Management Groups provides a level of scope above subscriptions. You organize subscriptions into containers called "management groups" and apply your governance policies to the management groups. All subscriptions within a management group automatically inherit the policies applied to the

management group. Each directory is given a single top-level management group called the "root" management group. This root management group is built into the hierarchy to have all management groups and subscriptions fold up to it. This root management group allows for global policies and RBAC assignments to be applied at the directory level. To set up management groups for use with Azure Security Center, follow the instructions in the [Gain tenant-wide visibility for Azure Security Center](#) article.

NOTE

It's important that you understand the hierarchy of management groups and subscriptions. See [Organize your resources with Azure Management Groups](#) to learn more about management groups, root management, and management group access.

Customize OS security configurations in Azure Security Center (Preview) – Red Subscription

What are OS security configurations?

Azure Security Center monitors security configurations by applying a set of [over 150 recommended rules](#) for hardening the OS, including rules related to firewalls, auditing, password policies, and more. If a machine is found to have a vulnerable configuration, Security Center generates a security recommendation.

By customizing the rules, organizations can control which configuration options are more appropriate for their environment. You can set a customized assessment policy and then apply it on all applicable machines in the subscription.

NOTE

- Currently, customization of the OS security configuration is available for Windows Server versions 2008, 2008 R2, 2012, and 2012 R2 operating systems only.
- The configuration applies to all VMs and computers that are connected to all
- workspaces under the selected subscription.

OS security configuration customization is available only on the Security Center standard tier.

You can customize the OS security configuration rules by enabling and disabling a specific rule, changing the desired setting for an existing rule, or adding a new rule that's based on the supported rule types (registry, audit policy, and security policy). Currently, the desired setting must be an exact value.

New rules must be in the same format and structure as other existing rules of the same type.

NOTE

To customize OS security configurations, you must be assigned the role of *Subscription Owner*, *Subscription Contributor*, or *Security Administrator*.

Customize the default OS security configuration

To customize the default OS security configuration in Security Center, do the following:

1. Open the **Security Center** dashboard.
2. In the left pane, select **Security policy**.

The screenshot shows the 'Security Center - Security policy' interface. The left-hand navigation pane is visible, with 'Security policy' selected and highlighted in a red box. The main content area is titled 'Policy Management' and contains a table of subscriptions. The table has columns for 'NAME', 'POLICY INITIATIVE ASSIGNMENT(S)', 'COMPLIANCE', 'COVERAGE', and 'SETTINGS'. A single row is visible for 'contoso-1' with a 'Settings' column containing an 'Edit settings >' link, which is also highlighted in a red box.

NAME	POLICY INITIATIVE ASSIGNMENT(S)	COMPLIANCE	COVERAGE	SETTINGS
contoso-1	[Preview]: Enable Monitoring in Azure Secu	50%	Standard	Edit settings >

3. In the row of the subscription you want to customize, click **Edit settings**.

4. Select **Edit security configurations**.

The screenshot shows the 'Settings - Edit security configurations' page. The left-hand navigation pane is visible, with 'Edit security configurations' selected. The main content area contains instructions for downloading, editing, and uploading the OS Security Configuration file. The instructions are as follows:

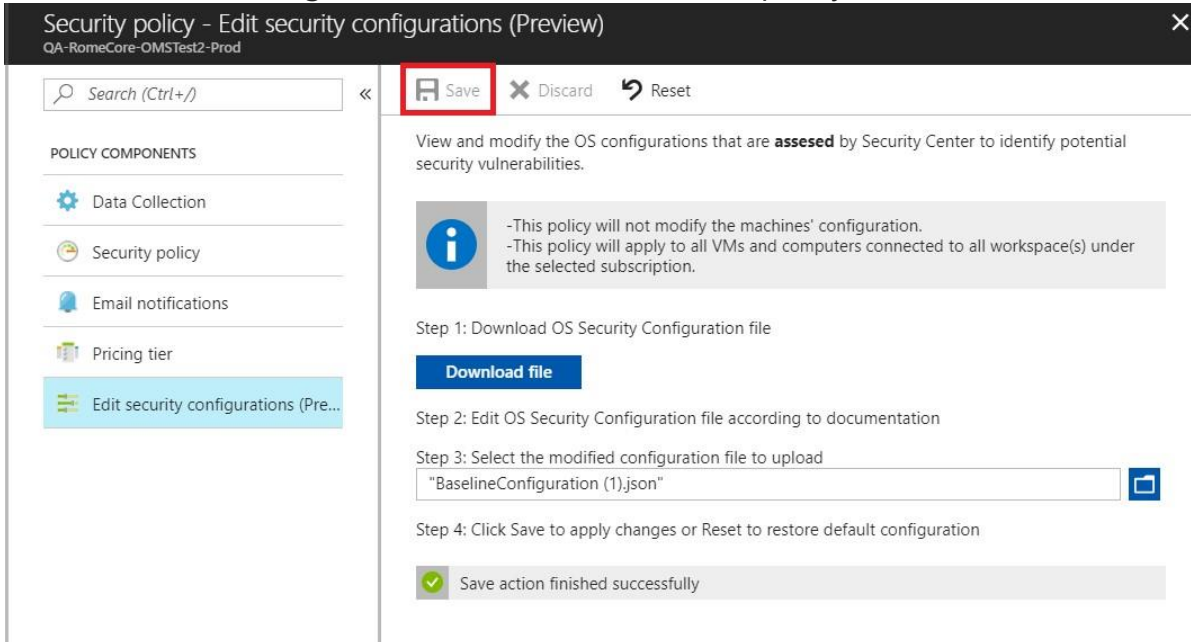
- Step 1: Download OS Security Configuration file
- Step 2: Edit OS Security Configuration file according to documentation
- Step 3: Select the modified configuration file to upload
- Step 4: Click Save to apply changes or Reset to restore default configuration

5. Follow the steps to download, edit, and upload the modified file.

NOTE

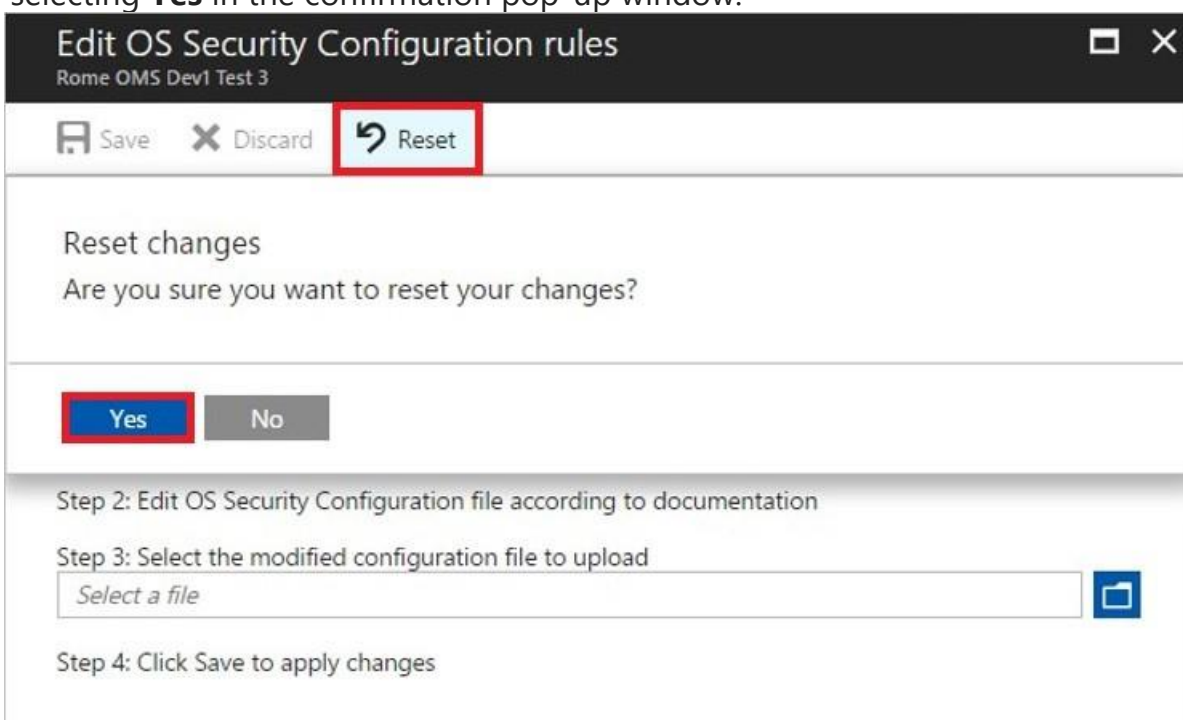
By default, the configuration file that you download is in *json* format. For instructions about modifying this file, go to [Customize the configuration file](#).

6. To commit the change, select **Save**. Otherwise, the policy is not stored.



After you've successfully saved the file, the configuration is applied to all VMs and computers that are connected to the workspaces under the subscription. The process usually takes a few minutes but can take longer, depending on the infrastructure size.

At any point, you can reset the current policy configuration to its default state. To do so, in the **Edit OS security configuration rules** window, select **Reset**. Confirm this option by selecting **Yes** in the confirmation pop-up window.



Customize the configuration file

In the customization file, each supported OS version has a set of rules, or ruleset. Each ruleset has its own name and unique ID, as shown in the following example:

```
[{
  "rules": ...,
  "baselineId": "b65666d9-2769-4930-8929-95f39a08d1db",
  "baselineName": "WS2008SP2 Member Server Security Compliance"
},
{
  "rules": ...,
  "baselineId": "9ff96717-0c7f-4ed0-a7a0-22bdd9c7a75b",
  "baselineName": "WS2008R2SP1 Member Server Security Compliance"
},
{
  "rules": ...,
  "baselineId": "51e9a436-6790-4ea7-9ddd-c07b789fe228",
  "baselineName": "WS2012 Member Server Security Compliance"
},
{
  "rules": ...,
  "baselineId": "6ba1ce80-e4e5-4b8a-bc88-257612e72185",
  "baselineName": "WS2012R2 Member Server Security Compliance"
}]
```

NOTE

This example file was edited in Visual Studio, but you can also use Notepad if you have the JSON Viewer plug-in installed.

When you edit the customization file, you can modify one rule or all of them. Each ruleset includes a *rules* section that's separated into three categories: Registry, Audit Policy, and Security Policy, as shown here:

```
[{
  "rules": {
    "baselineRegistryRules": ...,
    "baselineAuditPolicyRules": [],
    "baselineSecurityPolicyRules": ...
  },
  "baselineId": "b65666d9-2769-4930-8929-95f39a08d1db",
  "baselineName": "WS2008SP2 Member Server Security Compliance"
},
```

Each category has its own set of attributes. You can change the following attributes:

- **expectedValue:** This attribute's field data type must match the supported values per *rule type*, for example:
 - **baselineRegistryRules:** The value should match the [regValueType](#) that's defined in that rule.
 - **baselineAuditPolicyRules:** Use one of the following string values:
 - *Success and Failure*
 - *Success*
 - **baselineSecurityPolicyRules:** Use one of the following string values:
 - *No one*
 - List of allowed user groups, for example: *Administrators, Backup Operators*

The string can contain the options *Disabled* or *Enabled*. For this private preview release, the string is case-sensitive.

These are the only fields that can be configured. If you violate the file format or size, you won't be able to save the change. You will receive an error telling you that you need to upload a valid JSON configuration file.

For a list of other potential errors, see [Error codes](#).

The following three sections contain examples of the preceding rules. The *expectedValue* and *state* attributes can be changed. **baselineRegistryRules**

```

{
  "hive": "LocalMachine",
  "regValueType": "Int",
  "keyPath":
    "System\\\\"CurrentControlSet\\\\"Services\\\\"LanManServer\\\\"Parameters",
  "valueName": "restrictnullsessaccess",
  "ruleId": "f9020046-6340-451d-9548-3c45d765d06d",
  "originalId": "0f319931-aa36-4313-9320-86311c0fa623",
  "cceId": "CCE-10940-5",
  "ruleName": "Network access: Restrict anonymous access to Named
Pipes and Shares",
  "ruleType": "Registry",
  "expectedValue": "1",
  "severity": "Warning",
  "analyzeOperation": "Equals",
  "source": "Microsoft",
  "state": "Disabled"
}

```

baselineAuditPolicyRules

```
{
  "auditPolicyId": "0cce923a-69ae-11d9-bed3-505054503030",
  "ruleId": "37745508-95fb-44ec-ab0f-644ec0b16995",
  "originalId": "2ea0de1a-c71d-46c8-8350-a7dd4d447895",
  "cceId": "CCE-11001-5",
  "ruleName": "Audit Policy: Account Management: Other Account Management Events",
  "ruleType": "AuditPolicy",
  "expectedValue": "Success and Failure",
  "severity": "Critical",
  "analyzeOperation": "Equals",
  "source": "Microsoft",
  "state": "Enabled"
}
```

baselineSecurityPolicyRules

```
{
  "sectionName": "Privilege Rights",
  "settingName": "SeIncreaseWorkingSetPrivilege",
  "ruleId": "b0ec9d5e-916f-4356-83aa-c23522102b33",
  "originalId": "b61bd492-74b0-40f3-909d-36b9bf54e94c",
  "cceId": "CCE-10548-6",
  "ruleName": "Increase a process working set",
  "ruleType": "SecurityPolicy",
  "expectedValue": "Administrators, Local Service",
  "severity": "Warning",
  "analyzeOperation": "Equals",
  "source": "Microsoft",
  "state": "Enabled"
}
```

Some rules are duplicated for the different OS types. Duplicate rules have the same *originalId* attribute.

Review the allowed Expected Values for the settings above and try to make a change and upload the file back into your Azure Subscription.

Reviewing security recommendations in Azure Security Center – Blue Subscription

This scenario walks you through how to use recommendations in Azure Security Center to help you protect your Azure resources.

NOTE

This document introduces the service by using an example deployment. This document is not a step-by-step guide.

What are security recommendations?

Security Center periodically analyzes the security state of your Azure resources. When Security Center identifies potential security vulnerabilities, it creates recommendations. The recommendations guide you through the process of configuring the needed controls.

Implementing security recommendations

Set recommendations

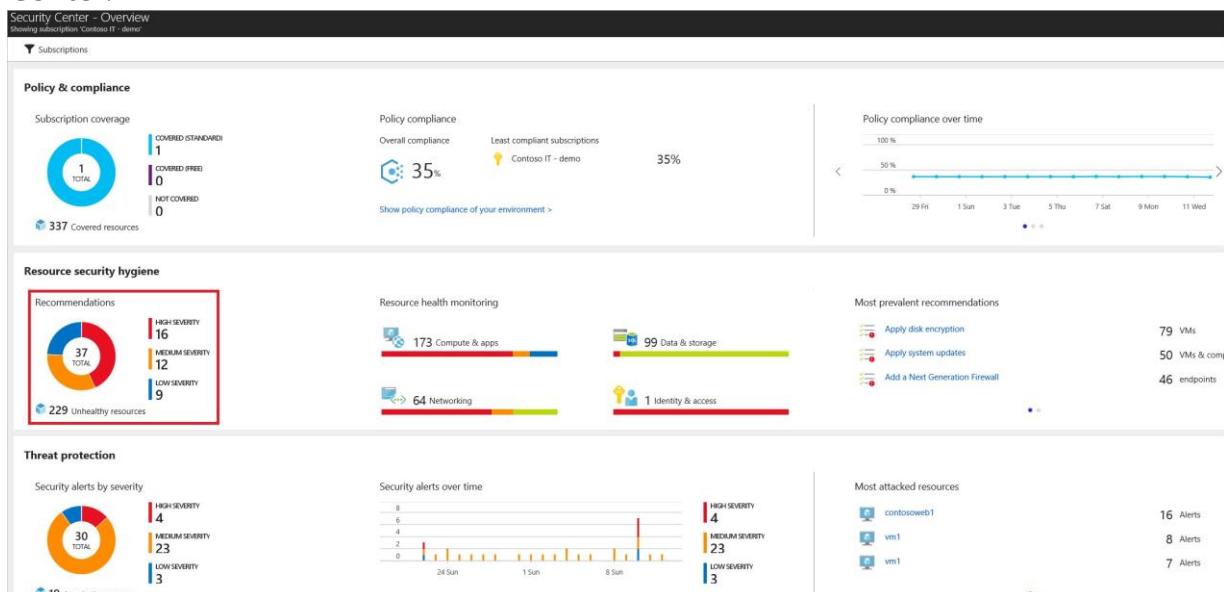
In [Setting security policies in Azure Security Center](#), you learn to:

- Configure security policies.
- Turn on data collection.
- Choose which recommendations to see as part of your security policy.

Current policy recommendations center around system updates, baseline rules, anti-malware programs, [network security groups](#) on subnets and network interfaces, SQL database auditing, SQL database transparent data encryption, and web application firewalls. [Setting security policies](#) provides a description of each recommendation option.

Monitor recommendations

After setting a security policy, Security Center analyzes the security state of your resources to identify potential vulnerabilities. The **Recommendations** tile under **Overview** lets you know the total number of recommendations identified by Security Center.



To see the details of each recommendation, select the **Recommendations tile** under **Overview. Recommendations** opens.

DESCRIPTION	RESOURCE	STATE	SEVERITY
Install Endpoint Protection	5 virtual mac...	Open	High
Add a web application firewall	2 web applic...	Open	High
Add a Next Generation Firewall	2 endpoints	Open	High
Enable Network Security Groups on subn..	2 subnets	Open	High
Enable Network Security Groups on virtu..	2 virtual mac...	Open	High
Enable Transparent Data Encryption	2 SQL databa..	Open	High
Apply disk encryption	5 virtual mac...	Open	High
Reboot after system updates	2 virtual mac...	Open	Medium
Provide security contact details	1 subscriptions	Open	Medium

You can filter recommendations. To filter the recommendations, select **Filter** on the **Recommendations** blade. The **Filter** blade opens and you select the severity and state values you wish to see.

The recommendations are shown in a table format where each line represents one particular recommendation. The columns of this table are:

- **DESCRIPTION:** Explains the recommendation and what needs to be done to address it.
- **RESOURCE:** Lists the resources to which this recommendation applies.
- **STATE:** Describes the current state of the recommendation:
 - **Open:** The recommendation hasn't been addressed yet.
 - **In Progress:** The recommendation is currently being applied to the resources, and no action is required by you.
 - **Resolved:** The recommendation has already been completed (in this case, the line is grayed out).
- **SEVERITY:** Describes the severity of that particular recommendation:
 - **High:** A vulnerability exists with a meaningful resource (such as an application, a VM, or a network security group) and requires attention.
 - **Medium:** A vulnerability exists and non-critical or additional steps are required to eliminate it or to complete a process.

Low: A vulnerability exists that should be addressed but does not require immediate attention. (By default, low recommendations aren't presented, but you can filter on low recommendations if you want to see them.)

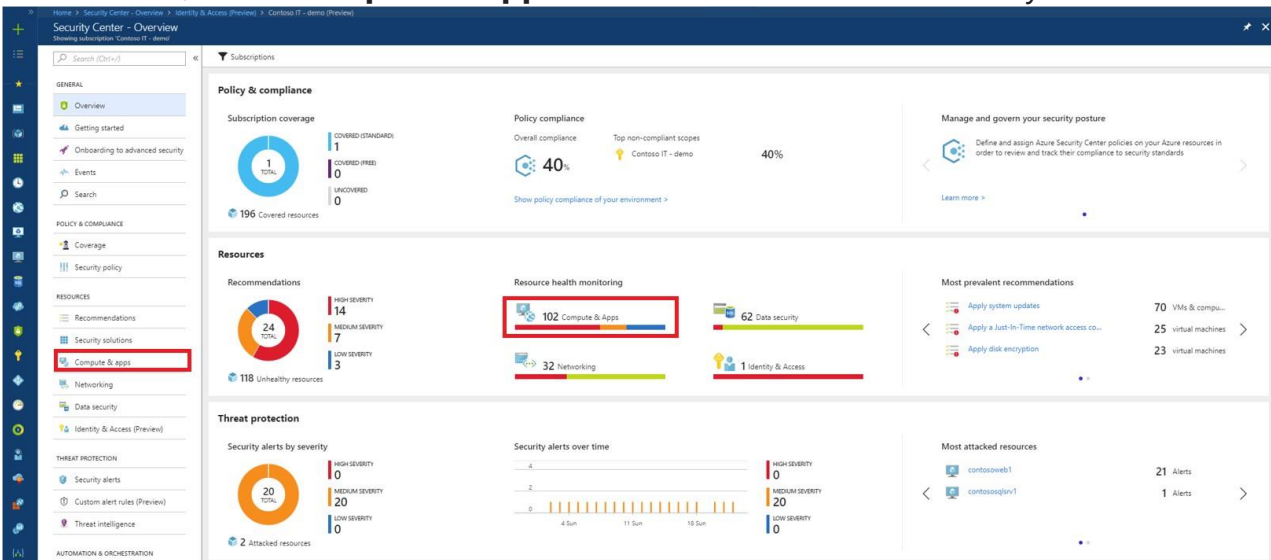
Monitoring security health

You can monitor the security state of your resources on the **Security Center – Overview** dashboard. The **Resources** section provides the number of issues identified and the security state for each resource type.

You can view a list of all issues by selecting **Recommendations**. For more information about how to apply recommendations, see [Implementing security recommendations in Azure Security Center](#).

For a complete list of Compute and App services recommendations, see [Recommendations](#).

