

RSA IDENTITY GOVERNANCE AND LIFECYCLE

**RSA Identity Governance and Lifecycle 7.5
Azure Installation Guide**

Contact Information

RSA Link at <https://community.rsa.com> contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA Conference Logo, RSA, and other trademarks, are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to <https://www.rsa.com/en-us/company/rsa-trademarks>. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability. This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on RSA Link. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright © 2021 RSA Security LLC or its affiliates. All Rights Reserved.

Contents

Preface	5
Documentation Set	5
Support and Service	5
Chapter 1: Overview	6
Deployment Overview	6
Requirements.....	6
Chapter 2: Deploy a Standalone Instance.....	7
Creating the Azure Virtual Machine.....	7
Basics.....	7
Disks	8
Networking.....	8
Management.....	8
Advanced.....	8
Tags	8
Review + Create	8
Pre-Installation Configuration.....	9
Post-Deployment Restart.....	9
Accessing the Virtual Machine.....	9
Login and root Access	9
Internal Hostname and DNS	9
NTP	10
Disk Partitioning.....	10
Disable Hardware Lock Elision	11
Install Required OS Packages	12
Inbound Port Rules	12
Local Firewall Port Forwarding	13
Remote Database.....	14
Installing RSA IG&L v7.5.0.....	14
Copy the installation packages	14
Run the Installer.....	14
Access the Application	15
Upgrade Standalone RSA Identity Governance and Lifecycle	15

Chapter 3: Performance Metrics 15

 Azure Configuration 15

 ACME Collections 16

 Volume tests with Single User 16

 Business Description Imports 17

 Entitlement Reports 18

 Reviews 18

 Annual Review Generation Performance 19

 Annual Review Refresh 19

 User Review Generation Performance 20

 User Review Refresh 20

 Rules and Roles 21

 Separation of Duty (SoD) Rule Processing 21

 Unauthorized Change Detection (UCD) Rule Processing 22

 Suggest Entitlements for Roles 22

 500 concurrent users test on Azure 23

 Access Requests on Azure with 500 concurrent users 23

 User Reviews on Azure with 500 concurrent users 24

 AFX Performance Results 24

 Joiner (Implicit) – Attribute Change Rule 25

 Mover (Implicit) – Attribute Change Rule 25

 Leaver – Provisioning Termination Rule 25

Chapter 4: ACME Specification 26

 Identity Information 26

 Account Information 26

 Application Information 26

Chapter 5: Knowledgebase 27

 Disk Partitioning 27

Preface

This guide provides recommendations and procedures to deploy RSA Identity Governance and Lifecycle (RSA IG&L) on the Microsoft Azure cloud platform. It also provides a summary of sample performance test results for RSA Identity Governance and Lifecycle 7.5.0 in an Azure environment.

Although each customer environment and deployment are unique, this document outlines a common set of use cases, tests and associated performance numbers, and serves as a guideline. Results may vary based on resources allocated, chosen instance type, and the location of the data centers as compared to that of the users. The Azure instance continues to evolve, and while this guide provides high level steps, some of the steps and instance types may evolve over time.

This document details the steps to configure and deploy a single Microsoft Azure Virtual Machine (VM). This VM can then be used to install a standalone RSA IG&L v7.5.0. Software Bundle. The same steps can be used to deploy multiple VMs for clustering purposes or to install the database and application server components on separate VMs.

Documentation Set

The latest product documentation is always available at:

<https://community.rsa.com/community/products/governance-and-lifecycle>.

Document	Description
Release Notes	What's new in the release, fixed issues, known issues and workarounds.
Installation Guide	Product installation instructions.
Upgrade and Migration Guide	Instructions for upgrading your product version and data.
Database Setup and Management Guide	Instructions for setting up and managing a customer-supplied Oracle database for RSA Identity Governance and Lifecycle.
Online Help	All concepts and instructions you need to configure and use the product.
Administrator's Guide	How to configure and manage RSA Identity Governance and Lifecycle. Contains a subset of the information provided in the Online Help.
Public Database Schema Reference	The public view of the database schema.

For information about how to create a Linux virtual machine in Azure, see:

<https://docs.microsoft.com/en-us/learn/modules/create-linux-virtual-machine-in-azure/>.

Support and Service

You can access community and support information on RSA Link at

<https://community.rsa.com/community/products/governance-and-lifecycle>. RSA Link contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

You can also access the RSA Identity Governance and Lifecycle community at <https://community.rsa.com/community/products/governance-and-lifecycle/client-partner-community>. This private community is only available to RSA Identity Governance and Lifecycle customers, partners, and internal RSA staff.

Chapter 1: Overview

This guide provides instructions for deploying a standalone RSA Identity Governance and Lifecycle deployment in a Microsoft Azure cloud platform and is intended for RSA Identity Governance and Lifecycle installers and administrators. To follow these procedures, users must have the appropriate network and administrative permissions to install and configure the RSA Identity Governance and Lifecycle application and familiarity with Azure.

Note: This document is not intended to suggest optimal installations or configurations.

An Azure deployment of RSA Identity Governance and Lifecycle comprises the following components:

- An RSA Identity Governance and Lifecycle node deployed on a virtual machine.
- The database component is deployed in a separate virtual machine and must adhere to the requirements specified in the *RSA Identity Governance and Lifecycle Database Setup and Maintenance Guide*.

Deployment Overview

To deploy RSA Identity Governance and Lifecycle in an Azure environment, you must:

1. Provision an Azure virtual machine. For more information, see [Deploy a Standalone Instance](#).
2. Configure the Azure environment for a standalone/single-node RSA Identity Governance and Lifecycle installation. For more information, see [Deploy a Standalone Instance](#).

To deploy RSA Identity Governance and Lifecycle with a remote database, provision an Azure virtual machine for the RSA Identity Governance and Lifecycle database. For more information, see [Remote Database](#).

Requirements

Ensure that you meet the following requirements before proceeding:

- Microsoft Azure subscription to create Virtual Machines.
- Access to the RSA IG&L 7.5.0 installation packages.

Azure VM Requirements

OS	SUSE Linux Enterprise Server 12 SP4+ or Red Hat Enterprise Linux 7.6+ or Red Hat Enterprise Linux 8.0+
Size	Standard_E4s_v3-4cpus (32GB RAM) or better
Storage	128GB

Database Requirements

Database engine version	Oracle 12c or Oracle 19c
-------------------------	--------------------------

These requirements and recommendations are based on internal performance testing. For more information, see Chapter 3: Performance Metrics.

Required Software Access

The following additional software applications or equivalent are required during the RSA IG&L deployment, installation and configuration process.

Software	Use
Putty	Open source terminal emulator supporting SCP, SSH, Telnet and raw socket connection.
SFTP Client (e.g. FileZilla or WinSCP)	Cross-platform SFTP application used to transfer RSA IG&L installation files from local to remote machine.
Firefox/Chrome	A second browser allows two concurrent sessions and is useful for testing purposes.

Chapter 2: Deploy a Standalone Instance

Creating the Azure Virtual Machine

Log in to the Azure Portal.

1. Go to <https://portal.azure.com/>, login, then click on the Virtual Machines icon under Azure Services.
2. Click the Add button to start the new Virtual Machine wizard. There are seven tabs in the new Virtual Machine wizard:
 - Basics
 - Disk
 - Networking
 - Management
 - Advanced
 - Tags
 - Review + create

Basics

1. Select the subscription based on your account type.
2. Create a new resource group to organize your resources. A resource group is a collection of resources that share the same lifecycle, permissions, and policies.
3. Set a virtual machine name. Virtual machines in Azure have two distinct names: virtual machine name used as the Azure resource identifier, and the guest host name.

