

HOL-2401-03-CMP

# Becoming A Power User (Advanced)

## Table of contents

Lab Overview - HOL-2401-03-CMP - Aria Operations - Becoming A Power User (Advanced)	4
Lab Description.....	4
Lab Guidance .....	4
Module 1 - Configuring and Managing Alert Notifications (35 minutes)	
Advanced	6
Introduction.....	6
Log in to Aria Operations.....	6
Notification using Webhooks .....	9
Notifications using email.....	19
Payload templates .....	27
Conclusion.....	41
Module 2 - Creating a Custom Alert Definition (25 minutes) Intermediate	43
Introduction.....	43
Log in to Aria Operations.....	43
Custom Alert Definition .....	46
Custom Notifications .....	56
Conclusion.....	79
Module 3 - Application Monitoring with VMware Aria Operations (35 minutes)	
Advanced	80
Introduction.....	80
Introduction to Actions .....	81
Log in to Aria Operations.....	82
Script Execution and Top Processes.....	85
Conclusion.....	111
Module 5 - Workload Placement – Running Host Based Optimization (35 minutes)	
Intermediate	112
Introduction.....	112
Log in to Aria Operations.....	113
Business Intent .....	117
Conclusion.....	149
Module 6 - Report Generation in VMware Aria Operations (15 minutes) Basic	150
Introduction.....	150

Log in to Aria Operations.....	150
Introduction to Reports .....	152
Creating Custom Reports .....	168
Conclusion.....	193
Module 7 - Creating Custom Dashboards for VMware Aria Operations (15 minutes) Basic	194
Introduction.....	194
Log in to Aria Operations.....	194
Clone and Modify Existing Dashboards.....	196
Creating a New Custom Dashboard .....	248
Importing New Dashboards.....	289
Sharing Dashboards.....	307
Conclusion.....	320
Module 8 - Enhancing depth of VMware Aria Operations with Super Metrics (15 minutes) Basic	322
Introduction.....	322
Log in to Aria Operations.....	322
Understanding Hierarchical Relationships in Aria Operations .....	324
Create Your First Super Metric .....	332
Using 'This' and Negative Depth Parameters in Super Metric Formulas .....	363
Handling Sets of Data Points in a Super Metric Formula .....	396
Using String Operators and the "Where" Clause in a Super Metric Formula.....	417
Verifying Super Metric Calculation.....	439
Conclusion.....	443
Conclusion	444
Learning Path Next Steps! .....	444

## Lab Overview - HOL-2401-03-CMP - Aria Operations - Becoming A Power User (Advanced)

### Lab Description

[2]

Get familiar with Custom Alerts and Notifications. Use Application Monitoring features and Workload Placement for optimization. Explore Custom Dashboards, Reports, and Super Metrics.

### Lab Guidance

[3]

Welcome! This lab is available for you to repeat as many times as you want. To start somewhere other than the beginning, use the Table of Contents in the upper right-hand corner of the Lab Manual or click on one of the modules below.

- [Module 1 - Configuring and Managing Alert Notifications](#) (30 minutes) (Advanced)
- [Module 2 - Creating a Custom Alert Definition](#) (30 minutes) (Advanced)
- [Module 3 - Application Monitoring with VMware Aria Operations](#) (30 minutes) (Advanced)
- [Module 5 - Workload Placement – Running Host Based Optimization](#) (30 minutes) (Advanced)
- [Module 6 - Report Generation in VMware Aria Operations](#) (15 minutes) (Basic)
- [Module 7 - Creating Custom Dashboards for VMware Aria Operations](#) (15 minutes) (Basic)
- [Module 8 - Enhancing depth of VMware Aria Operations with Super Metrics](#) (15 minutes) (Basic)

#### Lab Captains:

- Module 1 - Bengt Grønås, Senior Specialist Solution Engineer, Norway
- Module 2 - Bengt Grønås, Senior Specialist Solution Engineer, Norway
- Module 3 - Bengt Grønås, Senior Specialist Solution Engineer, Norway
- Module 5 - Bengt Grønås, Senior Specialist Solution Engineer, Norway
- Module 6 - Greg Sylvestre, Senior Solution Engineer, US
- Module 7 - Greg Sylvestre, Senior Solution Engineer, US
- Module 8 - Greg Sylvestre, Senior Solution Engineer, US

This lab manual can be downloaded from the Hands-on Labs document site found here:

<http://docs.hol.vmware.com>

This lab may be available in other languages. To set your language preference and view a localized manual deployed with your lab, utilize this document to guide you through the process:

<http://docs.hol.vmware.com/announcements/nee-default-language.pdf>

### First time using Hands-on Labs?

[4]

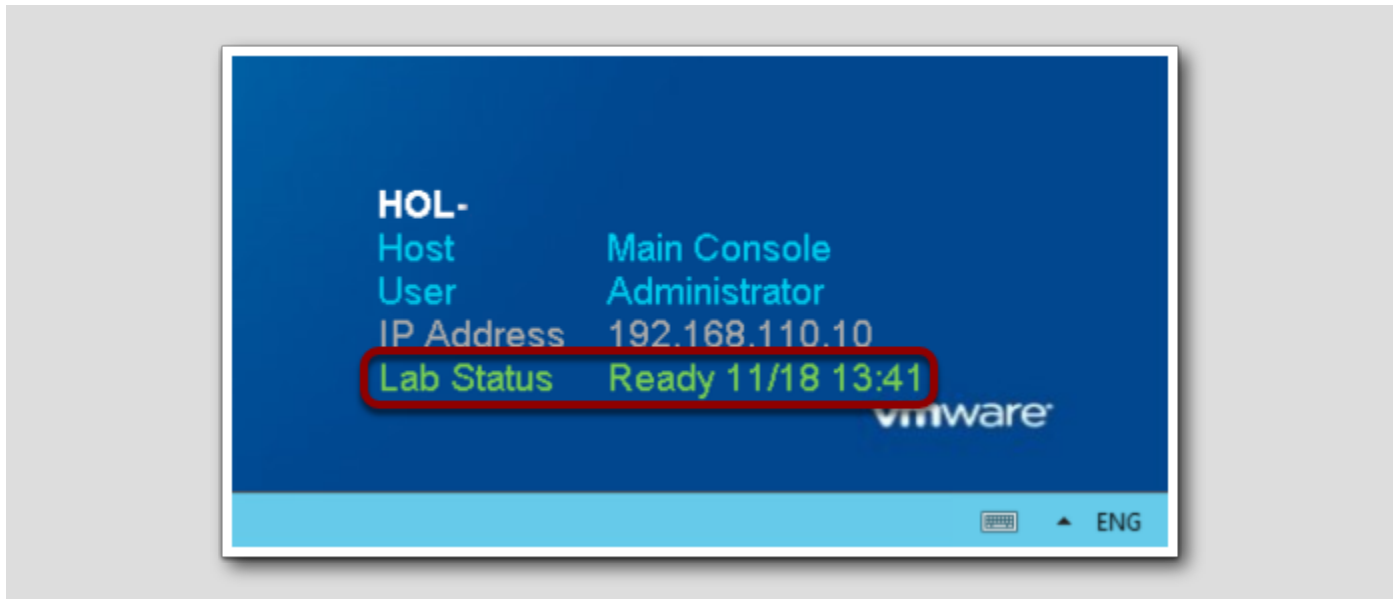
Welcome! If this is your first time taking a lab review the VMware Learning Platform interface and features before proceeding.

For returning users, feel free to start your lab by clicking next in the manual.



You are ready....is your lab?

[5]



The lab console will indicate when your lab has finished all the startup routines and is ready for you to start. If you see anything other than "Ready", please wait for the status to update. If after 5 minutes your lab has not changed to "Ready", please ask for assistance.

## Module 1 - Configuring and Managing Alert Notifications (35 minutes) Advanced

### Introduction

[7]

Upon completing this lab, you will be able to:

- Create a Notification using a Webhook and email.
- Customize a payload template.

**Creating notifications** using WebHooks or email serves many purposes. By leveraging notifications, you stay informed, collaborate effectively, respond promptly, and continuously improve your IT operations. It's the key to a proactive, integrated, and optimized environment.

**Real-time Notifications:** Webhooks facilitate real-time notifications, ensuring that relevant parties receive alerts promptly. This timely delivery helps facilitate faster incident response, reducing downtime and minimizing the impact on your environment.

A **payload template** acts as a customization blueprint for webhook notifications. It allows us to design the structure and content of the payload sent to external systems or applications. You can populate the payload with key information from the triggering event, ensuring it meets the recipient's requirements and enables informed actions. We can include specific information, such as alert details, impacted objects, timestamps, or any other relevant data, providing comprehensive context to aid in incident investigation and resolution.

Think of it as your creative tool to curate valuable and tailored data, empowering external systems to make swift and effective responses.

Note: This Lab is mostly on how to prepare for creating Alerts

### Log in to Aria Operations

[8]

We will log in to a live instance of Aria Operations running in this lab.

### Open the Firefox Browser from the Windows Task Bar

[9]

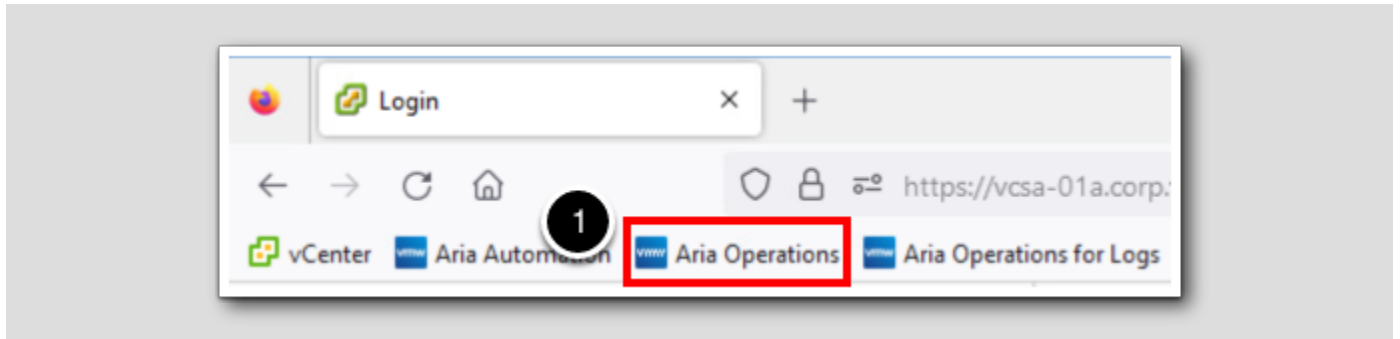


If the browser is not already open, launch **Firefox**.

1. Click the **Firefox** icon in the Windows Quick Launch Task Bar at the bottom of the screen.

## Navigate to Aria Operations


[10]



1. Click the Aria Operations bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[1]




Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login

[12]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

1. Click **Sign in**

## Notification using Webhooks

[13]

**Webhooks** enhance Aria Operations by providing real-time, customizable alert mechanisms for promote proactive incident management

When a specified alert condition is met, the webhook notifications trigger an HTTP POST request containing a JSON payload. This payload encapsulates alert details, name, criticality, and other relevant data. The destination of these notifications is set as a URL

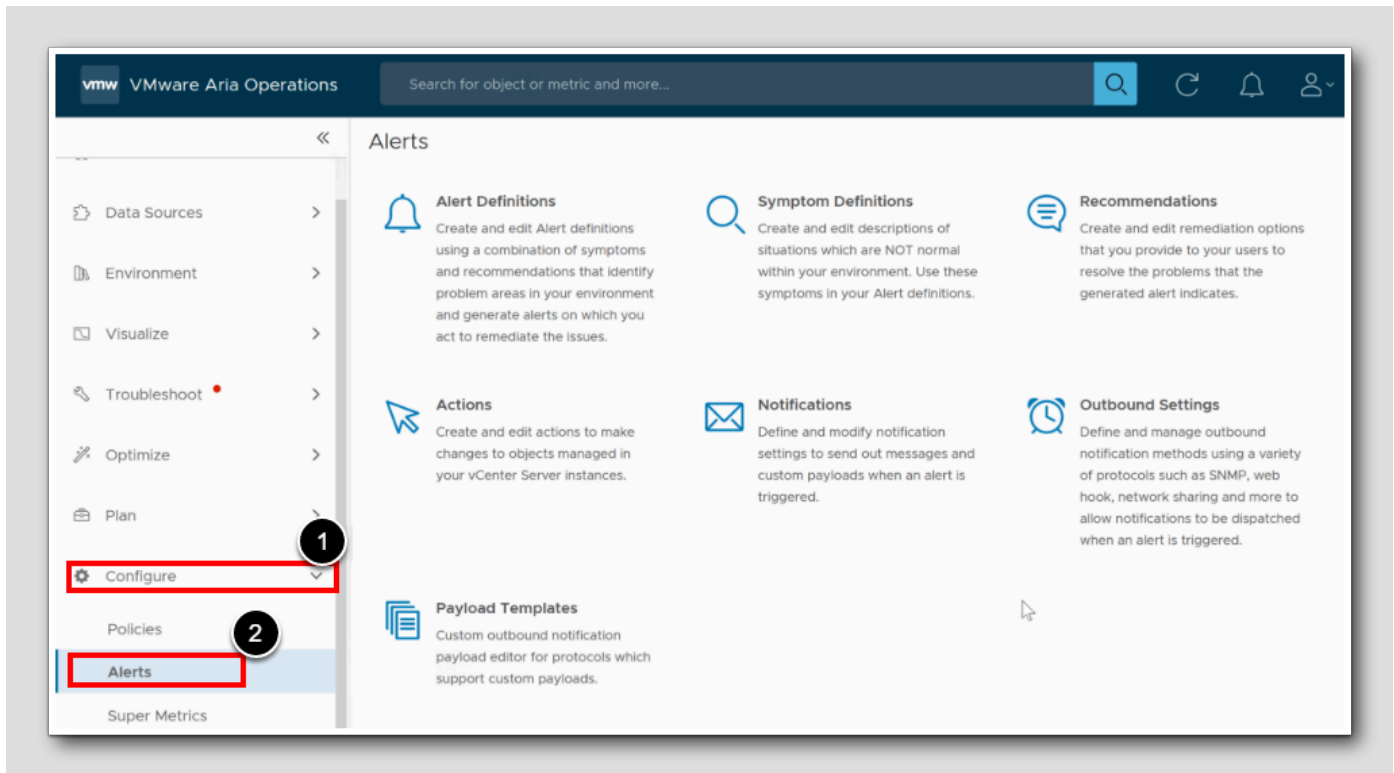
endpoint.

Webhooks provide substantial flexibility, we can configure webhooks to trigger remediation actions, integrate with ticketing systems, or notify on-call staff via communication tools.

Let's have a look on how to set up outgoing notifications with webhooks

## The Alerts Page

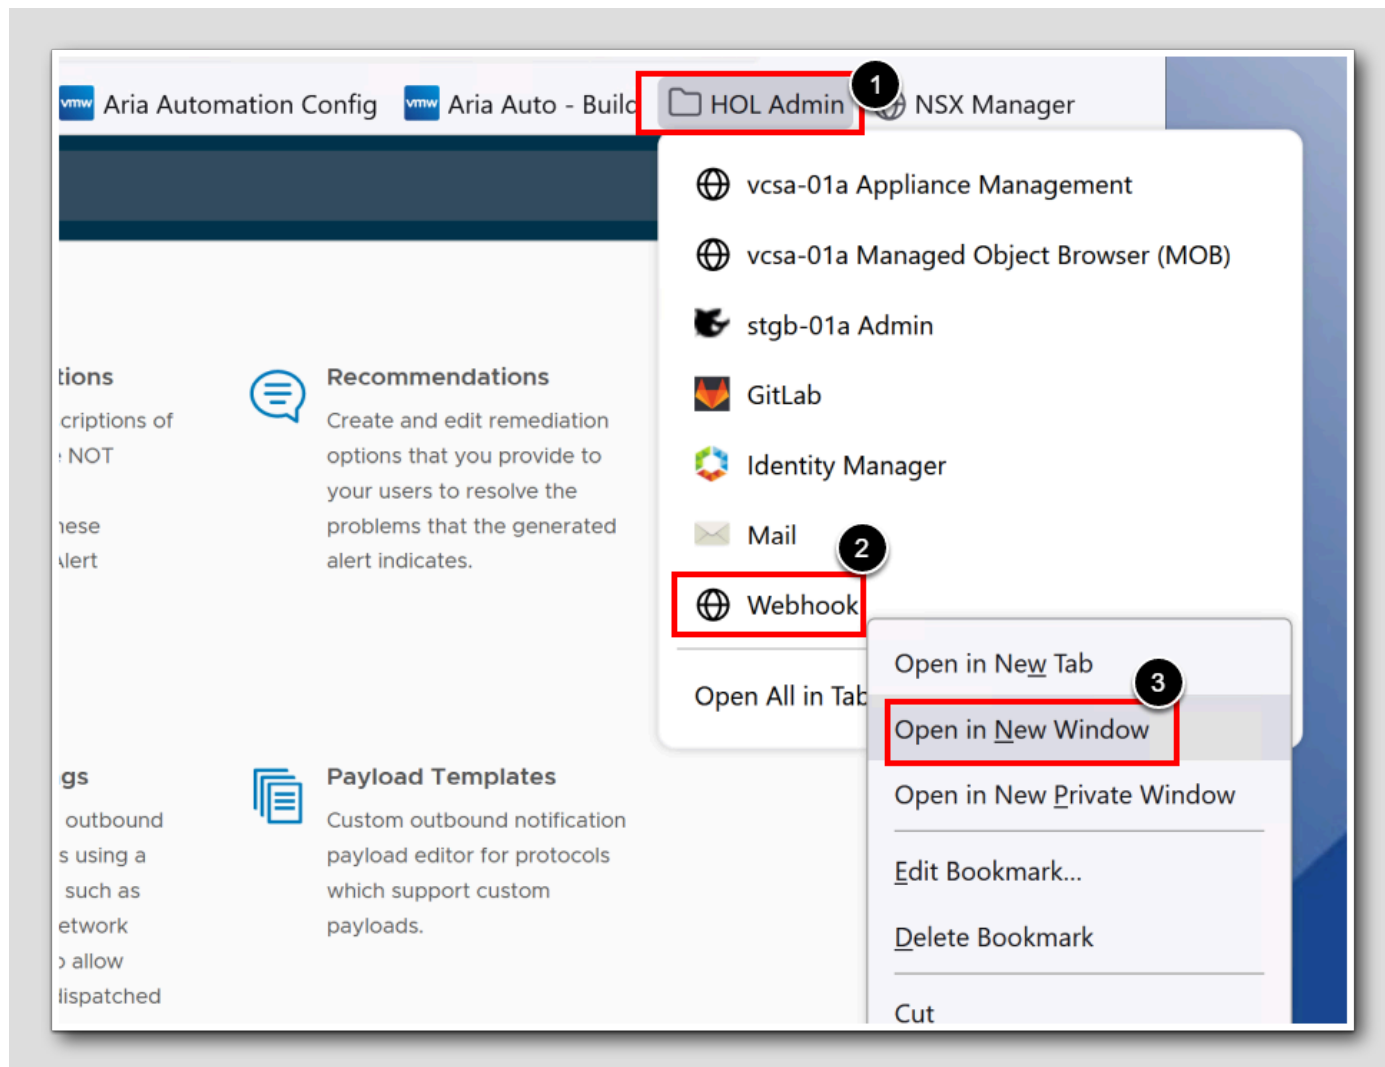
[14]



Let's find our alerts page where we will be configuring notifications using Webhook

1. In the left menu, click **Configure**
2. Click **Alerts**

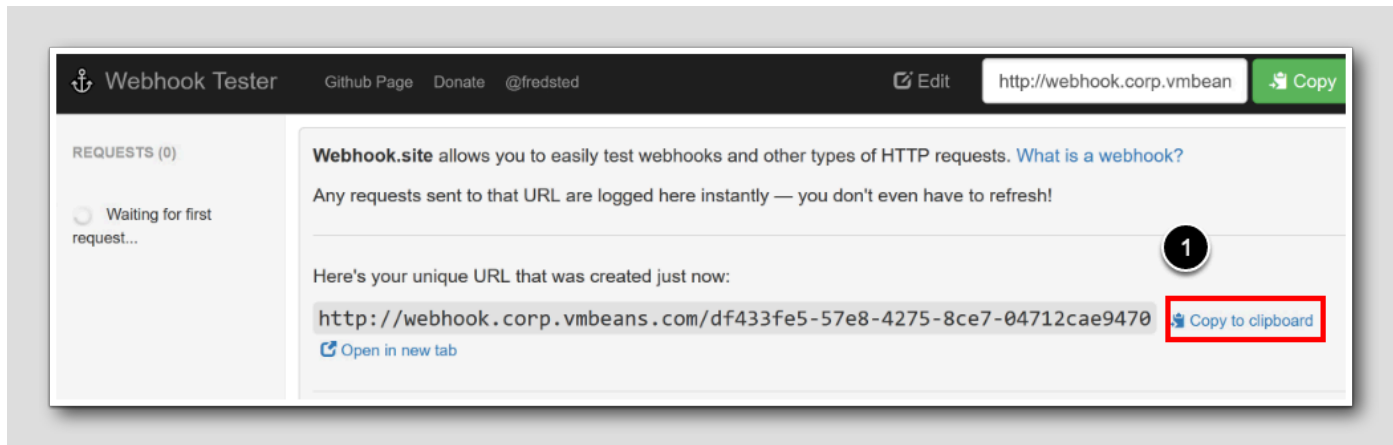
## The Webhook receiver



We are going to find the address for the webhook receiver we have enabled in the lab. We will have to use that address when we are going to configure Aria Operations to send webhooks to it.

1. On the Firefox menu, click **HOL Admin**
2. Next, let's Right-Click **Webhook**
3. Select **Open in New Window**

## Webhook Tester unique URL



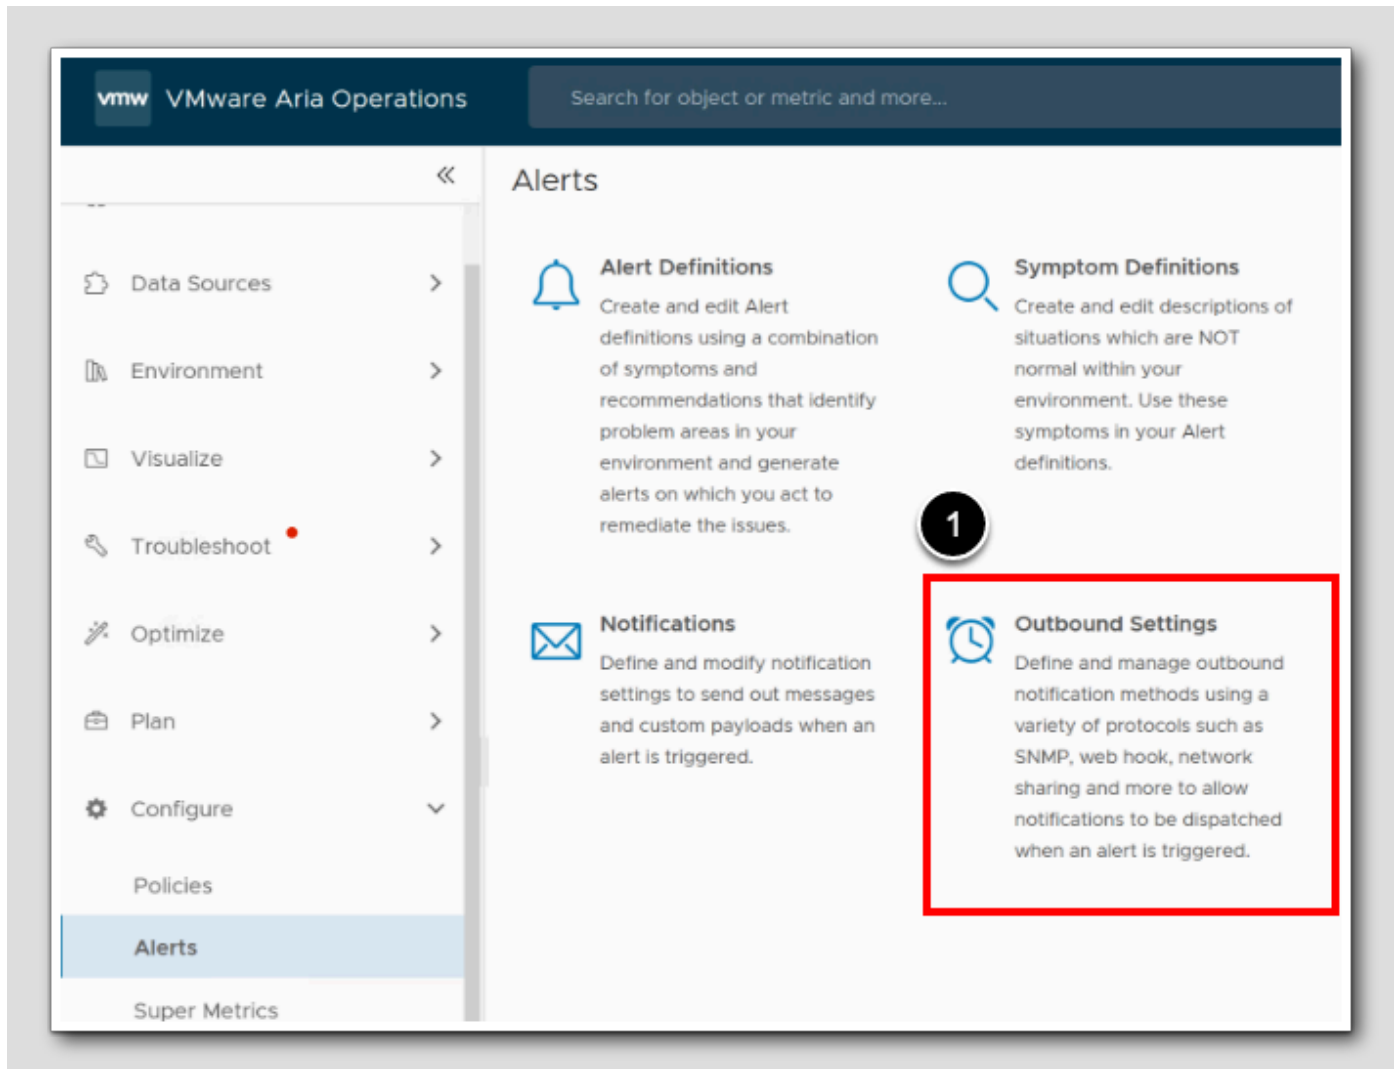
1. On the URL that was created for us, click **copy to clipboard**

Note: Do NOT close the window,

Use **ALT+TAB** or the taskbar to get back to the Aria Operations window (*not shown*)



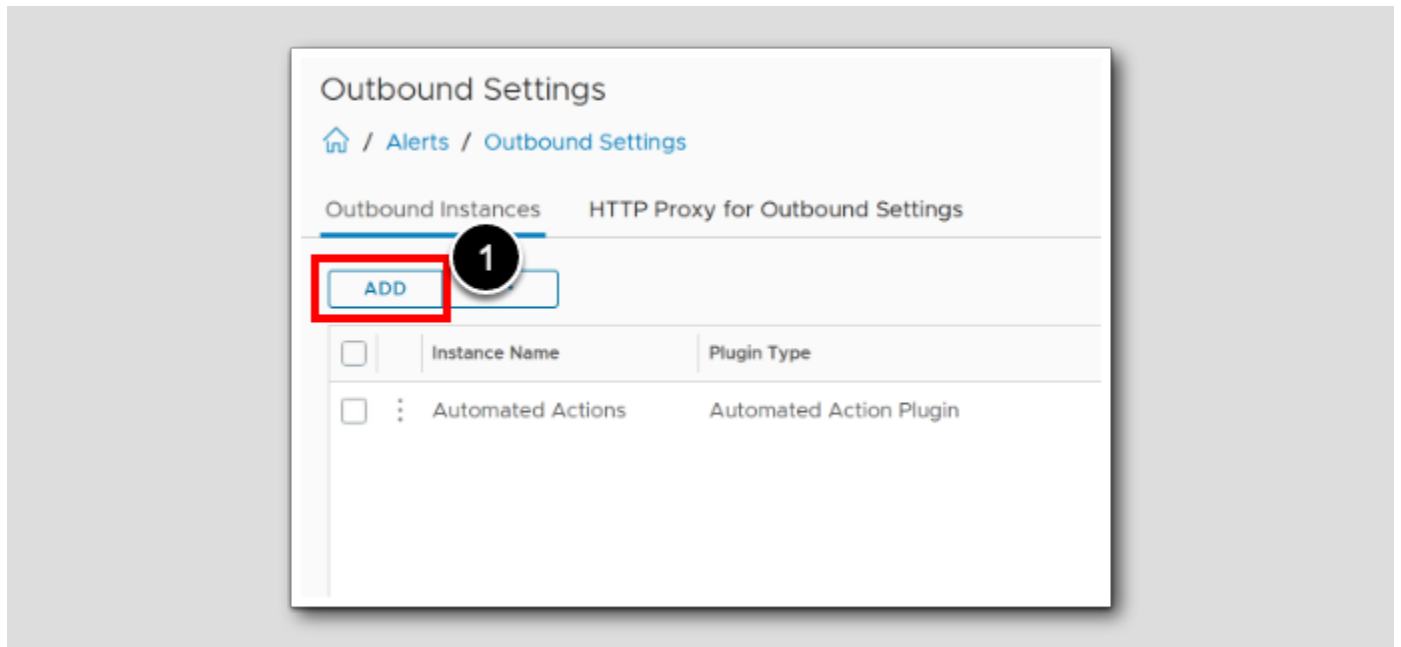
## The Outbound settings



1. In the Alerts page in Aria Operations, click Outbound Settings

## Outbound Settings

[18]



1. In the Outbound settings, Under Outbound Instances, click ADD

## New Outbound Instance

[19]

The screenshot shows a web form titled "Create New Outbound Instance" with a breadcrumb trail: "Alerts / Outbound Settings". The form contains several fields:

- Plugin Type:** A dropdown menu with "Webhook Notification Plugin" selected. A red box highlights the text, and a callout "1" points to the dropdown arrow.
- Instance Name:** A text input field containing "Outbound WebHook Instance". A red box highlights the text, and a callout "2" points to the input field.
- Url:** A text input field containing "http://webhook.corp.vmbeans.com/df433fe5-57e8-4275-8ce7-04712". A red box highlights the text, and a callout "3" points to the input field.
- Connection count:** A text input field containing "20".
- HTTP Proxy:** A dropdown menu with "Select Your HTTP Proxy" selected. A red "x" icon is visible on the right side.
- Credential type:** A dropdown menu with "No Credential" selected.

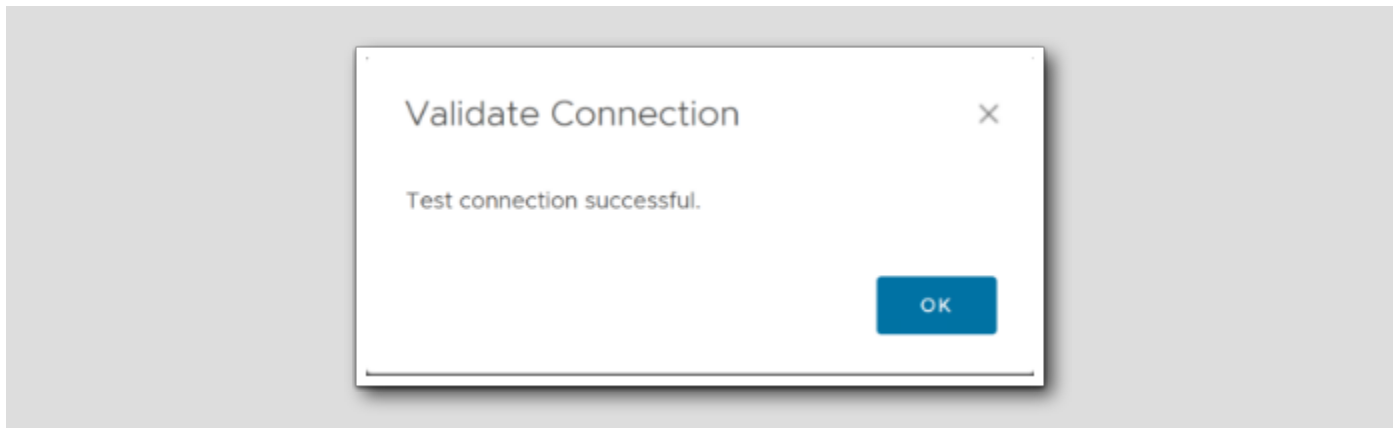
At the bottom of the form, there are three buttons: "TEST", "SAVE", and "CANCEL". A red box highlights the "TEST" button, and a callout "4" points to it.

We are about to create a new Outbound Instance using Webhook

1. From the drop-down menu for Plugin Type, choose **Webhook Notification Plugin**.
2. For the Instance Name, type something you can recognize, simply **Outbound WebHook Instance**.
3. For the URL, we previously copied the link, **Paste the link**.
4. To test this configuration, click **TEST**.

Validation successful

[20]

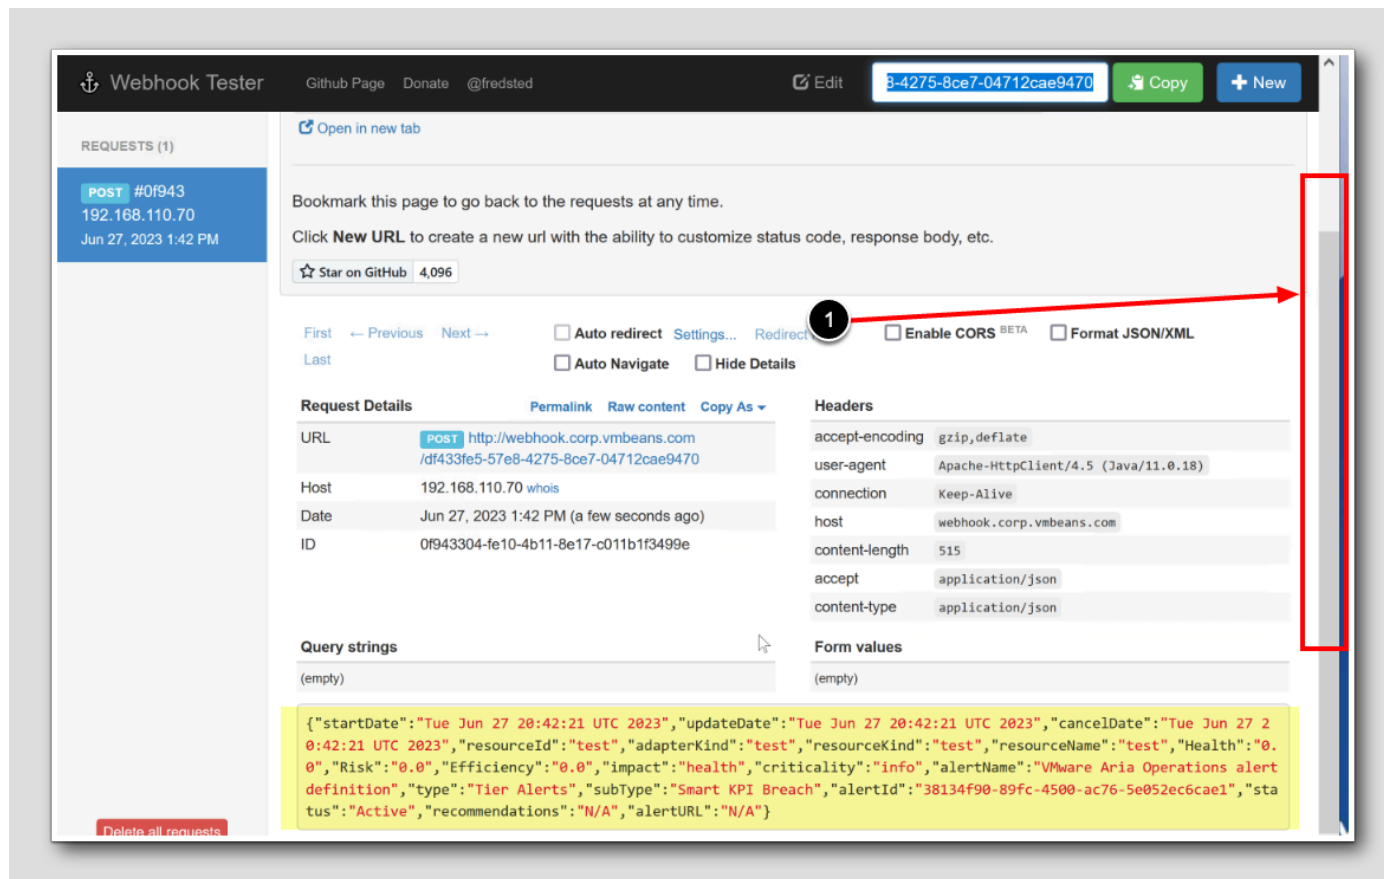


1. Click **OK**

NOTE: Do not close the Aria Operations window, use **ALT+TAB** or the taskbar to go to the Webhook Receiver.

## The Webhook tester

[21]



1. Scroll down to see the results properly

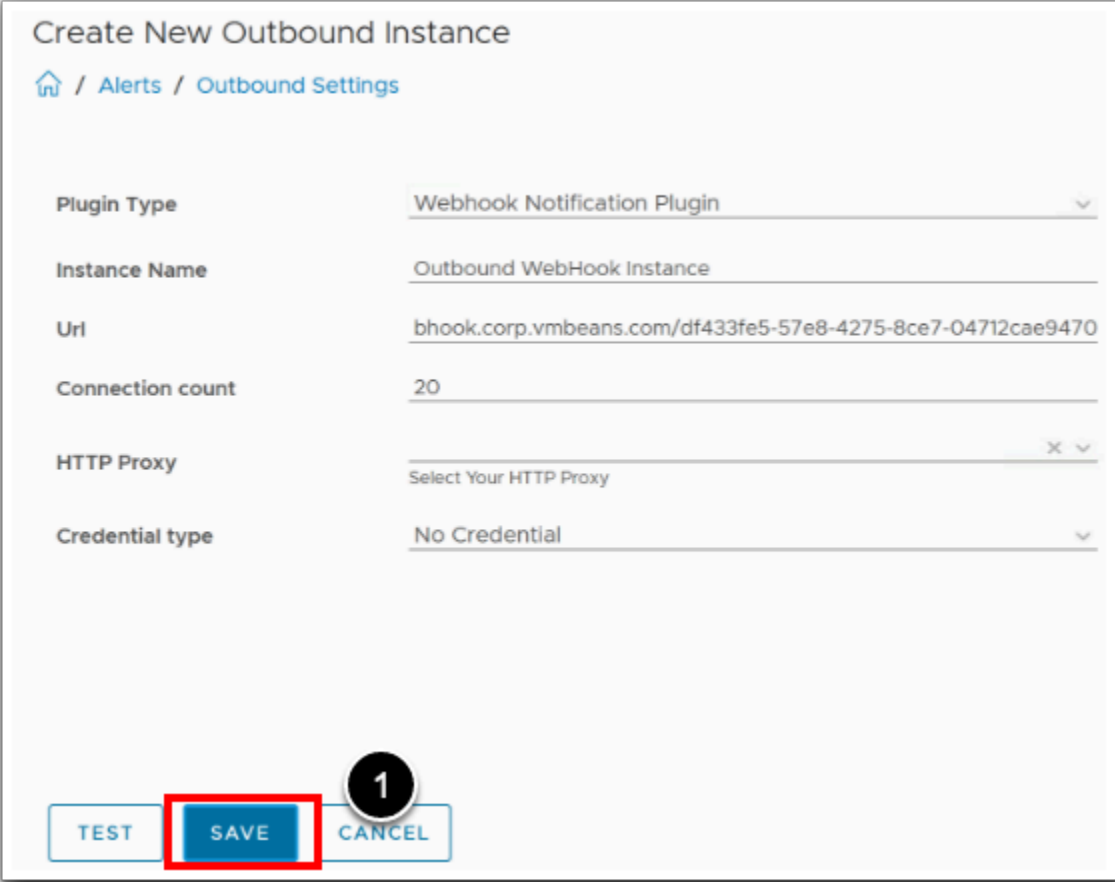
Congratulations! We have received our first webhook from Aria Operations. As you can see from the payload structure and content, several parameters and values are coming across.

In a real life use case we would probably have to review any documentation or specifications provided by the endpoint service to ensure we understand how the data should be formatted. Using Payload Templates will help us stay within these boundaries.

NOTE: Again Do not close this window, Again use ALT+TAB or the taskbar to get back to Aria Operations window.

## Save the Outbound Webhook Instance

[22]



Create New Outbound Instance

[Home](#) / Alerts / Outbound Settings

Plugin Type	Webhook Notification Plugin
Instance Name	Outbound WebHook Instance
Url	bhook.corp.vmbeans.com/df433fe5-57e8-4275-8ce7-04712cae9470
Connection count	20
HTTP Proxy	Select Your HTTP Proxy
Credential type	No Credential

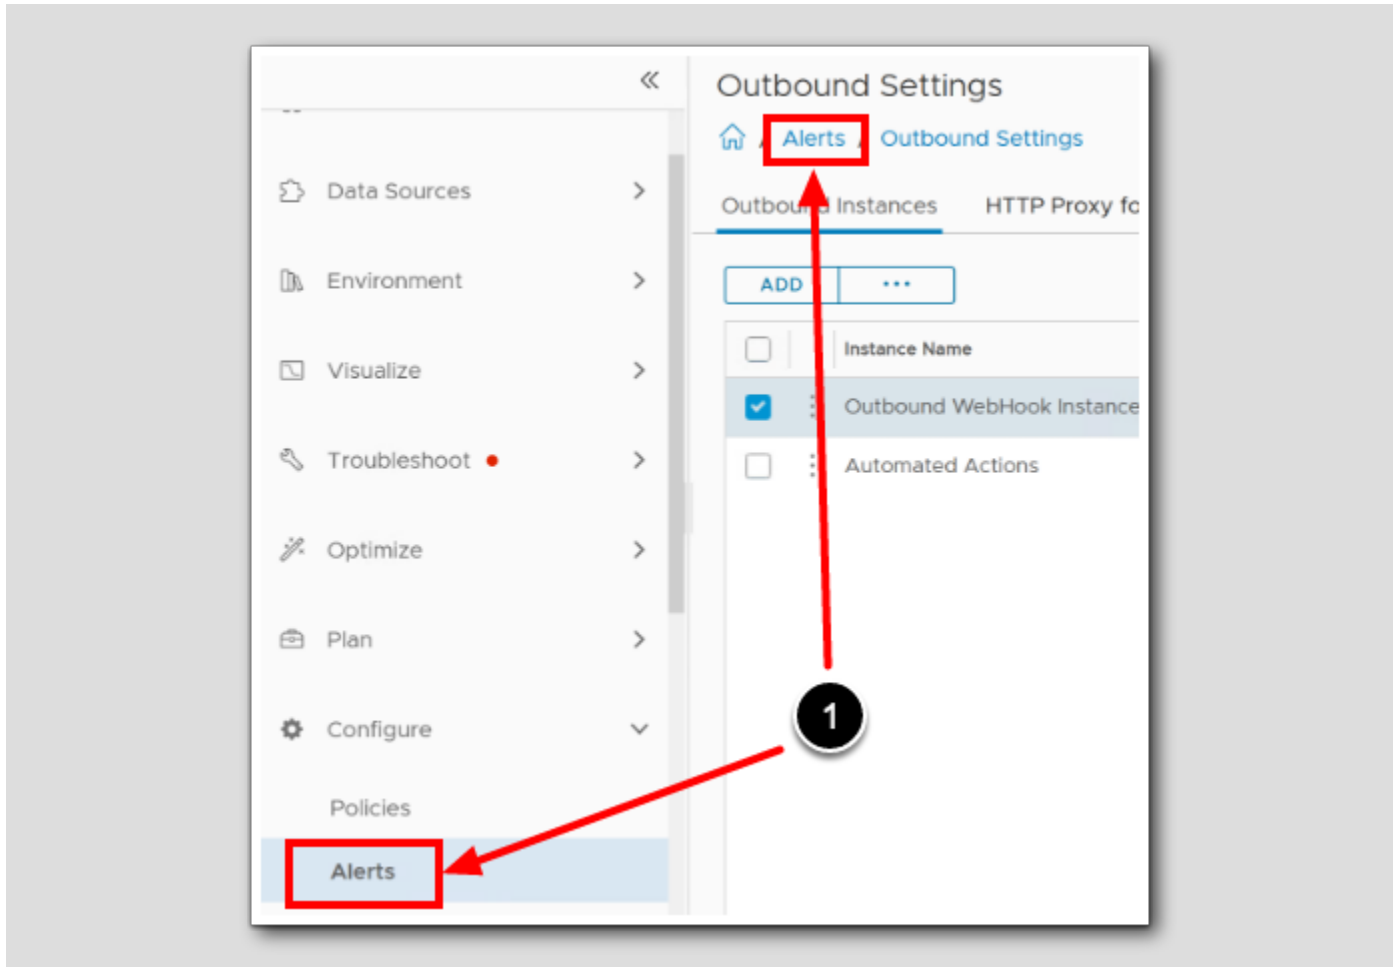
TEST SAVE CANCEL

1

1. To save our new Outbound Webhook Instance, just click SAVE

## Back to Alerts

[23]



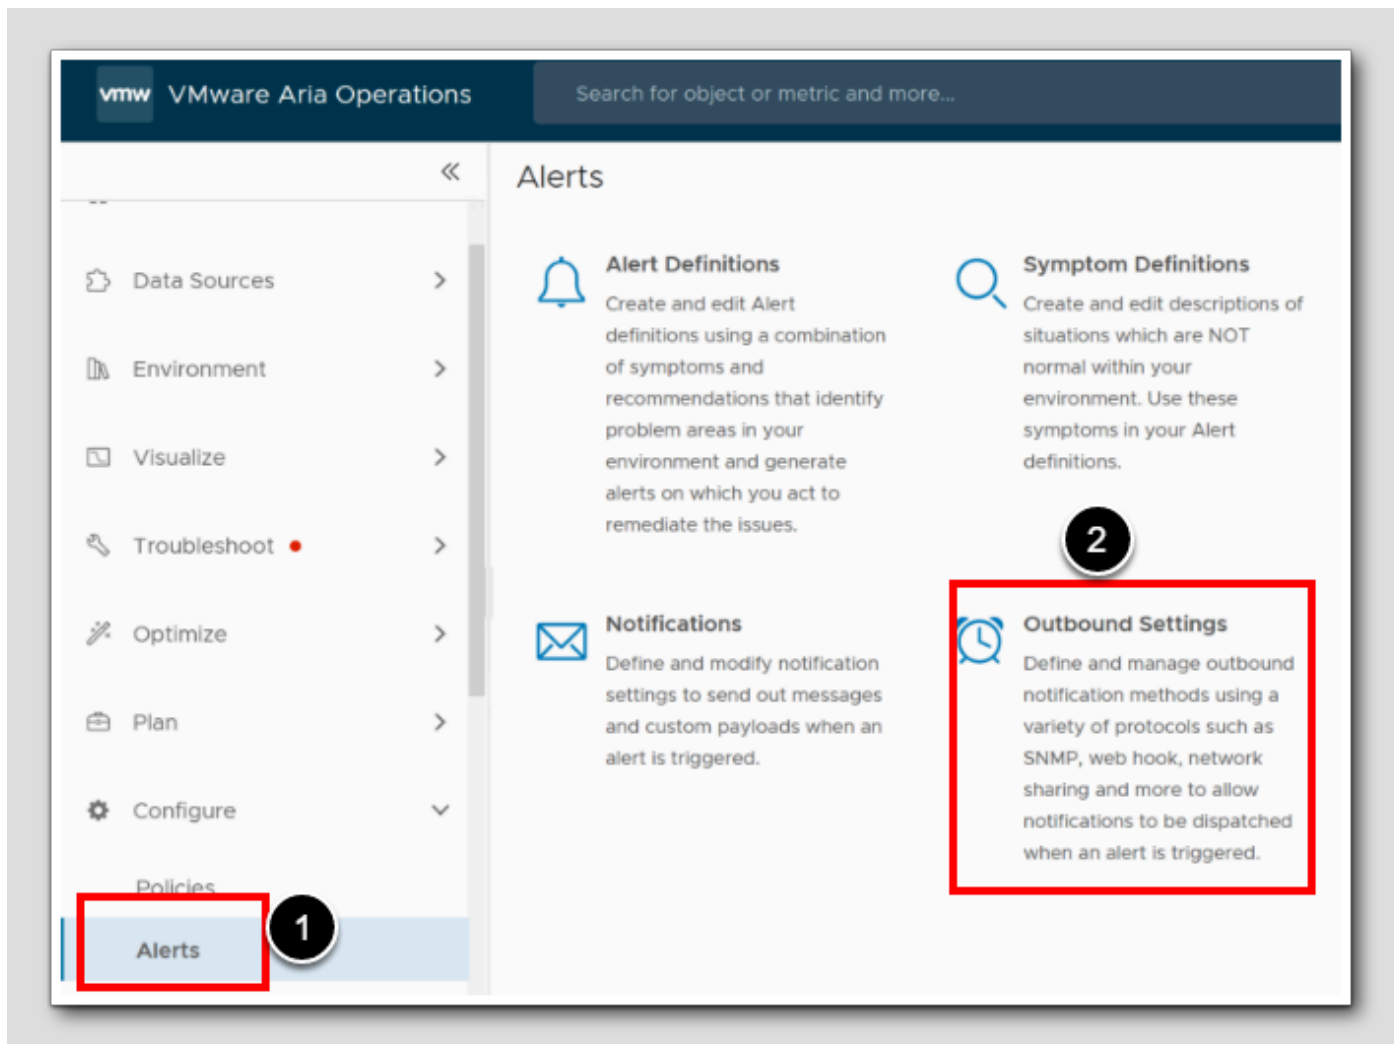
1. To get back to the Alerts, Click **Alerts** (any of the two places)

## Notifications using email

[24]

Creating notifications using email serves many purposes. By leveraging notifications, you stay informed, collaborate effectively, respond promptly, and continuously improve your IT operations. It's the key to a proactive, integrated, and optimized environment.

## Outbound settings



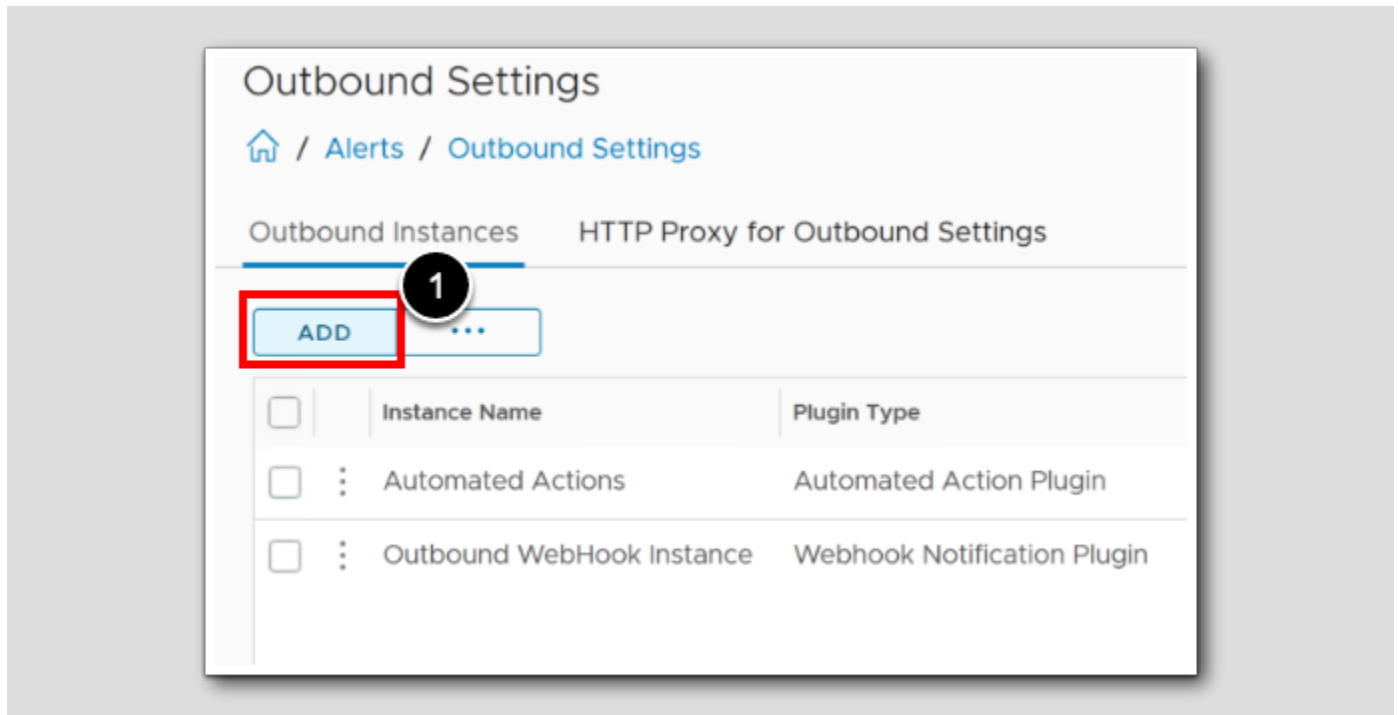
We are also going to set up the outbound settings for an email server.

1. Click Configure> Alerts
2. Click Outbound settings



## Add outbound instance

[26]



1. To add a new outbound instance, Click Add

## Create New Outbound Instance

[27]

Create New Outbound Instance

[Home](#) / [Alerts](#) / [Outbound Settings](#)

Plugin Type **1** Standard Email Plugin

Instance Name **2** Outbound email Instance

Use Secure Connection

Requires Authentication

SMTP Host **3** mail.corp.vmbeans.com

SMTP Port **4** 25

Secure Connection Type \_\_\_\_\_ x v

Sender Email Address **5** AriaOps@corp.vmbeans.com

Sender Name **6** YourNameGoesHere

Credential type No Credential v

Receiver Email Address **7** holadmin@corp.vmbeans.com

**TEST** SAVE CANCEL

Let's add a receiving E-mail server and a group of professional Operations Administrators to receive alerts via email from Aria Operations.

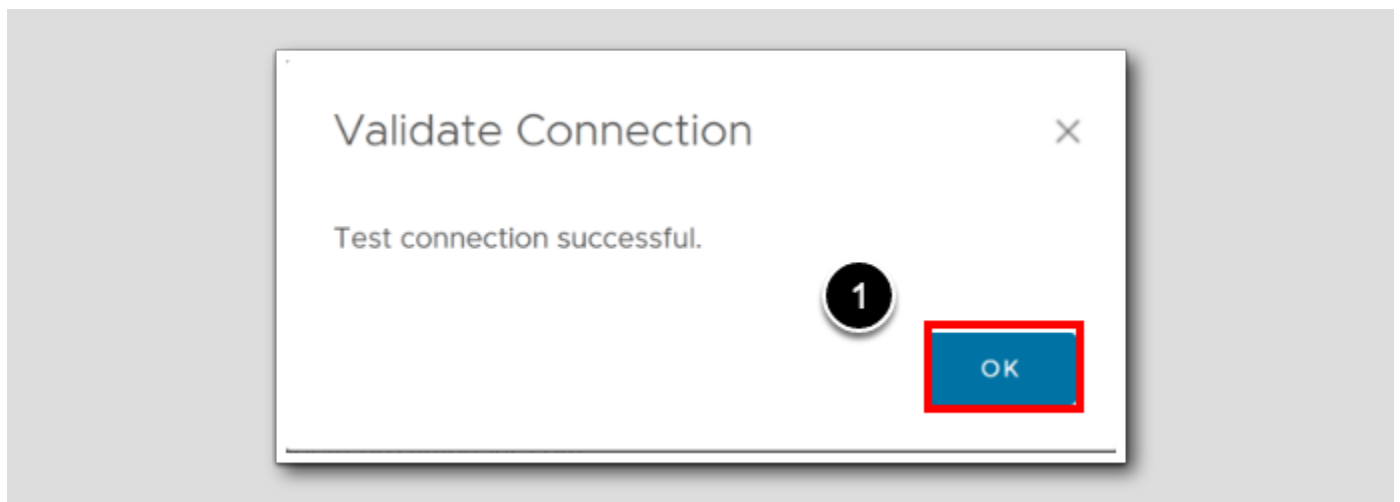
1. For Plugin type, choose **Standard Email Plugin**
2. Add an Instance name **Outbound email Instance**
3. Our SMTP Host **mail.corp.vmbeans.com**
4. SMTP port **25**
5. Sender Email Address **AriaOps@corp.vmbeans.com**

Note: Our sender email address needs to be with a valid domain with an MX record, and since this is a closed off installation with no internet access or access to external systems, we need to use the corp.vmbeans.com domain.

6. As the Sender Name, fill in your own name instead of **YourNameGoesHere**
7. In the Received Email Address field, add the group of admins to receive emails from Aria Operations, type **holadmin@corp.vmbeans.com**
8. Click **TEST**

## Validate Connection

[28]



After you have clicked TEST to validate the connection to our internal mailservers with the credentials we have provided

1. Click **OK**

Save instance

**Create New Outbound Instance**

[Home](#) / [Alerts](#) / [Outbound Settings](#)

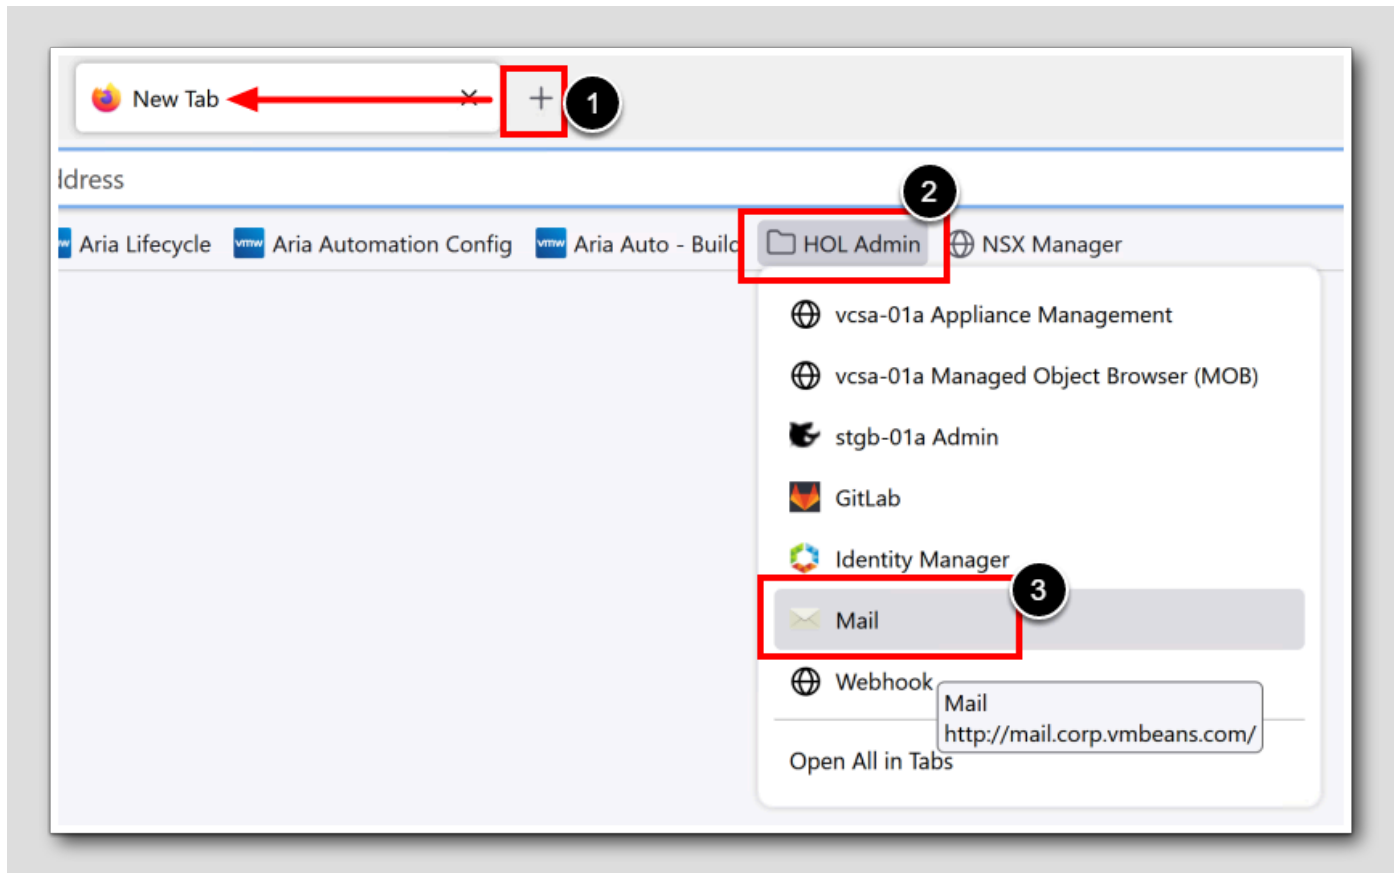
Plugin Type	Standard Email Plugin
Instance Name	Outbound email Instance
Use Secure Connection	<input type="checkbox"/>
Requires Authentication	<input type="checkbox"/>
SMTP Host	mail.corp.vmbeans.com
SMTP Port	25
Secure Connection Type	
Sender Email Address	AriaOps@corp.vmbeans.com
Sender Name	YourNAMEGoesHere
Credential type	No Credential
Receiver Email Address	holadmin@corp.vmbeans.com

**1**

TEST SAVE CANCEL

1. Click Save

## Start the email application




We have a local email server from poste.io that receives all emails we send in our lab.

1. Add a new TAB in the browser by pressing CTRL+T or click the '+'
2. Click the menu bar HOL Admin
3. Choose Mail

## Email Server Login

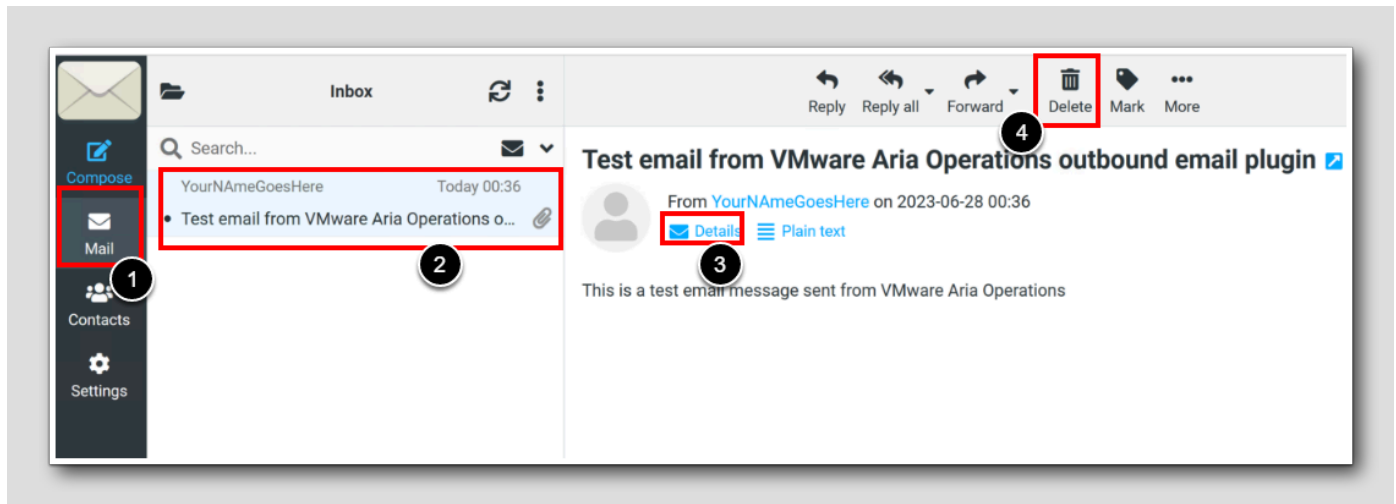
[31]



1. Username holadmin@corp.vmbeans.com
2. password VMware!
3. Click LOGIN

## Confirm email reception

[32]



1. In the left menu, click **Mail**
2. Click the first email in the Inbox received today **Test email from ..**
3. Click **Details**
4. Delete this email by clicking **Delete**

Note: Do not close this tab, just leave it, go back to the previous tab with Aria Operations (not shown)

## Payload templates

[33]

Let's continue to prepare for a better environment for Alerts by creating custom Payload template.

## Why Payload Templates

The screenshot displays the 'Payload Templates' configuration interface. On the left, a list of templates is shown, with 'Default Email Template' selected. The right pane shows the configuration for this template, including a description, a 'Payload Details' section with tabs for 'New Alert', 'Updated Alert', and 'Canceled Alert', and a 'Subject' field containing a template string: `[VMware Aria Operations] new alert Type:${ALERT_TYPE}, Sub-Type:${ALERT_SUBTYPE}, State:${ALERT_CRITICALITY}, Object Type:${RESOURCE_KIND}, Name:${RESOURCE_NAME}`. The 'Body' section contains a multi-line template string: `New alert was generated at ${CREATE_TIME}: Info:${RESOURCE_NAME} ${RESOURCE_KIND} is acting abnormally since ${CREATE_TIME} and was last updated at ${UPDATE_TIME} Alert Definition Name: ${ALERT_DEFINITION} Alert Definition Description: ${ALERT_DEFINITION_DESCRIPTION} Object Name : ${RESOURCE_NAME} Object Type : ${RESOURCE_KIND}`.

By leveraging notifications, you stay informed, collaborate effectively, respond promptly, and continuously improve your IT operations. It's the key to a proactive, integrated, and optimized environment.

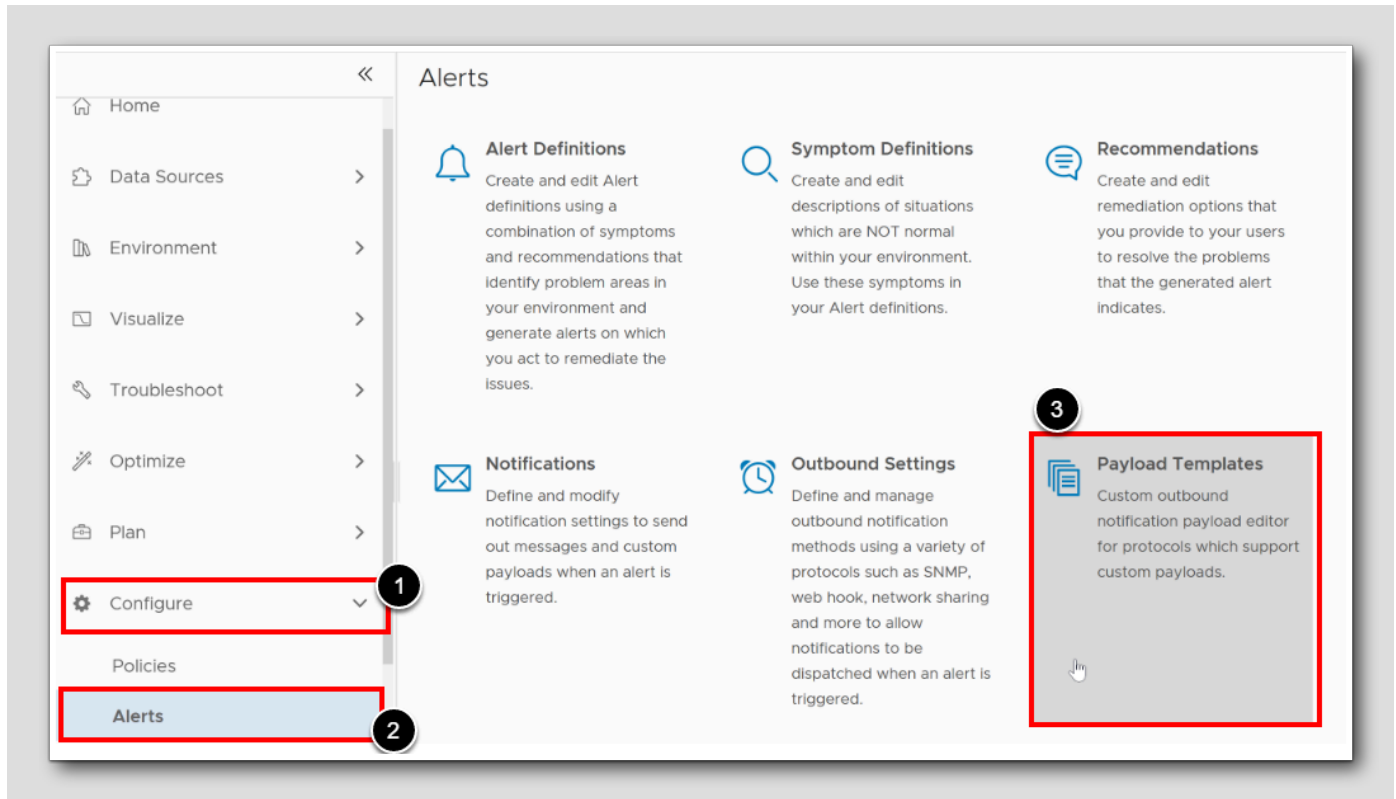
A **payload template** acts as a customization blueprint for notifications. It allows us to design the structure and content of the payload sent to external systems or applications. You can populate the payload with key information from the triggering event, ensuring it meets the recipient's requirements and enables informed actions. We can include specific information, such as alert details, impacted objects, timestamps, or any other relevant data, providing comprehensive context to aid in incident investigation and resolution.

Think of it as your creative tool to curate valuable and tailored data, empowering external systems to make swift and effective responses.

With this in mind let's customize two payload templates to use with email and webhooks.



## Opening Payload Templates



To get to the Payload templates page

1. Click Configure
2. Click Alerts
3. Click Payload Templates

## The default Email Template

The screenshot shows the 'Payload Templates' management interface. On the left, a list of templates is displayed, including 'Default Email Template', 'Default Trap Template', 'Default eNow Template', 'Default Slack Template', and 'Default Webhook Template'. The 'Default Email Template' is selected, and its details are shown on the right. The 'Subject' field is highlighted with a red arrow, showing the template text: '[VMware Aria Operations] new alert Type State:\${ALERT\_CRITICALITY}, Object Type'. The interface includes a breadcrumb 'Alerts / Payload Templates', an 'ADD' button, and a list of templates with checkboxes and ellipses. A context menu is open for the 'Default Email Template', showing 'Edit', 'Clone', and 'Delete' options. The 'Clone' option is highlighted with a red box. The 'Subject' field is also highlighted with a red box and a red arrow.

1. Click the blue text **Default Email Template**

Note: You see some of the text and variables that can be used for a template.