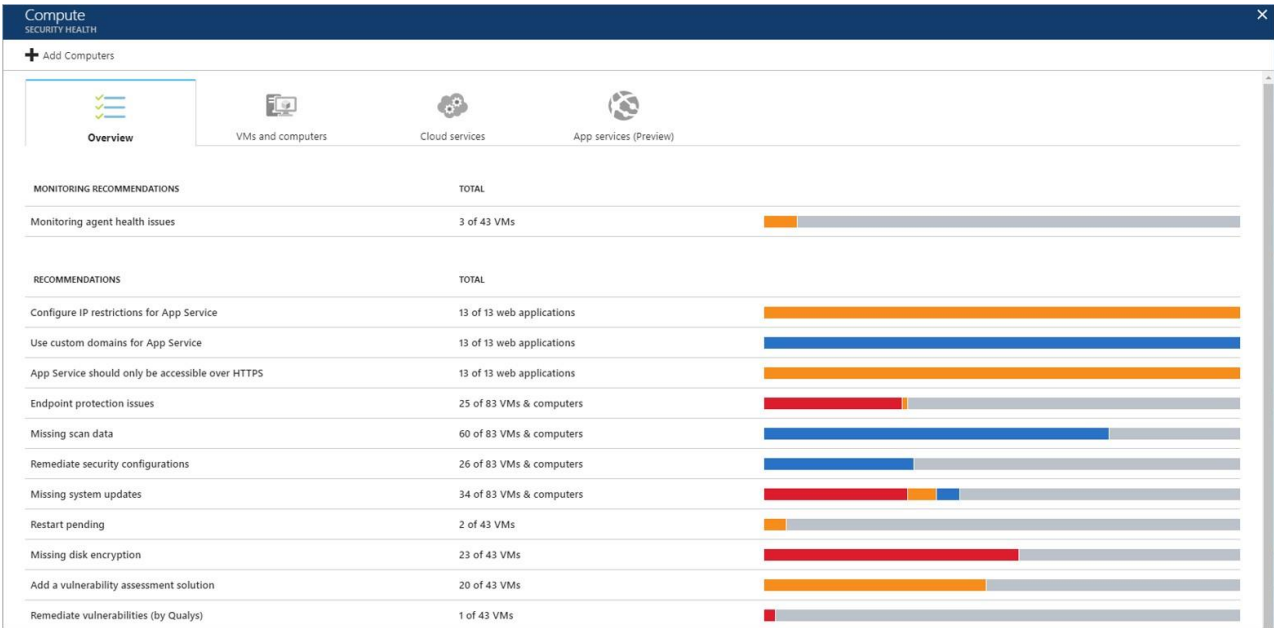
To continue, select **Compute & apps** under **Resources** or the Security Center main menu.



Monitor Compute and App services

Under **Compute**, there are four tabs:

- **Overview:** monitoring and recommendations identified by Security Center.
- **VMs and computers:** list of your VMs, computers, and current security state of each.
- **Cloud Services:** list of your web and worker roles monitored by Security Center.
- **App services (Preview):** list of your App service environments and current security state of each. To continue, select **Compute & apps** under **Resources** or the Security Center main menu.



In each tab you can have multiple sections, and in each section, you can select an individual option to see more details about the recommended steps to address that particular issue.

Monitoring recommendations

This section shows the total number of VMs and computers that were initialized for automatic provisioning and their current statuses. In this example there is one recommendation, **Monitoring agent health issues**. Select this recommendation.

The screenshot shows the detailed view of the 'Monitoring agent health issues' recommendation. It is split into two panes. The left pane provides a summary of the issue and a list of affected VMs. The right pane shows the specific details for the selected VM, 'ContosoWebBE3'.

Monitoring agent health issues

Security Center is unable to successfully monitor the following VMs, due to different reasons. To resolve this please follow the detailed instructions in the documentation.

[Learn how to resolve monitoring issues >](#)

NAME	INSTALLATION STATUS
ContosoWebBE3	Installation failed - local agent already installed
ContosoWebDC	Installation failed - local agent already installed
kncon	Agent not responsive or missing ID

ContosoWebBE3
Virtual machine security health

Filter

Virtual machine info

- VIRTUAL MACHINE: ContosoWebBE3
- RESOURCE GROUP: CONTOSORETAIL100
- WORKSPACE: contosoretail-it
- SUBSCRIPTION: Contoso IT - demo
- VIRTUAL IP: 104.215.90.235
- OPERATING SYSTEM: Windows
- VERSION: Compute
- STATUS: Running
- MONITORING STATE: No data collection : Installation failed - local agent already installed
- PREVENTION STATUS: High severity

Security Solutions

- SYSTEM UPDATES: Microsoft (Last scan time - Not applicable)
- SECURITY CONFIGURATIONS: Microsoft (Last scan time - Not applicable)

Recommendations

DESCRIPTION	SEVERITY
Endpoint Protection not installed on Azure VMs	High
Add a vulnerability assessment solution	Medium

Monitoring agent health issues opens. VMs and computers that Security Center is unable to successfully monitor are listed. Select a VM or computer for detailed information.

MONITORING STATE provides a reason

why Security Center is unable to monitor. See the [Security Center troubleshooting guide](#) for a list of **MONITORING STATE** values, descriptions, and resolution steps.

Unmonitored VMs and computers

A VM or computer is unmonitored by Security Center if the machine is not running the Microsoft Monitoring

Agent extension. A machine may have a local agent already installed, for example the OMS direct agent or the SCOM agent. Machines with these agents are identified as unmonitored because these agents are not fully supported in Security Center. To fully benefit from all of Security Center's capabilities, the Microsoft Monitoring Agent extension is required.

You can install the extension on the unmonitored VM or computer in addition to the already installed local agent. Configure both agents the same, connecting them to the same workspace. This enables Security Center to interact with the Microsoft Monitoring Agent extension and collect data. See [Enable the VM extension](#) for instructions on how to install the Microsoft Monitoring Agent extension.

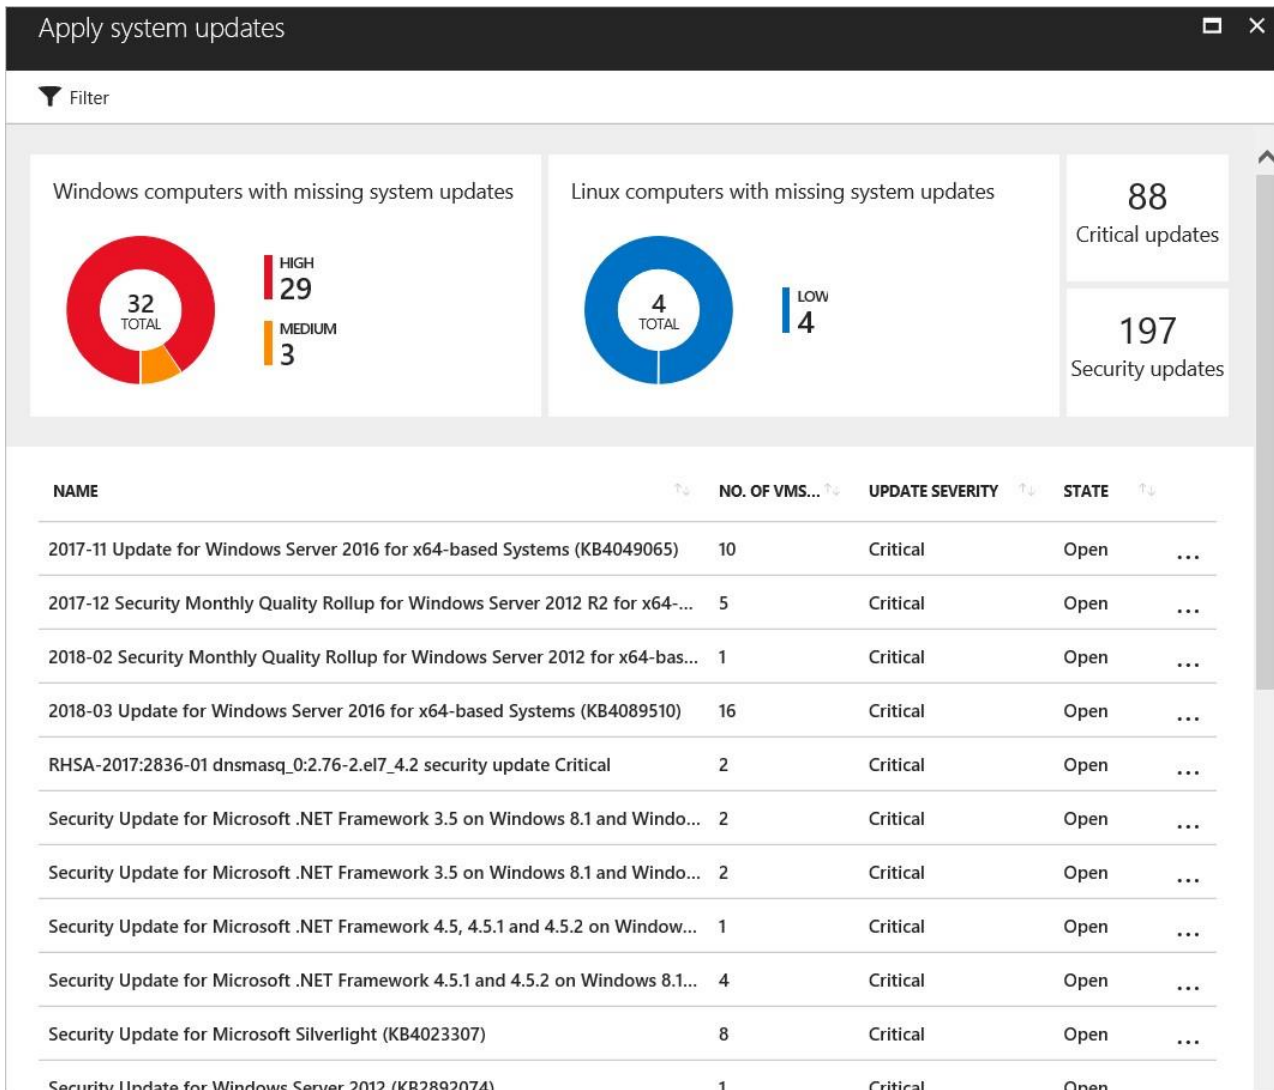
See [Monitoring agent health issues](#) to learn more about the reasons Security Center is unable to successfully monitor VMs and computers initialized for automatic provisioning.

Recommendations

This section has a set of recommendations for each VM and computer, web and worker roles, Azure App Service Web Apps, and Azure App Service Environment that Security Center monitors. The first column lists the recommendation. The second column shows the total number of resources that are affected by that recommendation. The third column shows the severity of the issue as illustrated in the following screenshot:

RECOMMENDATIONS	TOTAL	
App Service should only be accessible over HTTPS	13 of 13 web applications	
Use custom domains for App Service	13 of 13 web applications	
Configure IP restrictions for App Service	13 of 13 web applications	
Endpoint protection issues	26 of 81 VMs & computers	
Missing scan data	54 of 81 VMs & computers	
Remediate security configurations	21 of 81 VMs & computers	
Missing system updates	36 of 81 VMs & computers	
Missing disk encryption	23 of 47 VMs	
Add a vulnerability assessment solution	23 of 47 VMs	
Remediate vulnerabilities (by Qualys)	1 of 47 VMs	
Add web application firewall	9 of 9 web applications	

Each recommendation has a set of actions that you can perform after you select it. For example, if you select **Missing system updates**, the number of VMs and computers that are missing patches, and the severity of the missing update appears, as shown in the following screenshot:



Apply system updates has a summary of critical updates in a graph format, one for Windows, and one for Linux. The second part has a table with the following information:

- **NAME:** Name of the missing update.
- **NO. OF VMs & COMPUTERS:** Total number of VMs and computers that are missing this update.
- **UPDATE SEVERITY:** Describes the severity of that particular recommendation:
 - **Critical:** A vulnerability exists with a meaningful resource (application, virtual machine, or network security group) and requires attention.
 - **Important:** Non-critical or additional steps are required to complete a process or eliminate a vulnerability.

Moderate: A vulnerability should be addressed but does not require immediate attention. (By default, low recommendations are not presented, but you can filter on low recommendations if you want to view them.)

- **STATE:** The current state of the recommendation:
 - **Open:** The recommendation has not been addressed yet.
 - **In Progress:** The recommendation is currently being applied to those resources, and no action is required by you.
 - **Resolved:** The recommendation was already finished. (When the issue has been resolved, the entry is dimmed).

To view the recommendation details, click the name of the missing update from the list.



NOTE

The security recommendations here are the same as those under the **Recommendations** tile. See [Implementing security recommendations in Azure Security Center](#) for more information about how to resolve recommendations.

VMs and computers

The VMs and computers section gives you an overview of all VM and computer recommendations. Each column represents one set of recommendations as shown in the following screenshot:



































































Compute
SECURITY HEALTH

+ Add Computers Filter

Overview VMs and computers Cloud services App services (Preview)

Filtered By: Power State: Running

Search by name

NAME	MONITORED	SYSTEM UPDATES	ENDPOINT PROTECTION	VULNERABILITIES	DISK ENCRYPTION
 ContosoWebBE3					
 kncon					
 CDMNEBALTVM0397.smx.net					
 CDMNEBALTVM0406.smx.net					
 DEMOVM4.redmond.corp.microsoft.com					
 DemoVM7.redmond.corp.microsoft.com					
 DemoVM5.redmond.corp.microsoft.com					
 infoweb02.contoso.com					
 OpsInsights02.contoso.com					
 ContosoAzADDs1					
 ContosoAzADDs2					

There are four types of icons represented in this list:



Non-Azure computer.



Azure Resource Manager VM.



Azure Classic VM.



VMs that are identified only from the workspace that is part of the viewed subscription. This includes VMs from other subscriptions that report to the workspace in this subscription, and VMs that were installed with SCOM direct agent, and have no resource ID.

The icon that appears under each recommendation helps you to quickly identify the VM and computer that needs attention, and the type of recommendation. You can also use the Filter option to select which options you will see on this screen.

Filter ×

Environment

- Azure
- Non-Azure

Power State

- Running
- Stopped

Monitored

- High Severity
- Medium Severity
- Low Severity
- Healthy
- N/A

System updates

- High Severity
- Medium Severity
- Low Severity
- Healthy
- N/A

Endpoint Protection

- High Severity
- Medium Severity
- Low Severity
- Healthy
- N/A

Vulnerabilities

- High Severity
- Medium Severity
- Low Severity
- Healthy
- N/A

Disk encryption

- High Severity
- Medium Severity
- Low Severity
- Healthy
- N/A

In the previous example, one VM has a critical recommendation regarding endpoint protection. Select the VM to get more information about it:

ContosoWebBE3
☐ ✕

Virtual machine security health

▼ Filter

Virtual machine info

VIRTUAL MACHINE	ContosoWebBE3
RESOURCE GROUP	CONTOSORETAIL100
WORKSPACE	contosoretail-it
SUBSCRIPTION	Contoso IT - demo
VIRTUAL IP	104.215.90.235
OPERATING SYSTEM	Windows
VERSION	Compute
STATUS	Running
MONITORING STATE	No data collection : Installation failed - local agent already installed
PREVENTION STATUS	High severity

Security Solutions

SYSTEM UPDATES	Microsoft (Last scan time - Not applicable)
SECURITY CONFIGURATIONS	Microsoft (Last scan time - Not applicable)

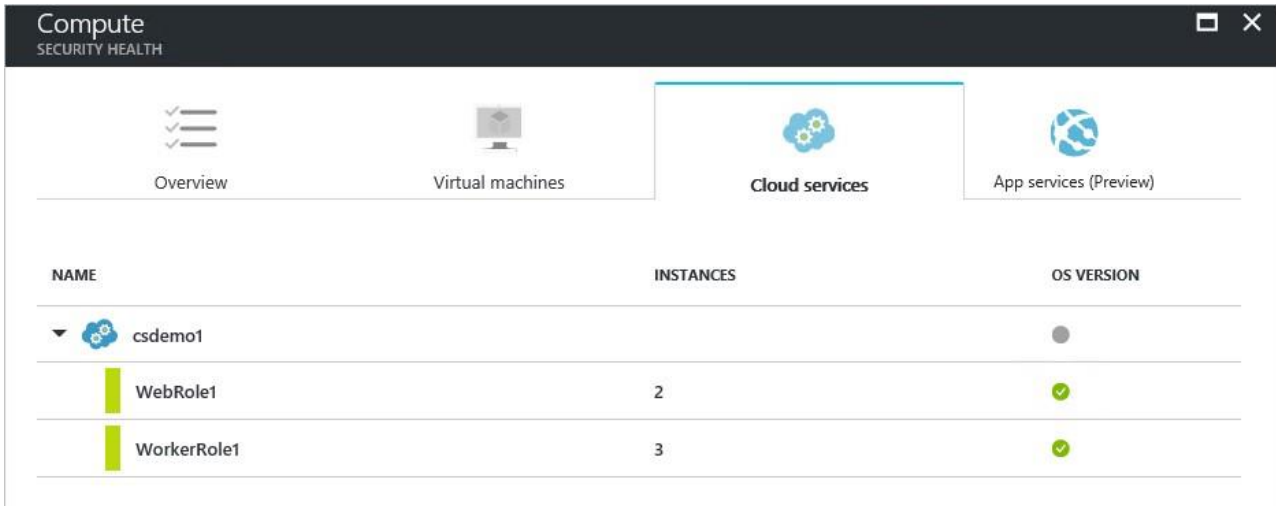
Recommendations

DESCRIPTION	SEVERITY
Endpoint Protection not installed on Azure VMs	🔴 High
Add a vulnerability assessment solution	🟡 Medium

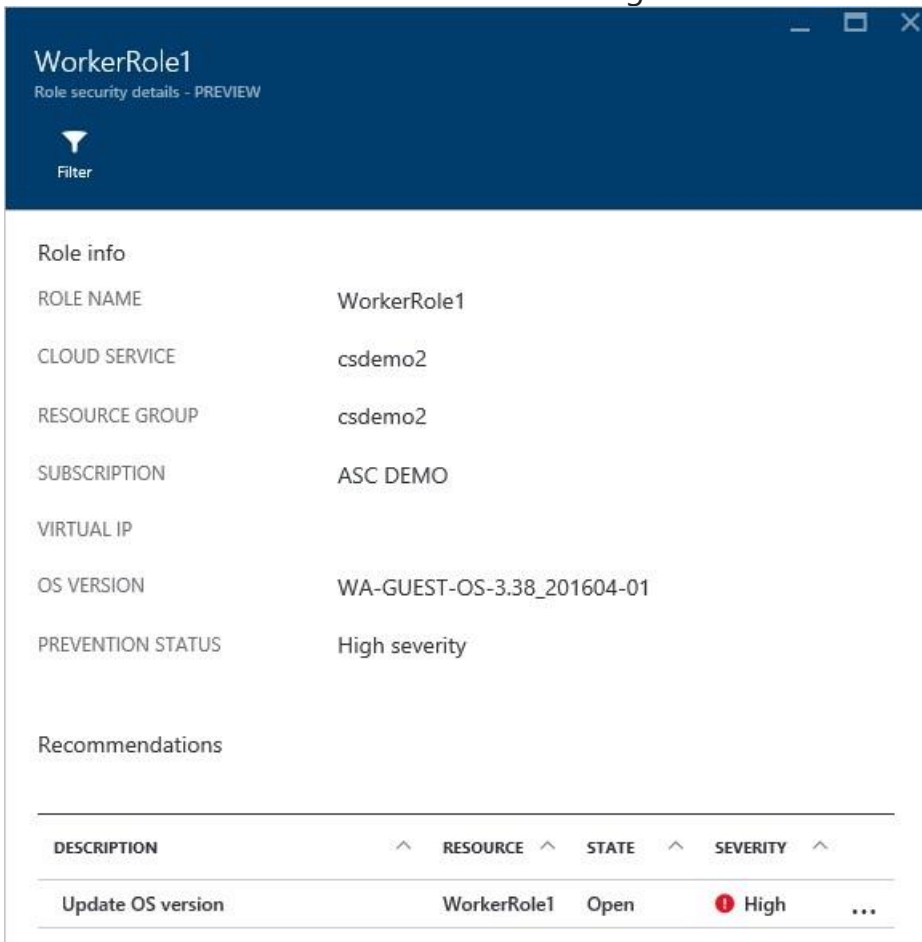
Here you see the security details for the VM or computer. At the bottom you can see the recommended action and the severity of each issue.

Cloud services

For cloud services, a recommendation is created when the operating system version is out of date as shown in the following screenshot:



In a scenario where you do have a recommendation (which is not the case for the previous example), you need to follow the steps in the recommendation to update the operating system version. When an update is available, you will have an alert (red or orange - depends on the severity of the issue). When you select this alert in the WebRole1 (runs Windows Server with your web app automatically deployed to IIS) or WorkerRole1 (runs Windows Server with your web app automatically deployed to IIS) rows, you see more details about this recommendation as shown in the following screenshot:



To see a more prescriptive explanation about this recommendation, click **Update OS version** under the **DESCRIPTION** column.

Update OS version (Preview)
PREVIEW

Filter

DESCRIPTION

Please update the operating system (OS) version for your Cloud Service to the most recent version available for your OS family. You can do this manually by editing the service config file, or by choosing the "Automatic" option for OPERATING SYSTEM VERSION in the Azure Classic Portal to enable automatic updates to the current OS version. In the current (not classic) Azure portal, you will need to update the service config file (cscfg). Changes to the OS version for a Cloud Service are applied to all role instances in the service.

MORE INFORMATION

<https://azure.microsoft.com/en-us/documentation/articles/cloud-services-how-to-configure>

ROLE	CLOUD SERVICE	OS VERSION	STATE	SEVERITY	
WorkerRole1	csdemo2	WA-GUEST-OS-3...	Open	High	...

App services (Preview)

NOTE

Monitoring App Service is in preview and available only on the Standard tier of Security Center. See [Pricing](#) to learn more about Security Center's pricing tiers.

Under **App services**, you find a list of your App service environments and the health summary based on the assessment Security Center performed.

