

Microsoft Azure Security Center

Second Edition



Yuri Diogenes
Dr. Thomas W. Shinder

Foreword by Hayden Hainsworth, Director of Engineering—Program Management—Microsoft Cybersecurity Engineering

Microsoft Azure Security Center

Second Edition

Yuri Diogenes
Dr. Thomas W. Shinder

Microsoft Azure Security Center

Second Edition

Published with the authorization of Microsoft Corporation by:
Pearson Education, Inc.

Copyright © 2020 by Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearsoned.com/permissions/. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

ISBN-13: 978-0-13-575203-6

ISBN-10: 0-13-575203-5

Library of Congress Control Number: 2019946956

ScoutAutomatedPrintCode

TRADEMARKS

Microsoft and the trademarks listed at <http://www.microsoft.com> on the "Trademarks" webpage are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

WARNING AND DISCLAIMER

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The author(s), the publisher, and Microsoft Corporation shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

SPECIAL SALES

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

EDITOR-IN-CHIEF

Brett Bartow

EXECUTIVE EDITOR

Loretta Yates

DEVELOPMENT EDITOR

Rick Kughen

MANAGING EDITOR

Sandra Schroeder

SENIOR PROJECT EDITOR

Tracey Croom

COPY EDITOR

Rick Kughen

INDEXER

Valerie Perry, Happenstance
Type-O-Rama

PROOFREADER

Vanessa Ta

TECHNICAL EDITOR

Mike Martin

ASSISTANT SPONSORING EDITOR

Charvi Arora

EDITORIAL ASSISTANT

Cindy Teeters

COVER DESIGNER

Twist Creative, Seattle

COMPOSITOR

Jeff Lytle, Happenstance
Type-O-Rama

GRAPHICS

Richard Sheppard, Happen-
stance Type-O-Rama

Contents at a Glance

	<i>Foreword</i>	<i>xv</i>
	<i>Introduction</i>	<i>xvii</i>
CHAPTER 1	The threat landscape	1
CHAPTER 2	Introduction to Azure Security Center	21
CHAPTER 3	Policy management	47
CHAPTER 4	Strengthen your security posture	55
CHAPTER 5	Accessing security alerts from API	93
CHAPTER 6	Threat detection	111
CHAPTER 7	Protect your IoT solution with Azure Security Center	135
CHAPTER 8	SIEM integration	145
CHAPTER 9	Integration with other Microsoft solutions	171
CHAPTER 10	Accessing Security Center Alerts from API	179
APPENDIX A	Managing Security Center at Scale	189
	<i>Index</i>	<i>195</i>

Contents

	<i>Foreword</i>	<i>xv</i>
	<i>Introduction</i>	<i>xvii</i>
Chapter 1	The threat landscape	1
	Understanding cybercrime.....	1
	Understanding the cyber kill chain	2
	Common threats	4
	Building a security posture	5
	Assume breach mentality	6
	Cloud threats and security	6
	Compliance	8
	Risk management	9
	Identity and access management	9
	Operational security	9
	Endpoint protection	10
	Data protection	10
	Azure Security.....	11
	VM protection	12
	Network protection	12
	Storage protection	15
	Identity	17
	Logging	17
Chapter 2	Introduction to Azure Security Center	21
	Deployment scenarios	21
	Understanding Azure Security Center	22
	Security Center architecture	23
	Security Center dashboard	26
	Planning adoption.....	27
	Roles and permissions	27
	Centralized management	28

Storage	29
Recommendations	29
Automation	29
Incorporating Security Center into your security operations	30
Onboarding resources	31
Onboard virtual machines located on-premises or in a different cloud provider	36
Creating workspace	40
E-mail notifications	43
Onboard resources using PowerShell	44
Chapter 3 Policy Management	47
Integration with Azure Policy	47
Fine-tuning security policies	49
Chapter 4 Strengthen your security posture	55
Secure Score	55
Fine-tuning your Secure Score	57
Compute and apps recommendations	59
Add a web application firewall	60
Endpoint protection	62
Networking recommendations	66
Network map	67
Adaptive network hardening	68
Enable Network Security Groups on subnets	71
Restrict access through Internet-facing endpoint	72
Data and storage	74
Enable auditing on SQL Server	75
Disable unrestricted network access to storage account	77
Identity and access	79
Subscriptions	84
Key vaults	85
App services	86
Containers	87
Virtual machine scale sets	88

Chapter 5	Advanced cloud defense	93
	Just-in-Time Virtual Machine Access	93
	Configuring JIT	96
	Requesting access	98
	File integrity monitoring	100
	Customizing your settings	100
	Visualizing changes	104
	Application whitelisting (adaptive application controls)	106
Chapter 6	Threat detection	111
	Methods of threat detection	111
	Understanding alerts	112
	Accessing security alerts	114
	Detection scenarios	116
	Detecting spam activity	116
	Crash dump analysis	117
	Windows detections	118
	Linux detections	119
	Containers detections	120
	Azure SQL detections	120
	Azure App Services detections	121
	Azure Resource Manager detections	122
	The cyber kill chain and fusion alerts	122
	Automating response with playbooks	127
	Creating a playbook	127
	Building the workflow	129
	Executing a playbook	131
	Auditing playbook execution	133
Chapter 7	Protect your IoT Solution with Azure Security Center	135
	Understanding Security Center for IoT solution	135
	Configuring Security Center for IoT	138
	Security recommendations	140
	Monitoring alerts	142
	Investigating a suspicious activity	143

Chapter 8	SIEM integration	145
	Streaming logs to a SIEM solution	145
	Azure Sentinel.....	146
	Integration with Azure Sentinel	149
	Accessing alerts in Azure Sentinel	151
	Integration with Splunk	154
	Confirming accessible logs in Azure Monitor	155
	Configuring the subscription for the Splunk SIEM pipe	155
	Creating and configuring a resource group for the Splunk SIEM pipe	155
	Setting up an Azure AD application to provide an access control identity	156
	Creating an Azure key vault	158
	Copying the app password into Key Vault	161
	Making an event hub	162
	Creating a shared access key for event hub access control	163
	Placing the event hub shared access key in Azure Key Vault	165
	Connecting the event hub to Azure Monitor	166
	Provision Virtual Machine used by Splunk	167
	Installing and configuring the Azure Monitor add-on for Splunk	169
Chapter 9	Integration with other Microsoft Solutions	171
	Increasing threat detection scope	171
	Integration with Windows Server Defender ATP	172
	Integration with Microsoft Cloud App Security	174
	Integration SQL ATP	176
	Integration with ATP for Azure Storage	177
Chapter 10	Accessing security alerts from API	179
	Understanding REST API	179
	Accessing alerts using Security Center REST API	180
	Accessing alerts using the Security Graph API	184
	Using the Security Graph API.....	186

Appendix A Managing Security Center at scale	189
The importance of management at scale	189
The three cornerstones	189
Security Center and Azure Policy, better together	192
Scoping your deployment of ARM templates and Azure Policy	193
Best practices for managing Security Center at scale	194
<i>Index</i>	195

Acknowledgments

The authors would like to thank Loretta Yates and the entire Microsoft Press/Pearson team for their support on this project and Hayden Hainsworth for writing the foreword. We'd also like to thank the other Microsoft colleagues who contributed by writing sidebars for this book: Nicholas DiCola, Oren Parag, Koby Koren, Sarah Fender, Ram Pliskin, Laura Machado de Wright, Miri Landau, Ben Kliger, Michelle Swafford, Ajeet Prakash, Anupam Vij, Michelle Swafford, Milad Aslaner, John Kernetz, Ben Nick, Laura Hunter, and Michael Bargury. Thanks to Tiander Turpijn for the great work writing Appendix A. We would also like to thank Mike Martin for reviewing this book.

Yuri would also like to thank: my wife and daughters for their endless support; my great God for giving me strength and guiding my path on each step of the way; my great friend and co-author Tom Shinder for another awesome partnership; the Azure Security Center Engineering / Dev teams (Gilad Elyasha, Meital Taran-Gutman, Itamar Falcon, Maya Herskovic, Ron Matchoro, Ronit Reger, Rotem Lurie, Yoav Francis, Yaniv Shasha, Mor Weinberger, Yaniv Zohar, Dotan Patrich, Arik Noyman, Yoav Daniely, Lior Becker, Eli Sagie, Liron Kachko, Omer Chechik, Eli Sagie, Vlad Korsunsky, Jonathan Gazit, and Tamer Salman). Thanks for the ongoing collaboration and contribution. Thanks to my manager, Nicholas DiCola, and the entire CxE ASC team for the great partnership. Last but not least, thanks to my parents for working hard to give me an education, which is the foundation I use every day to keep moving forward in my career.

Tom would also like to thank so many people that it's very difficult to name them all in the space allocated. Probably most important is Yuri Diogenes, who motivated me to partner up on another book with him. I don't know why he asks me because I know I drive him crazy each time we write a book together. Nevertheless, Yuri is a blessing to me and all those around him, and he keeps me from resting on my prodigious laurels. I want to thank David Cross, who brought me into Azure Security Engineering and all the fascinating opportunities it has provided. Major props go to Avi Ben-Menahem, who has always been supportive of my efforts and who is a model of the best that Microsoft has to offer. And never to be forgotten, I'm delivering sizeable acknowledgements to Karen Dahmen and Laurie Litwack, two of the best managers anyone could ever have and an inspiration to all they work with. Let's also stand up and give a huge shout out to the entire Azure Security Engineering PM team—the dedication, diligence, intelligence, and number of hours worked per week by this team is unmatched, and the results of these attributes show in the fact that Azure is the most secure public cloud service platform in the industry. Finally, eternal thanks to my wife—my lifetime love, partner, and confidant—and to God to whom I am grateful for giving me much more in life than I deserve.

About the Authors

Yuri Diogenes, MsC

Yuri Diogenes has a Master of Science in cybersecurity intelligence and forensics investigation (UTICA College) and is a senior program manager for the Microsoft CxE Security team, where he primarily helps customers onboard and deploy Azure Security Center as part of their security operations/incident response. Yuri has been working for Microsoft since 2006 in different positions; he spent five years as a senior support escalation engineer for the CSS Forefront Edge team. From 2011 to 2017, he worked for the content development team where he also helped create the Azure Security Center content experience since its launch in 2016. Yuri has published a total of 21 books, mostly covering information security and Microsoft technologies. Yuri also holds an MBA and many IT/Security industry certifications, such as CISSP, E|CND, E|CEH, E|CSA, E|CHFI, CompTIA Security+, CySA+, Cloud Essentials Certified, Mobility+, Network+, CASP, CyberSec First Responder, MCSE, and MCTS. You can follow Yuri on Twitter at @yuridiogenes or read his articles at his personal blog: <http://aka.ms/yuridio>.

Tom Shinder

Tom Shinder is a cloud security program manager in Azure Security Engineering. He is responsible for Azure security technical content and education, Azure security baselining, and public cloud security competitive analysis. He has presented at many of the largest security industry conferences on topics related to both on-premises and public cloud security and architecture. Tom earned a bachelor's degree in neuropsychobiology from the University of California, Berkeley, and an M.D. from the University of Illinois, Chicago. He was a practicing neurologist prior to changing careers in the 1990s. He has written more than 30 books on OS, network, and cloud security, including Microsoft Azure Security Infrastructure. Tom can be found hugging his Azure console when he's not busy hiding his keys and secrets in Azure Key Vault.

Foreword

I was so pleased to hear that Yuri and Tom were teaming up to write another book on Security. I found their first book about Azure core security, “Microsoft Azure Security Infrastructure” riveting. I read it cover to cover twice, recommending it to anyone interested in learning more about security in Azure.

This book extends that work. It will teach you all you want to know about how to use Azure Security Center—the security solution to get visibility and control and prevent and detect threats in your Azure subscriptions. Security Center is a critical solution for organizations using a cloud workload protection (CWP) solution, as indicated by Gartner in their CWP Magic Quadrant. And because the classic security perimeter we relied upon is gone with the migration of datacenter workloads into public clouds, which are a new security paradigm. Also, the integration with Log Analytics means you can use Azure Security Center for your machines on-premises, in a private datacenter, or in another cloud as long as the monitoring agent is installed on your machines. This will simplify your life, and I trust you’ll come to rely on Azure Security Center as your primary dashboard and alerting engine for years to come.

No industry is immune to cyberattacks. This book is relevant for everyone working with cloud computing and information security. Given the cybersecurity landscape as it exists today and the criticality of the information digital age, we need to assume breach as a mindset and think about what capabilities we used to detect adversarial activity or malicious insiders in our networks rather than over-relying on thwarting attackers at the front door. Gone are the days when it was acceptable to turn a blind eye to risk. If you don’t know what you don’t know, it’s impossible to take action. Prevention is of the utmost importance. However, the ability to detect and control is paramount. Simple, intuitive, and intelligent investigation capabilities are a must to support SecOps teams that are flooded in a sea of alerts, as well as Incident Response (IR) teams.