When you create a virtual machine in the portal, the same name is used for both the virtual machine name and the host name. The virtual machine name cannot be changed after the virtual machine is created; however, you can change the host name when you login to the virtual machine.

4. Choose the Azure region that is appropriate for you and your customers.

Note: Not all virtual machine sizes are available in all regions.

5. Select the required availability options. Azure offers a range of options for managing availability and resiliency ([Learn more](#)).
6. Select Operating System (SUSE Linux Enterprise Server 12 SP4+ or Red Hat Enterprise Linux 7.6+ or Red Hat Enterprise Linux 8.0+) image for the virtual machine.
7. Select **No** for Azure Spot Instance.
8. Select a virtual machine size **Standard E4s v3**.
9. Select the **Password** authentication type.

Note: For security reasons, do not choose the Password authentication type if you plan to expose SSH publicly.

10. Click **Next: Disks**.

Disks

1. Select the required OS disk type and size. For optimal performance in production environments, RSA recommends the following:
 - Use separate disks for the database and application server.
 - At least the database disk should be "Premium SSD".
2. Click "Next: Networking."

Networking

Networking settings vary based on the customer environments. Work with your Azure admin to determine the appropriate settings.

When you finish configuring your networking settings, click **Next: Management**.

Management

Monitoring and management settings vary based on customer environments. Work with your Azure admin to determine the appropriate settings.

When you finish configuring your monitoring and management settings, click **Next: Advanced**.

Advanced

Click **Next: Tags**.

Tags

Click **Next: Review + Create**.

Review + Create

1. Review all the previously selected settings on this summary page and click **Create**.
2. Wait until your virtual machine deployment is completed before moving on to the next section. If required, click **Refresh** to get the latest status.
3. Click **Go to resource** to start the next section.

Pre-Installation Configuration

Post-Deployment Restart

Before continuing, restart the virtual machine. This ensures that all settings are applied correctly and stops ExtHandler console updates that affects the access to the virtual machine.

Accessing the Virtual Machine

You can access the virtual machine through the Azure Portal's Serial console, or use any standard SSH client (for example: PuTTY).

On the left panel, scroll to "Support + troubleshooting" and click **Serial console**.

Login and root Access

Log in using the administrator account created in previous steps.

Use the command below along with your administrator account password to switch to root, which is required for all upcoming steps.

```
sudo -i
```

Internal Hostname and DNS

RSA Identity Governance & Lifecycle v7.5.0 requires both forward and reverse DNS to be configured for the internal hostname and IP address of the virtual machine. Azure does not allow configuring reverse DNS for internal hostnames and IP addresses. Do the following to force update the internal hostname and DNS settings:

1. Change the internal hostname to the required hostname using the below command:

For Example

```
hostnamectl set-hostname rsa-igl.internal.cloudapp.net
```

Note: "rsa-igl" is the internal hostname and "internal.cloudapp.net" is the domain name.

2. Use any text editor to update the default search domain in the `/etc/resolv.conf` to match the above. The following example uses vim:

```
vim /etc/resolv.conf
```

Change the search line to search internal.cloudapp.net.

3. List and note the internal virtual machine IP address using the following command:

```
ip address show
```

The IP address is displayed under the interface name after "inet."

4. Use any text editor to update the `/etc/hosts` file with the new internal hostname and IP address. The following example uses vim:

```
vim /etc/hosts
```

Add the below line before "#Added by SMT ..." using the IP address and hostname for your deployment.

```
10.0.0.6 rsa-igl.internal.cloudapp.net rsa-igl
```

5. Reboot the virtual machine for the changes to take effect using the following command:

```
reboot
```

6. When the virtual restarts, log in and confirm that both forward and reverse DNS lookups resolve to the same values (IP address lookup resolves to the hostname and vice versa) using the below commands:

```
nslookup 10.0.0.6
```

```
nslookup rsa-igl.internal.cloudapp.net
```

NTP

NTP is required for installing RSA IG&L v7.5.0. Ideally, you should have an internal NTP server. If not then you may use a public NTP server. This example uses the US zone public [NTP Pool Project](#) servers.

1. Use any text editor to update /etc/ntp.conf with the NTP servers. The following example uses vim:

```
vim /etc/ntp.conf
```

Add the following NTP server lines after the "# server 127.127.8.0 ..." comment line:

```
server 0.us.pool.ntp.org
```

```
server 1.us.pool.ntp.org
```

```
server 2.us.pool.ntp.org
```

```
server 3.us.pool.ntp.org
```

Important: In a Red Hat Enterprise Linux environment, run the following command before adding the above mentioned NTP server lines to ntp.conf:

```
yum install ntp
```

2. Restart the NTP service for changes to take effect using the following command:

```
systemctl restart ntpd.service
```

3. Verify that the NTP server is reachable and that the virtual machine is in sync using the below command. The delay, offset and jitter values should be relatively small but not zero as shown below:

```
ntpq -p
```

Disk Partitioning

By default, two virtual hard disks (VHDs) are created for your Linux VM:

- The **operating system disk**: This is primary drive, and it has a maximum capacity of 2048 GB. It is labelled as **/dev/sda** by default.
- A **temporary disk**: This provides temporary storage for the OS or any apps. On Linux virtual machines, the disk is **/dev/sdb** and is formatted and mounted to **/mnt** by the Azure Linux Agent. It is sized based on the VM size and is used to store the swap file.
- Any **data disks** created in the previous steps are labelled **/dev/sdc, /dev/sdb ...** and so on.

Below are minimum requirements, these may be increased based on individual deployment requirements.

- 16GB – SWAP
- 100GB – Application
- 140GB – Database

For additional information, see [Disk Partitioning](#).

1. Create the database install directory `/u01` to mount the logical volume using the following commands:

```
mkdir /u01
chmod a+rwX /u01
```

2. Copy all contents of `/home` to the temporary partition using the following commands:

```
mkdir /mnt/resource
mkdir /mnt/resource/home_backup
cp -pr /home/* /mnt/resource/home_backup
```

3. Use any text editor to update the `/etc/fstab` file with the new mount points. The following example uses vim:

```
vim /etc/fstab
```

Add the following lines to the end of the file:

```
/dev/VG_rsa/LV_app /home ext4 defaults 0 0
/dev/VG_rsa/LV_db /u01 ext4 defaults 0 0
/dev/VG_rsa/LV_swap swap swap defaults 0 0
```

4. Apply the new mount points using the following command:

```
mount -a
```

5. Copy the `/home` data back to the new mount point using the following command:

```
cp -pr /mnt/resource/home_backup /home/
```

Disable Hardware Lock Elision

1. Use any text editor to add `/lib64/noelision` to the beginning of the `/etc/ld.so.conf` file to disable Hardware Lock Elision. The following example uses vim:

```
vim /etc/ld.so.conf
```

Add the following line to the beginning of the file:

```
/lib64/noelision
```

2. Update the LD configuration using the following command:

```
ldconfig
```

Install Required OS Packages

1. In a SuSE Enterprise Linux environment, install all required OS packages using the following command:

```
zypper install binutils gcc gcc48 glibc glibc mksh libaio1
libaio-devel libcap1 libstdc++48-devel libstdc++6 libstdc++-devel
libgcc_s1 make sysstat xorg-x11-driver-video xorg-x11-server
xorg-x11-essentials xorg-x11-Xvnc xorg-x11-fonts-core xorg-x11
xorg-x11-server-extra xorg-x11-libs xorg-x11-fonts libcap2
oracleasm-kmp-default javapackages-tools lcms2 bea-stax-api rhino
xmlbeans syslinux zip unzip bzip2
```