Compute
SECURITY HEALTH

+ Add Computers

Overview VMs and computers Cloud services **App services (Preview)**

Search resources

NAME	TOTAL	Health Summary
Appservice-demo	1 of 1 recommendations	High (Red)
webapp-demo	6 of 6 recommendations	Medium (Orange)
ContosoRetailDataAccessAPI	6 of 6 recommendations	Medium (Orange)
contosoretail-usage-generator	6 of 6 recommendations	Medium (Orange)
contosoretailweb	6 of 6 recommendations	Medium (Orange)
contosobotservice	6 of 6 recommendations	Medium (Orange)
contosoit2	6 of 6 recommendations	Medium (Orange)
contosoit	6 of 6 recommendations	Medium (Orange)
contosoitnode	6 of 6 recommendations	Medium (Orange)
contosolinux	6 of 6 recommendations	Medium (Orange)
AdWebAppContoso1	6 of 6 recommendations	Medium (Orange)
AdWebAppContoso2	6 of 6 recommendations	Medium (Orange)
AdWebAppContoso3	6 of 6 recommendations	Medium (Orange)
ContosoStore001	6 of 6 recommendations	Medium (Orange)

There are three types of icons represented in this list:



App services environment.



Web application.



Function application.

1. Select a web application. A summary view opens with three tabs:

- **Recommendations:** based on assessments performed by Security Center that failed.
- **Passed assessments:** list of assessments performed by Security Center that passed.
- **Unavailable assessments:** list of assessments that failed to run due to an error or the recommendation is not relevant for the specific App service

Under **Recommendations** is a list of the recommendations for the selected web application and severity of each recommendation.

webapp-demo (Preview)
Web application security health

Details

NAME webapp-demo

RESOURCE GROUP Appservice-demo

SUBSCRIPTION Contoso IT - demo

Recommendations

4 4

[Recommendations](#)
[Passed assessments](#)
[Unavailable assessments](#)

DESCRIPTION	STATUS
Configure IP restrictions for App Service (Preview)	Medium
App Service should only be accessible over HTTPS (Preview)	Medium
Use the latest supported PHP version for App Service on Windows hosts (Preview)	Low
Use custom domains for App Service (Preview)	Low

2. Select a recommendation for a description of the recommendation and a list of unhealthy resources, healthy resources, and unscanned resources.

Configure IP restrictions for App Service (Preview)

^ Description
IP Restrictions allow you to define a list of IP addresses that are allowed to access your app. Use of IP Restrictions protects a web application from common attacks.

Remediation steps
To configure IP restrictions, we recommend the following steps:
1. Go to the app service networking page
2. Under the IP restrictions section select Add rule
3. Add all IP addresses that are allowed access to your app

UNHEALTHY RESOURCES: 8 HEALTHY RESOURCES: 0

LEARN MORE
[Learn more about recommendations](#)

Unhealthy resources Healthy resources Unscanned resources

Search web applications

NAME	SUBSCRIPTION	RESOURCE GROUP
webapp-demo	Contoso IT - demo	Appservice-demo
ContosoRetailDataAccessAPI	Contoso IT - demo	ContosoAzureHQ
contosoretailweb	Contoso IT - demo	ContosoAzureHQ
contosolinux	Contoso IT - demo	contosolinux90f
AdWebAppContoso1	Contoso IT - demo	ContosoRetailDev
AdWebAppContoso2	Contoso IT - demo	ContosoRetailDev
AdWebAppContoso3	Contoso IT - demo	ContosoRetailDev
ContosoStore001	Contoso IT - demo	ContosoRetailStore001

Under **Passed assessments** is a list of passed assessments. Severity of these assessments is always green.

webapp-demo (Preview)
Web application security health

Details

NAME: [webapp-demo](#)

RESOURCE GROUP: Appservice-demo

SUBSCRIPTION: Contoso IT - demo

Recommendations

4 4

Recommendations Passed assessments Unavailable assessments

DESCRIPTION	STATUS
Web Sockets should be disabled for App Service (Preview)	Healthy
Use the latest supported .NET Framework for App Service on Windows hosts (Preview)	Healthy
Remote debugging should be turned off for App Service (Preview)	Healthy
CORS should not allow every resource access to your application (Preview)	Healthy

3. Select a passed assessment from the list for a description of the assessment, a list of unhealthy and healthy resources, and a list of unscanned resources. There is a tab for unhealthy resources but that list is always empty since the assessment passed.

Web Sockets should be disabled for App Service (Preview)

^ Description

The Web Sockets protocol is vulnerable to different types of security threats. Use of Web Sockets within web applications must be carefully reviewed.

Remediation steps

To disable web sockets protocol, we recommend the following steps:

1. Go to the app service applications settings page
2. In the web sockets toggle select Off
3. Click Save

UNHEALTHY RESOURCES HEALTHY RESOURCES LEARN MORE
[Learn more about recommendations](#)

0 8

[Unhealthy resources](#) [Healthy resources](#) [Unscanned resources](#)

Search web applications

NAME	SUBSCRIPTION	RESOURCE GROUP
webapp-demo	Contoso IT - demo	Appservice-demo
ContosoRetailDataAccessAPI	Contoso IT - demo	ContosoAzureHQ
contosoretailweb	Contoso IT - demo	ContosoAzureHQ
contosolinux	Contoso IT - demo	contosolinux90f
AdWebAppContoso1	Contoso IT - demo	ContosoRetailDev
AdWebAppContoso2	Contoso IT - demo	ContosoRetailDev
AdWebAppContoso3	Contoso IT - demo	ContosoRetailDev
ContosoStore001	Contoso IT - demo	ContosoRetailStore001

LUNCH

Break

Azure Security Center Protect Scenario – Red Subscription

Security Center limits your exposure to threats by using access and application controls to block malicious activity. Just in time virtual machine (VM) access reduces your exposure to attacks by enabling you to deny persistent access to VMs. Instead, you provide controlled and audited access to VMs only when needed. Adaptive application controls help harden VMs against malware by controlling which applications can run on your VMs. Security Center uses machine learning to analyze the processes running in the VM and helps you apply whitelisting rules using this intelligence.

In this scenario you learn how to:

- Configure a just in time VM access policy
- Configure an application control policy

Manage VM access

Just in time virtual machine (VM) access can be used to lock down inbound traffic to your Azure VMs, reducing exposure to attacks while providing easy access to connect to VMs when needed.

NOTE

The just in time feature is available on the Standard tier of Security Center. See [Pricing](#) to learn more about Security Center's pricing tiers.

Attack scenario

Brute force attacks commonly target management ports as a means to gain access to a VM. If successful, an attacker can take control over the VM and establish a foothold into your environment.

One way to reduce exposure to a brute force attack is to limit the amount of time that a port is open.

Management ports do not need to be open at all times. They only need to be open while you are connected to the VM, for example to perform management or maintenance tasks.

When just in time is enabled, Security Center uses [network security group \(NSG\)](#) rules, which restrict access to management ports so they cannot be targeted by attackers.



How does just in time access work?

When just in time is enabled, Security Center locks down inbound traffic to your Azure VMs by creating an NSG rule. You select the ports on the VM to which inbound traffic will be locked down. These ports are controlled by the just in time solution.

When a user requests access to a VM, Security Center checks that the user has [Role-Based Access Control \(RBAC\)](#) permissions that provide write access for the VM. If they have write permissions, the request is approved and Security Center automatically configures the Network Security Groups (NSGs) to allow inbound traffic to the selected ports for the amount of time you specified. After the time has expired, Security Center restores the NSGs to their previous states. Those connections that are already established are not being interrupted, however.

NOTE

Security Center just in time VM access currently supports only VMs deployed through Azure Resource Manager. To learn more about the classic and Resource Manager deployment models see [Azure Resource Manager vs. classic deployment](#).

Using just in time access

1. Open the **Security Center** dashboard.
2. In the left pane, select **Just in time VM access**.

Security Center - Overview
Showing subscription 'Contoso IT - demo'

Search (Ctrl+/)

POLICY & COMPLIANCE

- Coverage
- Security policy

RESOURCE SECURITY HYGIENE

- Recommendations
- Compute & apps
- Networking
- Data & storage
- Identity & access (Preview)
- Security solutions

THREAT PROTECTION

- Security alerts
- Custom alert rules (Preview)
- Threat intelligence (Preview)

AUTOMATION & ORCHESTRATION

- Playbooks (Preview)

ADVANCED CLOUD DEFENSE

- Adaptive application controls
- Just in time VM access**
- File Integrity Monitoring (Pre)

Subscriptions

Policy & compliance

Subscription coverage

1 TOTAL

- COVERED (STANDARD) 1
- COVERED (FREE) 0
- NOT COVERED 0

325 Covered resources

Policy compliance

Overall compliance 36%

Least compliant subscriptions Contoso IT - demo 36%

Show policy compliance of your environment >

Resource security hygiene

Recommendations

36 TOTAL

- HIGH SEVERITY 15
- MEDIUM SEVERITY 12
- LOW SEVERITY 9

217 Unhealthy resources

Resource health monitoring

- 161 Compute & apps
- 99 Data & storage
- 64 Networking
- 1 Identity & access

Threat protection

Security alerts by severity

30 TOTAL

- HIGH SEVERITY 4
- MEDIUM SEVERITY 23
- LOW SEVERITY 3

18 Attacked resources

Security alerts over time

8
6
4
2
0

24 Sun 1 Sun 8 Sun

- HIGH SEVERITY 4
- MEDIUM SEVERITY 23
- LOW SEVERITY 3

The **Just in time VM access** window opens.

Just in time VM access
✕

Last week

✓ **What is just in time VM access?**

Just in time VM access enables you to lock down your VMs in the network level by blocking inbound traffic to specific ports. It enables you to control the access and reduce the attack surface to your VMs, by allowing access only upon a specific need.

✓ **How does it work?**

Upon a user request, based on Azure RBAC, Security Center will decide whether to grant access. If a request is approved, Security Center automatically configures the NSGs to allow inbound traffic to these ports, for the requested amount of time, after which it restores the NSGs to their previous states.

[For more information go to the documentation >](#)

Virtual machines

Configured
Recommended
No recommendation

VMs for which the just in time VM access control is already in place. Presented data is for the last week.

5 VMs
Request access

	VIRTUAL MACHINE	APPROVED	LAST ACCESS	LAST USER	
<input type="checkbox"/>	vm1	1 Requests	9/13/17 5:19 AM	<User ID>	...
<input type="checkbox"/>	vm2	1 Requests	9/14/17 5:07 AM	<User ID>	...
<input type="checkbox"/>	vm2WL	0 Requests	N/A	N/A	...
<input type="checkbox"/>	vm3WL	0 Requests	N/A	N/A	...
<input type="checkbox"/>	win2012dc	0 Requests	N/A	N/A	...

Just in time VM access provides information on the state of your VMs:

- **Configured** - VMs that have been configured to support just in time VM access. The data presented is for the last week and includes for each VM the number of approved requests, last access date and time, and last user.
 - **Recommended** - VMs that can support just in time VM access but have not been configured to. We recommend that you enable just in time VM access control for these VMs. See [Configuring a just in time access policy](#).
 - **No recommendation** - Reasons that can cause a VM not to be recommended are:
 - Missing NSG - The just in time solution requires an NSG to be in place.
 - Classic VM - Security Center just in time VM access currently supports only VMs deployed through Azure Resource Manager. A classic deployment is not supported by the just in time solution.
- Other - A VM is in this category if the just in time solution is turned off in the security policy of the subscription or the resource group, or that the VM is missing a public IP and doesn't have an NSG in place.

Configuring a just in time access policy

To select the VMs that you want to enable:

1. Under **Just in time VM access**, select the **Recommended** tab.

Just in time VM access
Last week

What is just in time VM access?
JIT network access enables you to lock down your VMs in the network level by blocking inbound traffic to specific ports. It enables you to control the access and reduce the attack surface to your VMs, by allowing access only upon a specific need.

How does it work?
Upon a user request, based on Azure RBAC, Security Center will decide whether to grant access. If a request was approved, Security Center will automatically configure the NSGs to allow inbound traffic to these ports, for only 3 hours, after which it will restore the NSGs to its previous state.
[For more information go to the documentation >](#)

Virtual machines
[Configured](#) **Recommended** [No recommendation](#)

Virtual machines for which we recommend you to apply the just-in-time network access control.

4 VMs [Enable JIT on 1 VMs](#)

<input type="checkbox"/>	VIRTUAL MACHINE	STATE	SEVERITY
<input checked="" type="checkbox"/>	vm3WL	Open	High
<input type="checkbox"/>	vm6WL	Open	High
<input type="checkbox"/>	vm1WL	Open	High
<input type="checkbox"/>	vm7WL	Open	High

2. Under **VIRTUAL MACHINE**, select the VMs that you want to enable. This puts a checkmark next to a VM.
3. Select **Enable JIT on VMs**.
4. Select **Save**.

Default ports

You can see the default ports that Security Center recommends enabling just in time.

1. Under **Just in time VM access**, select the **Recommended** tab.

Just in time VM access

What is just in time VM access?

Just in time VM access enables you to lock down your VMs in the network level by blocking inbound traffic to specific ports. It enables you to control the access and reduce the attack surface to your VMs, by allowing access only upon a specific need.

How does it work?

Upon a user request, based on Azure RBAC, Security Center will decide whether to grant access. If a request is approved, Security Center automatically configures the NSGs to allow inbound traffic to these ports, for only 3 hours, after which it restores the NSGs to their previous states.

For more information go to the documentation >

Virtual machines

Configured **Recommended** No recommendation

Virtual machines for which we recommend you to apply the just in time VM access control.

21 VMs Enable JIT on 1 VMs

Search to filter items...

VIRTUAL MACHINE	STATE	SEVERITY
BarNgfwJun3	Open	High
webapp3	Open	High
BarWafT2Jun3	Open	High
Amit-Windows	Open	High
WebApp2	Open	High
WinVM	Open	High
F5WafJun40	Open	High
webapp4	Open	High
impervajun3-MX	Open	High
<input checked="" type="checkbox"/> vm2	Open	High
vm1	Open	High

JIT VM access configuration

Configure the ports for which the just in time VM access will be applicable.

PORT	PROT...	ALLOWED SOUR...	IP RANGE	TIME RANGE
22 (Recommended)	Any	Per request	N/A	3 hours ...
3389 (Recommended)	Any	Per request	N/A	3 hours ...
5985 (Recommended)	Any	Per request	N/A	3 hours ...
5986 (Recommended)	Any	Per request	N/A	3 hours ...

2. Under **VMs**, select a VM. This puts a checkmark next to the VM and opens **JIT VM access configuration**. This blade displays the default ports.

Add ports

Under **JIT VM access configuration**, you can also add and configure a new port on which you want to enable the just in time solution.

1. Under **JIT VM access configuration**, select **Add**. This opens **Add port configuration**.

JIT VM access configuration
impervawaf-GW0, Snir-test

+ Add Save Discard

Configure the ports for which the just in time VM access will be applicable.

PORT	PROT...	ALLOWED SOUR...	IP RANGE	TIME RANGE
22	Any	Per request	N/A	3 hours ...
3389	Any	Per request	N/A	3 hours ...
5985	Any	Per request	N/A	3 hours ...
5986	Any	Per request	N/A	3 hours ...

* Port

Protocol
Any TCP UDP

Allowed source IPs
Per request CIDR block

IP range

Max request time
3 (hours)

Discard OK

2. Under **Add port configuration**, you identify the port, protocol type, allowed source IPs, and maximum request time.

Allowed source IPs are the IP ranges allowed to get access upon an approved request.

Maximum request time is the maximum time window that a specific port can be opened.

3. Select **OK**.

NOTE

When JIT VM Access is enabled for a VM, Azure Security Center creates deny all inbound traffic rules for the selected ports in the network security groups associated with it. The rules will either be the top priority of your Network Security Groups, or lower priority than existing rules that are already there. This depends on an analysis performed by Azure Security Center that determines whether a rule is secure or not.

Set just-in-time within a VM

To make it easy to roll out just-in-time access across your VMs, you can set a VM to allow only just-in-time access directly from within the VM.

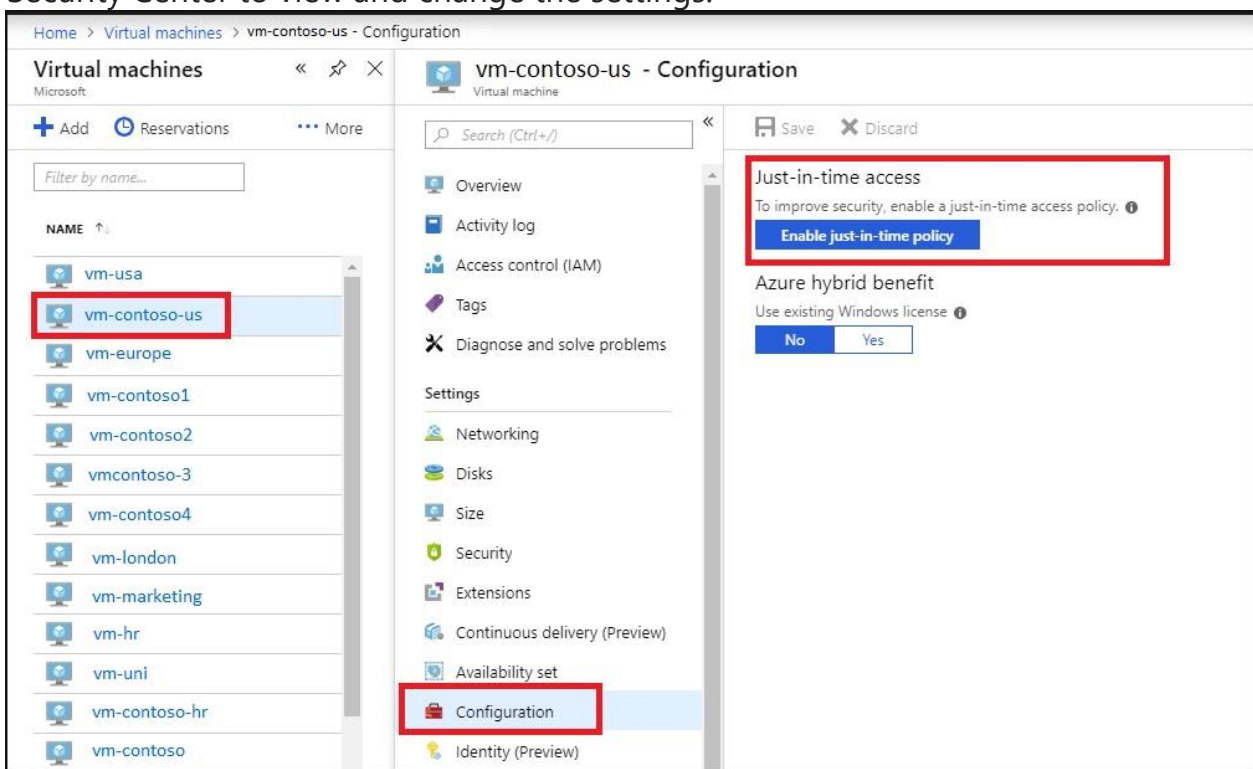
1. In the Azure portal, select **Virtual machines**.
2. Click on the virtual machine you want to limit to just-in-time access.

3. In the menu, click **Configuration**.
4. Under **Just-in-time-access** click **Enable just-in-time policy**.

This enables just-in-time access for the VM using the following settings:

- Windows servers:
 - RDP port 3389
 - 3 hours of access
 - Allowed source IP addresses is set to Per request
- request Linux servers:
 - SSH port 22
 - 3 hours of access
 - Allowed source IP addresses is set to Per request

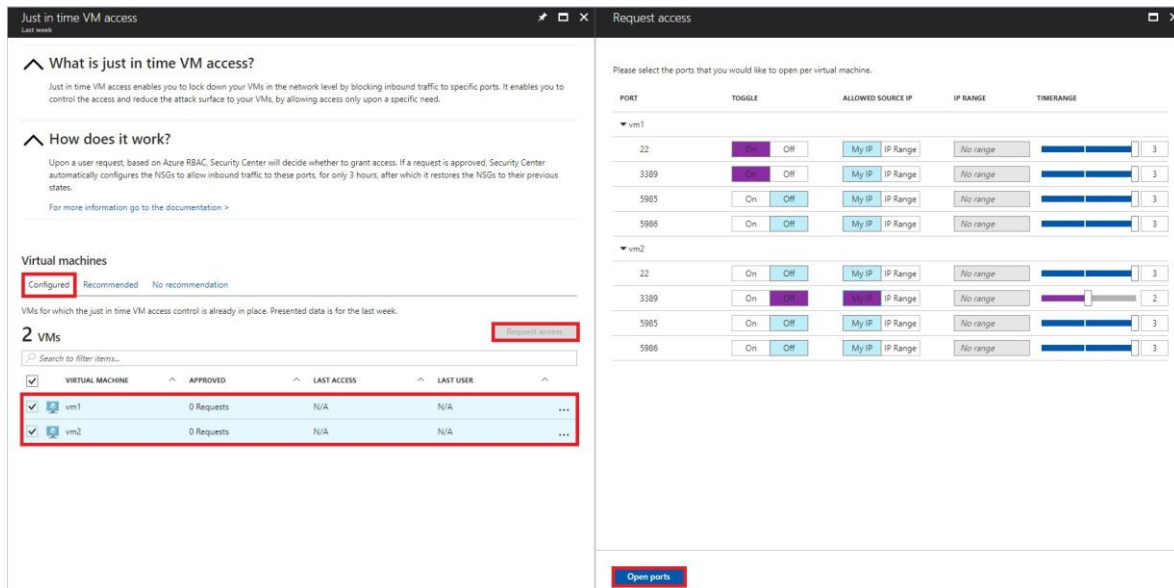
If a VM already has just-in-time enabled, when you go to its configuration page you will be able to see that justin-time is enabled and you can use the link to open the policy in Azure Security Center to view and change the settings.



Requesting access to a VM

To request access to a VM:

1. Under **Just in time VM access**, select the **Configured** tab.
2. Under **VMs**, select the VMs that you want to enable access. This puts a checkmark next to a VM.
3. Select **Request access**. This opens **Request access**.