This book will help you plan, onboard, and learn how to effectively use Security Center to detect and investigate threats in your Azure subscriptions (or alternately your datacenter workloads). You’ll also learn how to integrate with other solutions like Azure Active Directory Identity Protection Center. You’ll also learn how to export your logs to a SIEM should you choose to do so. I sincerely hope that you are energized by this book and that you will be spurred to action by following its best practices and recommendations.

After reading this book, you will have a better understanding of what Security Center is and how to incorporate it into your security operations center. Yuri and Tom were inspired to write this book because many customers have asked for a one-stop resource that teaches you how to install and operate Security Center. This book is written with the security analysts, architects, cloud operators, and IT professionals in mind.

If you've read Tom's work previously, you'll know he's a long-term, experienced, and seasoned security veteran and author. He's also a senior program manager on the Azure Security Engineering team. Follow him on Twitter and read his blogs if you'd like to learn more. He's a wealth of knowledge and wisdom from his days working on-premises and his journey to the cloud.

Yuri is another well-established writer in his own right, and he has published document after document in his former role as a content writer for Azure Security. He's recently joined my team working as a senior program manager, supporting customers' and partners' success using Microsoft's cloud and enterprise security products and services. He's a wealth of information and excels at simplifying the complex.

Dig in.

Hayden Hainsworth
Director of Engineering – Program Management
Microsoft Cybersecurity Engineering

Introduction

Welcome to *Microsoft Azure Security Center*, a book that was developed together with the Azure Security Center product group to provide in-depth information about Azure Security Center and to demonstrate best practices based on real-life experience with the product in different environments.

The purpose of this book is to introduce the wide array of security features and capabilities available in Azure Security Center. After being introduced to all these security options, you will dig in to see how they can be used in a number of operational security scenarios so that you can get the most out of the protect, detect, and respond skills provided only by Azure Security Center.

Who is this book for?

Microsoft Azure Security Center is for anyone interested in Azure security: security administrators, support professionals, developers, and engineers.

Microsoft Azure Security Center is designed to be useful for the entire spectrum of Azure users. You can have no security experience, some experience, or be a security expert and get value from Azure Security Center. This book provides introductory, intermediate, and advanced coverage on a large swath of security issues that are addressed by Azure Security Center.

The approach is a unique mix of didactic, narrative, and experiential instruction. Didactic covers the core introductions to the services. The narrative leverages what you already understand, and we bridge your current understanding with new concepts introduced in the book.

Finally, the experience component is presented in two ways—we share our experiences with Azure Security Center and how to get the most out of it by showing in a stepwise, guided fashion how to configure Azure Security Center to gain all the benefits it has to offer.

In this book, you will learn:

- How to secure your Azure assets no matter what your level of security experience
- How to save hours, days, and weeks of time by removing the trial and error
- How to protect, detect, and respond to security threats better than ever by knowing how to get the most out of Azure Security Center

System requirements

- Anyone with access to a Microsoft Azure subscription can use the information in this book.

Errata, updates & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

MicrosoftPressStore.com/Azuresecuritycenter/errata

If you discover an error that is not already listed, please submit it to us at the same page.

For additional book support and information, please visit
<https://MicrosoftPressStore.com/Support>.

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to *<http://support.microsoft.com>*.

Stay in touch

Let's keep the conversation going! We're on Twitter: *<http://twitter.com/MicrosoftPress>*.

Strengthen your security posture

Ransomware outbreaks, such as WannaCry and Petya, reinforced the importance of having a vulnerability-management system. However, these outbreaks also reinforced the fact that many computers are not fully updated and do not use the most secure configuration.

To enhance your overall security posture, you need to increase your protection. A security assessment is critical to identifying the current security state of your assets—and what you need to do to improve it. Azure Security Center can perform a security assessment for all major workloads: compute, network, storage, and applications. The result of this security assessment is a set of recommendations that will help you enhance the security posture of your workloads.

In this chapter, you will learn how to use Security Center to perform a security assessment for major workloads and how to use the result of this assessment to improve your defense system.

Secure Score

When working in a cloud environment, monitoring the security state of multiple workloads can be challenging. How do you know if your security posture across all workloads is at the highest-possible level? Are there any security recommendations that you are not meeting? These are hard questions to answer when you don't have the right visibility and tools to manage the security aspects of your cloud infrastructure.

Security Center reviews your security recommendations across all workloads, applies advanced algorithms to determine how critical each recommendation is, and calculates your Secure Score based on them. Secure Score helps you to assess your workload security posture from a single dashboard. You can view the overall Secure Score in the Overview page in Security Center dashboard, as shown in Figure 4-1.




FIGURE 4-1 Overall Secure Score of your workloads in Azure

The overall Secure Score shown in the main dashboard is an accumulation of all your recommendation scores. Keep in mind that this score can vary because it reflects the subscription that is currently selected and the resources that belong to that subscription. If you have multiple subscriptions selected, the calculation will be for all subscriptions. The active recommendations on the selected subscription also make this score change. For the example, as shown previously in Figure 4-1, the current Secure Score of this subscription is 397 out of 570. This means that to achieve 570, it is necessary to address all current recommendations. To access more details about your Secure Score, click the **Review Your Secure Score** option in the **Secure Score** tile (see Figure 4-2).

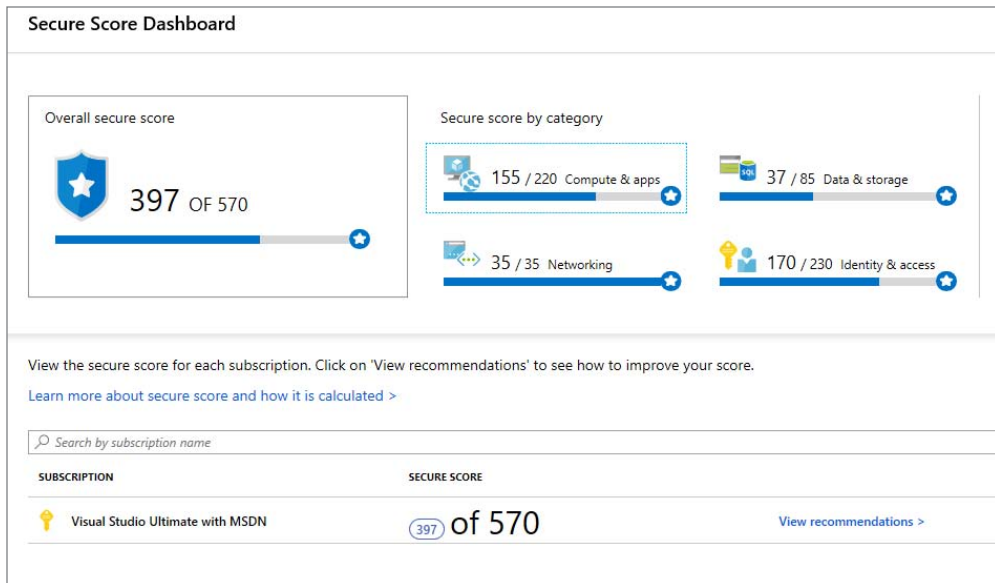


FIGURE 4-2 Details about your current Secure Score

From this dashboard, you have a better visualization of how your workloads impact your overall Secure Score. The example shown in Figure 4-2 has an interesting breakdown because the **Networking** workload is fully compliant, but the other workloads are still a long way from being fully compliant. From here on, you can either click on each workload to see the recommendations or click view recommendations on the subscription. For this example, click the **View Recommendations** option to see all recommendations, as shown in Figure 4-3.

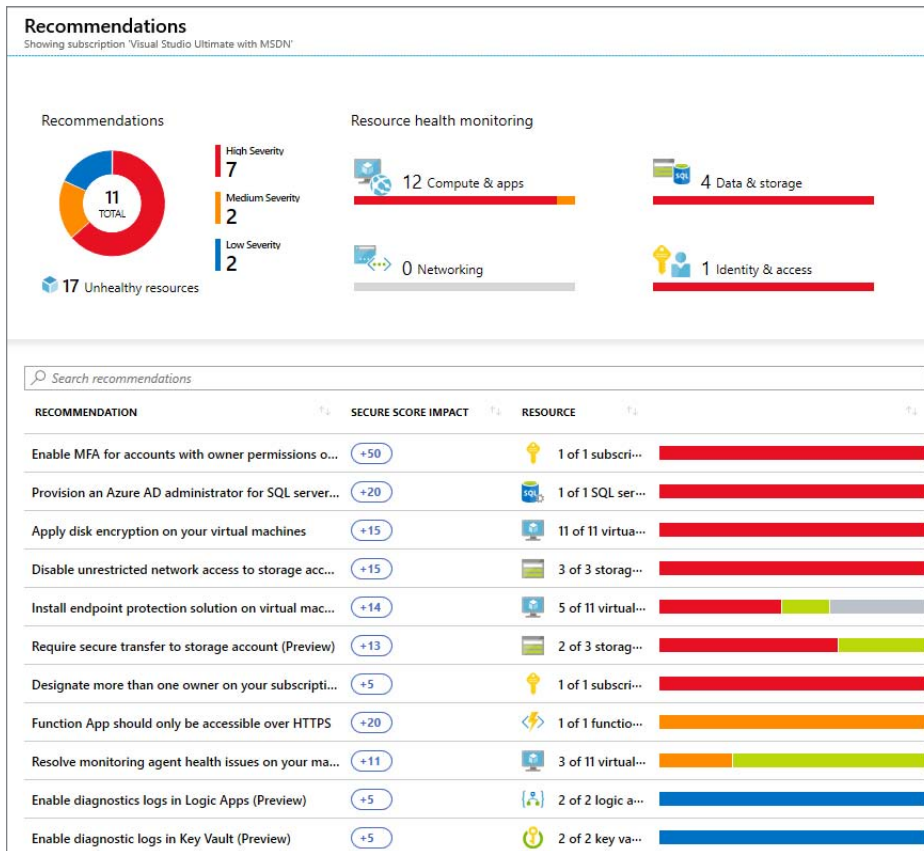


FIGURE 4-3 Recommendations that have a direct impact on the Secure Score

The recommendation Secure Score is a calculation based on the ratio between your healthy resources and your total resources. If the number of healthy resources is equal to the total number of resources, you get the maximum Secure Score of 50. To try to get your Secure Score closer to the max score, fix the unhealthy resources by following the recommendations. Notice that each recommendation has its Secure Score Impact; this number allows you to see how much your Secure Score will be impacted once you address this recommendation. For example, if your Secure Score is 50 and the recommendation impact is +5, performing the steps outlined in the recommendation will improve your score to 55.

Fine-tuning your Secure Score

While Secure Score can be utilized to assist your organization in enhancing its security posture, there will be some scenarios in which not all recommendations are applicable to your environment. It is common to have customers asking to fine-tune those recommendations; they ask because there are items they consider to be false positives.

Organizations commonly use a third-party MFA solution for subscription accounts with owner permissions, and the organizations believe they can safely ignore the *Enable MFA For Accounts With Owner Permissions On Your Subscription* option. However, because the organizations are not addressing this recommendation, there is a 50-point drop in their Secure Score. How can they safely disable this recommendation?

If you are absolutely sure that this recommendation has been addressed by implementing an external factor that is not being taken into consideration by Security Center, you can follow the steps described in **Chapter 3, “Policy Management,”** to disable the policy that reflects your desired recommendation. For the recommendation described in the previous paragraph, you need to choose **Disabled** from the **Monitor MFA for accounts with owner permissions policy** drop-down menu, as shown in Figure 4-4.




FIGURE 4-4 Disabling an Azure Policy to reflect a more accurate list of recommendations

Using Security Center recommendations to drive a better security posture

Misconfigurations, lack of security expertise, and lack of visibility are the most common reasons for attack vulnerabilities. Azure Security Center recommendations are based on security policies by which you choose to have your environment assessed. You can choose which security policies by which you want Azure Security Center to assess your environment and tell you where your environment is vulnerable. Security Center assesses your environment 24/7 and provides recommendations with a Secure Score Impact. This allows you to prioritize the work that your organization needs to do in order to reduce attack surface and harden your security posture.

ASC continuously assesses your environment—in Azure or in a hybrid configuration including on-premises or other clouds—based on the security policies you chose to enable. In addition to providing security recommendations, ASC also maps assessments to several regulatory compliance standards. For each assessment under regulatory compliance, ASC shows how many resources are not passing compliance, and it provides remediation steps. Because of the nature of regulatory compliance assessments, ASC cannot assess all controls. Controls not assessed by ASC are shown on the list as unavailable (dimmed).