2. Click the **ellipsis**
3. Choose **CLONE**

## Step 1 - cloning email template

Create Payload Template

🏠 / Alerts / Payload Templates

1 - Details      2 - Object Content      3 - Payload Details

Name: Company Email Template **1**

Description: Email template we use in our company **2**

Outbound method: Standard Email Plugin **3**

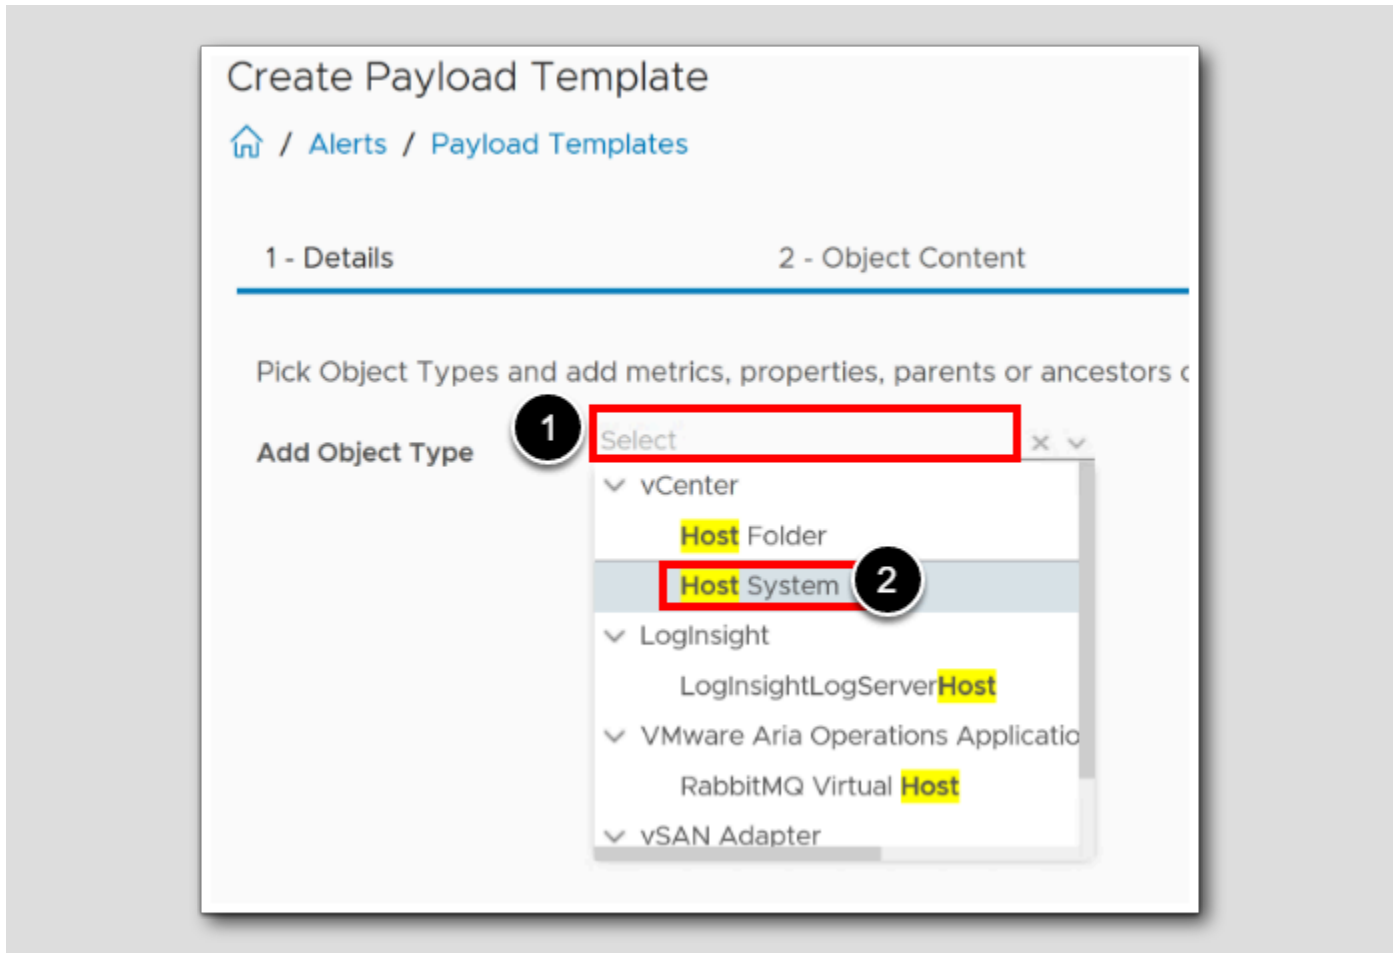
**4**

PREVIOUS   **NEXT**   CREATE   CANCEL

We are creating a new version for our company to use this as a template for information we need sent to us by email every time a notification is triggered from Aria Operations

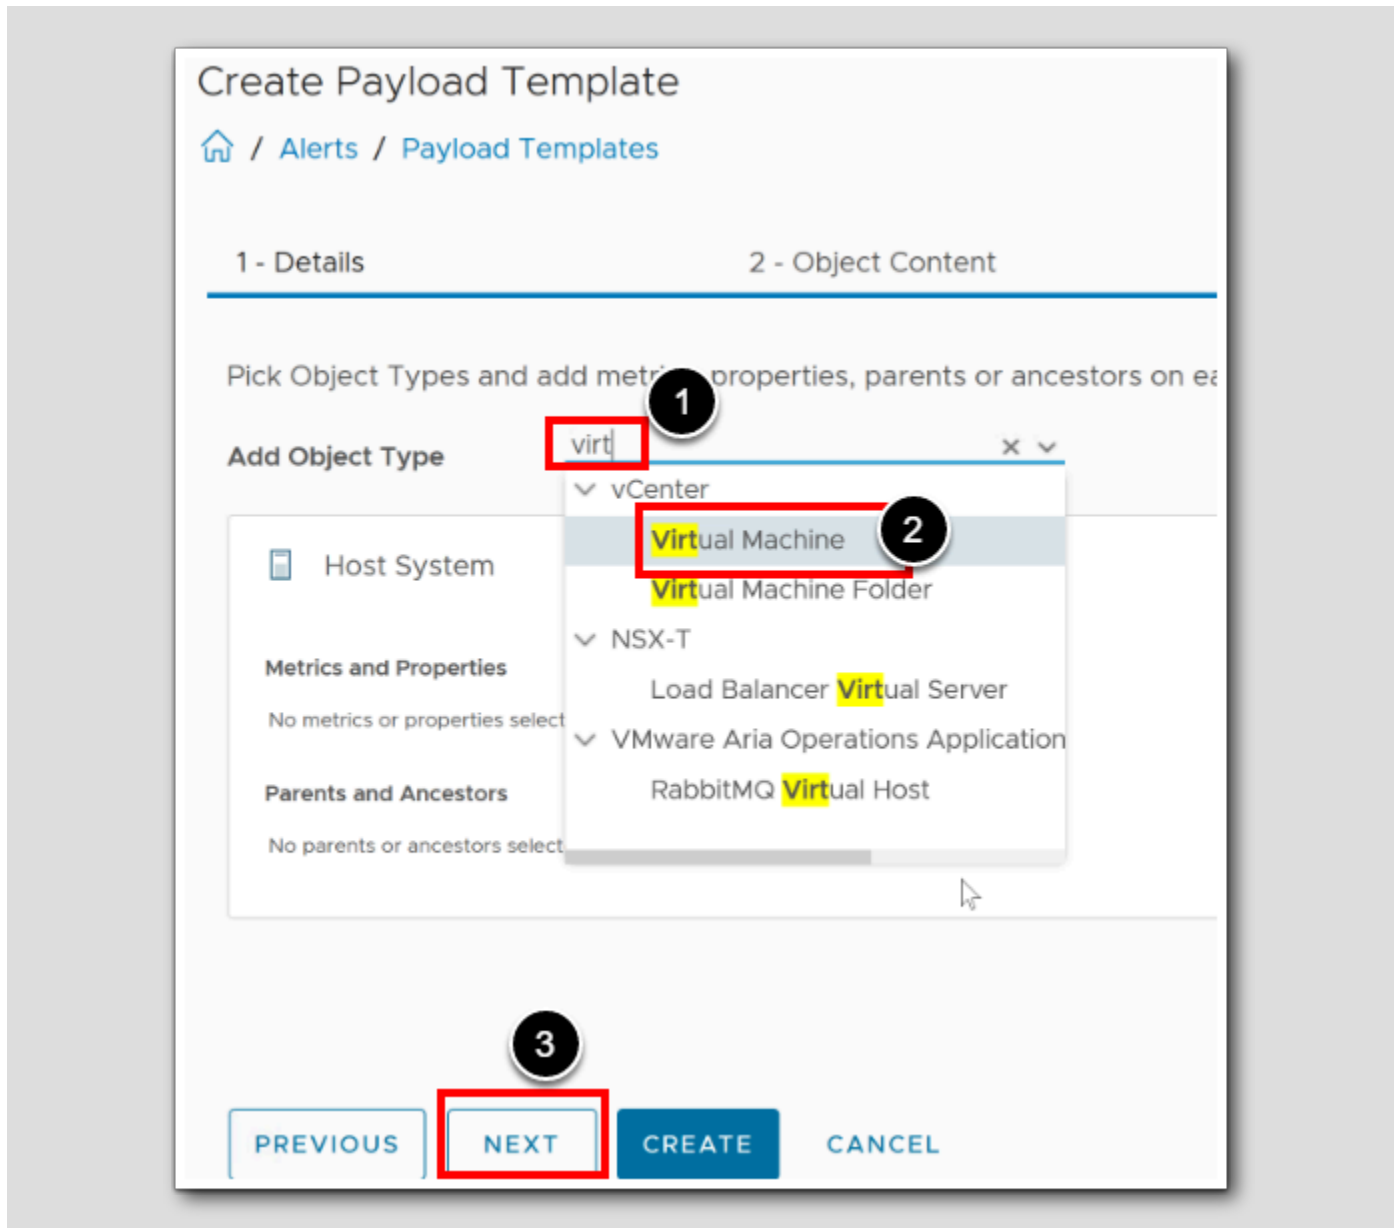
1. Give this payload template a meaningful name, for simplicity I've typed **Company Email Template**
2. For the description **Email template we use in our company**
3. Leave the value **Standard Email Plugin**
4. Click **Next**

## Add host objects



1. To easily find host object type. Start typing host
2. Select Host System

## Add VM Object



1. To easily find Virtual Machine object type. Start typing virt
2. Select Virtual Machine
3. Click Next

## Payload email Subject details

Create Payload Template

Alerts / Payload Templates

1 - Details      2 - Object Content      3 - Payload Details

Fill in the payload details below to be included in the notification.

Do you want to add template input properties?  Yes  No **1**

Do you want different payload details for new, updated, and canceled alerts?  Yes  No

**2** New Alert Updated Alert Canceled Alert

Subject **3**

`${ALERT_CRITICALITY} - New ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [`

Body

New alert was generated at `CREATE_TIME` :

Info: `RESOURCE_NAME` `RESOURCE_KIND` is acting abnormally since `CREATE_TIME` and was last updated at `UPDATE_TIME`

**4**

Parameters

Search Parameter	
<code>\$(CREATE_TIME)</code>	<input checked="" type="checkbox"/>
<code>\$(UPDATE_TIME)</code>	<input type="checkbox"/>
<code>\$(CANCEL_TIME)</code>	<input type="checkbox"/>
<code>\$(ALERT_STATUS)</code>	<input type="checkbox"/>

**5**

PREVIOUS NEXT CREATE CANCEL

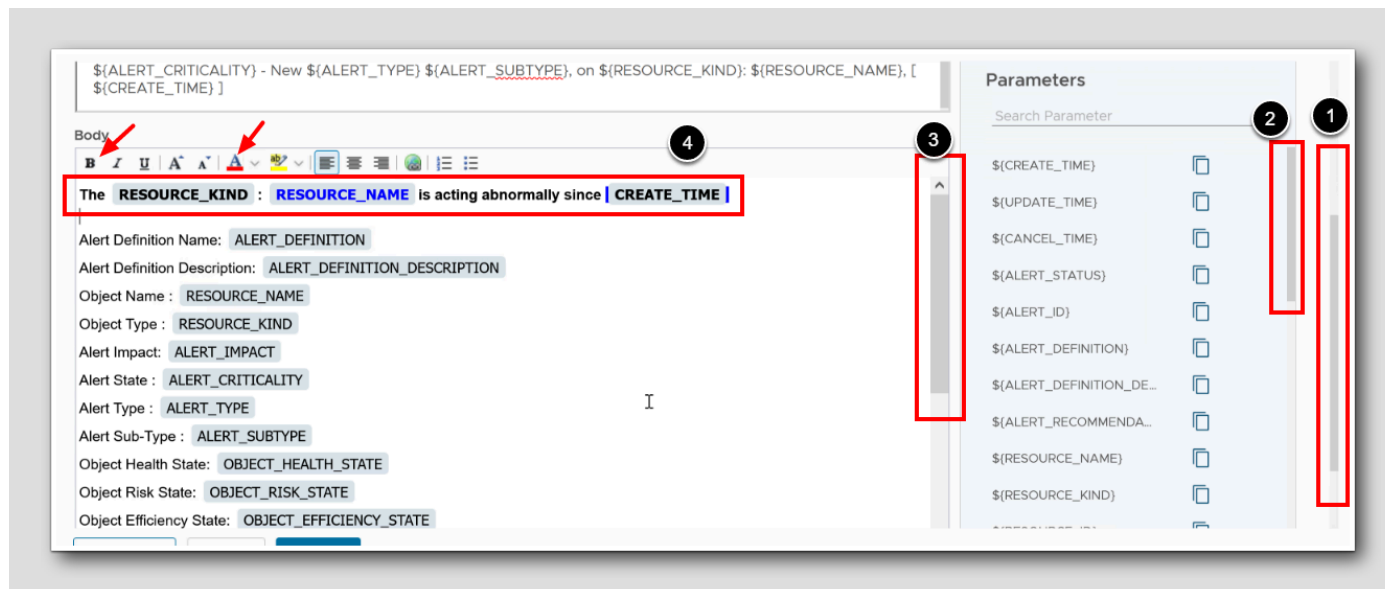
1. Choose the **default** radio buttons
2. Click **New Alert**  
We would like to clean up the Subject.
3. Edit or delete the subject, and replace with `${ALERT_CRITICALITY} - New ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [`
4. From the Parameters column, copy `$(CREATE_TIME)` by clicking the copy icon,
5. Click the empty space after `$(RESOURCE_NAME)` and **Paste** with CTRL+V into the Subject Window

The Complete Subject string could look like this: `${ALERT_CRITICALITY} - New ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [ at ${CREATE_TIME} ]`

Notice I added the word 'at' and a square bracket `[]` around the CREATE TIME tom make it readable for my company.

Up next is to change the body

## Payload Body details



1. To see expand the view of the Body and Parameters use the Window Scroll down
2. To see more Parameters use the Parameters Scroll Down/Up
3. To edit the complete Body, Use the Body Window Scroll Down/Up
4. Let's edit, cut'n'paste, and add Parameters to change the top of the body to what the picture above shows: **The RESOURCE\_KIND: RESOURCE\_NAME is acting abnormally since CREATE\_TIME**

Notice we made *resource\_name* and *create\_time* Blue! This means that the virtual machine name and the time of the alert will be shown in Blue color to force our eyes to focus on WHEN something happened to WHAT.

The result of this top part of the body would be something like this:

The VirtualMachine: **ubuntu-0008** is acting abnormally since **Wed Jun 28 13:34:15 UTC 2023**

This is your chance to showcase your abilities and give it your all, pouring in maximum effort, energy, and enthusiasm without holding back. However, for now, we will simplify our email alert template to avoid going overboard and keep it short, concise, and straight to the point. This will enable us to swiftly address and investigate any arising problems, ensuring clarity and ease of understanding for all IT administrators involved.

Let's continue:

The screenshot shows the 'New Alert' configuration page. The Subject field is populated with: `${ALERT_CRITICALITY} - New ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [ ${CREATE_TIME} ]`. The Body field contains the text: `The RESOURCE_KIND : RESOURCE_NAME is acting abnormally since CREATE_TIME`. A red box highlights a section of the body template containing the following fields: `# Alert:` (with sub-fields `ALERT_TYPE`, `ALERT_SUBTYPE`, `ALERT_IMPACT`, `ALERT_CRITICALITY`), `alert definition: ALERT_DEFINITION`, `# Symptoms:` (with sub-field `SYMPTOMS`), `# Recommendations:` (with sub-field `ALERT_RECOMMENDATIONS`), and `# Link: LINK_TO_ALERT`. A second red box highlights the text `Created by YourName at YourCompany`. A third red box highlights the `CREATE` button at the bottom of the form. A 'Parameters' list is visible on the right side of the interface.

1. We want a more simple and cleaner template, and here is One suggestion on how we could format our e-mail template to make it concise. Try to fill in **most of these values**, or edit as you see fit
2. Add your name and your company `Created By YourName at YourCompany`
3. Click `CREATE`



## The finished result

Standard Email Plugin

Company Email Template

Email template we use in our company

**Payload Details**

New Alert Updated Alert Canceled Alert

**Subject**

`${ALERT_CRITICALITY} - New ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}:  
${RESOURCE_NAME}, [ ${CREATE_TIME} ]`

**Body**

The `${RESOURCE_KIND}: ${RESOURCE_NAME}` is acting abnormally since `${CREATE_TIME}`

**# Alert:**  
`${ALERT_TYPE}${ALERT_SUBTYPE}${ALERT_IMPACT}${ALERT_CRITICALITY}`  
alert definition: `${ALERT_DEFINITION}`

**# Symptoms:**  
`${SYMPTOMS}`

**# Recommendations:**  
`${ALERT_RECOMMENDATIONS}`

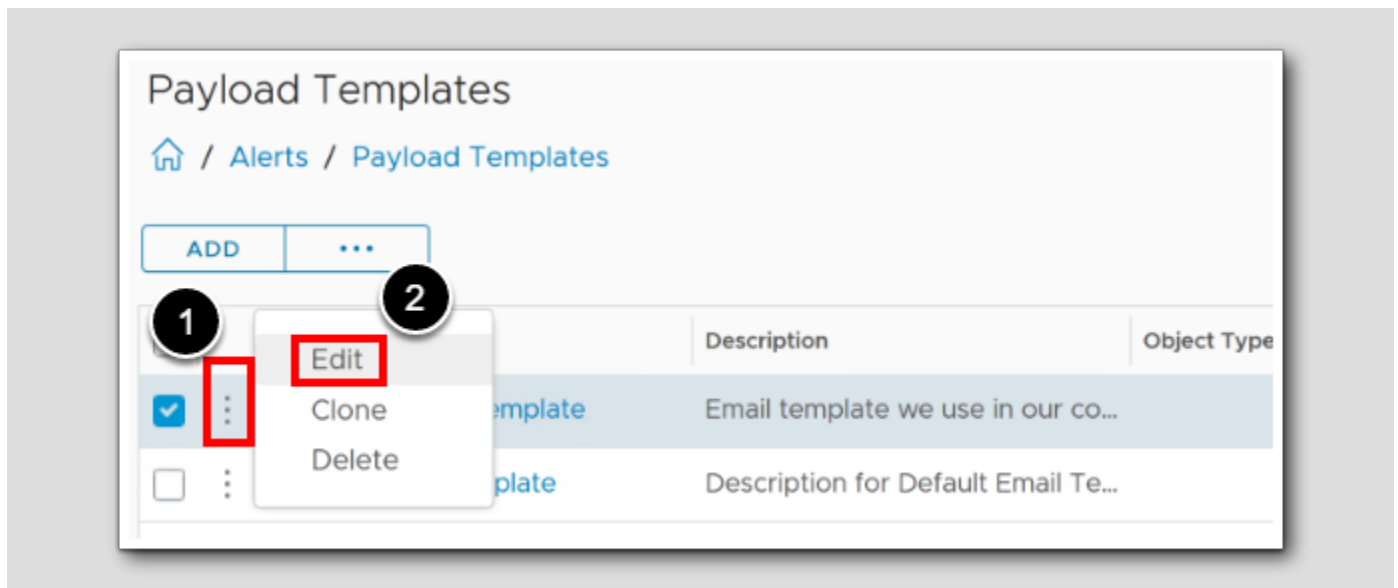
**# Link:** `${LINK_TO_ALERT}`

Created by YourName at YourCompany

Congratulations! You have now finished a cool new Payload template for email notifications.

1. To get an overview of the finished result, Click on Company Email Template

The next step is to edit this Payload template and make the Updated Alert and the Canceled Alert ju



1. On the *Company Email Template*, click the ellipsis
2. Choose **Edit**
3. In the *Edit Payload Template* page Click **Payload Details** (Not Shown)

## Changing the Updated Alert

The screenshot shows the alert configuration interface with three numbered callouts:

- 1**: The 'Updated Alert' tab is selected and highlighted with a red box.
- 2**: The 'Subject' field is highlighted with a red box, containing the text: `${ALERT_CRITICALITY} - UPDATE ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}; ${RESOURCE_NAME}, [${UPDATE_TIME}]`.
- 3**: The 'Body' field is highlighted with a red box, containing a rich text template:
 

```
Alert updated at UPDATE_TIME :
The RESOURCE_KIND : RESOURCE_NAME is acting abnormally since CREATE_TIME

# Alert:
ALERT_TYPE ALERT_SUBTYPE ALERT_IMPACT ALERT_CRITICALITY
alert definition: ALERT_DEFINITION

# Symptoms:
SYMPTOMS

# Recommendations:
ALERT_RECOMMENDATIONS

# Link: LINK_TO_ALERT

Created by YourName at YourCompany
```

1. Make sure you are at the Updated Alert
2. Change the Subject to `${ALERT_CRITICALITY} - UPDATED ALERT ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}; ${RESOURCE_NAME}, [${UPDATE_TIME}]`
3. Change the body by copying most of the Body from the *New Alert* to the *Updated Alert Body*. Edit the Body to mostly look like this image, it does not need to be perfect.

## Changing the Canceled Alert

The screenshot shows the configuration interface for a 'Canceled Alert'. The 'Canceled Alert' tab is selected (1). The subject line is edited to: `${ALERT_CRITICALITY} - CANCELLED ALERT ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [${CANCEL_TIME}]` (2). The body text is edited to include the following content (3):

**Alert was cancelled at** `CANCEL_TIME` :

The `RESOURCE_KIND` : `RESOURCE_NAME` is acting abnormally since `CREATE_TIME`

**# Alert:**  
`ALERT_TYPE` `ALERT_SUBTYPE` `ALERT_IMPACT` `ALERT_CRITICALITY`  
 alert definition: `ALERT_DEFINITION`

**# Symptoms:**  
`SYMPTOMS`

**# Recommendations:**  
`ALERT_RECOMMENDATIONS`

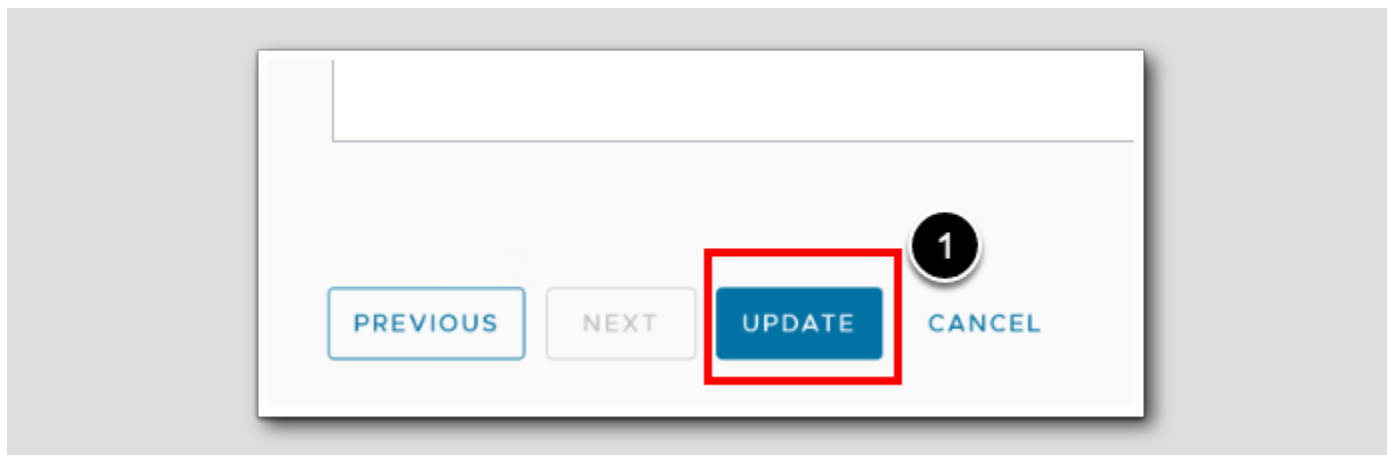
**# Link:** `LINK_TO_ALERT`

Created by `YourName` at `YourCompany`

1. Make sure you are at the Canceled Alert
2. Change the Subject to `${ALERT_CRITICALITY} - CANCELLED ALERT ${ALERT_TYPE} ${ALERT_SUBTYPE}, on ${RESOURCE_KIND}: ${RESOURCE_NAME}, [${CANCEL_TIME}]`
3. Change the body by copying most of the Body from the *Updated Alert Body* to the Canceled Alert Body. Edit the Body to mostly look like this image, it does not need to be perfect in this Lab

## Save, I mean Update

[47]



1. Remember to save the updates by clicking **UPDATE**

## Conclusion

[48]

### What did we do here?

Well, you are now prepared to make an alert system and send meaningful emails as New alerts (Something has happened to an object) and get updates (meaning the problem is still there, or has escalated) and see when the alert has been canceled (meaning the problem has gone away)

When you get back to the office, you can also edit and customize your own Payload template for the webhooks as well, depending if you would like more or less instant alerts or not.

### What happens in the next module?

In the next module we will create an **Alert Definition**, where we will define the conditions that indicate high CPU usage on a VM. For example, set a threshold of CPU usage exceeding 70-90% for a duration of 5 minutes and configure the severity level for the alert, such as "High" or "Critical," based on the impact of high CPU usage.

## Conclusion

[49]

- This module has equipped you with the essential skills to create notifications using webhooks and email. By mastering the customization through payload templates, you can design and deliver comprehensive information to external systems, enabling swift and informed actions.
- Real-time notifications facilitated by webhooks ensure timely alerts for prompt incident response, minimizing downtime and optimizing your IT operations.

- This lab emphasized the importance of leveraging notifications as a proactive measure, fostering collaboration, and continuously improving your environment. By embracing these techniques, you can create a proactive, integrated, and optimized system that enhances your overall productivity.

## You've finished Module 1

[50]

Congratulations on completing the lab module.

If you are looking for additional general information on vRealize Operations 8.4, try one of these:

- **VMware Product Public Page - Aria Operations:** <https://docs.vmware.com/en/VMware-Aria-Operations/index.html>
- **Aria Operations 8.12.1 - Release Notes:** <https://docs.vmware.com/en/VMware-Aria-Operations/8.12.1/rn/vmware-aria-operations-8121-release-notes/index.html>

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 2 - Creating a Custom Alert Definition (25 minutes) Intermediate

### Introduction

[52]

By the end of this guide, you'll have understanding of creating and managing custom alert definitions and notifications in Aria Operations, allowing for proactive monitoring, faster issue detection, and expedited response actions in your VMware infrastructure.

Upon completing this lab, you will be able to:

- Understand Alerts, Symptoms, Recommendations and Actions
- Build a custom Alert Definition
- Simulate issues in the environment to demonstrate how to customize the alerts
- Utilize the different ways alerts can be used based on the critical nature or other characteristics of the monitored infrastructure

### Log in to Aria Operations

[53]

We will log in to a live instance of Aria Operations running in this lab.

### Open the Firefox Browser from the Windows Task Bar

[54]




If the browser is not already open, launch **Firefox**.

1. Click the **Firefox** icon in the Windows Quick Launch Task Bar at the bottom of the screen.

## Navigate to Aria Operations

[55]




1. Click the Aria Operations bookmark in the bookmarks toolbar.



## Log in to Aria Operations

[56]



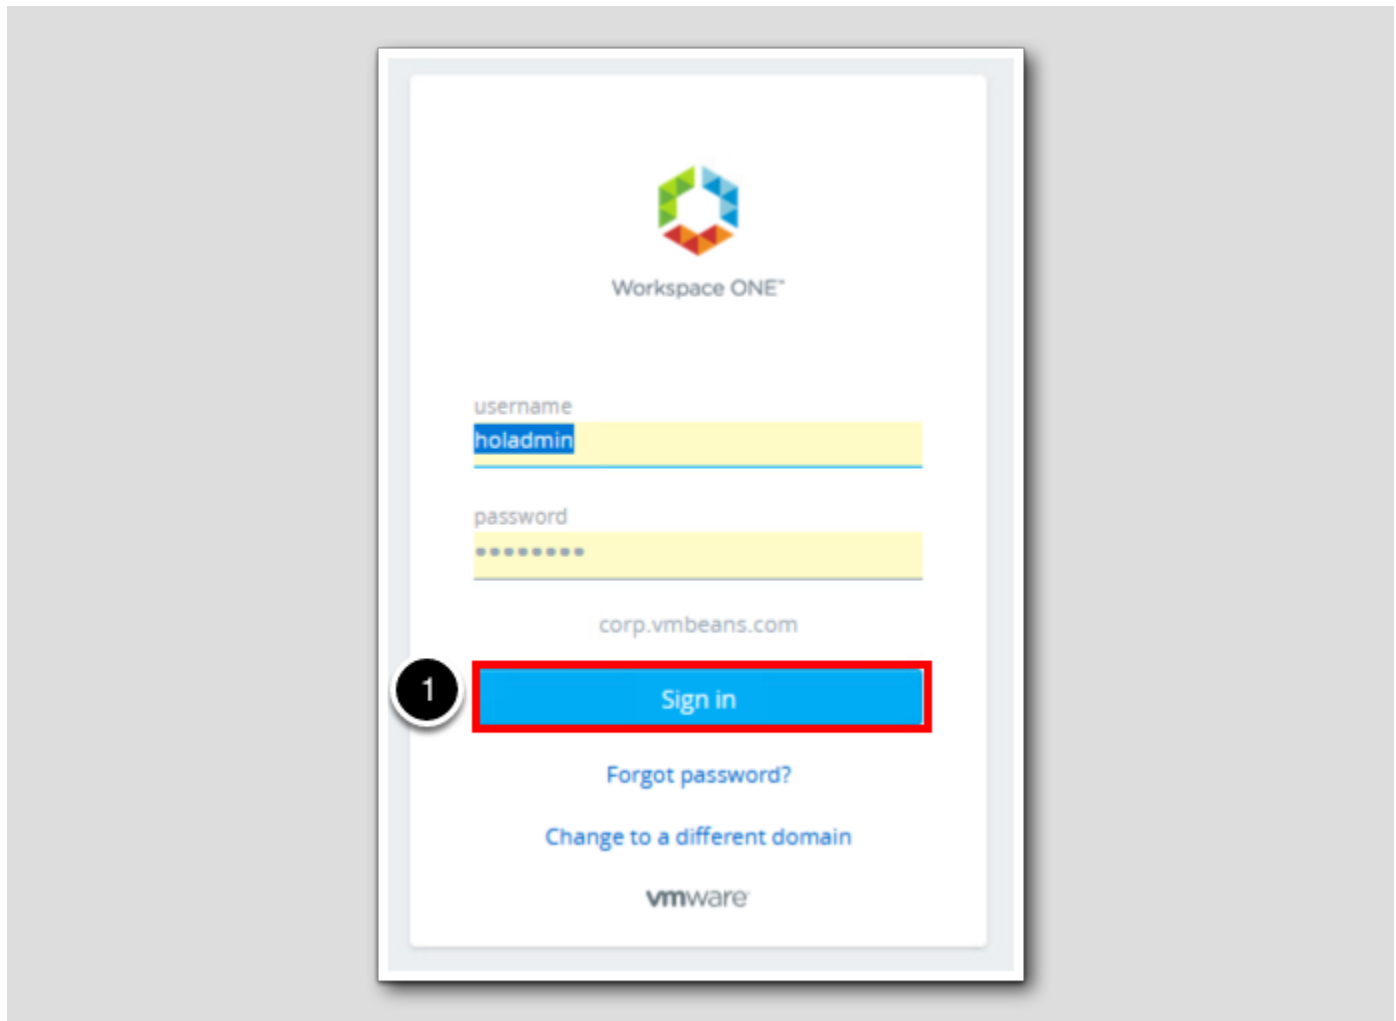
Aria Operations is integrated with VMware Workspace ONE Access (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login

[57]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

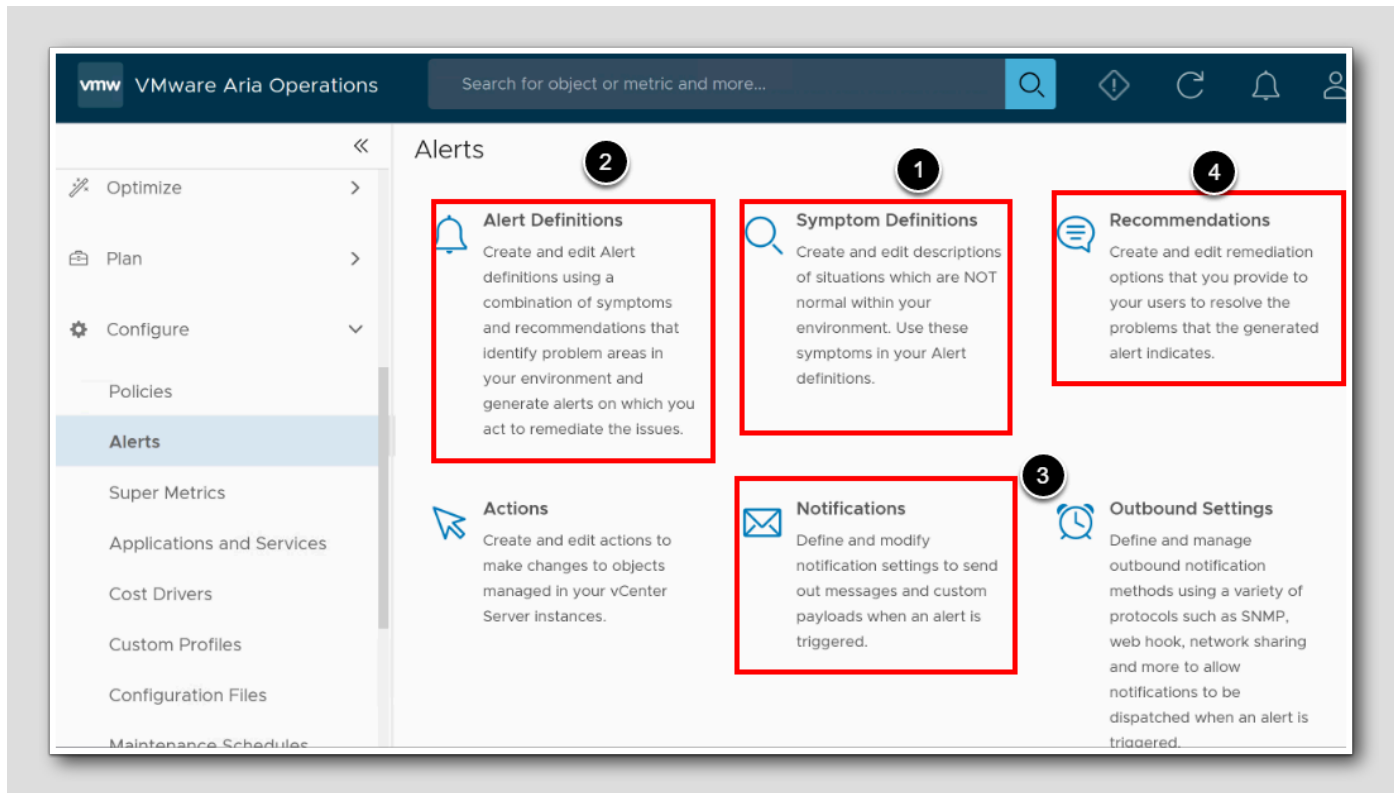
1. Click **Sign in**.

## Custom Alert Definition

[58]

We will create a custom Alert sent via email using our Custom Payload we have created earlier. We will create an Alert definition with symptoms

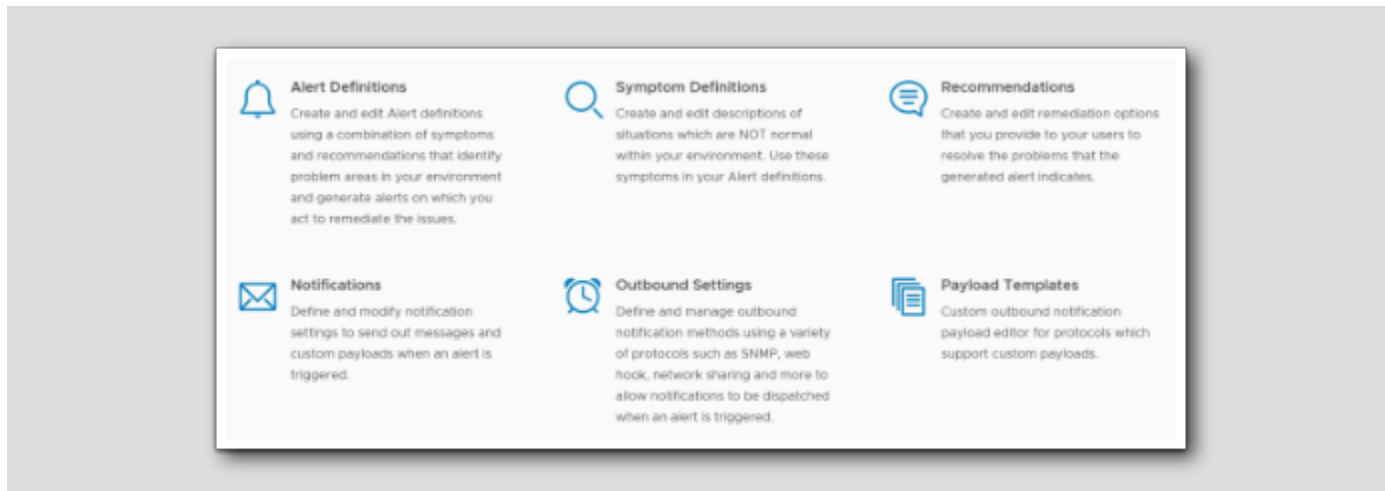
## The holistic approach



By leveraging Symptom Definitions, Alert Definitions, Notifications, and Recommendations, Aria Operations enables IT operations teams to monitor, detect, and respond to potential issues or deviations from desired states in their VMware infrastructure. This **holistic approach** helps ensure efficient resource utilization, proactive problem management, and improved overall operational efficiency within the IT environment.

1. **Symptom Definitions:** Define the specific metrics, conditions, or thresholds that indicate the presence of a problem or an abnormal behavior.
2. **Alert Definitions:** Rules or conditions that determine when an alert should be triggered based on the occurrence or combination of specific symptoms.
3. **Notifications** Can be delivered through email, and are the means through which relevant stakeholders are informed about triggered alerts or events. When an alert is generated based on the defined criteria, We send notifications to administrators, IT teams, or other designated recipients.
4. **Recommendations** Provide actionable insights and guidance to address or resolve the issues identified by the monitoring system. These recommendations are based on VMware's knowledge base, best practices, and experience with similar situations to help with remediation steps to mitigate risks or resolve problems within the environment.

## Our Scenario



We need to create a custom Alert sent via email using our Custom Payload template from [Module 1 - Configuring and Managing Alert Notifications](#).

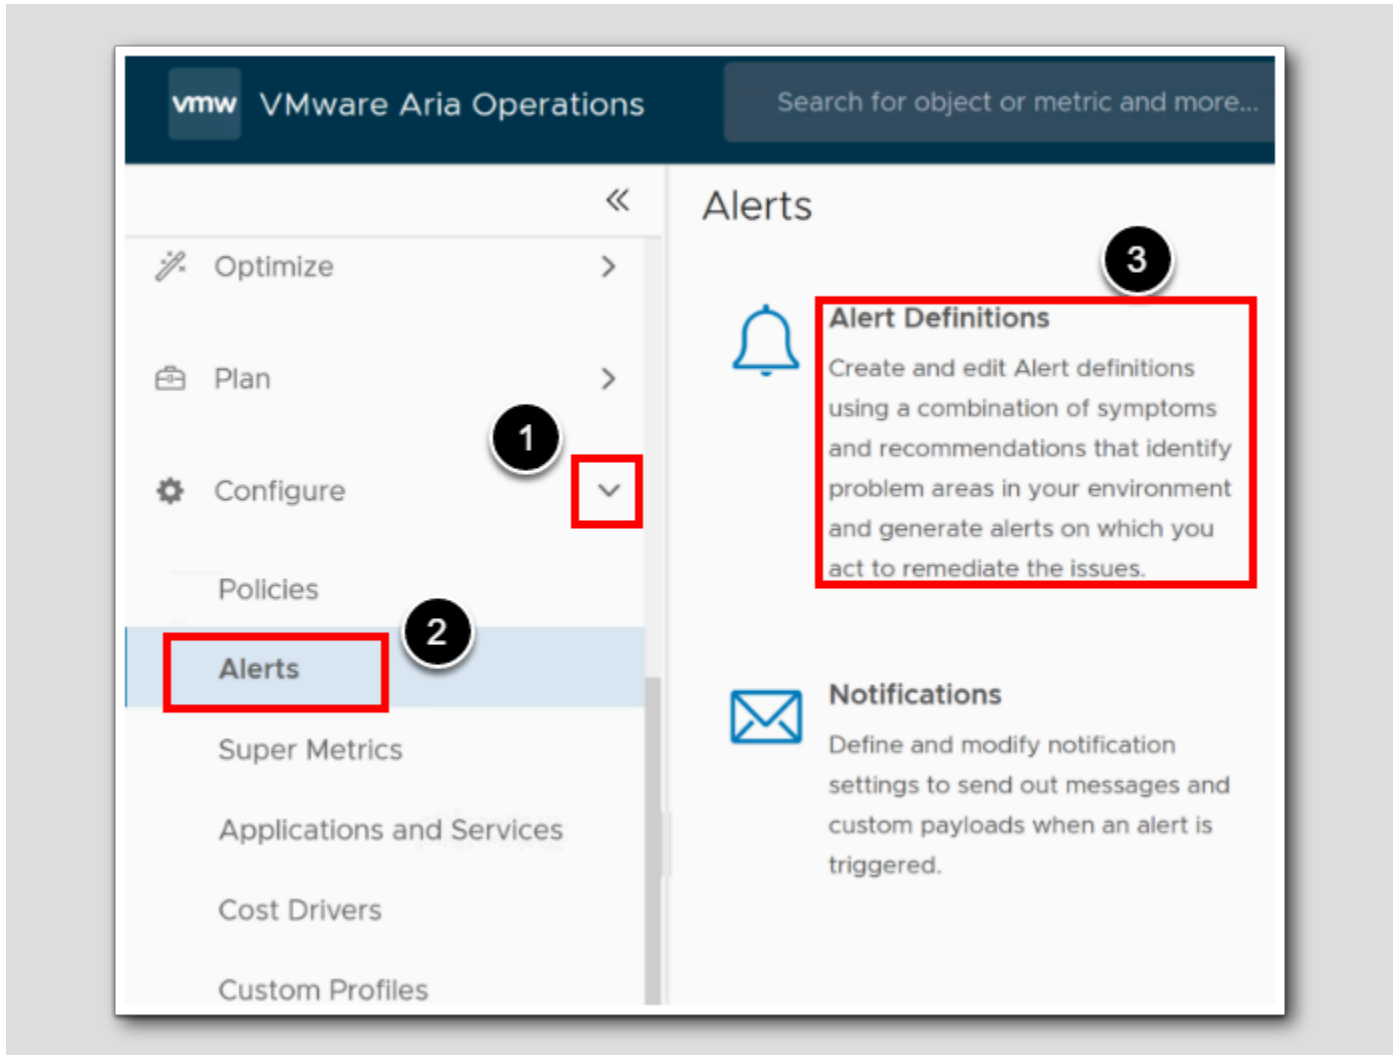
During this process we will use our own Symptoms and not utilize any Out of the box Symptom definitions. We will concentrate on the **CPU Demand** metric to tell us when Virtual Machines are working too hard. These are delicate servers and can not be overloaded, so by 'hard' we mean 70% is a warning, 80% is a immediate alert, and a 90% is a critical Alert. The VMs we're monitoring are connected to a specific policy.

**CPU Demand(%)** Is a metric that helps to understand how much of the CPU resources a virtual machine (VM) needs, regardless of any contention or limits. The key point to note here is that CPU Demand represents the 'need' for CPU resources, not the actual usage.

**Testing:** We will validate the configuration by simulating high CPU usage on the VMs. We will monitor the email system, as explained in the lesson [Start the email application](#) in module 1, to confirm that the alert is triggered, that our Payload Template works and that the defined symptom and conditions are correctly identified.

With this as a background, the next natural step is to create an Alert Definition!

Open Alert definition page



1. Click Configure
2. Click Alerts
3. Click Alert Definitions

## Review Out of the box content

Alert Definitions

Home / Alerts / Alert Definitions

ADD

1 demand

2

<input type="checkbox"/>	Name	Adapter Type	Object Type	Alert Type	Alert Subtype	Criticality	Impact
<input type="checkbox"/>	Pod is demanding more CPU than the configured limit	vCenter	Pod	Virtualization/Hyper...	Performance	Warning	Health
<input type="checkbox"/>	Virtual machine in a cluster is demanding more CPU than its entitlement	vCenter	Virtual M...	Virtualization/Hyper...	Performance	Warning	Health
<input type="checkbox"/>	Virtual machine is demanding more CPU than the configured limit	vCenter	Virtual M...	Virtualization/Hyper...	Performance	Warning	Health
<input type="checkbox"/>	Virtual machine in a DRS cluster is demanding more CPU than its entitlement	vCenter	Virtual M...	Virtualization/Hyper...	Performance	Warning	Health
<input type="checkbox"/>	Virtual machine is demanding more CPU than its entitlement	vCenter	Virtual M...	Virtualization/Hyper...	Performance	Warning	Health

1. Let's search for Demand, in the Filter box type **demand** and press **Enter**

Notice the name of the Alert Definitions that explains what these Alert Definitions are about

2. To add a New Alert definition, Click **ADD**

Edit Alert Definition

**Edit Alert Definition** Virtual Machine

[Home](#) / [Alerts](#) / [Alert Definitions](#)

1 - Alert 2 - Symptoms / Conditions

**Name**  1

**Description**  2

**Base Object Type**  3

**Advanced Settings** 4

**Impact**

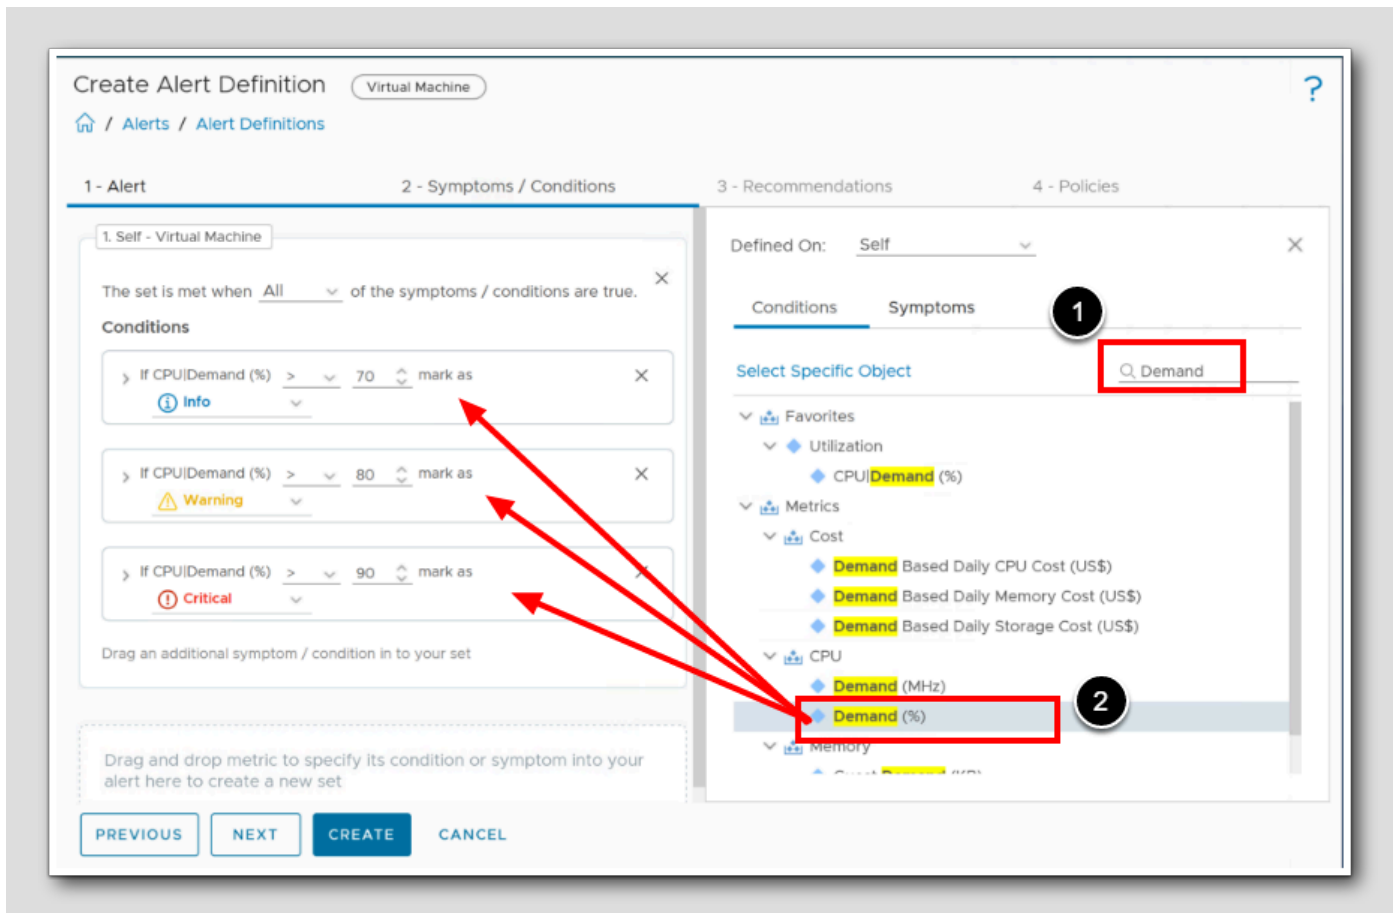
**Criticality**  5

**Alert Type & Subtype**

1. Enter a new Name for our Alert Definition, OurCompany Custom Alert Definition
2. Enter a Description: Our Company Custom Alert definition. Sets the conditions that determine when an alert should be triggered based on high CPU Demand(%) meaning a high 'need' for CPU resources
3. As the Base Object Type choose Virtual Machine
4. Expand the **Advanced** section by clicking the arrow
5. Make sure the *Criticality* is based on the symptoms, choose **Symptom Based**
6. For the *Alert Type* we will change from Application Performance to Virtualization Performance, Choose **Virtualization/Hypervisor Performance**
7. Click **NEXT**

Add the metrics

[64]





1. In the Search Box, Type **Demand**, and press **ENTER**
2. To create a set of symptoms **Drag and Drop** the *CPU Demand (%)* metric **3x times**

## Change the Symptoms conditions

**Edit Alert Definition** Virtual Machine

[Home](#) / Alerts / Alert Definitions

1 - Alert 2 - Symptoms / Conditions

1. Self - Virtual Machine

The set is met when **Any** of the symptoms / conditions are true.

**Conditions**

>	If CPU Demand (%) >	70	mark as	Warning	X
>	If CPU Demand (%) >	80	mark as	Immediate	X
>	If CPU Demand (%) >	90	mark as	Critical	X

Drag an additional symptom / condition in to your set

Drag and drop metric to specify its condition or symptom into your alert here to create a new set

**PREVIOUS** **NEXT** **UPDATE** **CANCEL**

1. Behind *The set is met when*, Change from All to Any
2. Set the *if CPU/Demand (%)* > than respective, 70, 80, and 90 percent
3. Behind *mark as*, choose **warning**, **immediate**, and **Critical** respectively
4. Click **Next**

## Adding Recommendations

[66]

The screenshot shows the 'Recommendations' step in the alert configuration process. A search for 'cpu' has been performed, resulting in a list of recommendations. Two specific recommendations are highlighted with red arrows and numbered circles: 'Power Off VM' (2) and 'Set CPU Count for VM' (3). The 'NEXT' button at the bottom is also highlighted with a red box and a '3'.

Description	Action	Defined By	Modified By
For best CPU performance set the ESXi power mana...		vCenter	admin
For Production Virtual Machines, please assess the tr...	Power Off VM	User	holadmin@c...
If the host has 1 CPU, upgrade the host or use a host ...		vCenter	admin
If the load balancer CPU utilization is higher than syst...		NSX-T	admin
If the virtual machine has 1 vCPU, add an additional v...	Set CPU Count f...	vCenter	admin
If the virtual machine has multiple vCPUs, add an add...		vCenter	admin
If virtual machine CPU reservation is set, decrease th...	Set CPU Resour...	vCenter	admin
In the NSX UI, navigate to System   NSX Application ...		NSX-T	admin

1. Search for cpu related recommendations, in the search field type cpu and press enter
2. Drag recommendations into sets into your alert and order them by priority.
3. Click **Next**

Note: Pay notice to which recommendations that has Actions attached to them

## Attach policies

[67]

Create Alert Definition Virtual Machine ?

Home / Alerts / Alert Definitions

1 - Alert      2 - Symptoms / Conditions      3 - Recommendations      4 - Policies

Select which policies you would like to activate this alert in. You may also customize thresholds per policy.

- HOL Policy
- HOL Test Policy
- vSphere Solution's Default Policy (May 12, 2023 10:12:11 AM)

PREVIOUS    NEXT    **CREATE**    CANCEL

Normally we would select a specific policy that were made for certain Application Servers, a business unit, or a grouping of objects. Policies are normally used for Resource Allocation, Compliance and SLA Requirements, Business Priorities, or Experimentation and Testing.

1. Tick the check mark on All policies
2. Click Create

CONGRATULATIONS!

We have now finished creating an alert definition with symptoms, the last task is to Create a Custom Notification that will use this Alert definition and also use the previously created [Payload templates](#) from Module 1, to send mail using the [Add outbound instance](#)

On to the next..

## Custom Notifications

[68]

Notifications can be delivered through email and are the means through which relevant stakeholders are informed about triggered

alerts or events. When an alert is generated based on the defined criteria, We send notifications to administrators, IT teams, or other designated recipients.

In this lesson we will utilize everything we have added to our Alerts until Now.

## Notifications Page

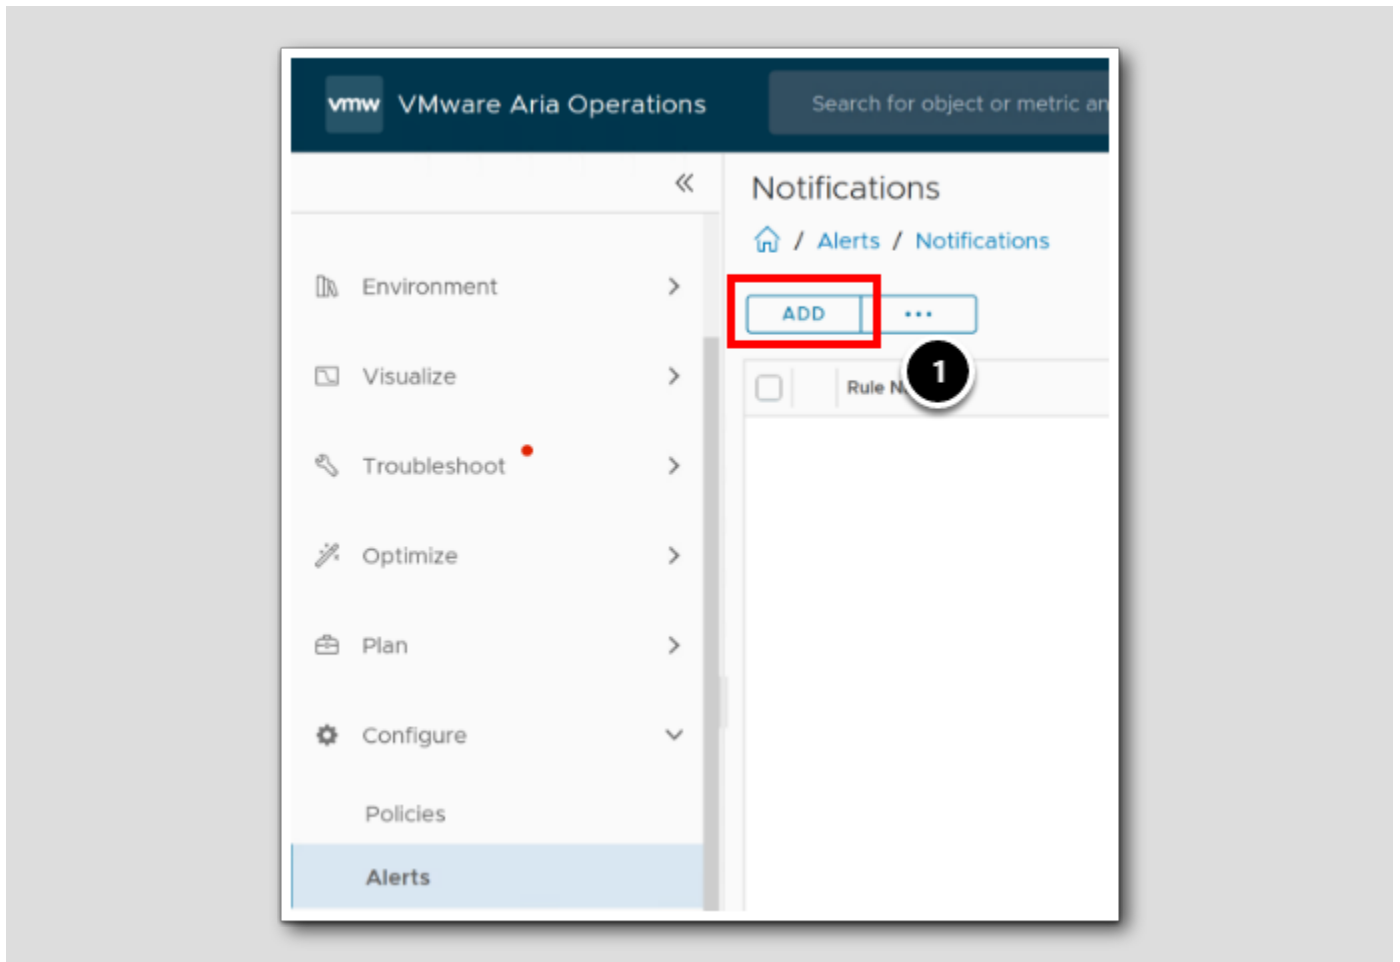
[69]

The screenshot shows the VMware Aria Operations interface. The top navigation bar includes the VMware logo and the text "VMware Aria Operations" on the left, and a search bar "Search for object or metric and more..." on the right. A left-hand navigation menu contains several items: "Optimize", "Plan", "Configure", "Policies", "Alerts", "Super Metrics", "Applications and Services", "Cost Drivers", "Custom Profiles", and "Configuration Files". The "Alerts" item is highlighted in blue and has a red box around it with a "2" in a circle. The "Configure" item has a red box around it with a "1" in a circle. The main content area is titled "Alerts" and contains three sections: "Alert Definitions" (with a bell icon), "Symptom Definitions" (with a magnifying glass icon), and "Actions" (with a mouse cursor icon). The "Alert Definitions" section describes creating and editing alert definitions. The "Symptom Definitions" section describes creating and editing descriptions of situations. The "Actions" section describes creating and editing actions to make changes to objects. A red box highlights the "Notifications" section, which includes an envelope icon and text describing how to define and modify notification settings. A "3" in a circle is placed next to this section.

1. Click Configure
2. Click Alerts
3. Click Notifications

## Add Notification

[70]



1. On the Notifications page, Click ADD

## Notification properties

Notifications

Home / Alerts / Notifications

1 - Notification      2 - Define Criteria      3 - Set Outbound Method      4 - Select Payload Template

Name: OurCompany Email Notification **1**

Description: Send Email Notifications using Company email payload template when Alert definition OurCompany Custom Alert Definition is triggered **2**

Notification Status:  **3**

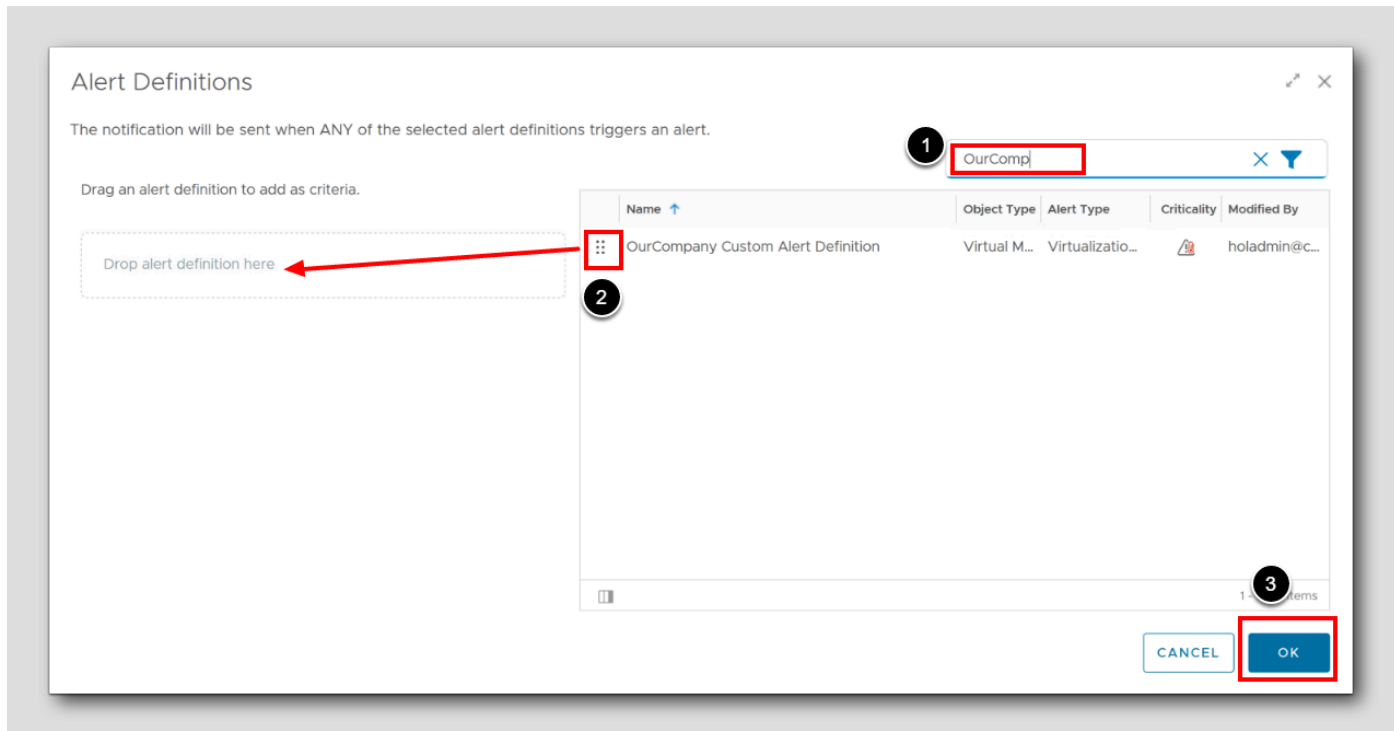
**4** PREVIOUS NEXT UPDATE CANCEL

1. Since we already have used "OurCompany", In the Name box, Type OurCompany Email Notification
2. In the Description, type Send Email Notifications using Company email payload template when Alert definition OurCompany Custom Alert Definition is triggered
3. Make sure Notifications are activated as shown
4. Click NEXT





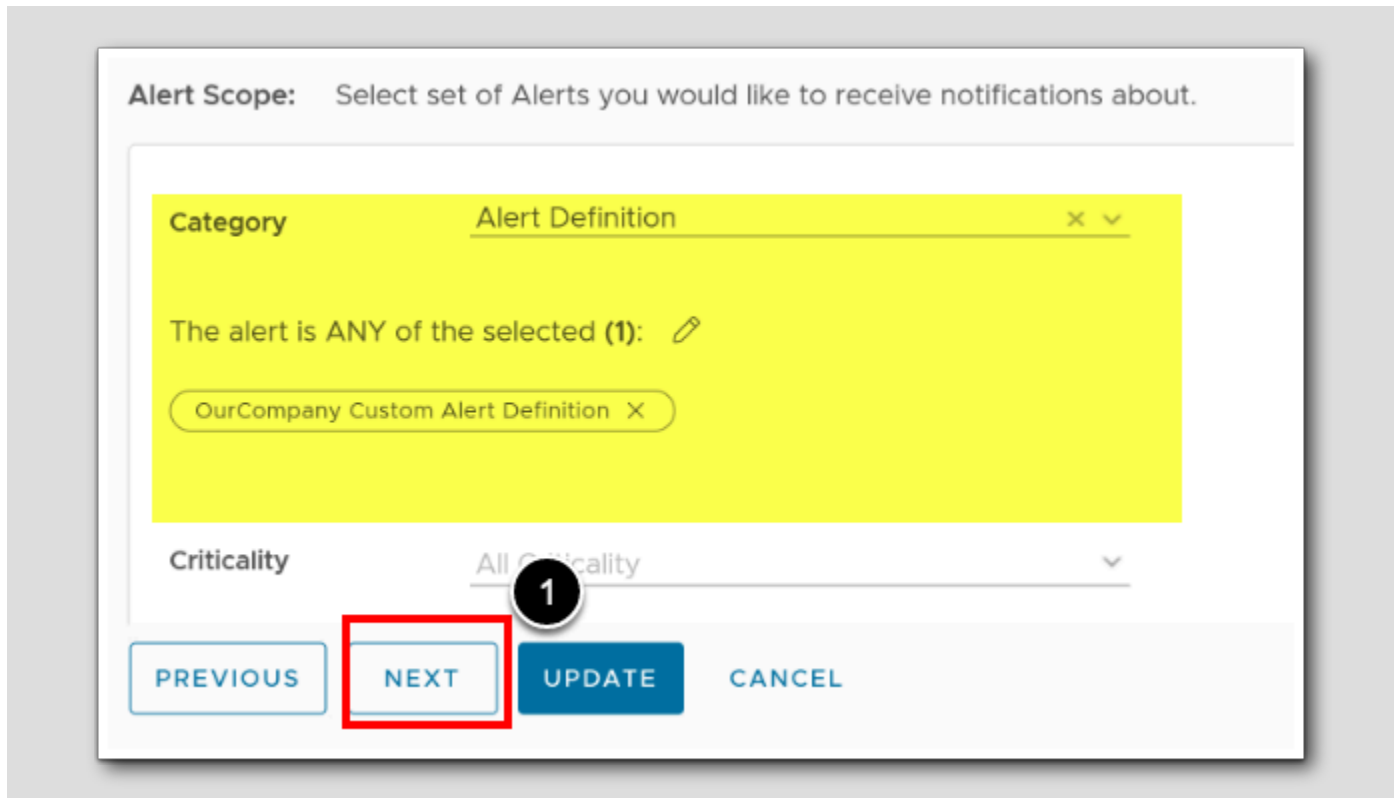
## Adding our Alert Definition



1. In the search box, type **OurComp** and Press **ENTER**
2. Drag and Drop the **OurCompany Custom Alert Definition**
3. Click **OK**

## Alert scope result

[74]



The screenshot shows a dialog box titled "Alert Scope: Select set of Alerts you would like to receive notifications about." The dialog has a yellow background for the main content area. At the top, it says "Category" and "Alert Definition" with a close button (X) and a dropdown arrow. Below this, it says "The alert is ANY of the selected (1):" followed by a pencil icon. Underneath, there is a single selection: "OurCompany Custom Alert Definition" with a close button (X). At the bottom, there is a "Criticality" dropdown menu set to "All Criticality". A red box highlights the "NEXT" button, and a black circle with the number "1" is placed over the "UPDATE" button.

1. Click NEXT

## Outbound Method

Notifications

Home / Alerts / Notifications

1 - Notification      2 - Define Criteria      3 - Set Outbound Method      4 - Select Payload Template

Pick the outbound method you would like to use to send your notification.

Outbound method: Standard Email Plugin (1)      Outbound email Instance (2)      + CREATE NEW INSTANCE

PREVIOUS      NEXT (3)      UPDATE      CANCEL

1. For Outbound Method, Choose Standard Email Plugin
2. Choose our previously created Outbound email instance from Create New Outbound Instance in Module 1 - Configuring and Managing Alert Notifications
3. Click NEXT

### Select Payload Template

**Notifications**  
Home / Alerts / Notifications

1 - Notification      2 - Define Criteria      3 - Set Outbound Method

Pick a payload template to include in the notification. The template includes additional content about the alert or object.

**Payload Template**  
Company Email Template (selected)  
Company Email Template  
Default Email Template

> Company Email Template

**Recipient(s)** holadmin@corp.vmbeans.com ⓘ  
**Cc Recipients** e.g. example@domain.com ⓘ  
**Bcc Recipients** e.g. example@domain.com ⓘ

**Notify again** 1 ⓘ  
**Max Notifications** 5 ⓘ  
**Delay to notify** e.g. 15 (Optional) ⓘ  
**Description** e.g. For Mr. Smith (Optional) ⓘ Deprecated

PREVIOUS    NEXT    **CREATE**    CANCEL

### 1. Under Payload Template, Select the Company Email Template

Note: We created The Custom Payload Template in the Chapter Payload templates in Module 1 - Configuring and Managing Alert Notifications

### 2. Set the Recipient(s) to holadmin@corp.vmbeans.com

### 3. Set the Notify again to 1 minute

### 4. Set the Max Notifications to 5

### 5. Click CREATE

## Prepare for Alerts

[77]

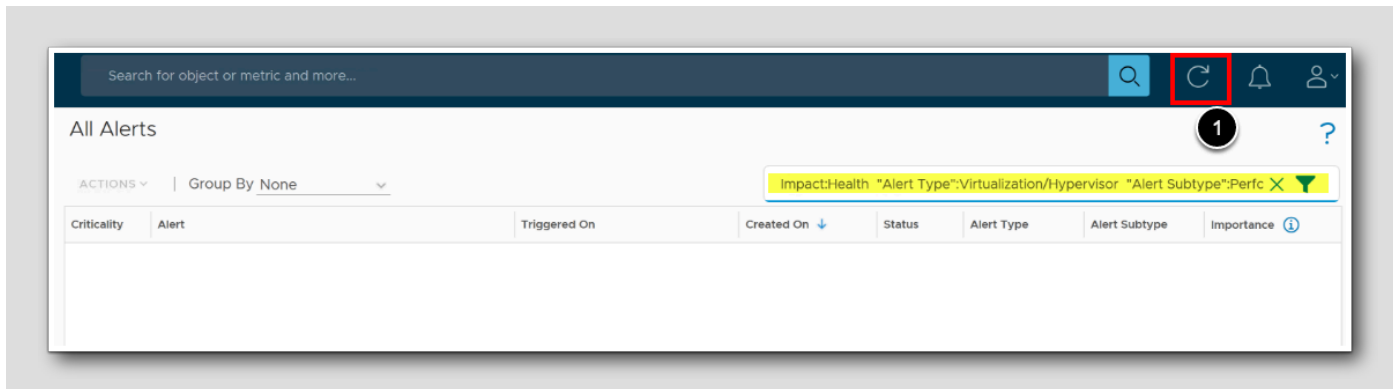
The screenshot shows the VMware vSphere Alerts configuration interface. The left sidebar has the 'Alerts' menu item highlighted. The main area displays a list of alerts under 'All Alerts' with columns for Criticality, Alert, Triggered On, and Created. A filter panel on the right is open, showing a search bar and several filter criteria: Impact (Health), Alert Type (Virtualization/Hypervisor), Alert Subtype (Performance), and Status (Active). The 'APPLY' button is highlighted in red.

We are going to test our alerts for the next 10 minutes, but to narrow down the results, we will filter out what we don't need.