4. Under **Request access**, you configure for each VM the ports to open along with the source IP that the port is opened to and the time window for which the port is opened. You can request access only to the ports that are configured in the just in time policy. Each port has a maximum allowed time derived from the just in time policy.

5. Select **Open ports**.

NOTE

When a user requests access to a VM, Security Center checks that the user has [Role-Based Access Control \(RBAC\)](#) permissions that provide write access for the VM. If they have write permissions, the request is approved.

NOTE

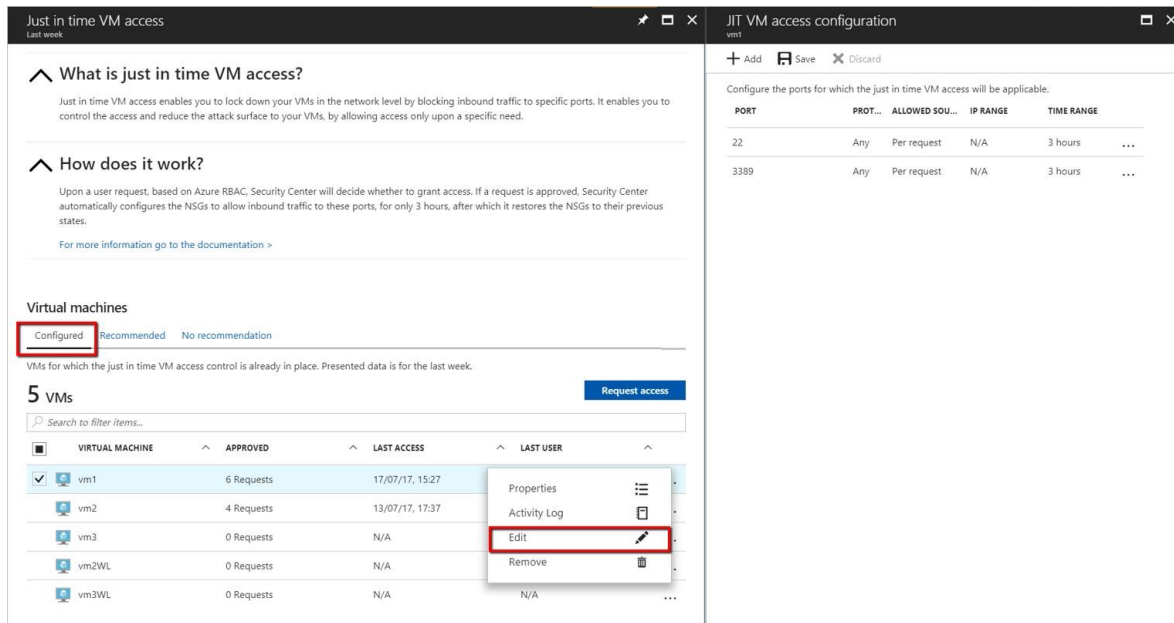
If a user who is requesting access is behind a proxy, the "My IP" option may not work. There may be a need to define the full range of the organization.

Editing a just in time access policy

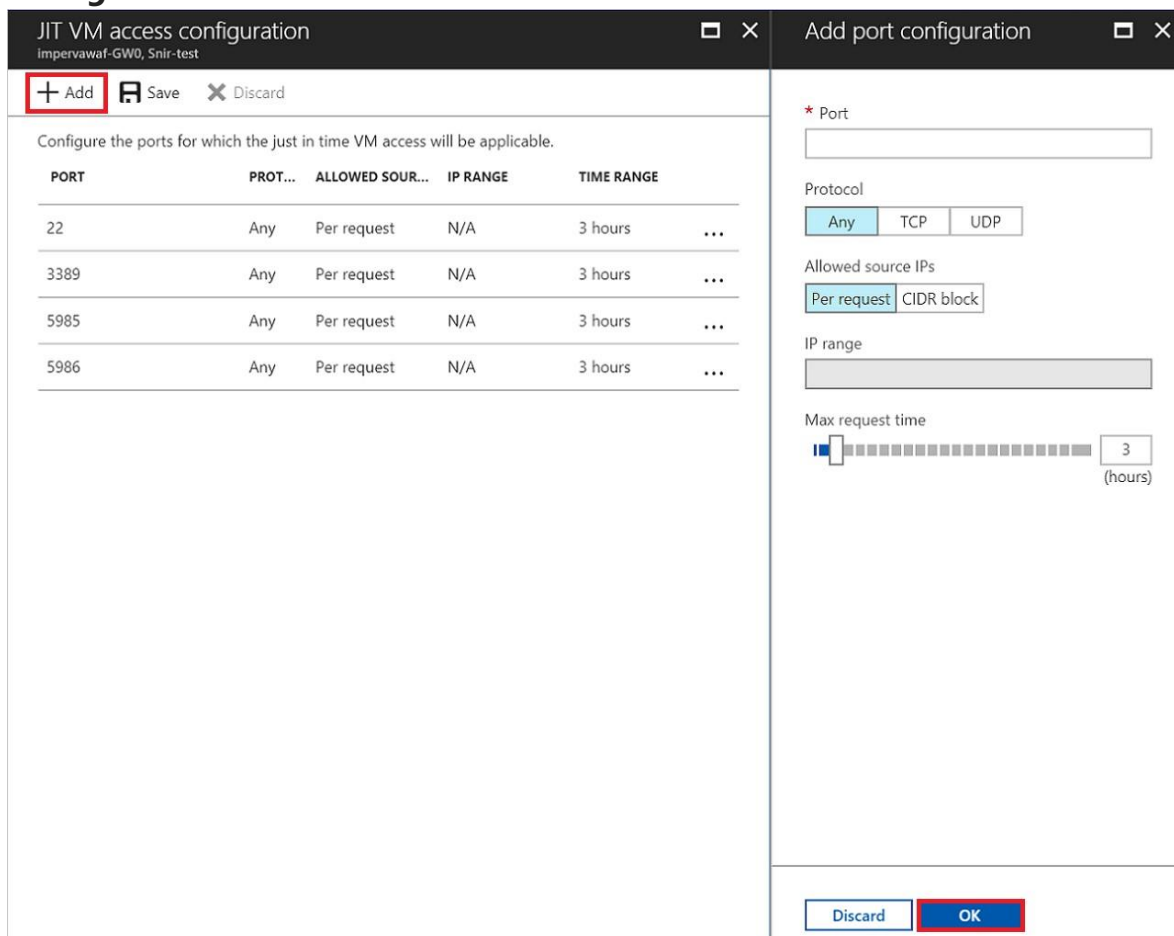
You can change a VM's existing just in time policy by adding and configuring a new port to open for that VM, or by changing any other parameter related to an already protected port.

In order to edit an existing just in time policy of a VM, the **Configured** tab is used:

1. Under **VMs**, select a VM to add a port to by clicking on the three dots within the row for that VM. This opens a menu.
2. Select **Edit** in the menu. This opens **JIT VM access configuration**.



3. Under **JIT VM access configuration**, you can either edit the existing settings of an already protected port by clicking on its port, or you can select **Add**. This opens **Add port configuration**.



4. Under **Add port configuration**, identify the port, protocol type, allowed source IPs, and maximum request time.

5. Select **OK**.

6. Select **Save**.

Auditing just in time access activity

You can gain insights into VM activities using log search. To view logs:

1. Under **Just in time VM access**, select the **Configured** tab.
2. Under **VMs**, select a VM to view information about by clicking on the three dots within the row for that VM. This opens a menu.
3. Select **Activity Log** in the menu. This opens **Activity log**.

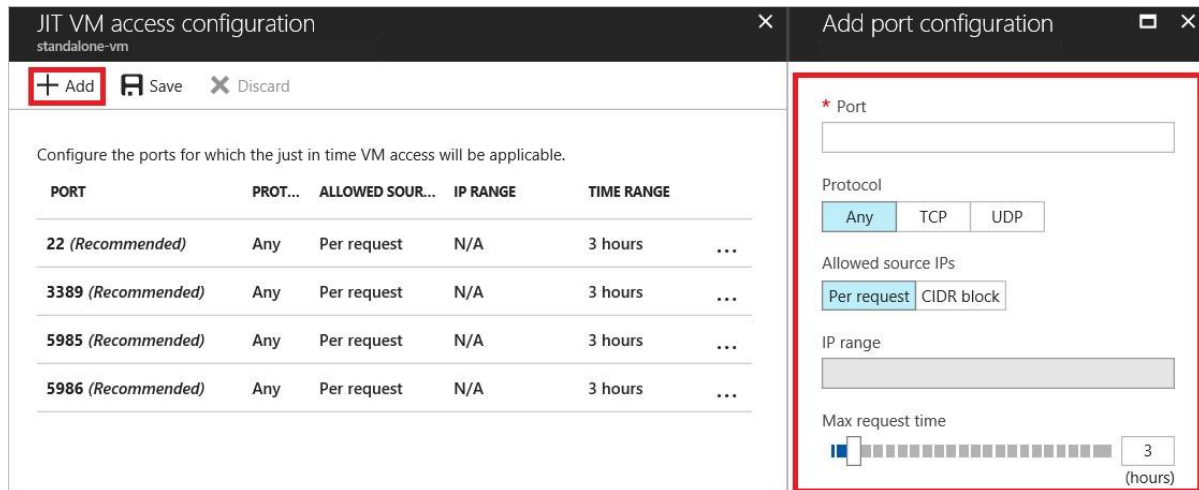
The screenshot shows the 'Just in time VM access' page in the Azure portal. The 'Configured' tab is selected and highlighted with a red box. Below the tabs, there is a section for 'Virtual machines' with a 'Request access' button. A table lists 5 VMs. The first row, 'vm1', is selected, and its context menu is open, with 'Activity Log' highlighted by a red box.

VIRTUAL MACHINE	APPROVED	LAST ACCESS	LAST USER
vm1	6 Requests	17/07/17, 15:27	
vm2	4 Requests	13/07/17, 17:37	
vm3	0 Requests	N/A	
vm2WL	0 Requests	N/A	
vm3WL	0 Requests	N/A	N/A

Activity log provides a filtered view of previous operations for that VM along with time, date, and subscription.

You can download the log information by selecting **Click here to download all the items as CSV**.

Modify the filters and select **Apply** to create a search and log.



1. Under **Add port configuration**, you identify:
 - The port
 - The protocol type
 - Allowed source IPs - IP ranges allowed to get access upon an approved request
 - Maximum request time - maximum time window that a specific port can be opened
2. Select **OK** to save.
3. Attempt to RDP to your VM. This should fail.
4. Enabled JIT for your VM. Attempt to RDP to your VM. This should work.

Harden VMs against malware

Learn how to configure application control in Azure Security Center using this walkthrough.

What are adaptive application controls in Security Center?

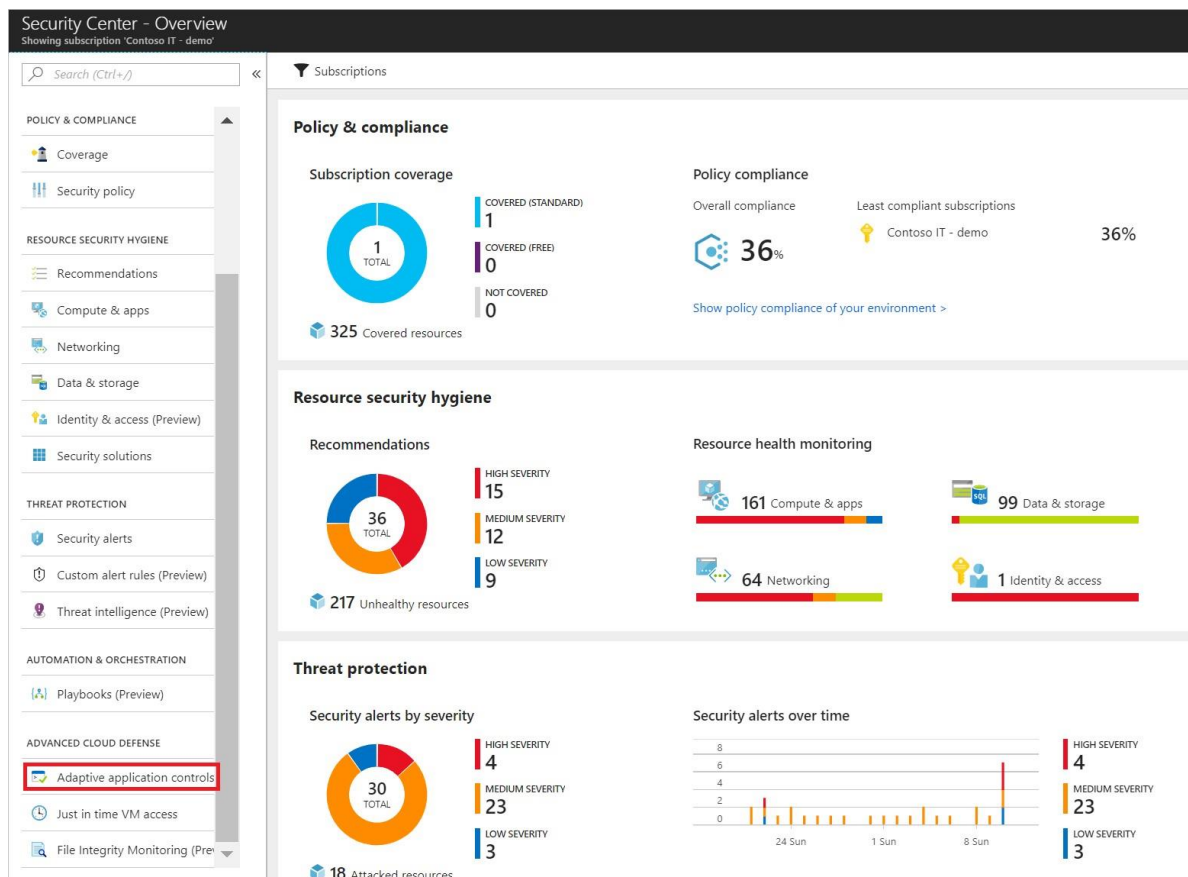
Adaptive application controls is an intelligent, automated end-to-end application whitelisting solution from Azure Security Center. It helps you control which applications can run on your VMs located in Azure, which among other benefits helps harden your VMs against malware. Security Center uses machine learning to analyze the applications running on your VMs and helps you apply the specific whitelisting rules using this intelligence. This capability greatly simplifies the process of configuring and maintaining application whitelisting policies, enabling you to:

- Block or alert on attempts to run malicious applications, including those that might otherwise be missed by antimalware solutions.
- Comply with your organization's security policy that dictates the use of only licensed software.
- Avoid unwanted software to be used in your environment.
- Avoid old and unsupported apps to run.
- Prevent specific software tools that are not allowed in your organization. Enable IT to control the access to sensitive data through app usage.

How to enable adaptive application controls?

Adaptive application controls help you define a set of applications that are allowed to run on configured groups of VMs. This feature is only available for Windows machines (all versions, classic, or Azure Resource Manager). The following steps can be used to configure application whitelisting in Security Center:

1. Open the **Security Center** dashboard.
2. In the left pane, select **Adaptive application controls** located under **Advanced cloud defense**.



The **Adaptive application controls** page appears.

▼ What is application control?

Application control helps you deal with malicious and/or unauthorized software, by allowing only specific applications to run on your VMs

▼ How does it work?

Security Center analyzes data of applications to find VMs for which there is a constant set of running applications. Security Center creates whitelisting rules for each group and presents the rules in the form of a recommendation. Once the recommendation is resolved, Security Center configures it by leveraging Applocker capabilities.

[For more information go to the documentation >](#)

Groups of VMs

Configured Recommended **No recommendation**

Groups of VMs in which the set of applications on the associated VMs keeps changing, and thus is not recommended for an application whitelist control.

NAME	VMs
▼  Contoso IT - demo	11
 CONTOSOA2ADEMO	1
 CONTOSOAZUREHQ-DR	1
 DRYRUN2	3
 DRYRUN40	3
 DRYRUN50	3

The **Groups of VMs** section contains three tabs:

- **Configured:** list of groups containing the VMs that were configured with application control.
- **Recommended:** list of groups for which application control is recommended. Security Center uses machine learning to identify VMs that are good candidates for application control based on whether the VMs consistently run the same applications.
- **No recommendation:** list of groups containing VMs without any application control recommendations. For example, VMs on which applications are always changing, and haven't reached a steady state.

NOTE

Security Center uses a proprietary clustering algorithm to create groups of VMs making sure that similar VMs get the optimal recommended application control policy.

Configure a new application control policy

1. Click on the **Recommended** tab for a list of groups with application control recommendations:

> What is application control?

> How does it work?

Configured Recommended No recommendation

Groups of VMs for which we recommend to apply the application whitelist control.

GROUP NAME	VMS	STATE	SEVERITY
ASC DEMO	2	Open	
CONTOSOWEB	1	Open	High
WL1	1	Open	High
Contoso IT - demo	10	Open	
GROUP1	2	Open	High
GROUP2	2	Open	High
GROUP3	3	Open	High
GROUP1-EU	2	Open	High
GROUP2-EU	1	Open	High

The list includes:

- **NAME:** the name of the subscription and group
- **VMs:** the number of virtual machines in the group
- **STATE:** the state of the recommendations
- **SEVERITY:** the severity level of the recommendations

2. Click on a group to open the **Create application control rules** option.

Create application control rules

Description
The steps below will guide you through the process of creating the rules that are unique to this specific group.

Select VMs

VIRTUAL MACHINE	STATE	SEVERITY
<input checked="" type="checkbox"/> contoso-1	Open	High
<input checked="" type="checkbox"/> contoso-2	Open	High
<input checked="" type="checkbox"/> contoso-3	Open	High

Recommended applications
The following applications are very frequent on the VMs within this group and are highly recommended for whitelisting rules.

NAME	FILE TYPES	EXPLOITABLE	USERS
<input checked="" type="checkbox"/> app-1	3 types		1 Users
<input checked="" type="checkbox"/> app-2	1 types		1 Users

More applications
The following applications have been seen on VMs within this group, but are recommended for your review.

NAME	FILE TYPES	EXPLOITABLE	USERS
<input type="checkbox"/> apps			

Notice: Adaptive application controls will be set on audit mode. You can later edit the policy and change it to enforce mode.

Create

3. In the **Select VMs**, review the list of recommended VMs and uncheck any you do not want to apply an application whitelising policy to. Next, you see two lists:

- **Recommended applications:** a list of applications that are frequent on the VMs within this group, and are recommended to be allowed to run.
 - **More applications:** a list of applications that are either less frequent on the VMs within this group or that are known as Exploitables (see more below), and recommended for review.
4. Review the applications in each of the lists, and uncheck any you do not want to apply. Each list includes:
- **NAME:** the certificate information or the full path of an application
 - **FILE TYPES:** the application file type. This can be EXE, Script, MSI, or any permutation of these types.
- EXPLOITABLE:** a warning icon indicates if a specific application could be used by an attacker to bypass an application whitelisting solution. It is recommended to review these applications prior to their approval.
- **USERS:** users that are recommended to be allowed to run an application
5. Once you finish your selections, select **Create**.
- After you select Create, Azure Security Center automatically creates the appropriate rules on top of the built-in application whitelisting solution available on Windows servers (AppLocker).

NOTE

- Security Center relies on a minimum of two weeks of data in order to create a baseline and populate the unique recommendations per group of VMs. New customers of Security Center standard tier should expect a behavior in which at first their groups of VMs appear under the *no recommendation* tab.
- Adaptive Application Controls from Security Center doesn't support VMs for which an AppLocker policy is already enabled by either a GPO or a local security policy.
- As a security best practice, Security Center will always try to create a publisher rule for applications that are selected to be allowed, and only if an application doesn't have a publisher information (aka not signed), a path rule will be created for the full path of the specific application.



Editing and monitoring a group configured with application control

1. To edit and monitor a group configured with an application whitelisting policy, return to the **Adaptive application controls** page and select **CONFIGURED** under **Groups of VMs**:

Groups of VMs

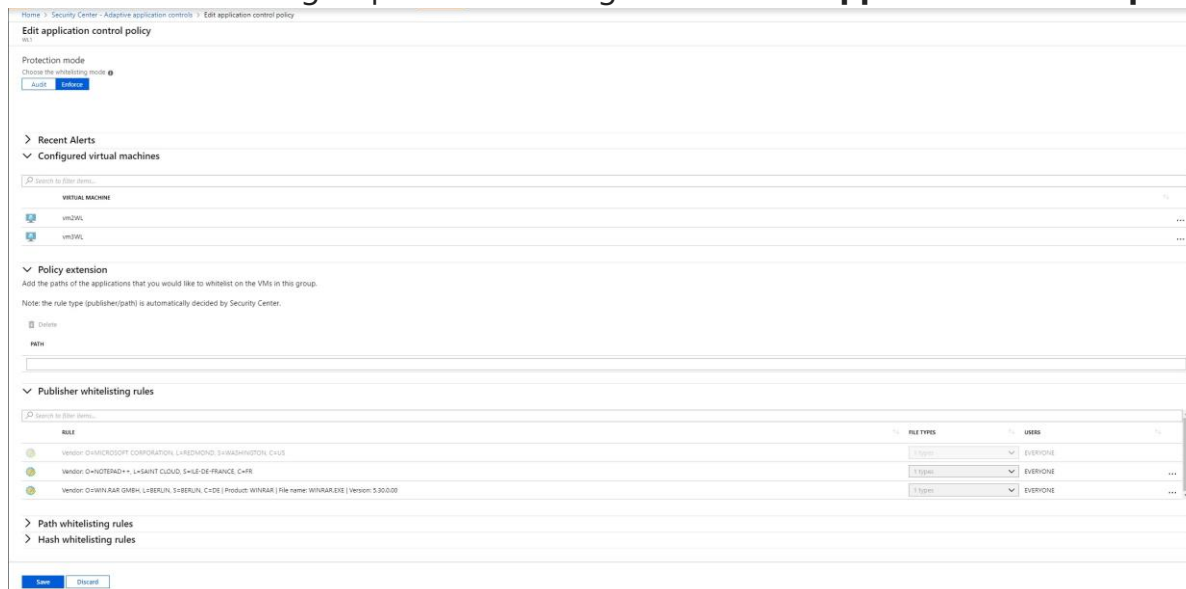
Configured Recommended No recommendation

Groups of VMs for which an application whitelist is already applied and can be centrally managed.