In a Red Hat Enterprise Linux environment, run the following prerequisite packages:

```
yum install binutils gcc gcc48 glibc glibc mksh libaio-devel
libstdc++-devel make sysstat javapackages-tools lcms2 bea-stax-
api rhino syslinux zip unzip bzip2 compat-libcap1-1.10-
7.el7.x86_64 compat-libstdc++-33-3.2.3-72.el7.x86_64 gcc-c++-
4.8.5-4.el7.x86_64 ksh iptables-services
```

2. Install the latest SuSE OS patches and packages by running following command:

```
zypper up
```

In a Red Hat Enterprise Linux environment, run the following command:

```
yum update
```

3. Download the following two OS packages from the Oracle's website:

<https://www.oracle.com/linux/downloads/linux-asmlib-sles12-downloads.html>

- a. oracleasm-support-2.1.8-1.SLE12.x86_64.rpm
- b. oracleasmlib-2.0.12-1.sle12.x86_64.rpm

Copy the two package files to any path on the virtual machine, then manually install them using the following commands as **root**:

```
rpm -i oracleasmlib-2.0.12-1.sle12.x86_64.rpm
```

```
rpm -i oracleasm-support-2.1.8-1.SLE12.x86_64.rpm
```

Inbound Port Rules

By default, the virtual machine is created with a limited number of permitted inbound ports. The following ports must be able to seamlessly access the RSA IG&L application services:

- 22 TCP for SSH access.
- 8443 TCP for secure UI and Web Services access.
- 8444 TCP for Remote Agents/AFX servers' access.
- 1555 TCP for database access.
- 8080 TCP for insecure UI and Web Services access.
- 443 TCP for secure UI and Web Services access after configuring the internal SuSE firewall for port forwarding.
- 444 TCP for Remote Agents/AFX servers' access after configuring the internal SuSE firewall for port forwarding.
- 80 TCP for insecure UI and Web Services access after configuring the internal SuSE firewall for port forwarding.

To allow additional port access:

1. On the left panel, click **Networking** under "Settings".
2. To add an extra inbound port rule, click **Add inbound port rule**.
3. On the left panel, click **Networking**, and then click on **NIC Public IP**.
4. On the left panel, under "Settings", click **Configuration**.
5. Add a DNS name label and click **Save**.

Local Firewall Port Forwarding

To use the default HTTP/HTTPS ports 80/443 to access the application, follow the steps under "Firewall Configuration for SUSE" in the *RSA Identity Governance and Lifecycle Installation Guide* for setting up port forwarding.

1. Use any text editor to modify the file `/etc/sysconfig/SuSEfirewall2` and create the required port forwarding settings.

```
vim /etc/sysconfig/SuSEfirewall2
```

Append the following port number string to `FW_SERVICES_EXT_TCP`:

```
FW_SERVICES_EXT_TCP="22 1555 8443 8444 8080"
```

Append the following port forwarding string to `FW_REDIRECT`:

```
FW_REDIRECT="0/0,0/0,tcp,443,8443 0/0,0/0,tcp,444,8444  
0/0,0/0,tcp,80,8080"
```

In a Red Hat Enterprise Linux environment, run the following:

```
firewall-cmd --zone=public --add-port=22/tcp --permanent  
firewall-cmd --zone=public --add-port=1555/tcp --permanent  
firewall-cmd --zone=public --add-port=8080/tcp --permanent  
firewall-cmd --zone=public --add-port=8443/tcp --permanent  
firewall-cmd --zone=public --add-port=8444/tcp --permanent  
firewall-cmd --reload
```

2. Run the following command(s) to apply these changes:

```
systemctl enable SuSEfirewall2 SuSEfirewall2_init
systemctl restart SuSEfirewall2 SuSEfirewall2_init
```

Remote Database

For remote database setup, create a virtual machine using the instructions in [Creating the Azure Virtual Machine](#) and [Pre-Installation configurations](#).

On the remote database machine, use **RSA_IGL_DatabaseOnly.tar.bz2** to install the database. For instructions, see the *RSA Identity Governance and Lifecycle Installation Guide*.

Add the following parameters in `/etc/sysctl.conf`.

```
net.ipv4.tcp_keepalive_time = 120
net.ipv4.tcp_keepalive_intvl = 80
net.ipv4.tcp_keepalive_probes = 8
```

Installing RSA IG&L v7.5.0

Copy the installation packages

1. Create the packages and staging directories using the following commands:

```
mkdir -p /tmp/aveksa/packages
mkdir -p /tmp/aveksa/staging
```

2. Use any application (for example: WinSCP) to copy the following install packages to `/tmp/aveksa/packages`:

- `adoptjdk_8u252b09.tar.gz`
- `aveksa-7.5.0.tar.bz2`
- `asmlib-008_x64.tar.bz2`
- `linuxx64_193000_db_home.zip`
- `linuxx64_193000_grid_home.zip`
- `oracle_19.0.0.0_patches_v001.zip`
- `wildfly-21.0.1.Final.tar.gz`

3. Extract the installation package to `/tmp/Aveksa/staging` using the following commands:

```
cd /tmp/aveksa/staging/
tar -jxvf ../packages/ aveksa-7.5.0.tar.bz2
```

Run the Installer

1. Run the installer using the following command. Press Enter to provide the default options for all decisions.

```
./install.sh -afx
```

2. Monitor the installation as normal until it completes successfully.

Note: If you are using a remote database, enter the remote database hostname details during the installation process.

Access the Application

Access the application through your browser using the public IP address or hostname.

Upgrade Standalone RSA Identity Governance and Lifecycle

To upgrade to RSA Identity Governance and Lifecycle 7.5.0:

1. Delete the files located in `/tmp/aveksa/staging` and `/tmp/aveksa/packages`
2. Use any application (for example: WinSCP) to copy the following install packages to `/tmp/aveksa/packages`:
 - `adoptjdk_8u252b09.tar.gz`
 - `aveksa-7.5.0.tar.bz2`
 - `asmlib-008_x64.tar.bz2`
 - `linuxx64_193000_db_home.zip`
 - `linuxx64_193000_grid_home.zip`
 - `oracle_19.0.0.0_patches_v001.zip`
 - `wildfly-21.0.1.Final.tar.gz`
3. Extract the installation package to `/tmp/Aveksa/staging` using the following commands:


```
cd /tmp/aveksa/staging/
tar -jxvf ../packages/ aveksa-7.5.0.tar.bz2
```
4. Run the installer using the following command. Press Enter to provide the default options for all decisions.


```
./install.sh -afx
```
5. Monitor the installation as normal until it completes successfully.

Note: If you are using a remote database, enter remote database hostname details during the installation process. The upgrade process indicates if you are still running on Oracle 12. If you are using an RSA license for the database, RSA recommends that you upgrade to 19C.

Chapter 3: Performance Metrics

This chapter summarizes performance tests run on the Microsoft Azure Cloud platform and the results for RSA Identity Governance and Lifecycle. ACME Corporation is a large, fictional enterprise that provides the basis for the test data and use cases. For more information about ACME, see [ACME Specification](#).