1. Click Troubleshoot
2. Click Alerts
3. Click the filter Icon
4. Impact select Health
5. Alert Type select Virtualization/Hypervisor
6. Alert Subtype select Performance
7. Status select Active
8. Click APPLY

## Ready for alerts

[78]

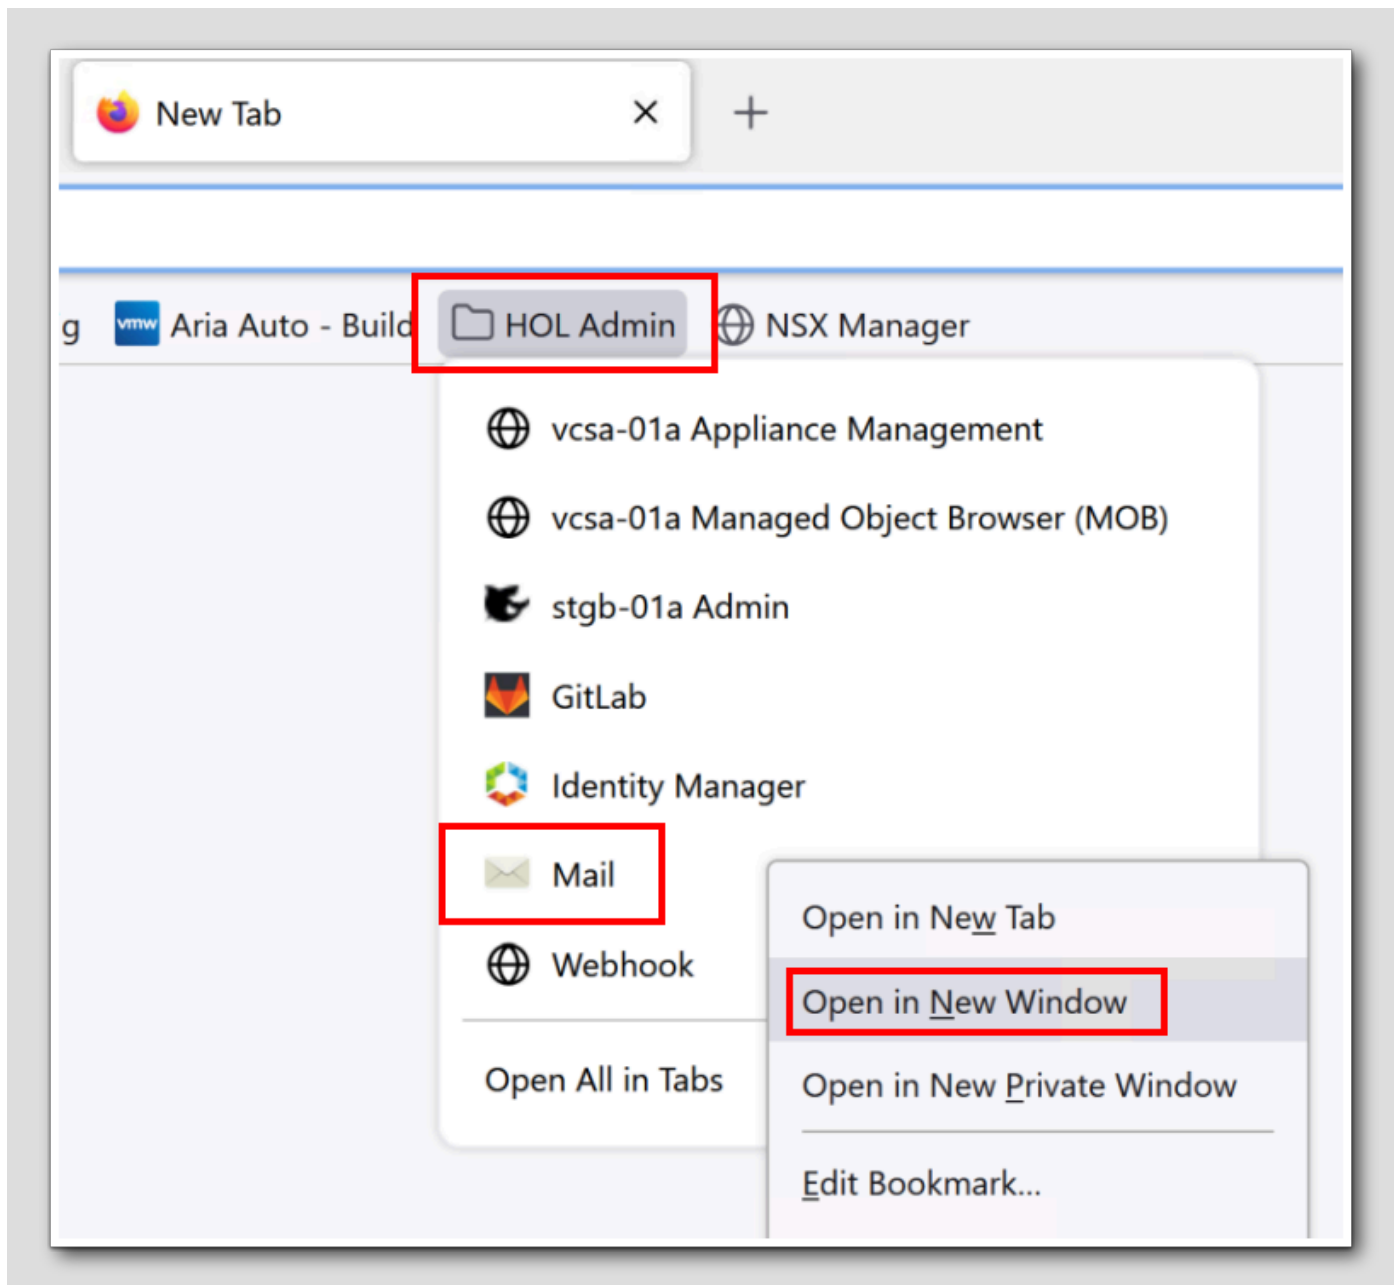


We are ready for our specific alerts. Remember that the alerts will take some 5-10 minutes to show up since Aria Operations have collection cycles every 5 minutes

1. During our test, Refresh this page

## Preparing the Email client

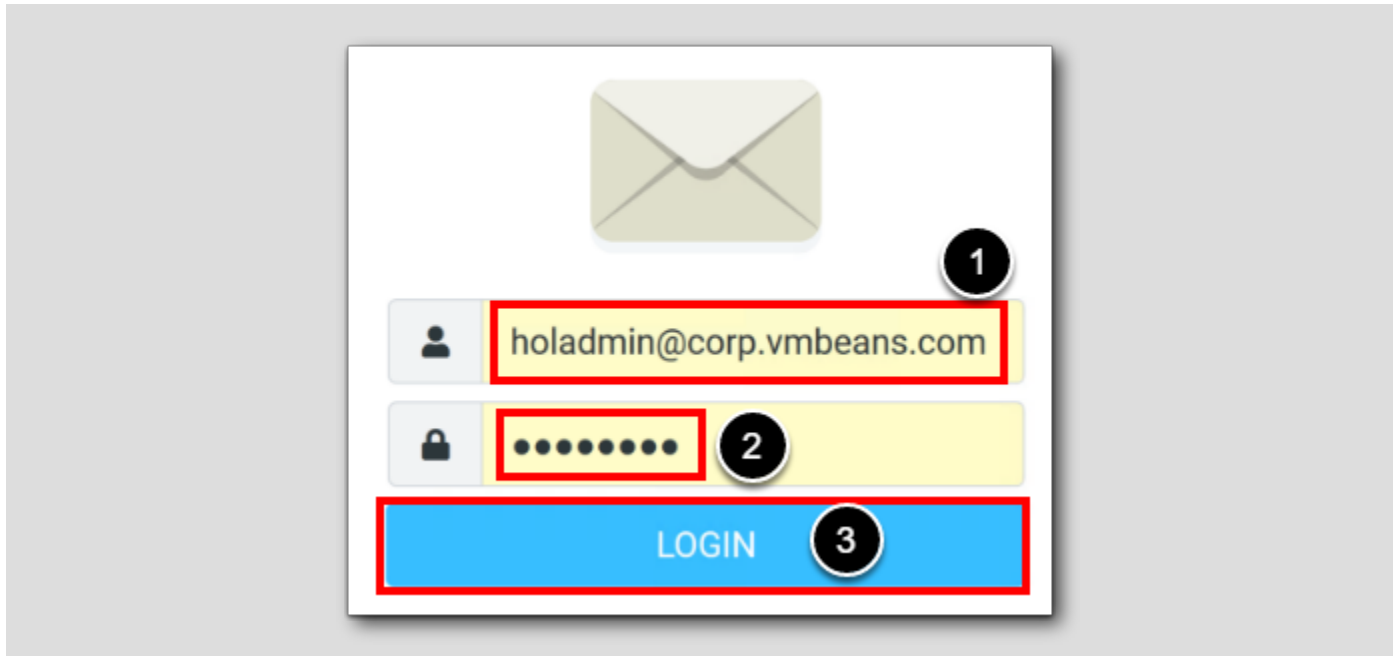
[79]



1. In the Browser, Click HOL Admin
2. Right Click on Mail
3. Choose Open In New Window

## Log In to the email client

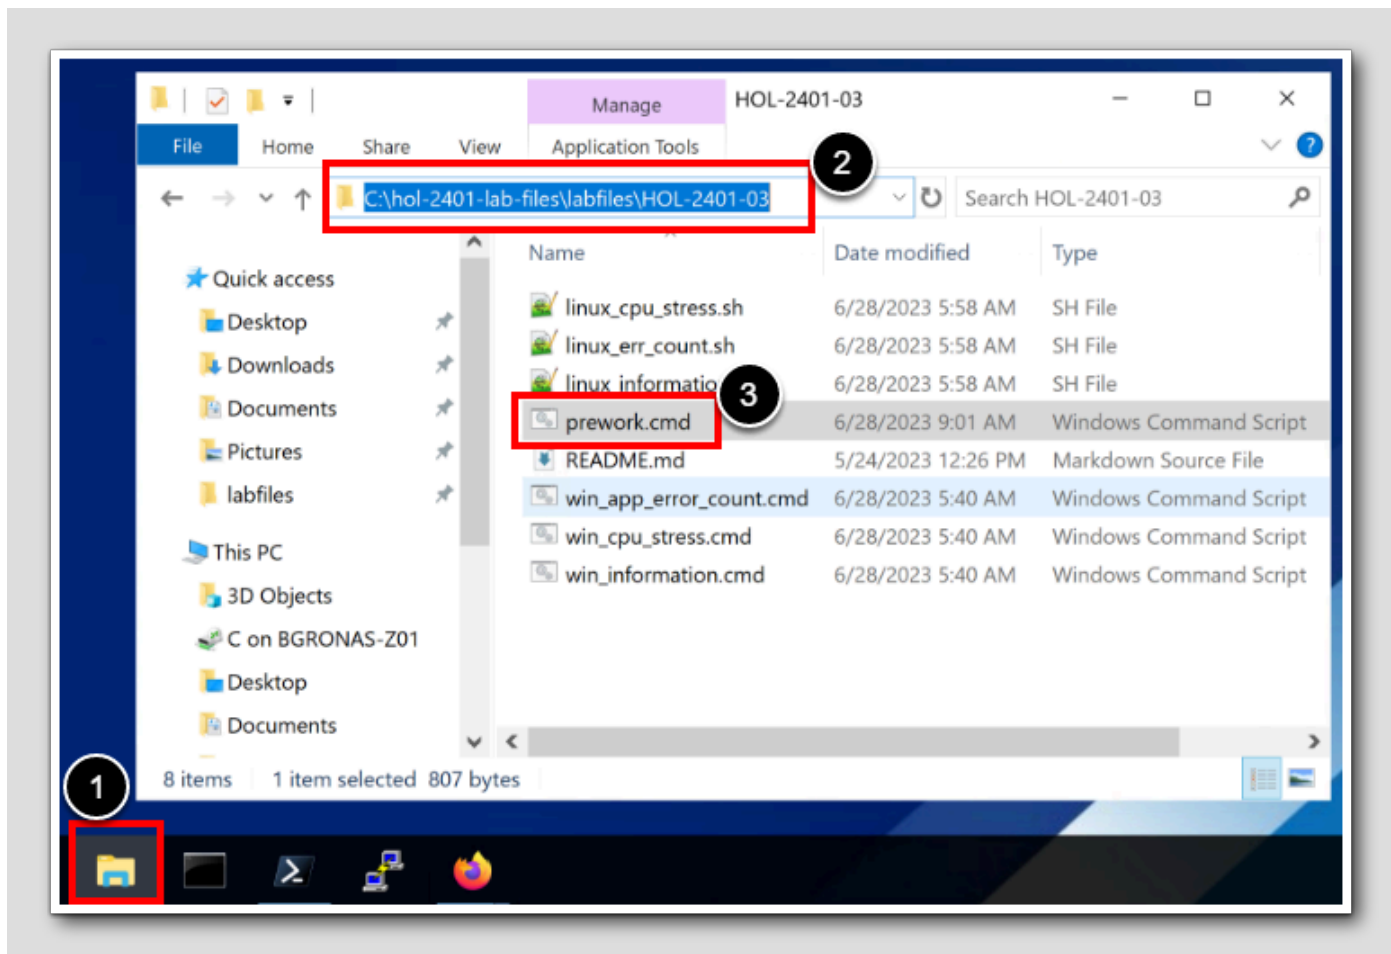
[80]



1. Type holadmin@corp.vmbeans.com
2. Type VMware!
3. Click LOGIN



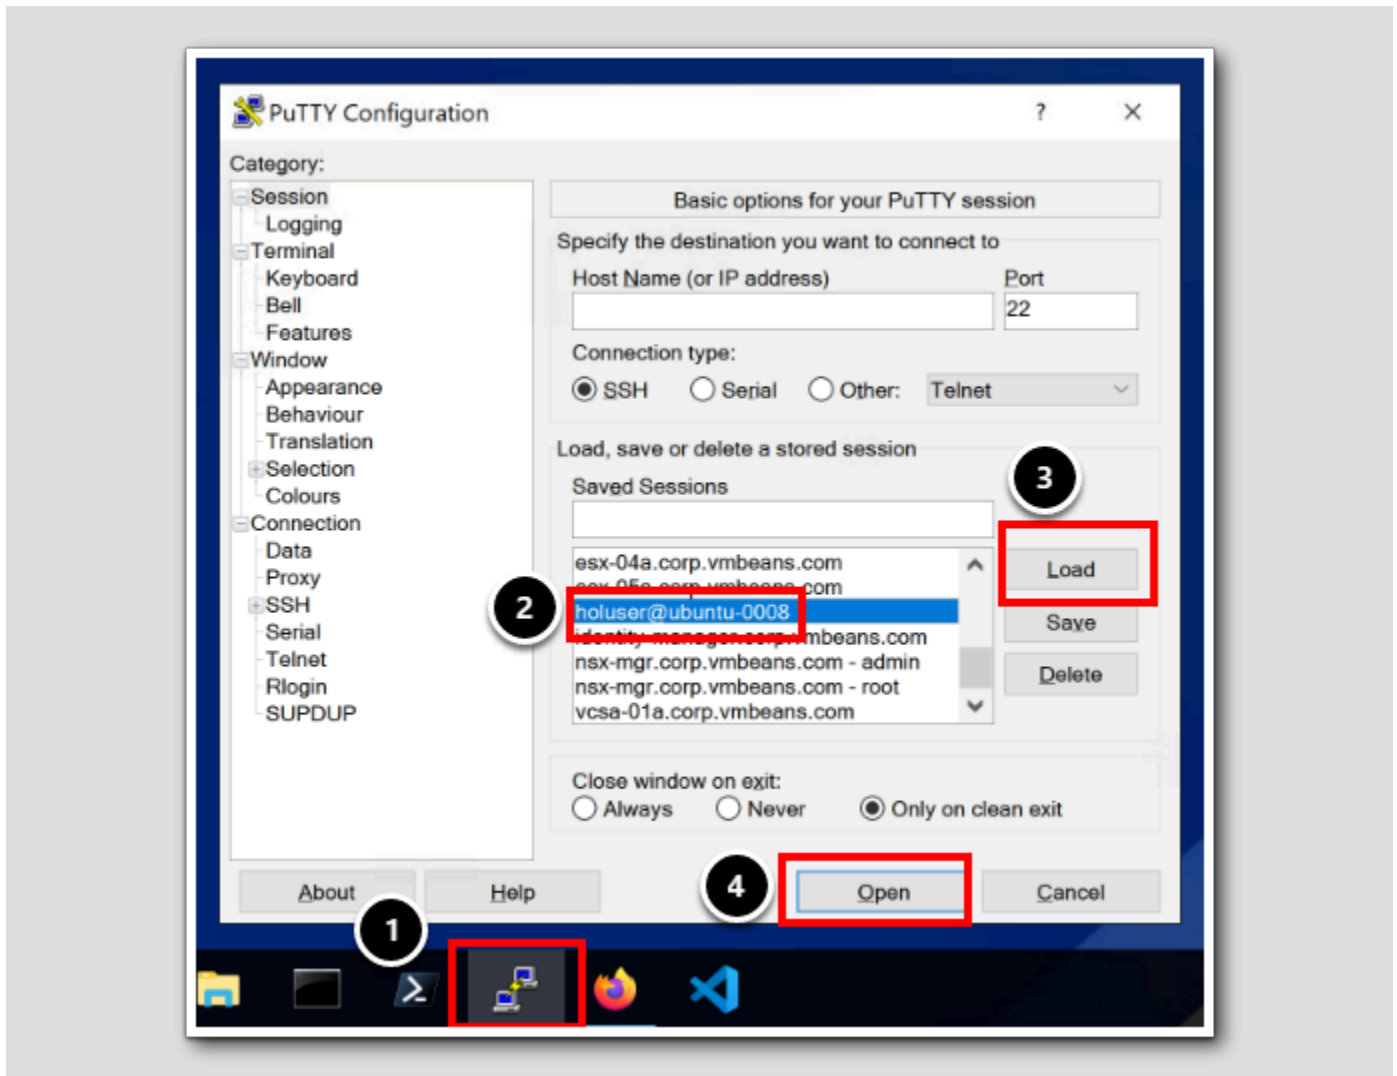
Prepare &quot;the blast&quot;



We will now copy some scripts over to an Ubuntu server to run a little CPU stress. The script: *linux\_cpu\_stress.sh* will help us create CPU load for approximately 10 minutes. During this time period, we will get e-mails and Aria Operations Alerts.

1. In the console, start a windows explorer window, click on the symbol on the taskbar
2. Go to the path; C:\hol-2401-lab-files\labfiles\HOL-2401-03
3. To start copying files, and transport them over to the /root folder on the Ubuntu VM called *ubuntu-008* Double-Click on the windows script file *prework.cmd*

## Use Putty



We will use putty to connect to the linux server to run the *CPU Stress* script.

1. In the Console, on the taskbar, start **putty**
2. Scroll to find **holuser@ubuntu-0008**
3. Click **Load**
4. Click **Open**

## Become root

```

root@ubuntu-0008: ~
Last login: Tue Aug  8 12:37:56 2023 from 192.168.110.10
holuser@ubuntu-0008:~$ sudo su -
[sudo] password for holuser:
root@ubuntu-0008:~# pwd
/root
root@ubuntu-0008:~# ll
total 60
drwx----- 6 root root 4096 Aug  8 12:39 ./
drwxr-xr-x 20 root root 4096 Jun 16 19:33 ../
-rw----- 1 root root 3690 Aug  8 12:40 .bash_history
-rw-r--r-- 1 root root 3582 Jan 12 2022 .bashrc
drwx----- 2 root root 4096 Oct 12 2021 .cache/
-rwxr-xr-x 1 root root  726 Aug  8 12:39 linux_cpu_stress.sh*
-rwxr-xr-x 1 root root  137 Aug  8 12:39 linux_err_count.sh*
-rwxr-xr-x 1 root root  939 Aug  8 12:39 linux_information.sh*
-rw-r--r-- 1 root root  161 Dec  5 2019 .profile
drwxr-xr-x 3 root root 4096 Oct 12 2021 snap/
drwx----- 2 root root 4096 Jun 16 19:36 .ssh/
drwxr-xr-x 2 root root 4096 Jun 19 23:34 .vim/
-rw----- 1 root root 10500 Jun 28 20:52 .viminfo
root@ubuntu-0008:~#

```

1. In the python console, type **sudo su -** and press ENTER
2. Enter the password **VMware!!** (hidden) and press ENTER
3. To ensure you are in the **/root** catalog, enter the command **pwd** and press ENTER
4. For a long listing (**ll**), Enter the command **ll** and press ENTER

Note: You should see the shell scripts we are going to use, listed in a horrifying green color (highlighted)

## Bring the noise

```

root@ubuntu-0008: ~
Last login: Wed Jun 28 09:30:14 2023 from 192.168.110.10
root@ubuntu-0008:~# pwd
/root
root@ubuntu-0008:~# ls
linux_cpu_stress.sh  linux_err_count.sh  linux_information.sh  snap
root@ubuntu-0008:~# ll
total 64
drwx----- 7 root root 4096 Jun 28 23:42 ./
drwxr-xr-x 20 root root 4096 Jun 16 19:33 ../
-rw----- 1 root root 3943 Jun 28 13:19 .bash_history
-rw-r--r-- 1 root root 3582 Jan 12 2022 .bashrc
drwx----- 2 root root 4096 Oct 12 2021 .cache/
drwx----- 3 root root 4096 Jun 28 08:18 .config/
-rwxr-xr-x 1 root root 726 Jun 28 23:42 linux_cpu_stress.sh*
-rwxr-xr-x 1 root root 259 Jun 28 23:42 linux_err_count.sh*
-rwxr-xr-x 1 root root 939 Jun 28 23:42 linux_information.sh*
-rw-r--r-- 1 root root 161 Dec 5 2019 .profile
drwxr-xr-x 3 root root 4096 Oct 12 2021 snap/
drwx----- 2 root root 4096 Jun 16 19:36 .ssh/
drwxr-xr-x 2 root root 4096 Jun 19 23:34 .vim/
-rw----- 1 root root 11056 Jun 28 13:03 .viminfo
root@ubuntu-0008:~# ./linux_cpu_stress.sh

```

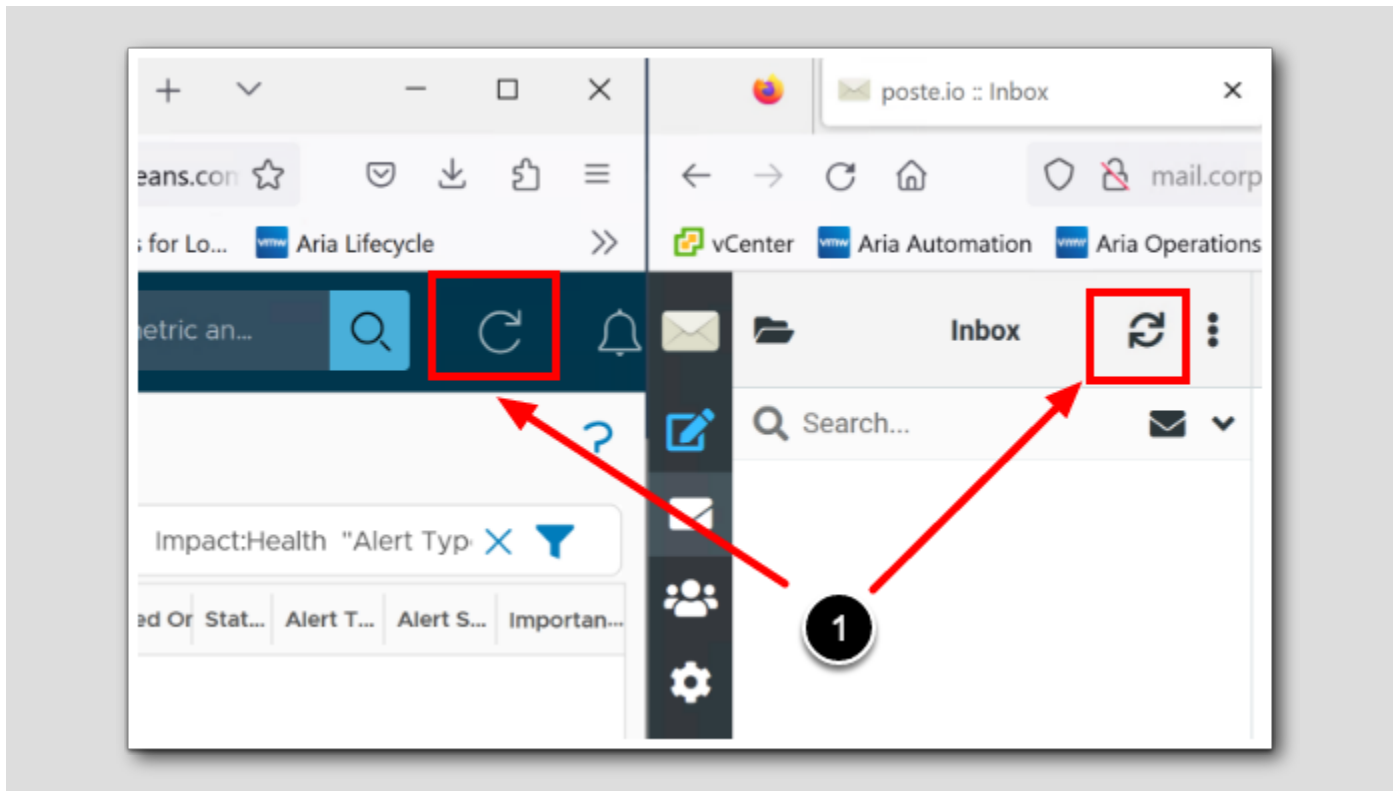
The terminal window shows a file listing with several files highlighted in green. A red arrow points to the file `linux_cpu_stress.sh*`. A red box highlights the command `./linux_cpu_stress.sh` at the bottom of the terminal. A circled number '1' is placed over the first column of the file listing.

Be ready with your mail client and Aria Operations alert page open. As soon as we start this script, it will generate CPU traffic for approximately 10 minutes, and then die.

1. Start the CPU Stress test by typing `./linux_cpu_stress.sh` and then hit ENTER

Note: You can use **Alt+Tab** to jump to your email client and Aria Operations to see what happens,

## The refreshing game



1. If we put both the ARIA Operations Alert Page and the mail client side by side, it is easier to click the Refresh button on both and monitor for alerts

NOTE: Since Aria Operations collects metrics every 5 minutes, we have to do a little waiting, but not for long. Now is a good time to take a sip of water, check your phone, send a nice text message to your loved ones, get update on the weather, and read the headlines on your favorite news page.

Using email alerts to tweak our Custom Payload Templates

The screenshot displays an email client interface with a list of messages on the left and a detailed view of a selected message on the right. The messages in the list are from 'YourNameGoesHere' and include subjects like 'warning - CANCELLED ALERT Virtualization/...', 'warning - UPDATED ALERT Virtualization/Hy...', and 'critical - New Virtualization/Hypervisor Performance...'. The detailed view shows the following information:

**critical - New Virtualization/Hypervisor Performance, on Virtu**  
**hu Jun 29 00:09:48 UTC 2023 ]**

From YourNameGoesHere on 2023-06-28 17:14

The VirtualMachine: **ubuntu-0008** is acting abnormally since **Thu Jun 29 00:09:48 UTC 2023**

**# Alert:**  
 ualization/HypervisorPerformancehealthcritical  
 Alert definition: OurCompany Custom Alert Definition

**# Symptoms:**  
 SYMPTOM SET - self  
 Conditions -

Object Name	Object ID	Metric	Message Info
ubuntu-0008	ca492859-bd5c-4b0d-a772-583516c5ce	b CPU Demand (%) > 90.0	
ubuntu-0008	ca492859-bd5c-4b0d-a772-583516c5ce	b CPU Demand (%) > 70.0	
ubuntu-0008	ca492859-bd5c-4b0d-a772-583516c5ce	b CPU Demand (%) > 80.0	

**# Recommendations:**

- Check the applications running on the virtual machines in the cluster to determine whether t
- Use vMotion to migrate some virtual machines with high CPU workload to other hosts that h
- Power Off this virtual machine to allow other virtual machines to use the CPU and memory t machine.
- Review the symptoms listed and remove the number of vCPUs from the virtual machine as r

**# Link:** <https://192.168.110.70/ui/index.action#environment/object-browser/hierarchy/ca49%2859-bd5c-4b0d-a772-583516c5ce-4102540-40400100-1000>

Our email looks and feel, will be exactly as we planned them, meaning they will look like the Custom Payload template we created in [Payload templates](#) in [Module 1 - Configuring and Managing Alert Notifications](#), here are some tips on how we could customize them:

1. First: We get 5 Critical alerts - We should tune the Payload Template to show the word critical with CAPITAL LETTERS. The second thing about this is "do we really want this to repeat 5 times?"
2. Second: We got updates as long as the error still were active and ongoing. But as the Linux server calmed down it changed from Critical to **Warning**, this is what actually fires off an update
3. Third: we get a CANCELED ALERT, this means the problem was either fixed or went away (the CPU load script ran for just 10 minutes)
4. Nicely underway, we got explanations and recommendations we needed to resolve the problem, or wait for the problem to settle down by itself. Which it does.

## Aria Operations Critical Alert

[87]

The screenshot shows the 'All Alerts' page in VMware Aria Operations. At the top, there are filters for 'Impact:Health' and '"Alert Type":Virtualization/'. Below the filters, there is a 'Group By Time' dropdown and a 'REFRESH' button. The main content is a table of alerts. The first alert is highlighted with a red box and a circled '1'. The alert name is 'OurCompany Custom Alert Definition'. The table columns include Criticality, Alert, Triggered On, Created On, Status, Alert Type, Alert Subtype, and Importance. The alert is triggered on 'ubuntu-0008' at '5:44 PM' with a status of 'Warning' and an importance of 'Very High (1..)'.

Criticality	Alert	Triggered On	Created On	Status	Alert Type	Alert Subtype	Importance...
Critical	OurCompany Custom Alert Definition	ubuntu-0008	5:44 PM	Warning	Virtualiza...	Performa...	Very High (1..)

So firstly we get a red Critical alert. Notice the name is "ourCompany Custom Alert Definition. Maybe the alert definition name could be a little smarter. Review what we did in the [Custom Alert Definition](#) chapter. Now that we know what this alert is for, we can just rename it.

1. We can click the alert to investigate more

## Recommendation with an Action

ubuntu-0008

Virtualization/Hypervisor Pe

Started on: 5:44:48 PM

VIEW I

Alert Details Related Alerts Potential Evidence

Recommendations < 3 of 4 > 1

Power Off this virtual machine to allow other virtual machines to use the CPU and memory that is being wasted by this virtual ma

POWER OFF VM

Alert Basis Acti

1. Self - Virtual Machine any

Conditions

✓ The **Critical** condition CPUIDemand (%) > 90 % has been met on ubuntu-0008

TROUBLESHOOT WITH LOGS

12:00 PM 01:00 PM 02:00 PM 03:00 PM 04:00 PM 05:00 PM 06:00 PM

— ubuntu-0008 - CPUIDemand (%)

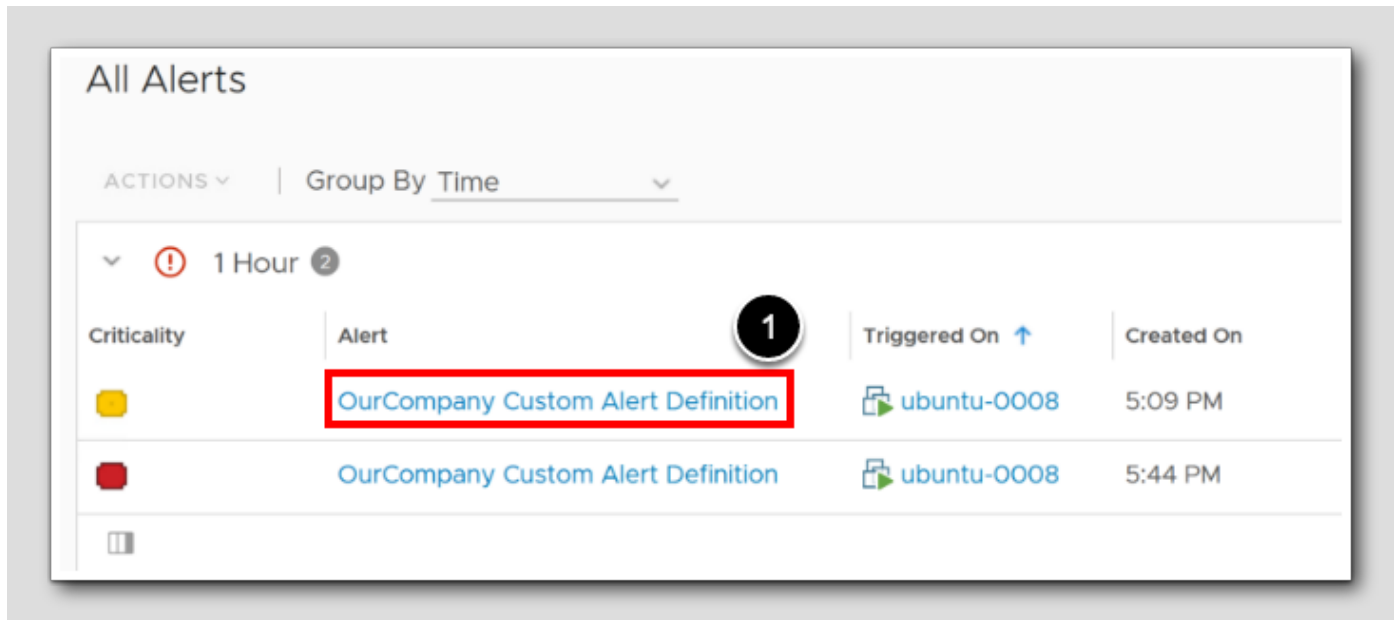
Remember when we set up that alert? We added four recommendations, and one of them lets us power off the VM with just a click! This capability enables us to swiftly respond and proactively address Alerts that arises, and stay ahead of the game.

1. To turn the pages of the four Recommendations, click the left and right Arrows

Go back to all the alerts (not shown)



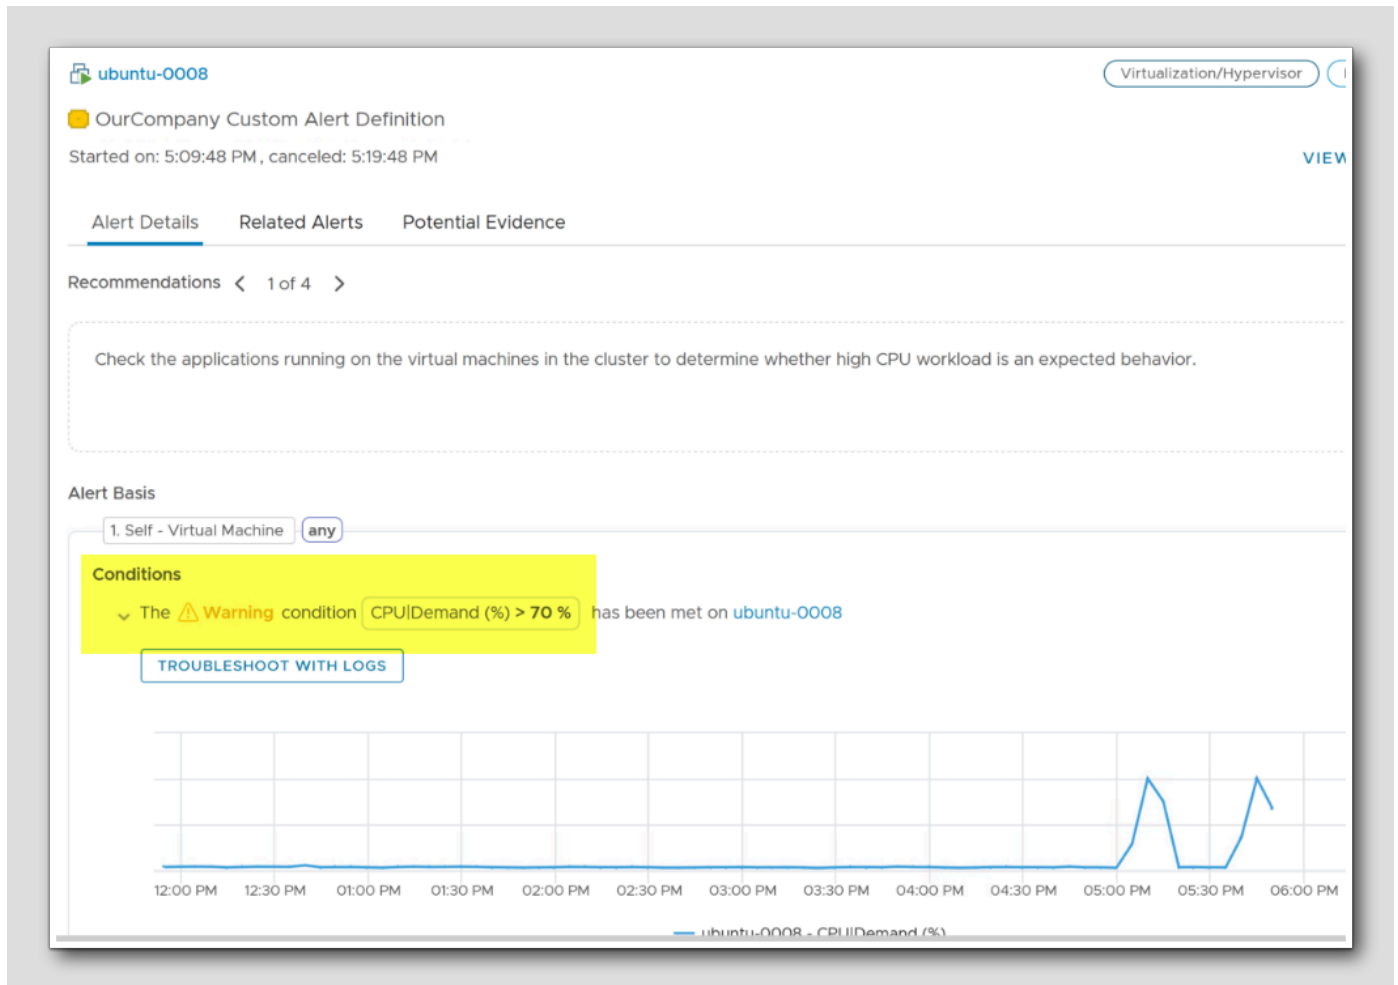
## Warnings



After a while, you will see the alert as yellow, meaning we went from a critical level to a warning level. In the picture I've deleted some of the filters to see more alerts from different create times.

1. Click the warning Alert

## Warning Alert



As you can see the **warning** alert is no different from the critical alert except for what it reports out of what we set as a symptom in our alert definition in the lesson [Change the Symptoms conditions](#).

## Conclusion

[91]

### Challenge

Try creating an Alert that will use the webhook instead of the email! It should be not be too hard.

### Alerts summarize

- While there are no ACTIVE alerts, The Alerts window is not showing anything since we
- When the Linux server CPU Load is a total ravage, the alert is showing as New, Red, Critical
- When the Linux server 'cools down' we get an updated **warning** state
- When the alert is canceled, the alert window will not show an ACTIVE alert, but if we filter out to show all alert, not just active, we will see previous alerts as well.

This is the end of this lab.

## Conclusion

[92]

Harnessing the capability to customize alert definitions and notifications in Aria Operations plays an essential role in optimizing the monitoring and incident response mechanisms of your VMware infrastructure. This segment offered critical insights that facilitate the proficient customization of these features in alignment with your specific operational requirements. The implementation of such customizations equips us to uphold proactive monitoring, rapid issue detection, and prompt resolution, leading to an enhancement in IT operations management and overall infrastructure stability.

## You've finished Module 2

[93]

Congratulations on completing the lab module.

If you are looking for additional general information on Aria Operations, try one of these:

- **VMware Product Public Page - Aria Operations:** <https://www.vmware.com/products/aria-operations.html>
- **Aria Operations - Documentation:** <https://docs.vmware.com/en/VMware-Aria-Operations/index.html>
- **Aria Operations - Configuring Alerts and Actions:** <https://docs.vmware.com/en/VMware-Aria-Operations/8.12/Configuring-Operations/GUID-62D6F047-7743-4B1A-90EF-F97B15D2E408.html>

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 3 - Application Monitoring with VMware Aria Operations (35 minutes) Advanced

### Introduction

[95]

Script execution enables task automation and data collection, while analyzing top processes enhances performance and resource efficiency. We will explore the significance of Aria Operations' script execution and top process analysis, providing insights for organizations aiming to achieve operational excellence.

- Script Execution: Automate tasks, collect data, and perform custom operations.
- Benefits: Streamline workflows, enhance productivity, and gain deeper insights.
- 
- Top Process Analysis: Identify resource consumption, optimize allocation, and troubleshoot performance issues.
- Benefits: Crucial for maintaining optimal performance and resource efficiency.

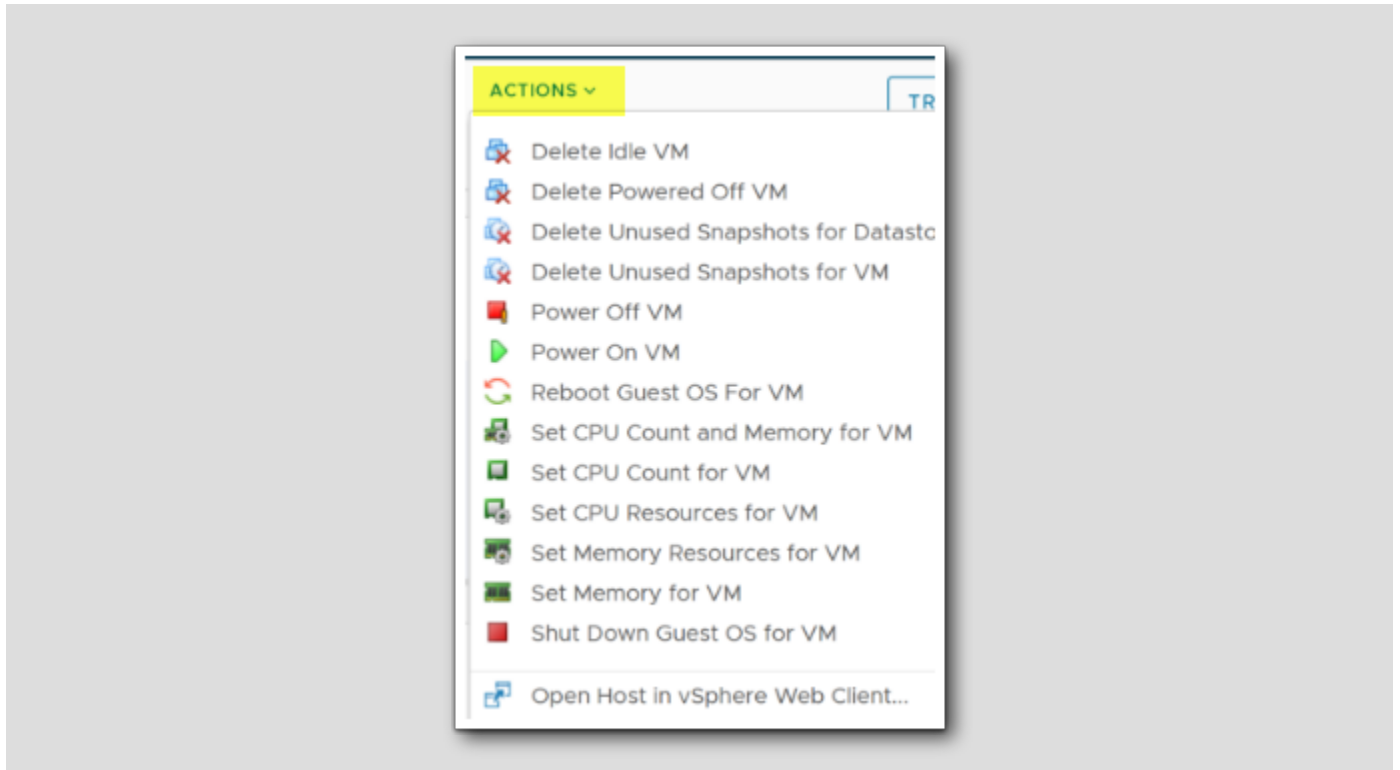
Lab Captain: Bengt Grønås, Senior Specialist Solution Engineer, Norway

Introduction to Actions

[96]

About Aria Operations Actions

[97]



"Actions" are tasks you can perform to either obtain data from, or modify, objects in the systems being monitored.

These tasks are generally added by "solutions" and are accessible from various locations within the user interface, including the object Actions menu, list and view menus, and some dashboard widgets. Actions can also be linked to alert definitions.

Actions can be divided into different categories.

**Read Actions:** These actions are used to extract data from the target objects.

**Update Actions:** These actions allow you to make changes to the target objects. For example, if an alert indicates a virtual machine is low on memory, you could set up a specific action to increase its memory, likely solving the problem.

**Ad Hoc Actions:** On-demand manual tasks such as powering a virtual machine on/off, snapshot deletion, or guest OS reboot.

**Recommendations/Automated Actions:** Suggestions for action based on vROps analysis, manually executable or automated per policy. These may include workload balancing or VM memory enhancement.

**Custom Actions:** User-defined actions built on vRealize Orchestrator scripts or workflows, manually or automatically triggered, such as provisioning a new VM on reaching capacity thresholds.

**Scheduled Actions:** Actions timed to occur at specific periods, useful for out-of-business-hours operations like maintenance tasks or batch jobs.

**Remediation Actions:** Actions meant to resolve triggered alerts, either manually executed or automated.

Note: Remember that the actions you can perform will depend on the permissions granted to your user account and the integrations set up in your Aria Operations environment. Always take care when running actions, especially when automating tasks. A poorly planned action can unintentionally impact your IT environment."

## Objects and Actions

[98]

Actions can be performed on a wide variety of objects that are part of your virtual or physical infrastructure, including but not limited to:

- **Virtual Machines (VMs):** for example powering on/off, migrating to a different host, increasing/decreasing memory or CPU, etc.
- **Hosts:** Entering/exit maintenance mode, power on/off, etc.
- **Datastores:** Enabling/disabling Storage DRS, changing the Storage DRS automation level, etc.
- **Clusters:** Enabling/disabling DRS or HA, changing the DRS automation level, etc.
- **Resource Pools:** Adjusting the CPU/memory resources.
- **Networks:** Modifying network settings.

The types of actions you can perform depend on your Aria Operations configuration and the solutions you have installed, and they require proper permissions, also, the actions might vary based on the versions of both Aria Operations and vSphere you are using.

## Log in to Aria Operations

[99]

We will log in to a live instance of Aria Operations running in this lab.

## Open the Firefox Browser from the Windows Task Bar

[100]




If the browser is not already open, launch Firefox.

1. Click the Firefox icon in the Windows Quick Launch Task Bar at the bottom of the screen.

## Navigate to Aria Operations

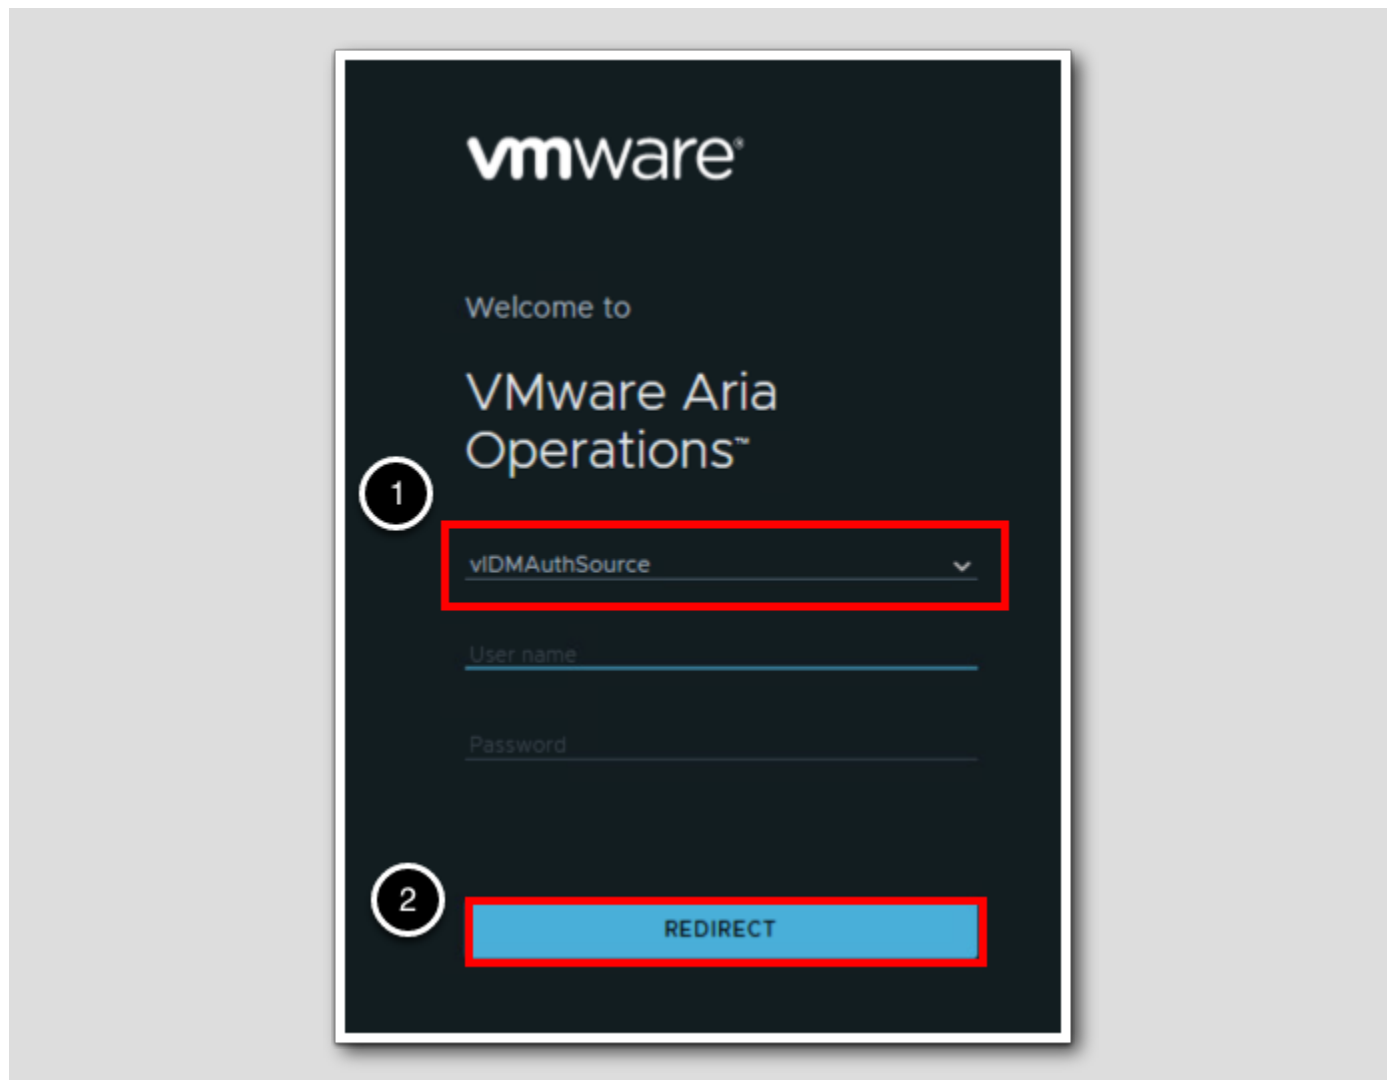
[101]



1. Click the Aria Operations bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[102]



Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

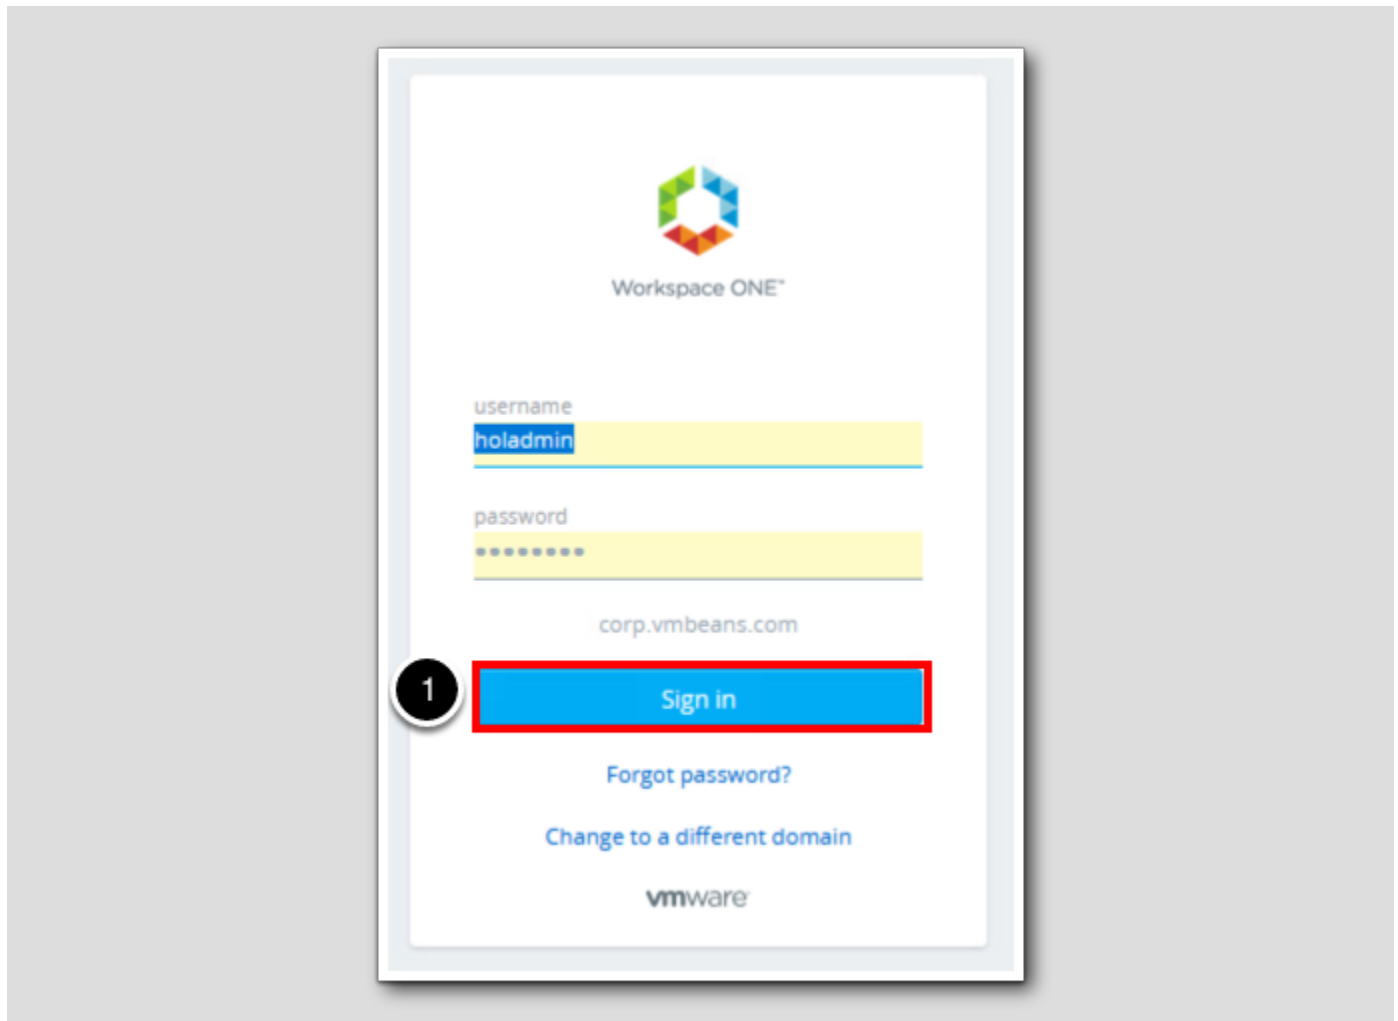
vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.



## VMware Identity Manager Login

[103]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

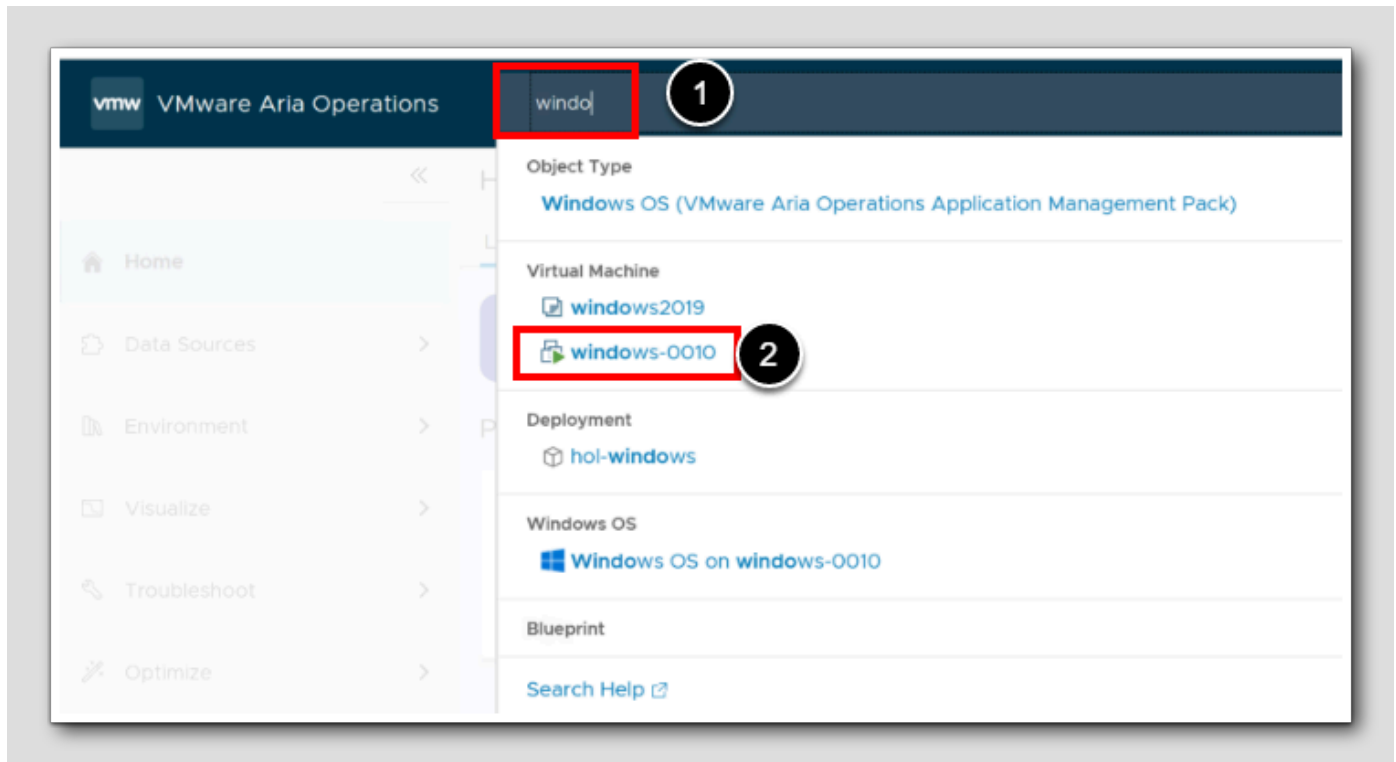
1. Click Sign in

## Script Execution and Top Processes

[104]

In this section we will run actions on ESXi hosts and on Virtual Machines. With Virtual machines with the Telegraf agent installed, we will run in-guest actions to perform tasks remotely. First off we will run a Built In action.

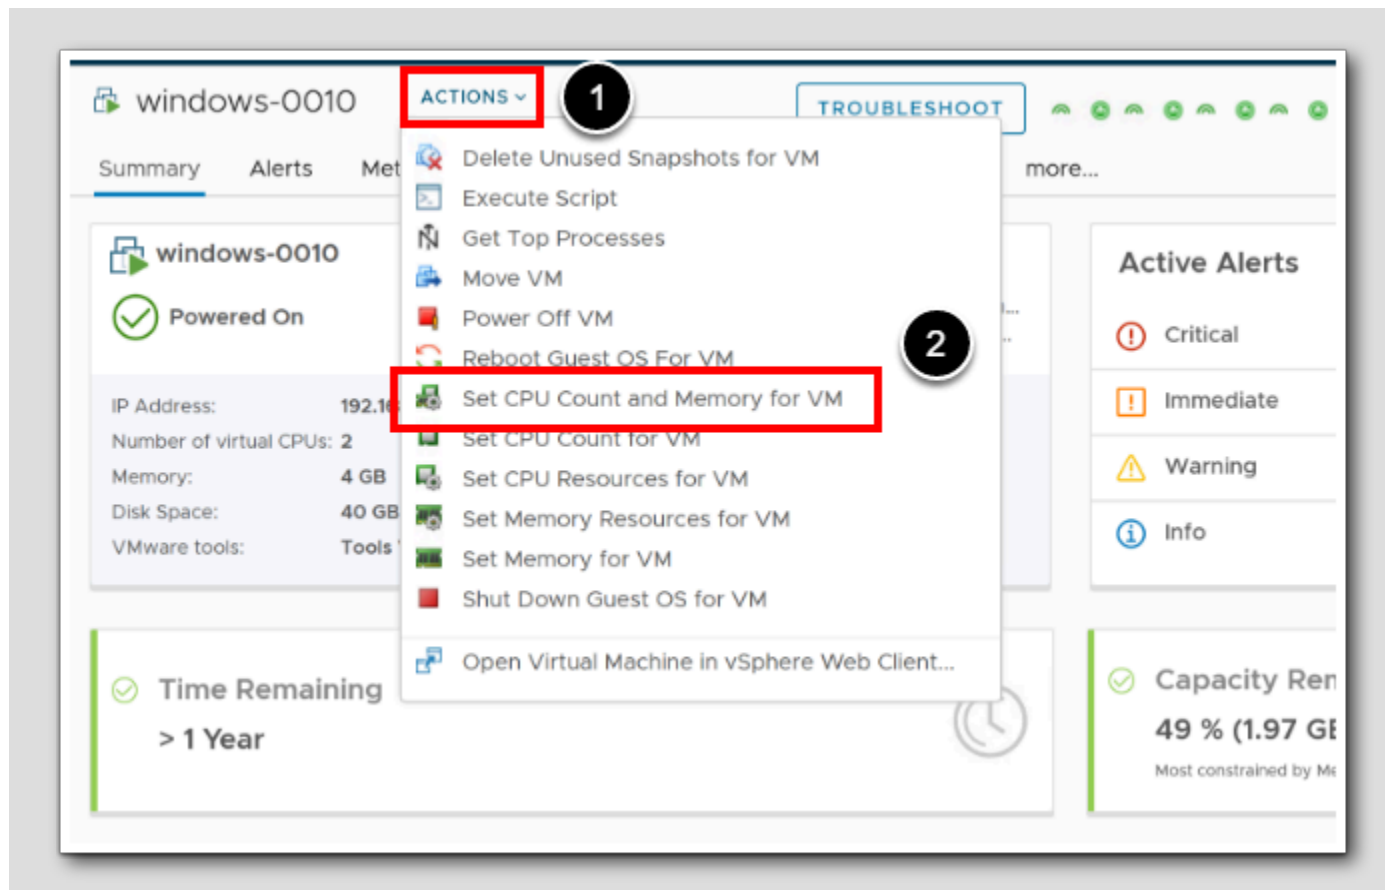
## Find a virtual machine



We are going to search for a Virtual machine to run actions on. We will sort out a windows machine we already know is online.

1. In the search field on top start searching by typing 'windo'
2. We will choose a windows machine that shows Powered On, choose windows-00xx (name can differ)

## Actions menu



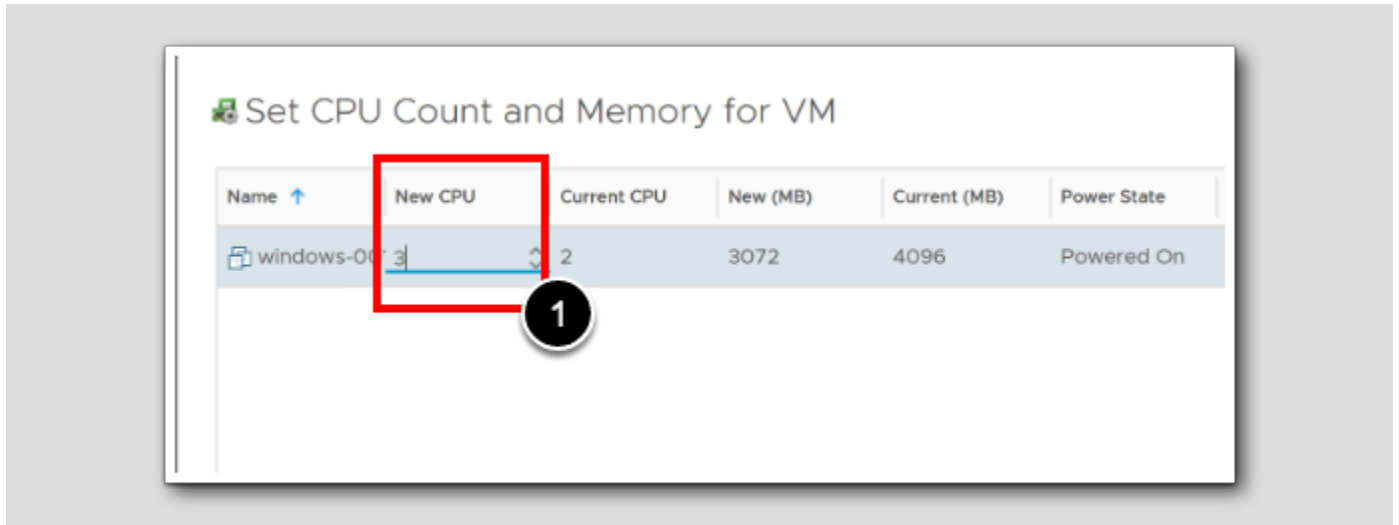
Note: Instead of searching for a specific VM, all the virtual machines would be located in Aria Operations under: Environment > Object Browser > All objects > vCenter > Virtual Machine.

During a performance troubleshooting and rightsizing scenario, we have found that this virtual machine needs more resources to perform efficiently. We know we have to add 1 more virtual CPU, and just 1 GB of RAM for this Virtual Machine to run without performance issues. (because Aria Operations told us so)

1. Pull down the Actions menu
2. Choose Set CPU Count and Memory for VM

## Add more CPU

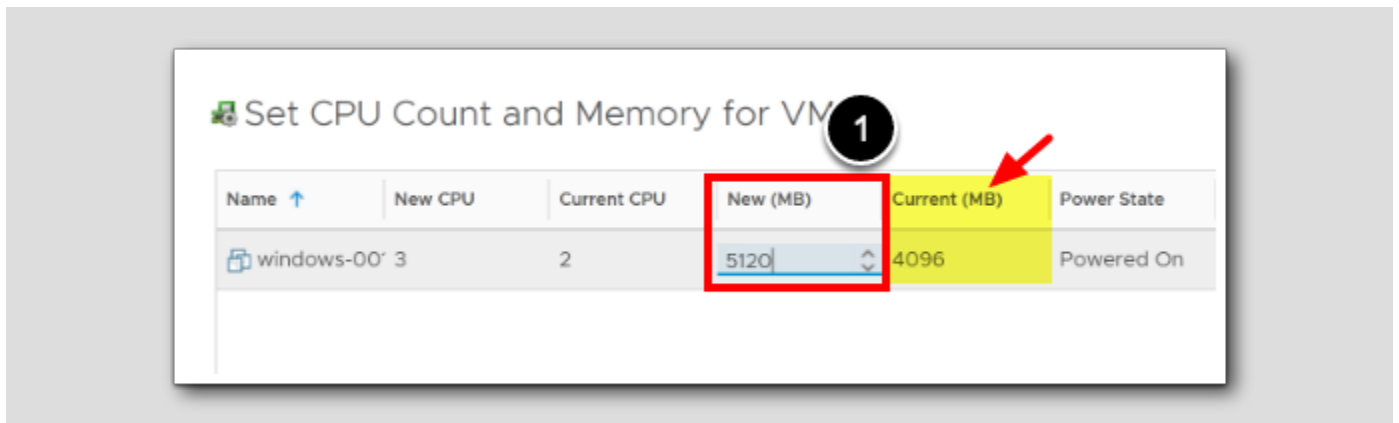
[107]



1. Change the New CPU count from '2' to 3

## Add more Memory

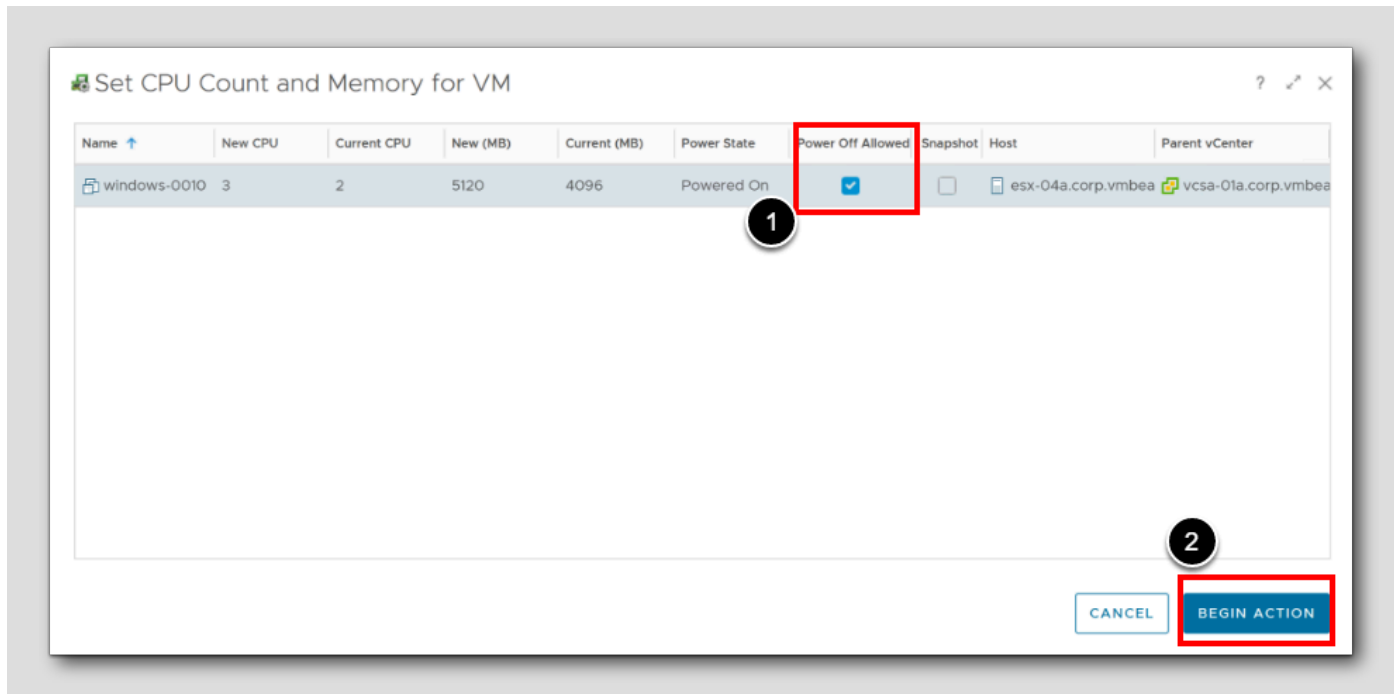
[108]



We are going to add 1024 MB memory to the Current (MB) memory. Which means  $4096 \text{ MB} + 1024 \text{ MB} = 5120 \text{ MB}$  in total.

1. Change New (MB) to 5120

## Allow power off



To complete this action, the machine would have to be powered off and on, in order to be reconfigured.

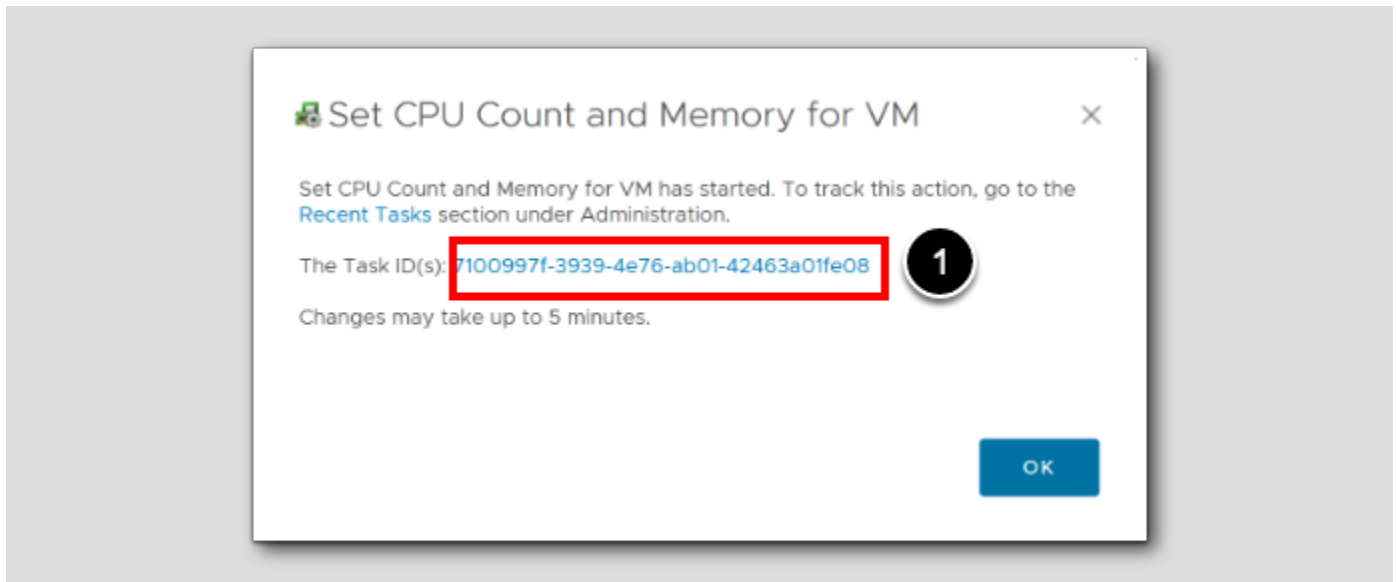
1. Check the checkmark **Power Off Allowed**

We will not take a snapshot before this action so leave it unchecked

2. Click **BEGIN ACTION**

Go to recent tasks

[110]



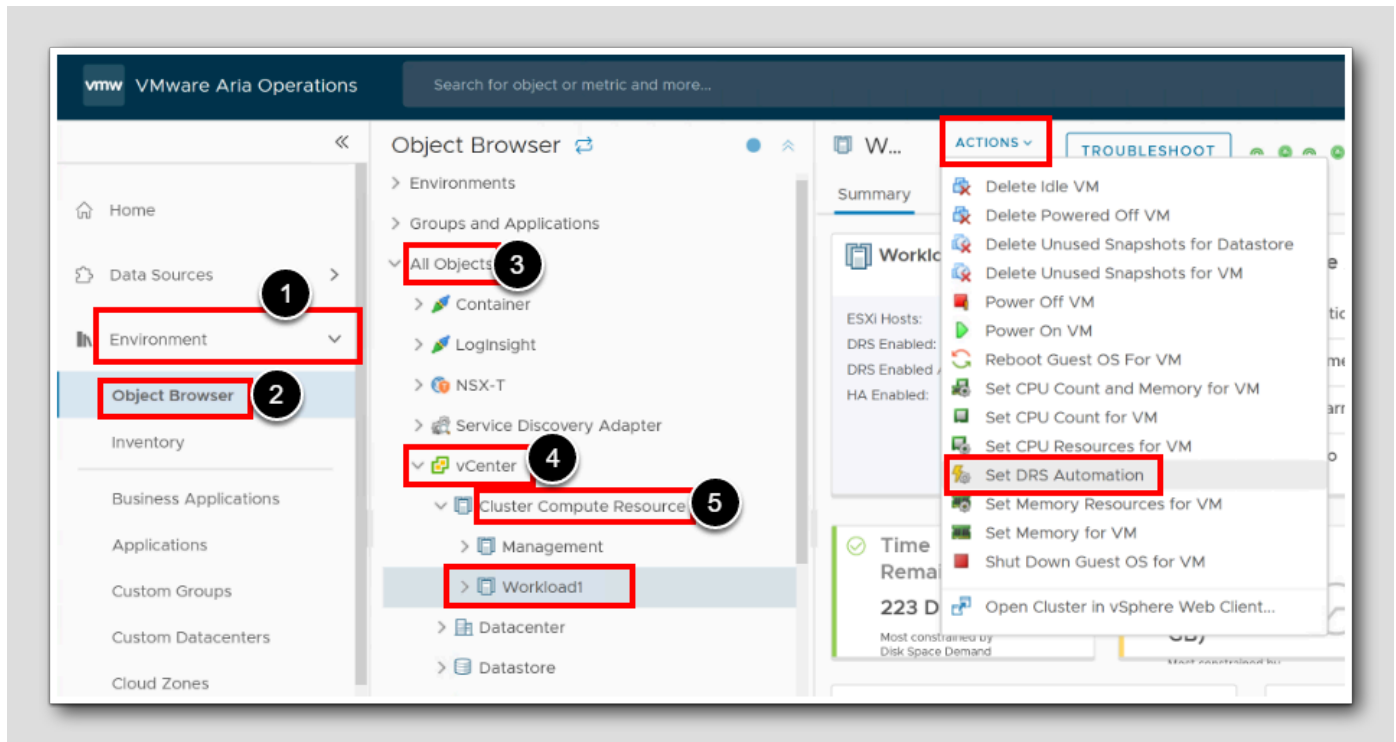
1. Click the Task ID

## Recent Tasks

The screenshot displays the 'Recent Tasks' page in the vSphere interface. At the top, there is a navigation breadcrumb 'Administration / Recent Tasks' and a search bar. Below this is a table of tasks. The first task, 'Set CPU Count and Memory for VM', is selected and highlighted with a red box and a circled '1'. The table columns include Task, Status, Started, Comp..., Auto..., Object Name, Objec..., Event..., Sourc..., Subm..., and Task ID. Below the table, the 'Details of Task Selected' pane shows the 'Associated Objects' table with one entry: 'windows-00...' of type 'Virtual Machine' with a status of 'Completed', highlighted with a red box and a circled '3'. To the right, the 'Messages' section shows a list of log entries with columns for Severity, Time, and Message. A red box and a circled '2' with an arrow point to the scroll bar on the right side of the messages list.