NAME	VMS	MODE	ISSUES
▼  Contoso IT - demo	6		
 A-MANAGEMENT	2	Audit	
 CONTOSOONPREMHQ	3	Audit	
 CONTOSORETAILDEV	1	Audit	

The list includes:

- **Name:** the name of the subscription and group
- **VMs:** the number of virtual machines in the group
- **Mode:** Audit mode will log attempts to run non-whitelisted applications; Enforce will not allow nonwhitelisted applications to run **Alerts:** any current violations
- 2. Click on a group to make changes in the **Edit application control policy** page.



RULE	FILE TYPES	USERS
Vendor: O=MICROSOFT CORPORATION, L=FREDMOND, S=WASHINGTON, CN=US	1 types	EVERYONE
Vendor: O=HOTELEMI+, L=SAINT CLOUD, S=ILE-DE-FRANCE, C=FR	1 types	EVERYONE
Vendor: O=WINRAR GMBH, L=BERLIN, S=BERLIN, C=DE Product: WINRAR File name: WINRAR.EXE Version: 5.00.00	1 types	EVERYONE

3. Under **Protection mode**, you have the option to select between the following:

- **Audit:** in this mode, the application control solution does not enforce the rules, and only audits the activity on the protected VMs. This is recommended for scenarios where you want to first observe the overall behavior before blocking an app to run in the target VM.
- **Enforce:** in this mode, the application control solution does enforce the rules, and makes sure that applications that are not allowed to run are blocked.

NOTE

As previously mentioned, by default a new application control policy is always configured in *Audit* mode.

4. Under **Policy extension**, you can add any application path that you want to allow. After you add these paths, Security Center updates the application whitelisting policy on the VMs within the selected group of VMS and creates the appropriate rules for these applications, in addition to the rules that are already in place.

5. Review the current violations listed in the **Recent alerts** section. Click on each line to be redirected to the **Alerts** page within Azure Security Center, and view all the alerts that were detected by Azure Security Center on the associated VMs.

- **Alerts:** any violations that were logged.
- **No. of VMs:** the number of virtual machines with this alert type.

6. Under **Publisher whitelisting rules**, **Path whitelisting rules**, and **Hash whitelisting rules** you can see which application whitelisting rules are currently configured on the VMs within a group, according to the rule collection type. For each rule you can see:

- **Rule**The specific parameters according to which an application is examined by AppLocker to determine if an application is allowed to run.
- **File type**The file types that are covered by a specific rule. This can be any of the following: EXE, MSI, or any permutation of those file types.
- **Users**Name or number of users who are allowed to run an application that is covered by an application whitelisting rule.

RULE	USERS
Vendor: O=MICROSOFT CORPORATION, L=REDMOND, S=WASHINGTON, C=US	5 Users allowed ...

7. Click on the three dots at the end of each line if you want to delete the specific rule or edit the allowed users.

8. After making changes to an **Adaptive application controls** policy, click **Save**.

Not recommended list

Security Center only recommends application whitelisting policies for virtual machines running a stable set of applications. Recommendations are not created if applications on the associated VMs keep changing.

Groups of VMs

Configured Recommended **No recommendation**

Groups of VMs in which the set of applications on the associated VMs keeps changing, and thus is not recommended for an application whitelist control.

NAME	↑↓ VMs
▼ 🔑 Contoso IT - demo	11
🖥️ CONTOSOA2ADEMO	1
🖥️ CONTOSOA2UREHQ-DR	1
🖥️ DRYRUN2	3
🖥️ DRYRUN40	3
🖥️ DRYRUN50	3

The list contains:

- **NAME:** the name of the subscription and group
- **VMs:** the number of virtual machines in the group

Azure Security Center enables you to define an application whitelisting policy on non-recommended groups of VMs as well. Follow the same principles as were previously described, to configure an application whitelisting policy on those groups as well.

File Integrity Monitoring

Learn how to configure File Integrity Monitoring (FIM) in Azure Security Center using this walkthrough.

What is FIM in Security Center?

File Integrity Monitoring (FIM), also known as change monitoring, examines files and registries of operating system, application software, and others for changes that might indicate an attack. A comparison method is used to determine if the current state of the file is different from the last scan of the file. You can leverage this comparison to determine if valid or suspicious modifications have been made to your files.

Security Center's File Integrity Monitoring validates the integrity of Windows files, Windows registry, and Linux files. You select the files that you want monitored by enabling FIM. Security Center monitors files with FIM enabled for activity such as:

- File and Registry creation and removal
- File modifications (changes in file size, access control lists, and hash of the content)
- Registry modifications (changes in size, access control lists, type, and the content)

Security Center recommends entities to monitor, which you can easily enable FIM on. You can also define your own FIM policies or entities to monitor. This walkthrough shows you how.

NOTE

The File Integrity Monitoring (FIM) feature works for Windows and Linux computers and VMs and is available on the Standard tier of Security Center. See [Pricing](#) to learn more about Security Center's pricing tiers. FIM uploads data to the Log Analytics workspace. Data charges apply, based on the amount of data you upload. See [Log Analytics pricing](#) to learn more.

NOTE

FIM uses the Azure Change Tracking solution to track and identify changes in your environment. When File Integrity Monitoring is enabled, you have a **Change Tracking** resource of type Solution. If you remove the **Change Tracking** resource, you disable the File Integrity Monitoring feature in Security Center.

Which files should I monitor?

You should think about the files that are critical for your system and applications when choosing which files to monitor. Consider choosing files that you don't expect to change without planning. Choosing files that are frequently changed by applications or operating system (such as log files and text files) create a lot of noise which make it difficult to identify an attack.

Security Center recommends which files you should monitor as a default according to known attack patterns that include file and registry changes.

Using File Integrity Monitoring

1. Open the **Security Center** dashboard.
2. In the left pane under **Advanced Cloud Defense**, select **File Integrity Monitoring**.

Security Center - Overview
Showing 10 subscriptions

Search (Ctrl+*/*)

Subscriptions

Policy & compliance

Subscription coverage

3 TOTAL

- Covered (standard) 2
- Covered (free) 1
- Not covered 0

166 Covered resources

Policy compliance

Overall compliance 47%

Least compliant subscriptions

- QA-RomeCore-OMSTest2-Prod 46%
- OMS Security Nir Gafni 47%

Show policy compliance of your environment >

Resource security hygiene

Secure score

800 SCORE

Secure score 800 of 1.2K

37 Active recommendations

Resource health monitoring

- 60 Compute & apps
- 69 Data & storage
- 35 Networking
- 2 Identity & access

Threat protection

Security alerts by severity

66 TOTAL

- High Severity 36
- Medium Severity 23
- Low Severity 7

5 Attacked resources

Security alerts over time

36 High severity

23 Medium severity

7 Low severity

19 Sun 26 Sun 2 Sun

File Integrity Monitoring

File Integrity Monitoring opens.

File Integrity Monitoring

Choose a workspace to view its File Integrity Monitoring dashboard

WORKSPACE NAME	TOTAL CHANGES	TOTAL COMPUTERS	LOCATION	SUBSCRIPTION	
testingworkspacecmdlet	0	0	East US	Contoso IT - demo	UPGRADE PLAN
a-mgmtworkspace	0	2	East US	Contoso IT - demo	ENABLE
contosoretail-it	351	56	East US	Contoso IT - demo	
defaultworkspace-e4272367-5645-4c4e-9c67-3b...	0	1	East US	Contoso IT - demo	
oms-experience-center-2016	0	21	East US	Contoso IT - demo	ENABLE

The following information is provided for each workspace:

- Total number of changes that occurred in the last week (you may see a dash "-" if FIM is not enabled on the workspace)
- Total number of computers and VMs reporting to the workspace
- Geographic location of the workspace
- Azure subscription that the workspace is under

The following buttons may also be shown for a workspace:

- **ENABLE** Indicates that FIM is not enabled for the workspace. Selecting the workspace lets you enable FIM on all machines under the workspace.
- **UPGRADE PLAN**

Indicates that the workspace or subscription is not running under Security Center's Standard tier. To use the FIM feature, your subscription must be running Standard. Selecting the workspace enables you to upgrade to Standard. To learn more about the Standard tier and how to upgrade, see [Upgrade to Security Center's Standard tier for enhanced security](#).

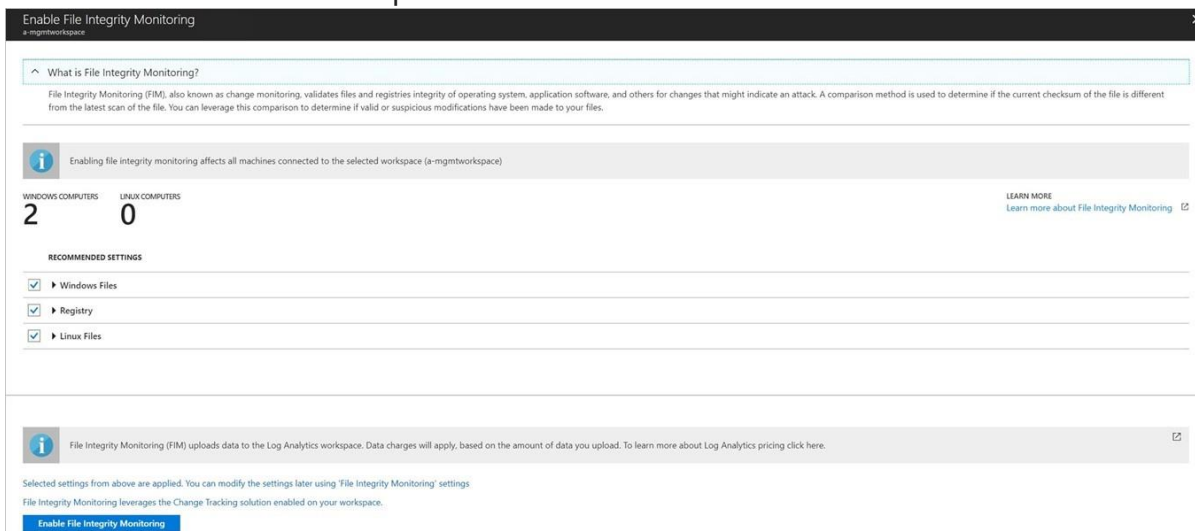
- A blank (there is no button) means that FIM is already enabled on the workspace.

Under **File Integrity Monitoring**, you can select a workspace to enable FIM for that workspace, view the File Integrity Monitoring dashboard for that workspace, or [upgrade](#) the workspace to Standard.

Enable FIM

To enable FIM on a workspace:

1. Under **File Integrity Monitoring**, select a workspace with the **Enable** button.
2. **Enable file integrity monitoring** opens displaying the number of Windows and Linux machines under the workspace.



The recommended settings for Windows and Linux are also listed. Expand **Windows files**, **Registry**, and **Linux files** to see the full list of recommended items.

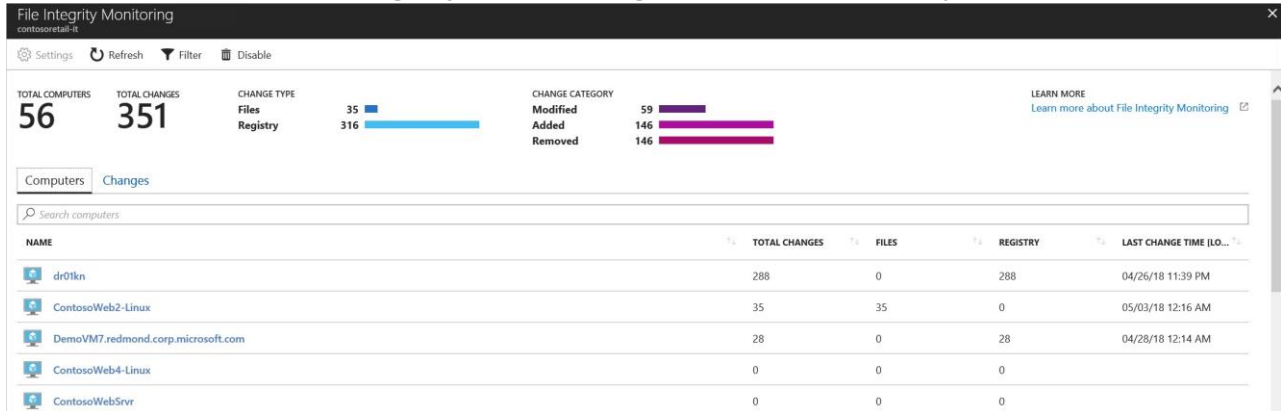
3. Uncheck any recommended entities you do not want to apply FIM to.
4. Select **Apply file integrity monitoring** to enable FIM.

NOTE

You can change the settings at any time. See [Edit monitored entities](#) below to learn more.

View the FIM dashboard

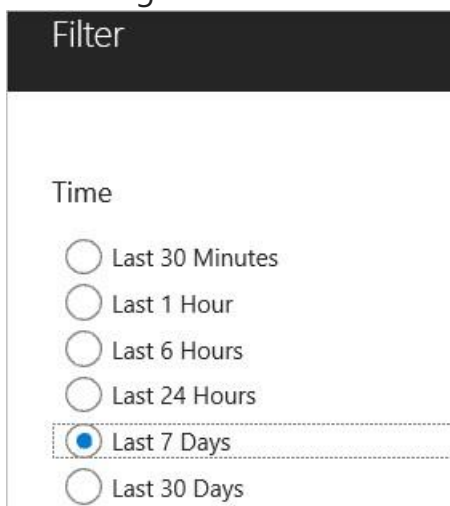
The **File integrity monitoring** dashboard displays for workspaces where FIM is enabled. The FIM dashboard opens after you enable FIM on a workspace or when you select a workspace in the **File Integrity Monitoring** window that already has FIM enabled.



The FIM dashboard for a workspace displays the following:

- Total number of machines connected to the workspace
- Total number of changes that occurred during the selected time period
- A breakdown of change type (files, registry)
- A breakdown of change category (modified, added, removed)

Selecting Filter at the top of the dashboard lets you apply the period of time that you want to see changes for.



The **Computers** tab (shown above) lists all machines reporting to this workspace. For each machine, the dashboard lists:

- Total changes that occurred during the selected period of time
- A breakdown of total changes as file changes or registry changes

Log Search opens when you enter a machine name in the search field or select a machine listed under the Computers tab. Log Search displays all the changes made during the selected time period for the machine. You can expand a change for more information.

The screenshot shows the Log Search interface with a query: `ConfigurationChange | where Computer == "ContosoWeb1.ContosoRetail.com" | where ConfigChangeType in("Files", "Registry") | order by TimeGenerated | render table`. The results table is as follows:

TimeGenerated	Computer	ConfigChangeType	ChangeCategory	Name	FileSystemPath	Size	DateCreated
2/27/2018 7:44:24.763 PM	ContosoWeb1.ContosoRetail.com	Files	Modified	web.config	C:\Windows\Microsoft.NET\Framework\v4.0.30319\Config\web.config	43,133	7/16/2016
2/27/2018 6:44:25.227 PM	ContosoWeb1.ContosoRetail.com	Files	Modified	web.config	C:\Windows\Microsoft.NET\Framework\v4.0.30319\Config\web.config	43,133	7/16/2016
2/27/2018 5:44:35.450 PM	ContosoWeb1.ContosoRetail.com	Files	Modified	web.config	C:\Windows\Microsoft.NET\Framework\v4.0.30319\Config\web.config	43,133	7/16/2016
2/27/2018 4:44:25.603 PM	ContosoWeb1.ContosoRetail.com	Files	Modified	web.config	C:\Windows\Microsoft.NET\Framework\v4.0.30319\Config\web.config	43,133	7/16/2016

The **Changes** tab (shown below) lists all changes for the workspace during the selected time period. For each entity that was changed, the dashboard lists the:

- Computer that the change occurred on
- Type of change (registry or file)
- Category of change (modified, added, removed)
- Date and time of change

The File Integrity Monitoring dashboard shows the following summary statistics:

- TOTAL COMPUTERS: 56
- TOTAL CHANGES: 351
- CHANGE TYPE: Files (35), Registry (316)
- CHANGE CATEGORY: Modified (59), Added (146), Removed (146)

The 'Changes' tab is active, displaying a table of the latest 100 changes:

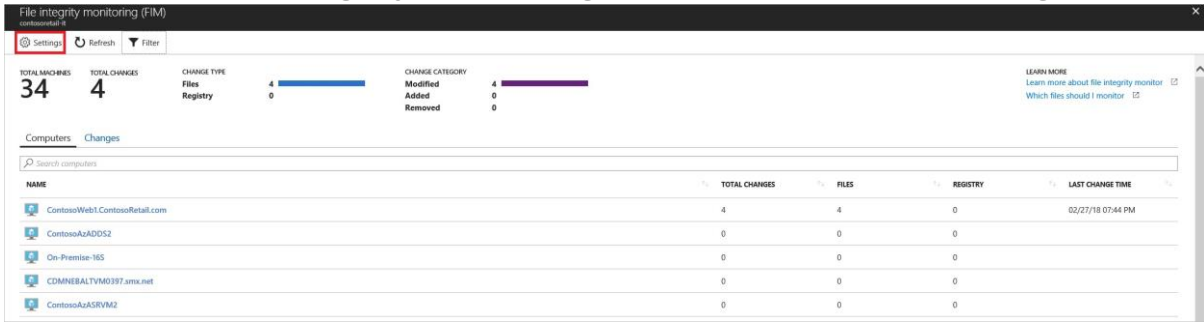
ENTITY	COMPUTER	TYPE	CATEGORY	CHANGE TIME (LOCAL)
/etc/webserver.conf	ContosoWeb2-Linux	Files	Modified	05/03/18 12:16 AM
/etc/webserver.conf	ContosoWeb2-Linux	Files	Modified	05/02/18 11:16 PM
/etc/webserver.conf	ContosoWeb2-Linux	Files	Modified	05/02/18 09:44 PM
/etc/webserver.conf	ContosoWeb2-Linux	Files	Modified	05/02/18 08:06 PM

Change details opens when you enter a change in the search field or select an entity listed under the **Changes** tab.

PROPERTY	VALUE BEFORE	VALUE AFTER
AcIs	[{ "Name": "owner" "Value": "NT AUTHORITY\SYSTEM" } { "...	[{ "Name": "owner" "Value": "NT AUTHORITY\SYSTEM" } { "...
ValueData	C:\Windows\System32\WINTRUST.DLL	C:\Windows\SysWOW64\WINTRUST.DLL
▼ Unchanged prope...		
SourceComput...	cda27197-3886-4fbb-8720-1f2304d1ae1d	No Change
RegistryKey	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Cryptograph...	No Change
Hive	HKEY_LOCAL_MACHINE	No Change
ValueName	Dll	No Change
ValueType	REG_SZ	No Change
Size	32	No Change
SourceSystem	OpsManager	No Change
MG	00000000-0000-0000-0000-000000000001	No Change
ManagementG...	AOI-5bc7c24b-ad6c-4f7f-905c-3ac3f3259938	No Change
TenantId	5bc7c24b-ad6c-4f7f-905c-3ac3f3259938	No Change
VMUUID	c14ded1b-054b-4483-9fd1-1829eb76c406	No Change

Edit monitored entities

1. Return to the **File Integrity Monitoring dashboard** and select **Settings**.



Workspace Configuration opens displaying three tabs: **Windows Registry**, **Windows Files**, and **Linux**

Files. Each tab lists the entities that you can edit in that category. For each entity listed, Security Center identifies if FIM is enabled (true) or not enabled (false). Editing the entity lets you enable or disable FIM.

GROUP	ENABLED	REGISTRY KEY	RECURSIVE
	true	HKEY_LOCAL_MACHINE\software\contoso	false
	true	HKEY_LOCAL_MACHINE\software\contoso\	false
Recommended	true	HKEY_LOCAL_MACHINE\Software\Classes\Directory\Background\ShellEx\ContextMenuHandlers	true
Recommended	true	HKEY_LOCAL_MACHINE\Software\Classes\Directory\ShellEx\ContextMenuHandlers	true
Recommended	true	HKEY_LOCAL_MACHINE\Software\Classes\Directory\Shellex\CopyHookHandlers	true

2. Select an identityprotection. In this example, we selected an item under Windows Registry. **Edit for Change Tracking** opens.