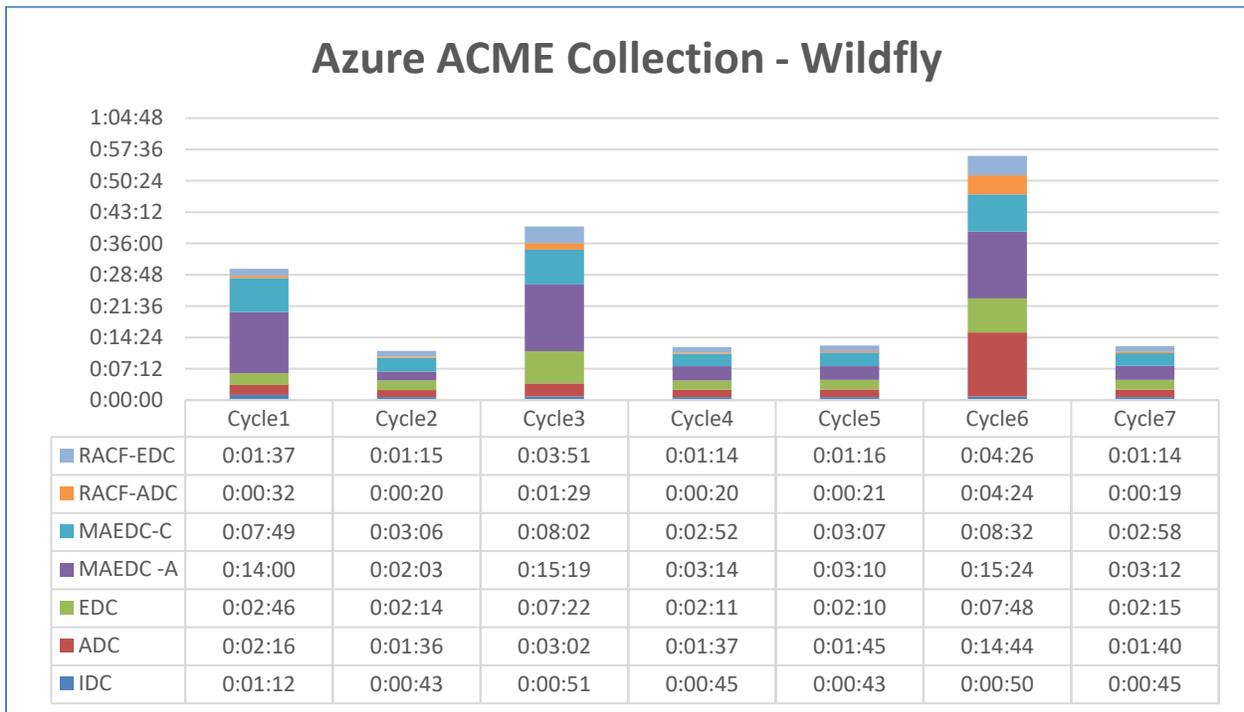
Azure Configuration

App Server Configuration	Remote Database Configuration
---------------------------------	--------------------------------------

OS : 12-sp4-gen2 Size : Standard D4as_v4 vCPUs : 4 RAM : 16 GiB Storage : 256 GB (Premium SSD)	OS : 12-sp4-gen2 Size : Standard E8s v3 vCPUs : 8 RAM : 64 GiB Storage : 512 GB (Premium SSD)
--	---

ACME Collections

The following chart shows the collection time for Azure configurations in a WildFly environment.



Performance Notes for Collections

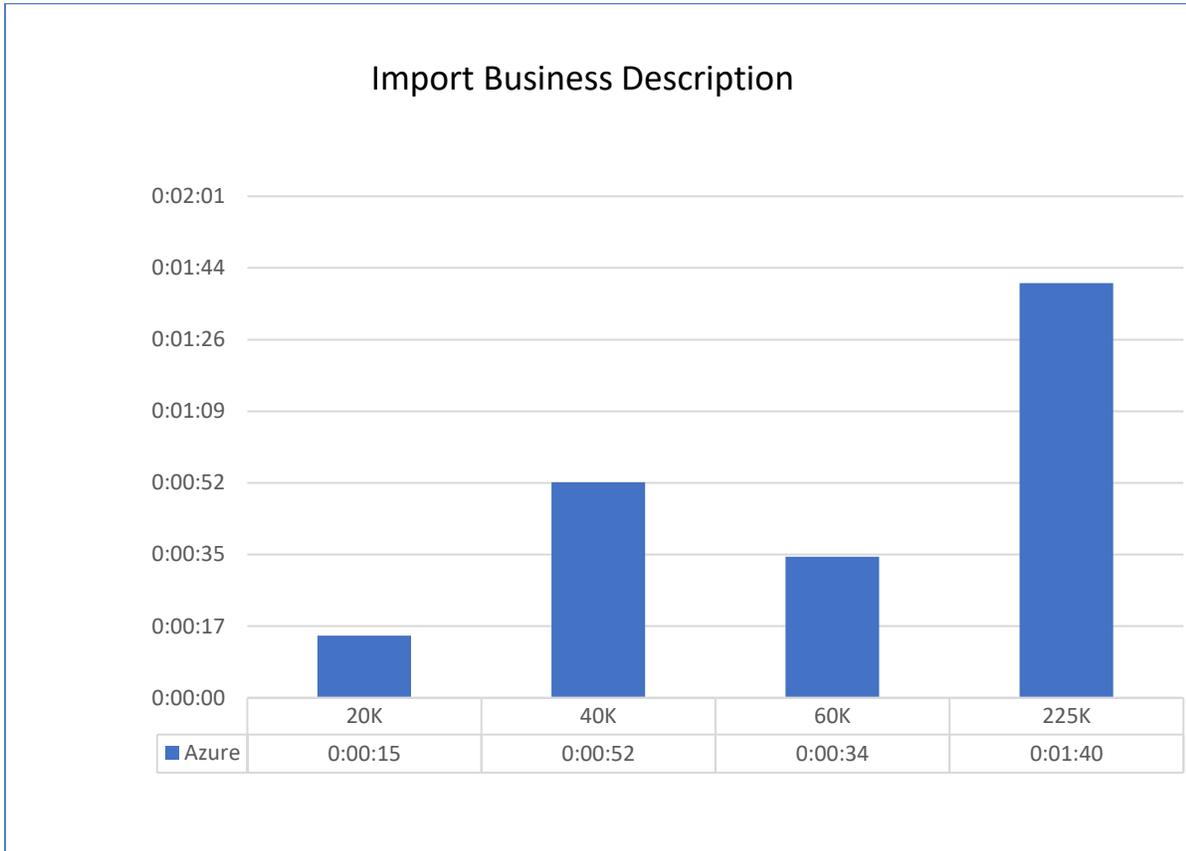
- IDC took almost 25% percent more time in the change cycles compared to the average of the other cycles. The unification process took the most time.
- ADC took over 70% more time in the change cycles compared to the average of the other cycles. Account data pre-processing took the most time.
- EDC and MAEDC each took almost 30 to 40% more time in the change cycles compared to the average of the other cycles.
- ADC and EDC time included RACF-ADC and RACF-EDC time respectively.
- EDCs and MAEDCs had the highest share in total collection time.
- The collection time remained consistent for the five cycles without the change.

Volume tests with Single User

The assorted performance scenarios for ACME involve a high volume of data processing from a single user. These tests ran on an appliance with a WildFly application server.