In order to be able to manage and prioritize ASC recommendations and the work required to address those recommendations, ASC provides you with an overall Secure Score per subscription as well as the Score Impact for each recommendation. Score Impact is based on the severity and best practices, and it allows you to identify the top recommendations you should address in order to tighten your security posture and protect your environment.

Michelle Swafford, Principal Program Manager, Azure Security Center Team

Compute and apps recommendations

In Chapter 2, “Introduction to Azure Security Center,” you learned about the Security Center agent, and how it performs the initial security assessment. As part of your onboarding process, you should make sure to address all recommendations that have a higher impact on your Secure Score. First, evaluate all other recommendations, and apply the recommendations according to your environment’s needs.

Some recommendations may require system downtime—for example, to apply certain security updates. This means that after you identify the changes that need to be made in the target system, you may need to start a change-control process to maintain compliance with the security assessment.

TIP Recommendations are applicable only for operating systems that are supported in Security Center. Visit the latest version of supported operating systems at <https://aka.ms/ASCSupportedOS>.

Compute and apps recommendations include a collection of recommendations for Azure VMs, non-Azure computers, App Services, Containers, and VM Scale Set (VMSS). New recommendations may be introduced without previous notice, and they appear in the dashboard appended with **(Preview)**, as shown in Figure 4-5.

Enable diagnostics logs in Logic Apps (Preview)
Enable diagnostics logs in Event Hub (Preview)
Enable diagnostic logs in Batch accounts (Preview)
Enable diagnostics logs in Virtual Machine Scale Sets (Preview)
Configure metric alert rules on Batch account (Preview)

FIGURE 4-5 New recommendations shown as (Preview)

For some recommendations, you can apply the remediation directly from the Security Center dashboard; others will require you to go through an external process. For example, the ***Install System Updates On Your Machines*** recommendation will only give you the list of computers that are missing system updates. However, you can't apply the updates from Security Center; you need to use an external solution, such as Windows Update Services (WSUS). Some other recommendations will give you remediation steps, as shown in Figure 4-6.

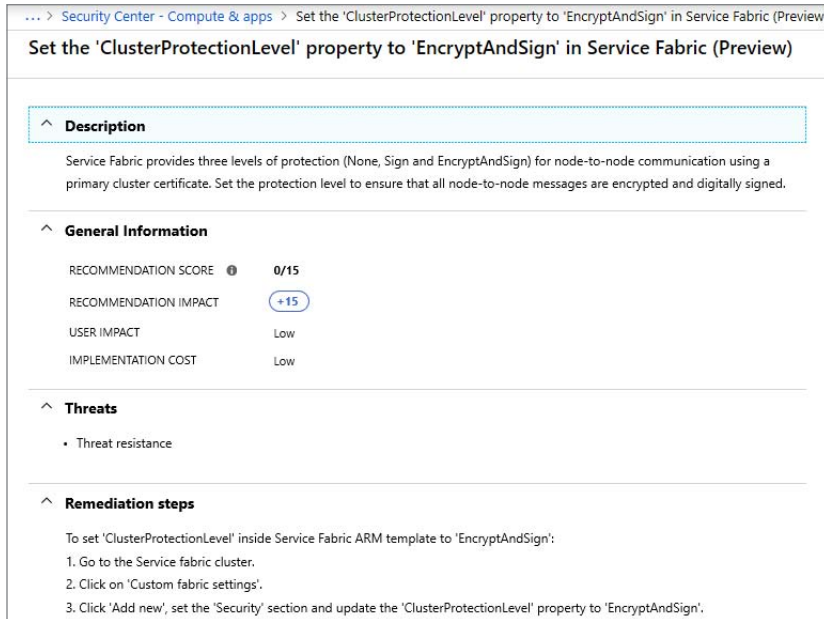


FIGURE 4-6 Remediation steps to be followed to address this recommendation

Many customers will host their web applications in VMs running on Azure, which means they need to ensure that the security states of these VMs are well configured. This should be done not only from the operating system side (which was done under the **Compute** recommendation), but also from the web application level.

The application recommendations will suggest security recommendations for Internet Information Services (IIS) web applications running on Azure VMs. Applications recommendations will vary according to the environment. The next section provides examples of recommendations that can be remediated directly through the Security Center dashboard.

Add a web application firewall

Web applications are increasingly targeted by cyberattacks, such as SQL injection, cross-site scripting (XSS), and many other attacks that are documented at the OWASP Top 10. While many of these attacks are preventable via well-secure code, it is hard to prevent some of these attacks in the application code itself because it may require rigorous maintenance, patching,

and monitoring at multiple layers. A web application firewall can enhance the protection of your web application from web vulnerabilities and attacks without modifying the application code. Follow these steps to address this recommendation:

1. Open **Azure Portal** and sign in with a user who has **Security Admin** privileges.
2. In the left navigation pane, click **Security Center**.
3. In the Security Center left navigation pane, under **Resource Security Hygiene**, click **Compute & Apps**.
4. In the **Overview** tab, type **Firewall** in the search field and click **Add a web application firewall**, as shown in Figure 4-7.



FIGURE 4-7 Application recommendations page

5. Click **Add a web application firewall** and select the systems that you want to remediate by installing the web application firewall. The **Add a web application firewall** blade appears, as shown in Figure 4-8.

NOTE For more information about the OWAS Top 10 Project, see https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project.

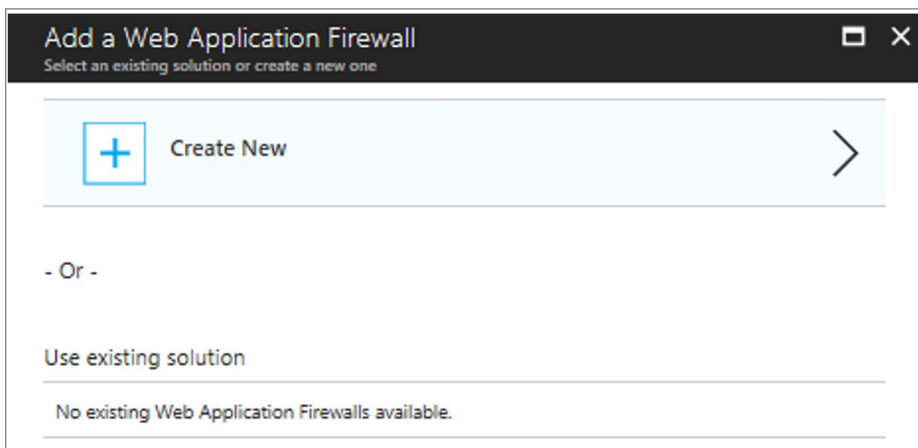


FIGURE 4-8 Adding a new web application firewall

6. Click **Create New**, and the **Create a new Web Application Firewall solution** page appears, as shown in Figure 4-9.

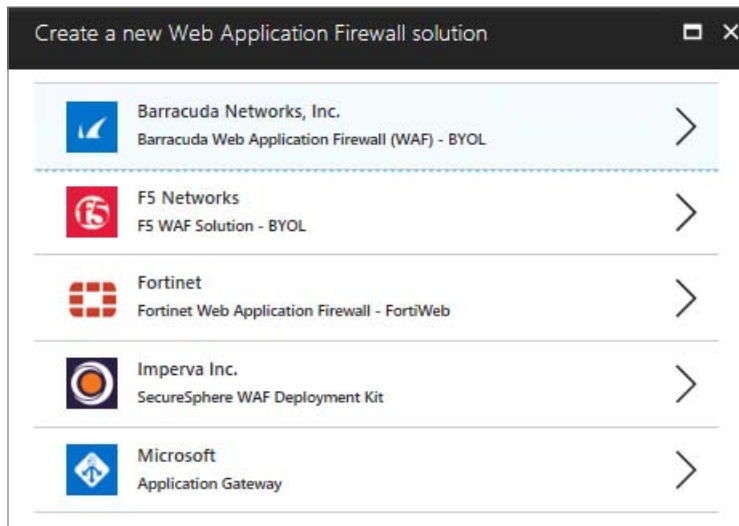


FIGURE 4-9 Adding a new web application firewall

From this point forward, the steps may vary according to the solution that you choose. You should consult the partner's solution documentation for further reference.

Endpoint protection

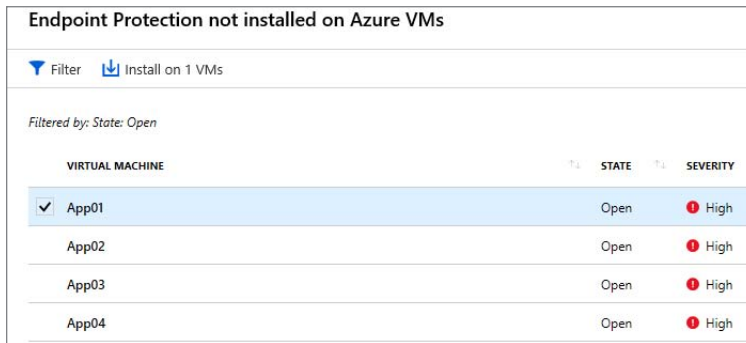
Security Center can detect whether your VM or computer has endpoint protection installed and whether that endpoint protection is up to date. However, this capability will detect only a certain number of supported endpoint protection partners. At the time of this writing, the supported partners are as follows:

- Windows Defender (Microsoft Antimalware)
- System Center Endpoint Protection (Microsoft Antimalware)
- Trend Micro (all versions)
- Symantec v12.1.1100+
- McAfee v10+

TIP The list of partners is always in revision, and new partners may be included without further notice. For the latest list of supported endpoint protection, see <https://aka.ms/ASCPartners>.

If a recommendation suggests the installation of endpoint protection in non-Azure VMs and computers, you will need to do this manually, following the instructions from the anti-malware vendor of your choice. For VMs in Azure, System Center will guide you through the installation process, and you can choose the endpoint protection based on the options available in the Azure marketplace. Follow these steps to remediate a compute recommendation by deploying Windows Defender (Microsoft Antimalware) in an Azure VM:

1. Open **Azure Portal** and sign in with a user that has **Security Admin** privileges.
2. In the left navigation pane, click **Security Center**.
3. In the Security Center left navigation pane, under **Resource Security Hygiene**, click **Compute & Apps**.
4. In the **Overview** tab, click the **Install Endpoint Protection Solution On Virtual Machines** recommendation.
5. In the **Endpoint Protection not installed on Azure VMs** blade, select the VM that you want to install, as shown in Figure 4-10. If you have multiple workspaces, you may see the option to select the workspace first.



Endpoint Protection not installed on Azure VMs		
Filter Install on 1 VMs		
Filtered by: State: Open		
VIRTUAL MACHINE	STATE	SEVERITY
<input checked="" type="checkbox"/> App01	Open	High
<input type="checkbox"/> App02	Open	High
<input type="checkbox"/> App03	Open	High
<input type="checkbox"/> App04	Open	High

FIGURE 4-10 Installing endpoint protection in the target system

6. Click the **Install On 1 VMs** button, and the **Select Endpoint Protection** page will appear, as shown in Figure 4-11.



FIGURE 4-11 The Select Endpoint Protection solutions that are integrated with Security Center

7. For this example, click **Microsoft Antimalware**. (Microsoft Antimalware is a free solution. However, if you choose to install Deep Security Agent from TrendMicro, you need to provide your license information.)

8. The Microsoft Antimalware page appears with the description of this service; click **Create** to add this extension to your VM, and the **Install Microsoft Antimalware** page appears, as shown in Figure 4-12.

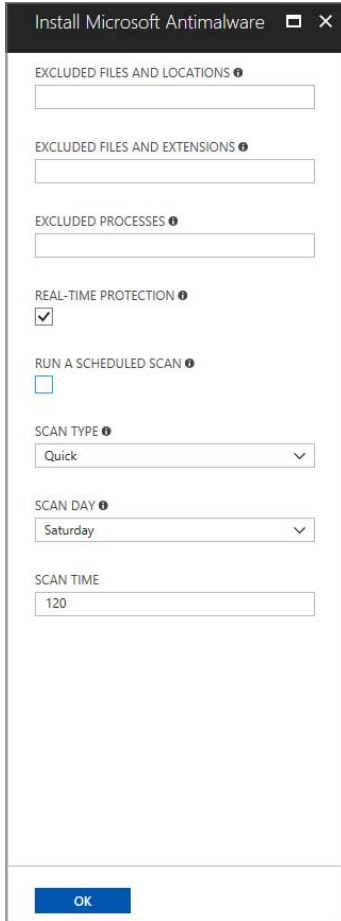


FIGURE 4-12 Options to customize the Microsoft Antimalware installation

The settings shown in the Install Microsoft Antimalware dialog box are explained here:

- **Excluded Files And Locations:** Here, you can specify any paths or locations to exclude from the scan. To add multiple paths or locations, separate them with semicolons. This is an optional setting.
- **Excluded Files And Extensions:** This box lets you specify file names or extensions to exclude from the scan. Again, to add multiple names or extensions, you separate them with a semicolon. Note that you should avoid using wildcard characters.