Edit Windows Registry for Change Tracking

Save
Delete
Discard

Enabled

True
False

* Item Name

Group

* Windows Registry Key

Under **Edit for Change Tracking** you can:

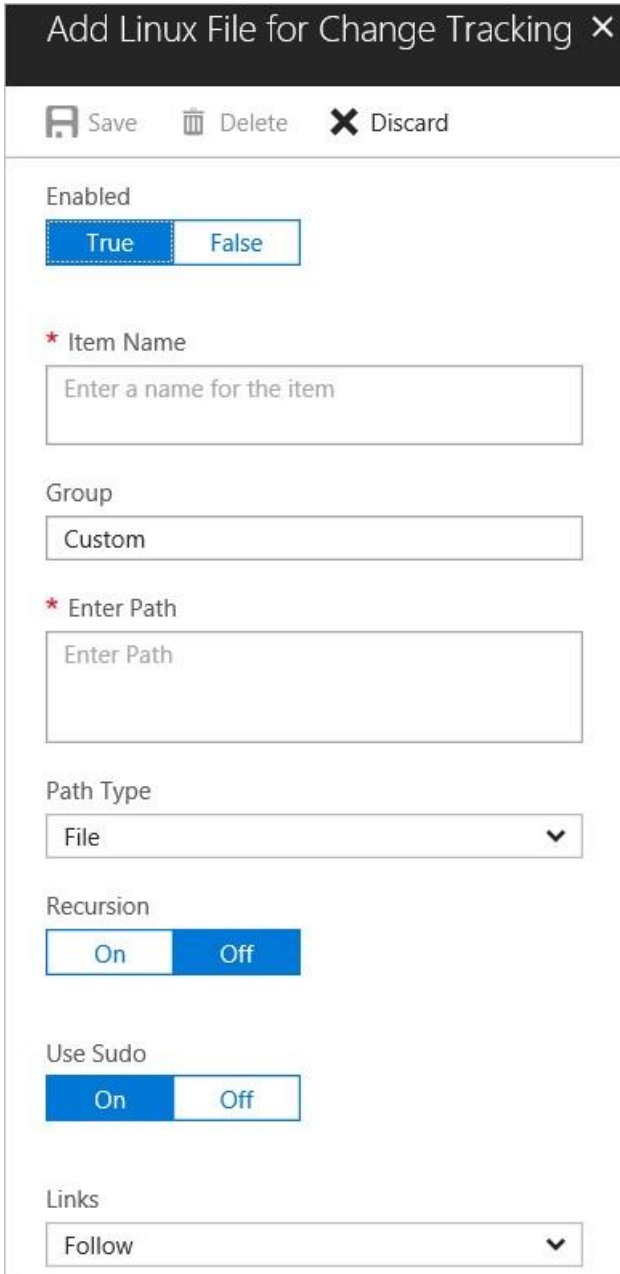
- Enable (True) or disable (False) file integrity monitoring
- Provide or change the entity name
- Provide or change the value or path
- Delete the entity, discard the change, or save the change

Add a new entity to monitor

1. Return to the **File integrity monitoring dashboard** and select **Settings** at the top. **Workspace Configuration** opens.
2. Under **Workspace Configuration**, select the tab for the type of entity that you want to add: Windows Registry, Windows Files, or Linux Files. In this example, we selected **Linux Files**.

GROUP	ENABLED	PATH	TYPE	LINKS	RECURSIVE	SUDO
ChangeDemo	true	/etc/webserver.conf	File	Follow	false	true
OMSAgent	false	/etc/rsyslog.d/95-omsagent.conf	File	Follow	false	true
Recommended	false	/etc/.conf	File	Follow	true	true

3. Select **Add**. **Add for Change Tracking** opens.



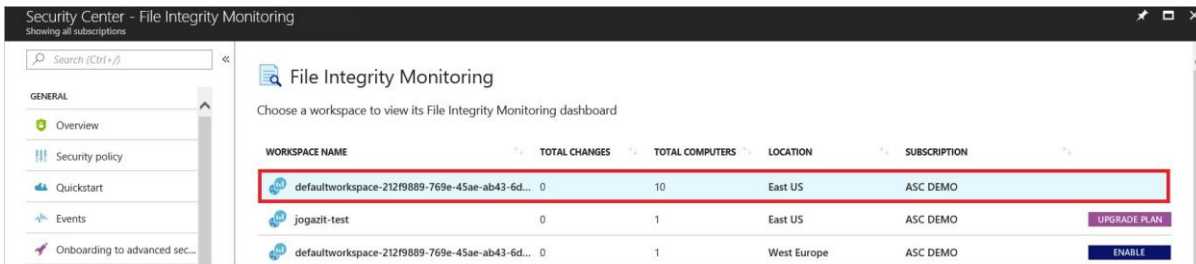
The screenshot shows a dialog box titled "Add Linux File for Change Tracking" with a close button (X) in the top right corner. Below the title bar, there are three buttons: "Save" (with a floppy disk icon), "Delete" (with a trash can icon), and "Discard" (with an X icon). The main content area contains several form fields and controls:

- Enabled:** A toggle switch with "True" selected (highlighted in blue) and "False" unselected.
- * Item Name:** A text input field with the placeholder text "Enter a name for the item".
- Group:** A text input field with the value "Custom".
- * Enter Path:** A text input field with the placeholder text "Enter Path".
- Path Type:** A dropdown menu with "File" selected and a downward arrow.
- Recursion:** A toggle switch with "Off" selected (highlighted in blue) and "On" unselected.
- Use Sudo:** A toggle switch with "On" selected (highlighted in blue) and "Off" unselected.
- Links:** A dropdown menu with "Follow" selected and a downward arrow.

4. On the **Add** page, type the requested information and select **Save**.

Disable monitored entities

1. Return to the **File Integrity Monitoring** dashboard.
2. Select a workspace where FIM is currently enabled. A workspace is enabled for FIM if it is missing the Enable button or Upgrade Plan button.



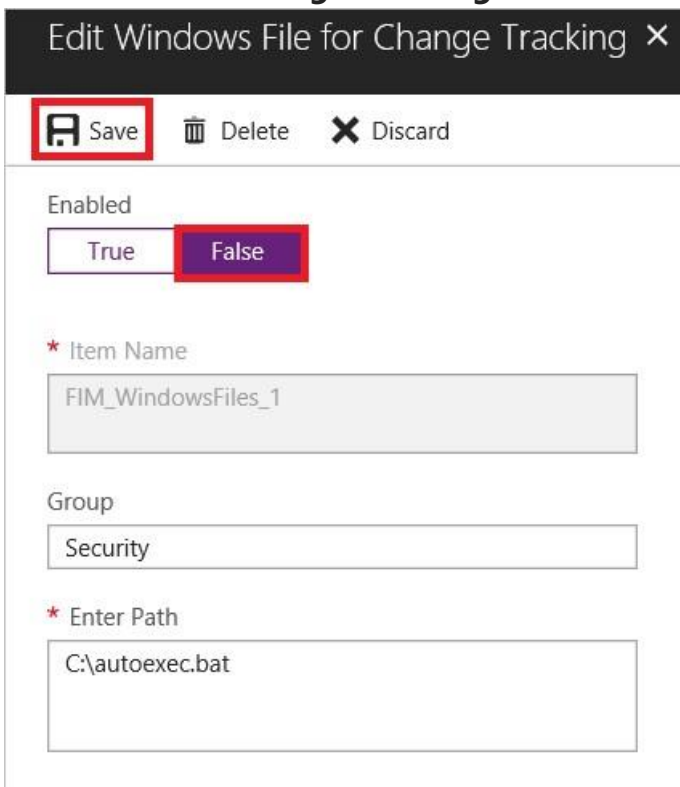
3. Under File Integrity Monitoring, select **Settings**.



4. Under **Workspace Configuration**, select a group where **Enabled** is set to true.



5. Under **Edit for Change Tracking** window set **Enabled** to False.



6. Select **Save**.

Folder and path monitoring using wildcards

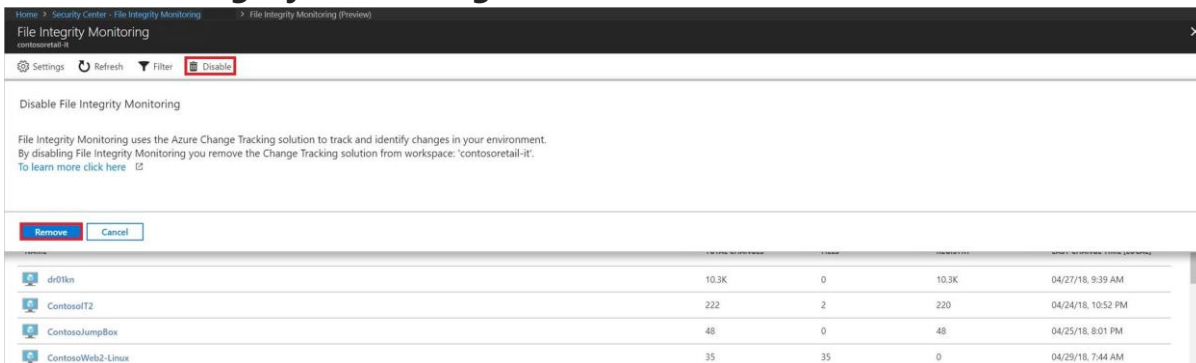
Use wildcards to simplify tracking across directories. The following rules apply when you configure folder monitoring using wildcards:

- Wildcards are required for tracking multiple files.
- Wildcards can only be used in the last segment of a path, such as C:\folder\file or /etc/*.conf
- If an environment variable includes a path that is not valid, validation will succeed but the path will fail when inventory runs.
- When setting the path, avoid general paths such as c:*. * which will result in too many folders being traversed.

Disable FIM

You can disable FIM. FIM uses the Azure Change Tracking solution to track and identify changes in your environment. By disabling FIM, you remove the Change Tracking solution from selected workspace.

1. To disable FIM, return to the **File Integrity Monitoring** dashboard.
2. Select a workspace.
3. Under **File Integrity Monitoring**, select **Disable**.



4. Select **Remove** to disable.

Azure Security Center – Detect and Respond Scenario – Blue Subscription

Security Center continuously analyzes your hybrid cloud workloads using advanced analytics and threat intelligence to alert you to malicious activity. In addition, you can integrate alerts from other security products and services into Security Center, and create custom alerts based on your own indicators or intelligence sources. Once an alert is generated, swift action is needed to investigate and remediate. In this tutorial, you will learn how to:

- Triage security alerts
- Investigate further to determine the root cause and scope of a security incident Search
- security data to aid in investigation

What are security alerts?

Security Center automatically collects, analyzes, and integrates log data from your Azure resources, the network, and connected partner solutions, like firewall and endpoint protection solutions, to detect real threats and reduce false positives. A list of prioritized security alerts is shown in Security Center along with the information you need to quickly investigate the problem and recommendations for how to remediate an attack.


NOTE

For more information about how Security Center detection capabilities work, read [Azure Security Center Detection Capabilities](#).

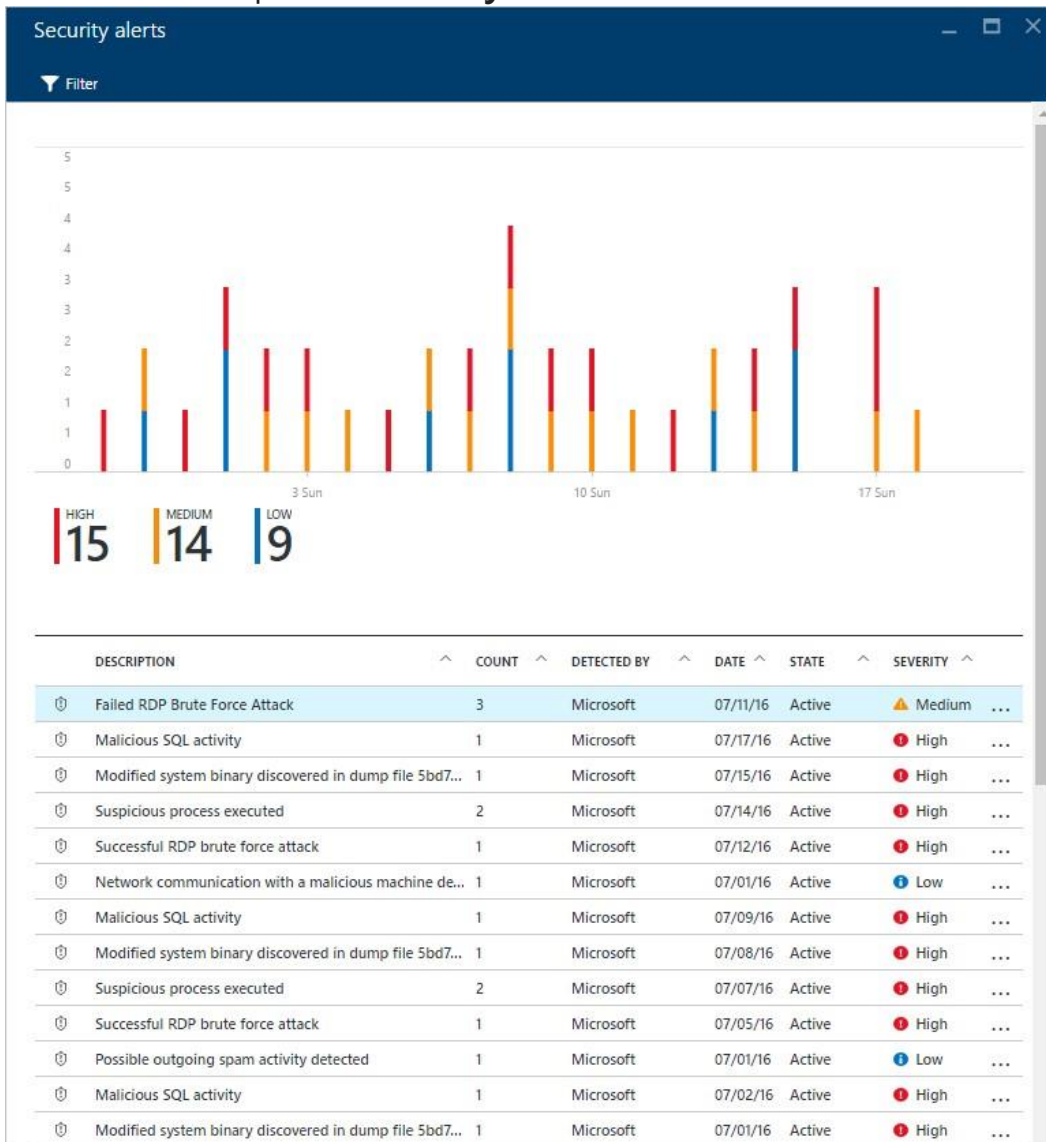
Managing security alerts

You can review your current alerts by looking at the **Security alerts** tile. Follow the steps below to see more details about each alert:

1. On the Security Center dashboard, you see the **Security alerts** tile.



2. Click the tile to open the **Security alerts** to see more details about the alerts.



In the bottom part of this page are the details for each alert. To sort, click the column that you want to sort by. The definition for each column is given below:

- **Description:** A brief explanation of the alert.
- **Count:** A list of all alerts of this specific type that were detected on a specific day.
- **Detected by:** The service that was responsible for triggering the alert.
- **Date:** The date that the event occurred.
- **State:** The current state for that alert. There are two types of states:
 - **Active:** The security alert has been detected.
 - **Dismissed:** The security alert has been dismissed by the user. This status is typically used for alerts that were investigated and either mitigated or found not to be an actual attack.
- **Severity:** The severity level, which can be high, medium or low.

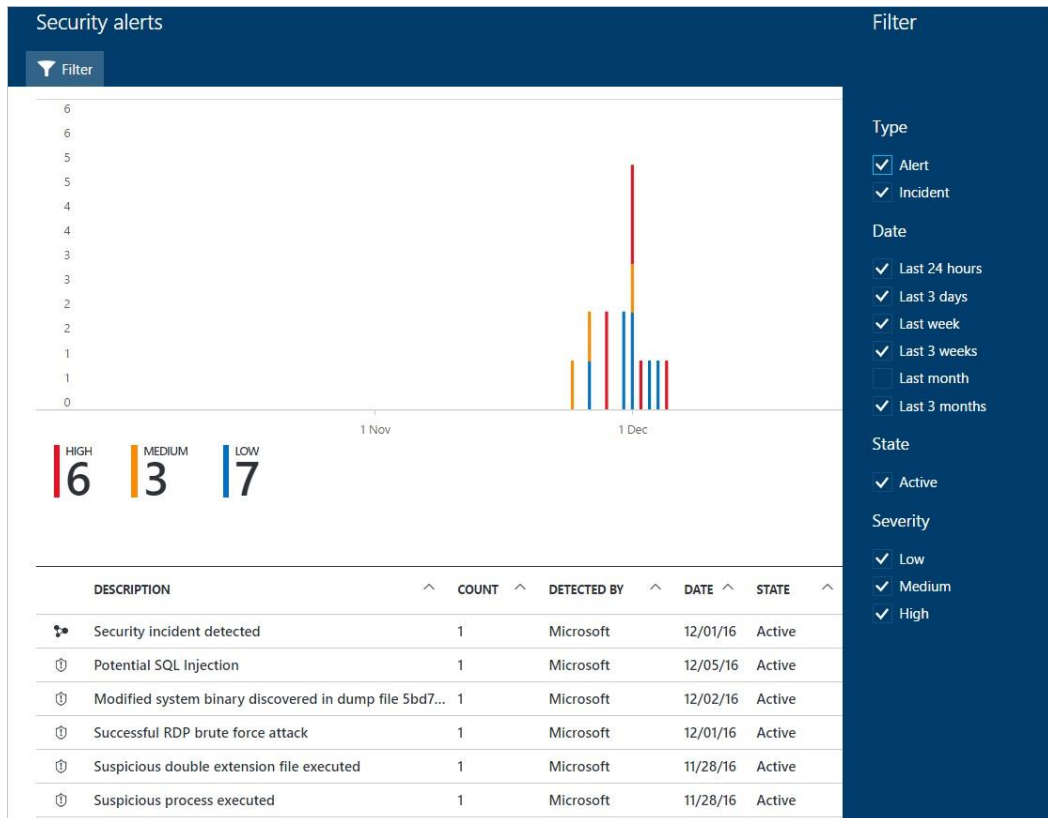
NOTE

Security alerts generated by Security Center will also appear under Azure Activity Log. For more information about how to access Azure Activity Log, read [View activity logs to audit actions on resources](#).

Filtering alerts

You can filter alerts based on date, state, and severity. Filtering alerts can be useful for scenarios where you need to narrow the scope of security alerts show. For example, you might you want to address security alerts that occurred in the last 24 hours because you are investigating a potential breach in the system.

1. Click **Filter** on the **Security Alerts**. The **Filter** opens and you select the date, state, and severity values you wish to see.



Respond to security alerts

Select a security alert to learn more about the event(s) that triggered the alert and what, if any, steps you need to take to remediate an attack. Security alerts are grouped by type and date. Clicking a security alert opens a page containing a list of the grouped alerts.

The screenshot shows the 'Failed RDP Brute Force Attack' alert details page. It features a 'Filter' button and a table with columns for Attacked Resource, Count, Detection Time, State, and Severity.

ATTACKED RESOURCE	COUNT	DETECTION TIME	STATE	SEVERITY
vm1classic	1	8:41:37 PM	Active	Medium
VM2	1	11:14:38 AM	Active	Low
VM1	1	11:14:37 AM	Active	Low

In this case, the alerts that were triggered refer to suspicious Remote Desktop Protocol (RDP) activity. The first column shows which resources were attacked; the second shows how many times the resource was attacked; the third shows the time of the attack; the fourth shows state of the alert; and the fifth shows the severity of the attack. After reviewing this information, click the resource that was attacked.

Failed RDP Brute Force Attack
_ □ ×

DESCRIPTION	Several Remote Desktop login attempts were detected from Windows7, none of them succeeded. Event logs analysis shows that in the last 4 minutes there were 294 failed attempts. 133 of the failed login attempts aimed at non-existent users. 1 of the failed login attempts aimed at existing users.
DETECTION TIME	Monday, July 11, 2016, 8:41:37 PM
SEVERITY	Medium
STATE	Active
ATTACKED RESOURCE	vm1classic
DETECTED BY	Microsoft
ACTION TAKEN	Detected
SOURCE	Windows7
ALERT START TIME (UTC)	07/12/2016 01:37:32
NON-EXISTENT USERS	133
EXISTING USERS	1
FAILED ATTEMPTS	294
SUCCESSFUL LOGINS	0
ATTACK DURATION	4 minutes
FAILED USER LOGONS	quest
REMEDIATION STEPS	<ol style="list-style-type: none"> 1. If available, add the source IP to NSG block list for 24 hours (see https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-nsg/) 2. Enforce the use of strong passwords and do not re-use them across multiple VMs and services (see http://windows.microsoft.com/en-us/Windows7/Tips-for-creating-strong-passwords-and-passphrases) 3. Create an allow list for RDP access in NSG (see https://azure.microsoft.com/en-us/documentation/articles/virtual-networks-nsg/)

In the **Description** field you find more details about this event. These additional details offer insight into what triggered the security alert, the target resource, when applicable the source IP address, and recommendations about how to remediate. In some instances, the source IP address is empty (not available) because not all Windows security events logs include the IP address.

The remediation suggested by Security Center vary according to the security alert. In some cases, you may have to use other Azure capabilities to implement the recommended remediation. For example, the remediation for this attack is to blacklist the IP address that is generating this attack by using a [network ACL](#) or a [network security group](#) rule. For more information on the different types of alerts, read [Security Alerts by Type in Azure Security Center](#).