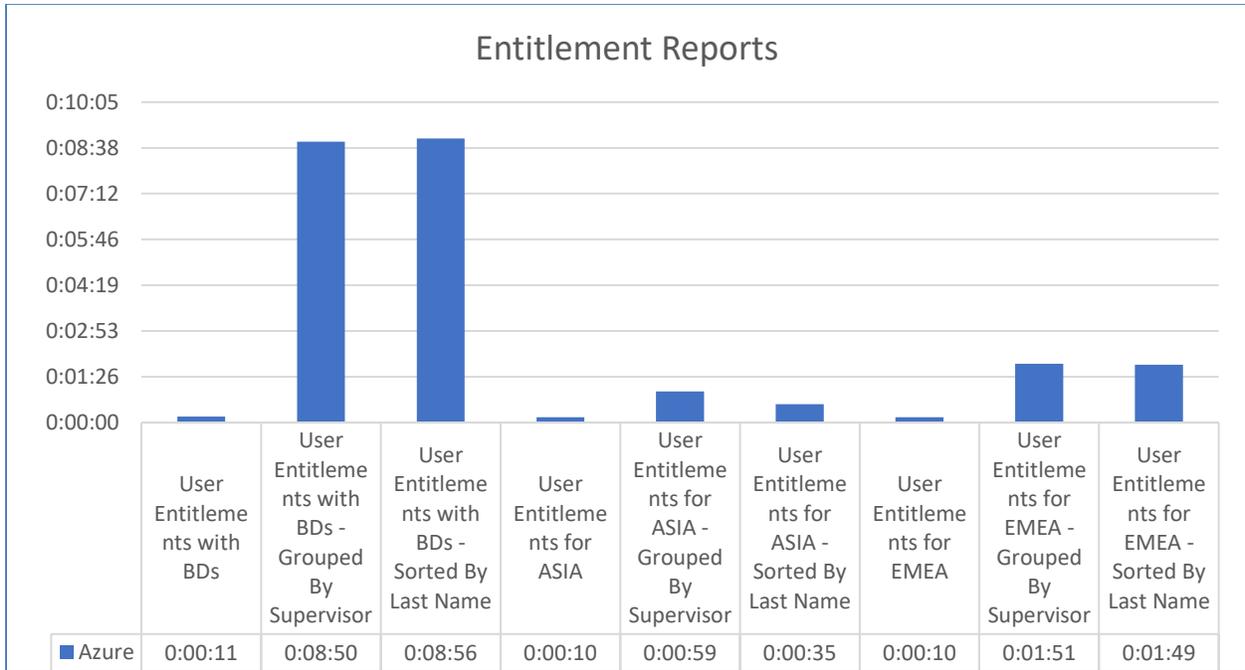
Business Description Imports

The following graph plots ACME's import time in Azure against the number of business descriptions in a text file.



Entitlement Reports

After unification, orphan account cleanup, and import of business descriptions, ACME runs a few basic reports on user entitlements and archives them in a CSV file. The following report shows the entitlement reports that ran against Azure environment.

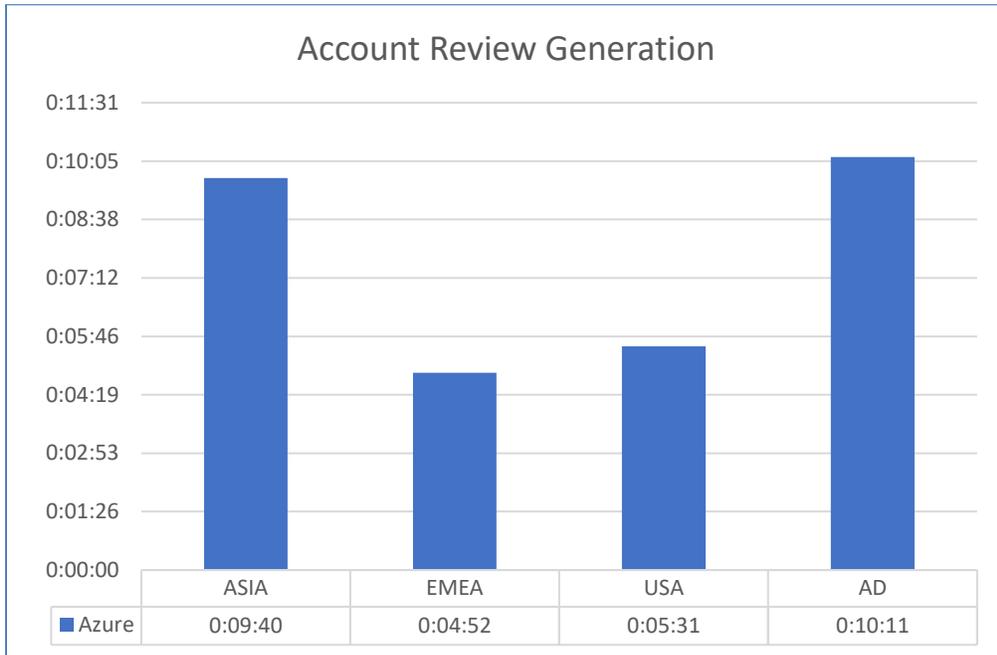


Reviews

After collecting the necessary access information, ACME runs multiple recertification processes to verify access and remove it as needed. ACME needs to plan its maintenance window accordingly.

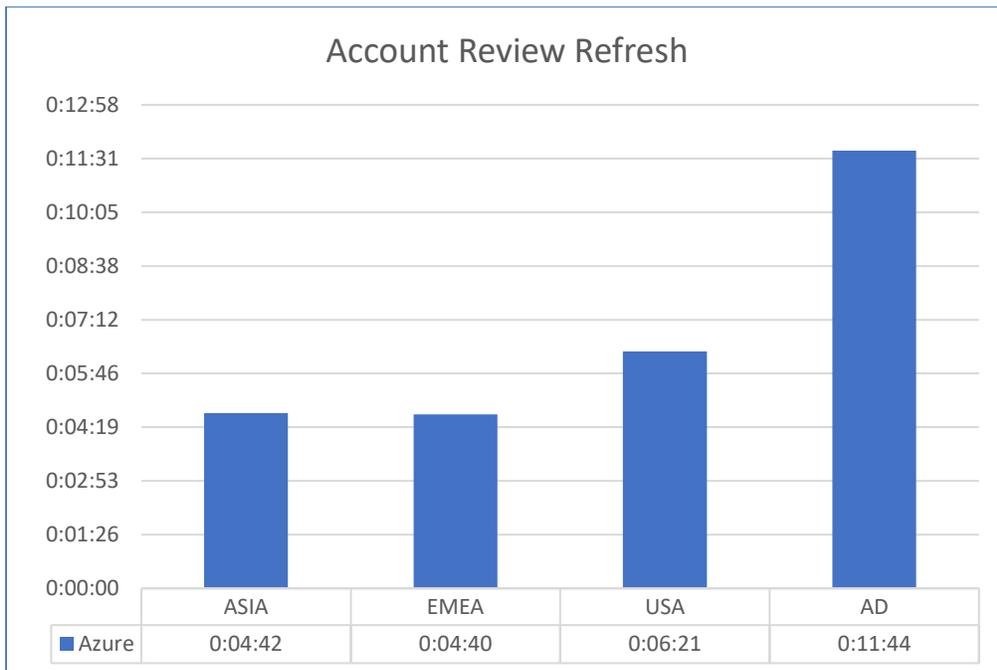
Annual Review Generation Performance

The ACME system generates three annual reviews for four different business units.