1. Unless it's already selected, from the top of the list, select our windows server Windows-00xx and our task Set CPU and Memory for VM
2. To see all the messages Scroll down
3. Observe that the messages shows the wanted results and that the status is Completed

## Cluster action



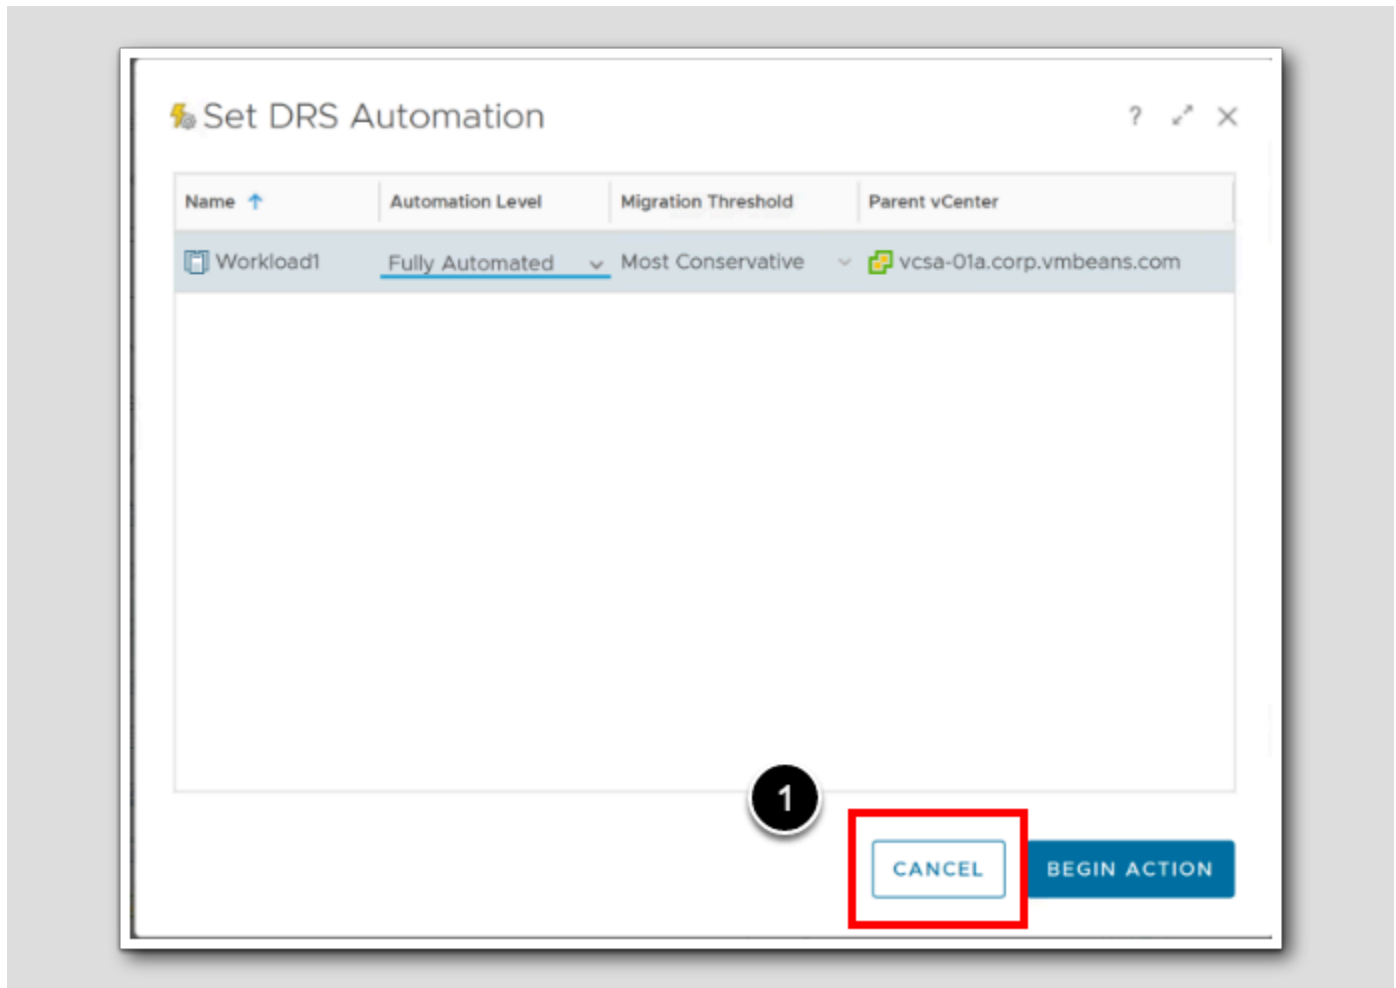
Let us investigate what options we have for actions against another object, not just VM's.

1. Select Environment
2. Select Object Browser
3. Select All Objects
4. Select vCenter
5. Select Cluster Compute Resource
6. Select our workload cluster Workload1
7. Pull down the Action Menu
8. Choose Set DRS Automation



## Set DRS Automation

[113]

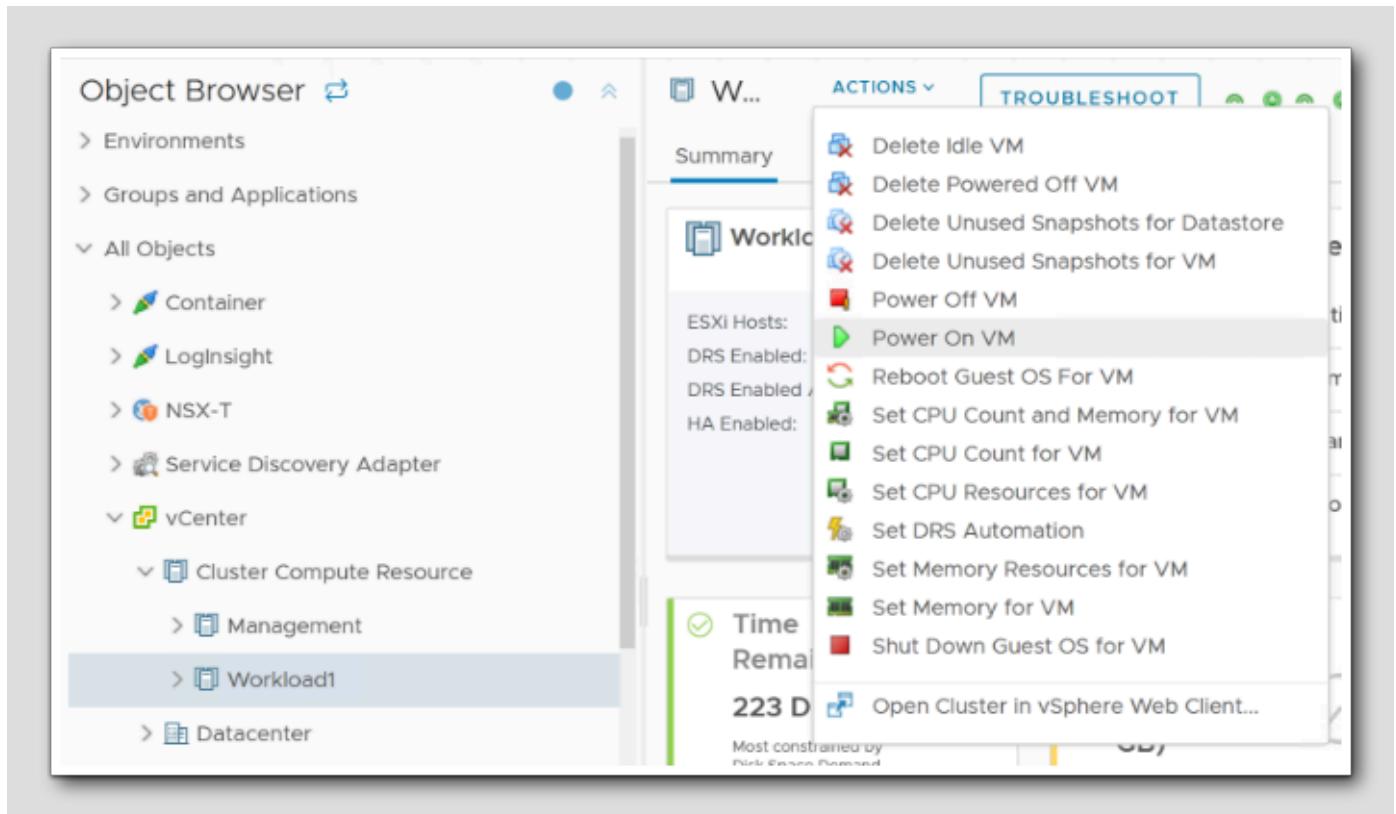


Right now we just wanted to explore that Actions can be run on many other objects

1. Click CANCEL

For more information about objects, have a look at the [Objects and Actions](#) in the [Introduction to Actions](#) in this module

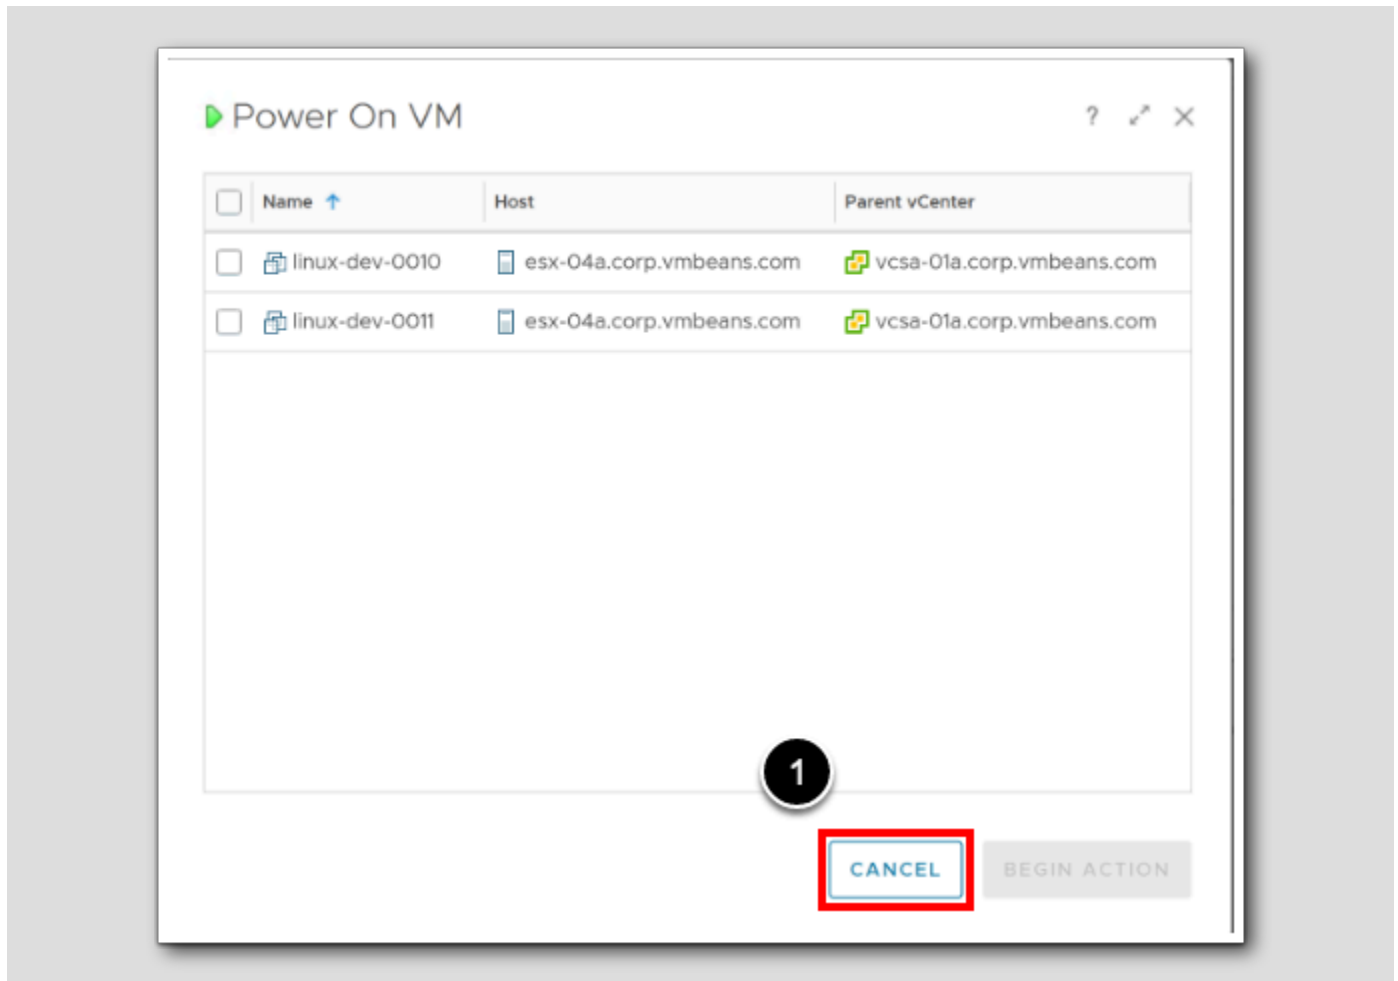
## Power On VMs



Just to see what happens if we choose a VM action while we are located on a Cluster or on a Host

1. Pull down ACTIONS
2. Choose Power On VM

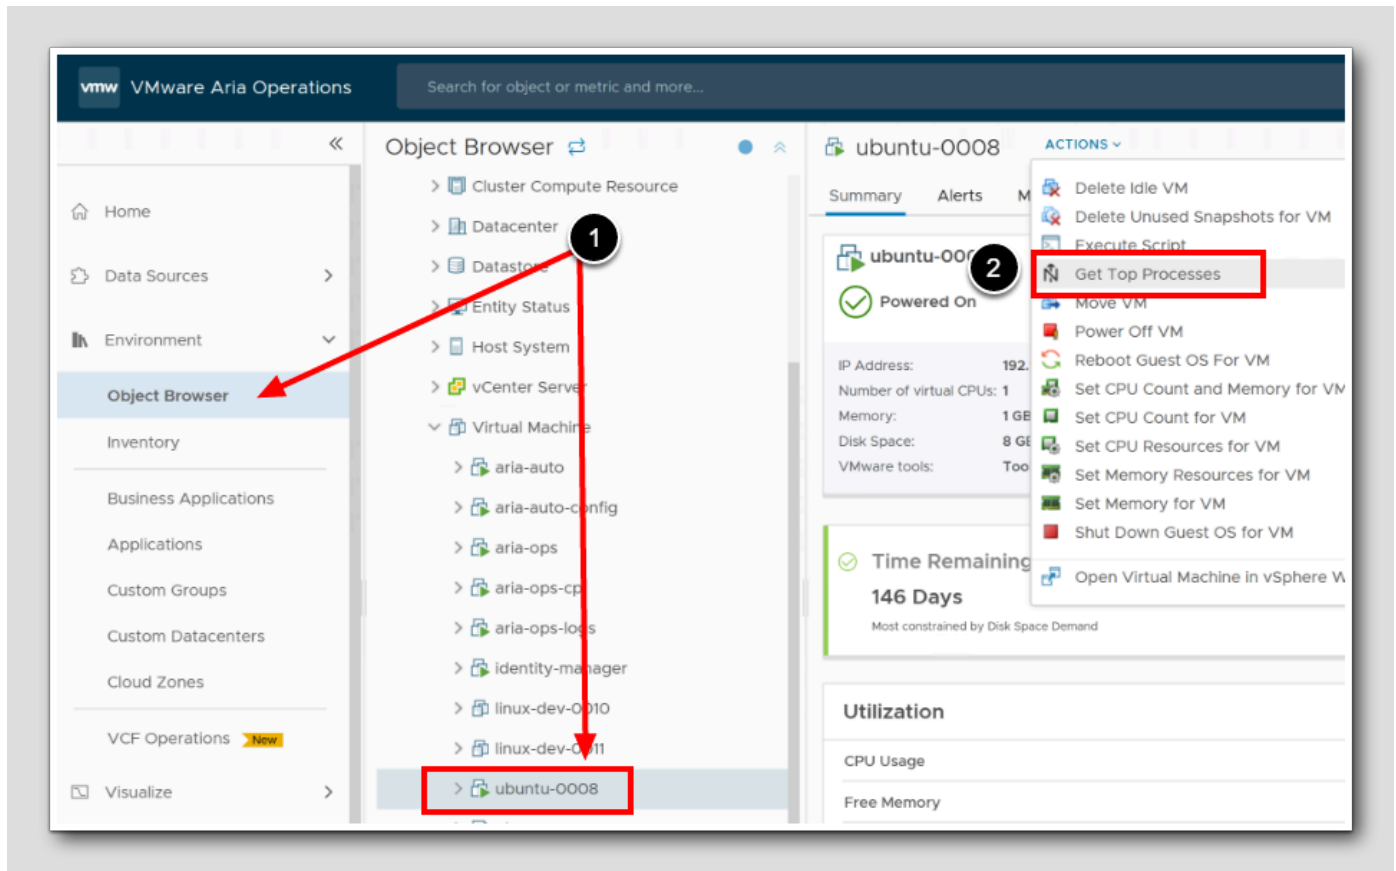
## Power on VMs in the cluster



NOTE: In this scenario, if we choose a cluster or host, the actions we take on the VMs will be applied to all VMs in that group. Specifically, if we select the esx-04 host instead of the Workload1 Cluster, any action performed will be applied to all VMs residing on that host. As a result, the situation would appear the same as if we had selected the Workload1 Cluster, since the two VMs in question are located on the esx-04 host in that cluster.

1. Click **Cancel**

## Get Top Processes

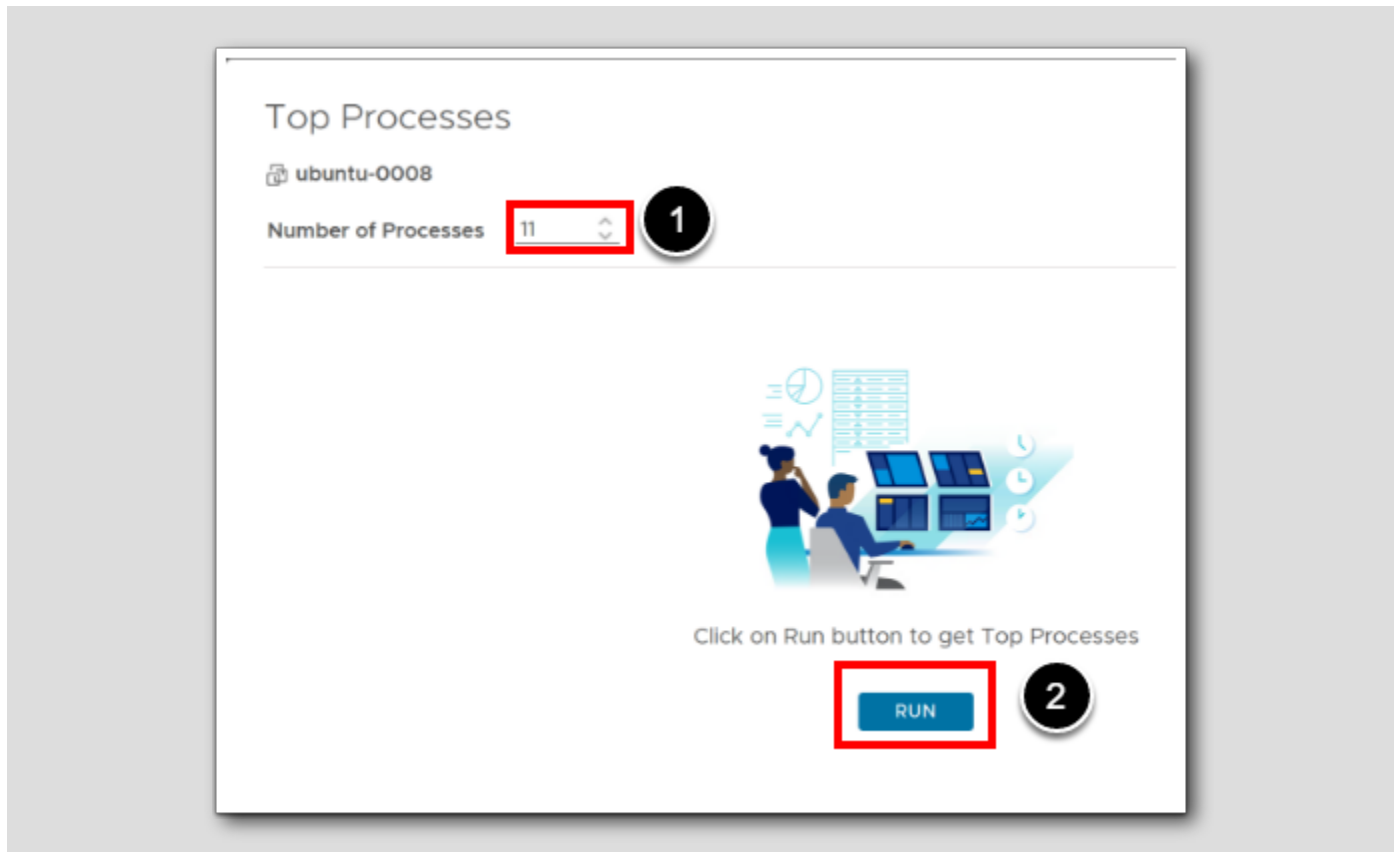


Note: For the next section to work, a VM must be powered on and connected, have the VMware Tools installed and running, and Service discovery must be activated.

1. In the object browser, find and click **ubuntu-0008**
2. Pull down the Actions menu, and choose **Get Top Processes**

## Number of processes

[117]



1. Change the number of processes, type 11
2. Click Run

Interpret 'top' results

[118]

### Top Processes

ubuntu-0008
 1

3 X

Number of Processes

Last time updated: Jul 5, 2023 10:12:46 PM
 REFRESH

```

top --: 05:12:48 up 7 days, 3:55, 0 users, load average: 0.02, 0.02, 0.00
Tasks: 191 total, 1 running, 190 sleeping, 0 stopped, 0 zombie
%Cpu(s): 6.2 us, 6.2 sy, 0.0 ni, 87.5 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
MiB Mem : 976.8 total, 74.6 free, 217.3 used, 684.9 buff/cache
MiB Swap: 0.0 total, 0.0 free, 0.0 used. 609.9 avail Mem
            
```

2

COMMAND	PID	CPU (%) ↓	MEM (%) ○	USER	STATUS
vmtoolsd	681	12.5	0.9	root	S
top	271391	6.2	0.4	holuser	R
rcu_gp	3	0	0	root	I
salt-mi+	783	0	3.5	root	S
kthreadd	2	0	0	root	S
kworker+	6	0	0	root	I
snaped	785	0	4.1	root	S

1. Change the number of processes to 7
2. Click REFRESH

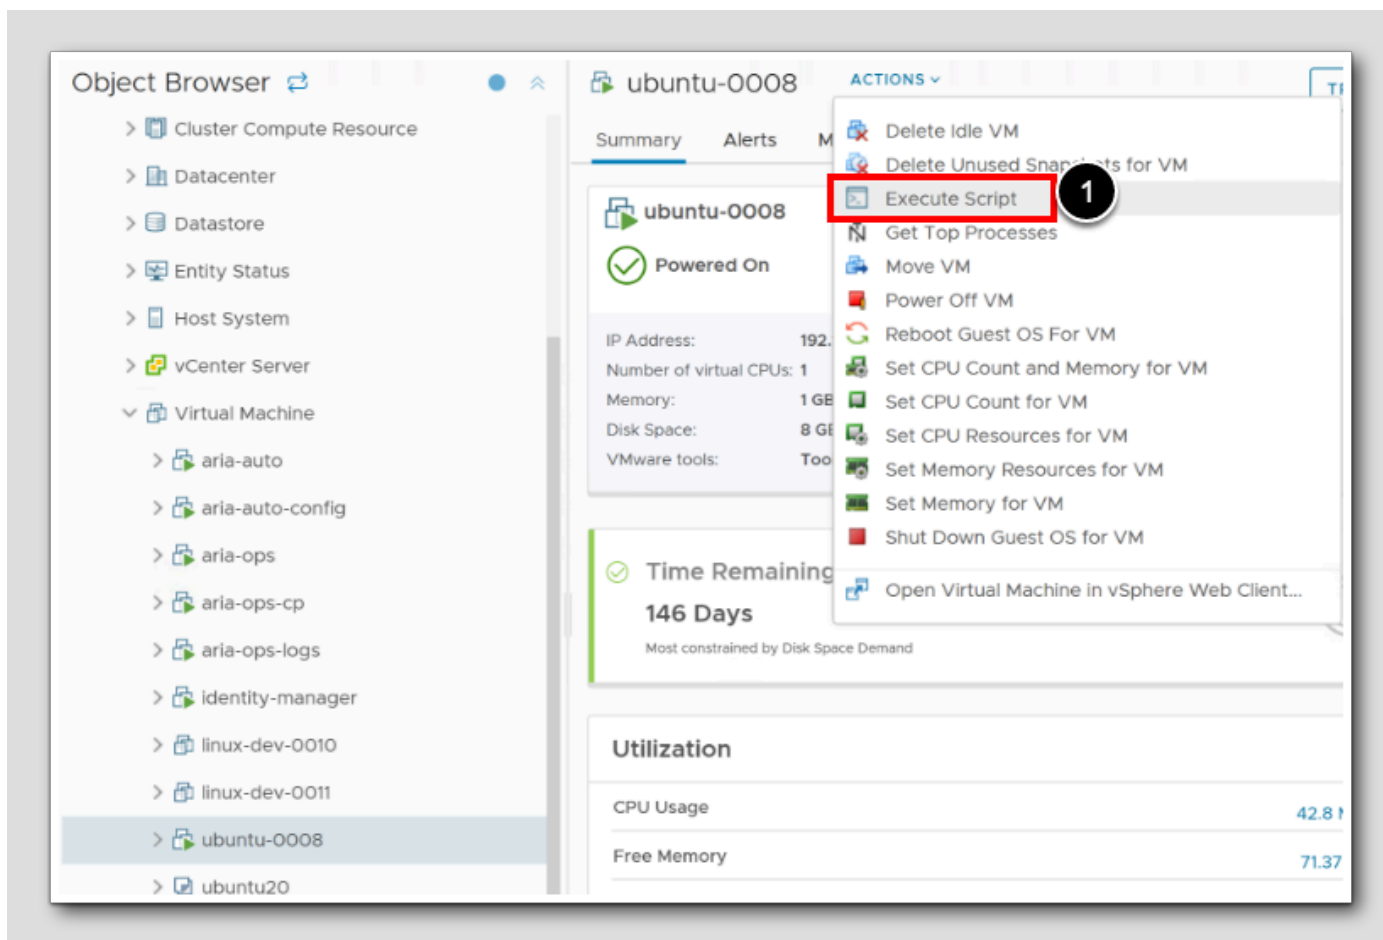
The Get Top Processes action provides the status of our top 7 processes for the selected virtual machine. Observe that funny enough, the two processes causing this idling virtual machine to increase CPU (%) for a short while, is the VMware Tools that's providing the results remotely, and the 'Top' command that produces this list.

Note: Normally we would troubleshoot issues related to the resources that are affecting the applications in the virtual machine. We have the option to view the processes based on CPU and Memory.

3. Click the 'X' to Close

## Execute script

[119]



1. From the actions menu click Execute Script

Running commands

[120]

### Execute Script

ubuntu-0008

⚠ Please exercise caution before executing any critical commands on the virtual machine, system will not check or restrict any kind of script execution. You must provide credentials before you execute any critical commands on the virtual machine.

Username  Password

Command  Upload File

/bin/ps aux --sort=-%cpu | head -n 8

Timeout  Seconds ⓘ

Exit Code: 0

**STDOUT**

USER	PID	%CPU	%MEM	VSZ	RSS	TTY	STAT	START	TIME	COMMAND
root	681	0.2	0.8	312004	8964	?	Ss1	Jun29	25:42	/usr/bin/vmtoolsd
root	995	0.1	5.6	461980	56964	?	Sl	Jun29	11:05	/usr/bin/python3 /usr/bin/salt-minion
root	265461	0.1	0.0	0	0	?	I	01:35	0:20	[kworker/0:0-events]
root	1	0.0	1.1	169396	11584	?	Ss	Jun29	0:47	/sbin/init maybe-ubiquity
root	2	0.0	0.0	0	0	?	S	Jun29	0:00	[kthreadd]
root	3	0.0	0.0	0	0	?	I<	Jun29	0:00	[rcu_gp]
root	4	0.0	0.0	0	0	?	I<	Jun29	0:00	[rcu_par_gp]

**STDERR**



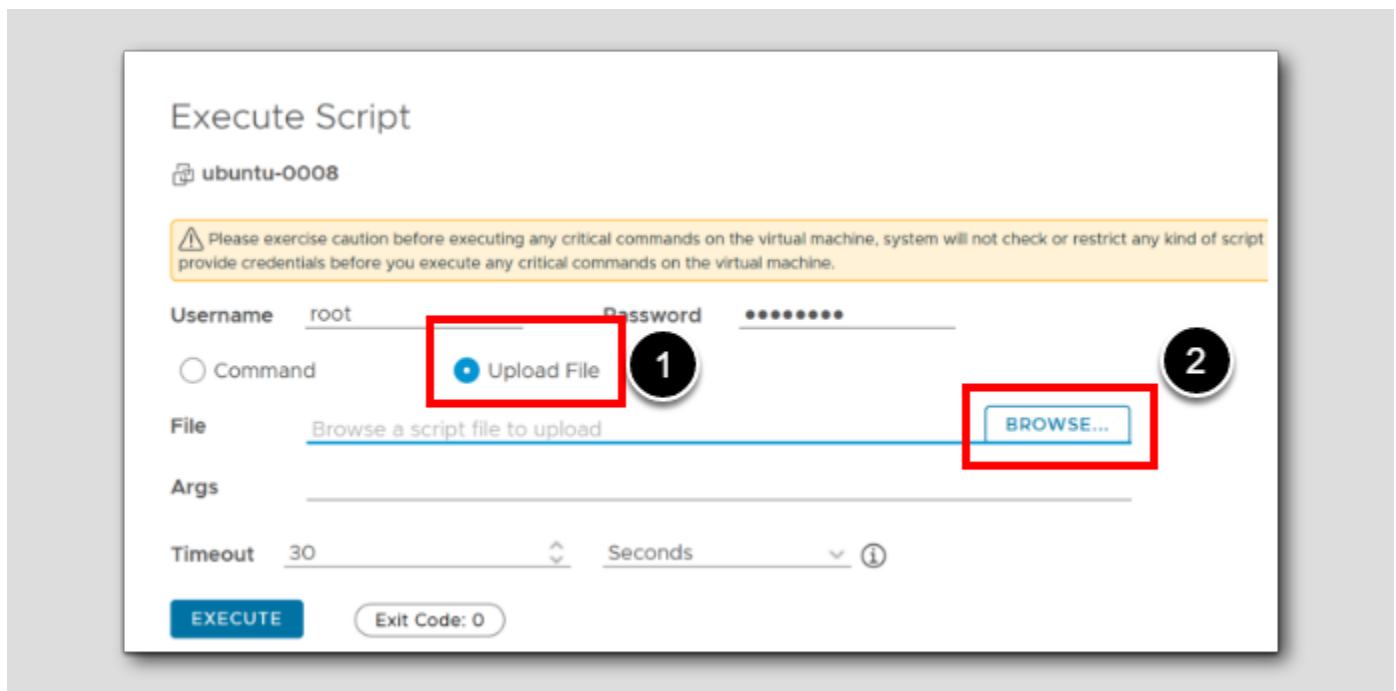
You can run a script by entering it directly or by uploading a script file. We will try to run directly by providing a Linux command. We need to Enter the VM credentials to authenticate.

1. username root
2. password VMware!
3. Type `/bin/ps aux --sort=-%cpu | head -n 8`
4. Click EXECUTE

Not unlike the 'top' command this will give us the top 7 processes sorted on CPU%.

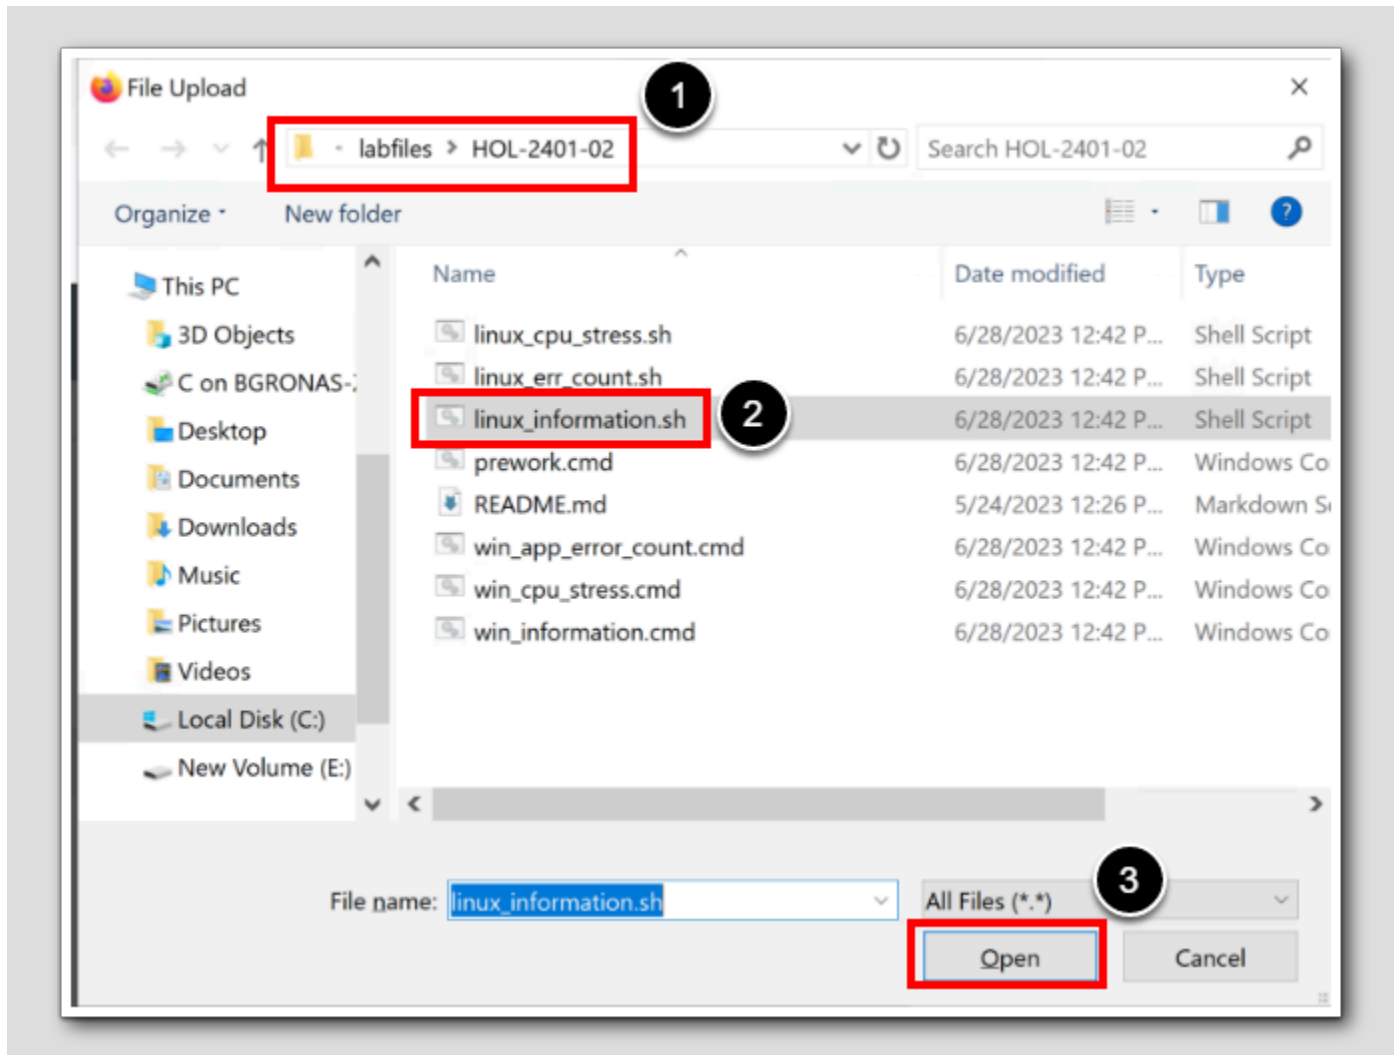
### Browse for a command file

[12]



1. Leave the credentials as is, but click Upload File
2. Click BROWSE

## Select Command file



1. Browse to C:\hol-2401-lab-files\labfiles\HOL-2401-02
2. Select linux\_information.sh
3. Click Open

## Execute

**Execute Script**

ubuntu-0008

⚠ Please exercise caution before executing any critical commands on the virtual machine. system will not check or restrict any ki provide credentials before you execute any critical commands on the virtual machine.

Username  Password

Command  Upload File

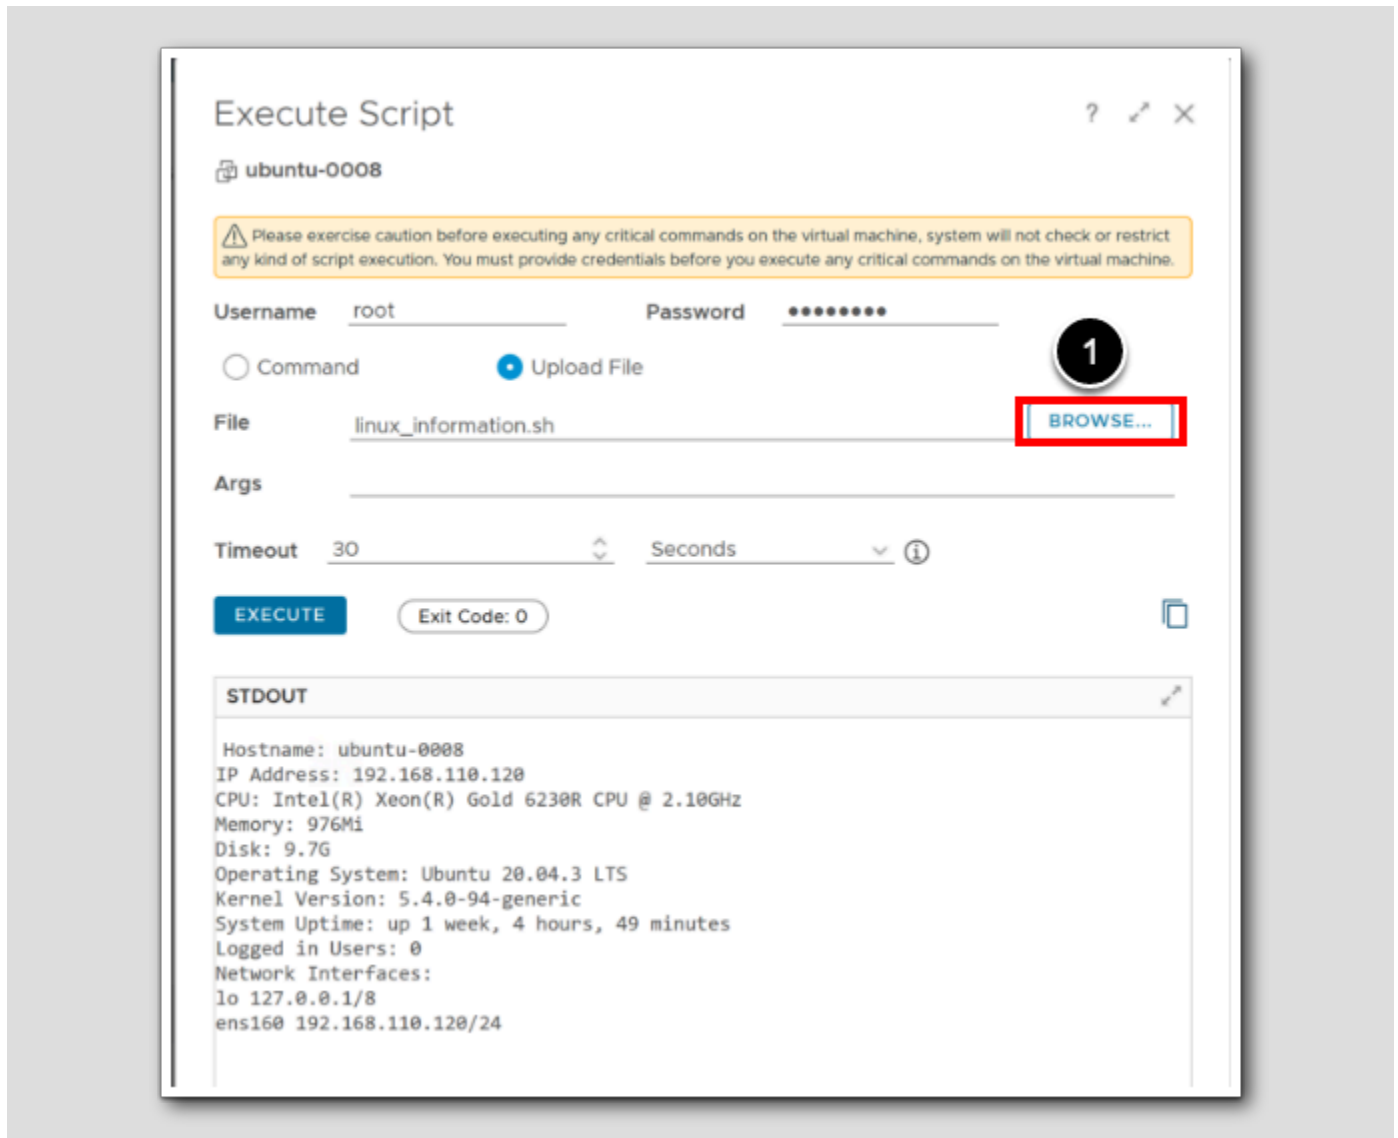
File

Args

Timeout

1. Leave the Args blank, just click EXECUTE

## Useful information



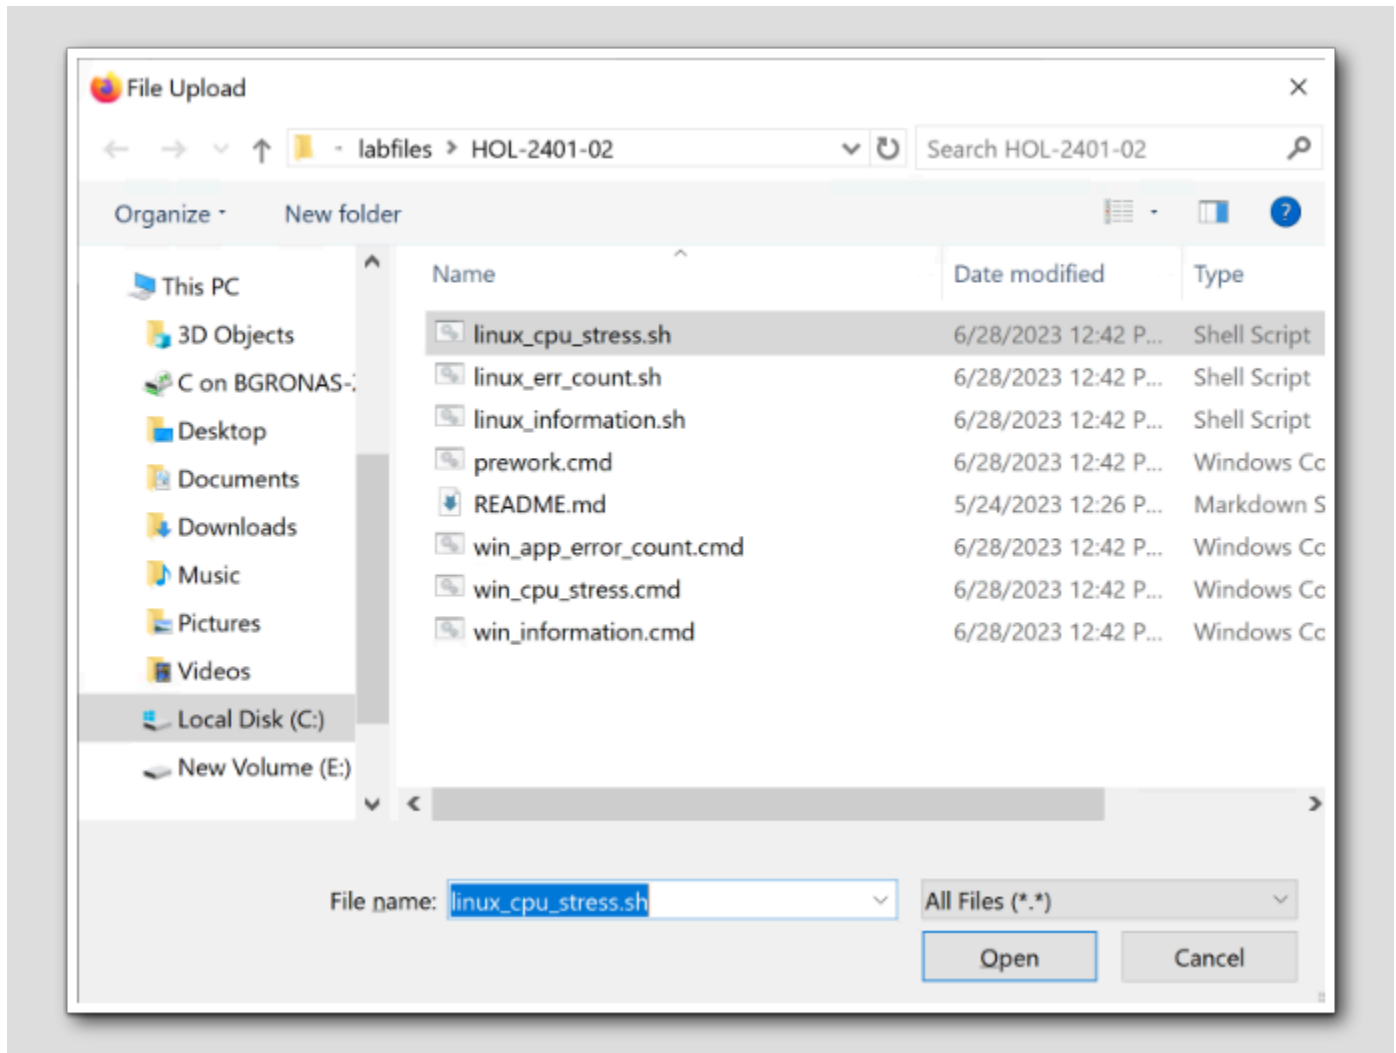
The script we provided returns useful information about the VM.

1. Let's choose another cool script-file we have made ready, click **Browse**

We have listed the Information script; *linux\_information.sh* here for your reference:

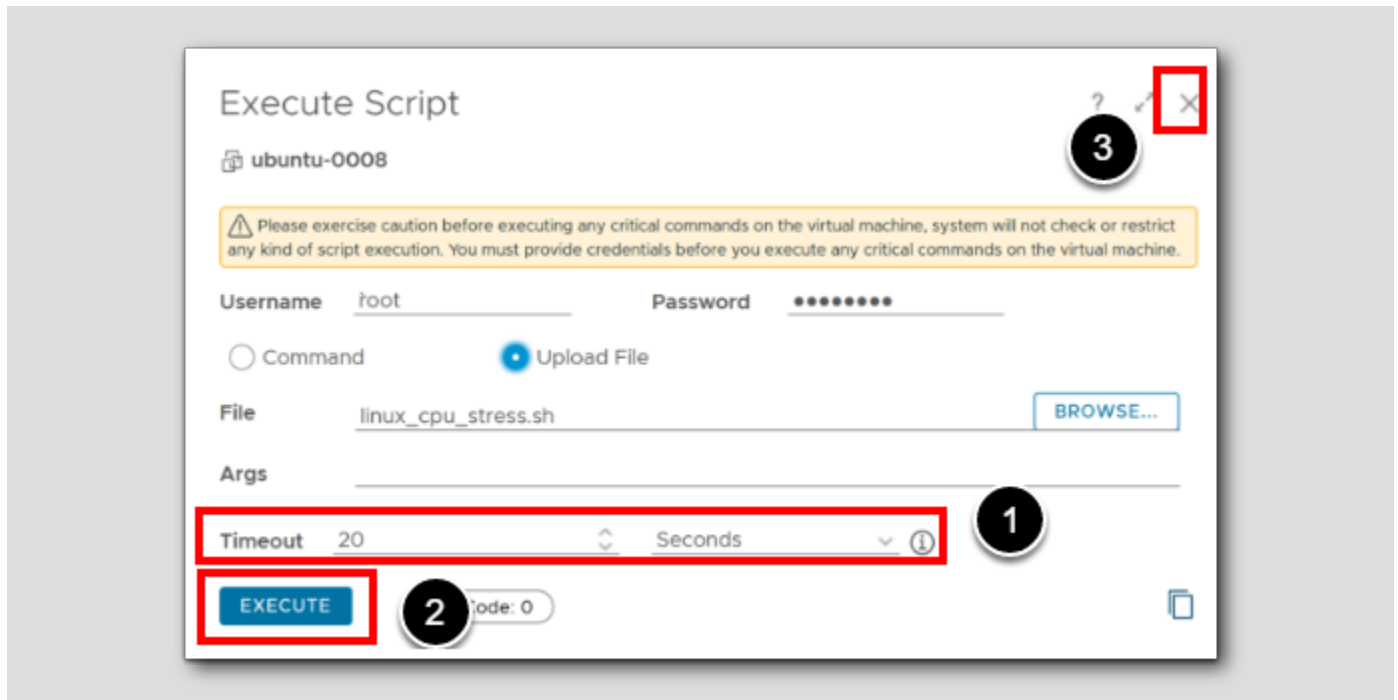
```
#!/usr/bin/bash
# Hostname
echo "Hostname: $(hostname)"
# IP Address
ip_address=$(hostname -I)
echo "IP Address: $ip_address"
# CPU
cpu_info=$(lscpu | grep "Model name" | awk -F ':' '{print $2}' | sed 's/^[ \t]*//')
echo "CPU: $cpu_info"
# Memory
memory_info=$(free -h | awk '/Mem:/{print $2}')
echo "Memory: $memory_info"
# Disk
disk_info=$(df -h --total | awk '/total/{print $2}')
echo "Disk: $disk_info"
# Operating System
os_info=$(lsb_release -d | awk -F ':' '{print $2}' | sed 's/^[ \t]*//')
echo "Operating System: $os_info"
# Kernel Version
kernel_info=$(uname -r)
echo "Kernel Version: $kernel_info"
# System Uptime
uptime_info=$(uptime -p)
echo "System Uptime: $uptime_info"
# Logged in Users
users_info=$(who | awk '{print $1}' | sort | uniq | wc -l)
echo "Logged in Users: $users_info"
# Network Interfaces
interfaces_info=$(ip -o -4 addr show | awk '{print $2, $4}')
echo -e "Network Interfaces:\n$interfaces_info"
```

## Stress script



1. Browse to C:\hol-2401-lab-files\labfiles\HOL-2401-02
2. Select linux\_cpu\_stress.sh
3. Click Open

## Run the stress script

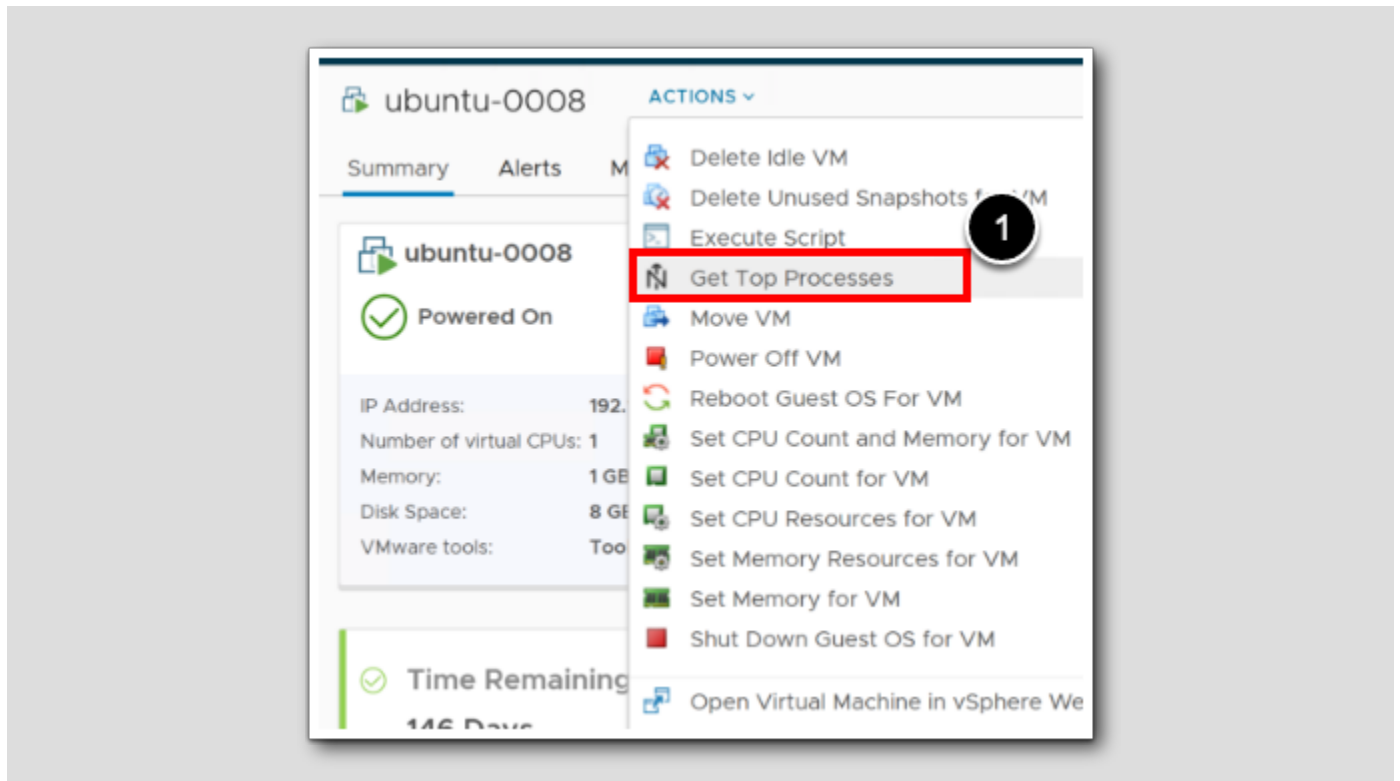


1. Set the timeout to 20 seconds
2. Click EXECUTE

After 20 seconds you will receive a "Task Execution failed. See details in Recent Tasks" - We will just accept the failure. The script will actually continue processes in the background.

3. Click 'X' to exit

## Check top processes



1. From the Actions menu, select Get Top Processes



### Review Top Processes

Top Processes

ubuntu-0008

Number of Processes  Last time updated: Jul 5, 2023 11:45:33 PM [REFRESH](#)

```
top -: 06:45:32 up 7 days, 5:27, 0 users, load average: 16.82, 14.15, 12.91
Tasks: 235 total, 23 running, 212 sleeping, 0 stopped, 0 zombie
%Cpu(s): 86.0 us, 14.0 sy, 0.0 ni, 0.0 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
MiB Mem : 976.8 total, 119.3 free, 246.3 used, 611.2 buff/cache
MiB Swap: 0.0 total, 0.0 free, 0.0 used. 385.9 avail Mem
```

COMMAND	PID	CPU (%)	MEM (%)	USER	STATUS
script_+	274248	4.8	0	root	R
script_+	274244	4.8	0	root	R
script_+	274247	4.8	0	root	R
dd	274255	4.8	0.3	root	R
script_+	274233	2.4	0	root	R
script_+	274235	2.4	0	root	R
script_+	274232	2.4	0	root	R
script_+	274234	2.4	0	root	R
script_+	274231	2.4	0	root	R
script_+	274230	2.4	0	root	R

As we can see from the image, the Linux server in the Virtual Machine is busy with our script.

1. After we are done with our review, Click 'X' to close

The `linux_cpu_stress.sh` script creates 20 infinite loops for CPU stress and after 10 minutes it kills all the background processes and the stress test. Listed here for your reference:

```
#!/usr/bin/bash
# Function to create CPU load
generate_cpu_load() {
  # infinite loop
  while :
  do
    :
  done &
}
# Function to create memory load
generate_mem_load() {
  # Allocate 256MB memory in /dev/shm
  dd if=/dev/zero of=/dev/shm/stress_test bs=1M count=256 &
}
# Create CPU load
for i in `seq 1 20`;
do
  generate_cpu_load
done
# Create Memory load
for i in `seq 1 20`;
do
  generate_mem_load
done
# Sleep for 10 minutes
sleep 600
# Kill all background jobs when we're done
kill $(jobs -p)
```

## Summary

[129]

With Aria Operations we can run built-in Out of the box actions for such as Power On/Power Off, and resizing of VMs. Actions can be performed on different type of object types. We went through how to do it on a Cluster, Host, and VMs.

We carried out actions remotely to VMs without opening any Terminal Window or remote desktop on that VM. we checked the condition of a VM, and Executed scripts that we had stored locally by copying and running it on the VM.

## Conclusion

[130]

In this module, we examined Aria Operations ability to monitor processes, services, and applications, leveraging the Telegraf agent.

Native application monitoring is facilitated via the Aria Operations Telegraf Agent.

Discover Services employs the VMware Tools agent to monitor processes and services, while Monitor Applications utilizes an open-source Telegraf agent for metric collection from managed VMs.

Discover Services offers more configuration information, whereas Monitor Applications provides a wider range of performance metrics.

## You've finished the module

[131]

Congratulations on completing this lab module.

If you are looking for additional information, please visit the [Aria Operations Documentation](#)

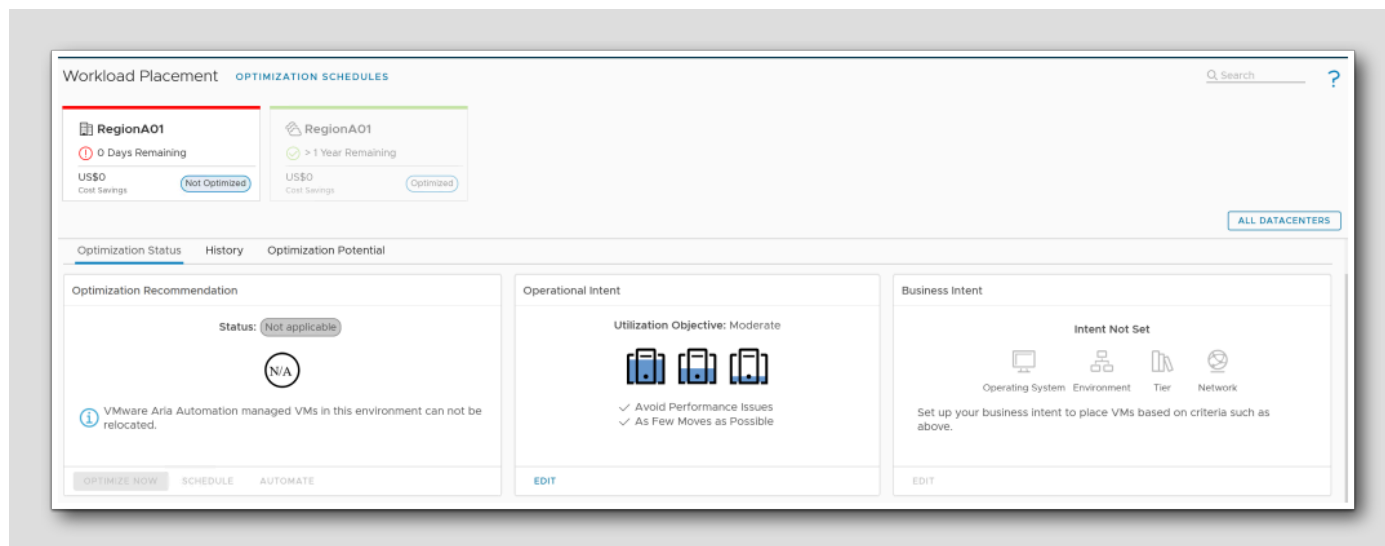
From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 5 - Workload Placement – Running Host Based Optimization (35 minutes) Intermediate

### Introduction

[133]



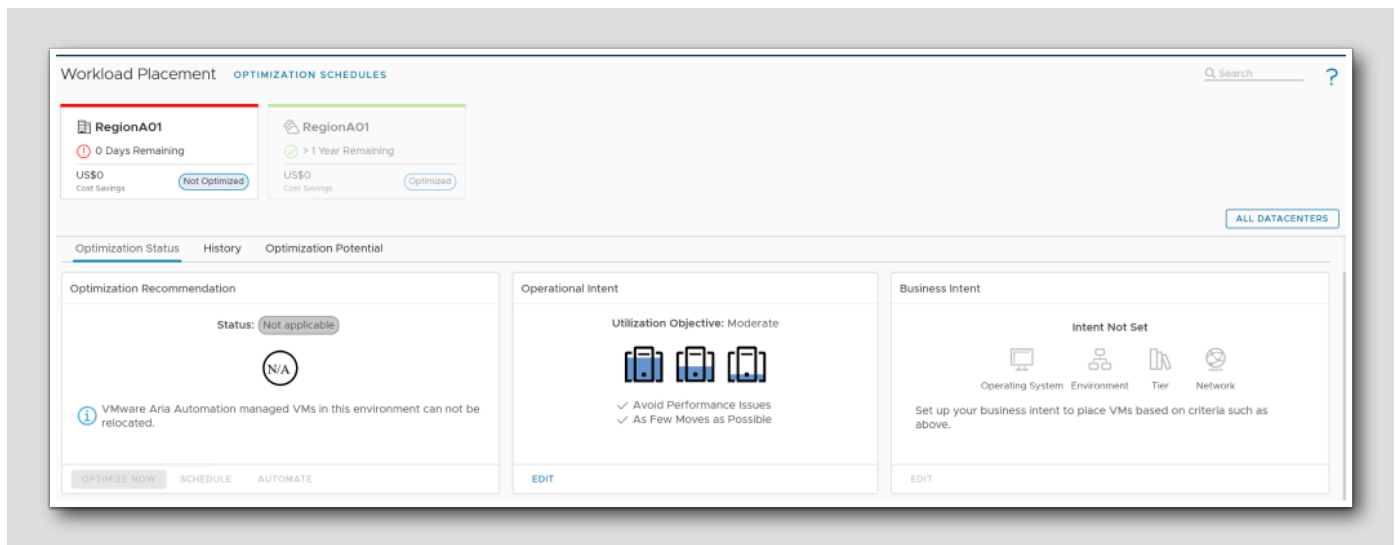
### Workload placement with Business intent

**Business Intent** is a powerful feature that allows you to align your workload placement and balancing decisions with your organization's specific business needs. It is essentially a set of user-defined rules or policies that guide the automation and decision-making processes within the software.

Business Intent can help with;

- Compliance and Licensing
- Tag-Based Placement
- Separation of Workloads
- Optimizing Resource Utilization

The effectiveness of Business Intent in Aria Operations is dependent on how accurately the user-defined rules reflect the organization's operational needs and goals. Therefore, it's essential to plan and consider your specific business requirements before defining these rules.



## Log in to Aria Operations

[134]

We will log in to a live instance of Aria Operations running in this lab.

## Open the Firefox Browser from the Windows Task Bar

[135]




If the browser is not already open, launch **Firefox**.

1. Click the **Firefox** icon in the Windows Quick Launch Task Bar at the bottom of the screen.

## Navigate to Aria Operations


[136]



1. Click the Aria Operations bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[137]




Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

1. Click Sign in

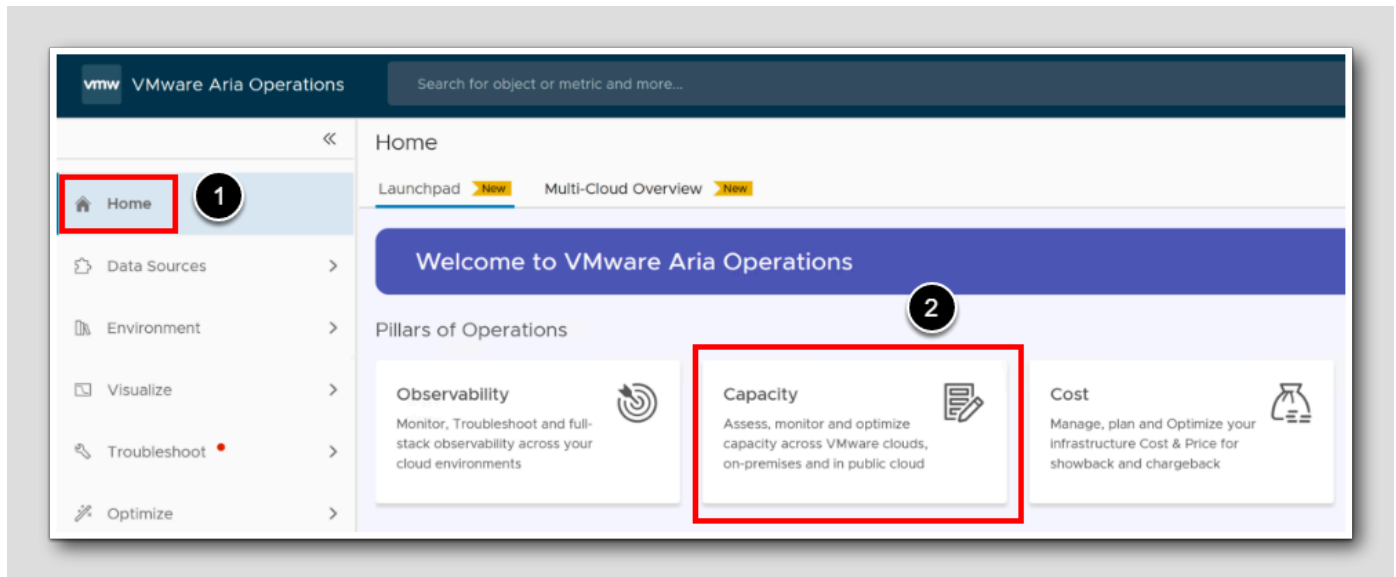


## Business Intent

[139]

## Locating Workload Placement

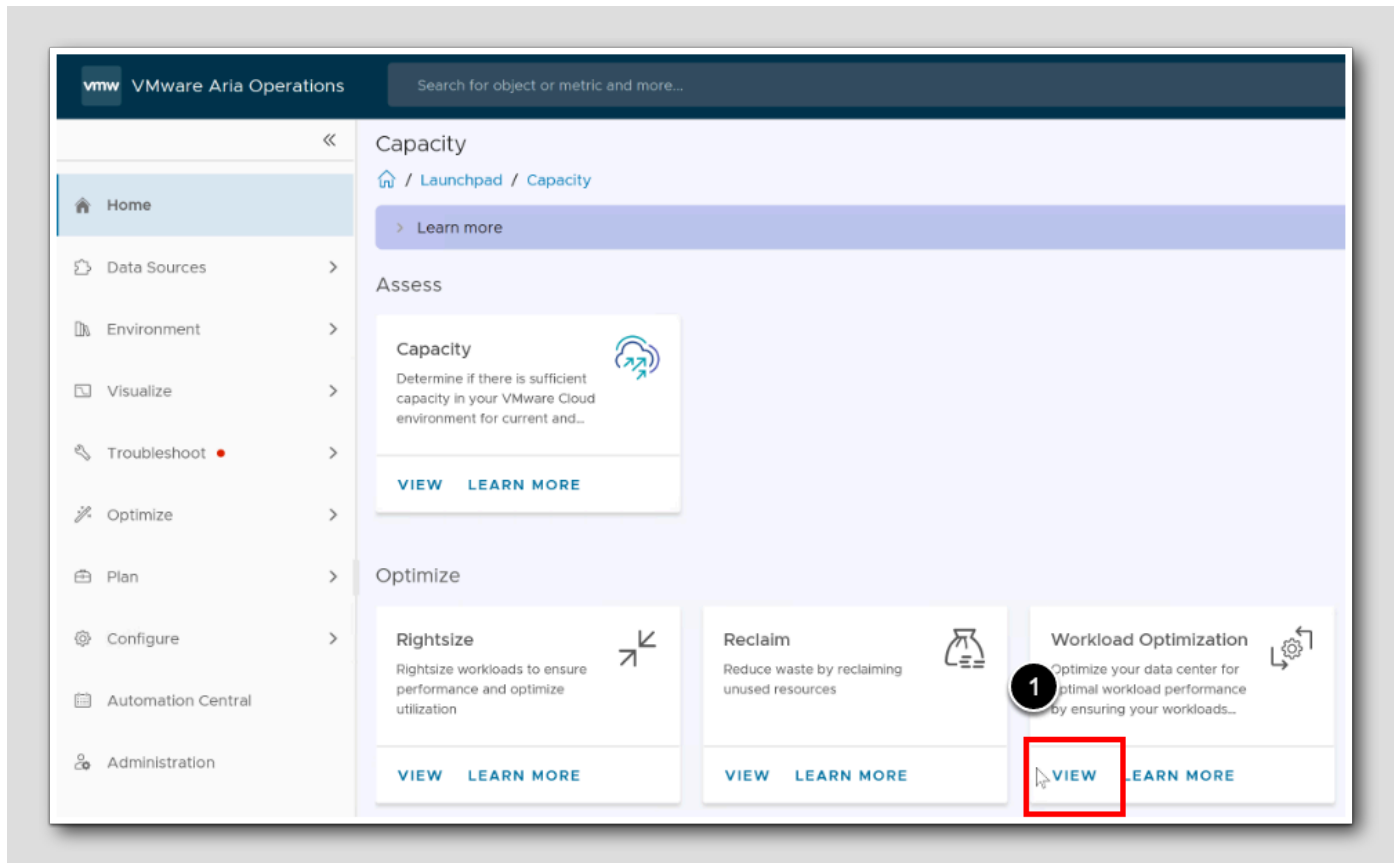
[140]



There are several ways to take you to the Workload placement and the Business Intent. One way is to Select Workload Placement under Optimize in the left menu, or we can go via the Capacity:

1. Click on Home
2. Click on Capacity

## From Capacity to Workload Placement



Workload Optimization is a part of the process of optimizing datacenter for optimal workload performance. It works closely with DRS to ensure applications get required resources.

VMware Aria Operations monitors the environment and, when the datacenter deviates from its desired state, it will recommend the optimization actions to move it back to a desired state.

1. On the Workload Optimization, Click VIEW

## Workload Placement Page

The screenshot displays the VMware Workload Placement interface. At the top, there are two cards for 'RegionA01' showing optimization status: 'Not Optimized' (0 Days Remaining) and 'Optimized' (> 1 Year Remaining). Below these are tabs for 'Optimization Status', 'History', and 'Optimization Potential'. The main content area is divided into three cards: 'Optimization Recommendation' (Status: Not applicable), 'Operational Intent' (Utilization Objective: Moderate), and 'Business Intent' (Intent Not Set). A section titled 'Are your clusters meeting your utilization objective?' contains a table with two clusters: 'Workload1' and 'Management'. The 'Management' cluster is highlighted with a red box, and the 'SET DRS AUTOMATION' link is circled with a red box and a '2'. A '1' is placed over the 'Management' cluster name in the table.

Name	Model	CPU Workload %	Memory Workload %	DRS Settings	Migration Threshold	Violated Tags	VM Name
Workload1	Demand Allocation	30%	83%	Deactivated	--	-	
Management	Demand Allocation	33%	78%	Partial Automated	Default	-	

On the Workload Placement Page, we see three different cards:

- The Optimization Recommendation Card
- The Operational Intent Card
- The Business Intent Card

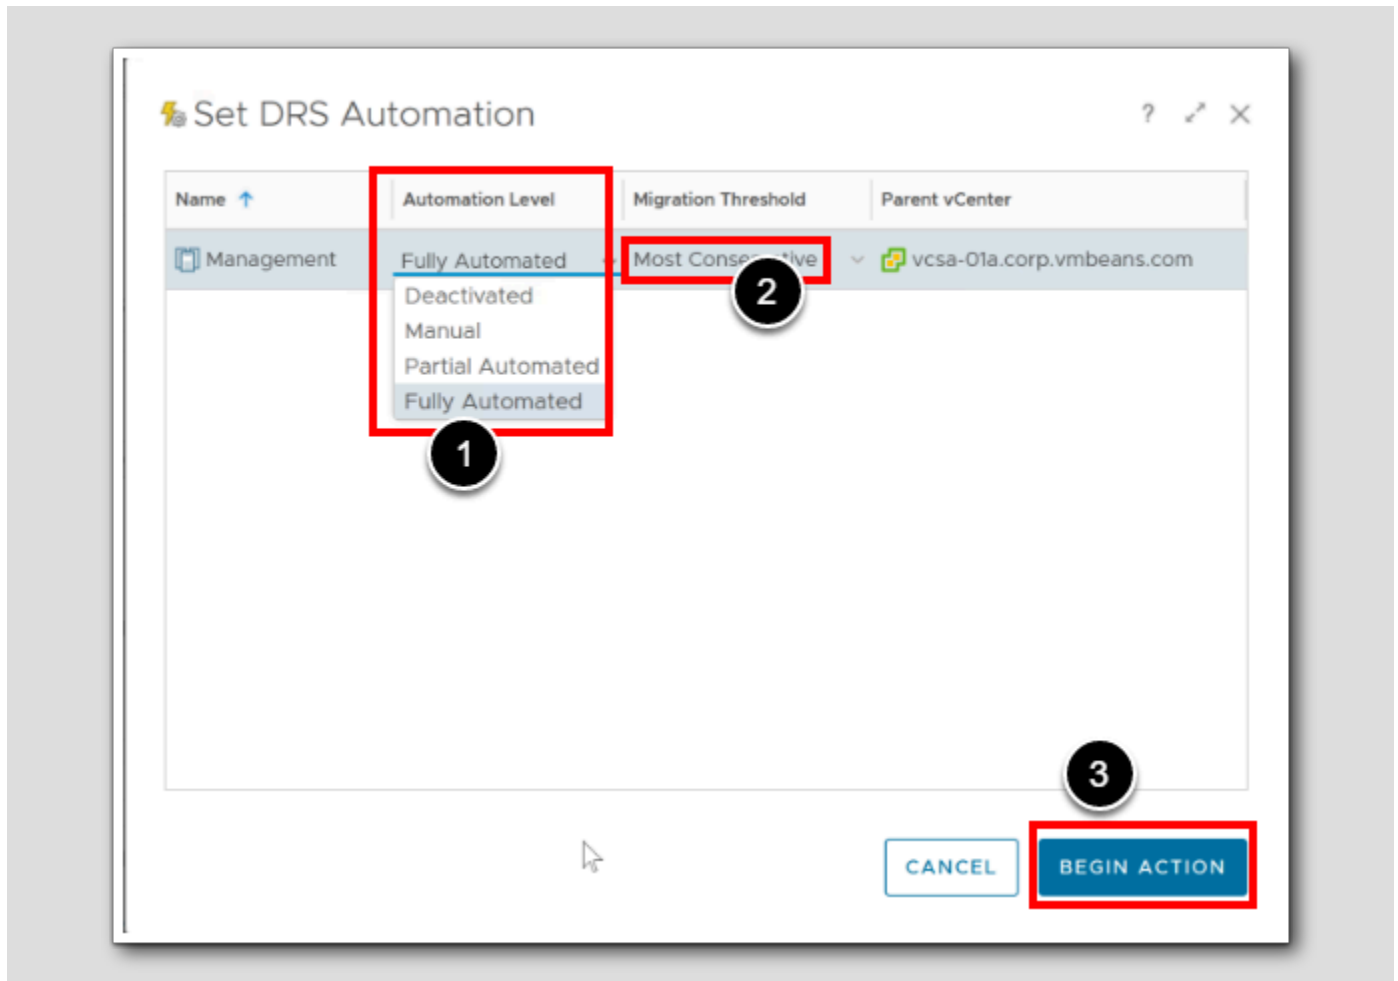
In this session we will concentrate on the Business Intent Card. But before we do that, we will to set the DRS (Dynamic Resource Scheduling) Automation on both of these clusters.