NOTE

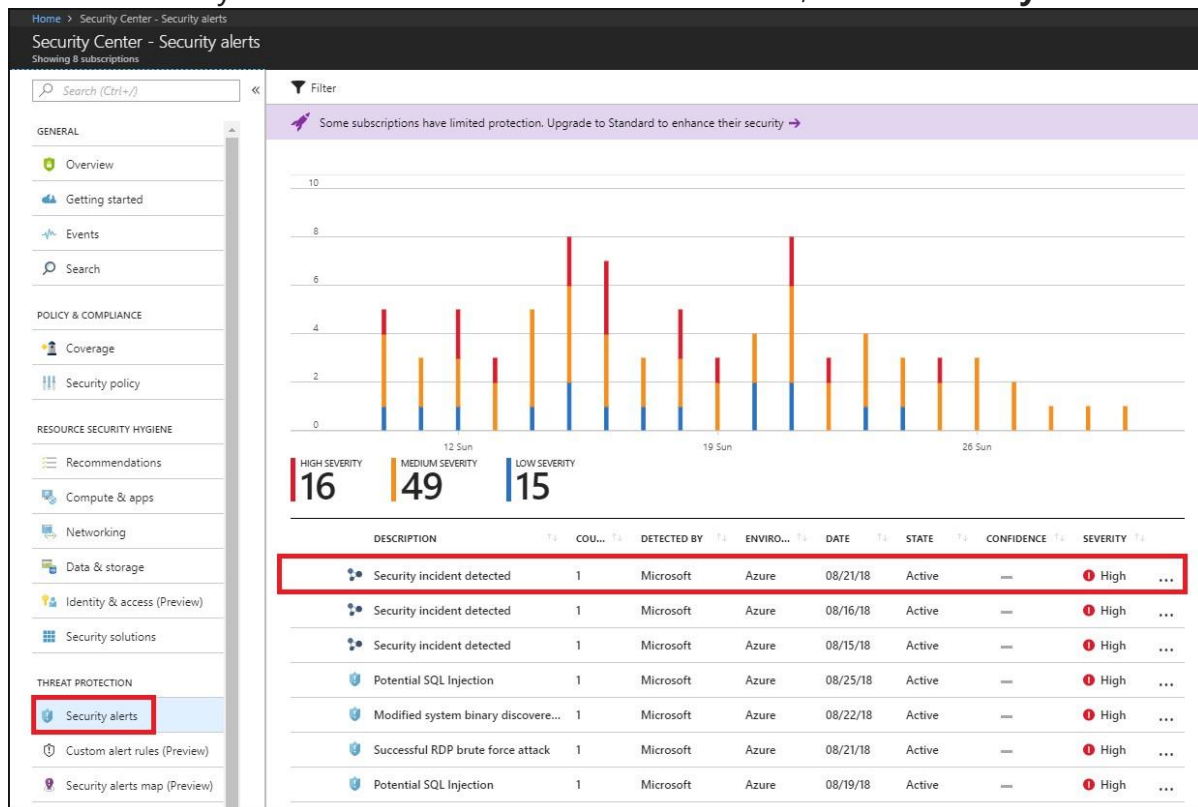
Security Center has released to limited preview a new set of detections that leverage auditd records, a common auditing framework, to detect malicious behaviors on Linux machines. Please send an email with your subscription IDs to [us](#) to join the preview.

Triage security alerts

Security Center provides a unified view of all security alerts. Security alerts are ranked based on the severity and when possible related alerts are combined into a security incident. When triaging alerts and incidents, you should:

- Dismiss alerts for which no additional action is required, for example if the alert is a false positive
- Act to remediate known attacks, for example blocking network traffic from a malicious IP address
- Determine alerts that require further investigation

1. On the Security Center main menu under **DETECTION**, select **Security alerts**:



2. In the list of alerts, click on a security incident, which is a collection of alerts, to learn more about this incident. **Security incident detected** opens.

Security incident detected
□ ×

Incident Detected

DESCRIPTION The incident which started on 2018-01-01T12:00:00.000Z and most recently detected on 2018-01-02T19:00:00.000Z indicate that an attacker has attacked other resources from your virtual machine ContosoWebFE1

DETECTION TIME Thursday, January 4, 2018, 3:02:00 AM

SEVERITY ! High

STATE Active

ATTACKED RESOURCE ContosoWebFE1

SUBSCRIPTION <Subscription ID>

DETECTED BY Microsoft

ENVIRONMENT Azure

REMIEDIATION STEPS

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

Alerts included in this incident

DESCRIPTION	COUNT	DETECTION TIME	ATTACKED RESOURCE	SEVERITY
Successful RDP brute force attack	1	01/04/18, 4:20 AM	ContosoWebFE1	! High
Suspicious SVCHOST process executed	1	01/04/18, 5:19 AM	ContosoWebFE1	i Low
Multiple Domain Accounts Queried	1	01/04/18, 5:21 AM	ContosoWebFE1	i Low

[Continue investigation](#)

3. On this screen you have the security incident description on top, and the list of alerts that are part of this incident. Click on the alert that you want to investigate further to obtain more information.

Successful RDP brute force attack
□ ×

ContosoWebFE1

Learn more

General information

DESCRIPTION	Several Remote Desktop login attempts were detected from FreeRDP (96.81.218.10), some of which were able to successfully login to the machine. Event logs analysis shows that in the last 30 minutes there were 60 failed attempts. 20 of the failed login attempts aimed at non-existent users. 1 of the failed login attempts aimed at existing users.
DETECTION TIME	Thursday, January 4, 2018, 4:20:00 AM
SEVERITY	! High
STATE	Active
ATTACKED RESOURCE	ContosoWebFE1
SUBSCRIPTION	<Subscription ID>
DETECTED BY	Microsoft
ENVIRONMENT	Azure
RESOURCE TYPE	Virtual Machine
TIMEGENERATEDOFFSETMIN	30
SOURCE	FreeRDP (96.81.218.10)
SUCCESSFUL LOGINS	1
ATTACK DURATION	30 minutes
FAILED ATTEMPTS	60
NON-EXISTENT USERS	20
EXISTING USERS	1
REPORTS	Report: RDP Brute Forcing
END TIME UTC	1/4/2018 1:21:00 PM

Remediation steps

Continue investigation

Run playbooks

The type of alert can vary, read [Understanding security alerts in Azure Security Center](#) for more details about the type of alert, and potential remediation steps. For alerts that can be safely dismissed, you can right click on the alert and select the option **Dismiss**:

12/02/17	Ac			
12/01/17	Ac			
11/30/17	Ac			
11/28/17	Ac			
11/28/17	Active	i Low		...
11/28/17	Active	i Low		...

4. If the root cause and scope of the malicious activity is unknown, proceed to the next step to investigate further.


What is a security incident?

In Security Center, a security incident is an aggregation of all alerts for a resource that align with [kill chain](#) patterns. Incidents appear in the [Security Alerts](#) tile and blade. An Incident will reveal the list of related alerts, which enables you to obtain more information about each occurrence.

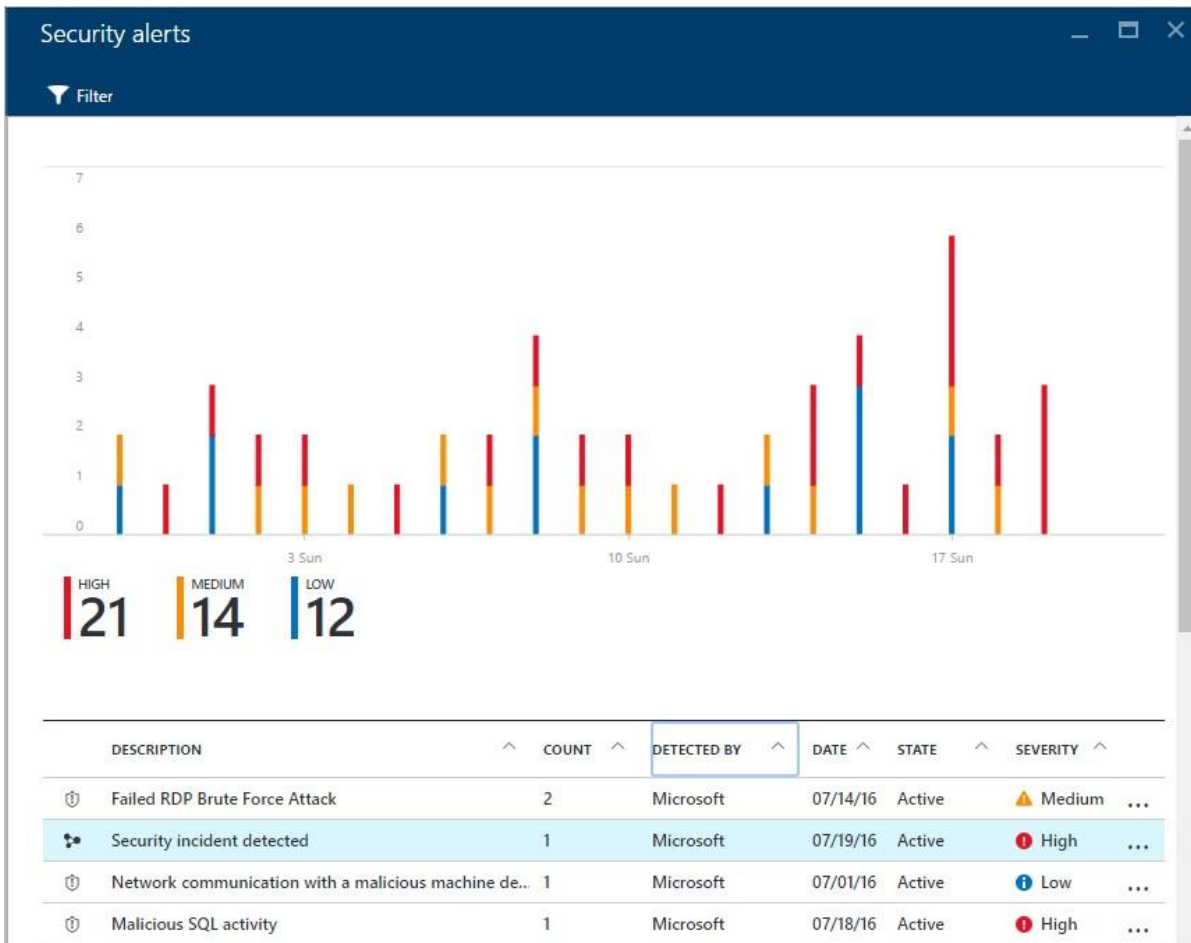
Managing security incidents

You can review your current security incidents by looking at the security alerts tile. Access the Azure Portal and follow the steps below to see more details about each security incident:

1. On the Security Center dashboard, you will see the **Security alerts** tile.



2. Click on this tile to expand it and if a security incident is detected, it will appear under the security alerts graph as shown below:



3. Notice that the security incident description has a different icon compared to other alerts. Click on it to view more details about this incident.

Security incident detected
— □ ×

Incident Detected - Preview

DESCRIPTION The incident which started on 2016-07-16 12:06:19 UTC and most recently detected on 2016-07-20 14:06:23 UTC indicate that an attacker has attacked other resources from your virtual machine VM1

DETECTION TIME Wednesday, July 20, 2016, 9:06:23 AM

SEVERITY ! High

STATE Active

ATTACKED RESOURCE VM1

DETECTED BY Microsoft

ACTION TAKEN Detected



REMEDATION STEPS

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

Alerts included in this incident

DESCRIPTION	COUNT	DETECTION TIME	SEVERITY
Multiple Domain Accounts Queried	1	07/16/16, 7:06 AM	i Low
Suspicious process executed	1	07/17/16, 7:06 AM	! High

4. On the **incident** blade you will see more details about this security incident, which includes its full description, its severity (which in this case is high), its current state (in this case it is still *active*, which implies the user hasn't taken an action to it - this can be done by right clicking on the incident in the **Security alerts** blade), the attacked resource (in this case *VM1*), the remediation steps for the incident, and in the bottom pane you have the alerts that were included in this incident. If you want to obtain more information on each alert, just click on it and another blade will open, as shown below:

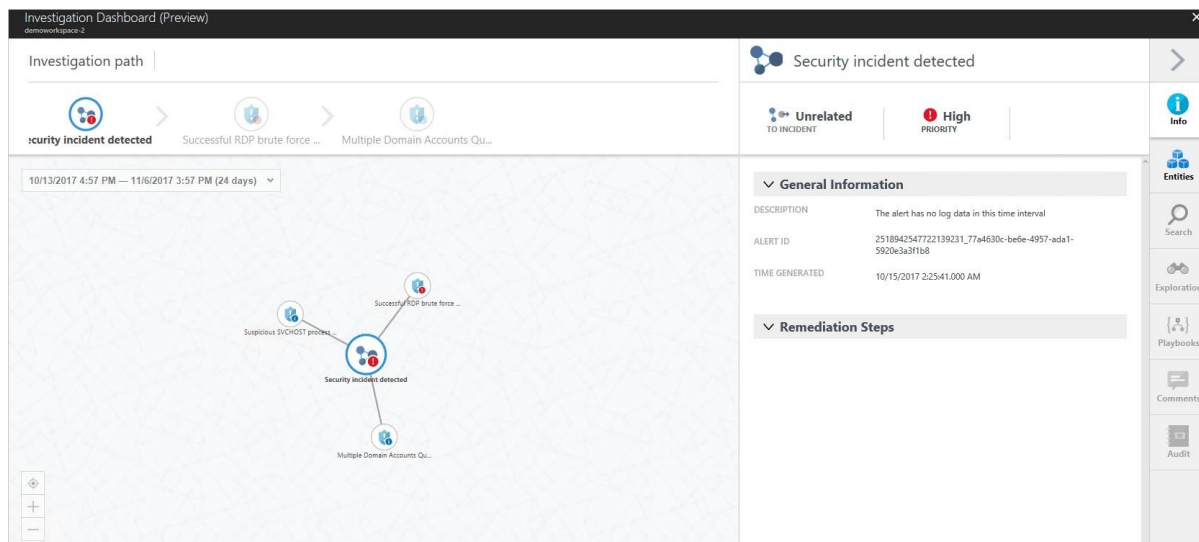
Multiple Domain Accounts Queried	
DESCRIPTION	Analysis of host data has determined that an unusual number of distinct domain accounts are being queried within a short time period from VM1 in your subscription. This kind of activity could be legitimate, but can also be an indication of compromise.
DETECTION TIME	Saturday, July 16, 2016, 7:06:19 AM
SEVERITY	 Low
STATE	Active
ATTACKED RESOURCE	VM1
DETECTED BY	 Microsoft
ACTION TAKEN	Detected
NUMBER OF QUERIED ACCOUNTS OBSERVED IN 24 HOURS	8
LATEST ACCOUNT QUERIED	ContosoUser
COMPROMISED HOST	VM1
FULL COMMAND	net user ContosoUser /do
QUERIED BY	admin
REMEDiation STEPS	Review the details of the accounts queried in this alert to see if you recognise this as legitimate administrative activity.

The information on this blade will vary according to the alert. Read [Managing and responding to security alerts in Azure Security Center](#) for more information on how to manage these alerts. Some important considerations regarding this capability:

- A new filter enables you to customize your view to Incident only, Alerts only, or both.
- The same alert can exist as part of an Incident (if applicable), as well as to be visible as a standalone alert.

Investigate an alert or incident

1. On the **Security alert** page, click **Start investigation** button (if you already started, the name changes to **Continue investigation**).



The investigation map is a graphical representation of the entities that are connected to this security alert or incident. By clicking on an entity in the map, the information about that entity will show new entities, and the map expands. The entity that is selected in the map has its properties highlighted in the pane on the right side of the page. The information available on each tab will vary according to the selected entity. During the investigation process, review all relevant information to better understand the attacker's movement.

2. If you need more evidence, or must further investigate entities that were found during the investigation, proceed to the next step.

Search data for investigation

You can use search capabilities in Security Center to find more evidence of compromised systems, and more details about the entities that are part of the investigation.

To perform a search open the **Security Center** dashboard, click **Search** in the left navigation pane, select the workspace that contains the entities that you want to search, type the search query, and click the search button.

What is a threat intelligence report?

Security Center threat detection works by monitoring security information from your Azure resources, the network, and connected partner solutions. It analyzes this information, often correlating information from multiple sources, to identify threats. This process is part of the Security Center [detection capabilities](#).

When Security Center identifies a threat, it will trigger a [security alert](#), which contains detailed information regarding a particular event, including suggestions for remediation. To assist incident response teams investigate and remediate threats, Security Center includes a threat intelligence report that contains information about the threat that was detected, including information such as the:

- Attacker's identity or associations (if this information is available)
- Attackers' objectives
- Current and historical attack campaigns (if this information is available)
- Attackers' tactics, tools and procedures
- Associated indicators of compromise (IoC) such as URLs and file hashes
- Victimology, which is the industry and geographic prevalence to assist you in determining if your Azure resources are at risk
- Mitigation and remediation information

NOTE

The amount of information in any particular report will vary; the level of detail is based on the malware's activity and prevalence.

Security Center has three types of threat reports, which can vary according to the attack. The reports available are:

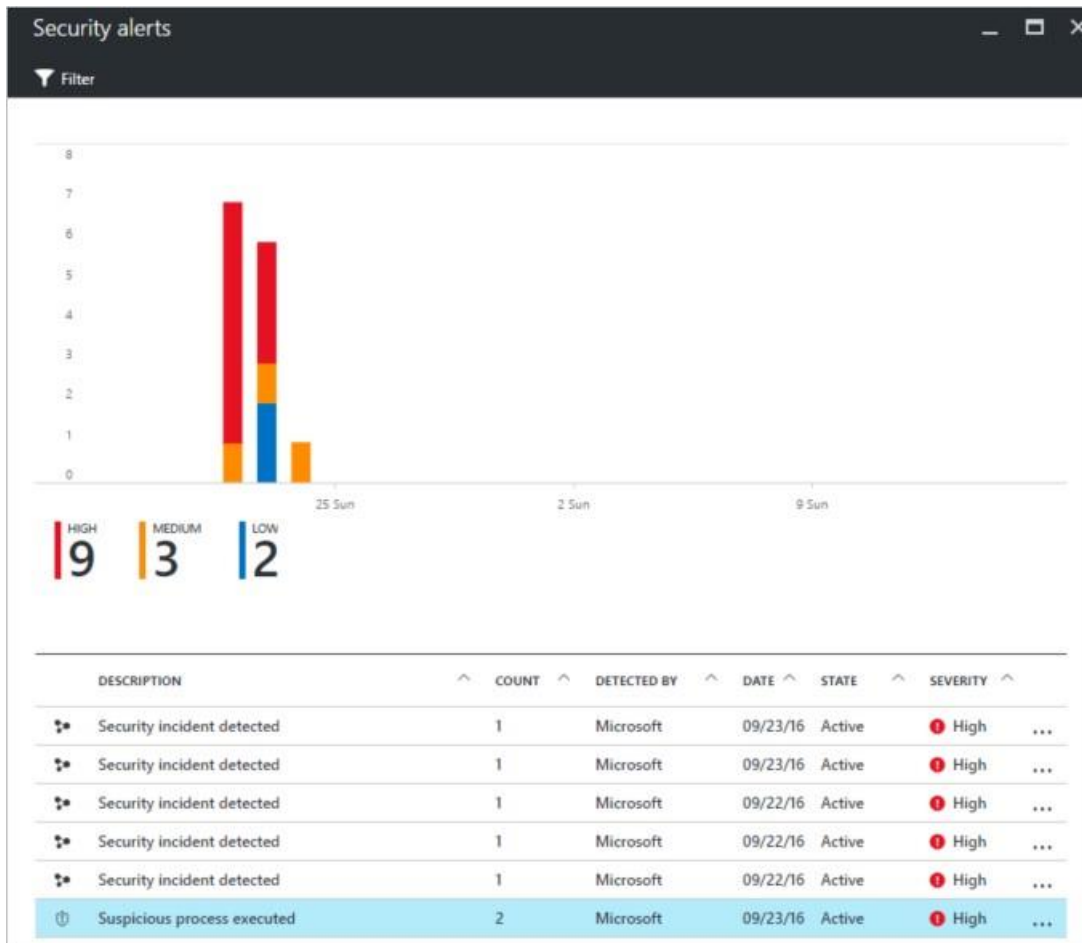
- **Activity Group Report:** provides deep dives into attackers, their objectives and tactics.
- **Campaign Report:** focuses on details of specific attack campaigns.
- **Threat Summary Report:** covers all of the items in the previous two reports.

This type of information is very useful during the [incident response](#) process, where there is an ongoing investigation to understand the source of the attack, the attacker's motivations, and what to do to mitigate this issue moving forward.

How to access the threat intelligence report?



You can review your current alerts by looking at the **Security alerts** tile. Open the Azure Portal and follow the steps below to see more details about each alert:

1. On the Security Center dashboard, you will see the **Security alerts** tile.
2. Click the tile to open the **Security alerts** blade that contains more details about the alerts and click in the security alert that you want to obtain more information about.