- **Exclude Processes:** Use this box to specify any processes that should be excluded from the scan—again, using semicolons to separate multiple processes.
 - **Real-Time Protection:** By default, this check box is selected. Unless you have a good business reason to do otherwise, you should leave it that way. By leaving this option selected, the Microsoft Antimalware will monitor events such as
 - Processes that are making unusual changes to existing files
 - Processes that are modifying or creating automatic startup registry keys or startup locations
 - Other changes to the file system or file structure
 - **Run A Scheduled Scan:** Selecting this check box enables you to run a scheduled scan. If you choose to run a Scheduled Scan, you set the Day and Time for the scan in the Scan Day and Scan Time boxes, respectively.
 - **Scan Type:** If you selected the Run A Scheduled Scan check box, you can use this drop-down menu to specify the type of scan. A quick scan is run by default, but you could also have full scan or custom.
 - **Scan Day:** If you selected the Run A Scheduled Scan check box, you can use this drop-down menu to specify the day the scan will run.
 - **Scan Time:** If you selected the Run A Scheduled Scan check box, you can use this drop-down menu to specify what time the scan will run. The time is indicated in increments of 60 minutes (60 = 1 AM, 120 = 2 AM, and so on).
9. Leave the default selections and click the **Create** button. In the upper-right corner of the dashboard, you will see a notification that the installation has started on the target system, as shown in Figure 4-13.




FIGURE 4-13 Notification that the endpoint protection installation has started

10. Close all blades and go back to the main **Security Center** dashboard.
11. To verify whether the Microsoft Antimalware extension was installed, open the VM properties in Azure and click the **Extensions** options (see Figure 4-14).

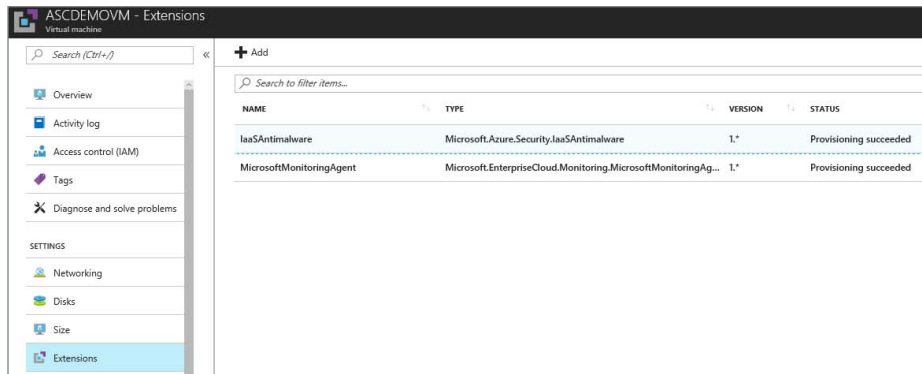


FIGURE 4-14 Validation that the antimalware extension was installed

Networking recommendations

An Azure virtual network is a logical isolation of the Azure cloud dedicated to your subscription. Security Center will identify the Azure virtual networks available in your subscription, and it will provide recommendations to improve the overall security. In the Security Center dashboard, under **Resource Security Hygiene**, click **Networking**, and the **Security Center – Networking** dashboard appears, as shown in Figure 4-15.

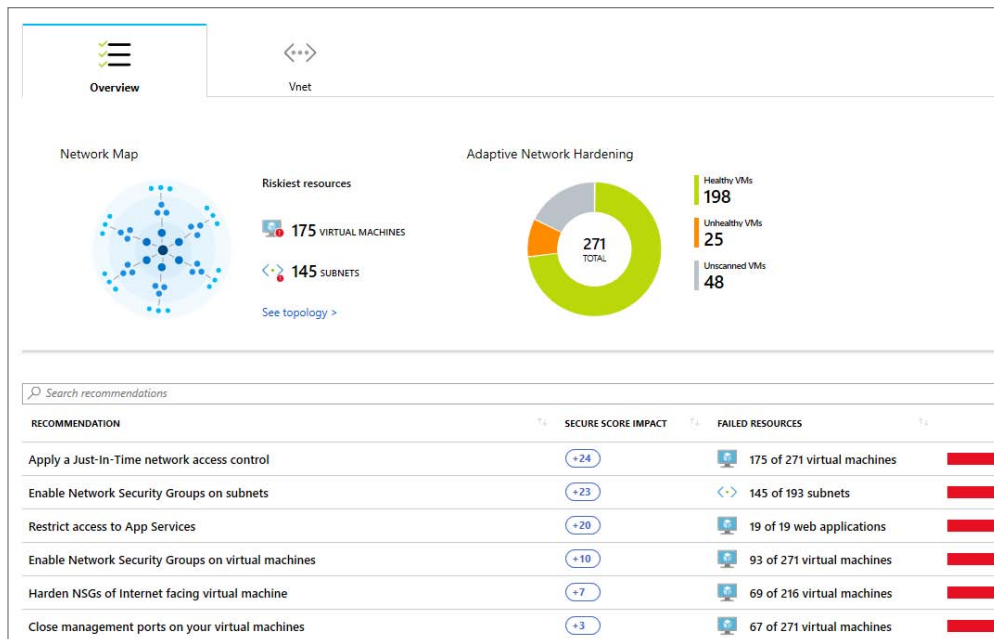


FIGURE 4-15 Networking recommendations blade

Network map

The **Network Map** appears in the first tile of this page. This option allows you view the topology of your Azure network and the traffic pattern, as shown in Figure 4-16.

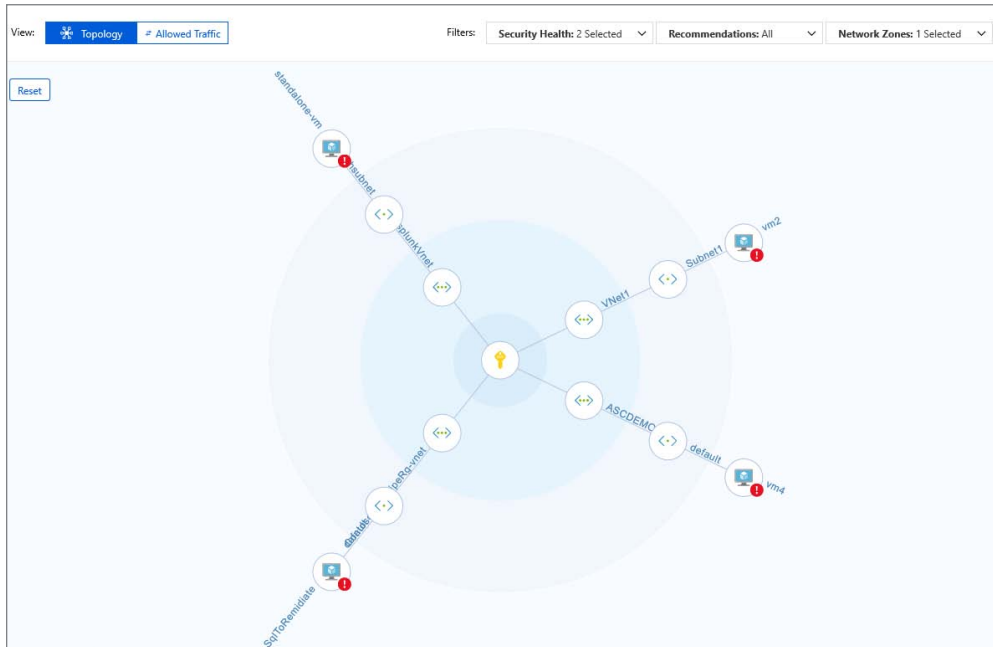


FIGURE 4-16 Network Map showing the Azure network topology for the selected subscription

TIP If there are too many resources being displayed in the map, Security Center will cluster the resources by highlighting the ones in the most critical state.

This map is organized from inside-out, with the subscription shown in the middle, then it goes to the Azure VNet, followed by the subnet and the Virtual Machine (VM) connected to that subnet. The default visualization includes some filters that can be adjusted by:

- **Security Health:** This is based on severity level (high, medium, or low) of your Azure resources.
- **Recommendations:** These are based on the active recommendations for those resources.
- **Network Zones:** This is based on internal or Internet-facing resources (or both).

If you hover the mouse over each icon, you will see more information about that object, and if you click on one of those VMs, you will see more details about the VM itself, as shown in Figure 4-17.

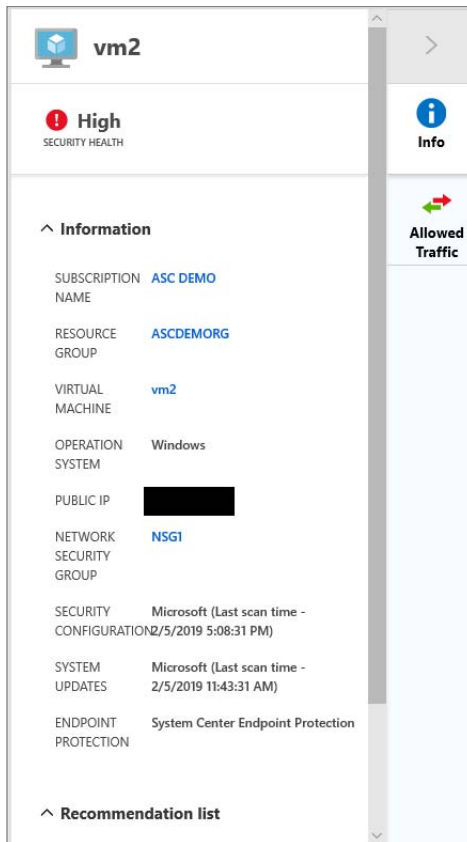


FIGURE 4-17 VM properties in the Network Map

In this blade, you have important attributes about the selected VM, and a list of relevant recommendations is shown at the bottom of the page. On the right side, you can switch to the **Allowed Traffic** page. In this tab, you have two tables with the list of TCP and UDP ports open for outbound traffic; a third tab shows the TCP and UDP ports for inbound traffic. This provides better visibility of the traffic allowed for that specific machine. If you want a broader view of the current possible traffic between your resources, you can switch to the traffic map by clicking the **Allowed Traffic** button in the left corner next to the **Topology** button. By switching to this view, you see the same map, but now it shows your configured rules defining which resources can communicate with other resources. Once you finish reviewing this blade, close it and go back to **Security Center - Networking** blade.

Adaptive network hardening

The intent of this feature is to enable you to easily harden your network traffic by using Network Security Group (NSG) rules. Security Center will use machine learning to gather information about the network traffic and after learning about the traffic pattern, it will suggest a list of rules to harden the known traffic. Click the **Adaptive Network Hardening** tile and you will

see the **Harden Network Security Group rules of internet facing virtual machines** blade, as shown in Figure 4-18.

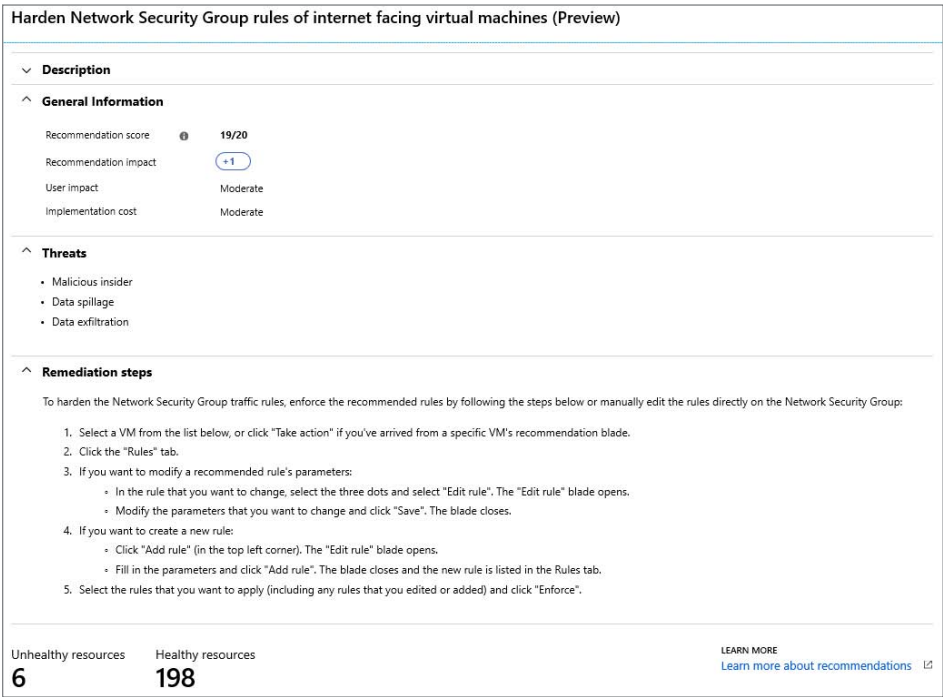


FIGURE 4-18 NSG Hardening dashboard showing the configured VMs

Unhealthy resources shows the number of VMs with network-hardening recommendations that should be applied. You can click on the VM and see the current rules and alerts, as shown in Figure 4-19.