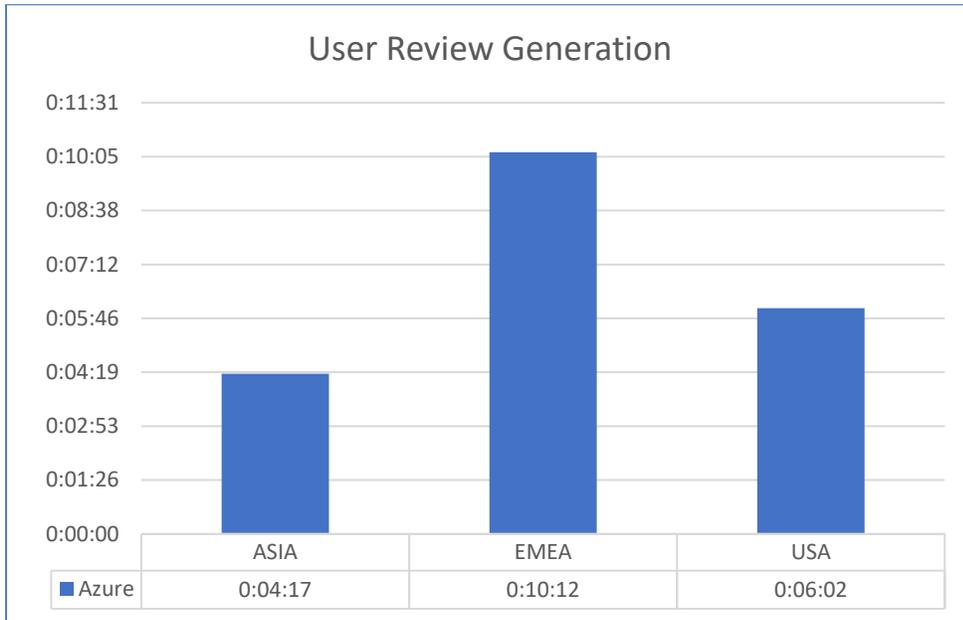
Annual Review Refresh

ACME refreshes the comprehensive, annual reviews once a week during the eight-week cycle. ACME requires to know the time duration taken if 5% of the entitlements change on every refresh.



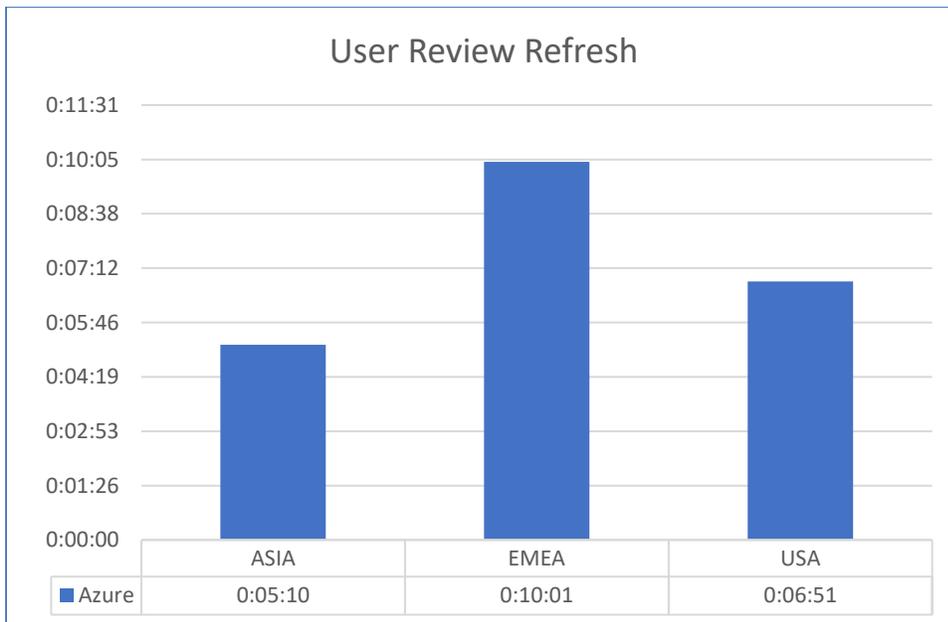
User Review Generation Performance

The ACME system generates three user reviews for three different business units.



User Review Refresh

ACME refreshes the comprehensive user reviews once a week during the eight-week cycle. ACME requires to know the time duration taken, if 5% of the entitlements change on every refresh.

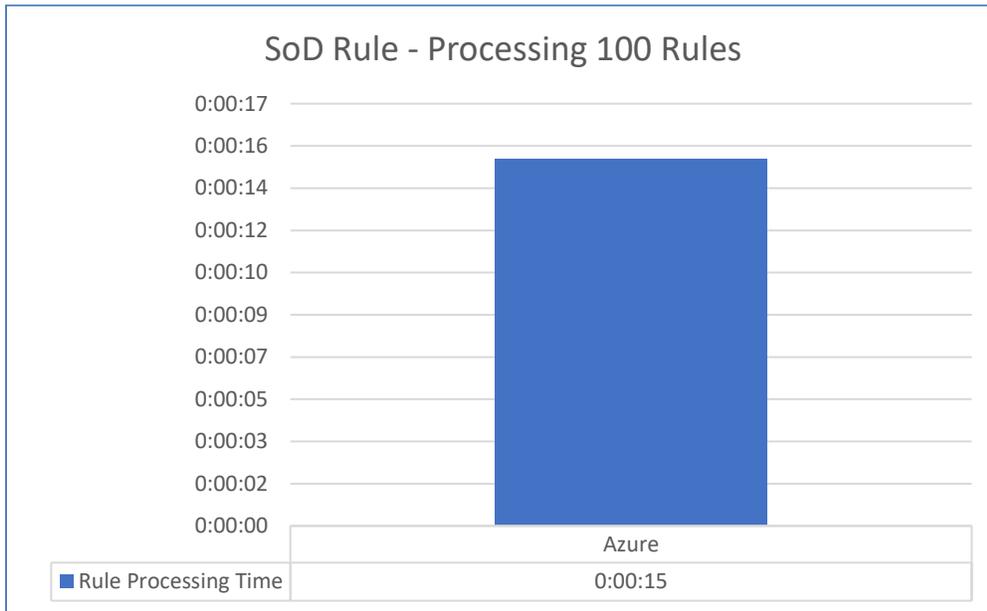


Rules and Roles

Separation of Duty (SoD) Rule Processing

The following chart shows the rule processing time for 100 SoD rules.

Processing time of 100 SOD rules



SoD Rule Processing with correlation condition

When processing SOD, an optional constraint is defined for business source attributes. For SOD Rule definition, the user can choose not to select any Correlation Specification Options or can select any of the three options. The logic associated with the options is as follows:

- Match for exact values between the set1 and set2.
- One of the values matches (Match for any of the common value between set1, set2 and rule definition value).
- System detects at least one common value (Match for any of the common value between set1 and set). Performance testing of the feature consists of defining SoD rule with the correlation condition and processing that rule.

The following table shows SoD with different correlation condition options:

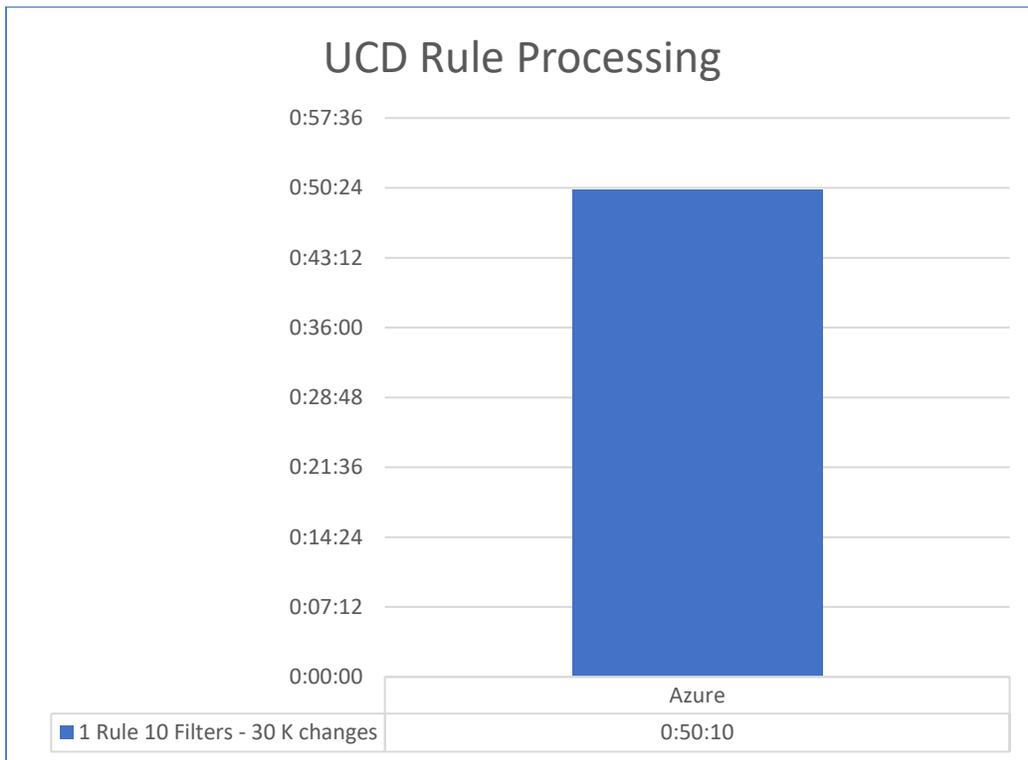
Correlation Specification	Number of Violations	Azure Instance
Name - Contains identical values	298709	0:08:21
Classification - At least one of the value matches	249793	0:04:04
Classification - System detects at least one	249793	0:04:14

Unauthorized Change Detection (UCD) Rule Processing

UCD checks the changes in account access found through collections against the change requests in RSA Identity Governance and Lifecycle. Changes without a change request on record are flagged as unauthorized.

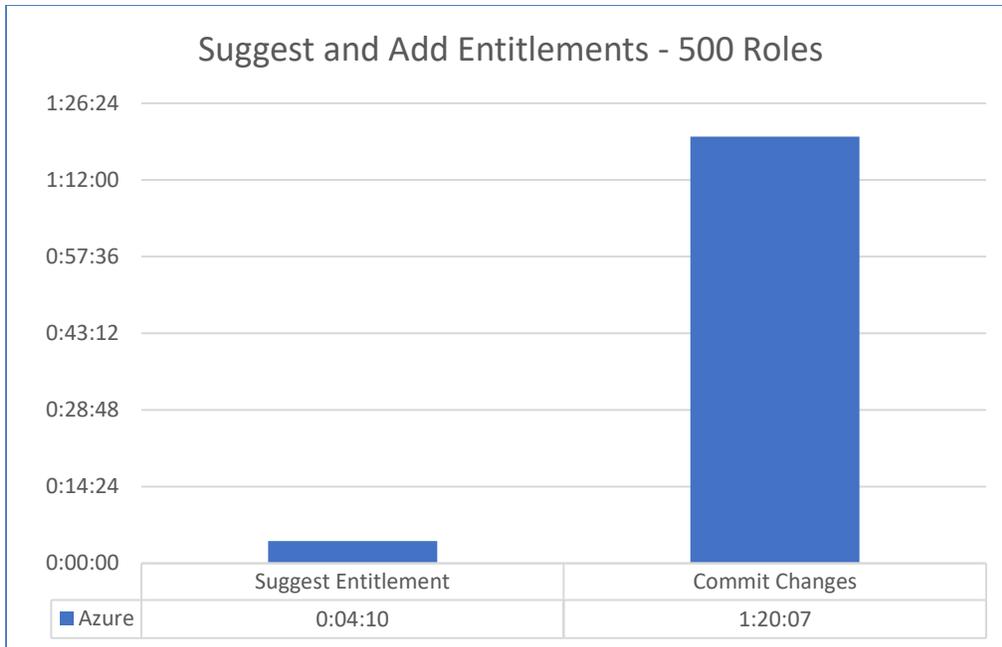
1 Rule – 10 Filters

The ACME system applies one UCD rule with 10 filters to detect unauthorized access in any of the 10 applications. Collections run to detect additional unauthorized access. The ACME system generates a change request for each unauthorized change to revoke access. Tests for 30,000 unauthorized changes run with no other change requests in queue.



Suggest Entitlements for Roles

An ACME Role Analyst performs the ‘Suggest and add entitlements’ operation for each of the roles that were created for each business unit. The following table records the time taken to commit the changes to 500 roles.



500 concurrent users test on Azure

The following sections describe the results of performance tests using 500 concurrent users on the Azure deployment.

Test Details: The load pattern of this test is as follows:

- 300 users gradually ramped up over the course of 30 minutes, followed by 20 minutes of steady state.
- Added another 100 users with a ramp up of 10 minutes, followed by 20 minutes of steady state.
- Added the final group of 100 users with a ramp up of 10 minutes, followed by 30 minutes of steady state with the total of 500 users.

Access Requests on Azure with 500 concurrent users

Transaction	Response Time (seconds)			
	Minimum	Average	Maximum	95 th Percentile
ACM Home Page	0.012	0.046	13.285	0.025
Login	0.144	0.569	24.673	1.069
Click Manage Access	0.037	0.103	14.473	0.08
Click Add Others Access	0.085	0.291	15.911	0.634
Group By Subordinates	0.077	0.297	15.807	0.662
Select Users	0.009	0.122	30.171	0.023
Click Add Entitlements	0.436	1.798	38.837	4.996
Group Entitlements By	0.601	2.007	15.451	6.364

Goto Random Page	0.25	0.774	19.208	2.311
Select An Entitlement	0.115	0.397	11.125	1.205
Submit Access Request	0.064	0.557	19.432	0.924
Click Ok Request	0.02	0.195	16.577	0.193
Click Home Tab	0.008	0.05	15.488	0.017
Click Logout	0.005	0.067	15.048	0.01
Confirm Logout	0.028	0.064	15.107	0.077

Observations:

- The average response time for all access request transactions was less than three seconds.
- The user CPU usage of the ACM host with a 500-user load was approximately 23%, and the usage of the database host at the same load was approximately 77%.

User Reviews on Azure with 500 concurrent users

Response Time (seconds)				
Transaction Name	Minimum	Average	Maximum	95th Percentile
ACM Home Page	0.007	0.029	13.684	0.02
Login	0.066	0.23	17.025	0.298
Click Review Link	0.012	0.045	16.796	0.031
Click Perform Review	0.192	0.339	16.13	0.551
Navigate to Random Entitlement Page	0.035	0.292	15.462	0.539
Single Entitlement Review Maintain	0.047	0.103	16.774	0.119
Single Entitlement Review Revoke	0.002	0.032	16.659	0.009
Logout	0.016	0.053	15.527	0.047
Confirm Logout	0.174	0.319	16.963	0.386

Observations:

- The average response time for all user review transactions was less than one second.
- The user CPU usage of the ACM host with a 500-user load was approximately 36%, and the usage of the database host at the same load was approximately 19%.

AFX Performance Results

AFX performance was evaluated with an Oracle database (DB) connector. All AFX tests had a single database connector created in AFX, and the operations generated thousands of fulfillment requests to help test AFX scalability for a large customer reference profile.

AFX Test Scenarios

- Joiner (Implicit) –Attribute Change Rule - Large Data - 18,000 new users
- Mover (Implicit) –Attribute Change Rule - Large Data - 18,000 users

- Leaver (Implicit) – Attribute Change Rule - Large Data - 18,000 users

Joiner (Implicit) – Attribute Change Rule

This scenario processes a Joiner attribute change defined to create an account on the Oracle DB application and then adds two entitlements for every new user detected. The test simulates 18,000 new users in a 250,000-user system triggered by an attribute change.