Note that clusters must be **fully automated** in order for workload optimization alerts to run actions set in the policies.

1. Select the Management Cluster, and click the link SET DRS AUTOMATION

## DRS Automation, first cluster

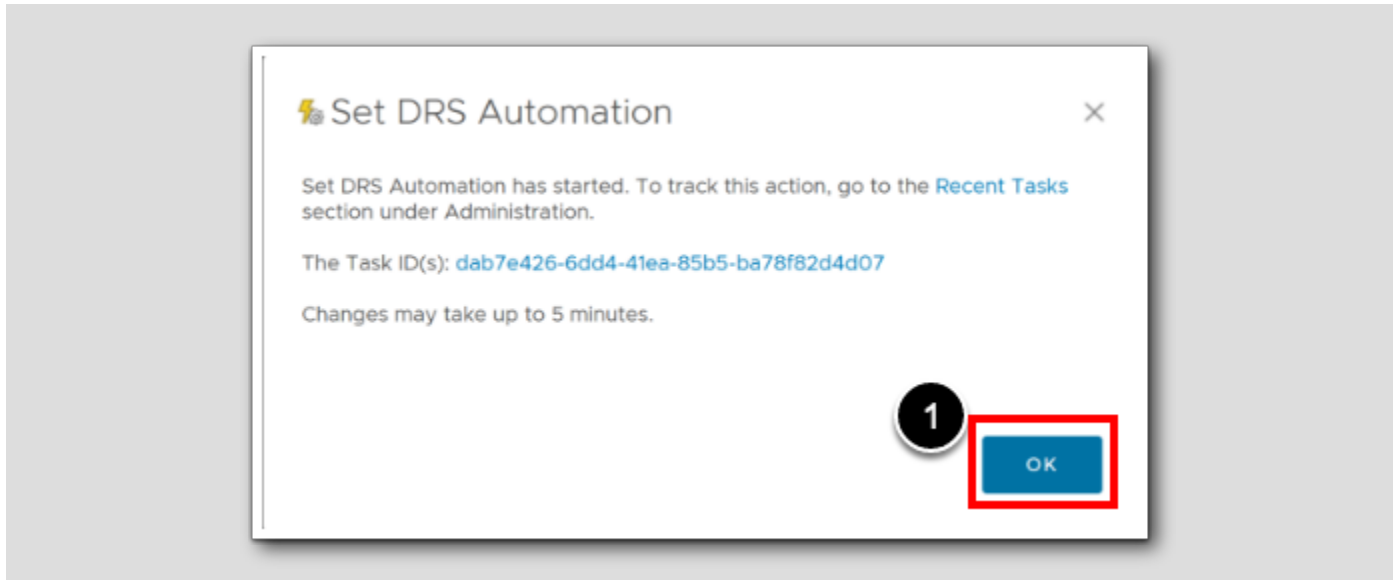
[143]



1. Set the Automation level to **Fully Automated**
2. Set the Migration Threshold to **Most Conservative**
3. Click **BEGIN ACTION**

Confirming The action

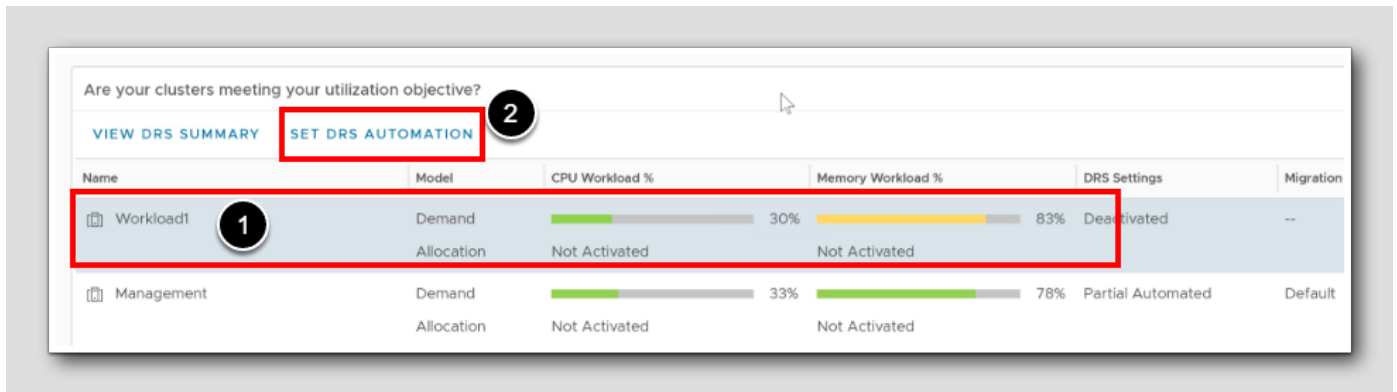
[144]



1. To confirm this action, Click OK

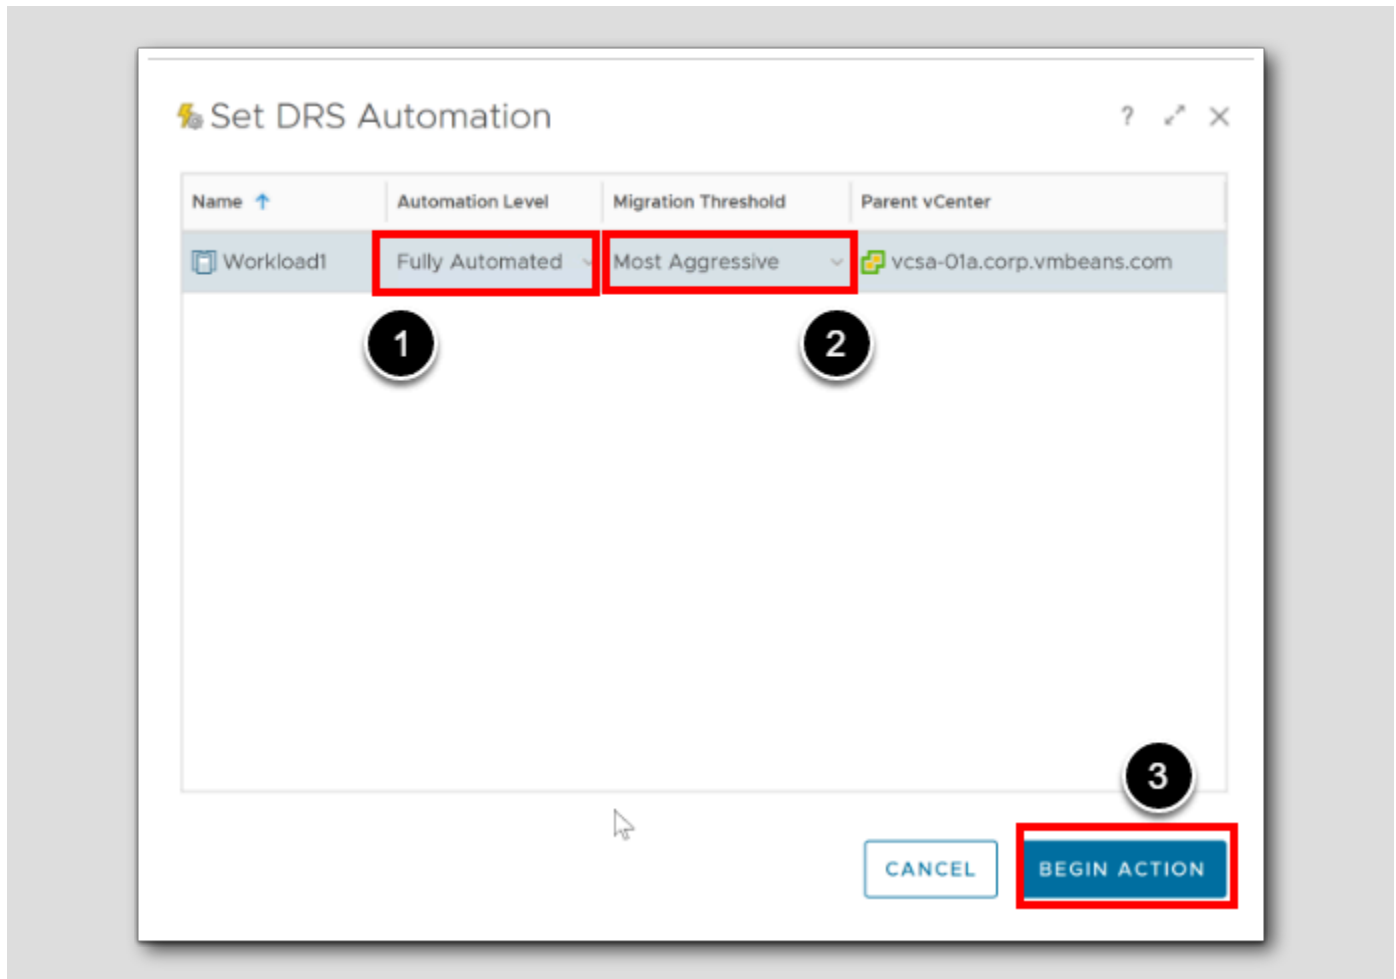
DRS Automation, Second Cluster

[145]



1. Select the Cluster named Workload1
2. Click SET DRS AUTOMATION

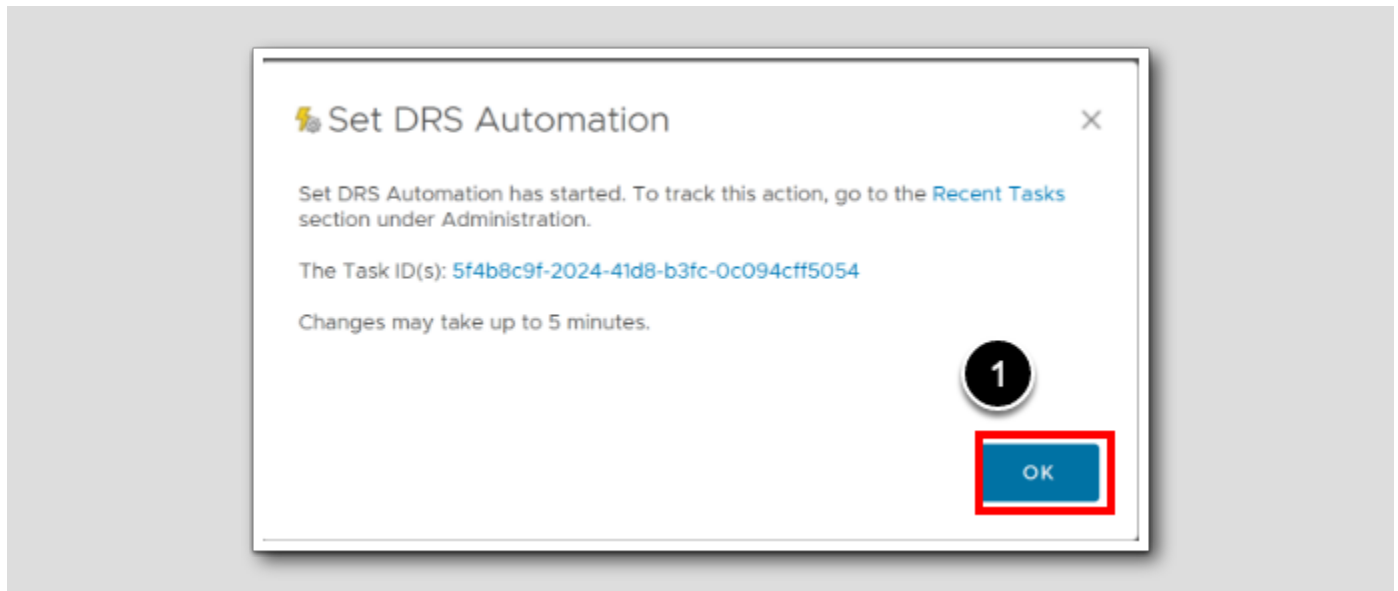
## Setting the DRS Automation



1. Set the Automation level to **Fully Automated**
2. Set the Migration Threshold to **Most Aggressive**
3. Click **BEGIN ACTION**

## Confirm the action

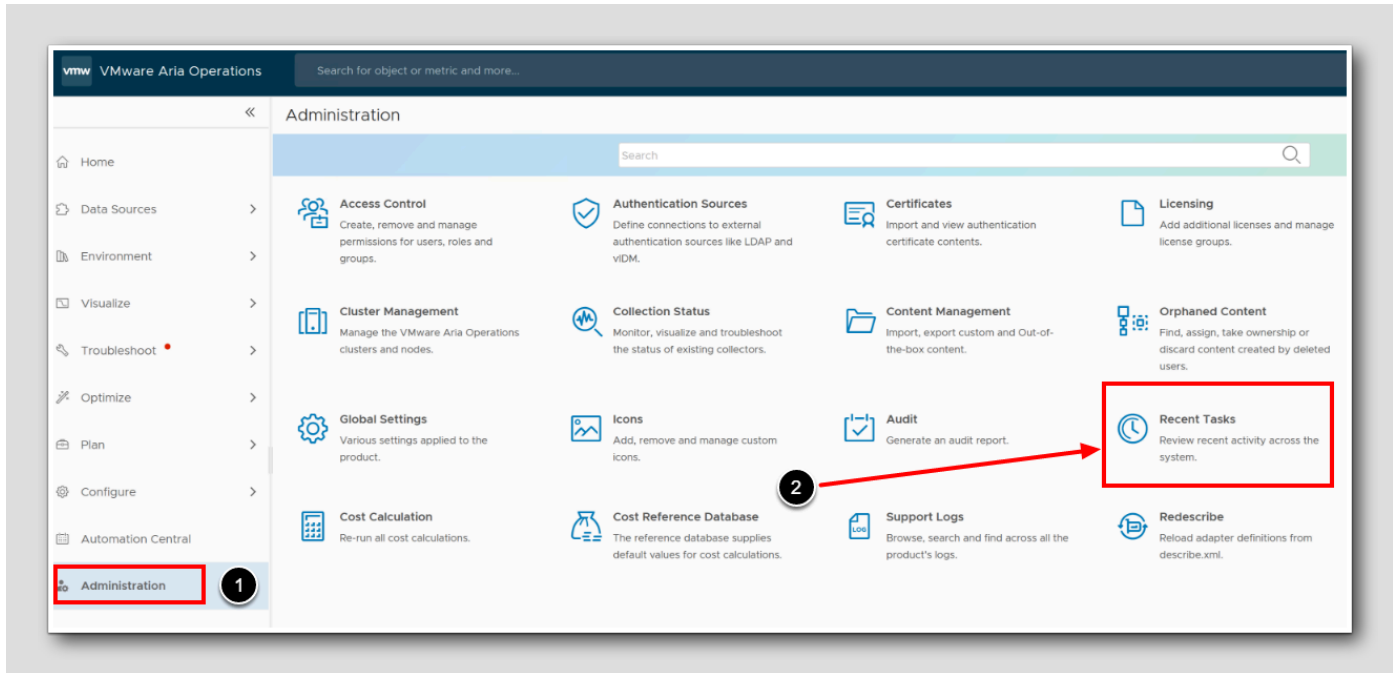
[147]



1. In the Dialog, Click OK

There has been a reason why we haven't clicked on Recent Tasks or the Task ID in this dialog box. Next we will confirm that the DRS settings have been set.

## Recent tasks



1. In the left Pane, Click Administration
2. Click Recent Tasks



## DRS Automation Status Completed

The screenshot displays the 'Recent Tasks' page in vSphere. The table below shows the tasks performed:

Task	Status	Started Time	Completed Time	Automated	Object Name
Set DRS Automation	Completed	12:54 PM	12:55 PM	No	Workload1
Set DRS Automation	Completed	12:47 PM	12:47 PM	No	Management
Set DRS Automation	Completed	12:23 PM	12:23 PM	No	Management

The 'Details of Task Selected' section for the Management cluster shows the following information:

Object Name	Object Type	Status
Management	Cluster Compute Resource	Completed

The 'Messages' section shows the following log entries:

Severity	Time	Message
Informa...	2023-06-26 12:47:29...	Task Id: dab7e426-6dd4-41ea-i
Informa...	2023-06-26 12:47:35...	Executing 'ModifyDRSConfig'
Informa...	2023-06-26 12:47:35...	Parameters: mOR: domain-c12C
Informa...	2023-06-26 12:47:35...	Params: MethodParam(targetR
Informa...	2023-06-26 12:47:53...	Return status: COMPLETED

1. Select The Management Cluster
2. Make sure the status is Completed
3. Note: Also make sure the Workload1 has it's DRS Automation setting to Completed

## Getting back to the Workload placement

The screenshot displays the VMware Aria Operations interface for Workload Placement. The left-hand navigation pane shows the 'Optimize' menu item highlighted with a red box and a circled '1', and the 'Workload Placement' menu item highlighted with a red box and a circled '2'. The main content area shows the 'Workload Placement' page for 'RegionA01'. It includes two 'RegionA01' cards at the top, each showing '0 Days Remaining' and 'US\$0 Cost Savings' with a 'Not Optimized' button. Below these is a detailed view for 'RegionA01' with tabs for 'Optimization Status', 'History', and 'Optimization Potential'. The 'Optimization Recommendation' section shows a status of 'Not Optimized' and a diagram of workload movement. The 'Operational Intent' section shows a 'Utilization Objective: Moderate' with three server icons and two checkmarks: 'Avoid Performance Issues' and 'As Few Moves as Possible'. The 'Business Intent' section shows 'Operating Sys' and 'Set up your business intent above.' Below these are buttons for 'OPTIMIZE NOW', 'SCHEDULE', 'AUTOMATE', and 'EDIT'. A section titled 'Are your clusters meeting your utilization objective?' includes a table with DRS settings.

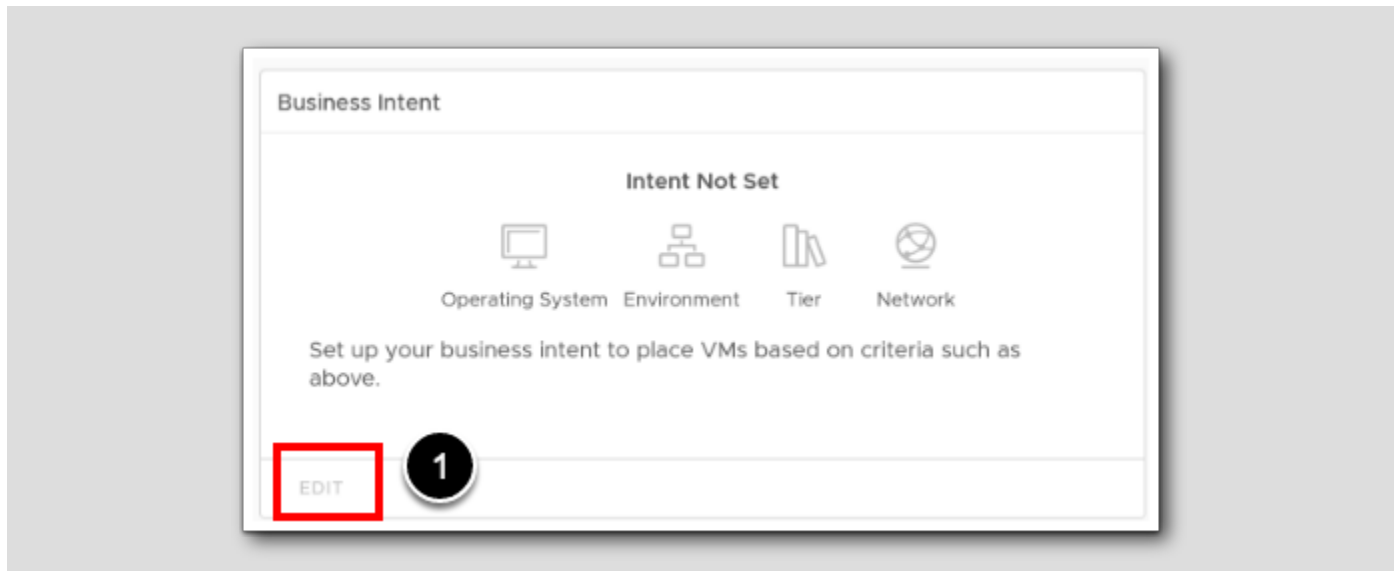
Name	Model	CPU Workload %	Memory Workload %	DRS Settings	Migration Threshold
Workload1	Demand	29%	82%	Fully Automated	Most Aggressive
	Allocation	Not Activated	Not Activated		
Management	Demand	31%	78%	Fully Automated	Most Conservative

Here is the other way of getting back to Workload placement

1. Click Optimize
2. Click Workload Placement

## The 'Grayed out' Business Intent

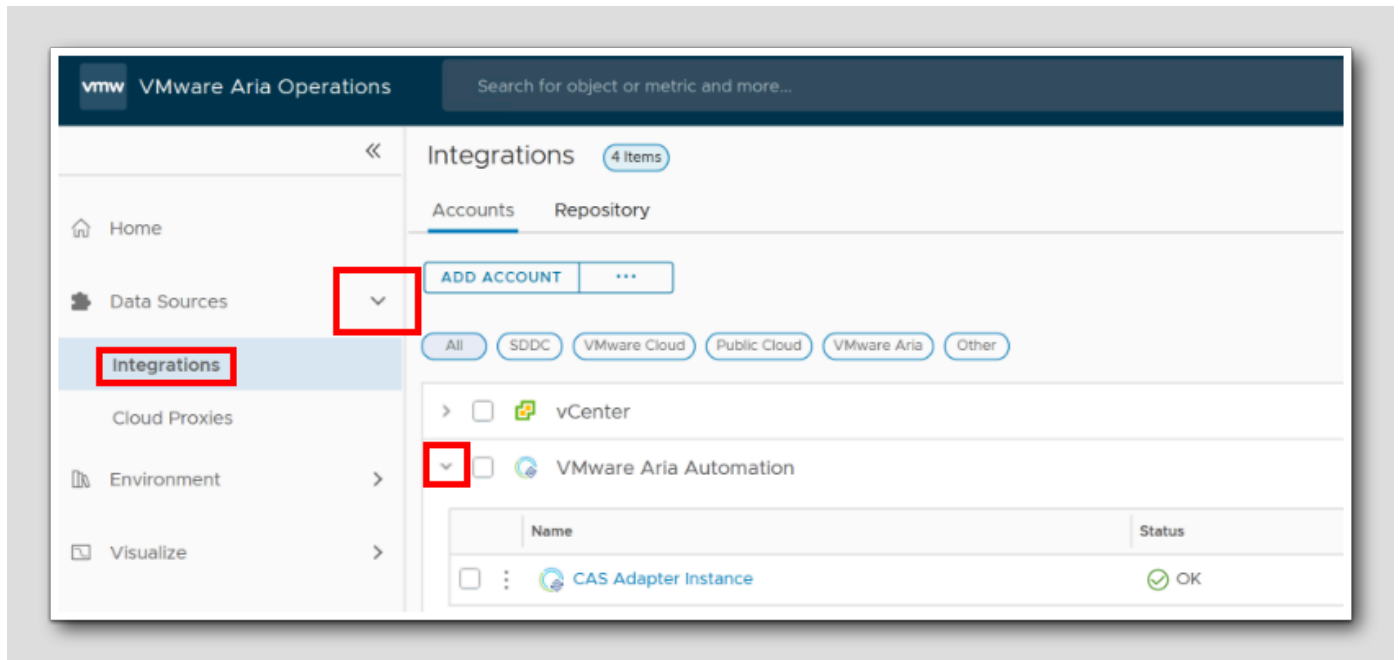
[15]



1. Our next step to configure a Business Intent is to click EDIT

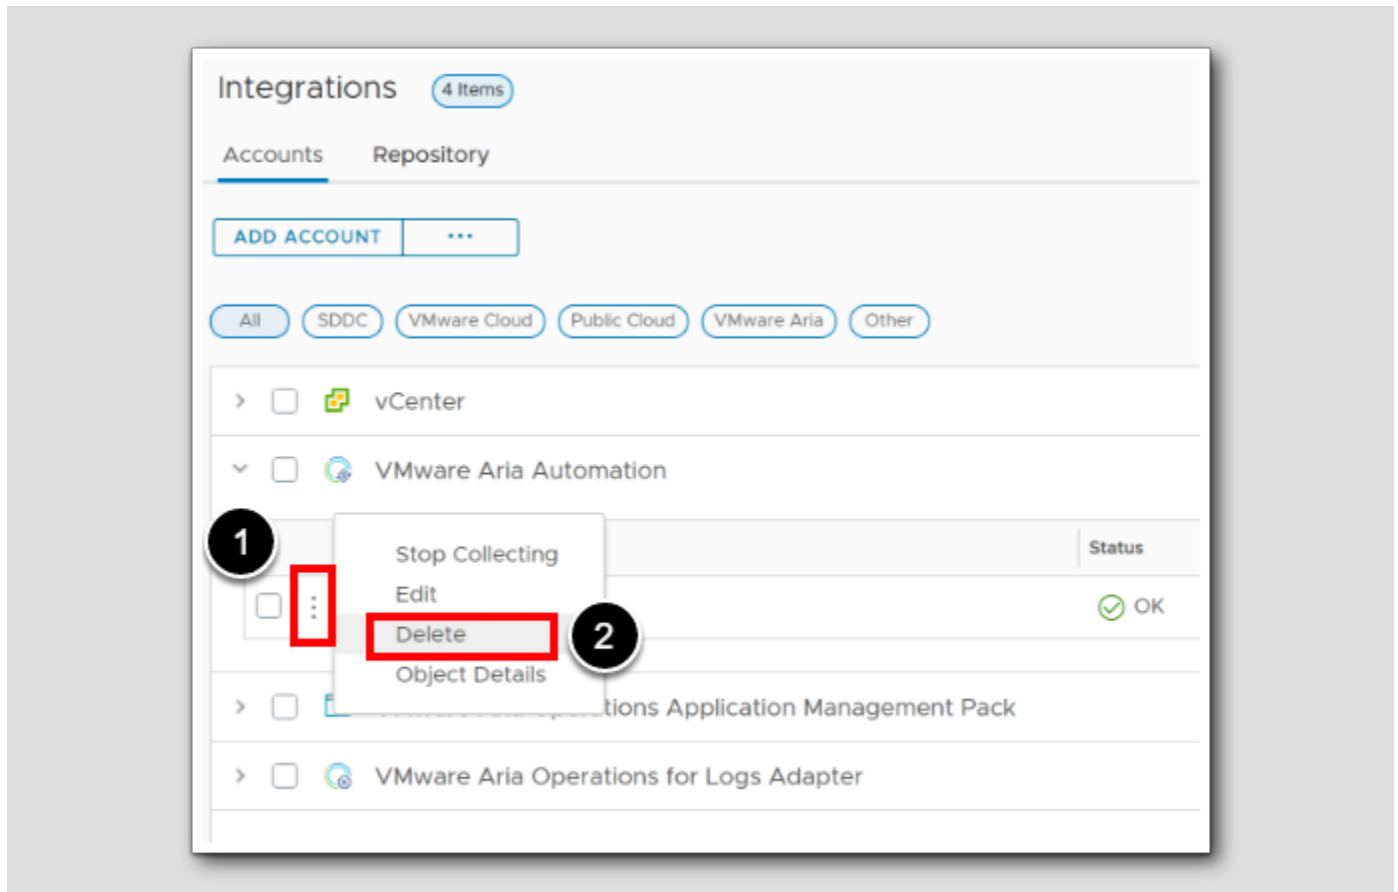
As you can see The EDIT button is grayed out and we can't continue. The reason for this is that we We need to break the co-operation between the two State of the art tools to be able to use the Plac

## Find the Automation integration



1. Expand Data Sources
2. Click on Integrations
3. Expand VMware Aria Automation

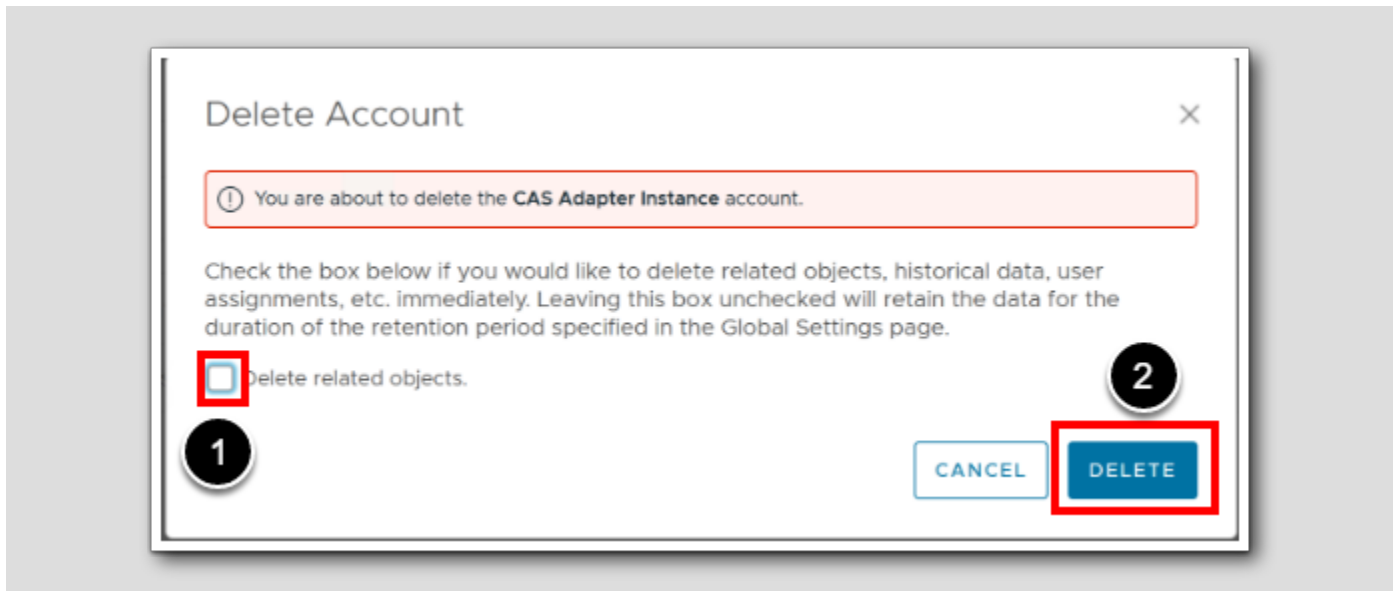
## Deleting the integration



1. Click the ellipsis
2. Select Delete

Confirm the delete

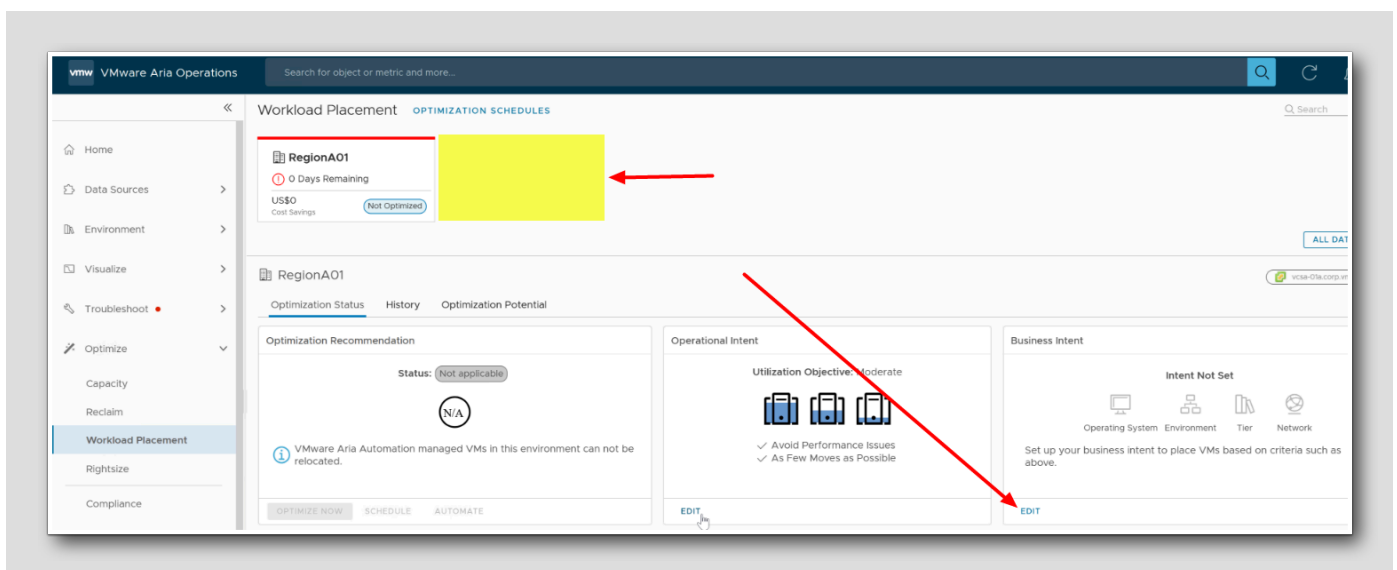
[154]



1. Make sure the Delete Related Objects is unchecked
2. Click Delete

Check the result

[155]



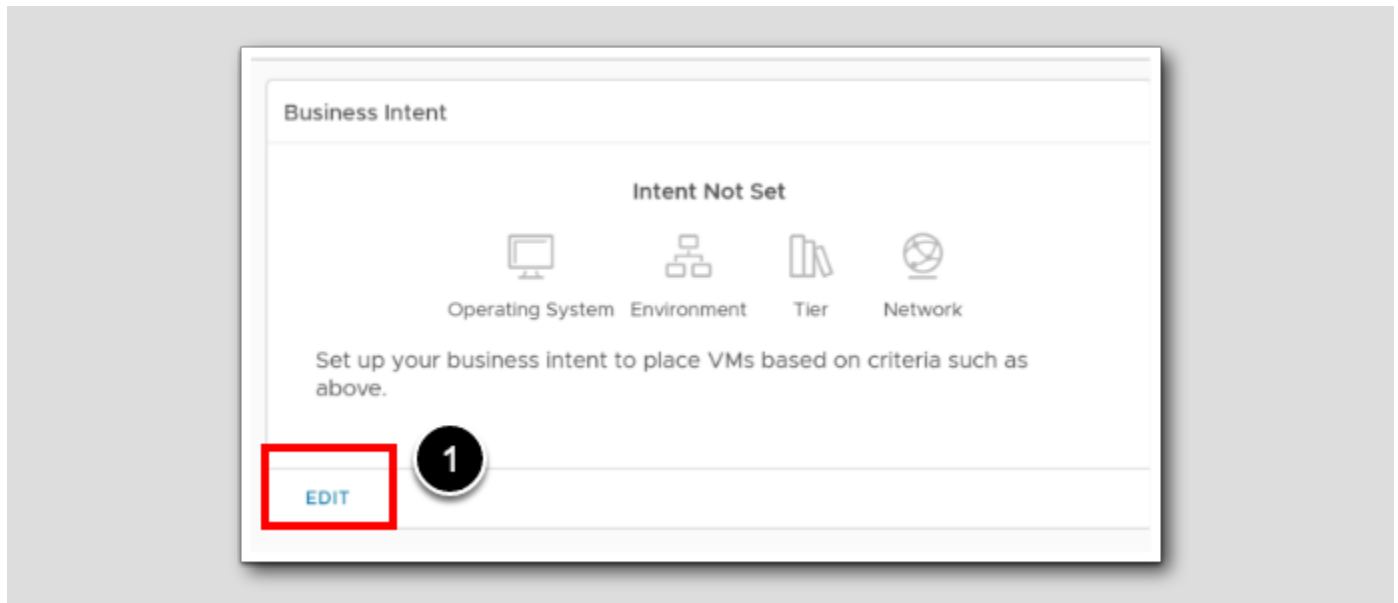
What happened?

First of all the representation of our RegionA01 datacenter through Aria Automation disappeared

Second The Business Intent opened up, and the Grayed out Edit button became available

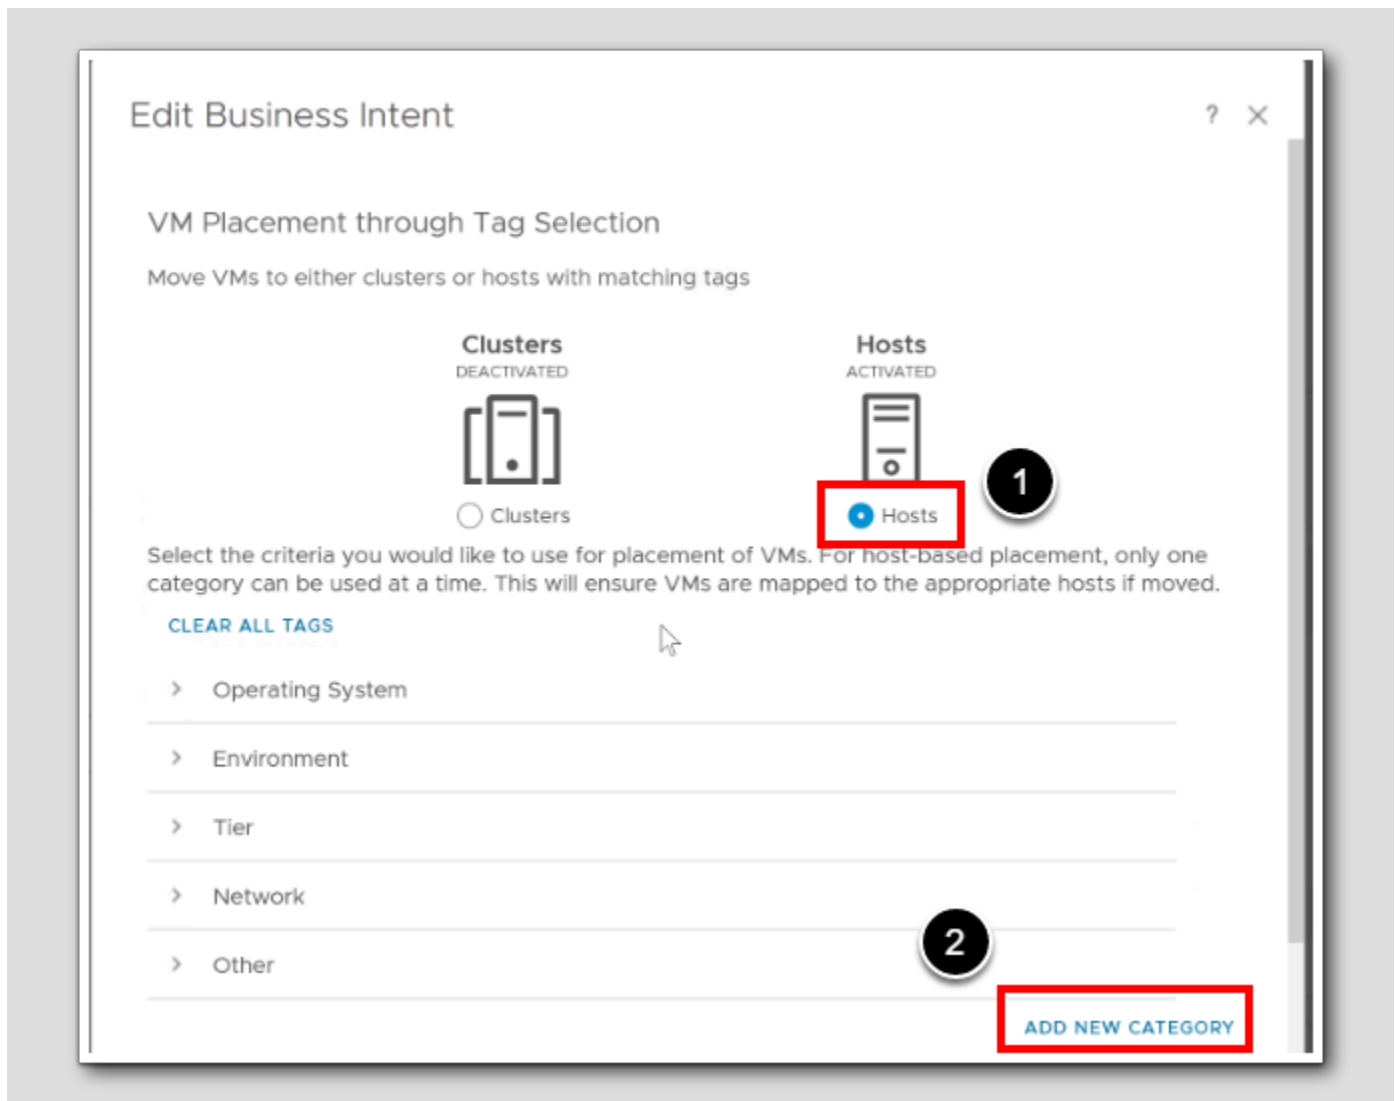
## Edit Business Intent

[156]



1. In the Business Intent Card, Click EDIT

## Cluster Based, or Host Based?



**Cluster Based**, will Move VMs to clusters with matching tags. We could select the criteria we would like to use for placement of VMs to ensure VMs are mapped to the appropriate clusters if moved. Only one category can be prioritized at a time. VM with higher priority tags will be moved last.

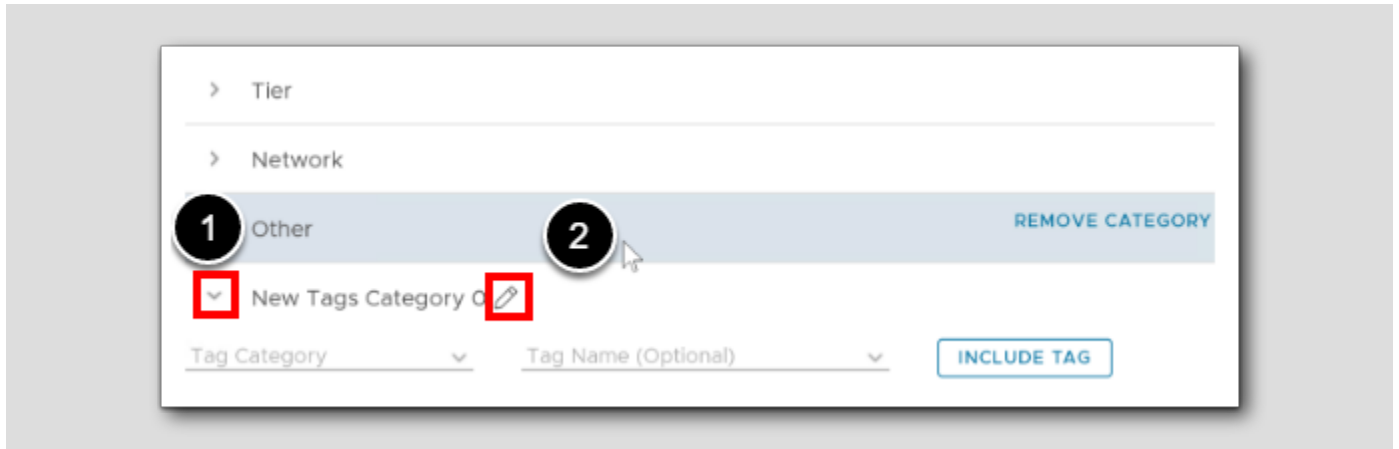
**Host based**, will place VMs through Tag Selection to hosts with matching tags. Select the criteria you would like to use for placement of VMs. Also for host-based placement, only one category can be used at a time. This will ensure VMs are mapped to the appropriate hosts if moved.

1. Select Hosts
2. Click Add New Category



## Renaming the category

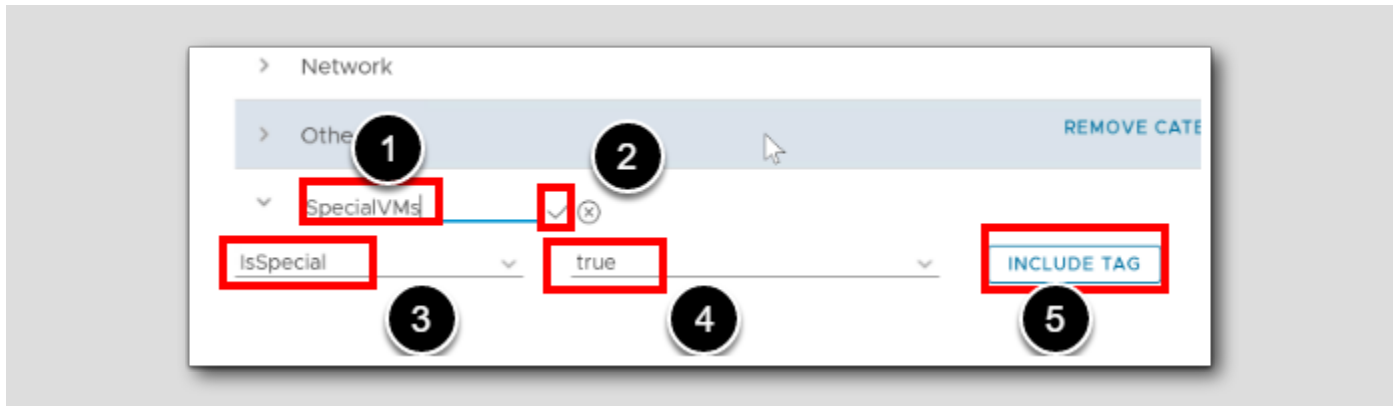
[158]



1. Click to expand the New Category
2. Edit the name by Clicking the pencil

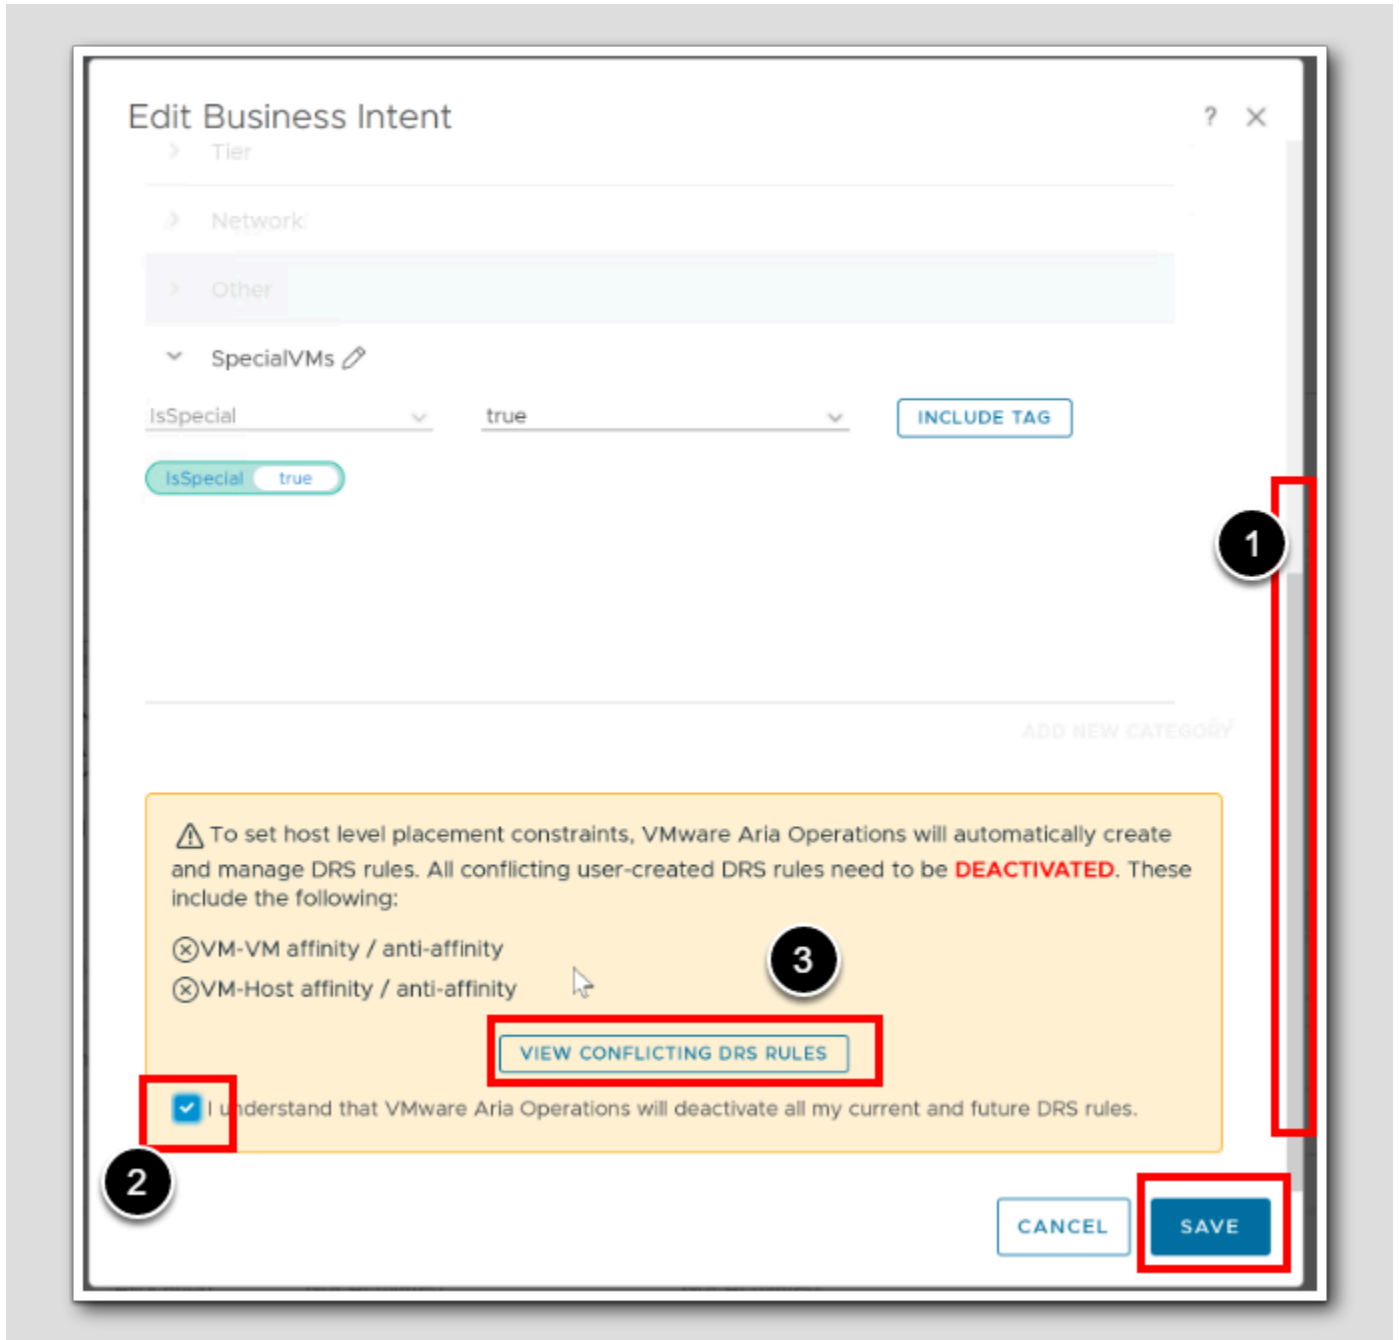
## Adding the name

[159]



1. Rename the category to SpecialVMs
2. Click the confirm icon
3. Add a Tag Category called IsSpecial
4. Set the Tag Name to true
5. Click the Include Tag

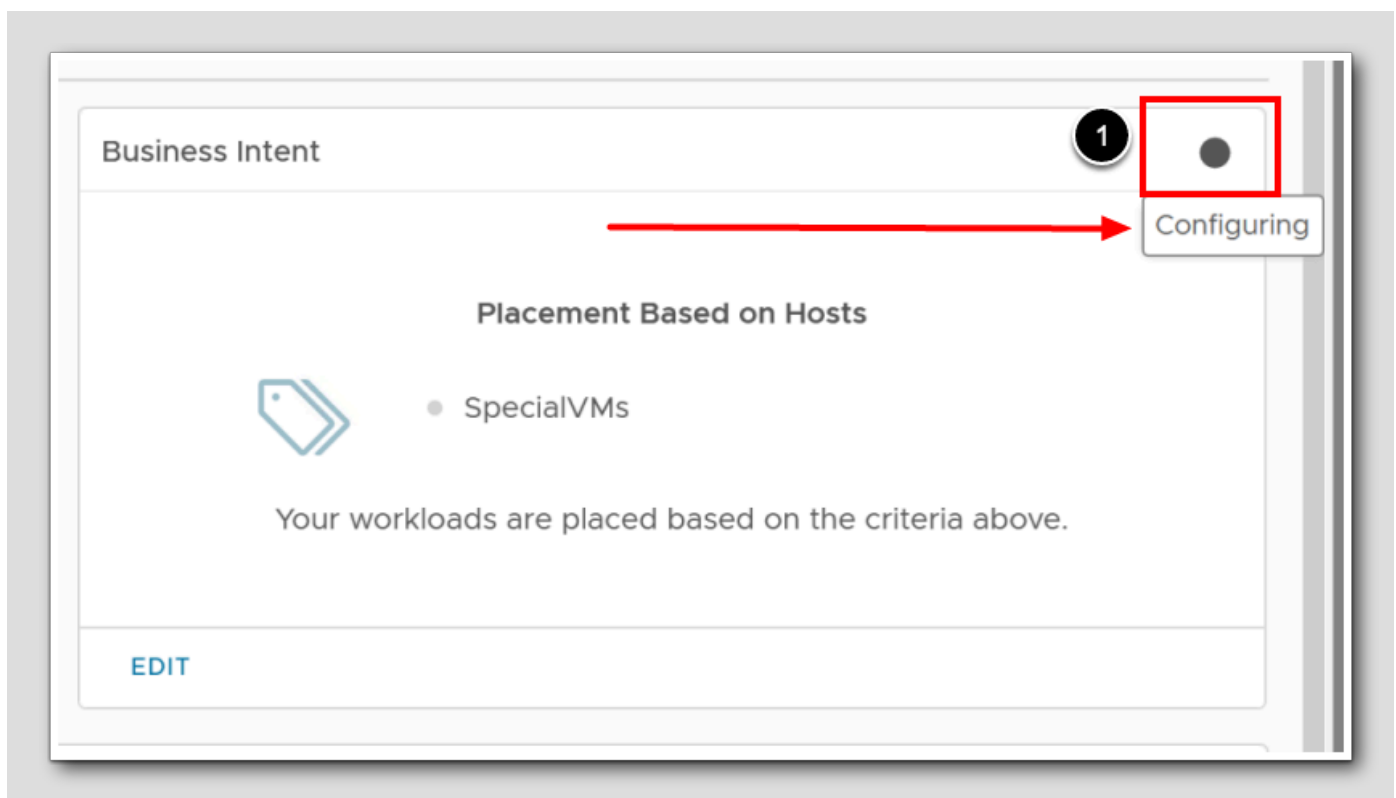
### Confirm Affinity



1. In the dialog, Scroll down to the bottom of the page
2. Click "I understand that VMware Aria Operations will deactivate all my current and future DRS rules."
3. Click VIEW CONFLICTING DRS RULES
  - Observe if there are any. Note: It will state that "0 user-created DRS Rule(s) will be deactivated across clusters in this datacenter"
4. Click SAVE

## Configuring and Configured states

[16]



1. First the state will be "Configuring", you will see this if you hover with the mouse over the black dot

When you refresh a couple of times with the top right refresh button, this black dot will turn green, at it will say "Configured" when you hover the button. (not shown)

Next step is to test. We will Tag a host and VMs in vSphere using vCenter.

Start and log in to vCenter

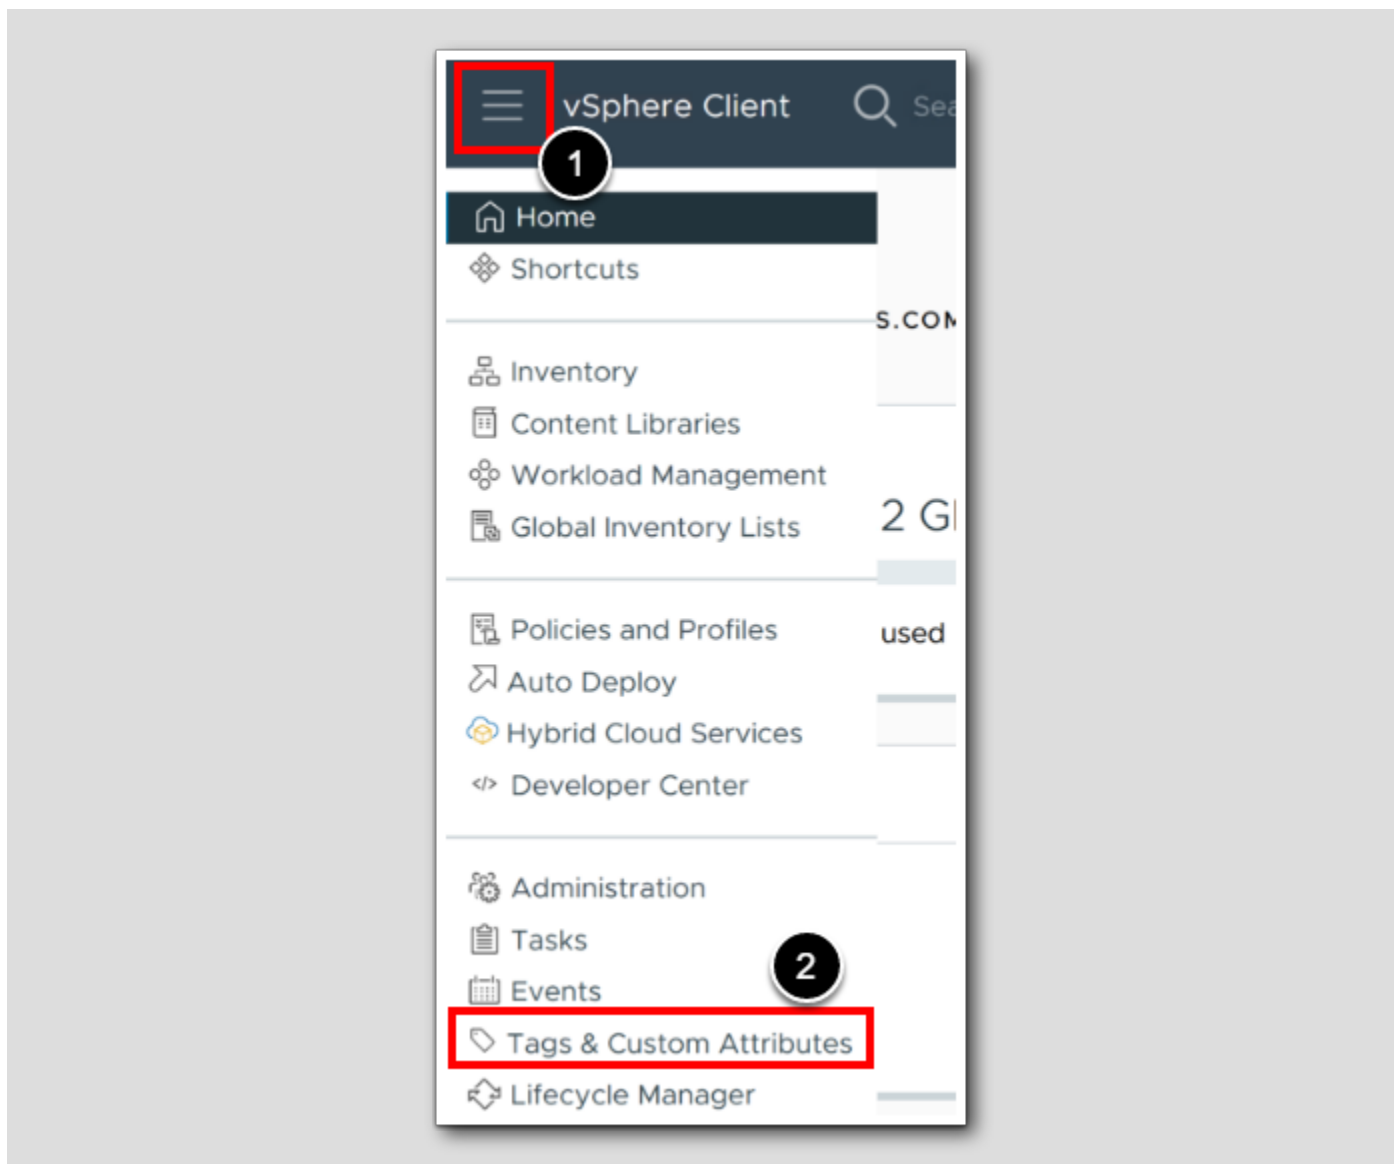
[162]



1. Hold CTRL and press T (ctrl+T) to open a New Tab (not shown)
2. Click vCenter
3. Check the "Use Windows session authentication" checkbox
4. Click Login

## Find Tags

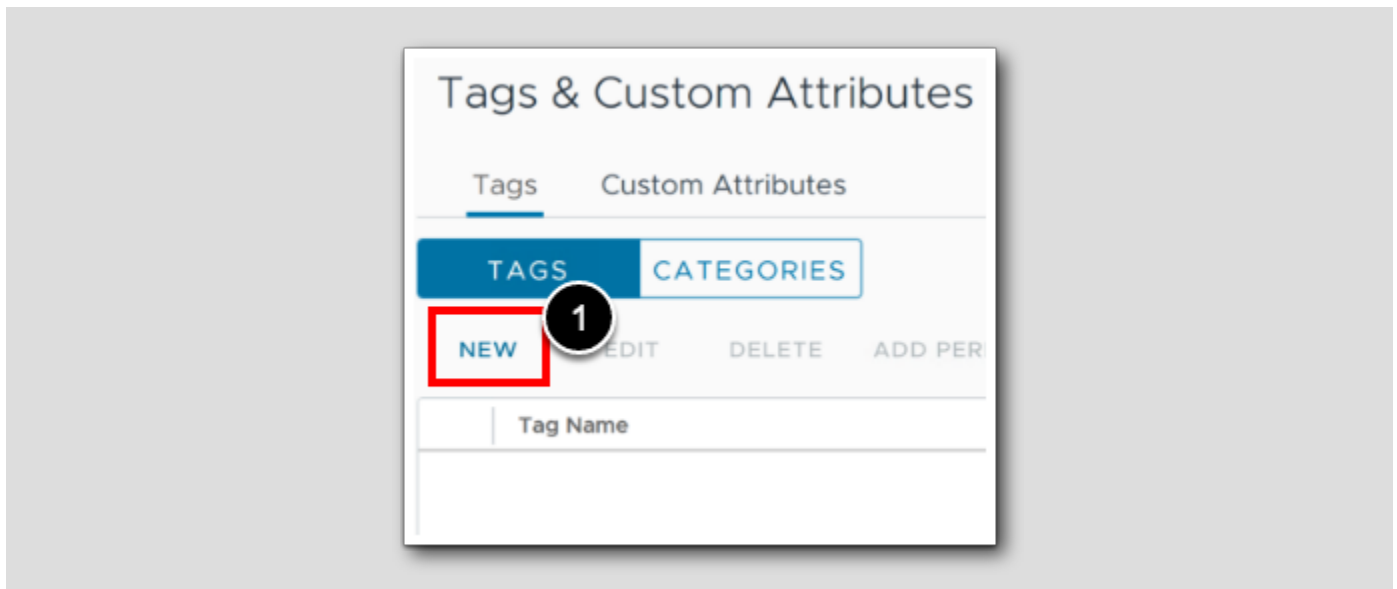
[163]



1. In the vSphere Client click the Ellipsis Menu
2. Choose Tags & Custom Attributes

### Create a new tag

[164]



Nearly there.

1. To create a Tag and a Tag category, Click New

## A new Tag and Tag Category

The screenshot shows a 'Create Tag' dialog box with the following fields and actions:

- Name:** A text input field containing the value 'true'. A red box highlights this field, and a circled '1' points to it.
- Description:** A text area containing the text 'Set the Tag Category IsSpecial = true'.
- Category:** A dropdown menu with the selected option 'Create New Category'. A red box highlights this option, and a circled '2' points to it.
- Buttons:** 'CANCEL' and 'CREATE' buttons are located at the bottom right of the dialog.

1. Set the tag Name to true
2. Click Create New Category

**Create Category** [X]

**Category Name:**  1

**Description:**

**Tags Per Object:**  One tag  Many tags 2

**Associable Object Types:**

- All objects
- Folder
- Datacenter
- Datastore Cluster
- Distributed Switch
- Content Library
- Network
- vApp
- Cluster
- Datastore
- Distributed Port Group
- Host
- Library Item
- Resource Pool
- Virtual Machine

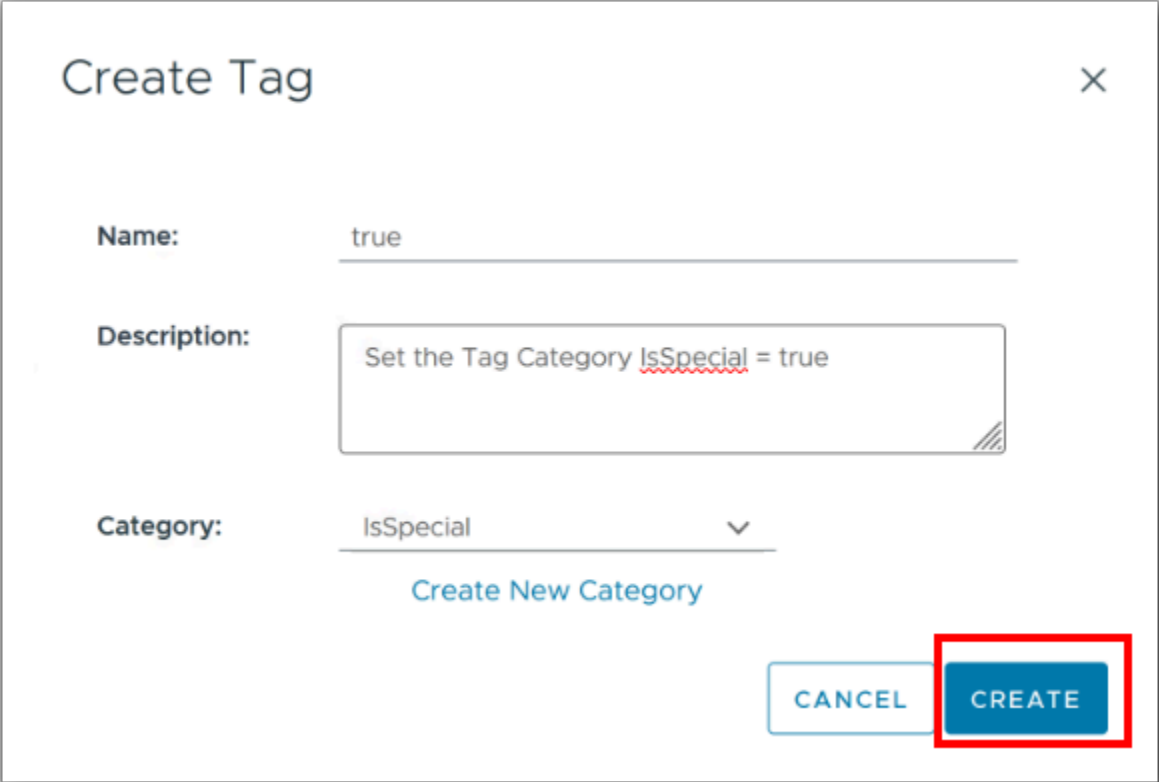
3

1. Set the Category Name to IsSpecial
2. Click on Many Tags
3. Click Create



## Finish the Tag Creation

[167]



**Create Tag** [X]

**Name:** true

**Description:** Set the Tag Category IsSpecial = true

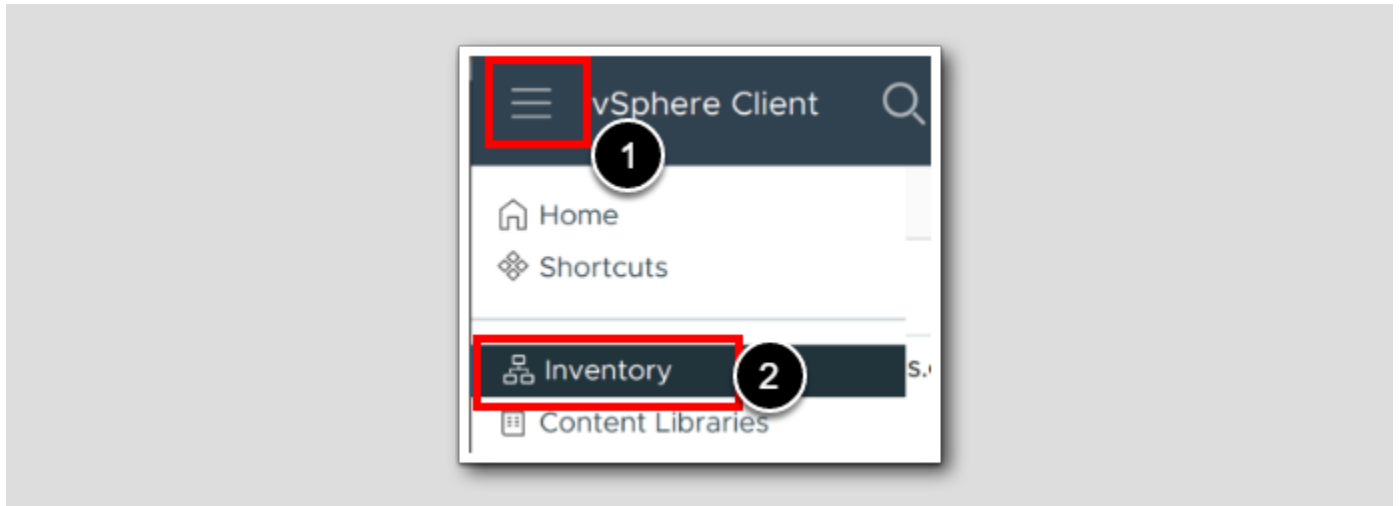
**Category:** IsSpecial [v]  
[Create New Category](#)

CANCEL CREATE

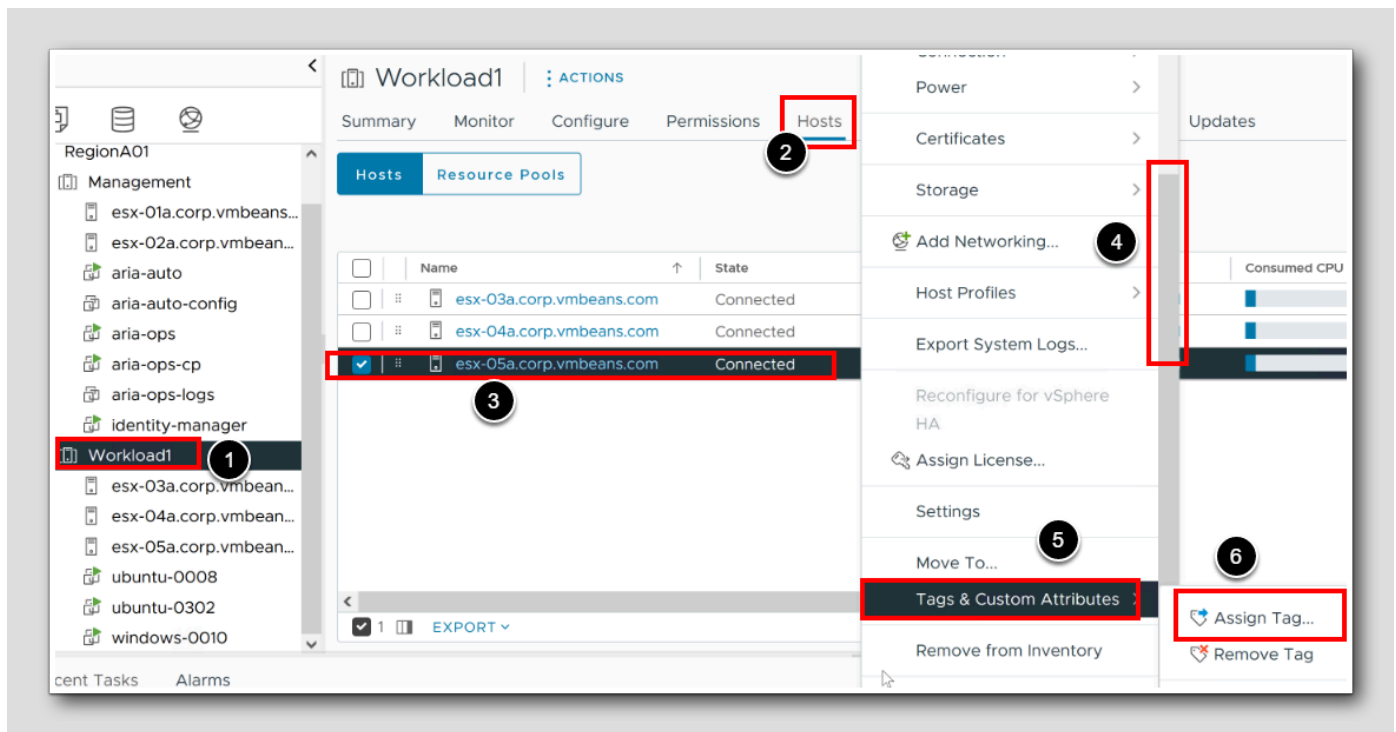
Click Create

## Go to inventory

[168]



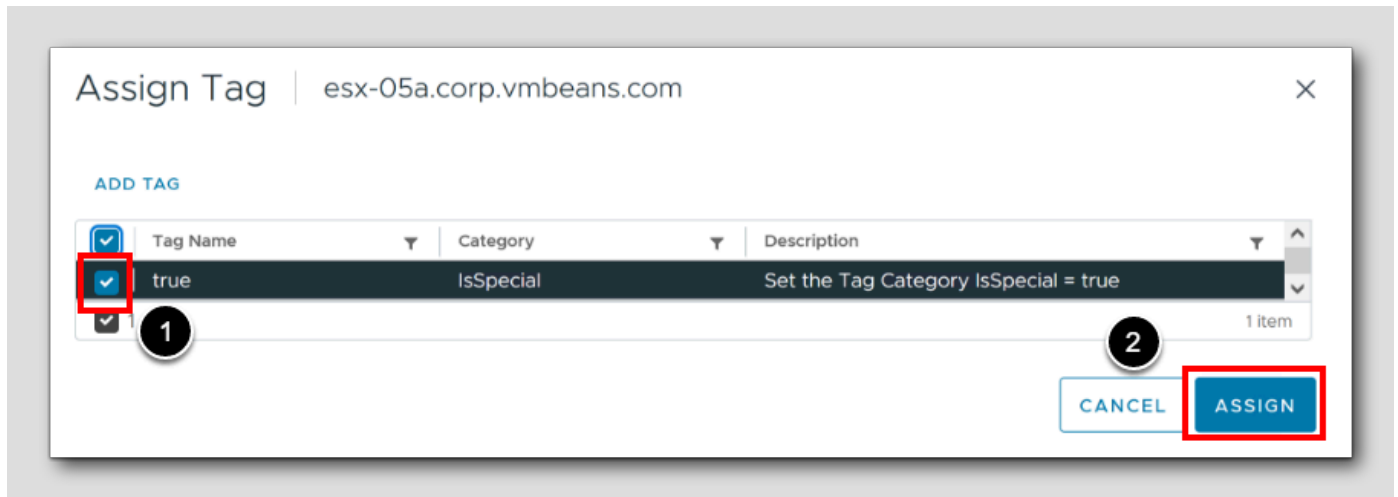
1. Click the ellipsis menu
2. Choose Inventory



We are going to assign our new tag to one specific host: `esx-05a.corp.vmbeans.com`

1. Select the **Workload1** Cluster
2. Click on **Hosts**
3. Right Click the ESXi Host `esx-05a.corp.vmbeans.com`
4. Scroll down
5. Click **Tags & Custom Attributes**
6. Click **Assign Tag**

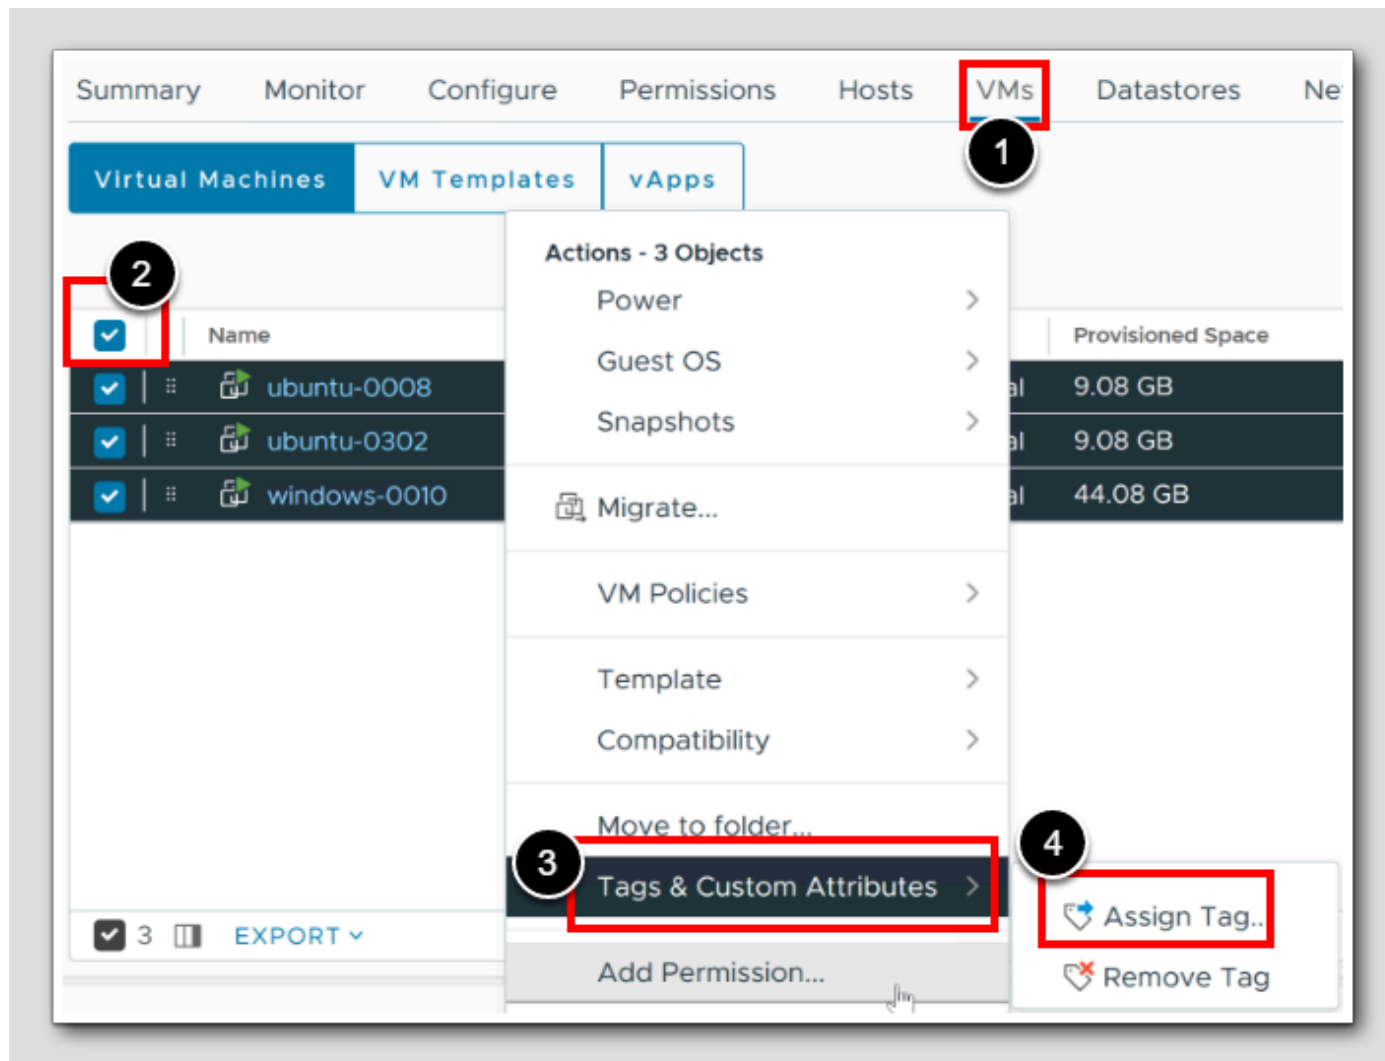
## Assign the Tag



1. Click on the checkbox
2. Click Assign

Now we need to do the same with some VMs!