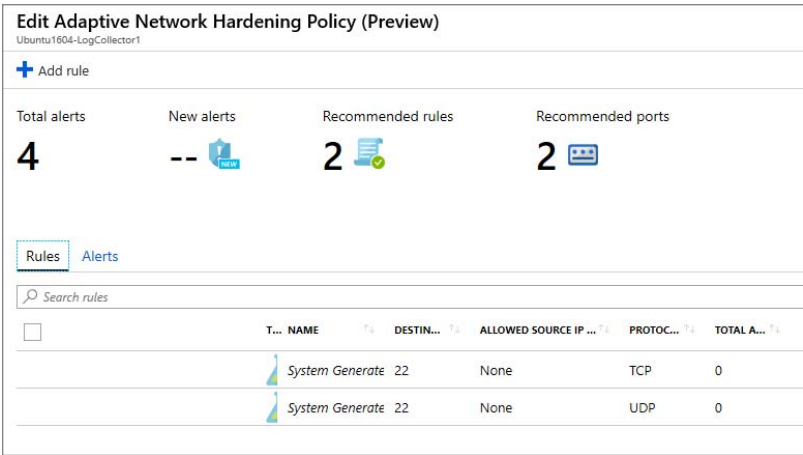


FIGURE 4-19 Hardening policy with visibility for alerts and rules

Under the **Rules** tab, you see the rules automatically created by the Adaptive Network Hardening capability based on the learning process. You can click the **Add Rule** button to create a new rule, or you can select the suggested rule and click the **Enforce** button. To see the alerts, you can click the **Alerts** tab, and select the alert itself. An example of an alert is shown in Figure 4-20.

Traffic from forbidden IP addresses was detected
vm2

[Learn more](#)

^
General information

DESCRIPTION

Azure security center has detected allowed inbound traffic from the below IP addresses/ranges although they were not allowed by the audit control

ACTIVITY TIME

Monday, February 4, 2019, 6:00:06 PM

SEVERITY

Medium

STATE

Active

ATTACKED RESOURCE

vm2

SUBSCRIPTION

DETECTED BY

Microsoft

ACTION TAKEN

Detected

ENVIRONMENT

Azure

RESOURCE TYPE

Virtual Machine

DESTINATION PORT

3389

PROTOCOL

TCP

FIGURE 4-20 Inbound traffic alert

There will be some cases in which the Adaptive Network Hardening algorithm won't be able to run, and the **Unscanned Resources** tab will appear. Following are the scenarios under which the Adaptive Network Hardening algorithm won't be able to run:

- VMs are Classic VMs: Only Azure Resource Manager VMs are supported.
- Not enough data is available: To produce accurate traffic-hardening recommendations, Security Center requires at least 30 days of traffic data.
- VM is not protected by ASC standard: This feature is only available in the Standard tier.

The following section includes examples of recommendations that can be remediated directly through the Security Center dashboard.

Leveraging machine learning for hardening network traffic

One of the biggest attack surfaces for workloads running in the public cloud is connections to and from the public internet. Our customers find it hard to know which network security group (NSG) rules should be in place to make sure that Azure workloads are only available to required source ranges.

With the Adaptive Network Hardening feature, Azure Security Center learns the network traffic and connectivity patterns of your Azure workloads and provides NSG rule recommendations for your internet-facing virtual machines. This helps you better configure your network access policies and limit your exposure to attacks, even when there are already filtering rules in place, as the filtering rules may be too permissive or the actual traffic flowing through the NSG is a subset of the NSG rules defined. In this case, you can further improve the security posture by hardening the NSG rules, based on the actual traffic patterns.

For example, let's say the existing NSG rule is to allow traffic from 140.20.30.10/24 on port 22. Based on the analysis, the adaptive network hardening's recommendation would be to narrow the range and allow traffic from 140.23.30.10/29 – which is a narrower IP range, and deny all other traffic to that port.

Azure Security Center uses machine learning to fully automate this process, including an automated enforcement mechanism, enabling customers to better protect their internet-facing virtual machines with only a few clicks. These recommendations also use Microsoft's extensive threat intelligence reports to make sure that known bad actors are blocked.

Additionally, Azure Security Center will alert users when traffic is identified from IPs that aren't recommended by the algorithm.

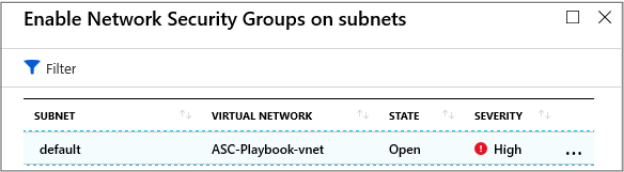
Oren Parag, Senior Program Manager, Azure Security Center Team

Enable Network Security Groups on subnets

A Network Security Group (NSG) contains a list of security rules that allow or deny network traffic to resources connected to Azure Virtual Networks (VNet). NSGs can be associated with VMs, NICs, and subnets. Security Center will verify if your subnet needs to have an NSG to be more secure, and if it does, Security Center will create a network recommendation for it.

In the main networking blade, you will see a list of recommendations. Follow the steps below to apply a recommendation:

1. Click **Enable Network Security Groups on subnets**, and you will see a new blade with all virtual networks that can have an NSG enabled, as shown in Figure 4-21.



SUBNET	VIRTUAL NETWORK	STATE	SEVERITY
default	ASC-Playbook-vnet	Open	High

FIGURE 4-21 Option to enable NSG for a subnet.

2. Click the subnet on which you want to enable NSG, and if you already have an NSG, you will see it listed on the **Choose network security group** blade. If you don't have an NSG, click **Create new** (see Figure 4-22) and follow the wizard to create a new Network Security Group.




FIGURE 4-22 Option to create a new NSG from the Security Center dashboard

Restrict access through Internet-facing endpoint

Security management and control is imperative when dealing with cloud resources. When new VMs are provisioned in your IaaS environment, you need to ensure that these VMs are not fully exposed to the Internet. This can be a hard task to do in a large environment, but with Security Center, you have full visibility of Internet-facing endpoints in a single location.

In the main networking blade, you will see a list of recommendations. Follow the steps below to apply this recommendation:

1. 1. Click **Access should be restricted for permissive Network Security Groups with Internet-facing VMs**, and you will see a new blade with all Internet facing machines, as shown in Figure 4-23:

Home > Security Center - Networking > Restrict access of Internet-facing VMs' permissive Network Security Groups

Restrict access of Internet-facing VMs' permissive Network Security Groups

Description

Azure Security center has identified some of your Network Security Groups' inbound rules to be too permissive. Inbound rules should not allow access from 'Any' or 'Internet' ranges. This can potentially enable attackers to easily target your resources.

General Information

Recommendation score	14/20
Recommendation impact	+6
User impact	High
Implementation cost	Low

Threats

- Malicious insider
- Data spillage
- Data exfiltration

Remediation steps

We recommend that you edit the inbound rules of some of your virtual machines, to restrict access to specific source ranges.

To restrict access to your virtual machines:

1. Select a VM to restrict access to.

FIGURE 4-23 List of Internet facing endpoints.

2. Scroll down to the bottom of this page, and you will see the list of unhealthy VMs, which are the ones you should address the recommendation, as shown in Figure 4-24:

Unhealthy resources	Healthy resources
80	177

Unhealthy resources (80) [Healthy resources \(177\)](#) [Unscanned resources \(10\)](#)

NAME
 standalone-vm
 UbuntuDVWA

FIGURE 4-24 List of virtual machines that need to be addressed by this recommendation

- Click on the VM that you want to address this recommendation and Networking page for that VM appears. Click **Add inbound security rule**, and the **Add inbound security rule** page appears as shown in Figure 4-25:

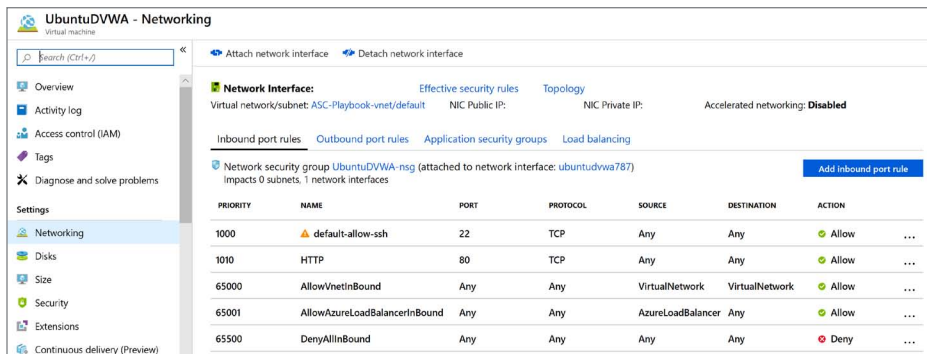


FIGURE 4-25 Current list of Inbound Security Rules

- Configure the new inbound rule according to your environment's needs, and click **Add** button to finish.

TIP The *SecurityCenterJITRule_<suffix>* is automatically created and managed by the Security Center just-in-time VM Access feature, and it should not be manually modified.

Data and storage

One of the ultimate goals of an attacker is to gain access to a target's data. Therefore, it is important to address all security recommendations for storage (where the data is located) and for the data itself.

Security Center meets these requirements by providing security recommendations for Azure SQL databases, Azure Storage, Redis, Data Lake Analytics, and Data Lake Store. Storage and data recommendations vary depending on the environment. When you first open data and storage recommendations, you will see an **Overview** tab with a summary of all recommendations, as shown in Figure 4-26.

Overview

SQL

Storage accounts

Redis

Data Lake Analytics

Data Lake Store

Search recommendations

RECOMMENDATION	SECURE SCORE IMPACT	FAILED RESOURCES	SEVERITY
Vulnerabilities on your SQL databases should be remediated (Preview)	+30	1 of 20 SQL databases	<div><div></div></div>
Secure transfer to storage accounts should be enabled	+19	168 of 180 storage accounts	<div><div></div></div>
Provision an Azure AD administrator for SQL server	+16	12 of 15 SQL servers	<div><div></div></div>
Vulnerability assessment should be enabled on your SQL servers	+15	15 of 15 SQL servers	<div><div></div></div>
Advanced data security should be enabled on your SQL servers	+10	5 of 15 SQL servers	<div><div></div></div>
Storage accounts should be migrated to new AzureRM resources	+10	5 of 5 classic storage accou...	<div><div></div></div>
Restrict access to storage accounts with firewall and virtual network configurations...	+5	179 of 180 storage accounts	<div><div></div></div>
Diagnostics logs in Azure Data Lake Store should be enabled	+5	1 of 1 data lake store accou...	<div><div></div></div>

FIGURE 4-26 Data and storage blade with the current list of recommendations

On the **Overview** tab, you have the list of all recommendations, and if you want, you can narrow the search by choosing the **SQL** or **Storage accounts** tabs. The following sections cover the implementation of some of these recommendations.

Enable auditing on SQL server

Security Center can perform a security assessment to verify whether you are leveraging Azure SQL auditing and threat detection security capabilities. Auditing and threat detection can assist you with the following tasks:

- Maintaining regulatory compliance
- Understanding database activity
- Gaining insight into discrepancies and anomalies
- Identifying security violations

Because of this integration with Azure SQL, this recommendation enables you to apply the remediation directly from Security Center. Follow these steps to address this recommendation:

1. On the **Data & Storage** blade, click the **Overview** tab, and choose **Enable Auditing On SQL server**.
2. On the **Enable auditing on SQL server** blade, you will see the detailed explanation about this recommendation, how it will affect your Secure Score, the threats to which you are exposed if you don't address this recommendation, and the remediation steps. See Figure 4-27.