KPI	Oracle DB Connector
Total Test Execution time (h:mm:ss)	2:44:24
Collection Time (h:mm:ss)	0:00:06
Unification Time (h:mm:ss)	0:00:18
Rule Processing Time (h:mm:ss)	0:15:22
Time for generation of 18,000 Change Requests (h:mm:ss)	0:17:15
Number of fulfillment requests (Complete/Total)	54000/54000
Total Time for fulfillment (h:mm:ss)	2:43:17
Throughput (requests fulfilled per second)	5.51
Average – Percent of CPU utilization on DB Server	5.1
Average - Available Memory on DB Server (GB)	138
Average – Percent of CPU utilization on ACM Server	4.2
Average - Available Memory on ACM Server (GB)	95.7

Mover (Implicit) – Attribute Change Rule

This scenario processes a Mover attribute change rule defined to add three entitlements for the account on the Oracle DB application for all existing users having the “Department” change. The test simulates a change in entitlements for 18,000 users.

KPI	Oracle DB Connector
Total Test Execution time (h:mm:ss)	2:24:43
Collection Time (h:mm:ss)	0:00:06
Unification Time (h:mm:ss)	0:00:20
Rule Processing Time (h:mm:ss)	0:29:33
Time for generation of 18,000 Change Requests (h:mm:ss)	0:29:31
Number of fulfillment requests (Complete/Total)	54000/54000
Total Time for fulfillment (h:mm:ss)	2:23:51
Throughput (requests fulfilled per second)	6.2
Average – Percent of CPU utilization on DB Server	23.7
Average - Available Memory on DB Server (GB)	29.7
Average – Percent of CPU utilization on ACM Server	41.9
Average - Available Memory on ACM Server (GB)	0.3

Leaver – Provisioning Termination Rule

This scenario processes the Provisioning Termination rule defined to delete the account from the Oracle database application for all terminated users. The test simulates account deletions for 18,000 terminated users.

KPI	Oracle DB Connector
Total Test Execution time (h:mm:ss)	1:58:54

Collection Time (h:mm:ss)	0:00:03
Unification Time (h:mm:ss)	0:02:02
Rule Processing Time (h:mm:ss)	0:00:00
Time for generation of 18,000 Change Requests (h:mm:ss)	0:12:55
Number of fulfillment requests (Complete/Total)	18000/18000
Total Time for fulfillment (h:mm:ss)	1:55:47
Throughput (requests fulfilled per second)	2.6
Average – Percent of CPU utilization on DB Server	18.6
Average - Available Memory on DB Server (GB)	31.7
Average – Percent of CPU utilization on ACM Server	22.5
Average - Available Memory on ACM Server (GB)	0.3

Chapter 4: ACME Specification

Identity Information

Total Business Units: 4 (USA, EMEA, ASIA, Latin America)

Total Users: 250,000 (225,000 users are employees. 25,000 users are contractors)

- USA -149,000 users (1 supervisor for every 10 users) - 200 job titles.
- ASIA – 10,000 users (1 supervisor for every 2 users) - 500 job titles.
- EMEA – 90,000 users (1 supervisor for every 11 users) - 2000 job titles.
- Latin America – 1,000 users (1 supervisor for every 10 users) - 50 job titles.

Account Information

ACME has 675,000 accounts distributed across the following account data collectors (ADCs):

- ADC_AD - Covers 250,000 users and 300,000 accounts (One account per each user and 50,000 orphan/system accounts), 60,000 groups. Group1 has 50,000 users. Groups 2-11 have 2,000 users. The remaining groups have three users per group.
- ASIA_ADC - Covers 10,000 users and 11,000 accounts (One account per user in the ASIA business unit and 1,000 orphan accounts).
- EMEA_ADC - Covers 90,000 users and 99,000 accounts (One account per user in the EMEA business and 9,000 orphan accounts).
- USA1_ADC - Covers 70,000 users and 77,000 accounts (One account per user and 7,000 orphan accounts).
- USA2_ADC - Covers 80,000 users and 88,000 Accounts (One account per user and 8,000 orphan Accounts).
- RACF_ADC - Covers 10,000 users across all business units and 100,000 accounts (10,000 orphan accounts).

Application Information

ACME has 1,000 applications under management:

- Across all 1,000 applications there are 15 million user entitlements, including 20,000 application roles. ACME application entitlements are managed as follows:

- **App Entitlement Database A:** Has entitlements for 500 applications and 5 million user entitlements (including application roles) across applications.
 - 100 of the applications have 100 application roles each.
 - Applications are marked as SOX (classification attribute).
 - 400 of the applications have 20 application roles each.
 - Exploded entitlements: 10,996,701.
 - Granular entitlements: 4,498,684.
- **App Entitlement Database C:** Has entitlements for 210 applications and 1 million user entitlements across applications.
 - Exploded entitlements: 808,000.
 - Granular Entitlements: 808,000.
- **10 Individual Application Databases:** Each has 100,000 user entitlements.
 - **HugeApplication1 – 4:** 183,750 Exploded entitlements, 87,500 granular entitlements (each).
 - **HugeApplication5 – 7:** 189,000 Exploded entitlements, 90,000 granular entitlements (each).
 - **HugeApplication8 – 9:** 190,910 Exploded entitlements, 90,910 granular entitlements (each).
 - **HugeApplication10:** 190,908 Exploded entitlements, 90,908 granular entitlements.

RACF Mainframe - ACME's mainframe environment has 7.03 million user entitlements. The mainframe environment has 10,000 user profiles that have access to 100,000 mainframe accounts.

- ACME maintains a separate database repository to manage associated business descriptions for all granular entitlements in the organization. This business description repository is updated quarterly by the application data stewards.
- App roles: 10,000
- App roles members: 222,500
- Exploded entitlements: 12,673,000
- Granular entitlements: 5,670,000

Chapter 5: Knowledgebase

Disk Partitioning

Run the following steps to partition the disk:

1. Run fdisk to partition the /dev/sdc data disk using the below command:

```
fdisk /dev/sdc
```

2. Create a new partition using the "n" command, then press Enter to accept the default inputs for the remaining options.
3. Write changes to the disk using the "w" command.
4. Create a physical volume with /dev/sdc1 using the below command:

```
pvcreate /dev/sdc1
```

5. Create a volume group with /dev/sdc1 using the below command:

```
vgcreate VG_rsa /dev/sdc1
```

6. Create three logical volumes from the volume group using the below commands:

Note: The swap size of 16777216KB is the exact value used in our installation inspect script:

```
/tmp/aveksa/staging/deploy/ENV-setup-scripts/InspectSystem.sh.
```

The final logical volume uses the 100%FREE size option to ensure using all remaining storage.

```
lvcreate -L 16777216K -n LV_swap VG_rsa
```

```
lvcreate -L 20G -n LV_app VG_rsa
```

```
lvcreate -l 100%FREE -n LV_db VG_rsa
```

7. Confirm all three created logical volumes using the below command:

```
lvs
```

8. Format the swap logical volume using the below command:

```
mkswap /dev/VG_rsa/LV_swap
```

9. Enable swap on the new swap logical volume using the below command:

```
swapon -v /dev/VG_rsa/LV_swap
```

10. Format the database and application logical volumes as ext4 using the below commands:

```
mkfs.ext4 /dev/VG_rsa/LV_app
```

```
mkfs.ext4 /dev/VG_rsa/LV_db
```