## Assign VMs

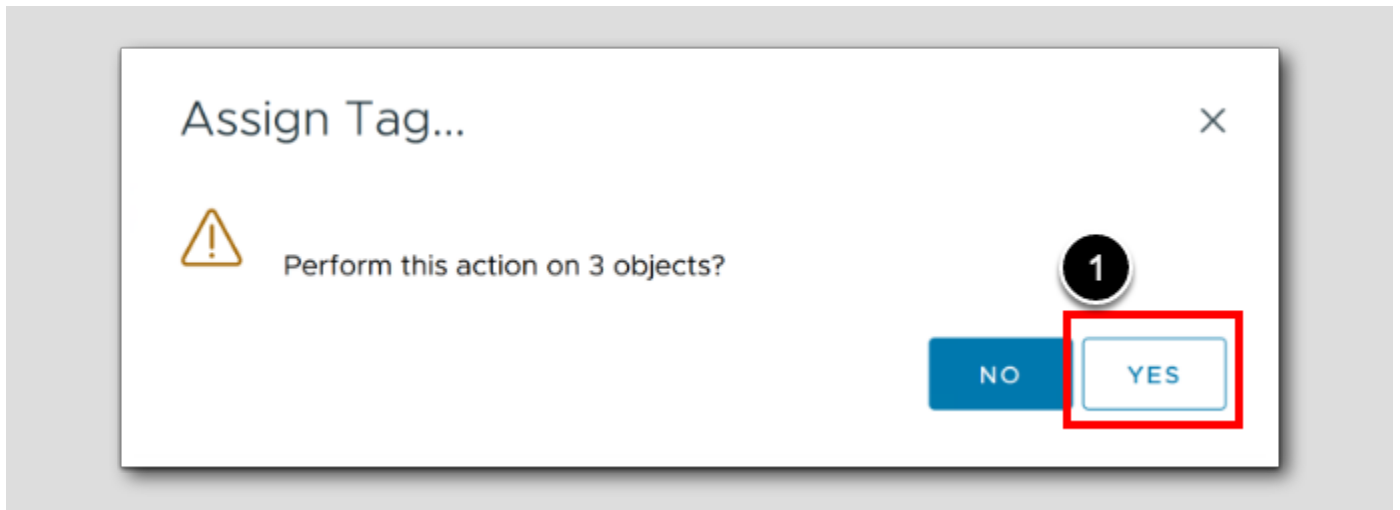


Using the same method we will assign tags to all visible VMs in this cluster. Note: Remember this is just a test

1. Click VMs
2. Select all VMs
3. Click on Tags & Custom Attributes
4. Click Assign Tag

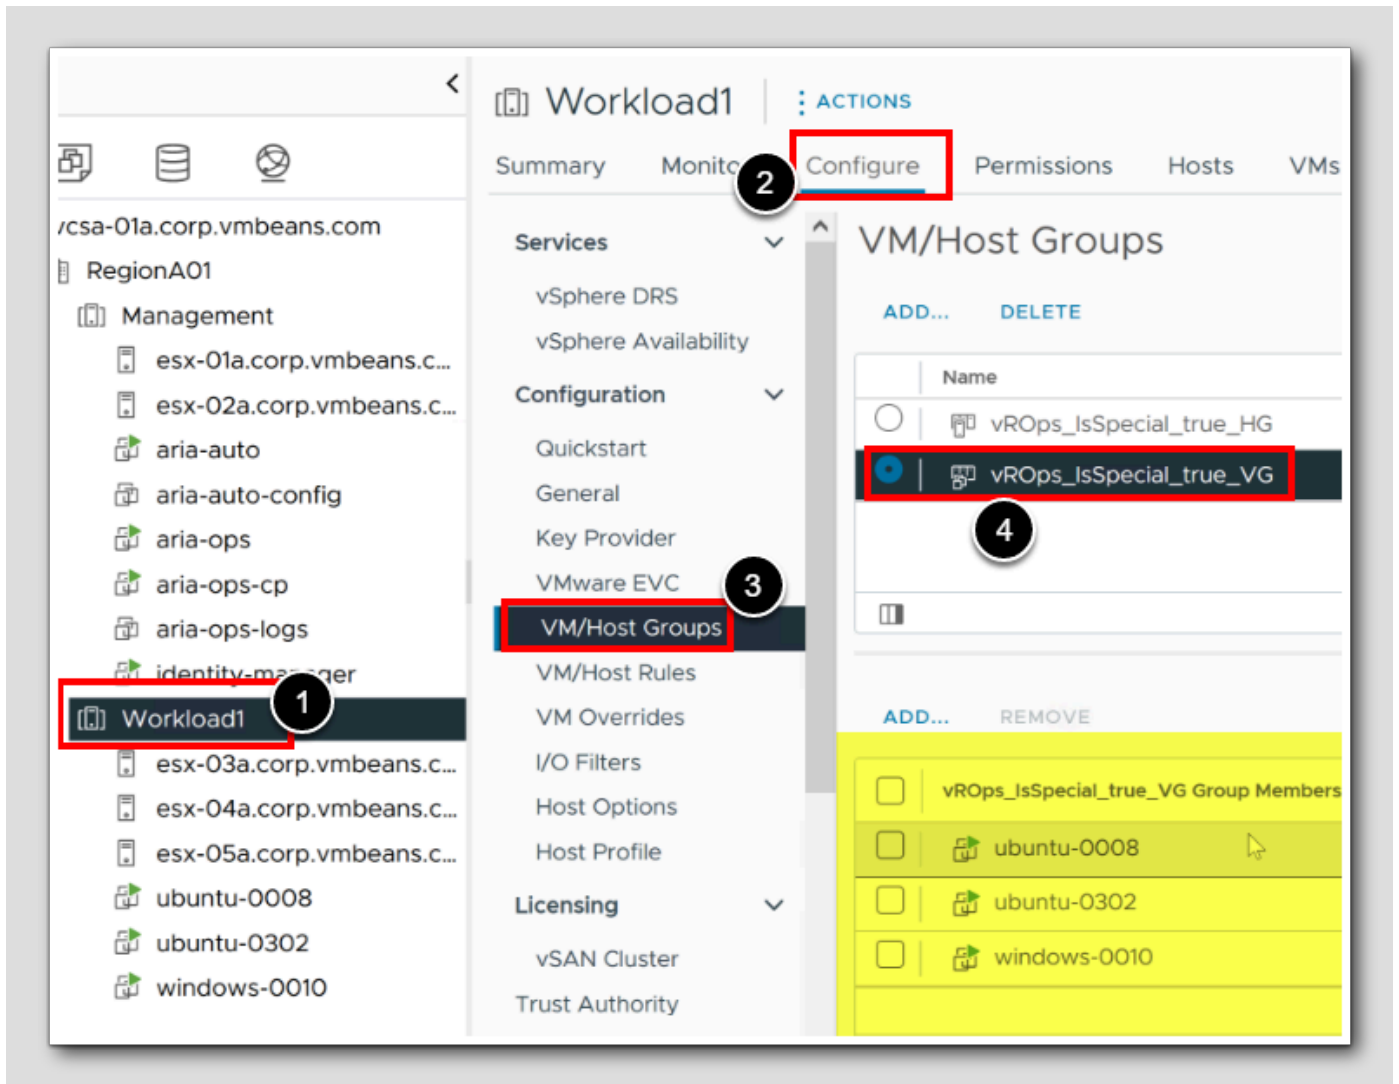
Confirm

[172]



1. Click YES

Confirming the VM/Hosts groups



Aria Operations will make sure that Virtual Machines that are members of a VM Group must run on hosts that are members of a Host Group. Let's Check the VM group.

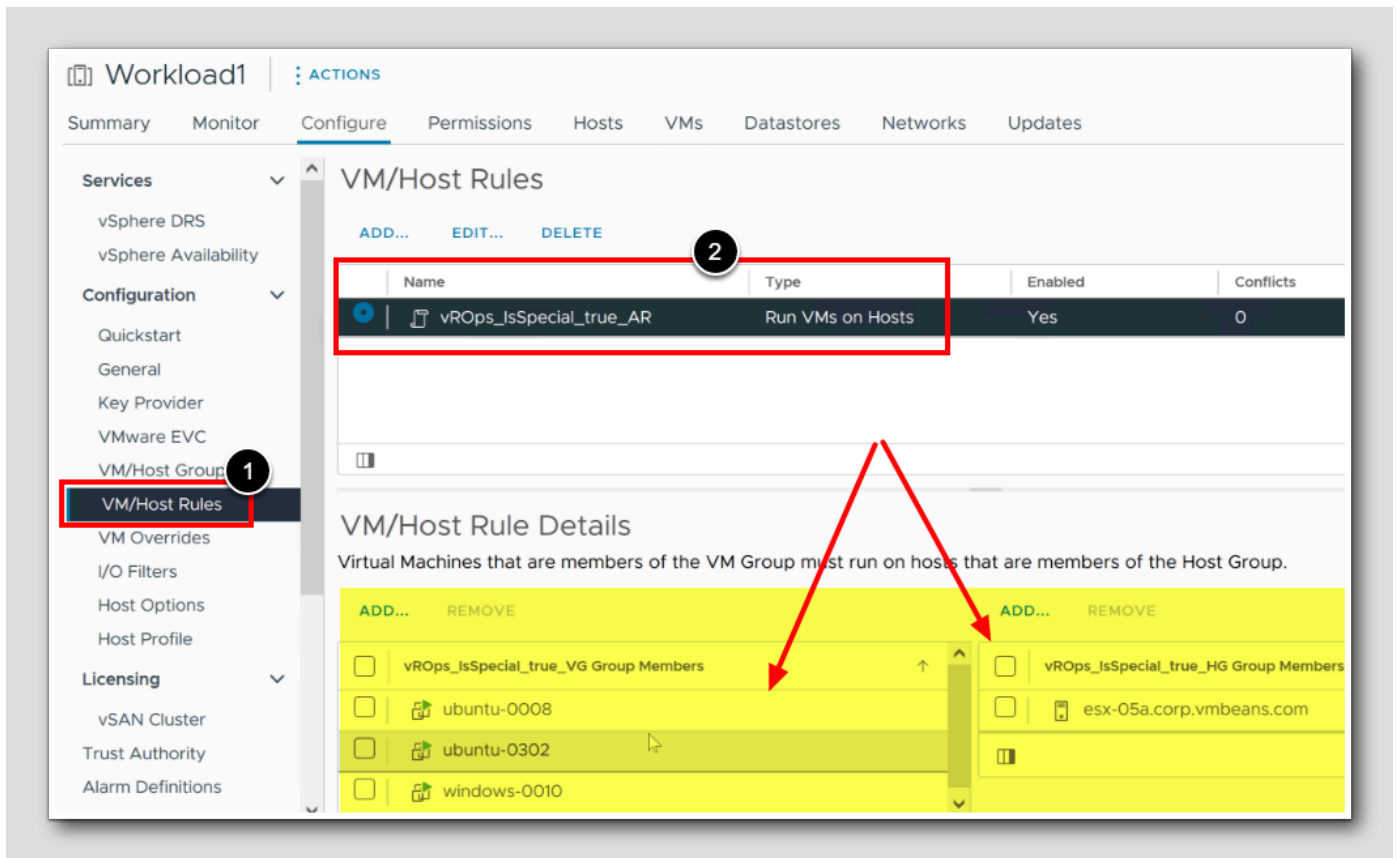
1. Click the Cluster **Workload1**
  2. Click Configure
  3. Click VM/Host Groups
  4. To check the members, Click VM group **vROps\_IsSpecial\_true\_VG**
- Observe that we've got all the VMs as members.

Note: Check the hosts group as well to see if the only member is **esx-05a.corp.vmbeans.com** (not shown).

Next we will confirm the VM/Hosts rules

### Confirm the Hosts/VM rules

[174]





1. Continue by clicking **VM/Host Rules**
2. Select the New rule `vROps_IsSpecial_true_AR`

Observe Both the the VM Group members and Host Group Members

*This Concludes this Module*

## Conclusion

[175]

In summary, Business Intent is the key to unlocking efficiency, streamlined operations, and software alignment with our business objectives.

By leveraging user-defined rules, we can drive automation and informed decision-making, ensuring workload decisions are in sync with our business needs. The advantages we gain include compliance, tag-based placement, workload separation, and optimized resource utilization.

## You've finished Module 5

[176]

Congratulations on completing the lab module.

If you are looking for additional general information on Aria Operations, try one of these:

- **VMware Product Public Page - Aria Operations:** <https://www.vmware.com/products/aria-operations.html>
- **Aria Operations - Documentation:** <https://docs.vmware.com/en/VMware-Aria-Operations/index.html>

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 6 - Report Generation in VMware Aria Operations (15 minutes) Basic

### Introduction

[178]

With the VMware Aria Operations reporting functions, you can generate a report to capture details related to current or predicted resource needs. You can download the report in a PDF or CSV file format for future and offline needs.

### Log in to Aria Operations

[179]

We will log in to a live instance of Aria Operations running in this lab.

### Open the Firefox Browser from the Windows Task Bar

[180]

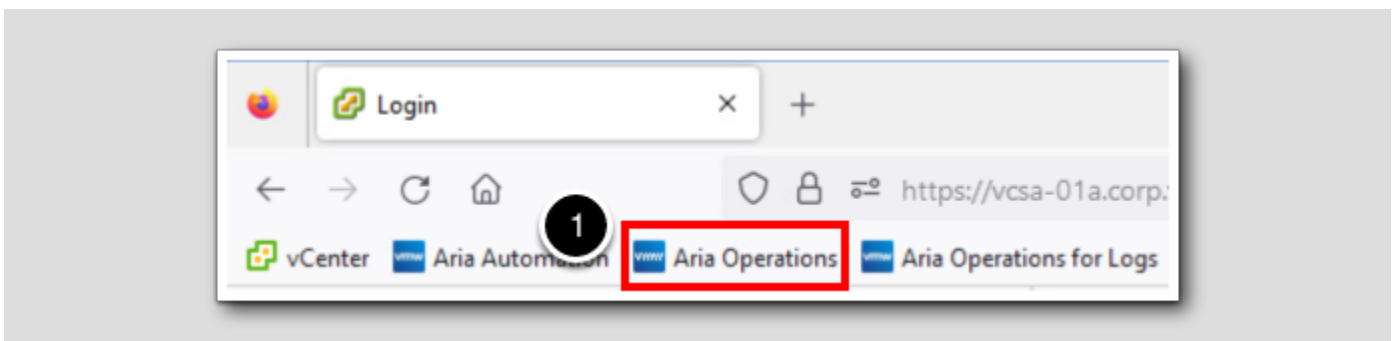


If the browser is not already open, launch Firefox.

1. Click the Firefox icon in the Windows Quick Launch Task Bar at the bottom of the screen.

### Navigate to Aria Operations


[181]



1. Click the Aria Operations bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[182]



Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login

[183]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

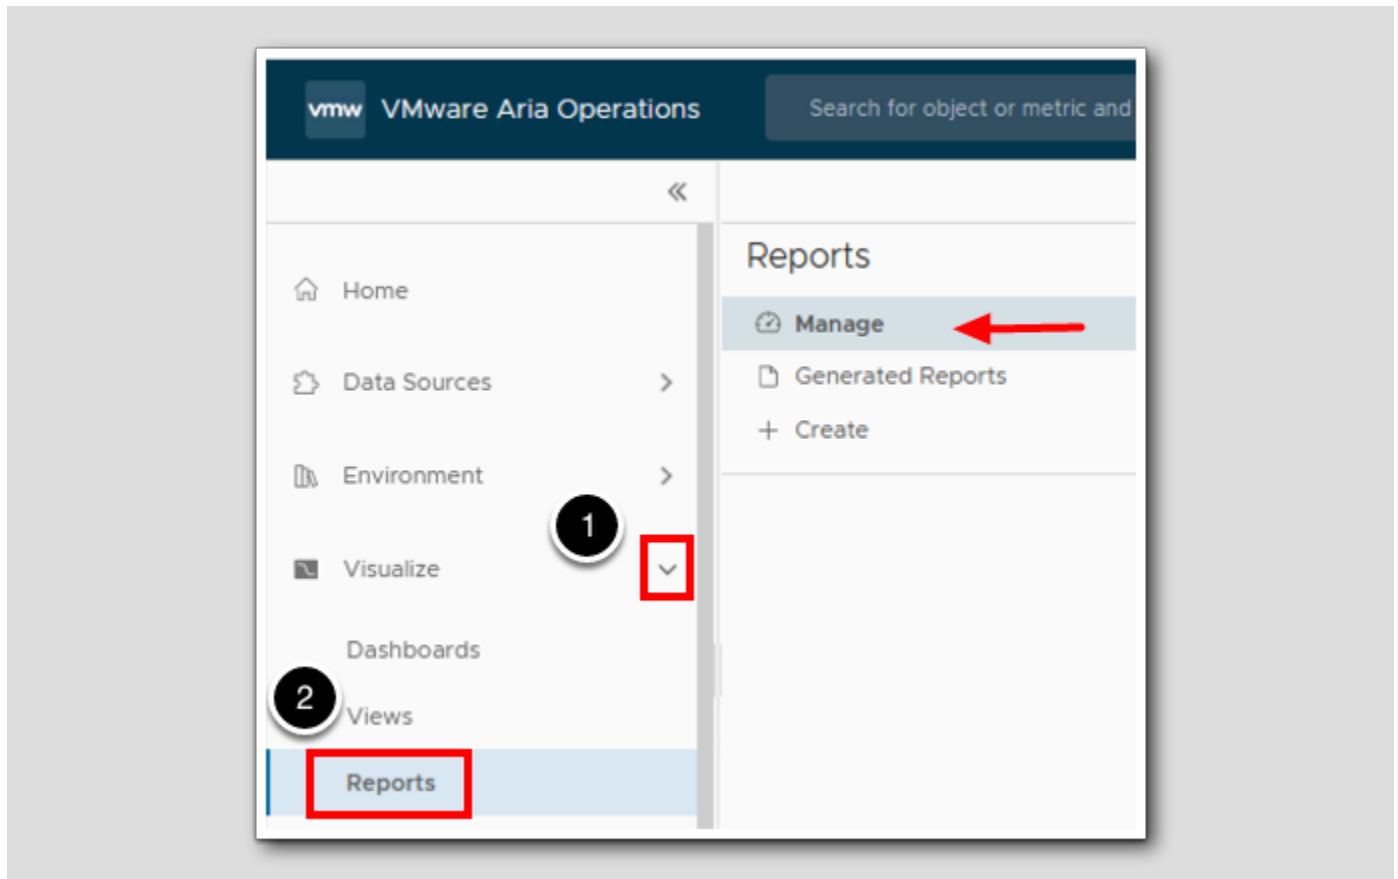
1. Click Sign in

## Introduction to Reports

[184]

In this lesson we will run some reports and build a custom report to familiarize how the reporting function works in Aria Operations.

## Open Reports



1. Expand Visualize.

2. Click on Reports.

Notice that **Manage** is selected by default.

## Report Templates

The screenshot displays the 'Report Templates' management interface. At the top, there is a search bar with the placeholder text 'Type here to apply filters' and a filter icon. Below the search bar is a table listing various report templates. The first row is selected, indicated by a blue checkmark in the selection column. The table columns are: Name (with an upward arrow), Description, Subject, Generation (0), Scheduling (0), Last Modified (6/16/23 10 -), Last run, and Modified By (admin). A red arrow points to the pagination control at the bottom right of the table, which displays '1 - 50 of 68 items' and navigation arrows.

<input type="checkbox"/>	Name ↑	Description	Subject	Generat...	Schedul...	Last Modi...	Last run	Modified ...
<input checked="" type="checkbox"/>	Capacity Report - Datastores	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - Distributed Port	This report provid...	vSphere Distribute...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - Distributed Swil	This report provid...	vSphere Distribute...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - Environment	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - Pods	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - Virtual Machine	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - vSphere Cluster	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - vSphere Cluster	This report provid...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Capacity Report - vSphere Hosts	This report provid...	Host System	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Cluster Cost Report	Report that contai...	Cluster Compute ...	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Compliance Report - vSphere Sec	This report shows ...	Symptom	0	0	6/16/23 10 -		admin
<input type="checkbox"/>	Configuration Report - Datastores	This report helos t...	Datastore	0	0	6/16/23 10 -		admin

1 - 50 of 68 items

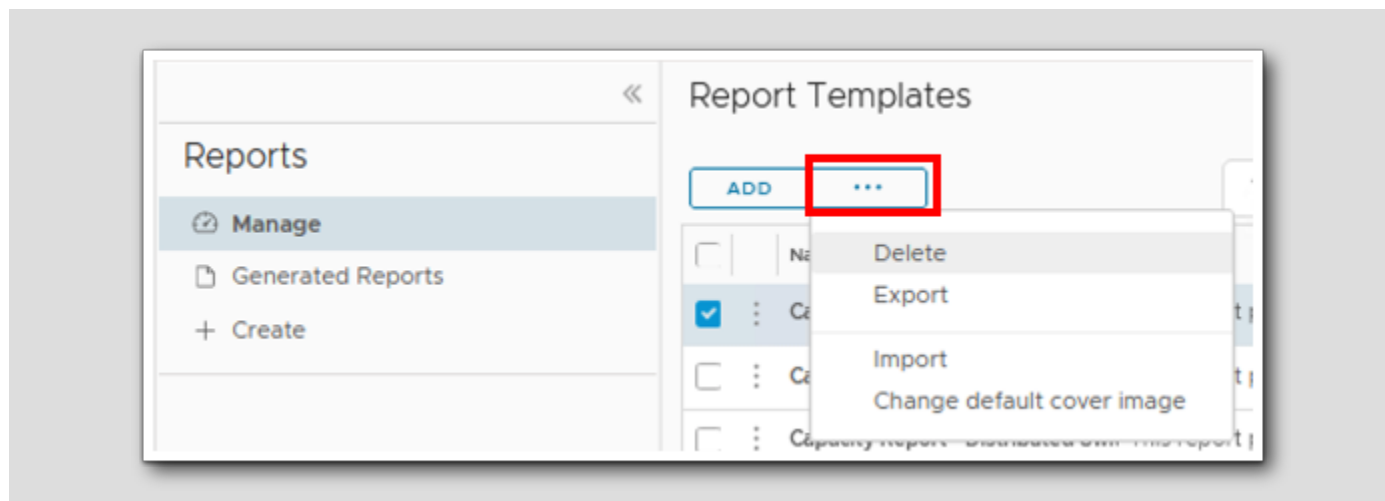
The Manage Reports list shows all of the available reports within vSphere. Notice that there are 68 reports delivered out of the box and are arranged by categories.

The delivered report types are:

- Capacity Reports
- Cluster Cost Reports
- Compliance Reports
- Configuration Reports
- Inventory Reports
- Optimization Reports
- Performance Reports
- Reclamation Reports
- Server Cost Reports
- Utilization Reports
- Virtual Machine Reports
- vSphere Optimization Assessment (VOA) Reports

## More Actions Menu

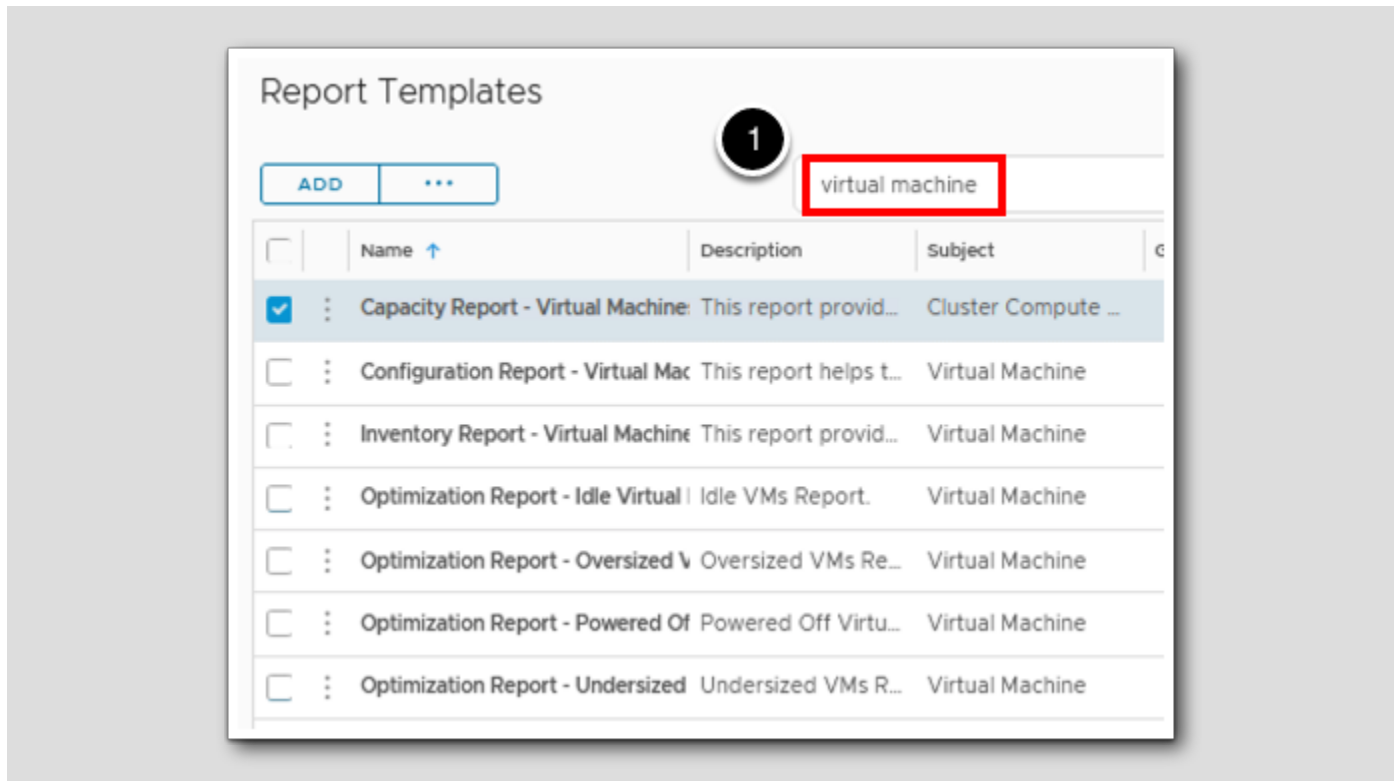
[187]



1. Click on the ... to open the **More Actions** menu.

Notice here, with a report selected (blue checkbox), you can Delete, Export, Import or Change default cover image. Changing the default cover allows you to personalize the report cover page reflecting your companies branding.

## Virtual Machine Reports



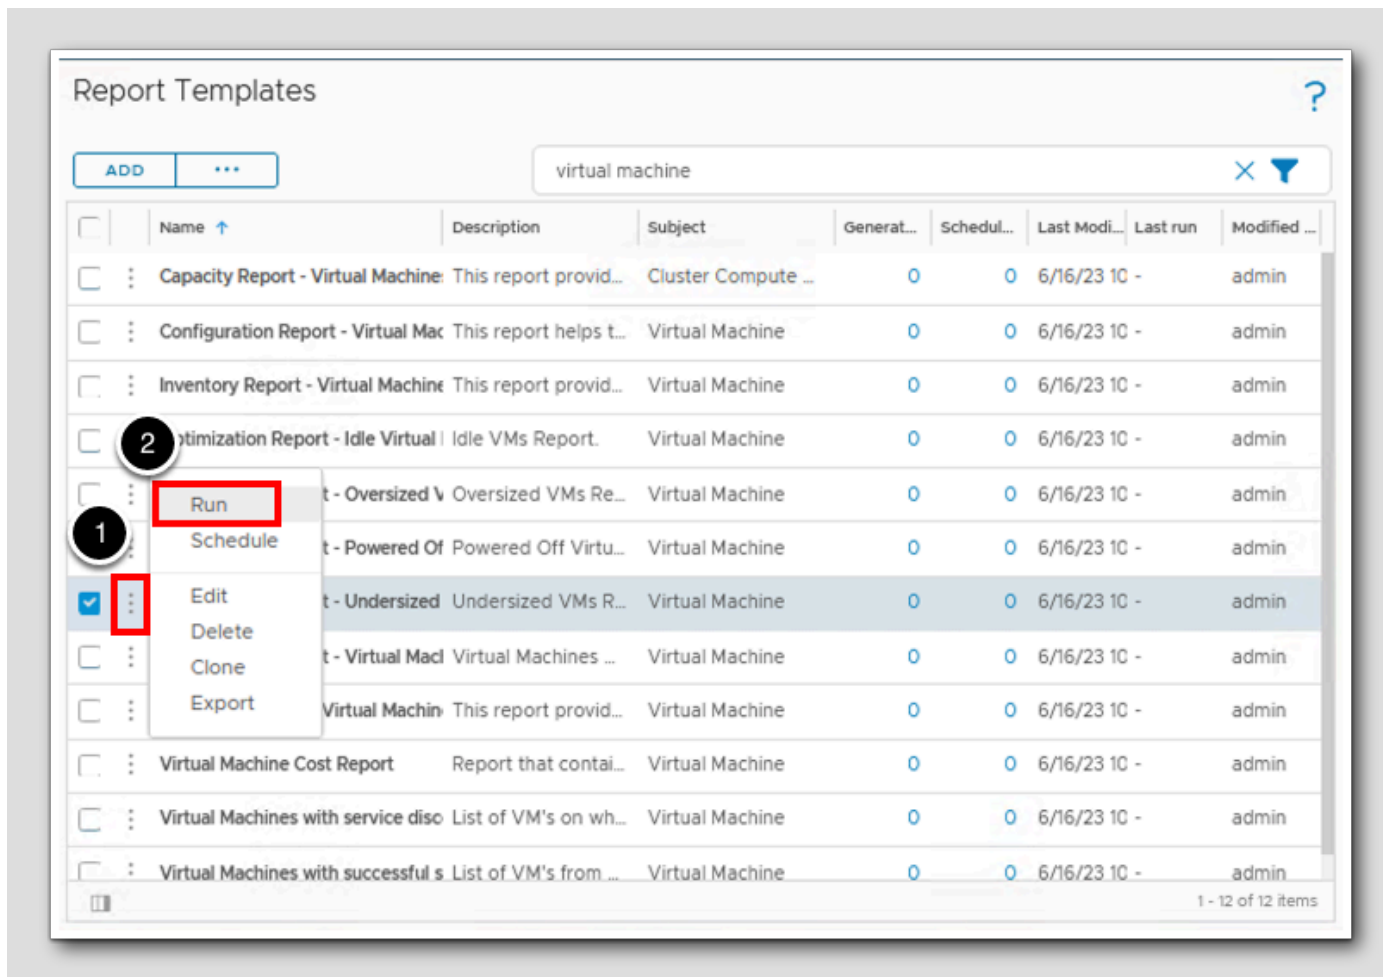
1. In the Search Bar type virtual machine and hit the Enter key.

Searching by virtual machine we filter down the report list to the 12 delivered virtual machine report types.



## The Undersized Report

[189]

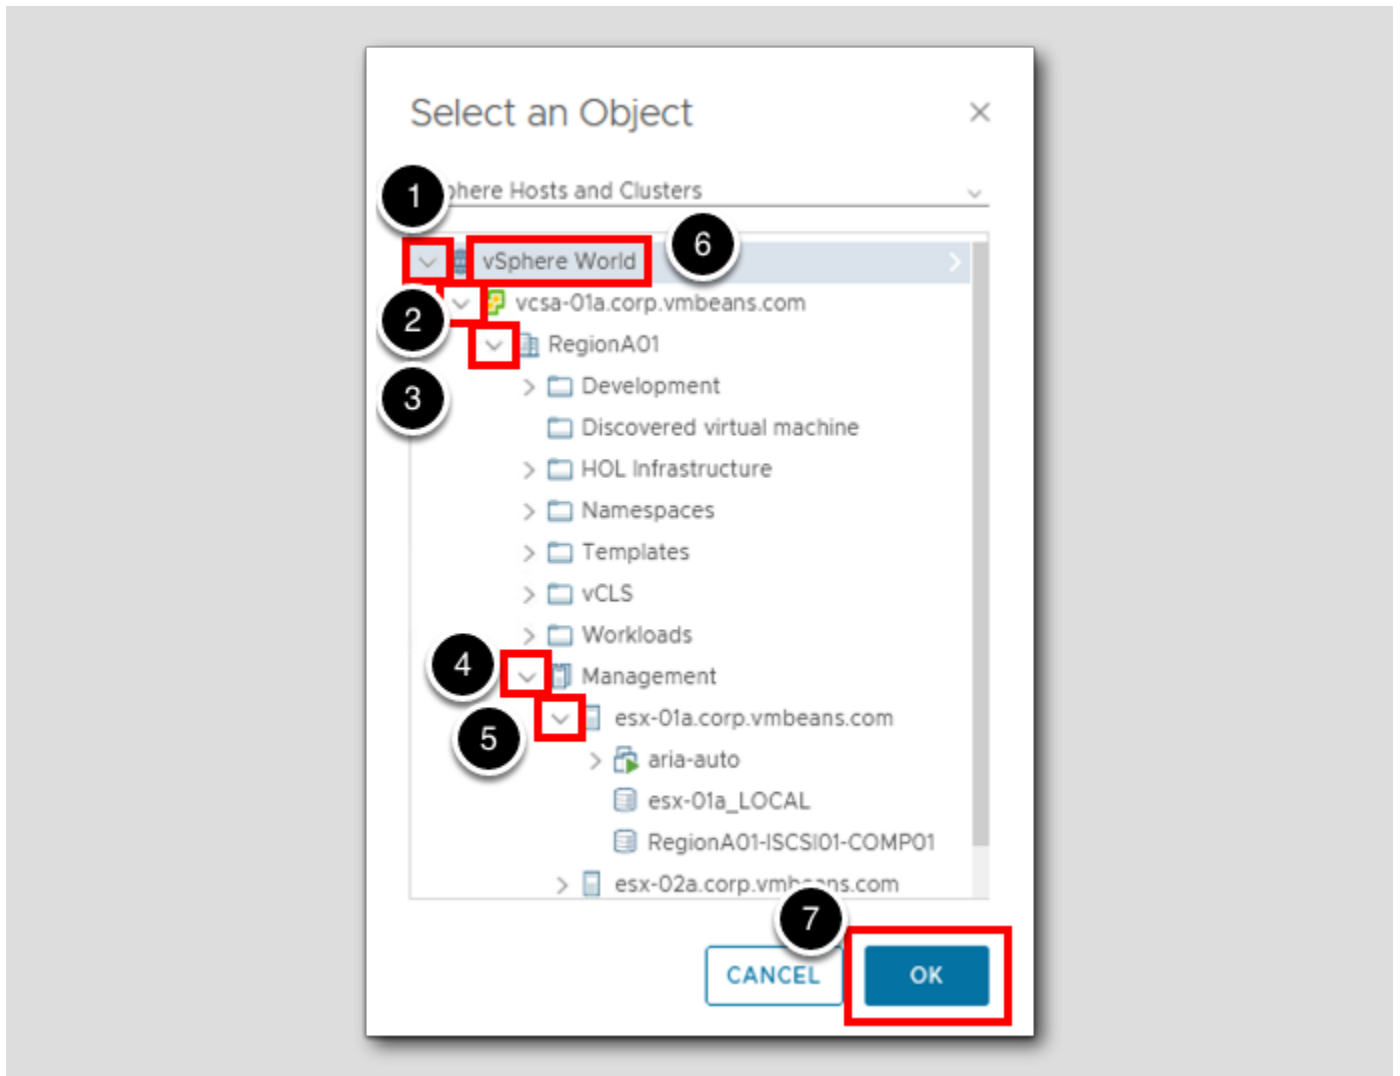


1. Click on the vertical 3 dots next to the Optimization Report - Undersized Virtual Machines.
2. Select Run.

## Select an Object

[190]

In the Select an Object window, you need to select what Object level you want your report to pull information from. By default, vSphere World is displayed unexpanded. If you expand vSphere World out you have the ability to choose a subordinate Object like a (in descending order) vCenter, Datacenter, Folder, Cluster, Host, VM or Datastore. For this lesson we will choose the top level, vSphere World.



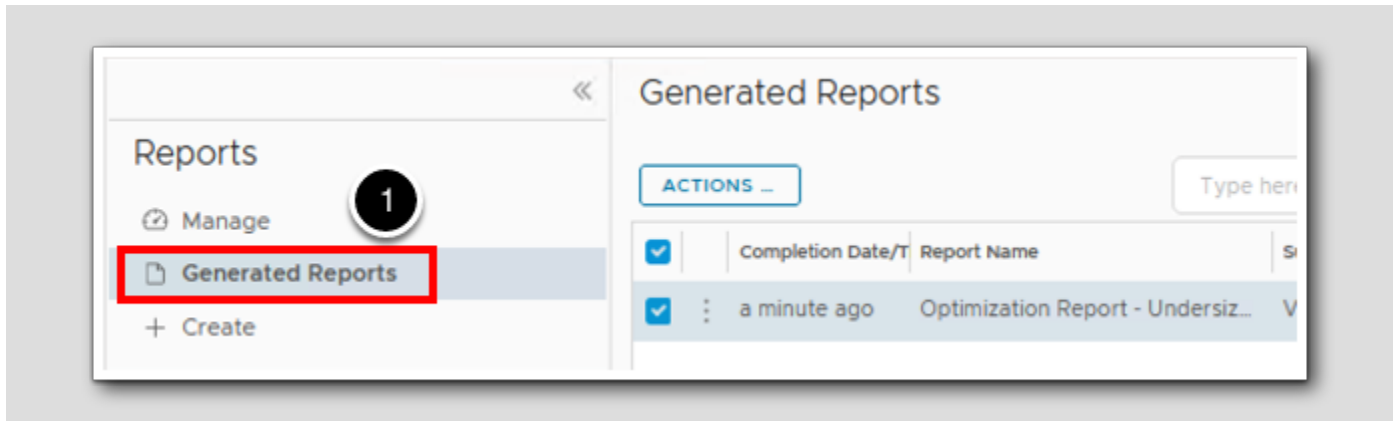
Steps 1-5 are expanding the Object list to show the hierarchy of Objects. If you are familiar with the hierarchy, skip to step 6.

1. Expand vSphere World.
2. Expand the vCenter, vcsa-01a.corp.vmbeans.com.
3. Expand the Datacenter, RegionA01.
4. Expand the Cluster, Management.
5. Expand the Host, esx-01a.corp.vmbeans.com.
6. Click vSphere World.
7. Click OK.

## The report ran, where did it go?

[191]


You are now returned to the Manage page and there is a subtle message next to the report that says running... and will disappear when the report is finished running. Many will expect the report to open by default and wonder what is happening. To see the report you just ran it is needed to switch to the Generated Reports tab.



1. Click on the Generated Reports tab.

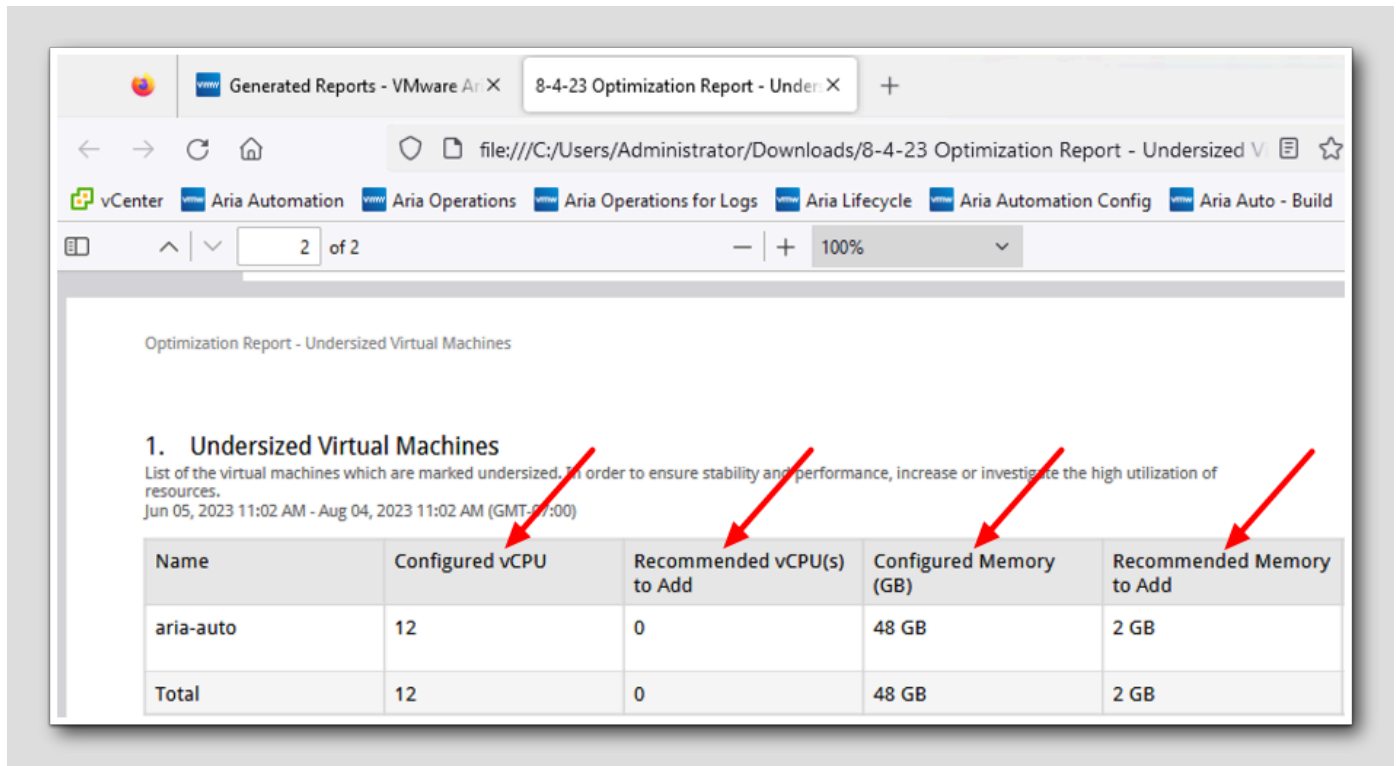
## Open the PDF

[192]



1. Click on the red PDF icon.

## Report Data



Optimization Report - Undersized Virtual Machines

### 1. Undersized Virtual Machines

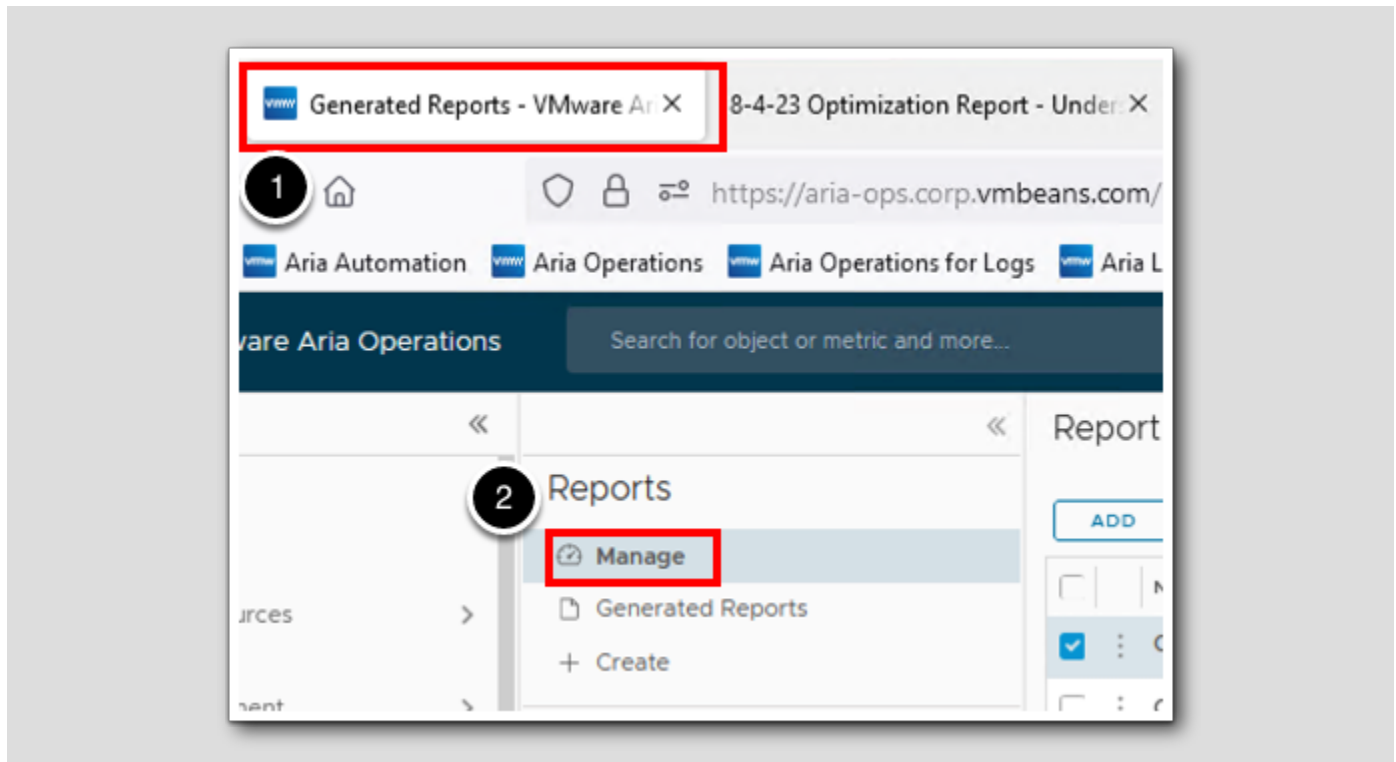
List of the virtual machines which are marked undersized. In order to ensure stability and performance, increase or investigate the high utilization of resources.  
Jun 05, 2023 11:02 AM - Aug 04, 2023 11:02 AM (GMT-07:00)

Name	Configured vCPU	Recommended vCPU(s) to Add	Configured Memory (GB)	Recommended Memory to Add
aria-auto	12	0	48 GB	2 GB
<b>Total</b>	<b>12</b>	<b>0</b>	<b>48 GB</b>	<b>2 GB</b>

The PDF should open in its own tab, let's take a look at the data. You will need to scroll down past the cover page. We only have one VM that made the Undersized report, aria-auto. It is configured with 12 virtual CPUs, the report is suggesting we don't need add any vCPUs so this means this VM is not CPU constrained. However, this VM has 48GB of virtual RAM and the report is suggesting that it gets another 2GB added to it. This indicates that this VM is Memory constrained.

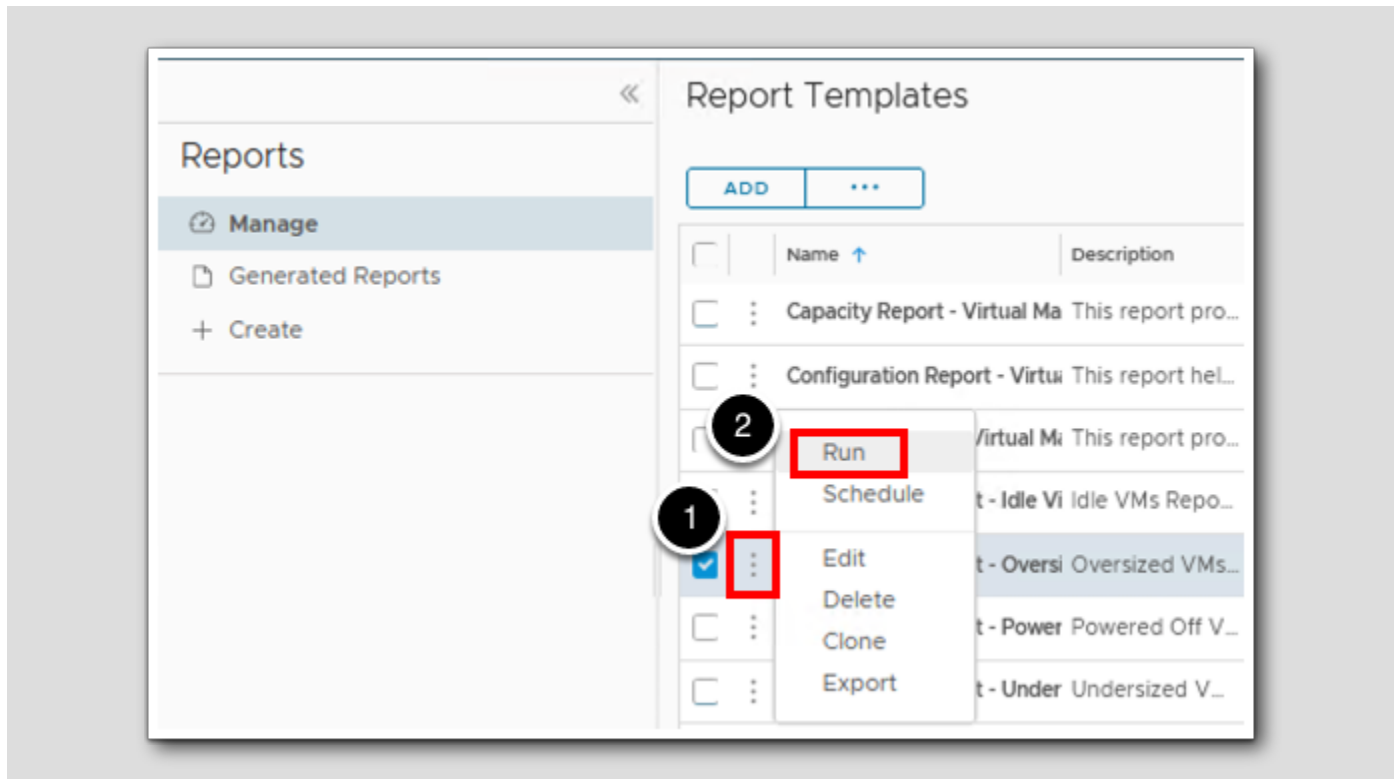
Now lets look at Oversized VMs

[194]



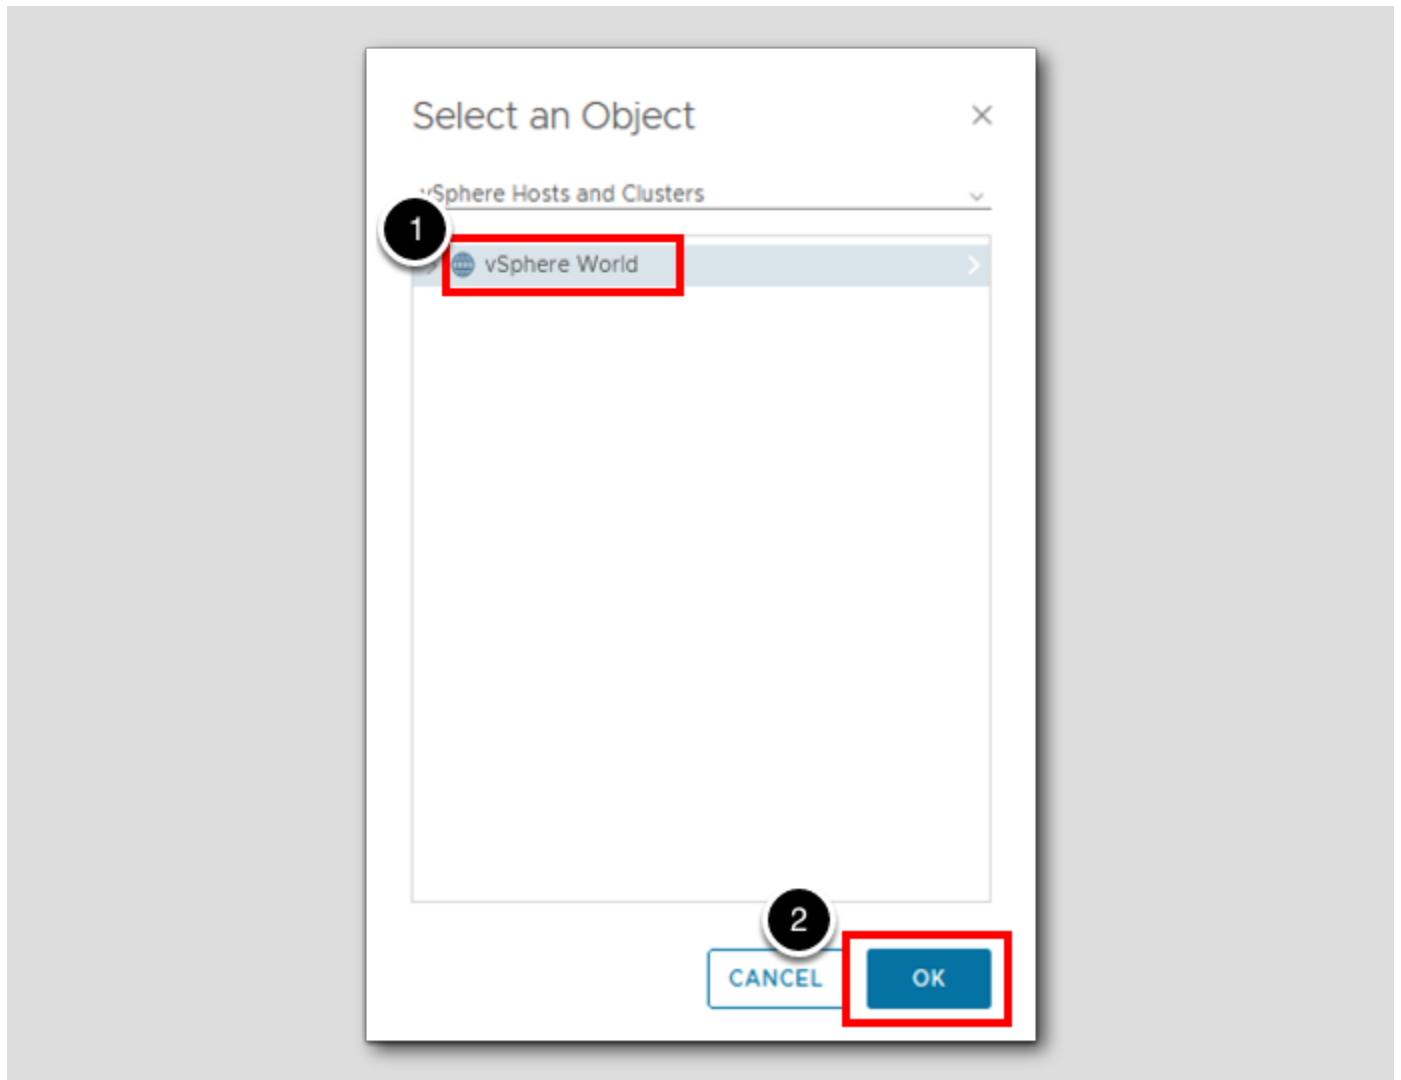
1. Click back into the Aria Operations browser tab.
2. Click on Manage.

## Run the Oversized report



1. Click on the vertical 3 dots next to the Optimization Report - Oversized Virtual Machines.
2. Select Run.

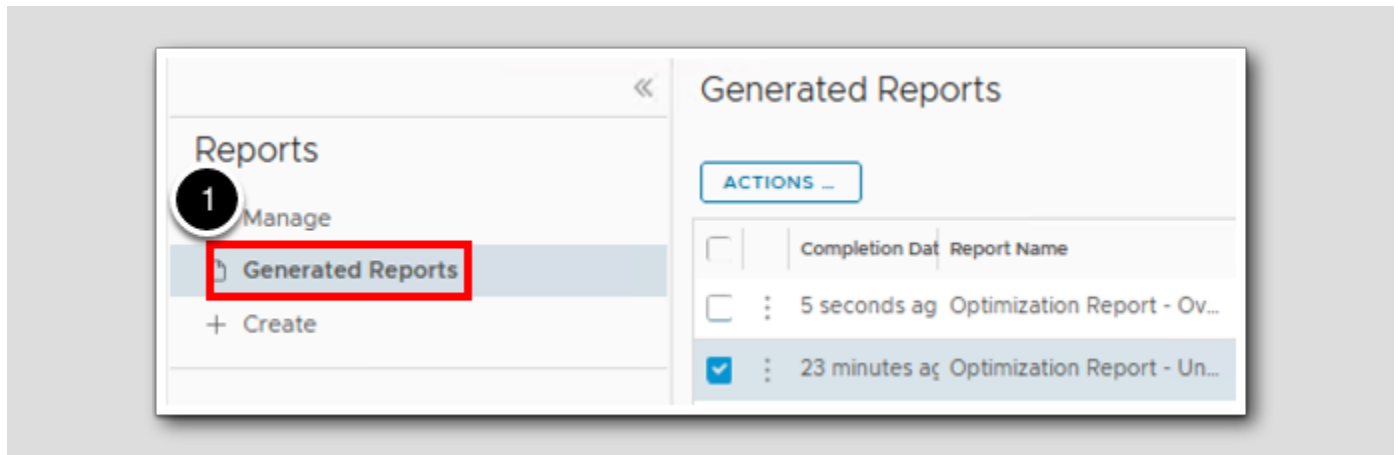
## Select vSphere World



1. Select vSphere World.
2. Click OK.

## Go to Generated Reports

[197]

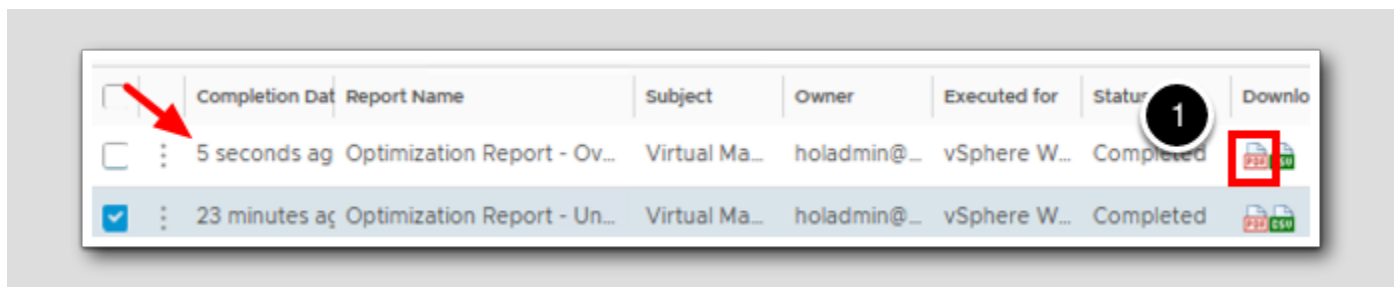


1. Click on the **Generated Reports** tab.

## Open the Oversized PDF report

[198]

In the Generated Reports tab, our first report may be highlighted by default. Please ensure to open the Oversized Report that was just run. You can see this by looking at the Completion Date/Time column.



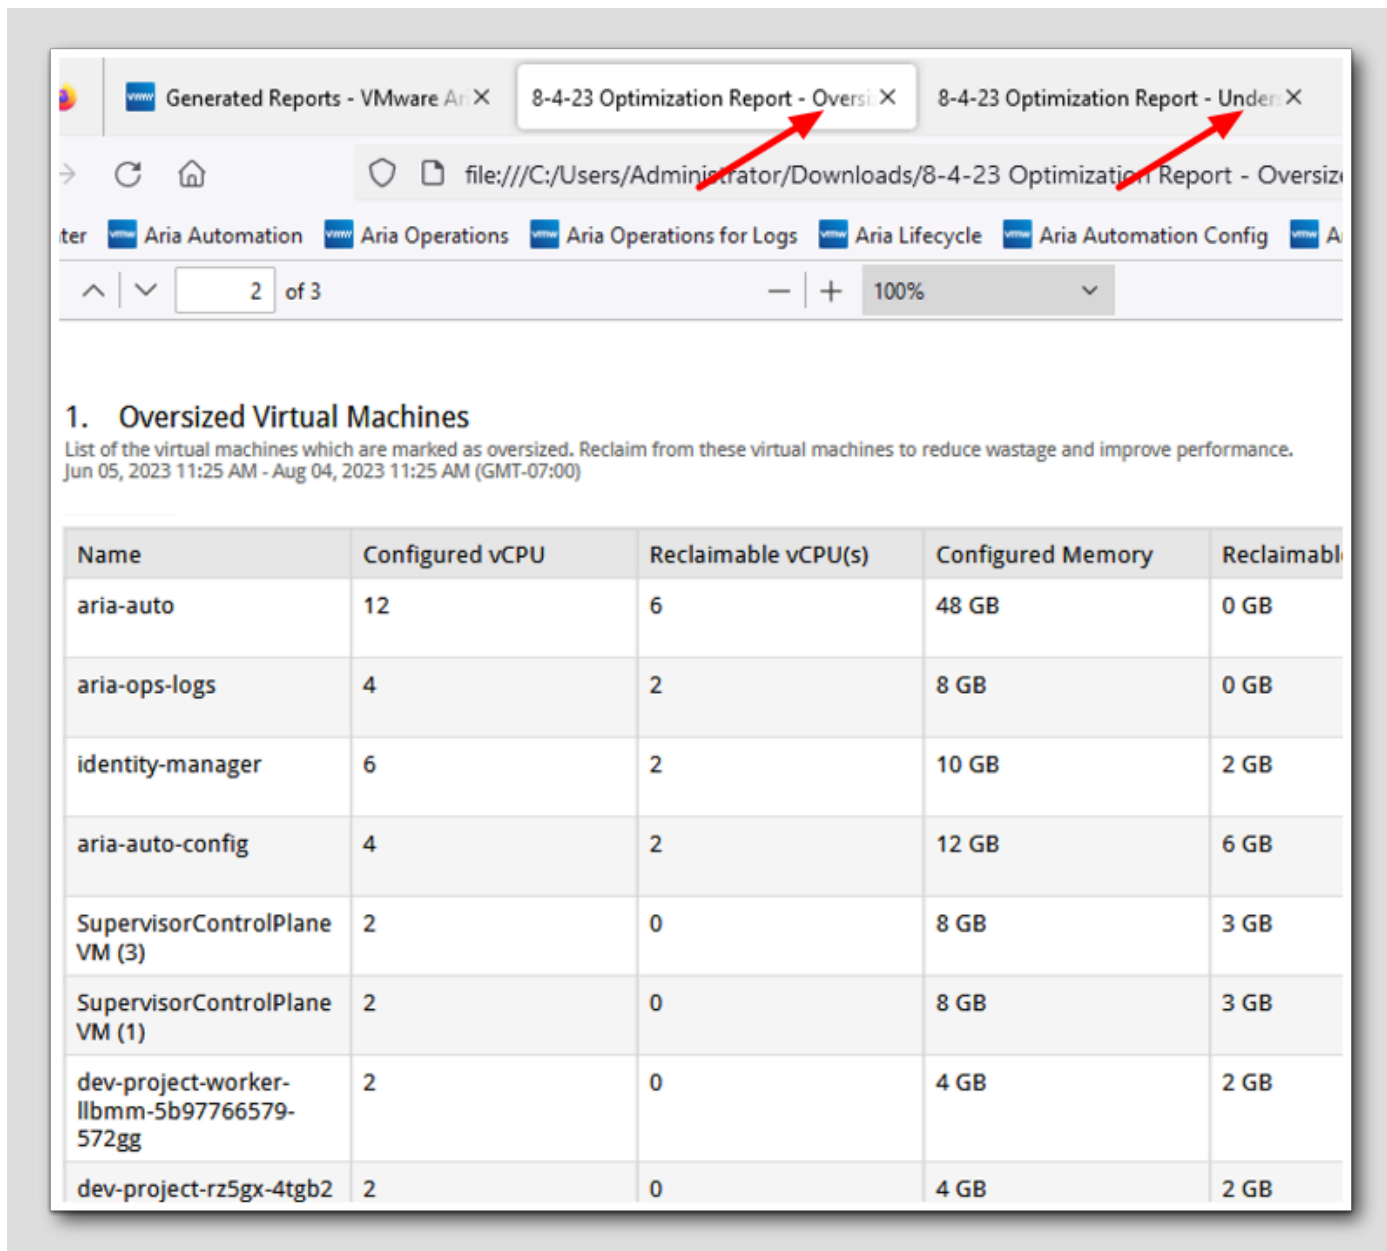
1. Click on the red PDF icon for the report we just ran.

## Many more VMs in this report

[199]

We should now have both the Undersized and Oversized reports open in two different tabs. Let's dive into this data starting out with, why is aria-auto in both reports. You will need to scroll down past the cover page again.





A VM that is in both undersized and oversized report

[200]

Below are isolated screenshots of the aria-auto VM showing up in both reports. This highlights that a VM can be oversized in either vCPU or vMemory and yet be undersized in its other resource of vMemory or vCPU. Below shows that aria-auto is oversized with vCPUs by 6 vCPUs and at the same time undersized with the amount of vMemory by 2GB.

1. Oversized Virtual Machines					
List of the virtual machines which are marked as oversized. Reclaim from these virtual machines to reduce wastage and improve performance. Jun 05, 2023 11:25 AM - Aug 04, 2023 11:25 AM (GMT-07:00)					
Name	Configured vCPU	Reclaimable vCPU(s)	Configured Memory	Reclaimable Memory	Parent vCenter
aria-auto	12	6	48 GB	0 GB	vcsa-01a.corp.vmbeans.com

1. Undersized Virtual Machines					
List of the virtual machines which are marked undersized. In order to ensure stability and performance, increase or investigate the high utilization of resources. Jun 05, 2023 11:02 AM - Aug 04, 2023 11:02 AM (GMT-07:00)					
Name	Configured vCPU	Recommended vCPU(s) to Add	Configured Memory (GB)	Recommended Memory to Add	Parent vCenter
aria-auto	12	0	48 GB	2 GB	vcsa-01a.corp.vmbeans.com
Total	12	0	48 GB	2 GB	-

Data in the moment vs your knowledge of your infrastructure.

[201]

Looking at the remaining data in this report we can see that there are places that we can reclaim vCPUs on some VMs and reclaim vMemory on others and even some VMs that we can reclaim both vCPUs and vMemory. Aria Operations will report on the historical performance of each VM underneath the chosen object (vSphere World in this case) however it cannot bring knowledge of the business to bare. For example, your company has a finance VM that runs numbers, and is stressed, once a quarter. There is a high chance that the Oversized Report will flag this VM for reclamation due to it being idle for most of the quarter. This is an example where knowledge of your environment combined with data from Aria Operations will be critical to dial in capacity efficiency.

The screenshot shows a web browser window displaying a VMware Aria Operations report. The report title is "8-4-23 Optimization Report - Oversized Virtual Machines". Below the title, there is a sub-header "1. Oversized Virtual Machines" and a brief description: "List of the virtual machines which are marked as oversized. Reclaim from these virtual machines to reduce wastage and improve performance." The report is dated from Jun 05, 2023 11:25 AM to Aug 04, 2023 11:25 AM (GMT-07:00).

Name	Configured vCPU	Reclaimable vCPU(s)	Configured Memory	Reclaimable Memory
aria-auto	12	6	48 GB	0 GB
aria-ops-logs	4	2	8 GB	0 GB
identity-manager	6	2	10 GB	2 GB
aria-auto-config	4	2	12 GB	6 GB
SupervisorControlPlane VM (3)	2	0	8 GB	3 GB
SupervisorControlPlane VM (1)	2	0	8 GB	3 GB
dev-project-worker-llbmm-5b97766579-572gg	2	0	4 GB	2 GB
dev-project-rz5gx-4tgb2	2	0	4 GB	2 GB

Lesson End

[202]

This lesson highlighted how to run and view reports and how to analyze the data Aria Operations provides.