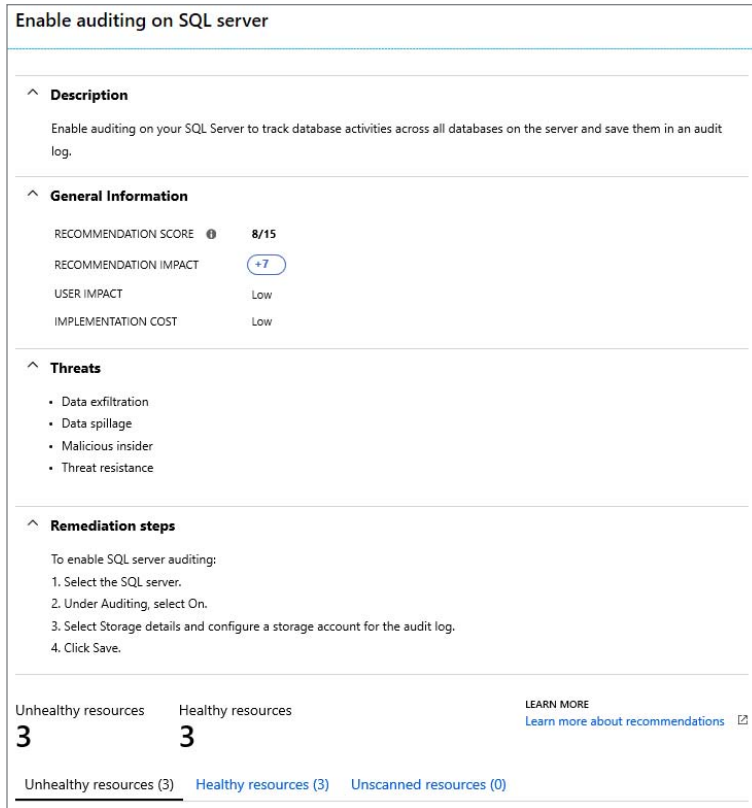


FIGURE 4-27 Details about the advantages of enabling auditing on SQL server

3. Click **Unhealthy Resources** tab at the bottom of this page to see the SQL servers that are affected by this recommendation.
4. Click the SQL server for which you want to remediate this recommendation, and the **Auditing** blade will open, as shown in Figure 4-28.

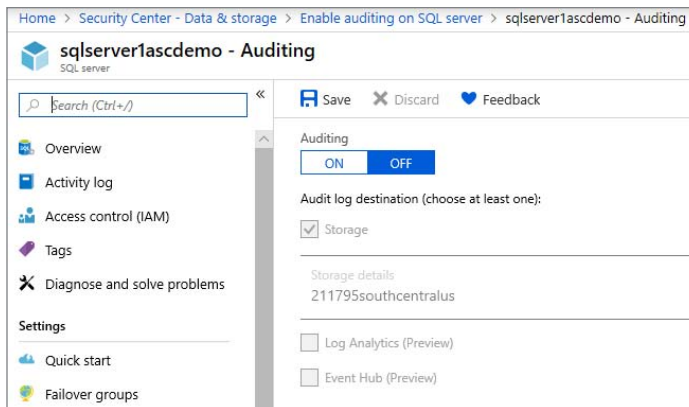


FIGURE 4-28 Remediating the recommendation by enabling auditing

5. Under **Auditing**, click **On** and click the **Save** button.
6. After the save has finished, close this blade.

Disable unrestricted network access to storage account

Security Center will also provide recommendations to harden the network access to storage accounts. This recommendation highlights that currently you have unrestricted network access to your storage account. The recommendation is to configure network rules so only applications from allowed networks can access the storage account. To allow connections from specific Internet or on-premise clients, access can be granted to traffic from specific Azure virtual networks or to public Internet IP address ranges. To address this recommendation, follow these steps:

1. On the **Data & Storage** blade's **Overview** tab, click **Disable Unrestricted Network Access To Storage Account**.
2. On the **Disable unrestricted network access to storage account** blade, you will see a detailed explanation about this recommendation, how it will affect your Secure Score, the threats that you are exposed if you don't address this recommendation, and the remediation steps, as shown in Figure 4-29.

Disable unrestricted network access to storage account (Preview)

^

Description

Disable the option of allowing unrestricted network access in your storage account firewall settings. Instead, configure network rules so only applications from allowed networks can access the storage account. To allow connections from specific internet or on-premise clients, access can be granted to traffic from specific Azure virtual networks or to public internet IP address ranges.

^

General Information

RECOMMENDATION SCORE	0/15
RECOMMENDATION IMPACT	+15
USER IMPACT	Low
IMPLEMENTATION COST	Low

^

Threats

- Data exfiltration
- Data spillage
- Account breach

^

Remediation steps

To disable unrestricted network access:

1. In your storage account, go to 'Firewalls and virtual networks'.
2. Under 'Allow access from', choose 'Selected networks'.
3. Configure the relevant virtual networks and IP ranges that should be allowed to access your account.

Unhealthy resources

81

Healthy resources

0

LEARN MORE

[Learn more about recommendations](#)

Unhealthy resources (81)

Healthy resources (0)

Unscanned resources (0)

FIGURE 4-29 Details about the storage account recommendation

3. Click the **Unhealthy resources** tab at the bottom of this page to see the storage accounts that are affected by this recommendation.
4. Click the storage account for which you want to remediate this recommendation, and the **Firewall and virtual networks** blade will appear, as shown in Figure 4-30.

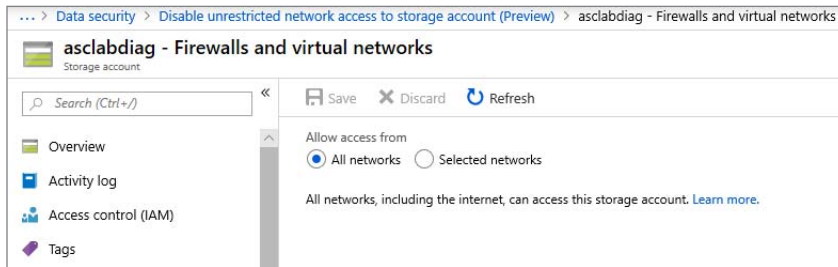


FIGURE 4-30 Hardening network access to storage accounts

5. Click the **Selected networks** option and configure the network restrictions based on the available options, as shown in Figure 4-31.

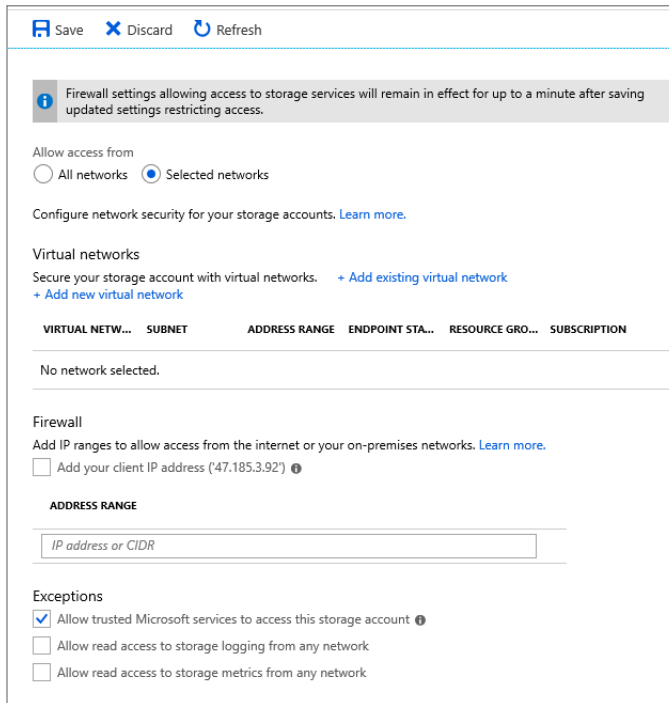


FIGURE 4-31 Configuring network security settings for storage account access

6. After you have finished configuring the appropriate options for your environment, click **Save** to close the blade.

TIP For more information about Azure Storage network security settings, visit <http://aka.ms/storagenetsec>.

Identity and access

Today, security monitoring goes beyond actively watching the security state of your workloads. Now, you must take a broader approach that also includes identity and access. When an adversary can compromise one user's credential, the attacker might leverage this legitimate account for a much larger attack campaign. Regardless of where your workload is located—in the cloud or on-premises—it becomes imperative to monitor your user's behaviors, access, and how their credentials are being used.

You can take advantage of Security Center to understand these patterns and monitor your identity posture. To access the identity dashboard, click **Identity & Access** under **Resource Security Hygiene**, in the left navigation pane. The **Identity & Access** blade appears, as shown in Figure 4-32.

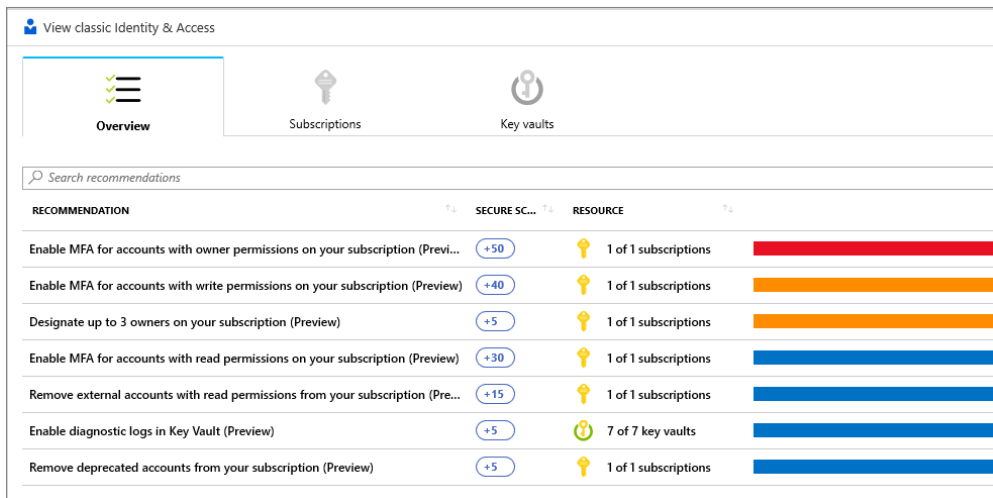


FIGURE 4-32 The Identity & Access overview page

The **View Classic Identity & Access** button allows you to switch to the legacy view in which only pure identity statistics were shown (see Figure 4-33). This button will be retired on July 31, 2019, but because many customers like this visualization, you can reproduce the data visualized in this dashboard using a query that the Azure Security Center team made available on GitHub. See <http://aka.ms/ASCLegacyIdentity>.

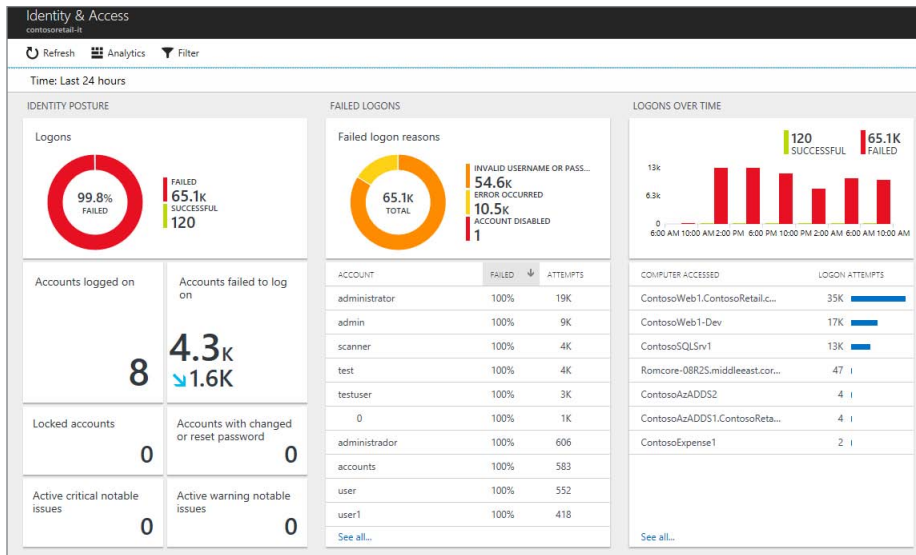


FIGURE 4-33 Legacy view of the Identity & Access page

This dashboard shows a summary of all identity-related activities monitored by Security Center. Security operations personnel should visit this dashboard multiple times throughout the day to quickly assess the current identity state. This dashboard contains three major sections:

- **Identity Posture**
- **Failed Logons**
- **Logons Over Time**

The **Identity Posture** section of the **Identity & Access** dashboard conveys various analytics representing the last 24 hours of data, which can help you better understand your users' authentication patterns. These analytics include the following:

- **Logons:** This pie graph shows the number of failed and successful logons.
- **Accounts Logged On:** This counter shows the number of user accounts that are currently logged on.
- **Accounts Failed To Log On:** This counter shows the total number of accounts that failed to log on, along with an arrow pointing up or down. In Figure 8-2, the arrow is pointing down, signaling that the number of failed logons has dropped by 1.6K since the last update.
- **Locked Accounts:** This counter conveys the number of accounts that are currently locked.
- **Accounts With Changed Or Reset Password:** This counter shows the total number of accounts whose passwords have been changed or reset.