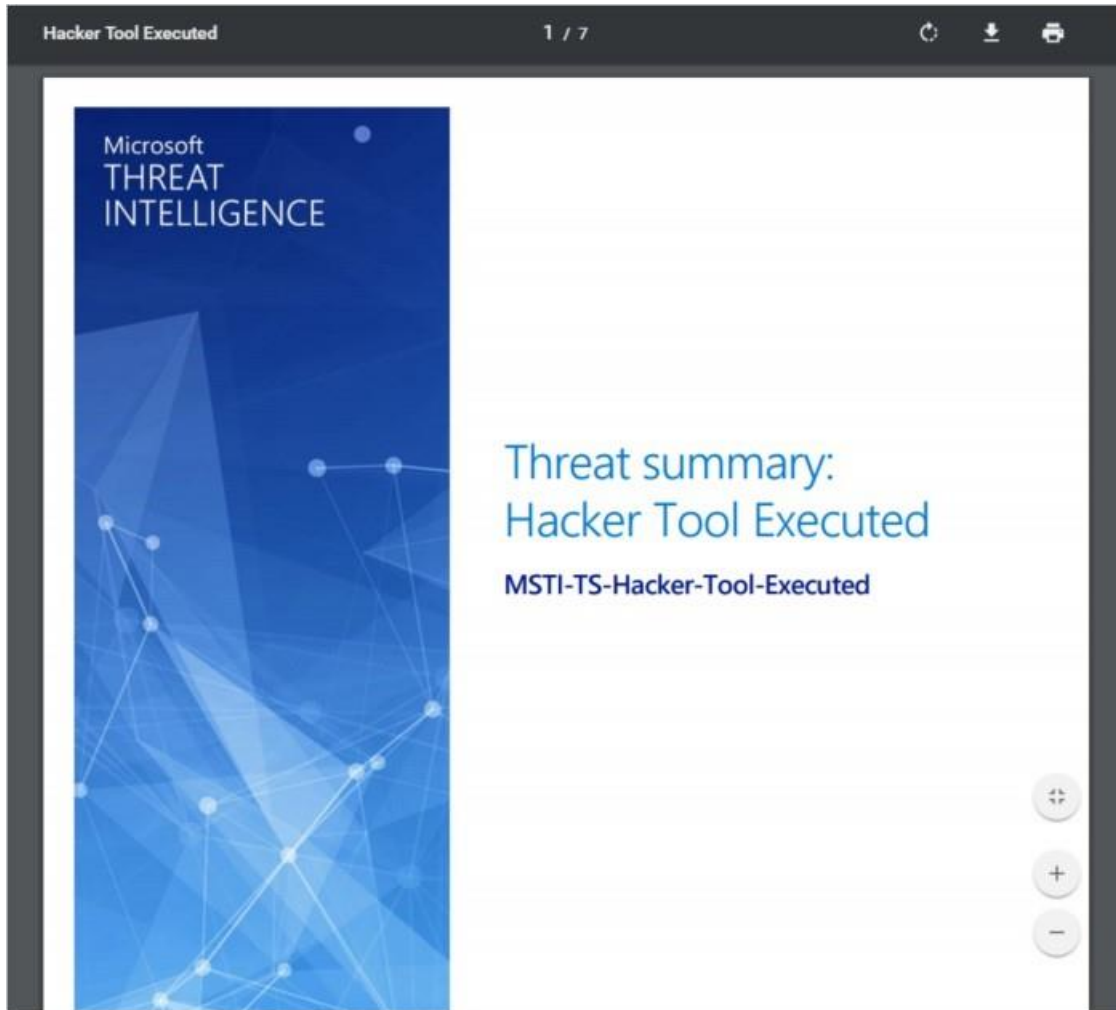


3. In this case the **Suspicious process executed** blade shows the details about the alert as shown in the figure below:

Suspicious process executed
IgniteOMS

DESCRIPTION	Machine logs indicate that the suspicious Process: 'C:\Temp\mimikatz.exe' was running with the command line: "C:\Temp\mimikatz.exe"
DETECTION TIME	Friday, September 23, 2016, 2:56:22 PM
SEVERITY	 High
STATE	Active
ATTACKED RESOURCE	IgniteOMS
SUBSCRIPTION	[REDACTED]
DETECTED BY	 Microsoft
ACTION TAKEN	Detected
DOMAIN NAME	IGNITEOMS
USER NAME	yuri
PARENT PROCESS	-
PROCESS ID	0xc1c
USER SID	S-1-5-21-1997640134-403717899-730413893-500
REPORTS	Report: Hacker tool executed
REMEDATION STEPS	<ol style="list-style-type: none"> 1. Run Process Explorer and try to identify unknown running processes (see https://technet.microsoft.com/en-us/sysinternals/bb896653.aspx) 2. Escalate the alert to the information security team 3. Make sure the machine is completely updated and has an updated anti-malware application installed 4. Run a full anti-malware scan and verify that the threat was removed 5. Install and run Microsoft's Malicious Software Removal Tool (see https://www.microsoft.com/en-us/download/malicious-software-removal-tool-details.aspx) 6. Run Microsoft's Autoruns utility and try to identify unknown applications that are configured to run at login (see https://technet.microsoft.com/en-us/sysinternals/bb896653.aspx)

4. The amount of information available for each security alert will vary according to the type of alert. In the **REPORTS** field you have a link to the threat intelligence report. Click on it and another browser window will appear with PDF file.

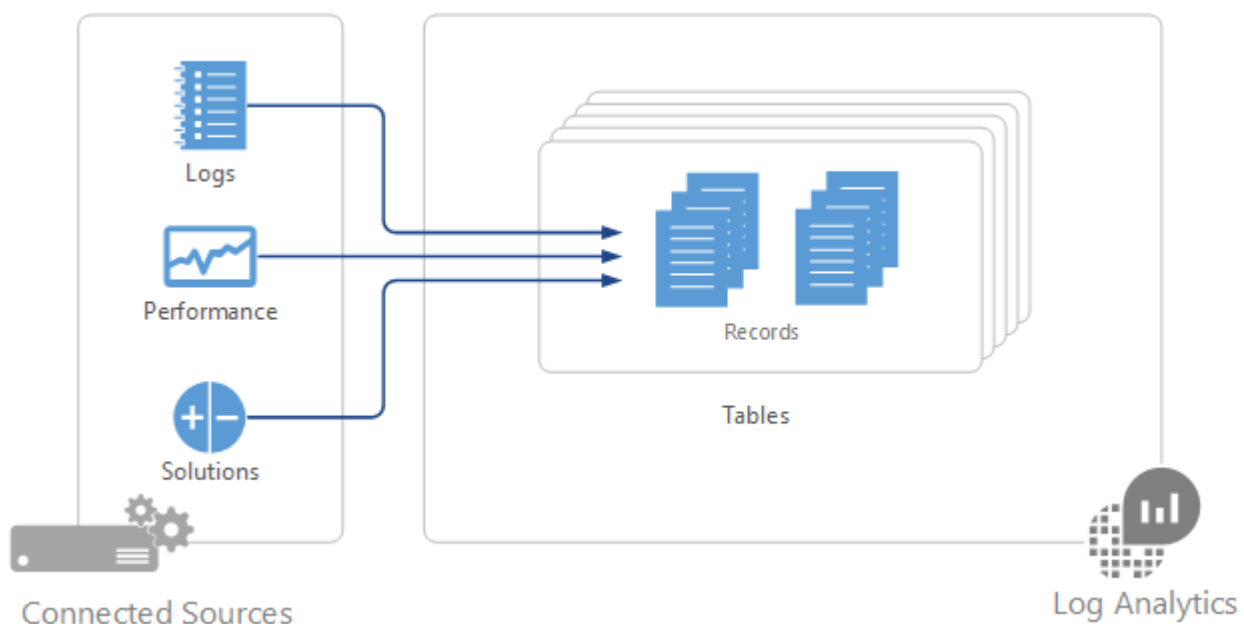


In the Blue Subscription please review the Windows registry persistence method detected Alert for this report.

Audit, Query, Report Scenario – Blue Subscription and Red Subscription

Security Center Audit Rules – Red Subscription

Log Analytics is used for most Audit scenarios and collects data from your Connected Sources and stores it in your Log Analytics workspace. The data that is collected from each is defined by the Data Sources that you configure. Data in Log Analytics is stored as a set of records. Each data source creates records of a particular type with each type having its own set of properties.



Data Sources are different than [management solutions](#), which also collect data from Connected Sources and create records in Log Analytics. In addition to collecting data, solutions typically include log searches and views to help you analyze the operation of a particular application or service.

Summary of data sources

The following table lists the data sources that are currently available in Log Analytics. Each has a link to a separate article providing detail for that data source. It also provides information on their method and frequency of data collection into Log Analytics. You can use the information in this article to identify the different solutions available and to understand the data flow and connection

requirements for different management solutions. For explanations of the columns, see [Data collection details for management solutions in Azure](#).

Data source	Platform	Microsoft monitoring agent	Operations Manager agent	Azure storage	Operations Manager required?	Operations Manager agent data sent via management group	Collection frequency
Custom logs	Windows	•					on arrival
Custom logs	Linux	•					on arrival
IIS logs	Windows	•	•	•			depends on Log File Rollover setting
Performance counters	Windows	•	•				as scheduled, minimum of 10 seconds
Performance counters	Linux	•					as scheduled, minimum of 10 seconds
Syslog	Linux	•					from Azure storage: 10 minutes ; from agent: on arrival

Data source	Platform	Microsoft monitoring agent	Operations Manager agent	Azure storage	Operations Manager required?	Operations Manager agent data sent via management group	Collection frequency
Windows Event logs	Windows	•	•	•		•	on arrival

Configuring data sources

You configure data sources from the **Data** menu in Log Analytics **Advanced Settings**. Any configuration is delivered to all connected sources in your workspace. You cannot currently exclude any agents from this configuration.

Advanced settings

Refresh Analytics Save Discard

Connected Sources >

Data >

Computer Groups >

- Windows Event Logs >
- Windows Performance Counters >
- Linux Performance Counters >
- IIS Logs >
- Custom Fields >
- Custom Logs >
- Syslog >

Collect events from the following event logs

Enter the name of an event log to monitor +

LOG NAME	ERROR	WARNING	INFORMATION	
Application	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Remove
Operations Manager	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Remove
System	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	Remove

1. In the Azure portal, select **Log Analytics** > your workspace > **Advanced Settings**.
2. Select **Data**.
3. Click on the data source you want to configure, such as Windows Event Logs
4. Add System Log Error, Warning, Information
5. Add Application Log Error, Warning, Information.

Data collection

Data source configurations are delivered to agents that are directly connected to Log Analytics within a few minutes. The specified data is collected from the agent and delivered directly to Log Analytics at intervals specific to each data source. See the documentation for each data source for these specifics.

For System Center Operations Manager agents in a connected management group, data source configurations are translated into management packs and delivered to the management group every 5 minutes by default. The agent downloads the management pack like any other and collects the specified data. Depending on the data source, the data will be either sent to a management server which forwards the data to the Log Analytics, or the agent will send the data to Log Analytics without going through the management server. See [Data collection details for management solutions in Azure](#) for details. You can read about details of connecting Operations Manager and Log Analytics and modifying the frequency that configuration is delivered at [Configure Integration with System Center Operations Manager](#).

If the agent is unable to connect to Log Analytics or Operations Manager, it will continue to collect data that it will deliver when it establishes a connection. Data can be lost if the amount of data reaches the maximum cache size for the client, or if the agent is not able to establish a connection within 24 hours.

Log Analytics records

All data collected by Log Analytics is stored in the workspace as records. Records collected by different data sources will have their own set of properties and be identified by their **Type** property. See the documentation for each data source and solution for details on each record type.

Azure Security Center Search – Blue Subscription

Azure Security Center uses [Log Analytics search](#) to retrieve and analyze your security data. Log Analytics includes a query language to quickly retrieve and consolidate data. From Security Center, you can leverage Log Analytics search to construct queries and analyze collected data.

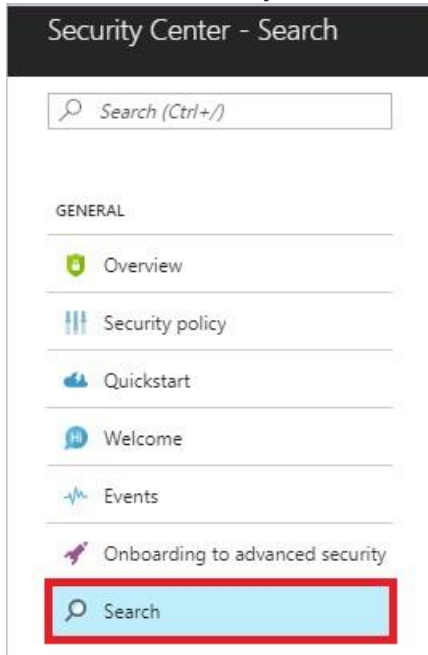
Search is available in both the Free tier and Standard tier of Security Center. The data available in your log searches is dependent on the tier level applied to your workspace. See the Security Center [pricing page](#) for more information.

NOTE

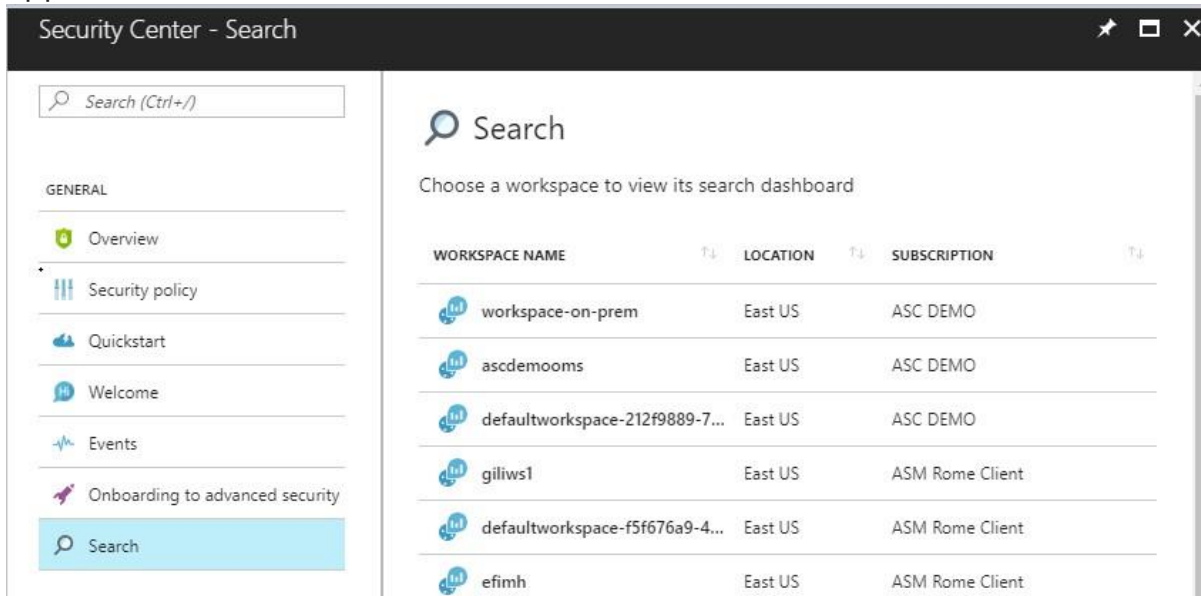
Security Center does not save security data for a workspace under the Free tier. You can send a variety of logs to a workspace under the Free tier and search on that data but search results do not include data from Security Center. Security Center only saves data to a workspace under the Standard tier.

Access search

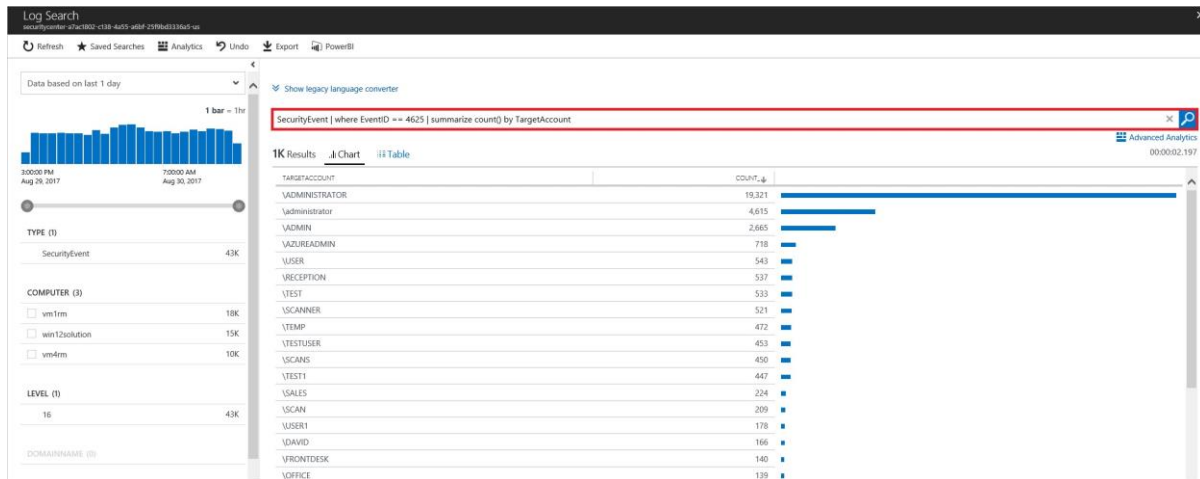
1. Under the Security Center main menu, select **Search**.



2. Security Center lists all workspaces under your Azure subscriptions. Select a workspace. (If you have only one workspace, this workspace selector does not appear.)



3. **Log Search** opens. To query for more data under the selected workspace, enter this example query: `SecurityEvent | where EventID == 4625 | summarize count() by TargetAccount` Result shows all accounts that failed to logon (event 4625).



4. Results can also be EXPORTED or you can use PowerBI to display results.

See [Log Analytics query language](#) for more information on how to query for data under the selected workspace.

Security Center Alert Rules – Red Subscription

What are custom alert rules in Security Center?

Security Center has a set of predefined [security alerts](#), which are triggered when a threat, or suspicious activity takes place. In some scenarios, you may want to create a custom alert to address specific needs of your environment.

Custom alert rules in Security Center allow you to define new security alerts based on data that is already collected from your environment. You can create queries, and the result of these queries can be used as criteria for the custom rule, and once this criteria is matched, the rule is executed. You can use computers security events, partner's security solution logs or data ingested using APIs to create your custom queries.

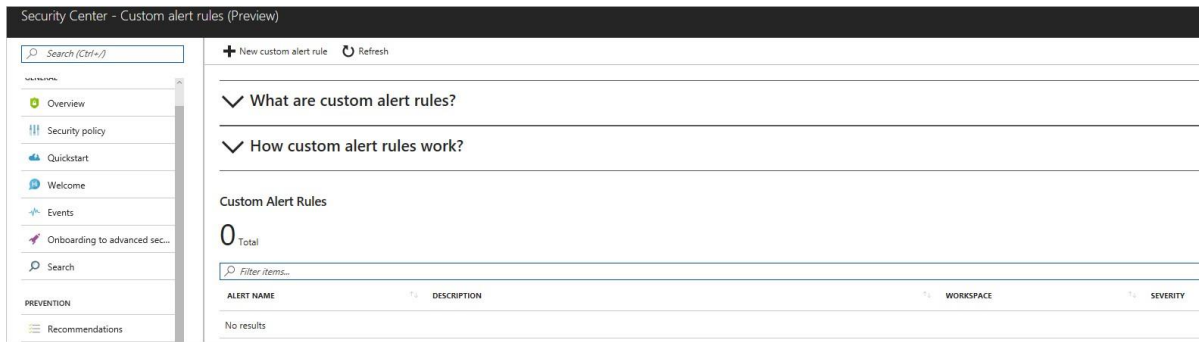
NOTE

Custom alerts are not supported in Security Center's [investigation feature](#).

How to create a custom alert rule in Security Center?

Open **Security Center** dashboard, and follow these steps to create a custom alert rule:

1. In the left pane, under **Detection** click **Custom alert rules (Preview)**.
2. In the **Security Center – Custom alert rules (Preview)** page click **New custom alert rule**.



3. The Create custom alert rule page appears with the following options:

Create custom alert rule
⊞ ×

*** Name** ⓘ

Description

Severity ⓘ

Medium
▼

Sources

Subscription

Contoso IT - demo
▼

Workspace

contosoretail-it
▼

Criteria

*** Search Query** ⓘ

Execute your search query now

Period ⓘ

Over the last 1 hours
▼

Evaluation

Evaluation Frequency

Every 1 hours
▼

Generate alert based on

Number of results

Greater than
▼

*** Threshold**

Suppress Alerts ⓘ

*** Suppress alerts for (in minutes)**

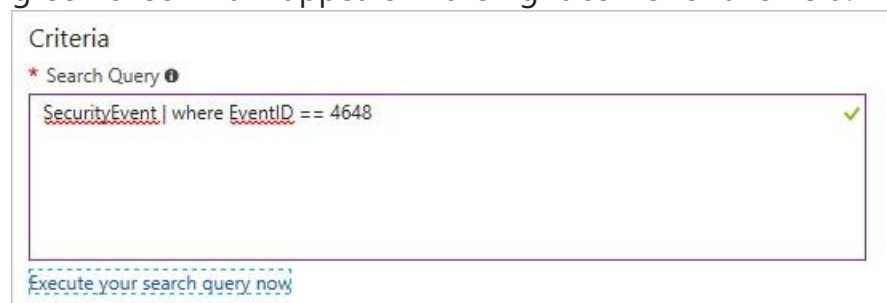
OK

4. Type the name for this custom rule in the **Name** field.
5. Type a brief description that reflects the intent of this rule in the **Description** field.
6. Select the severity level (High, Medium, Low) according to your needs in the **Severity** field.
7. Select the subscription in which this rule is applicable in the **Subscription** field.
8. Select the workspace that you want to monitor with this rule in the **Workspace** field, and in the **Search Query** field, the query that you want to use to obtain the results.

NOTE

You need write permission in the workspace that you select to store your custom alert.

The query's result triggers the alert. Notice that when you type a valid query, the green check mark appears in the right corner of this field:



Criteria

* Search Query ⓘ

SecurityEvent | where EventID == 4648 ✓

Execute your search query now

9. Select the time span in which the query above will be executed in the **Period** field. Notice that the search result in the bottom of this field will change according to the time span that you select.







Period ⓘ

Over the last 1 hours ▼

Your search returned 2 results for the time window selected.

10. In the **Evaluation** field select the frequency that this rule should be evaluated and executed.
11. In the **Number of results** field, select the operator (greater than, or lower than).
12. In the **Threshold** field type a number that will be used as reference for the operator that was previously selected.
13. **Enable Suppress Alerts** option if you want to set a time to wait before Security Center sends another alert for this rule.
14. Click **OK** to finish.

After you finish creating the new alert rule, it will appear in the list of custom alert rules. Once the conditions of that rule are met, a new alert will be triggered, and you can see in the **Security Alerts** dashboard.

AvihaTest2	
Various	
DESCRIPTION	wireData
DETECTION TIME	Friday, September 8, 2017, 1:05:40 PM
SEVERITY	 High
STATE	Active
ATTACKED RESOURCE	Various
SUBSCRIPTION	
DETECTED BY	Alert Rule
ENVIRONMENT	 Non-Azure
RESOURCE TYPE	 Non-Azure Resource
SEARCH QUERY	search ((Type == 'WireData' and Direction == 'Outbound') or (Type == 'WindowsFirewall' and CommunicationDirection == 'SEND') or (Type == 'CommonSecurityLog' and CommunicationDirection == 'Outbound')) and isnotempty(MaliciousIP) summarize by MaliciousIP
SEARCH QUERY RESULT COUNT	0
THRESHOLD OPERATOR	Less Than
THRESHOLD VALUE	1000
QUERY INTERVAL IN MINUTES	5
SUPPRESSION IN MINUTES	120

Notice that the parameters (search query, threshold, etc.) that were established during the rule creation are available in the alert for this custom rule.