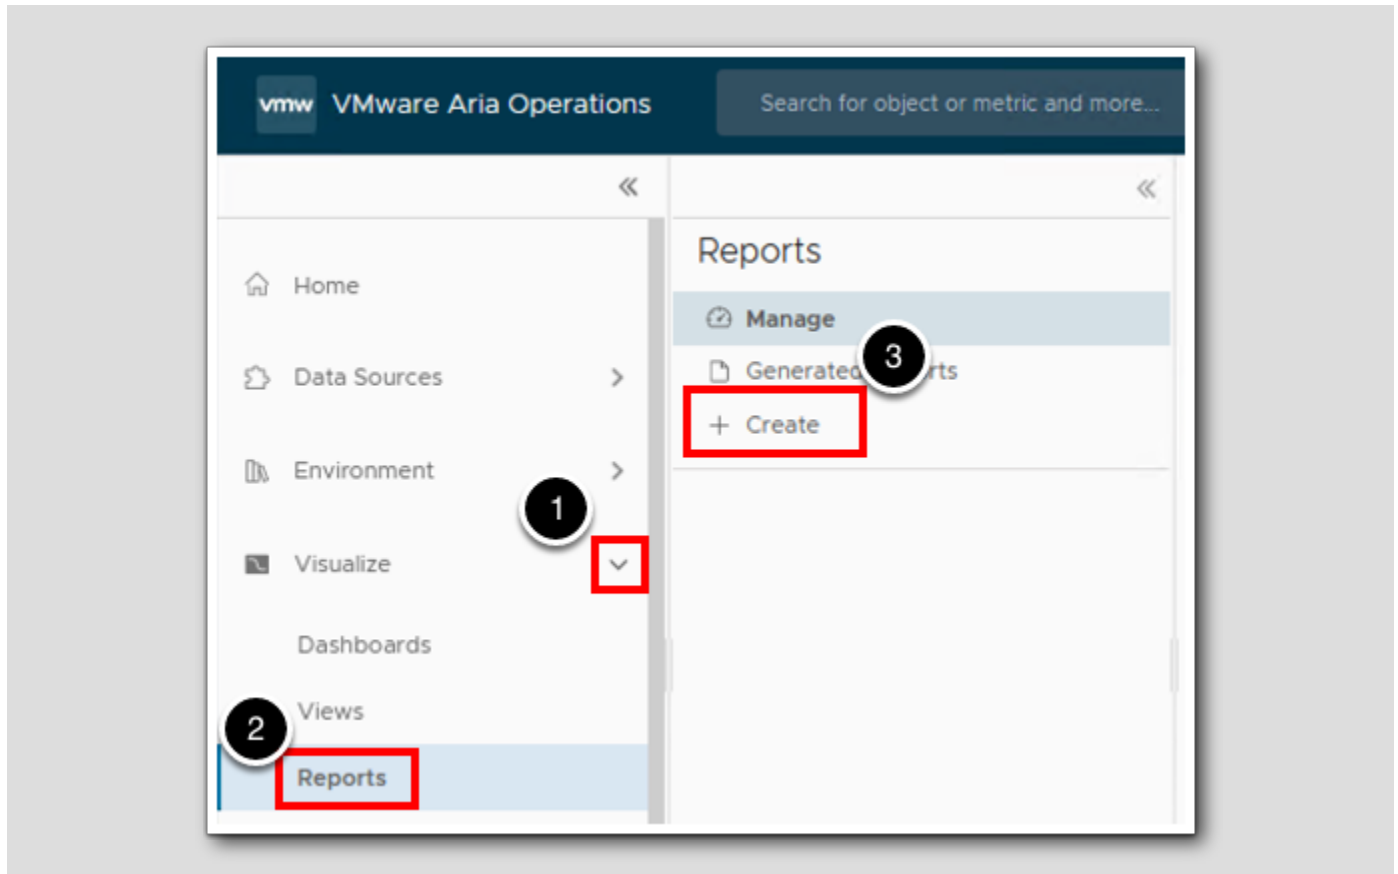
## Creating Custom Reports

[203]

In this lesson we will create a custom report and show how to setup the Standard Email Plugin so we can distribute the report

## Open Reports

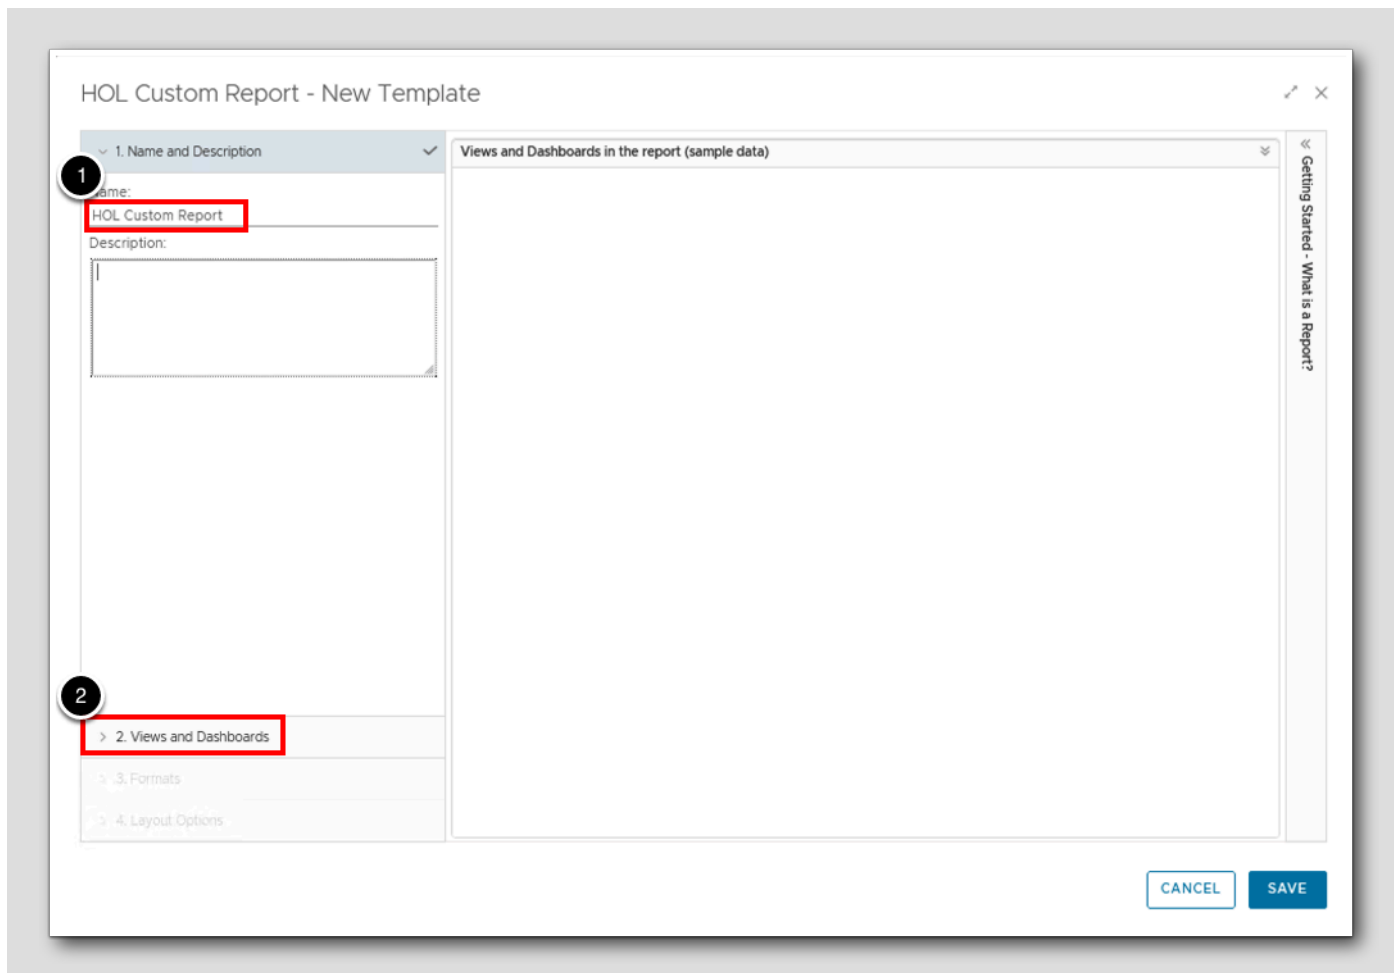
[204]



1. Expand Visualize.
2. Click Reports.
3. Click + Create.

## Name the Report

[205]

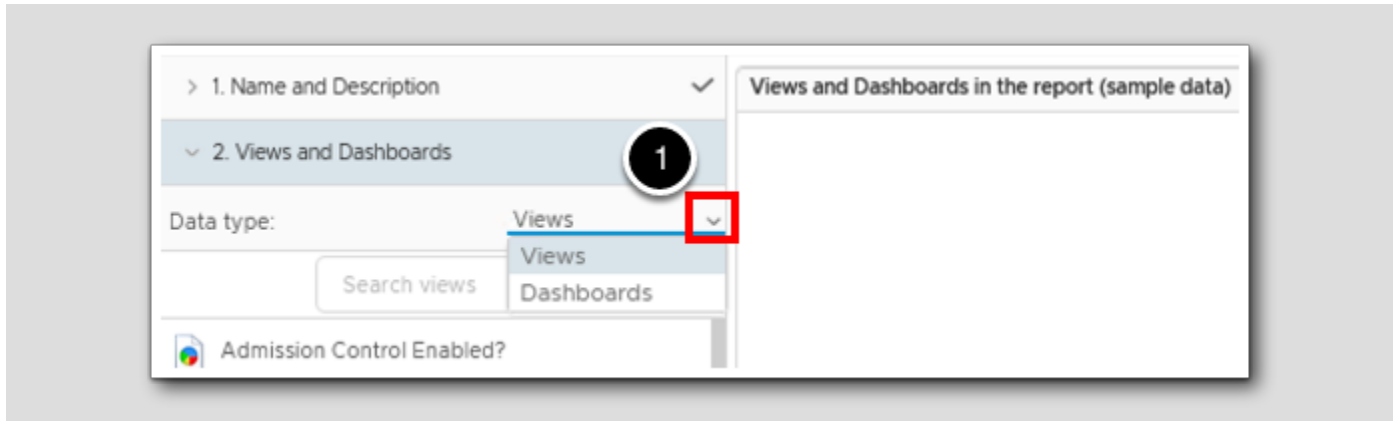


1. In the **Name:** field type HOL Custom Report.
2. Click **2. Views and Dashboards**.

## Data type

[206]

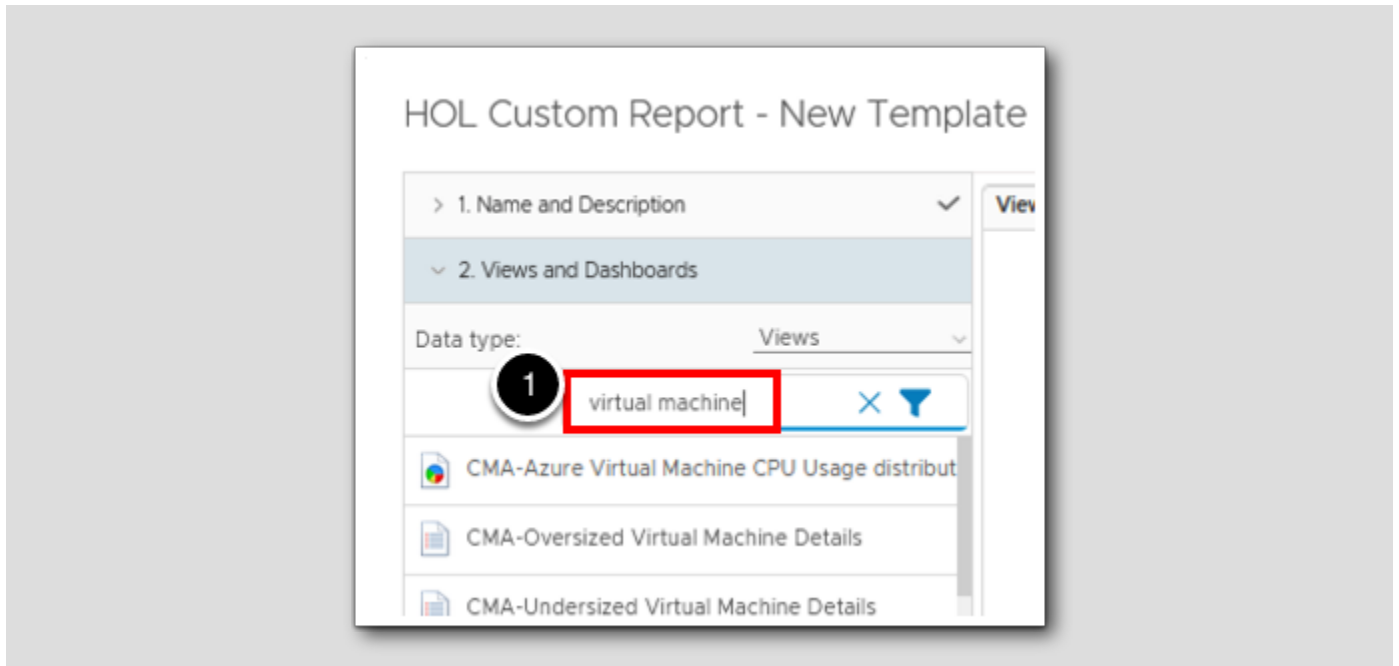
In the **2. Views and Dashboards** tab you can choose a Data type of either Views or Dashboards to build out your custom report.



1. Click on the > for the **Data type:** field. Notice this is where you can select to build a report with Views or with Dashboards. For this lesson we will use Views.

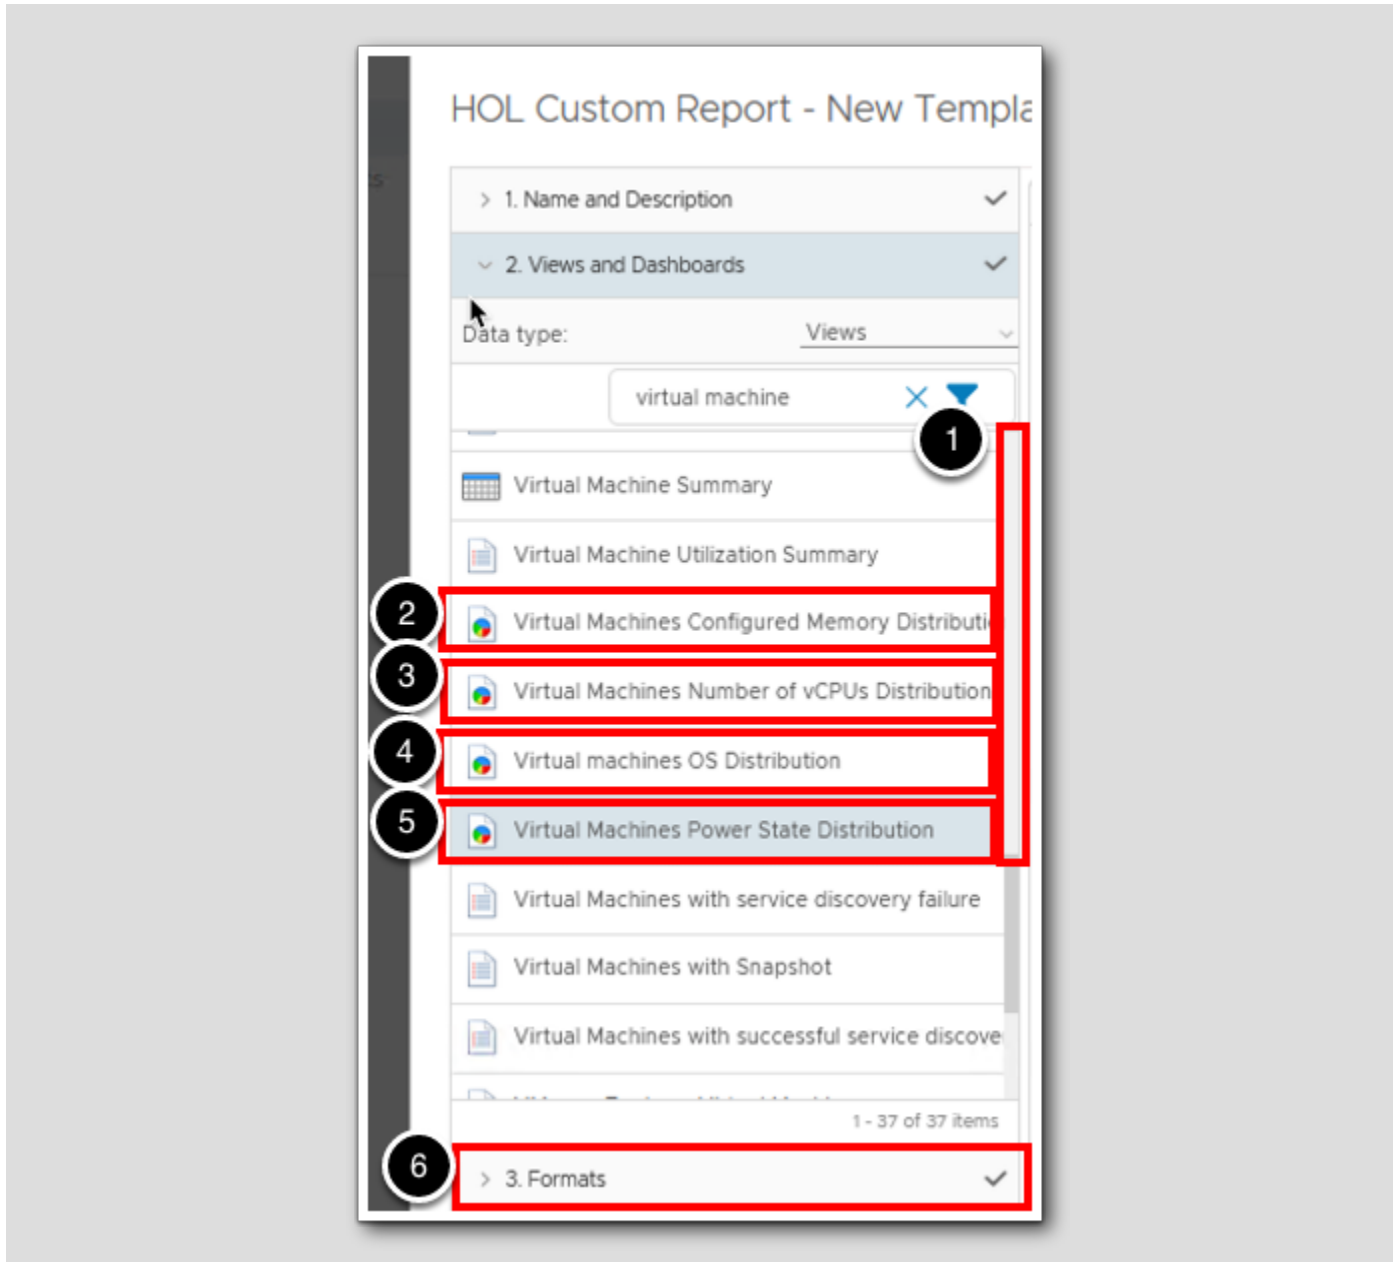
### Filter for Virtual Machines

[207]



1. Click into the filter search bar, type virtual machine and hit Enter.

### Distribution Pie Graphs



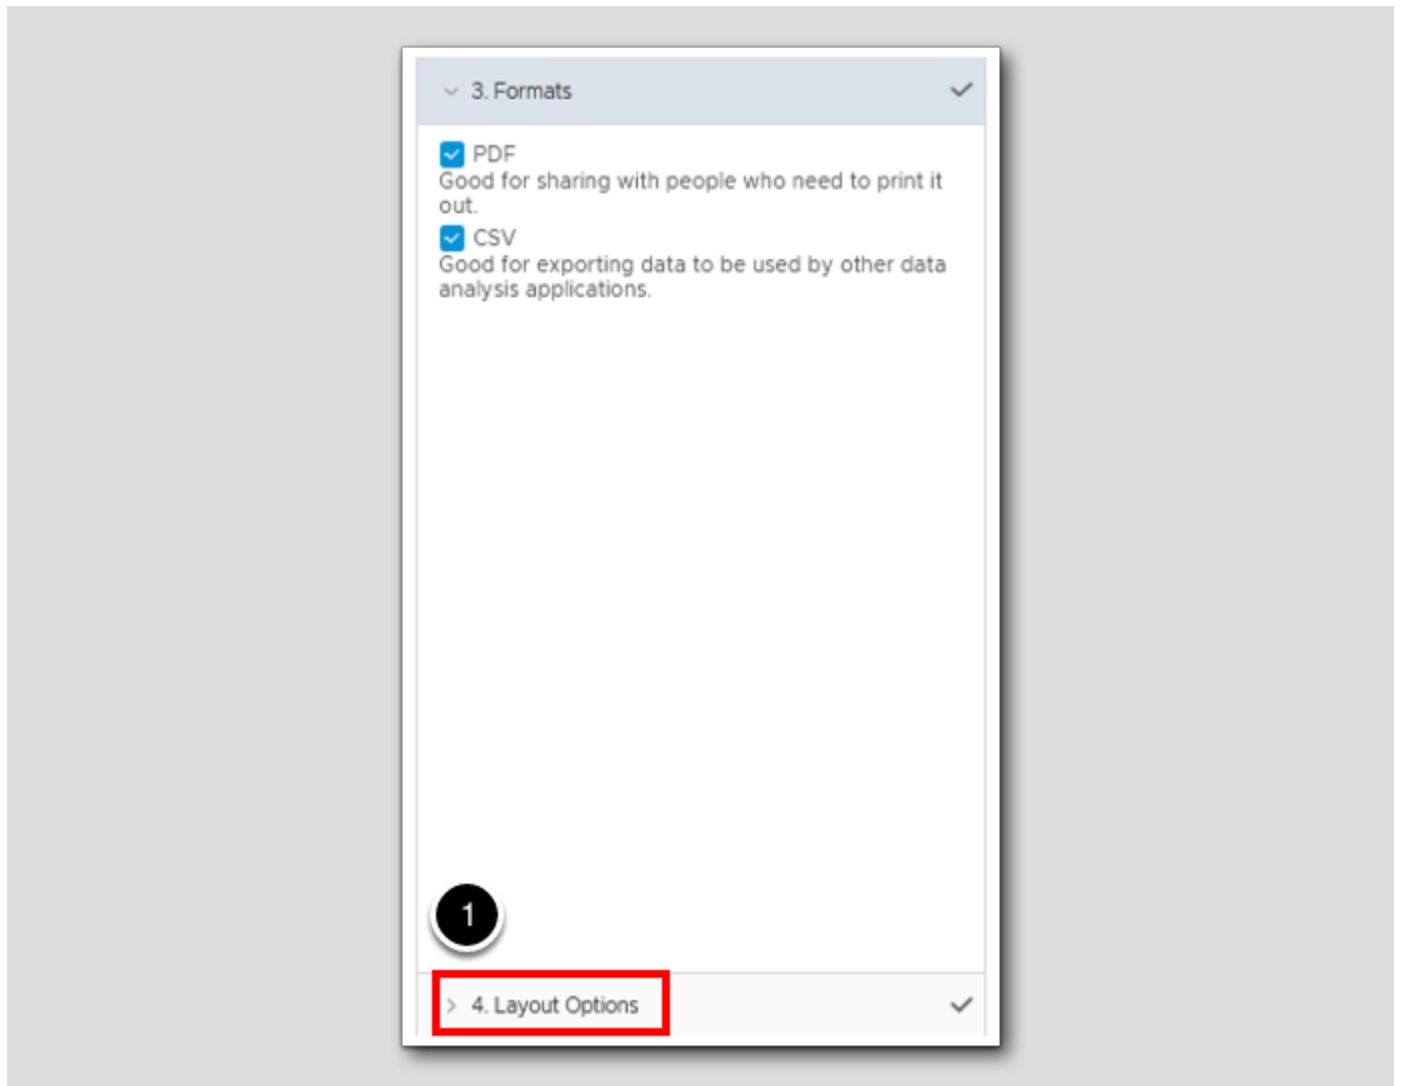
1. Scroll down until you see the 4 Distribution pie graph views shown above.
2. Double click on **Virtual Machines Configured Memory Distribution**.
3. Double click on **Virtual Machines Number of vCPUs Distribution**.
4. Double click on **Virtual Machines OS Distribution**.
5. Double click on **Virtual Machines Power State Distribution**.
6. Click on **3. Formats**.

## Formats

[209]

In the 3. Formats tab you have the ability to chose which export formats you would like to have for the customer reports. The choices are PDF and/or CSV. For this lesson we will leave both selected.



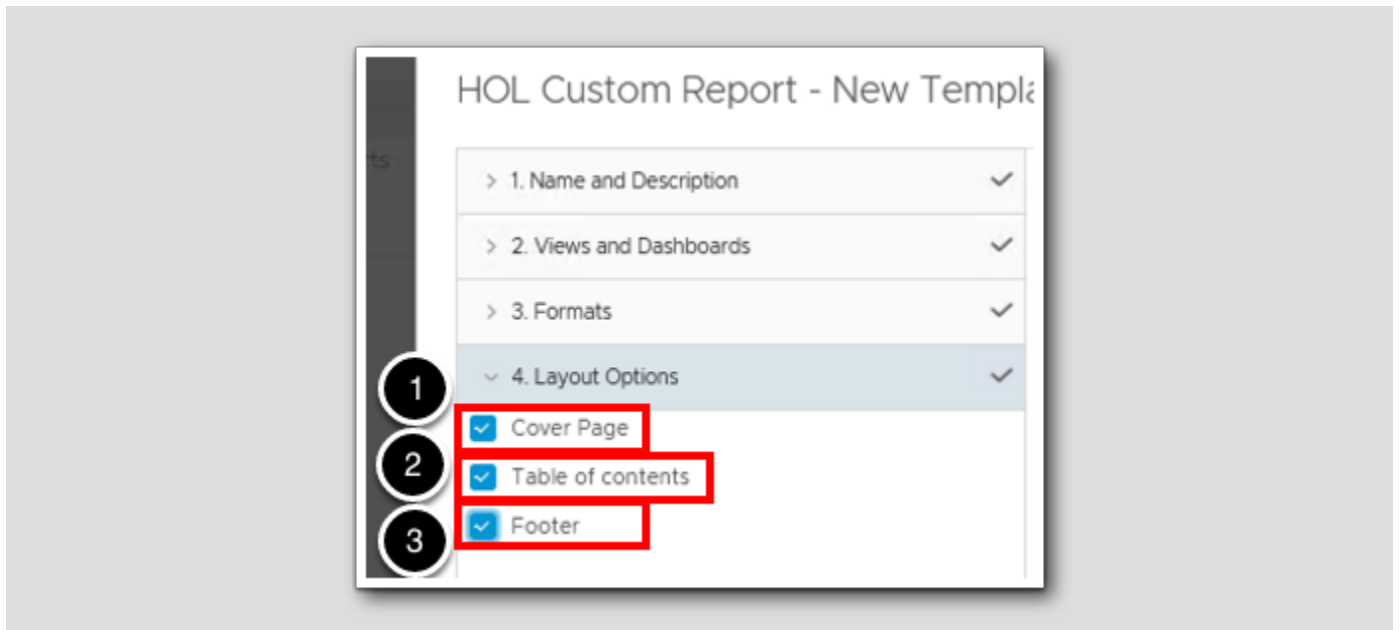


1. Select 4. Layout Options.

## Layout Options

[210]


In the 4. Layout Options you have the ability to add a Cover Page with the option to upload an image, a Table of contents and a Footer.



1. Click on the box for Cover Page.
2. Click on the box for Table of contents.
3. Click on the box for Footer.

## Save the Report


[21]



1. Click SAVE.

Lets take a look at what we just built

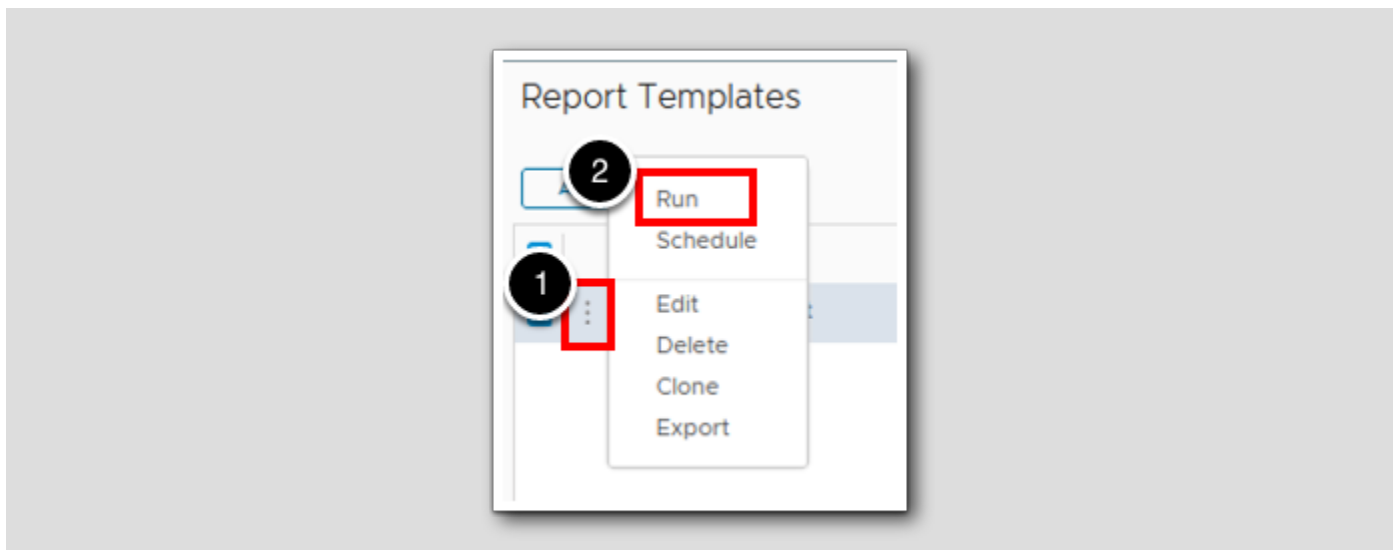
[212]



1. In the top right search bar type hol custom report and hit Enter.

Run the Report

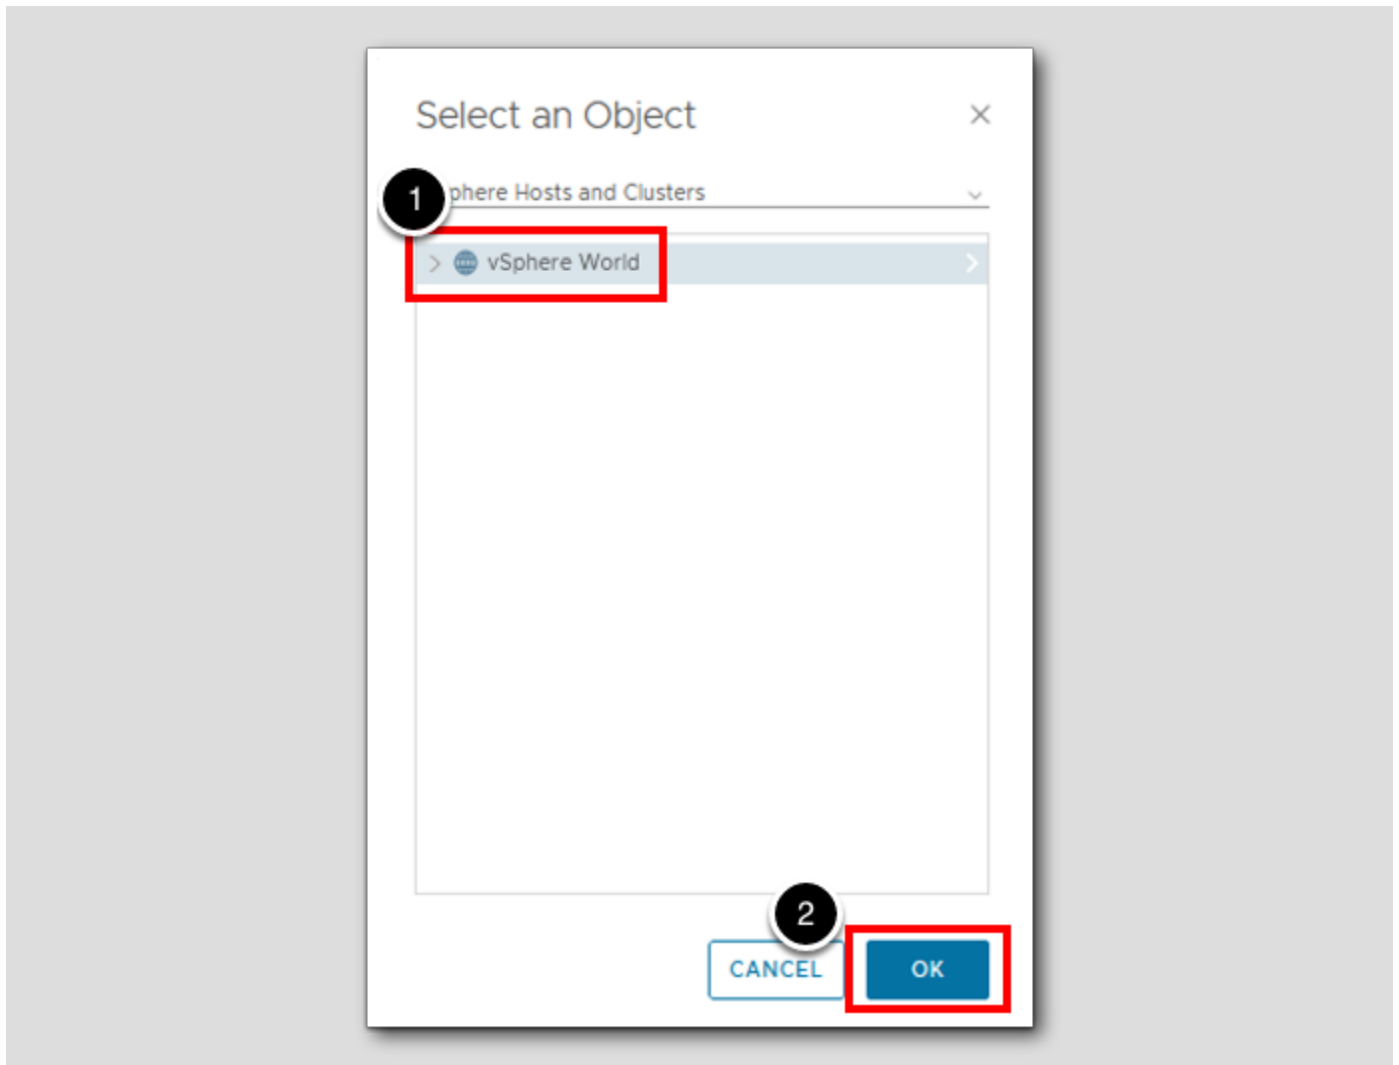
[213]



1. Click the 3 vertical dots next to the HOL Custom Report name.
2. Click Run.

## Select an Object

[214]

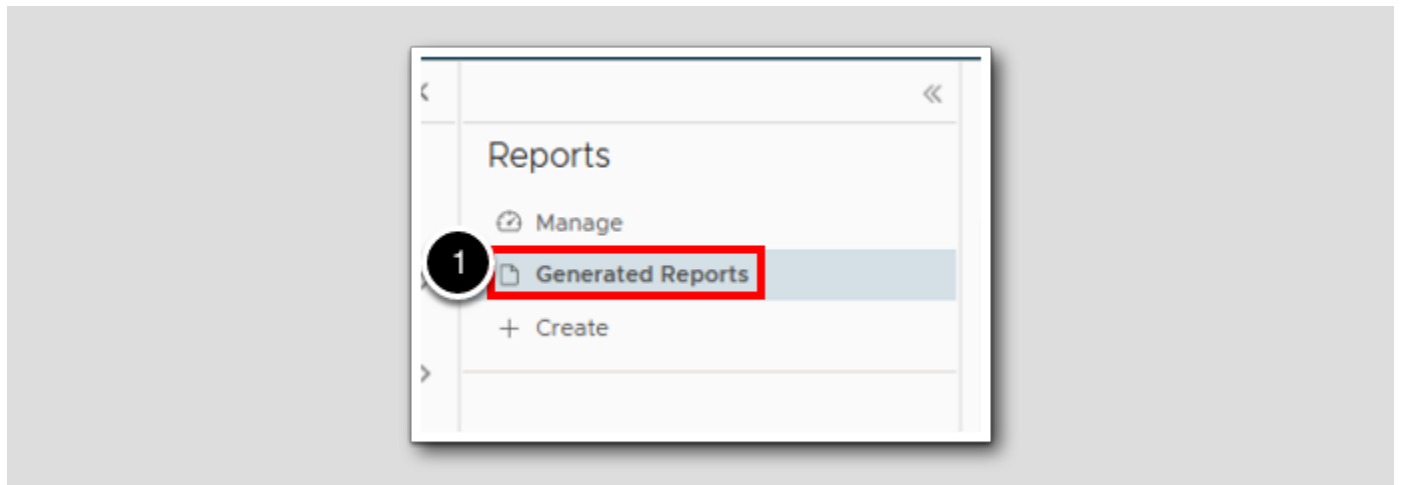


1. Click on vSphere World.
2. Click on OK.

## Generated Reports

[215]


This report may take longer to run than the Undersized and Oversized reports.



1. Click on Generated Reports.

## PDF

[216]

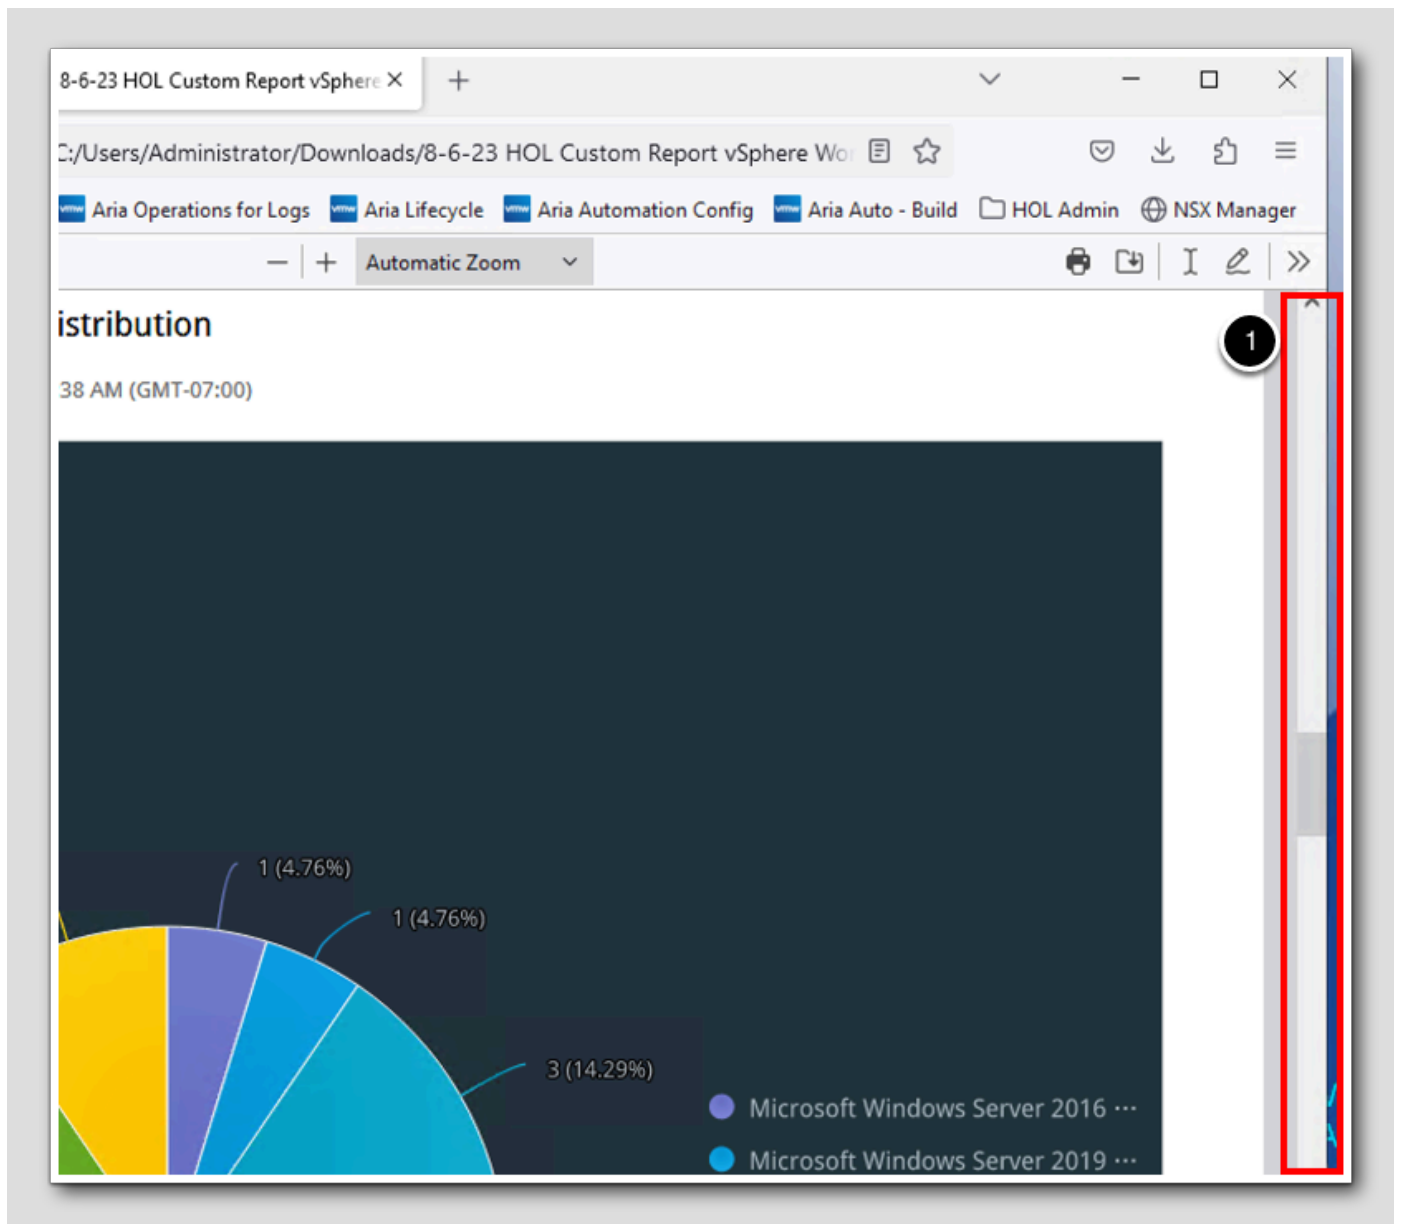


1. Click on the red PDF icon to open the report in a new browser tab.

## That was easy

[217]

As you can see it is very easy to create a custom report in Aria Operations with the help of using Views and Dashboards. With just a few clicks, we've created a virtual machine distribution report for OS, Memory, vCPUs and Power State.



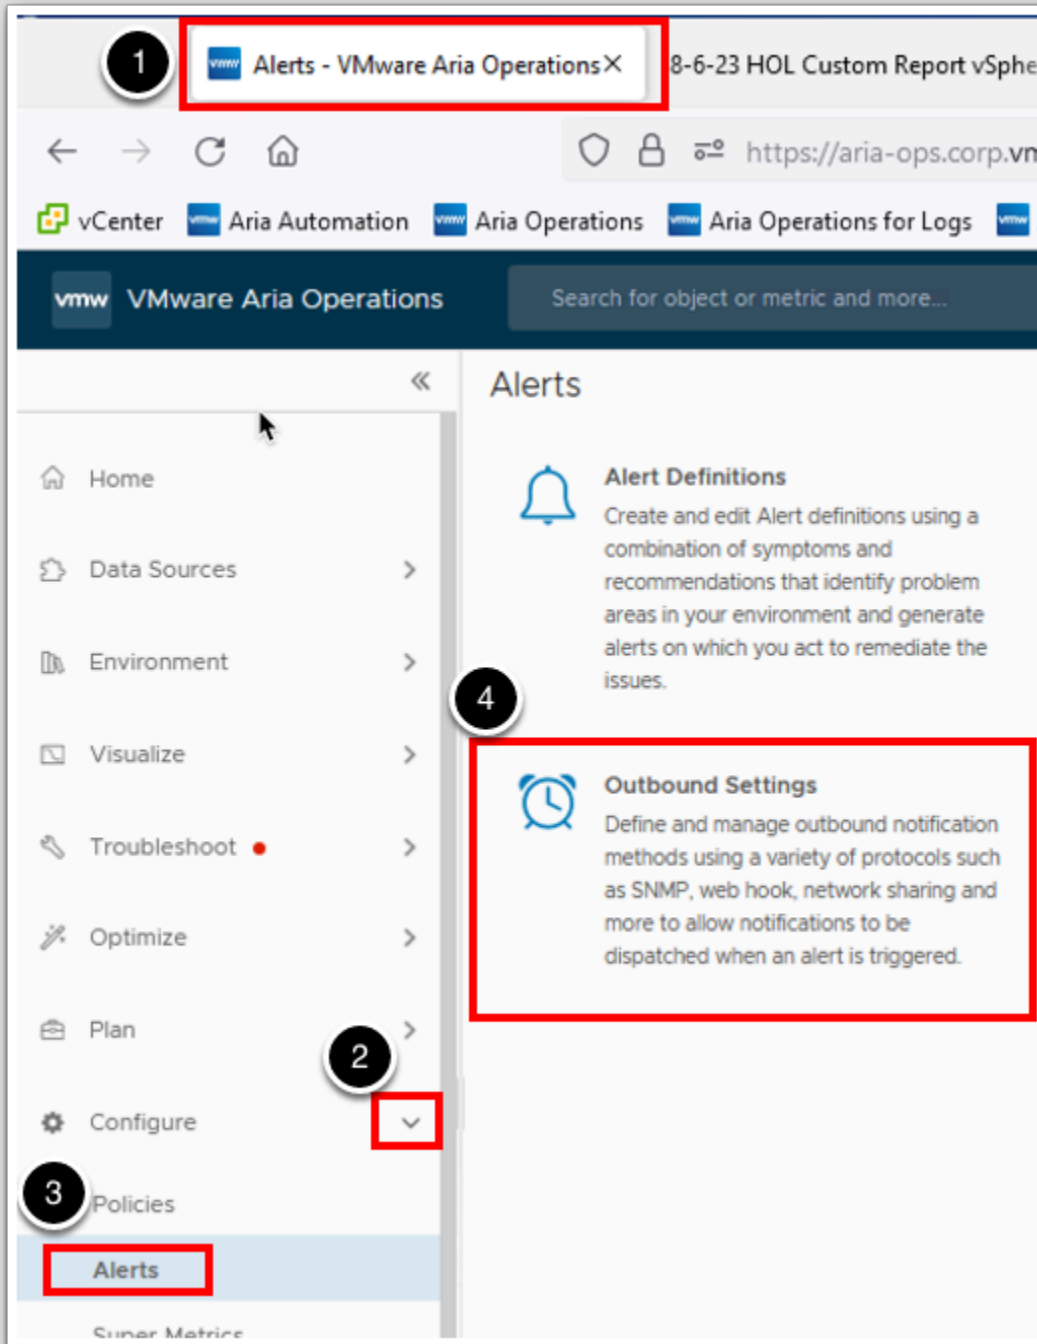
1. Scroll down on the right side of the new browser tab that opened by default to review the report results.

Notice that the Cover Page, Table of contents and that each page has a Footer.

## Configure SMTP Outbound

[218]

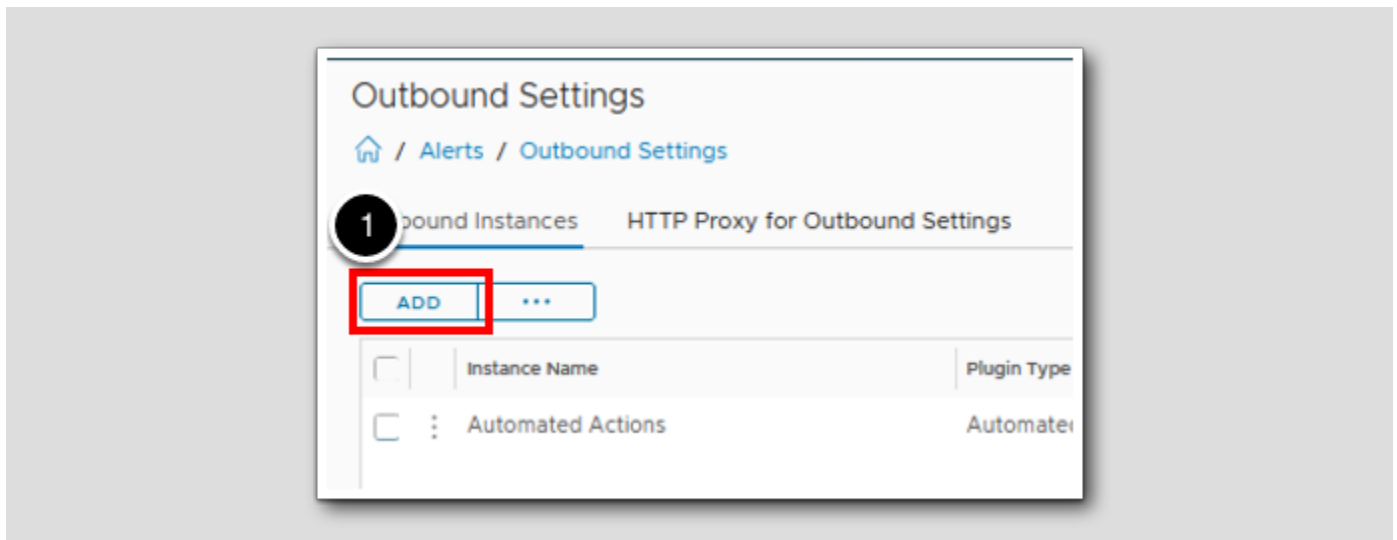
To have the ability to email reports out we need to setup a Standard Email Plugin.



1. Click on the Aria Operations browser tab.
2. Expand Configure.
3. Click on Alerts.
4. Click on Outbound Settings.

## Add an Outbound Instance

[219]

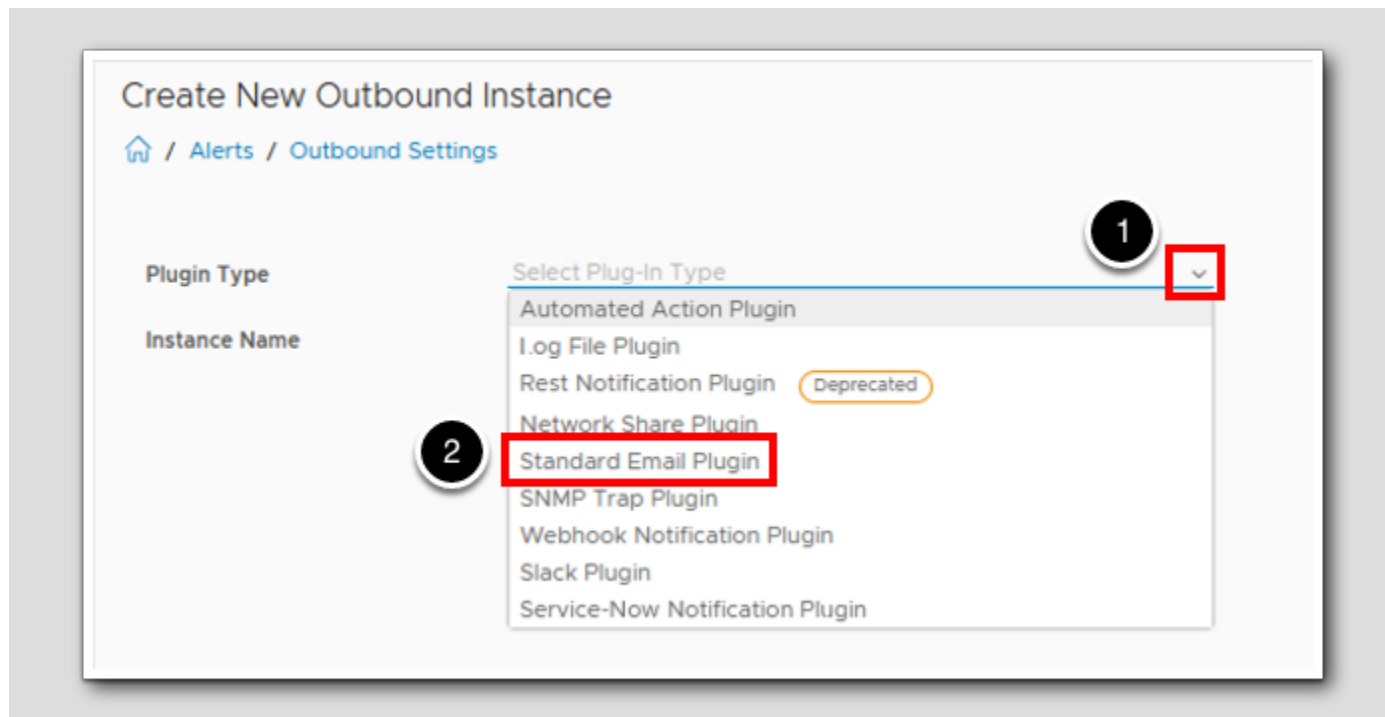


1. Click ADD.



## Standard Email Plugin

[220]



1. Expand the Plugin Type dropdown.
2. Click on Standard Email Plugin.

## SMTP Settings

[221]

These values are specific to the HOL environment. For your organization, you will need to gather your SMTP server details and requirements.

## Create New Outbound Instance

Home / Alerts / Outbound Settings

Plugin Type	1	Standard Email Plugin
Instance Name	2	HOL Email
Use Secure Connection	3	<input checked="" type="checkbox"/>
Requires Authentication	4	<input type="checkbox"/>
SMTP Host	5	mail.corp.vmbeans.com
SMTP Port	6	25
Secure Connection Type	7	
Sender Email Address	8	administrator@corp.vmbeans.com
Sender Name	9	Administrator
Credential type	10	No Credential
Receiver Email Address	11	holadmin@corp.vmbeans.com

8

TEST

SAVE

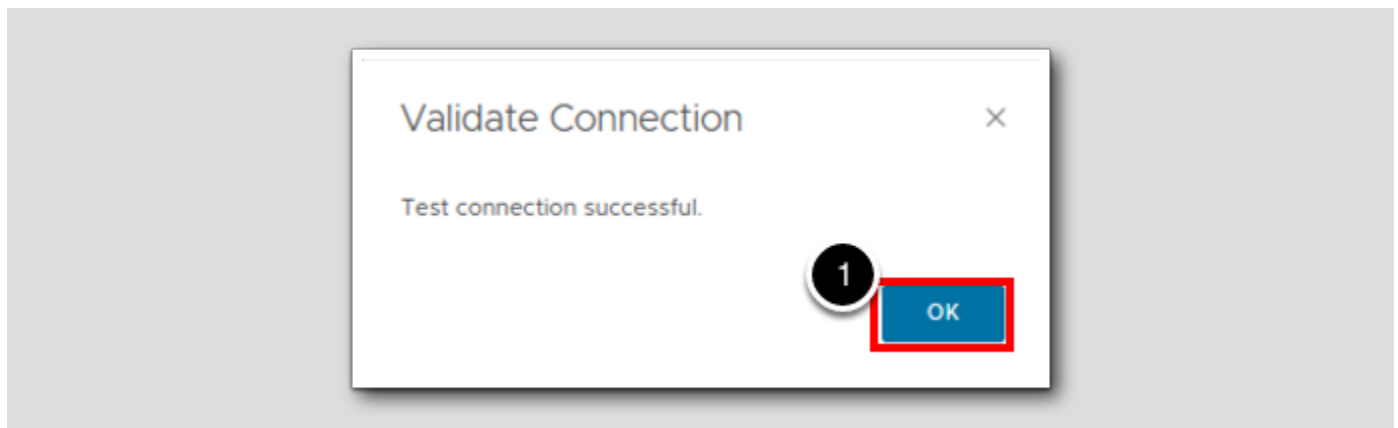
CANCEL

1. Instance Name - HOL Email
2. Check Use Secure Connection
3. SMTP Host - mail.corp.vmbeans.com
4. SMTP Port - 25
5. Sender Email Address - administrator@corp.vmbeans.com
6. Sender Name - Administrator
7. Receiver Email Address - holadmin@corp.vmbeans.com
8. Click TEST and ensure you get a Test connection successful response (Not Shown).

## Validate Connection

[222]

Hopefully you see the screenshot below, if not then re-verify that the above data has been entered correctly and re=test.



1. Click OK.

Save

[223]

### Create New Outbound Instance

[Home](#) / [Alerts](#) / [Outbound Settings](#)

Plugin Type	Standard Email Plugin
Instance Name	HOL Email
Use Secure Connection	<input checked="" type="checkbox"/>
Requires Authentication	<input type="checkbox"/>
SMTP Host	mail.corp.vmbeans.com
SMTP Port	25
Secure Connection Type	<input type="text"/> x
Sender Email Address	administrator@corp.vmbeans.com
Sender Name	Administrator
Credential type	No Credential
Receiver Email Address	holadmin@corp.vmbeans.com

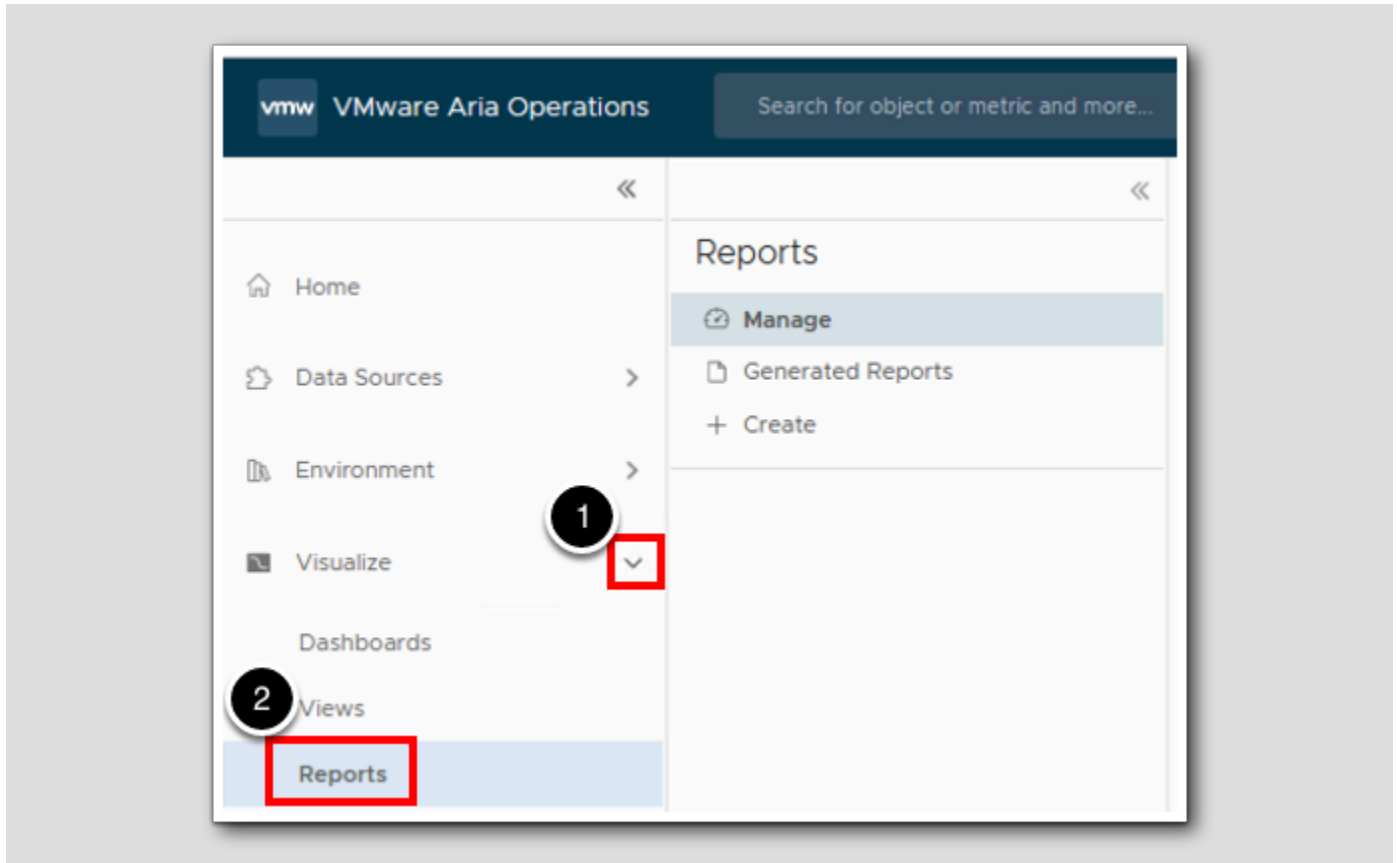
1

TEST SAVE CANCEL

1. Click SAVE.

## Return to Reports

[224]

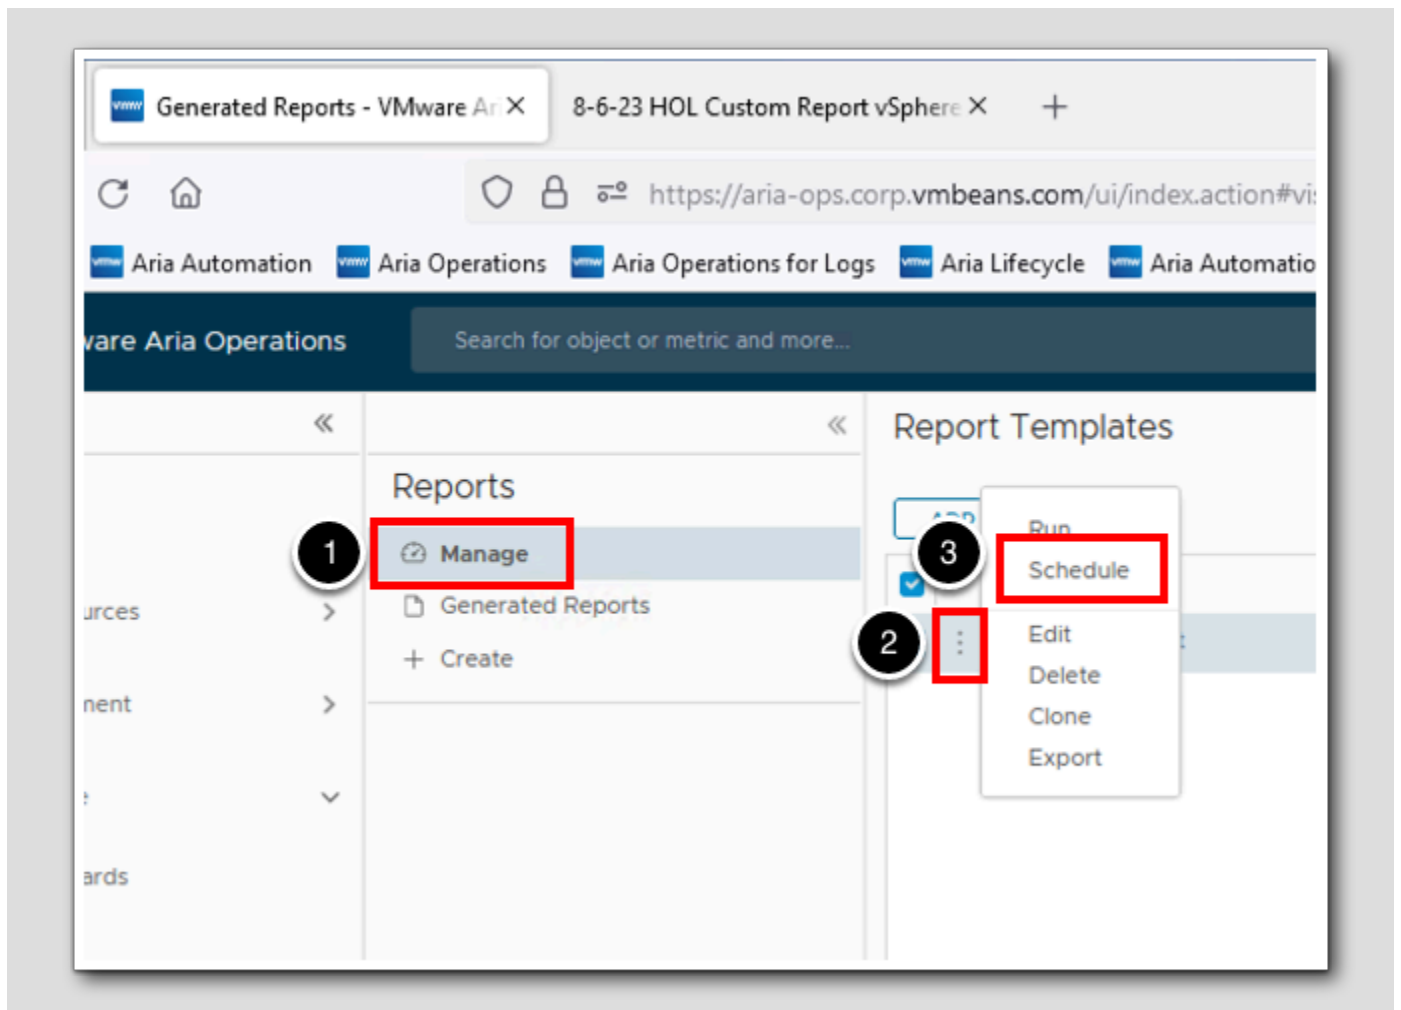


1. Expand Visualize.
2. Click on Reports.

## Scheduling a Report

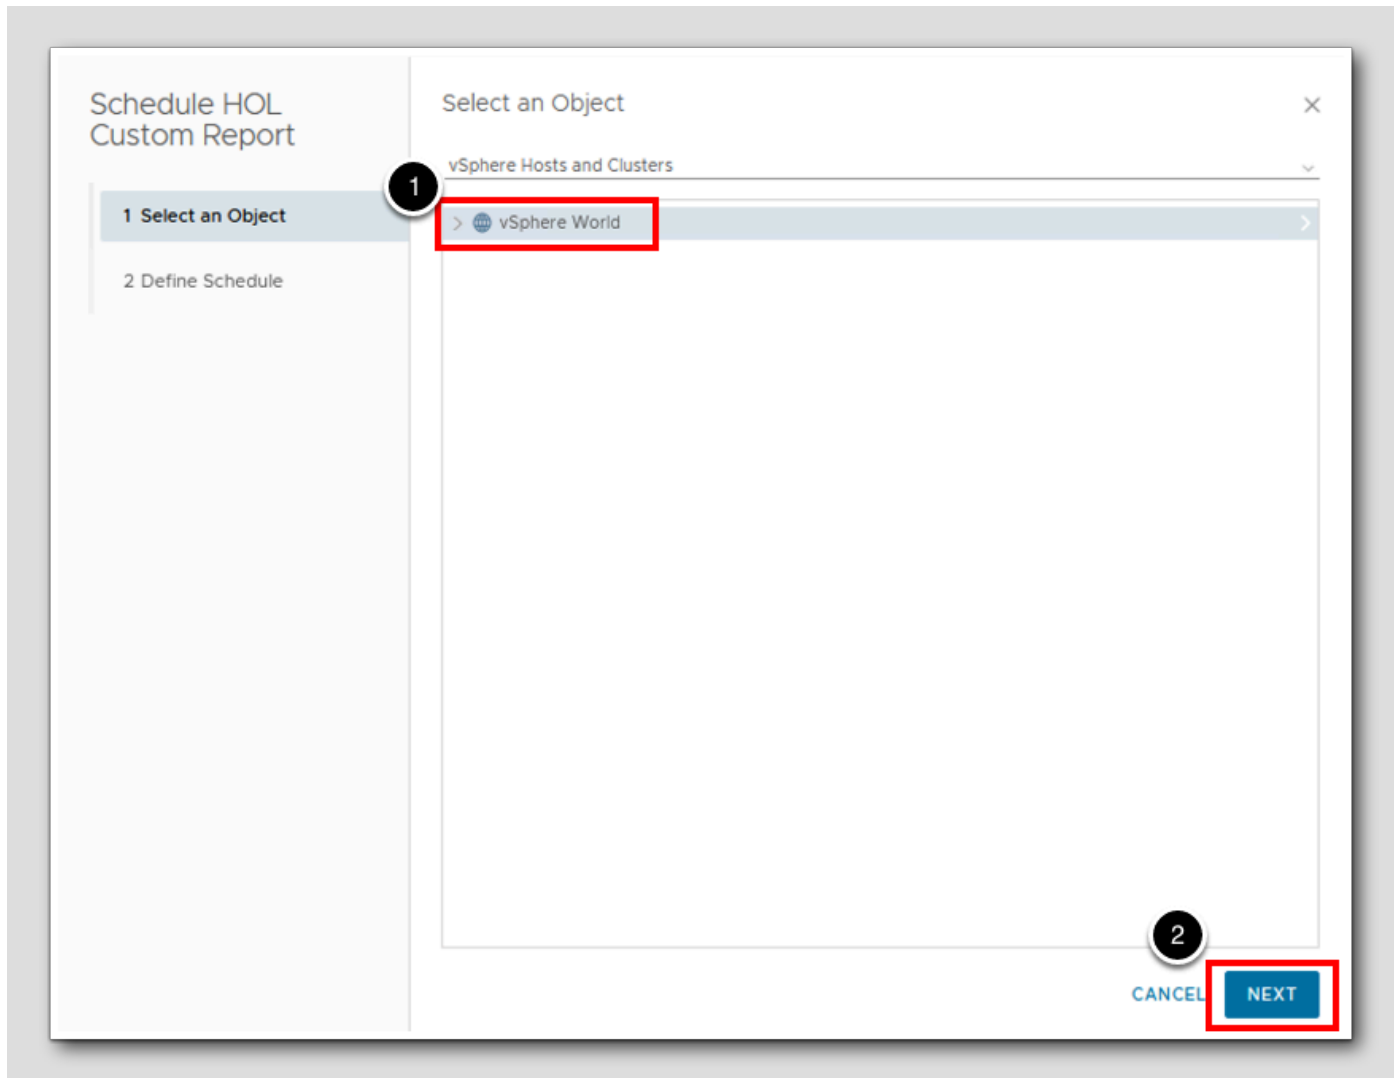
[225]

The hol custom report filter should still be active. If it is not, research for hol custom report and hit enter.



1. Click on Manage.
2. Click on the 3 vertical dots next to the HOL Custom Report.
3. Click on Schedule.

## Define the Object level



1. Click on the vSphere World object.
2. Click NEXT.

## Define Schedule

Here we can see how to schedule a report to run Daily, Weekly or Monthly. The Recurrence fields will change with what is selected in this dropdown.

### Schedule HOL Custom Report

- 1 Select an Object
- 2 Define Schedule**

### Define Schedule

Set the recurrence and publishing criteria for this report

**Recurrence**

Time zone: --Select--

Start hour: 8 00 AM

Start date: 8/7/23

Recurrence: Weekly

Monday  Tuesday  Wednesday  Thursday  Friday  Saturday

**Publishing**

Email report

Email addresses:  Cc Bcc

Select an outbound rule: --Select--

**⚠** There are no external locations defined, [click here](#) to configure a new external location.

Save to external location

Select a location: --Select--

Relative Path:

CANCEL BACK FINISH

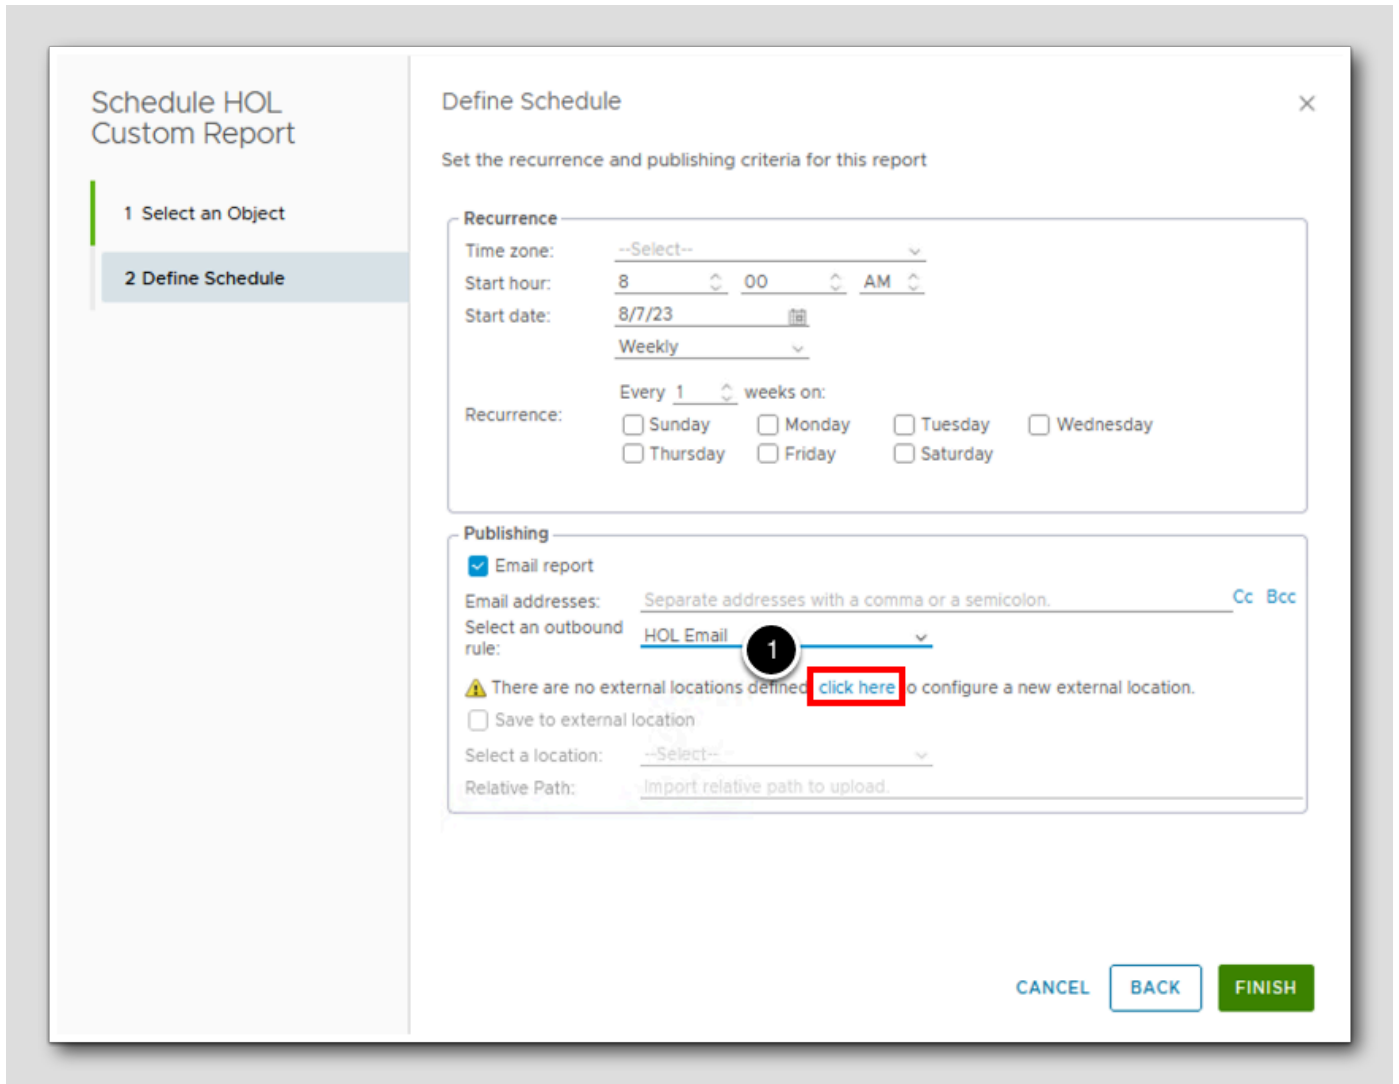


## Publishing

1. In the **Publishing** box, click the **Email report** checkbox.
2. Expand the **Select an Outbound rule:** dropdown.
3. Click on the newly created outbound rule **HOL Email**.