- **Active Critical Notable Issues:** Notable issues are critical events. In the context of identity, they are critical identity events. This counter reflects the active notable issues that require immediate attention.
- **Active Warning Notable Issues:** This is the same as the previous counter, but for medium-priority issues rather than critical ones.

The **Failed Logons** section of the **Identity & Access** dashboard conveys, at a glance, the main causes of failed logons.

When a user fails to authenticate, Windows generates an event called Event 4625. To investigate why a particular logon failed, click the logon in the **Failed Logon Reasons** section. As shown in Figure 4-34, an Event 4625 window opens that shows the reason for the failure. (See the **Failure Information** section.)

Keywords	Date and Time	Source	Event ID	Task Category
Audit Failure	12/29/2017 10:38:51 AM	Microsoft Windows security auditing.	4625	Logon
Event 4625, Microsoft Windows security auditing.				
<div>General</div> <div>Details</div> <div> <p>Failure Information:</p> <p>Failure Reason: Unknown user name or bad password.</p> <p>Status: 0xC000006D</p> <p>Sub Status: 0xC0000064</p> <p>Process Information:</p> <p>Caller Process ID: 0x0</p> <p>Caller Process Name: -</p> <p>Network Information:</p> <p>Workstation Name: YURIDBVT01</p> <p>Source Network Address: 192.168.1.18</p> <p>Source Port: 65121</p> <p>Detailed Authentication Information:</p> <p>Logon Process: NtLmSsp</p> <p>Authentication Package: NTLM</p> <p>Transited Services: -</p> <p>Package Name (NTLM only): -</p> <p>Key Length: 0</p> </div>				

FIGURE 4-34 A description of Event 4625

The information available in the Event 4625 dialog may vary depending on the reasons for the failure. Table 8-1 lists some of the most common reasons for failed logons.

TABLE 8-1 Common reasons for logon failures

Error code (status/substatus)	Description
0xC0000064	Unknown user name or bad password.
0xC0000234	Account locked out.
0xC0000072	Account currently disabled.
0xC0000193	The specified user account has expired.
0xC0000071	Expired password.

TIP For a complete list of logon failure reasons, see <https://aka.ms/AccountLogon>.

When reviewing the Event 4625 window, pay close attention to the *LogonTypeName* field. If this field is set to 3, it means the logon attempt came from the network. In this case, you'll also see an *IPAddress* field with the corresponding IP address. If *LogonTypeName* is set to 5, it means the logon attempt is coming from a service or process. A *Process* field will contain the process name. The code shown in Listing 4-1 contains the entire content of the Event ID 4625, which is generated when a logon request fails, and it appears on the computer where access was attempted.

LISTING 4-1 Event ID 4625

```
Log Name: Security
Source: Microsoft-Windows-Security-Auditing
Date: 2/19/2019 9:08:29 AM
Event ID: 4625
Task Category: Logon
Level: Information
Keywords: Audit Failure
User: N/A
Computer: ARGOS
Description:
An account failed to log on.
Subject:
Security ID: NULL SID
Account Name: -
Account Domain: -
Logon ID: 0x0
Logon Type: 3
Account For Which Logon Failed:
Security ID: NULL SID
Account Name:
YURIDBVT01$
    Account Domain: NORTHAMERICA
Failure Information:
Failure Reason:
Unknown user name or bad password.
Status:
0xC000006D
Sub Status: 0xC0000064
Process Information:
Caller Process ID: 0x0
Caller Process Name: -
Network Information:
Workstation Name: YURIDBVT01
Source Network Address: 192.168.1.254
Source Port: 53407
```

Detailed Authentication Information:

Logon Process: NtLmSsp

Authentication Package: NTLM

Transited Services: -

Package Name (NTLM only): -

Key Length: 0

- **Subject fields.** The Subject fields indicate the account on the local system that requested the logon. This is most commonly a service, such as the Server service, or a local process, such as Winlogon.exe or Services.exe.
- **Logon Type field.** The Logon Type field indicates the kind of logon that was requested. The most common types are 2 (interactive) and 3 (network).
- **Process Information fields.** The Process Information fields indicate which account and process on the system requested the logon.
- **Network Information fields.** The Network Information fields indicate where a remote logon request originated. Workstation Name is not always available and may be left blank in some cases.
- **Authentication Information fields.** The Authentication Information fields provide detailed information about this specific logon request:
 - **Transited Services.** Transited Services indicate which intermediate services have participated in this logon request.
 - **Package Name.** Package name indicates which subprotocol was used among the NTLM protocols.
 - **Key Length.** Key Length indicates the length of the generated session key. This will be 0 if no session key was requested.

LISTING 4-2 Event ID 4625

```
<Event xmlns="http://schemas.microsoft.com/win/2004/08/events/event">
  <System>
    <Provider Name="Microsoft-Windows-Security-Auditing" Guid="{54849625-5478-4994-A5BA-
3E3B0328C30D}" />
    <EventID>4625</EventID>
    <Version>0</Version>
    <Level>0</Level>
    <Task>12544</Task>
    <Opcode>0</Opcode>
    <Keywords>0x8010000000000000</Keywords>
    <TimeCreated SystemTime="2019-02-19T15:08:29.580037100Z" />
  <EventRecordID>60107528</EventRecordID>
  <Correlation />
  <Execution ProcessID="724" ThreadID="16256" />
```



```

<Channel>Security</Channel>
<Computer>ARGOS</Computer>
<Security />
</System>need
<EventData>
  <Data Name="SubjectUserSid">S-1-0-0</Data>
  <Data Name="SubjectUserName">-</Data>
  <Data Name="SubjectDomainName">-</Data>
  <Data Name="SubjectLogonId">0x0</Data>
  <Data Name="TargetUserSid">S-1-0-0</Data>
  <Data Name="TargetUserName">YURIDBVT01$</Data>
  <Data Name="TargetDomainName">NORTHAMERICA</Data>
  <Data Name="Status">0xc000006d</Data>
  <Data Name="FailureReason">%%2313</Data>
  <Data Name="SubStatus">0xc0000064</Data>
  <Data Name="LogonType">3</Data>
  <Data Name="LogonProcessName">NtLmSsp </Data>
  <Data Name="AuthenticationPackageName">NTLM</Data>
  <Data Name="WorkstationName">YURIDBVT01</Data>
  <Data Name="TransmittedServices">-</Data>
  <Data Name="LmPackageName">-</Data>
  <Data Name="KeyLength">0</Data>
  <Data Name="ProcessId">0x0</Data>
  <Data Name="ProcessName">-</Data>
  <Data Name="IpAddress">192.168.1.254</Data>
  <Data Name="IpPort">53407</Data>
</EventData>
</Event>

```

The **Failed Logon Reasons** section of the dashboard also contains a list of the top 10 accounts that failed to log on. Also, as with the **Logons** section of the dashboard, clicking a tile opens the **Log Search** dashboard with a relevant query result.

The **Logons Over Time** section contains a timeline of logon attempts, which can be very useful for understanding your users' authentication patterns. If you see an authentication spike every day at 6 AM—a time when your environment shouldn't be receiving a significant number of authentication requests—it certainly suggests a suspicious behavior requiring further investigation.

Subscriptions

Security Center will also manage your subscription to identify improvements in the identity ownership aspect of it. It will verify if Multifactor Authentication (MFA) is enabled for all subscription accounts with owner permissions, and it will also recommend that you designate up to three subscription owners in order to reduce the potential for breach by a compromised owner. Figure 4-35 shows an example of other subscription-level recommendations on the **Remove external accounts with read permissions from your subscription** page.

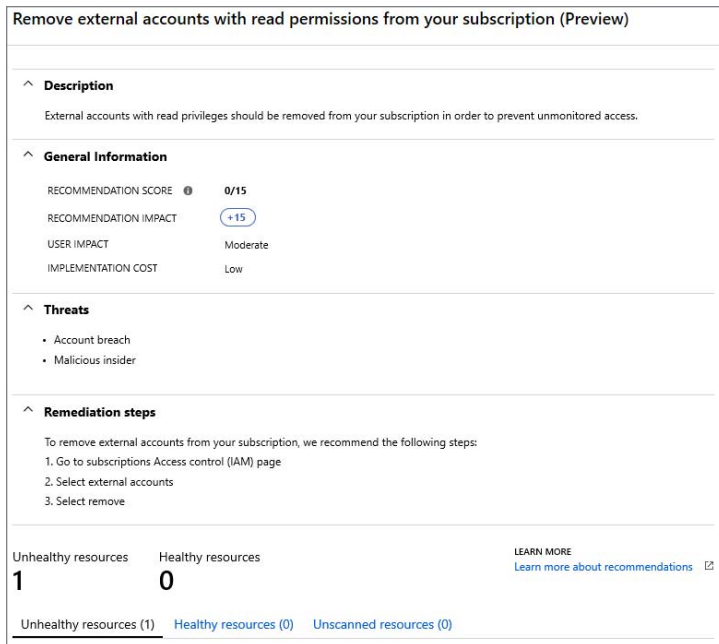


FIGURE 4-35 Identifying related recommendation for your subscription

Under **Unhealthy resources**, you will see a list of subscriptions that are affected by this recommendation. Click the subscription for which you want to remediate this recommendation.

Key vaults

Security Center monitors Azure Key Vault to identify potential security improvements in its configuration. When you click the **Key vaults** tab in the **Identity & Access** blade, you will see a list of recommendations, as shown in Figure 4-36.

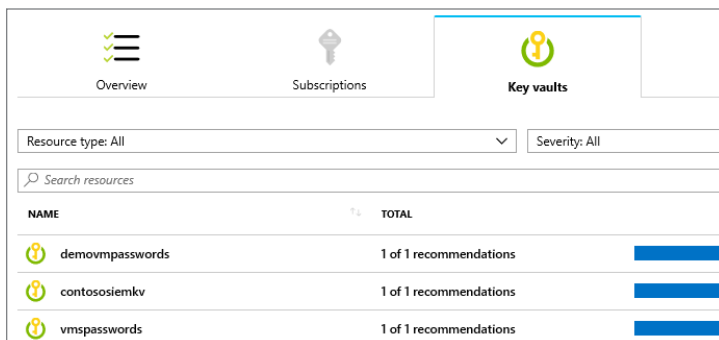


FIGURE 4-36 Recommendations for your key vaults

One of the available recommendations is to enable diagnostic logs and retain them for up to a year. This enables you to re-create activity trails for investigative purposes when a security incident occurs or your network is compromised.

App services

By leveraging the visibility that Azure has as cloud provider, Security Center analyzes App Service internal logs to identify attack methodology on multiple targets. For example, it identifies attempts to access the same Uniform Resource Identifiers (URI) on various web sites. This type of attacker typically exhibits a pattern of crawling the same web page on multiple web sites, searching for a particularly vulnerable page or plugin.

Security Center will also proactively create recommendations to enhance the security configuration of your App services. To access Apps Services, select the **Computer & Apps** option under **Resource Security Hygiene** and click the **App Services** tab. Relevant recommendations are shown (see Figure 4-37).




FIGURE 4-37 App services recommendation for an Azure function

To see the relevant recommendation, click the **App Service** item, which in Figure 4-37 is an Azure function, and a new blade appears (see Figure 4-38).

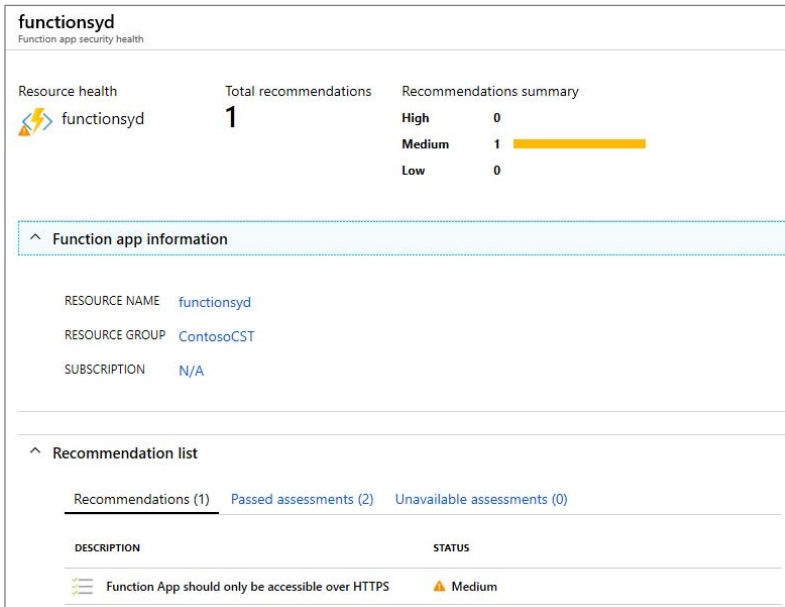


FIGURE 4-38 App Service recommendations for an Azure function

As you can see in Figure 4-38, there is only one medium-severity recommendation for this Azure function—**Function App should only be accessible over HTTPS**. The intent of this recommendation is to warn you that you currently have an Azure function that is accessible via clear text (HTTP). To protect the data in transit, the recommendation is to enable HTTPS on this function.

Containers

In Azure Security Center, the **Containers** tab displays all virtual machines deployed with Docker. Security Center provides additional information related to the containers on the machine, such as Docker version and the number of images running on the host. Security Center scans your Docker configurations and provides visibility into misconfigurations by providing a list of recommendations based on the Center for Information Security (CIS) benchmark for Docker.

NOTE For more information about the Center for Information Security (CIS) benchmark for Docker, see <https://www.cisecurity.org/benchmark/docker/>.

To see the Containers recommendation, on the **Computer & Apps** option, under **Resource Security Hygiene**, click the **Containers** tab, and you will see the relevant recommendations as shown in Figure 4-39.

Overview

VMs and Computers

Cloud services

App services

Containers (Preview)

Resource type: All

Severity: All

Search resources

NAME		TOTAL	
<div><div></div> ascdockercontainer</div>		1 of 1 recommendations	
<div><div></div> Containers-Demo-Ubutnu</div>		1 of 1 recommendations	

FIGURE 4-39 Containers tab showing the available recommendations

To get more information about a container or to remediate the recommendation, click the container. When the container shown in Figure 4-39 is clicked, the **Remediate Vulnerabilities In Container Security Configurations** recommendation appears, as shown in Figure 4-40.

Remediate vulnerabilities in container security configurations

^

Description

Remediate vulnerabilities in security configuration on machines with Docker installed to protect them from attacks.

^

General Information

USER IMPACT

Moderate

IMPLEMENTATION COST

Moderate

^

Threats

• Data exfiltration

• Data spillage

• Account breach

^

Remediation steps

To Remediate vulnerabilities in the container security configurations:

1. Review the list of failed rules.

2. Fix each rule according to the specified instructions.

Take action

FIGURE 4-40 Detailed explanation about the recommendation and remediation steps

Virtual machine scale sets

Azure Virtual Machine Scale Sets (VMSS) let you create and manage a group of identical, load-balanced VMs. One advantage of using this technology is that the number of VM instances can automatically increase or decrease in response to a business demand, or the number of VMs can also be defined by a schedule.

Security Center automatically discovers whether you have scale sets and recommends that you install the monitoring agent on these scale sets. Once the agent is installed, Security Center will perform an assessment to provide recommendations, such as enabling diagnostics logs in Virtual Machine Scale Sets, remediating vulnerabilities in security configuration on your virtual machine scale sets, installing system updates on virtual machine scale sets, and others.

The VMSS recommendation is located on the **Computer & Apps** option, under **Resource Security Hygiene**. After accessing this option, click the **VM Scale Sets** tab, and you will see the relevant recommendations as shown in Figure 4-41.




FIGURE 4-41 VMSS recommendations in Security Center

To access the recommendation for a particular VMSS, click it and you will see more details about the Scale Set, as well as a list of recommendations at the bottom of the page, as shown in Figure 4.42.

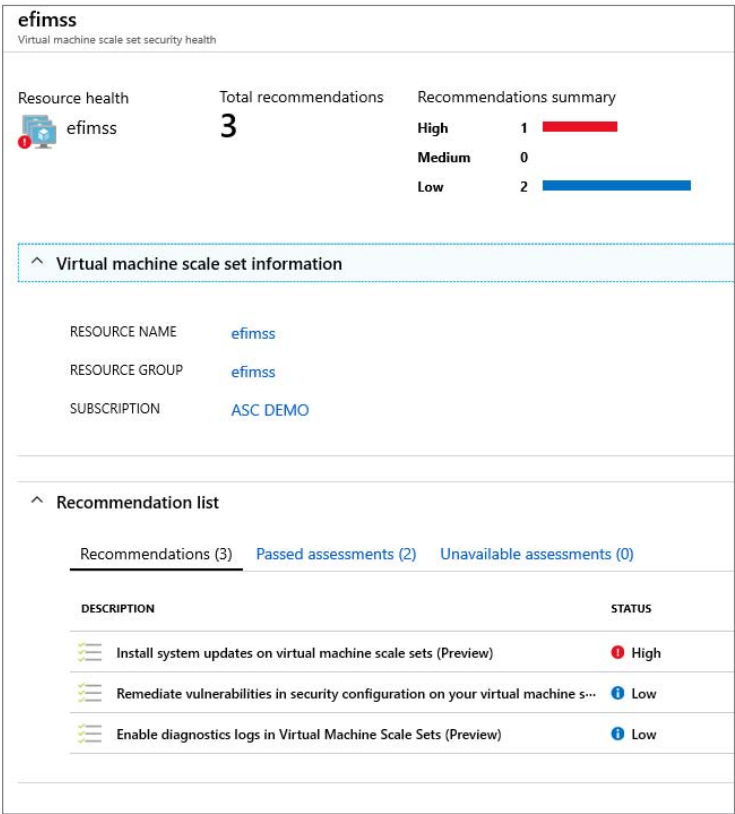


FIGURE 4-42 Details about a particular scale set and its recommendations

Single click remediation

Now that you know the scope of recommendations across different workloads, is possible to visualize some remediation scenarios and realize that it can be challenging to manually remediate multiple recommendations. To simplify the remediation process, Security Center launched a new feature in summer 2019 that enables you to remediate a recommendation with just a few steps.

By the time we were writing this chapter, the single-click remediation feature was in Private Preview, and for this release, only the following recommendations were available:

- Disk encryption should be applied on virtual machines.
- Web Application should only be accessible over HTTPS.
- Function App should only be accessible over HTTPS.
- API App should only be accessible over HTTPS.
- Remote debugging should be turned off for Function App.
- Remote debugging should be turned off for Web Application.
- Remote debugging should be turned off for API App.
- CORS should not allow every resource to access your Function App.
- CORS should not allow every resource to access your Web Application.
- CORS should not allow every resource to access your API App.
- Secure transfer to storage accounts should be enabled.
- Transparent Data Encryption on SQL databases should be enabled.
- Monitoring agent should be installed on your virtual machines.
- Auditing on SQL server should be enabled.

After applying the remediation using this feature, the recommendation will be refreshed according to the recommendation type. If the remediation was performed using Azure Policy (control plane), the refresh will take place in 10 minutes; in other words, the recommendation disappears from the list because it was remediated already.

TIP The Azure Security Center engineering team is constantly working to improve the dashboard freshness; therefore, it is recommended to visit the Security Center Frequent Asked Questions (<http://aka.ms/ascfaq>) article for the latest information regarding this latency.

Follow the steps below to use the single click remediation feature to remediate a recommendation in Security Center:

7. Open **Azure Portal** and sign in with a user that has **Security Admin** privileges.
8. In the left navigation pane, click **Security Center**.
9. In the Security Center left navigation pane, under **Resource Security Hygiene**, click **Recommendations**.
10. You will see that some recommendations will have the blue **1-Click Fix!** button, as shown Figure 4-43.

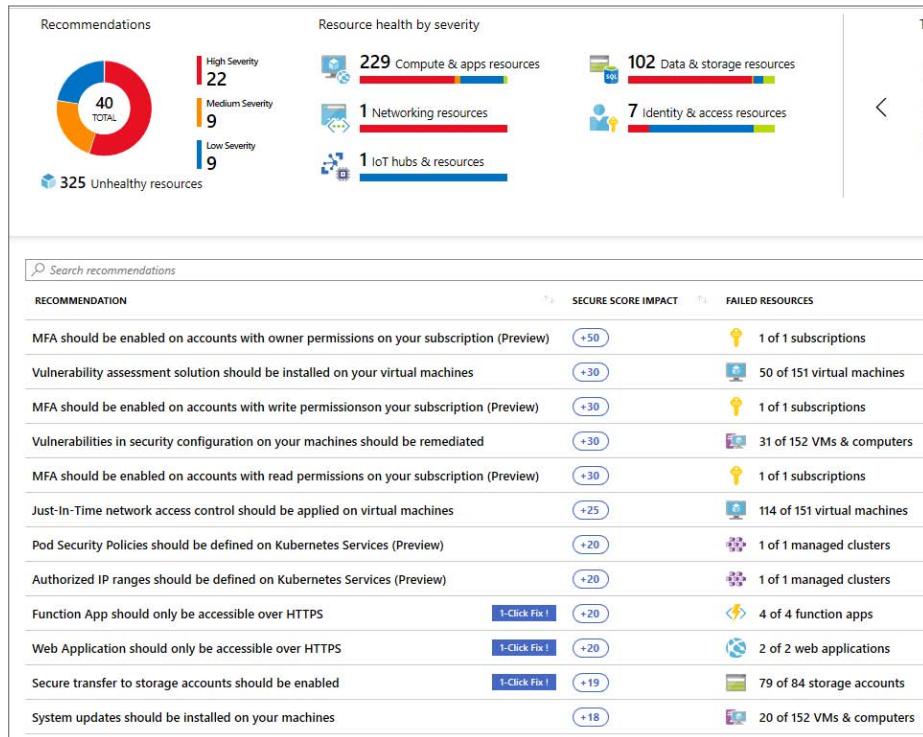


FIGURE 4-43 Some recommendations will have the single-click remediation feature enabled.

11. For this example, we will use the **Secure Transfer To Storage Accounts Should Be Enabled** recommendation. Click this recommendation, scroll down to see the storage account that is considered unhealthy, select the storage account, and click the **Remediate** button, as shown in Figure 4-44.




FIGURE 4-44 Remediating the recommendation directly from the recommendation page

12. The Remediate Resources page appears, as shown in Figure 4-45. Select the resource that you want to remediate (in this case, just one resource) and click **Remediate 1 Resource**. (Keep in mind that the button's name will vary to reflect the number of resources you selected.)

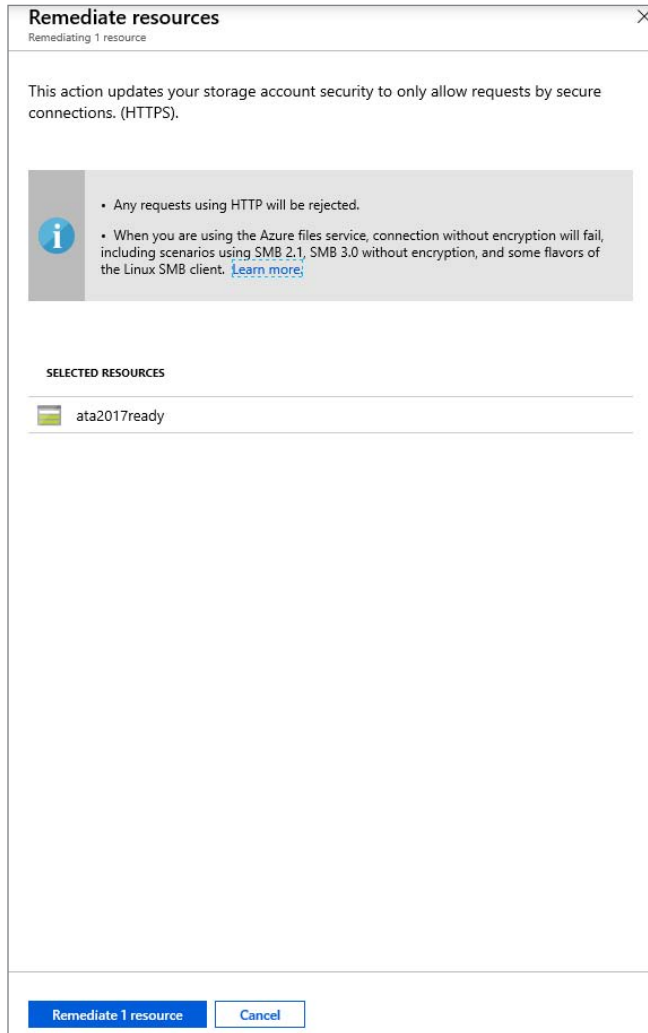


FIGURE 4-45 Resources available to apply the remediation

After you apply the remediation, notice that the resource will be unavailable in this list.