## Add Outbound Instance

If you would prefer to have the reports saved to a network share, you can configure a Network Share Plugin here. For this lesson we will only highlight the fields that are needed to configure this.



1. Click on click here.

## Network Share Plugin

[230]

Below you can see the details to configure a Network Share that Aria Operations can save reports to.

**Add/Edit Outbound Instance** ? X

**Plugin Type** Network Share Plugin

**Instance Name** Instance Name

**Domain** Required

**User Name** Required

**Password** \*\*\*\*\*

**Network share root** Required

TEST CANCEL SAVE

1. Click CANCEL.

## Cancel out of Define Schedule

[231]

Schedule HOL Custom Report

1 Select an Object

2 Define Schedule

### Define Schedule

Set the recurrence and publishing criteria for this report

**Recurrence**

Time zone: --Select--

Start hour: 8 00 AM

Start date: 8/7/23

Weekly

Every 1 weeks on:

Recurrence:  Sunday  Monday  Tuesday  Wednesday  Thursday  Friday  Saturday

**Publishing**

Email report

Email addresses: Separate addresses with a comma or a semicolon. Cc Bcc

Select an outbound rule: HOL Email

⚠ There are no external locations defined, [click here](#) to configure a new external location.

Save to external location

Select a location: --Select--

Relative Path: .import relative path to upload.

1 CANCEL BACK FINISH

1. Click CANCEL.

## Lesson End

[232]

In this lesson we created a custom report using 4 different pie chart views and learned how to setup a Standard Email Plugin so we can email the report out on a schedule.

## Conclusion

[233]

In this module, we reviewed how to run, create and schedule reports. How to setup a Standard Email Plugin and reviewed the Network Share Plugin.

## You've finished the module

[234]

Congratulations on completing the lab module.

For more information on getting started with Aria Operations, see the [VMware Aria Operations: Journey to Success](#) guide at the [VMware Apps & Cloud Management Tech Zone](#).

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 7 - Creating Custom Dashboards for VMware Aria Operations (15 minutes) Basic

### Introduction

[236]

Understanding how to create custom dashboards is a very powerful skill to have with Aria Operations. Being able to create a custom window into your infrastructure, purpose built for an OS admin or management or an application administrator... will help to run a lean and efficient infrastructure and save the company money in the long run.

### Log in to Aria Operations

[237]

We will log in to a live instance of Aria Operations running in this lab.

### Open the Firefox Browser from the Windows Task Bar

[238]

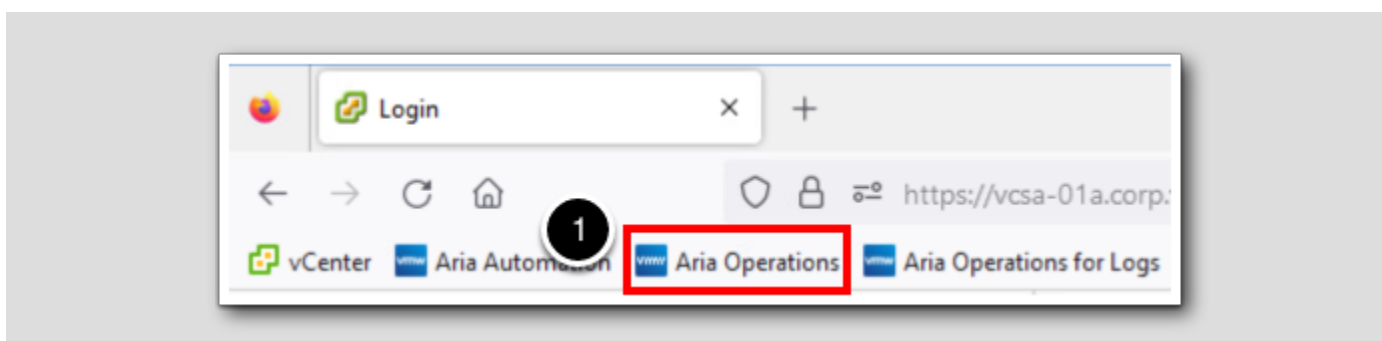


If the browser is not already open, launch Firefox.

1. Click the Firefox icon in the Windows Quick Launch Task Bar at the bottom of the screen.

### Navigate to Aria Operations

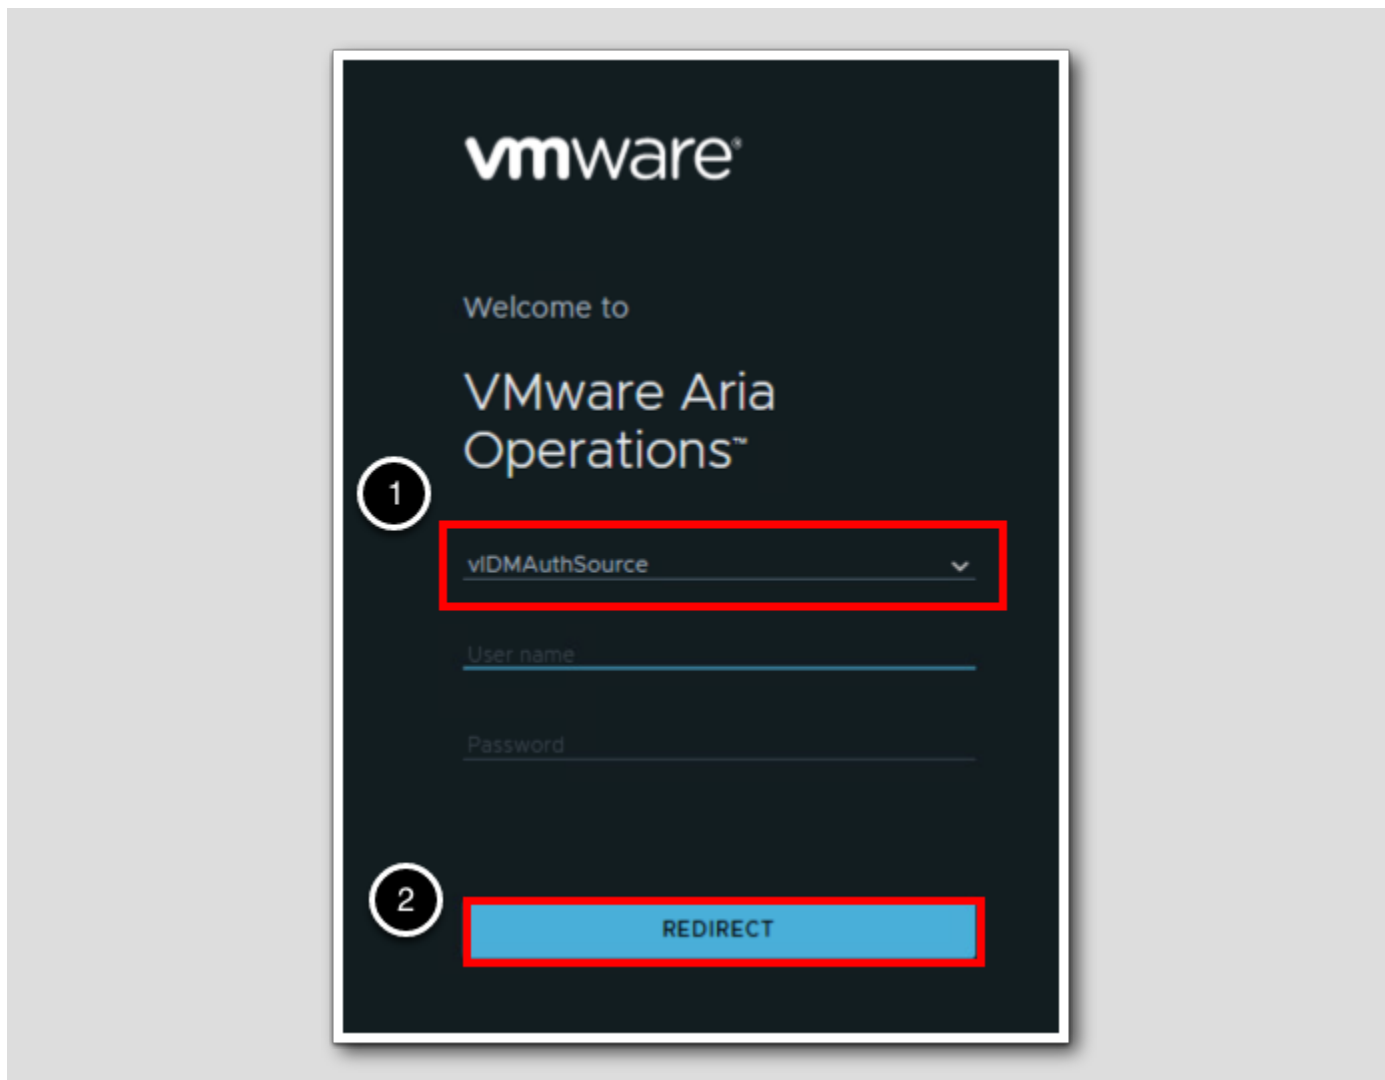
[239]



1. Click the Aria Operations bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[240]



Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login

[241]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

1. Click Sign in

## Clone and Modify Existing Dashboards

[242]

In this lesson we will clone an existing dashboard and modify it to add more views to enrich the existing dashboard.



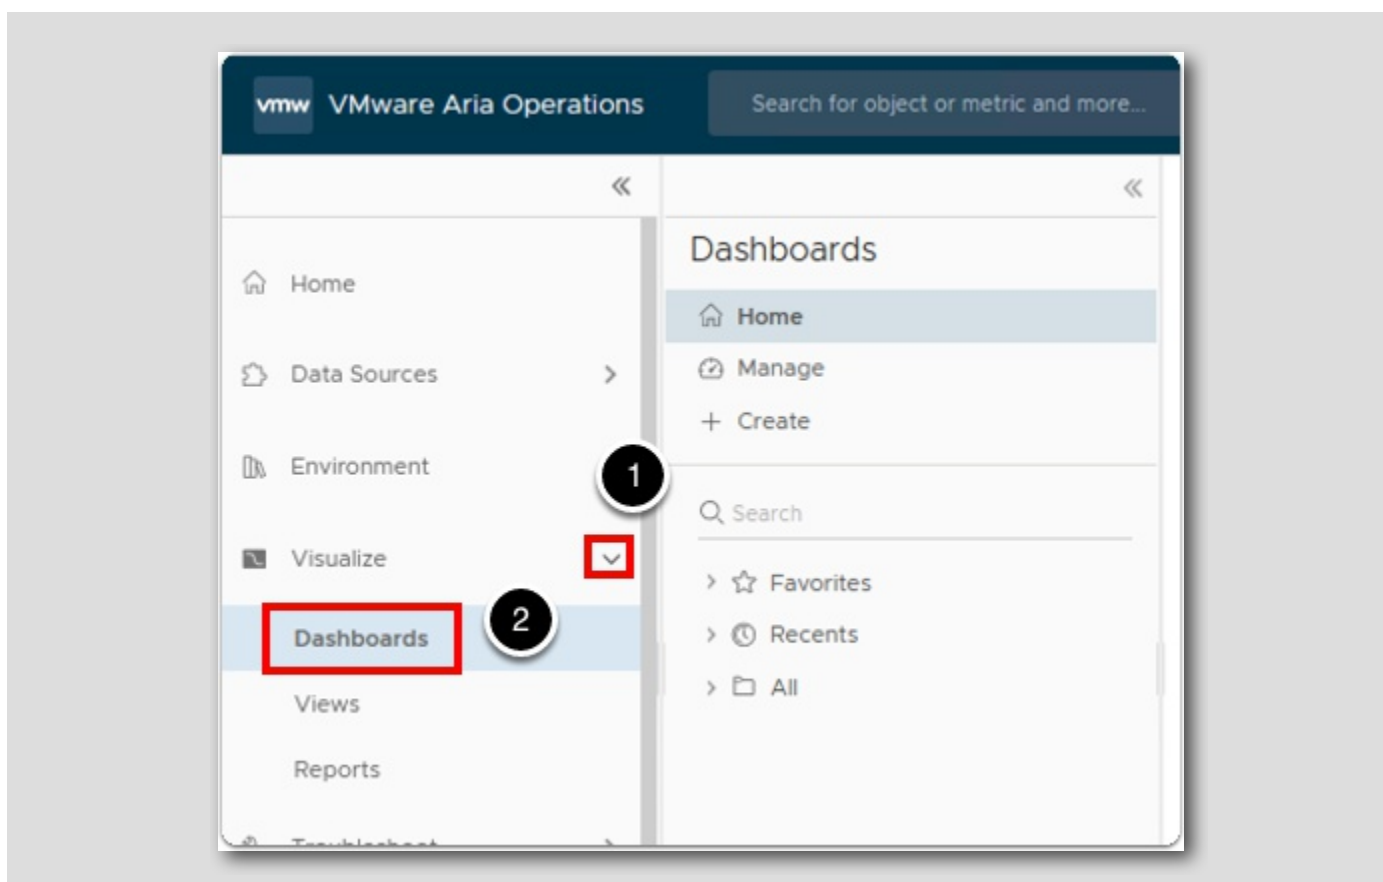
## Dashboards

In this lesson, we will learn how to clone an existing dashboard and modify it to make it our own.

Aria Operations has numerous out-of-the-box dashboards that were created by industry experts who have a deep understanding of Aria Operations as well as the characteristics and behavior of the underlying objects being managed. However, personalizing a Dashboard to fit a specific role or consolidate other information into a single view is a common use case for most administrators.

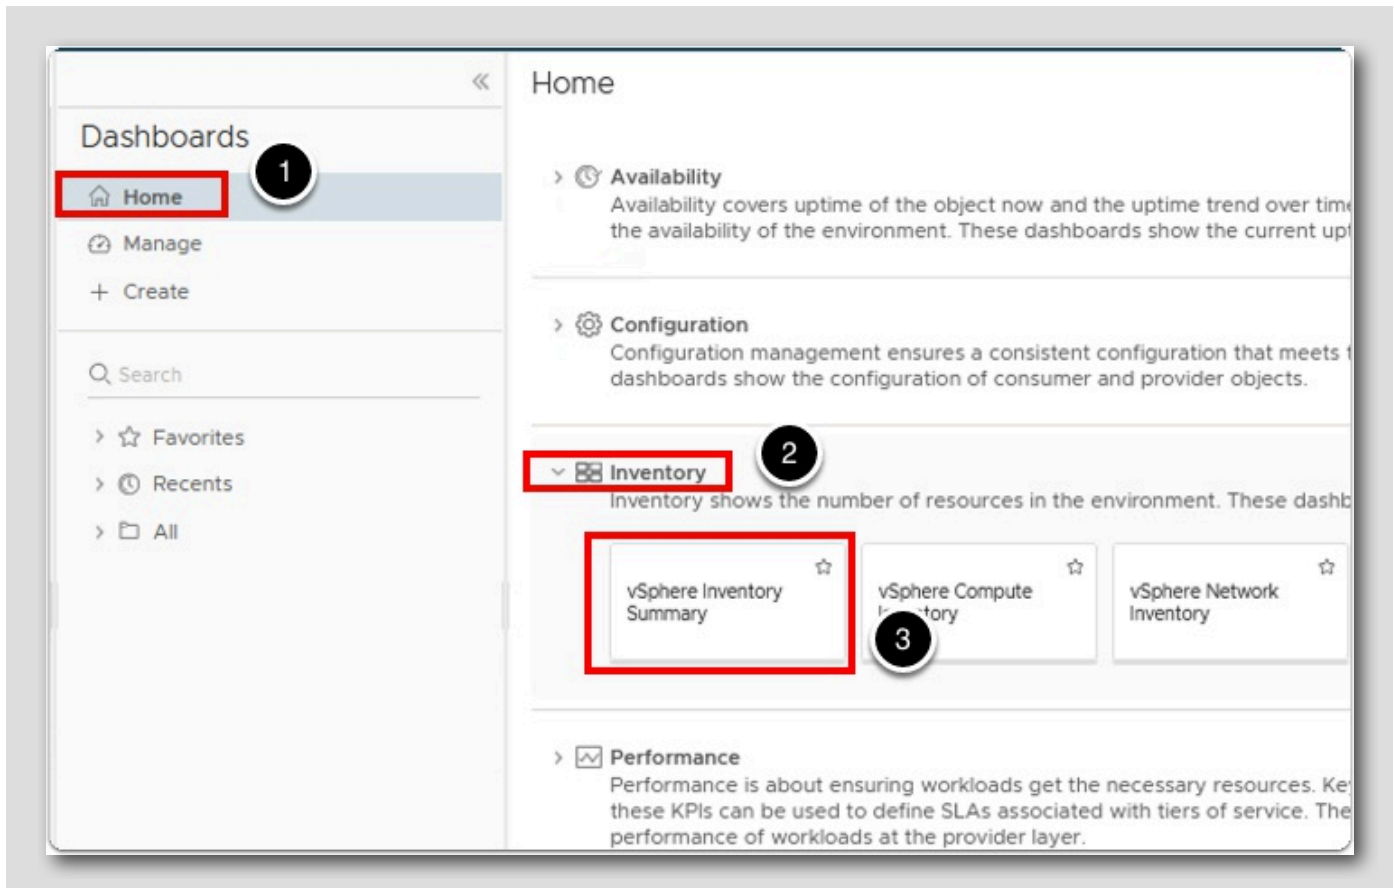
To start, we will clone and make some simple changes to create a custom Overview Dashboard for our administrators. For this example, we will clone an existing dashboard and add the Scoreboard Health, Object relationship and Top Alerts widgets. We will also minimize the three Top-15 widgets that are in the existing dashboard so we will have more screen real estate in the dashboard.

Cloning an existing dashboard to create a new or modified dashboard is considered a best practice to ensure your custom content is not affected during an upgrade of Aria Operations.



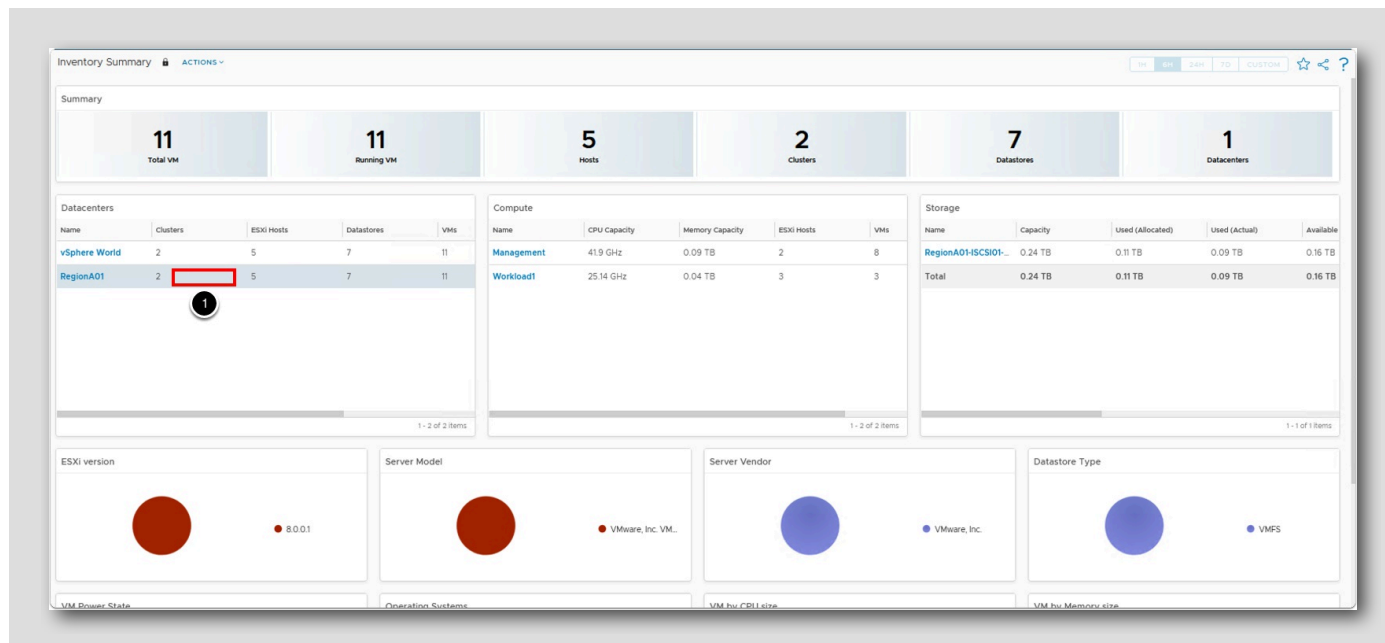
1. Expand **Visualize**.
2. Click on **Dashboards** in the menu bar at the left of the user interface.

## vSphere Inventory Summary



1. Under Dashboards, Click on Home. to expand the Dashboards menu.
2. Expand Inventory.
3. Click on vSphere Inventory Summary.

## vSphere Inventory Summary



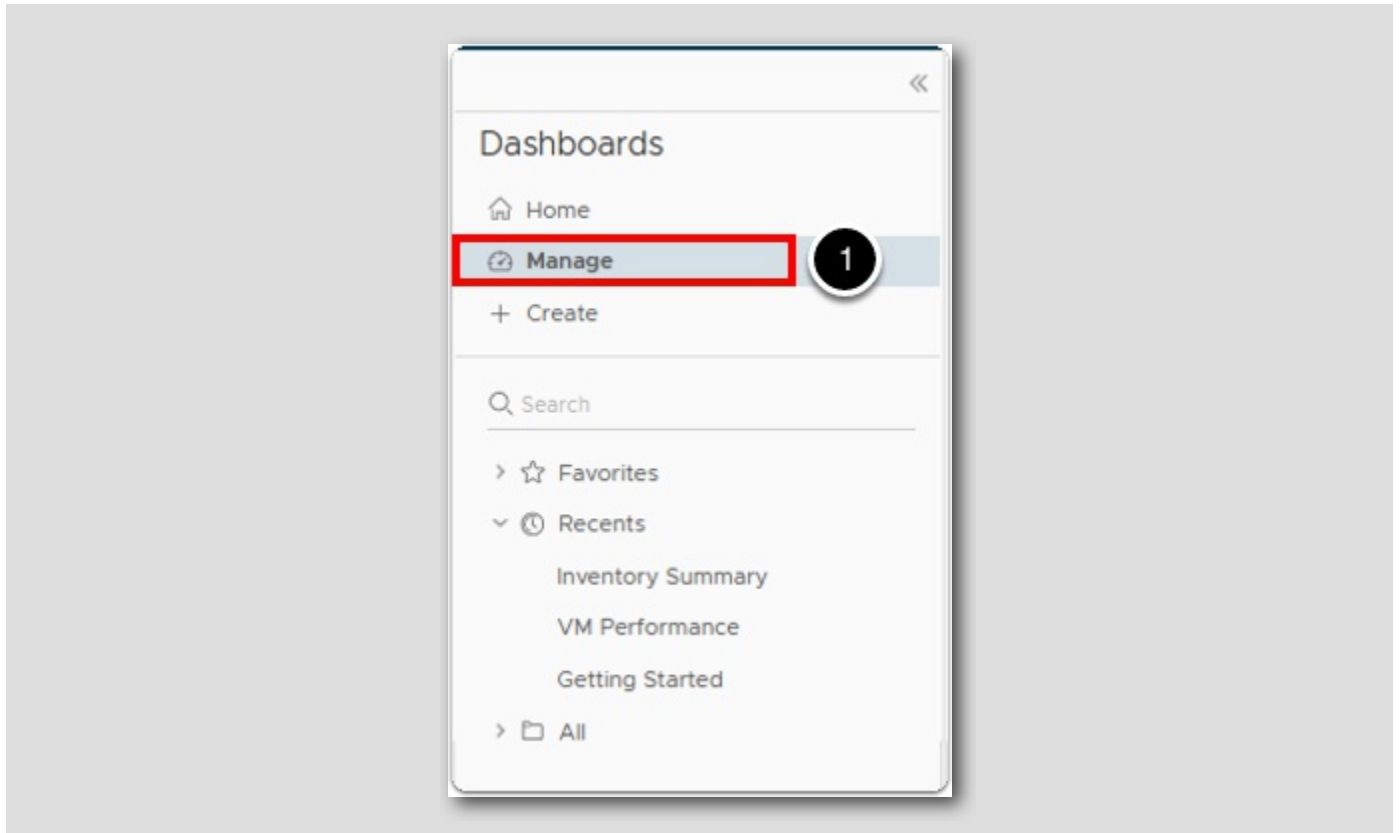
We can now see the Inventory Summary dashboard, which will be the basis for our own customized version of this dashboard. In order to modify this or any other existing dashboard, we will first want to "clone" it and then modify the cloned version. We do not want to edit any master or default out-of-the-box dashboards so we don't potentially break the content and flow. We ALWAYS want to clone a dashboard and edit the clone or just create a brand new custom dashboard from scratch as a best practice!

1. Click on an **empty space** in the **RegionA01** row. Notice how the boxes below populate based on what is selected in the Datacenter box.

If you click on the Datacenter text it will take you to that object in Object Browser.

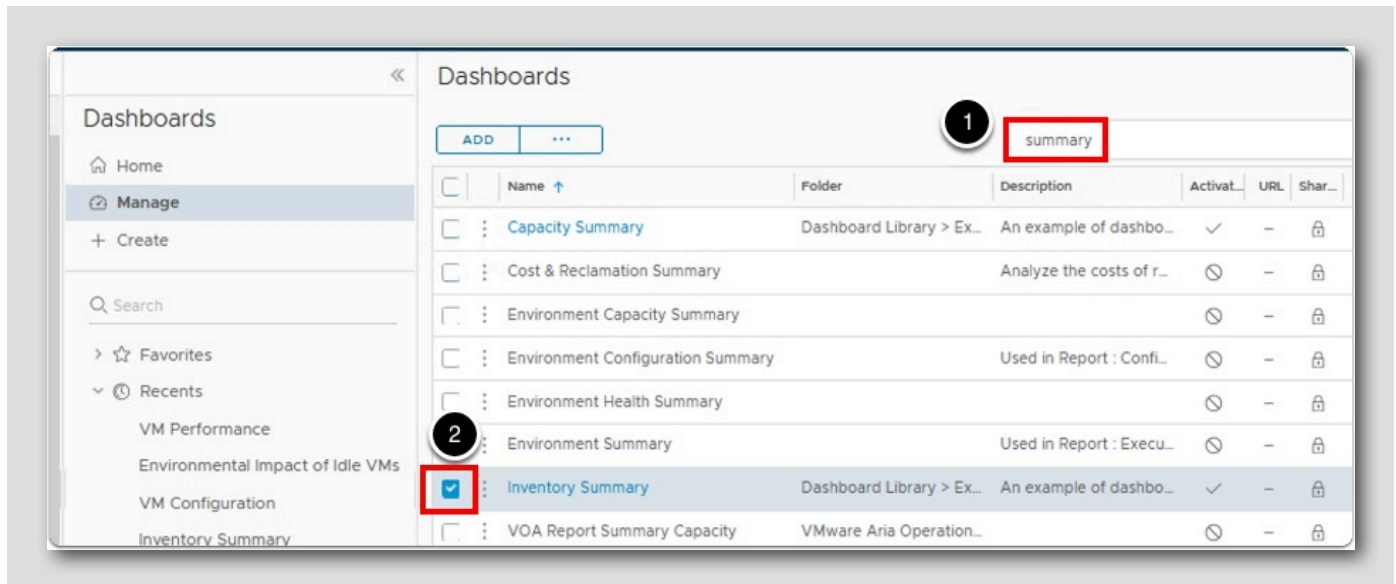
## Manage Dashboards

[246]



1. Click on **Manage** under the Dashboards .

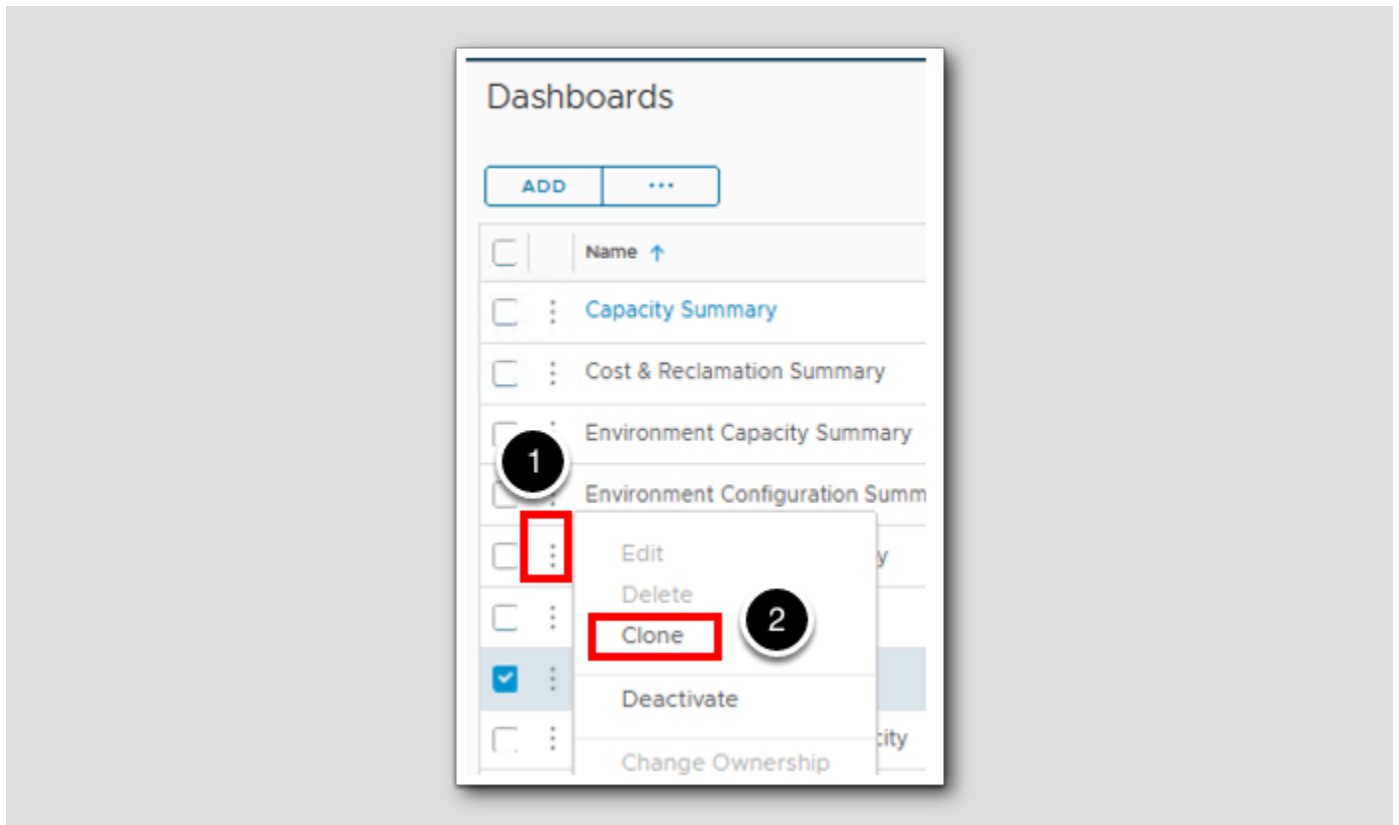
## Filter for Inventory Summary Dashboard



1. Type **summary** into the **Quick filter** field and then hit the **ENTER** key on our keyboard to search for dashboards with the word **environ** in the title.
2. Click the **checkbox** next to **Inventory Summary** report.

## Clone Dashboard

[248]

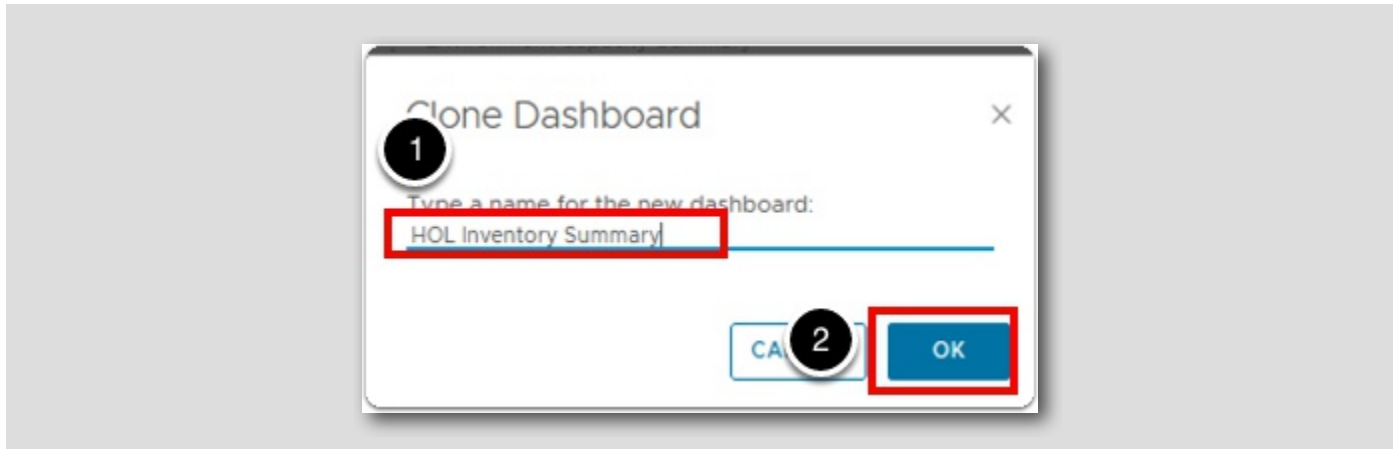


1. Click on the 3 dots to the right of the checkbox on the **Inventory Summary** Dashboard row to open the actions menu.
2. Then click **Clone** to start the cloning process.

To emphasize again, never modify a delivered Dashboard. ALWAYS Clone when you want to modify a Dashboard.

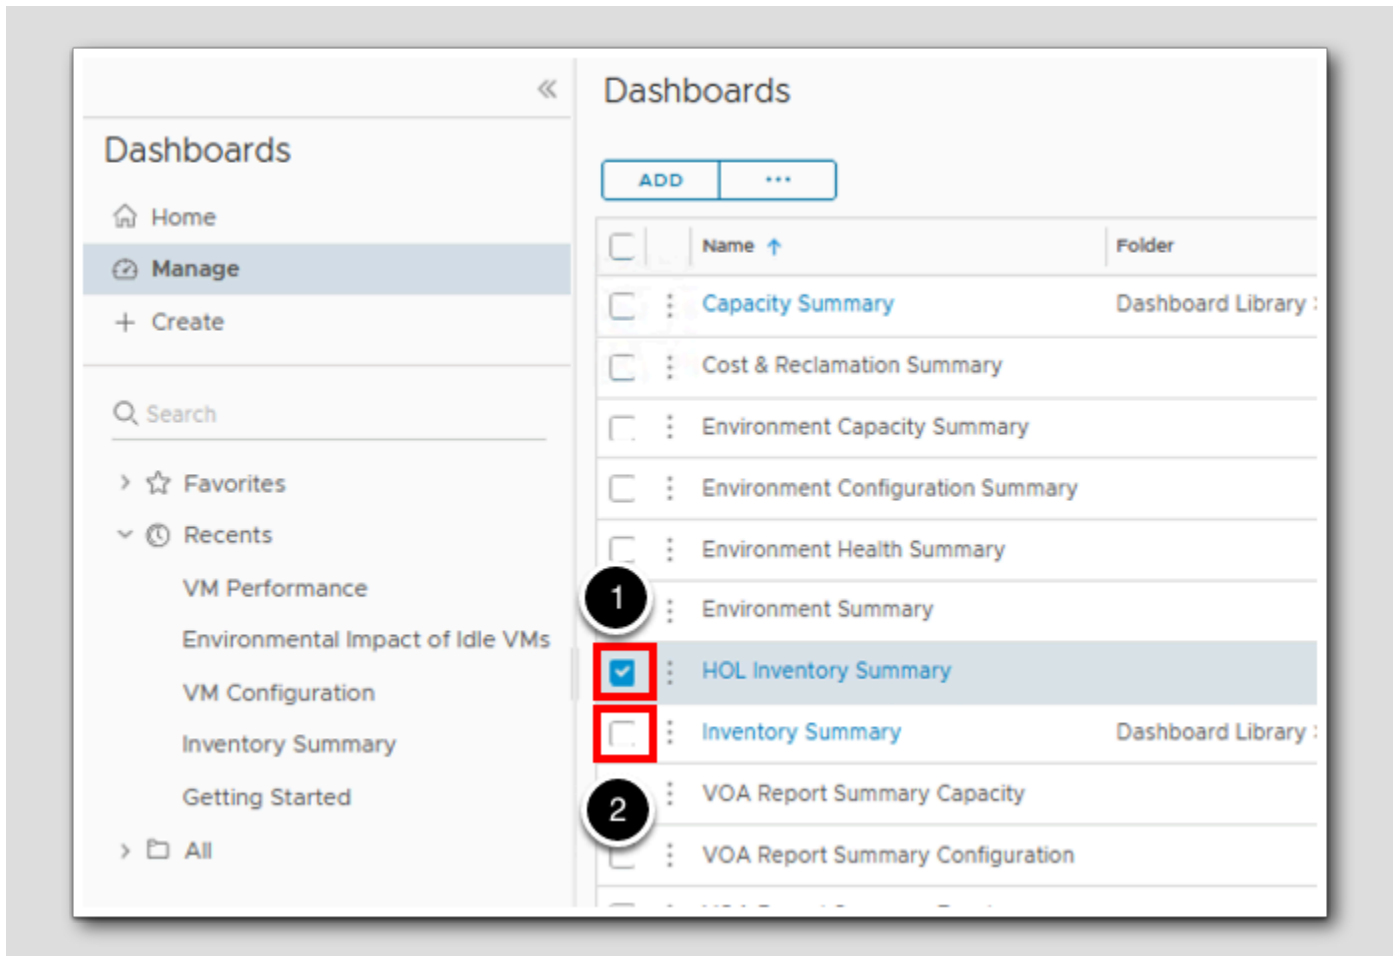
## Clone Dashboard - Name the New Dashboard

[249]



1. Type HOL Inventory Summary into the name text field.
2. Then click on the OK button to save the name.

## Change to the Cloned Dashboard

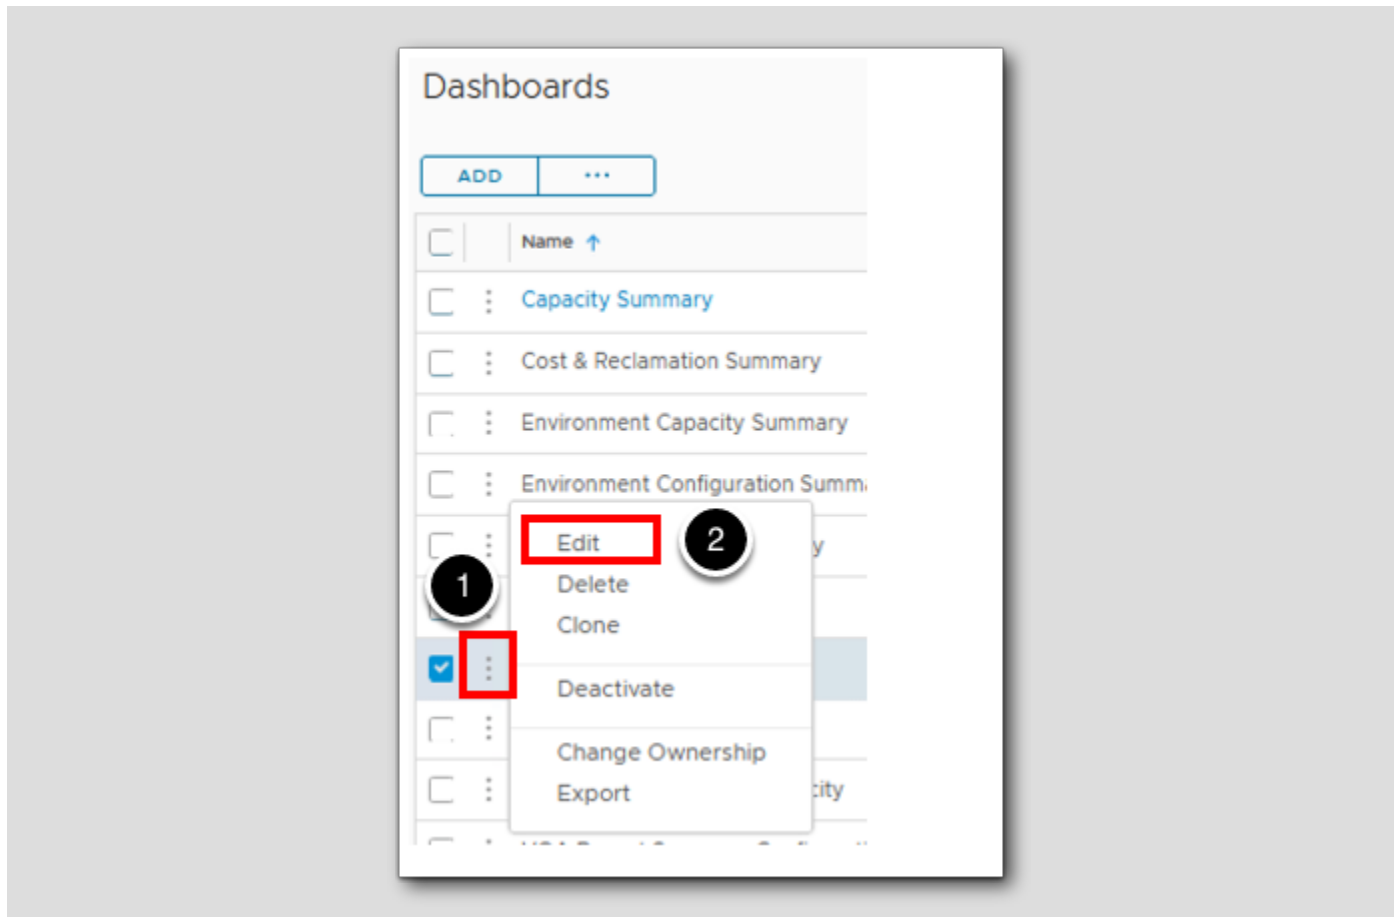


1. Check the box next to the new cloned dashboard HOL Inventory Summary.
2. Uncheck the box next to the original Inventory Summary Dashboard.



## Manage Dashboards - Edit Dashboard

[251]

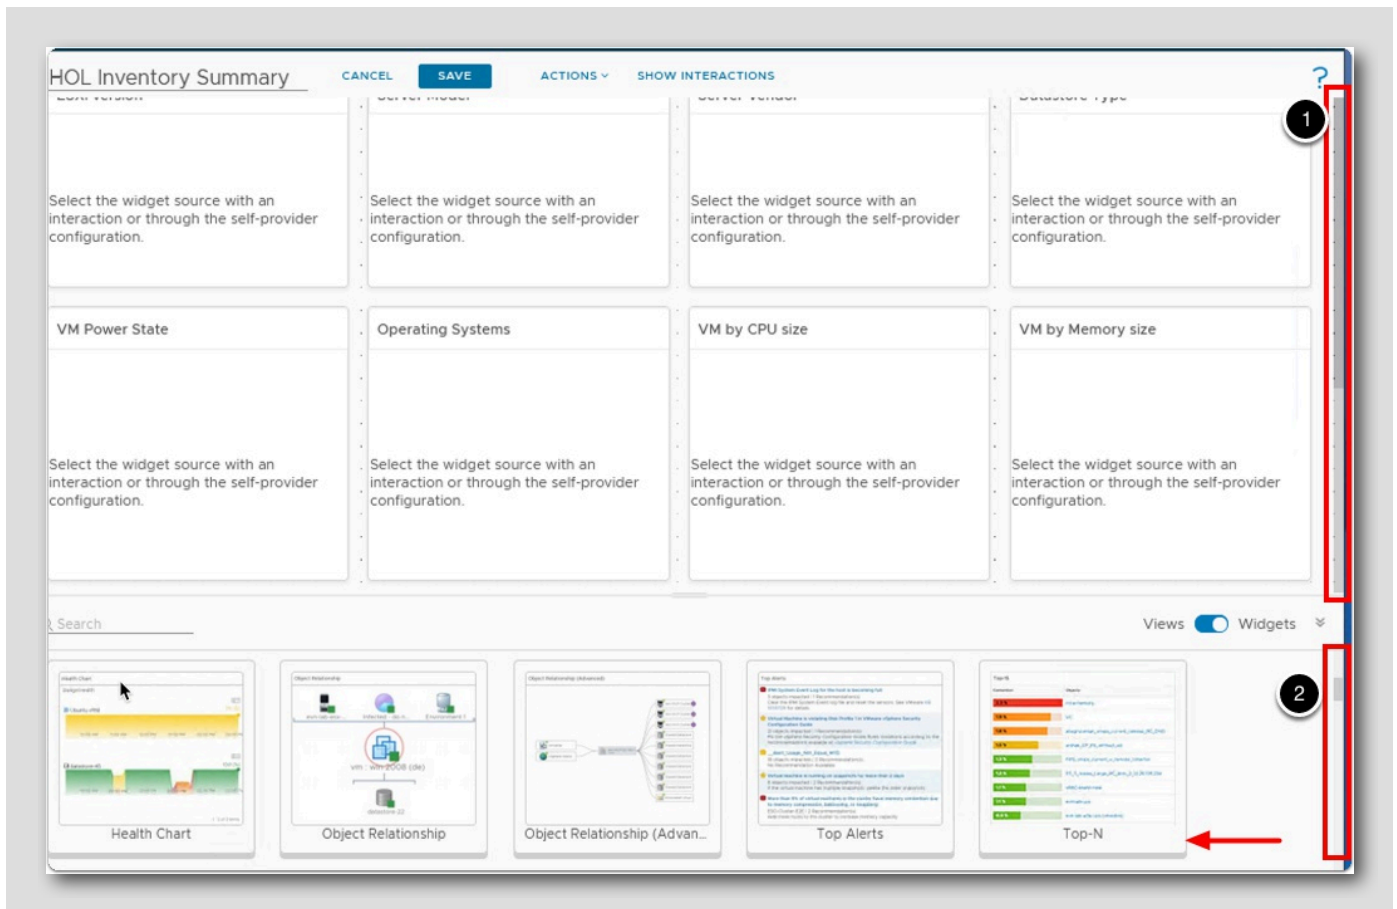


1. Click on the 3 dots next to the right of the checkbox on the HOL Inventory Summary Dashboard row to open the actions menu.
2. Then click Edit.

## Edit Dashboard - Add Widgets

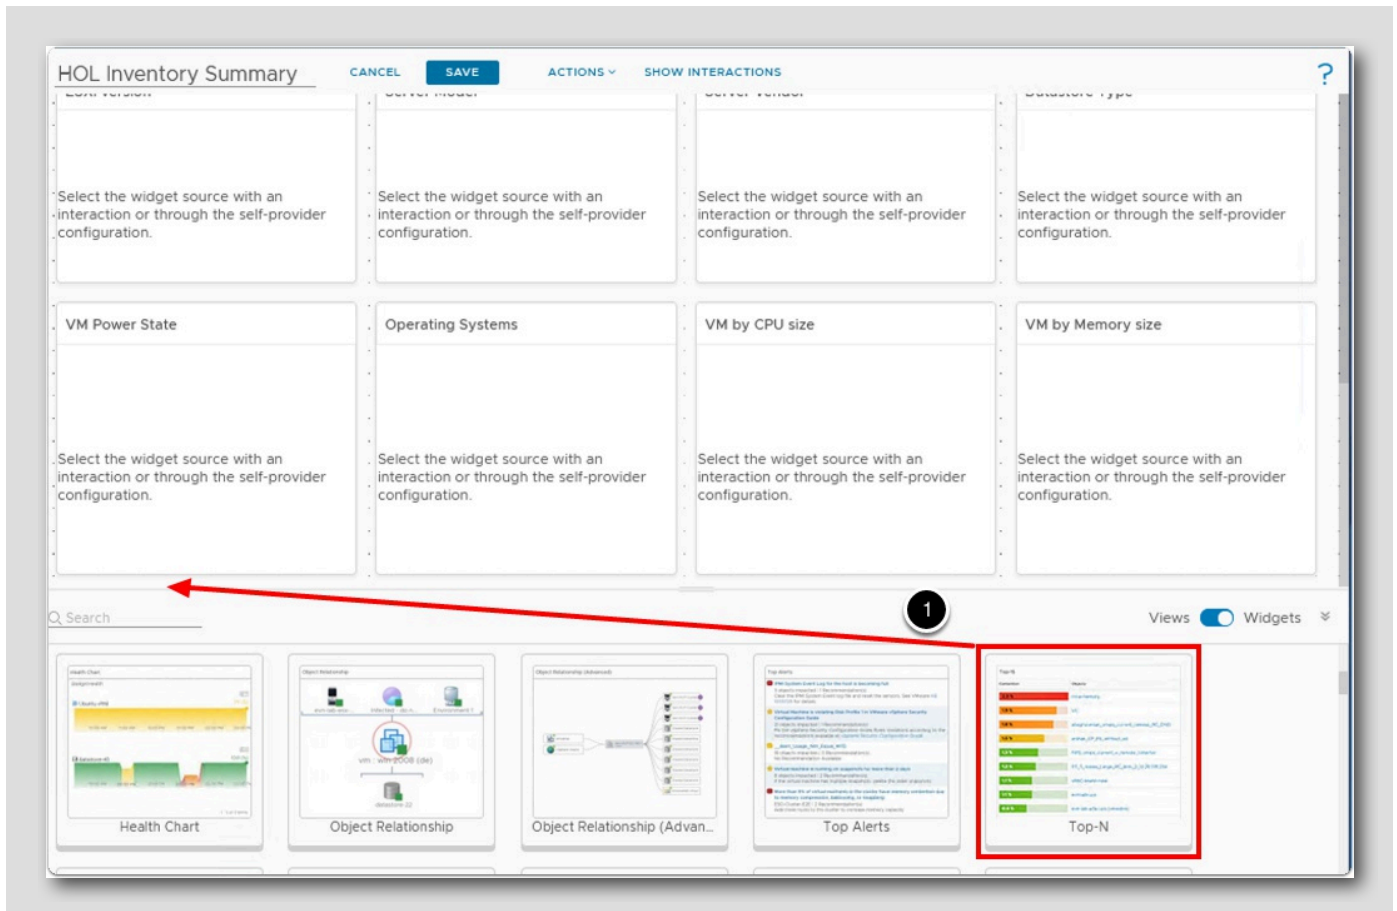
[252]

In this lesson we are going to add three Top-N widgets and configure them to show the top-5 CPU, Memory and Storage IOPS consumers.



1. Scroll down to the bottom of the top window.
2. Scroll down in the bottom window until you see the Top-N widget.

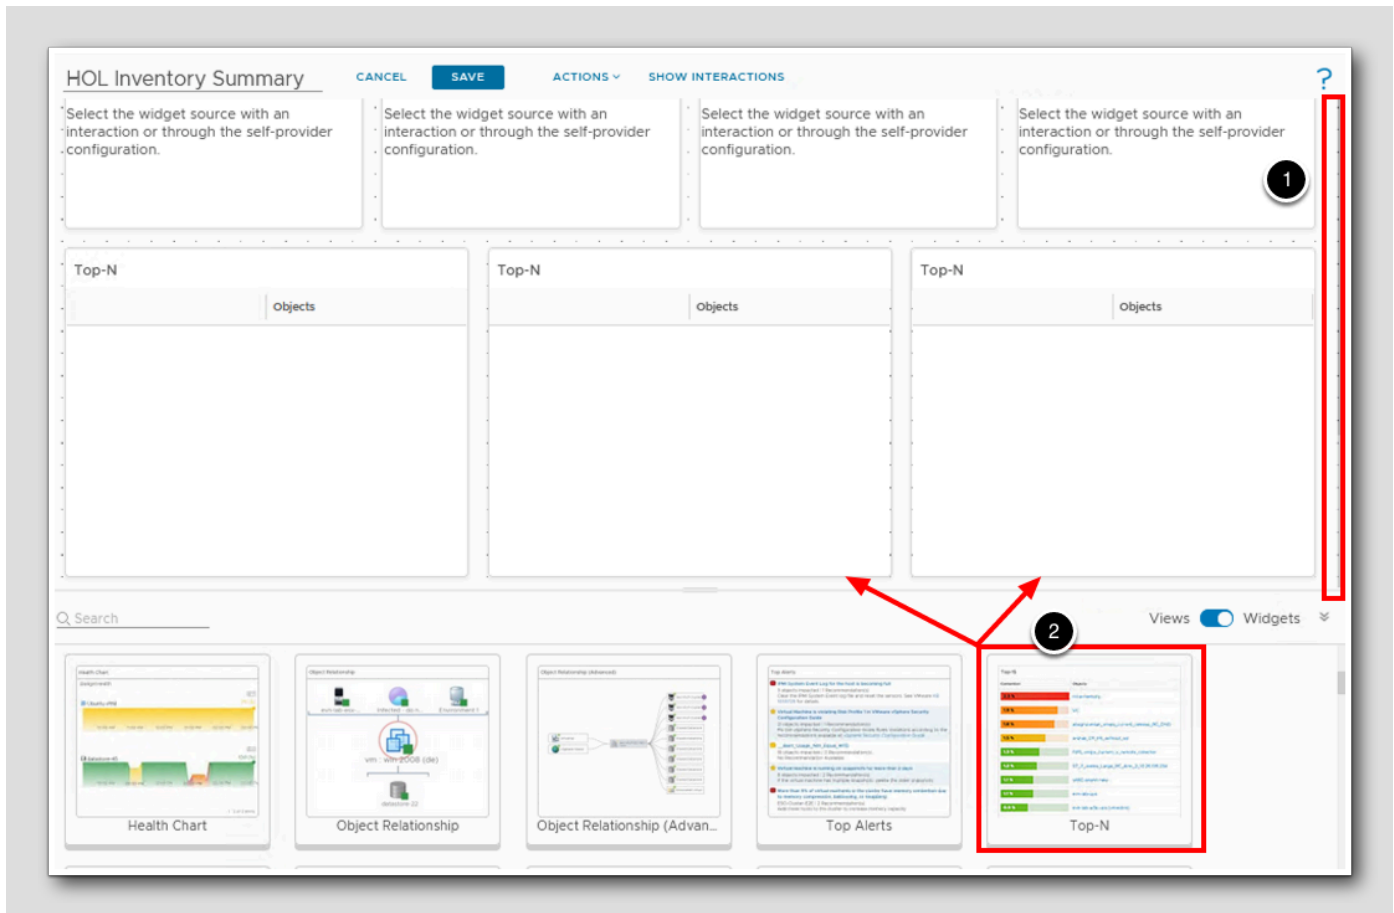
Edit Dashboard - Add Widgets (continued)



1. Drag and drop one Top-N widget into the space just below the VM Power State widget.

## Edit Dashboard - Add Widgets (continued)


[254]



1. Scroll down again in the top box to expose the Top-N widget we just added.
2. Drag and drop two more Top-N widgets. Adjust the widgets as needed so it looks like the screenshot above.

## Edit Dashboard - Show Interactions

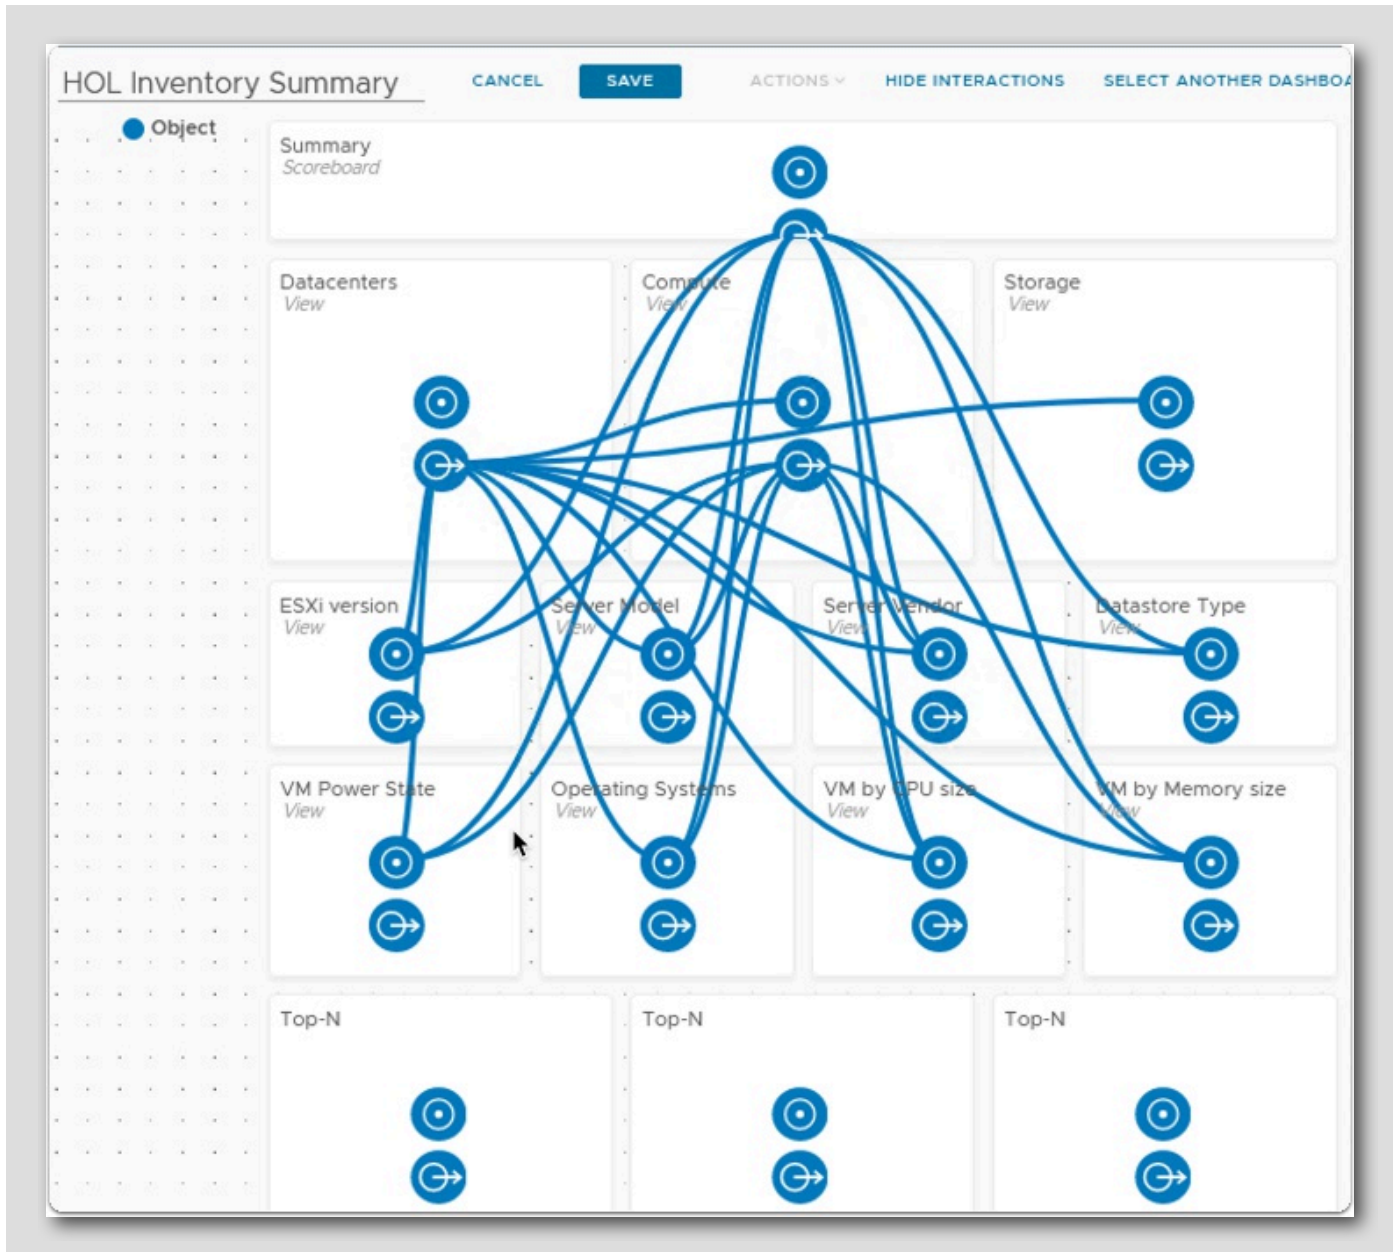
[255]



1. Click on the Show Interactions text link at the top of the window.

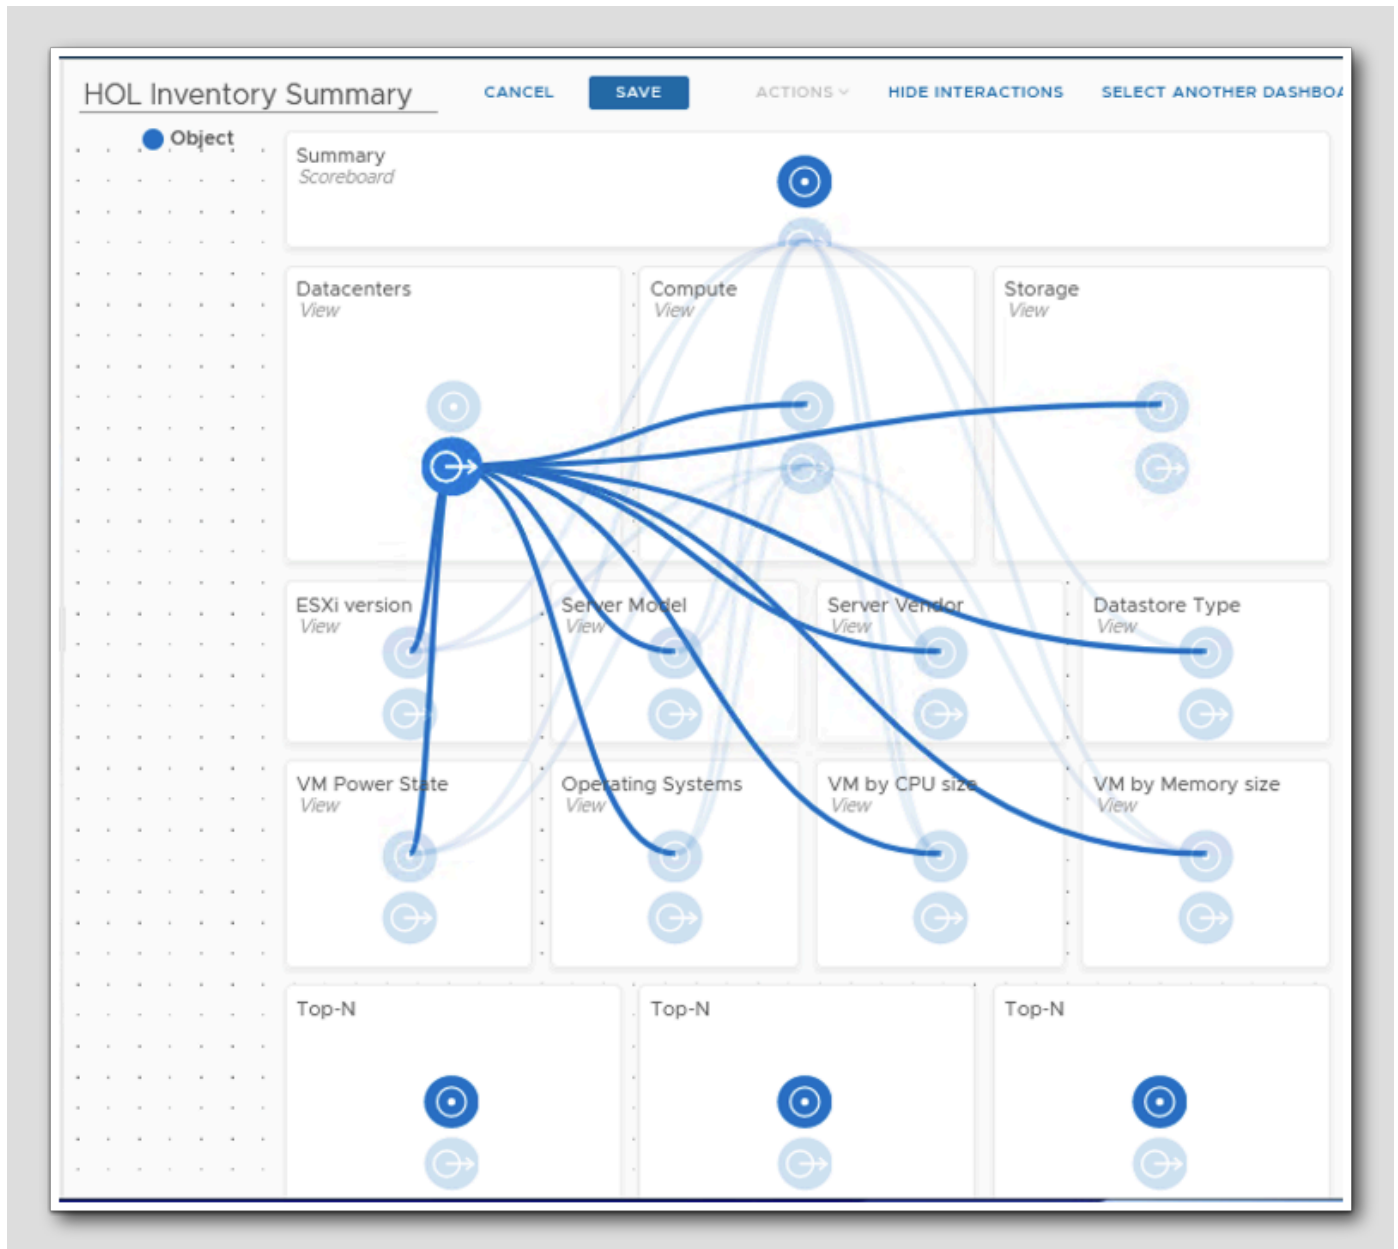
Edit Dashboard - Connected Relationships

[256]



We see that since we cloned an existing dashboard, there are already relationships created from the "Datacenters View". There are circles with an arrow and circles with a dot in the middle. The circles with an arrow are the driving control and connect to a circle with a dot in the middle, receiving control.

Hover over the Datacenters controlling object



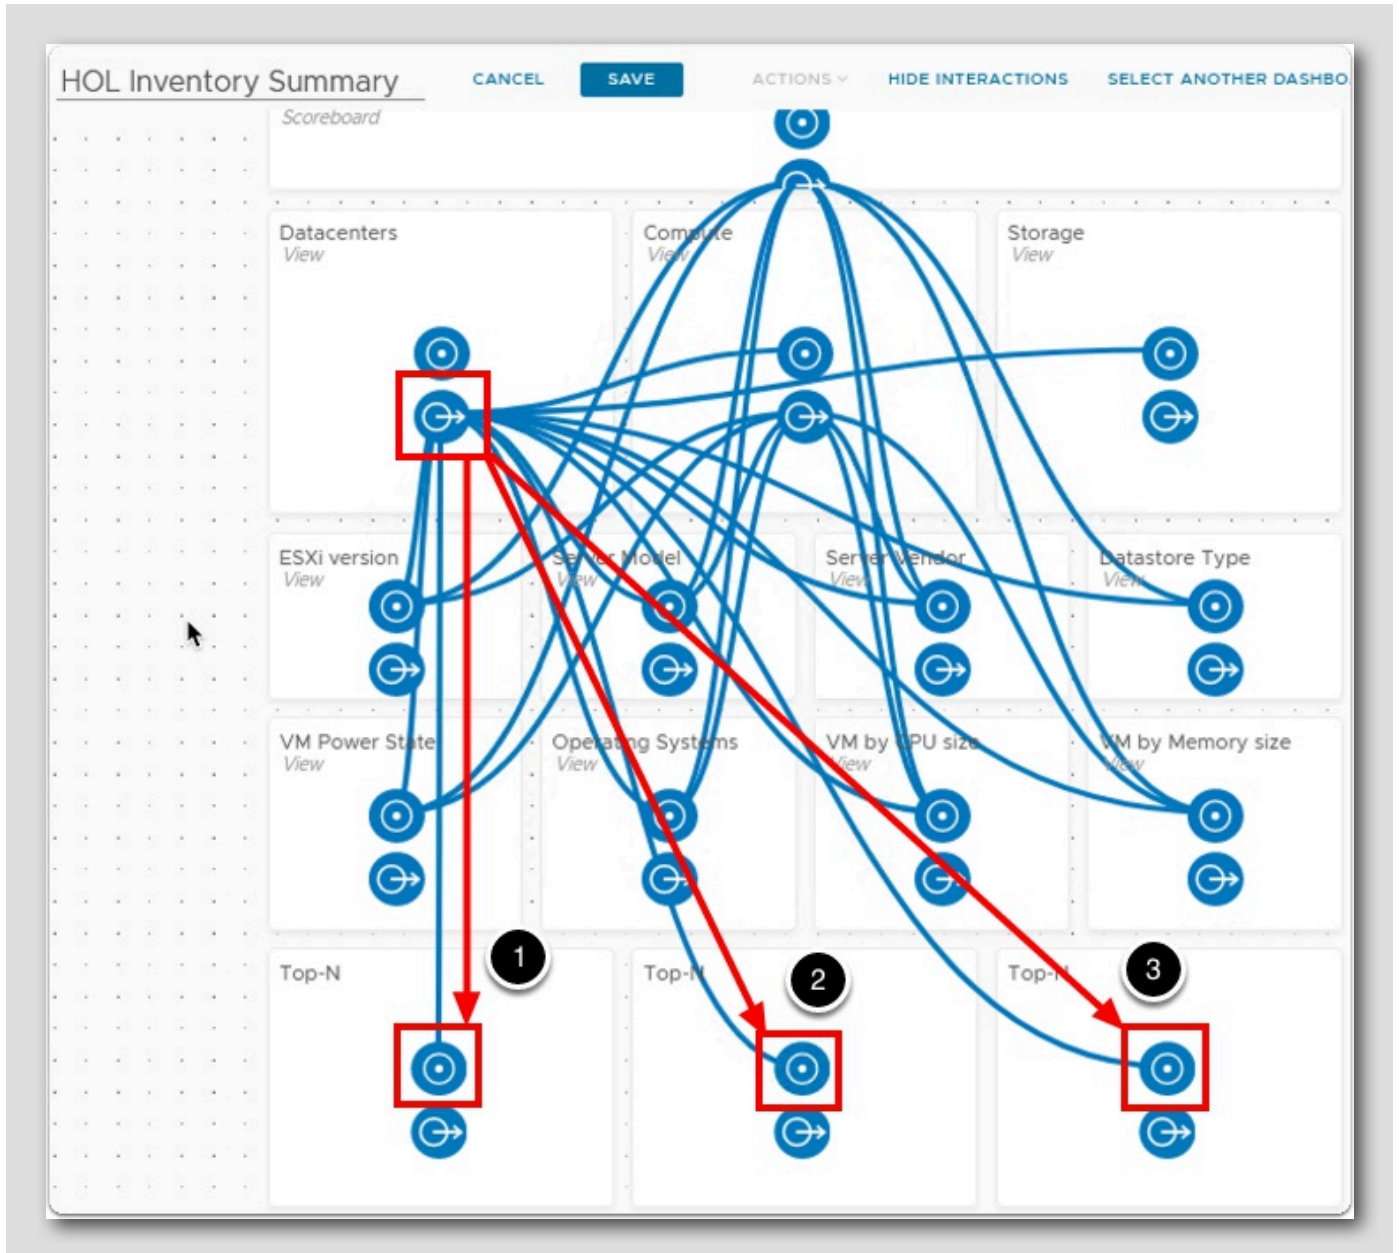
If you hover your mouse over the Datacenters driving control, you can see that it controls every widget in this dashboard except the top Summary Scoreboard. It is critical to understand widget interactions to get your dashboards to display the correct information.

Now we will get the Datacenters View to drive our new Top-N widgets we added.



### Edit Dashboard - Connecting Relationships

[258]



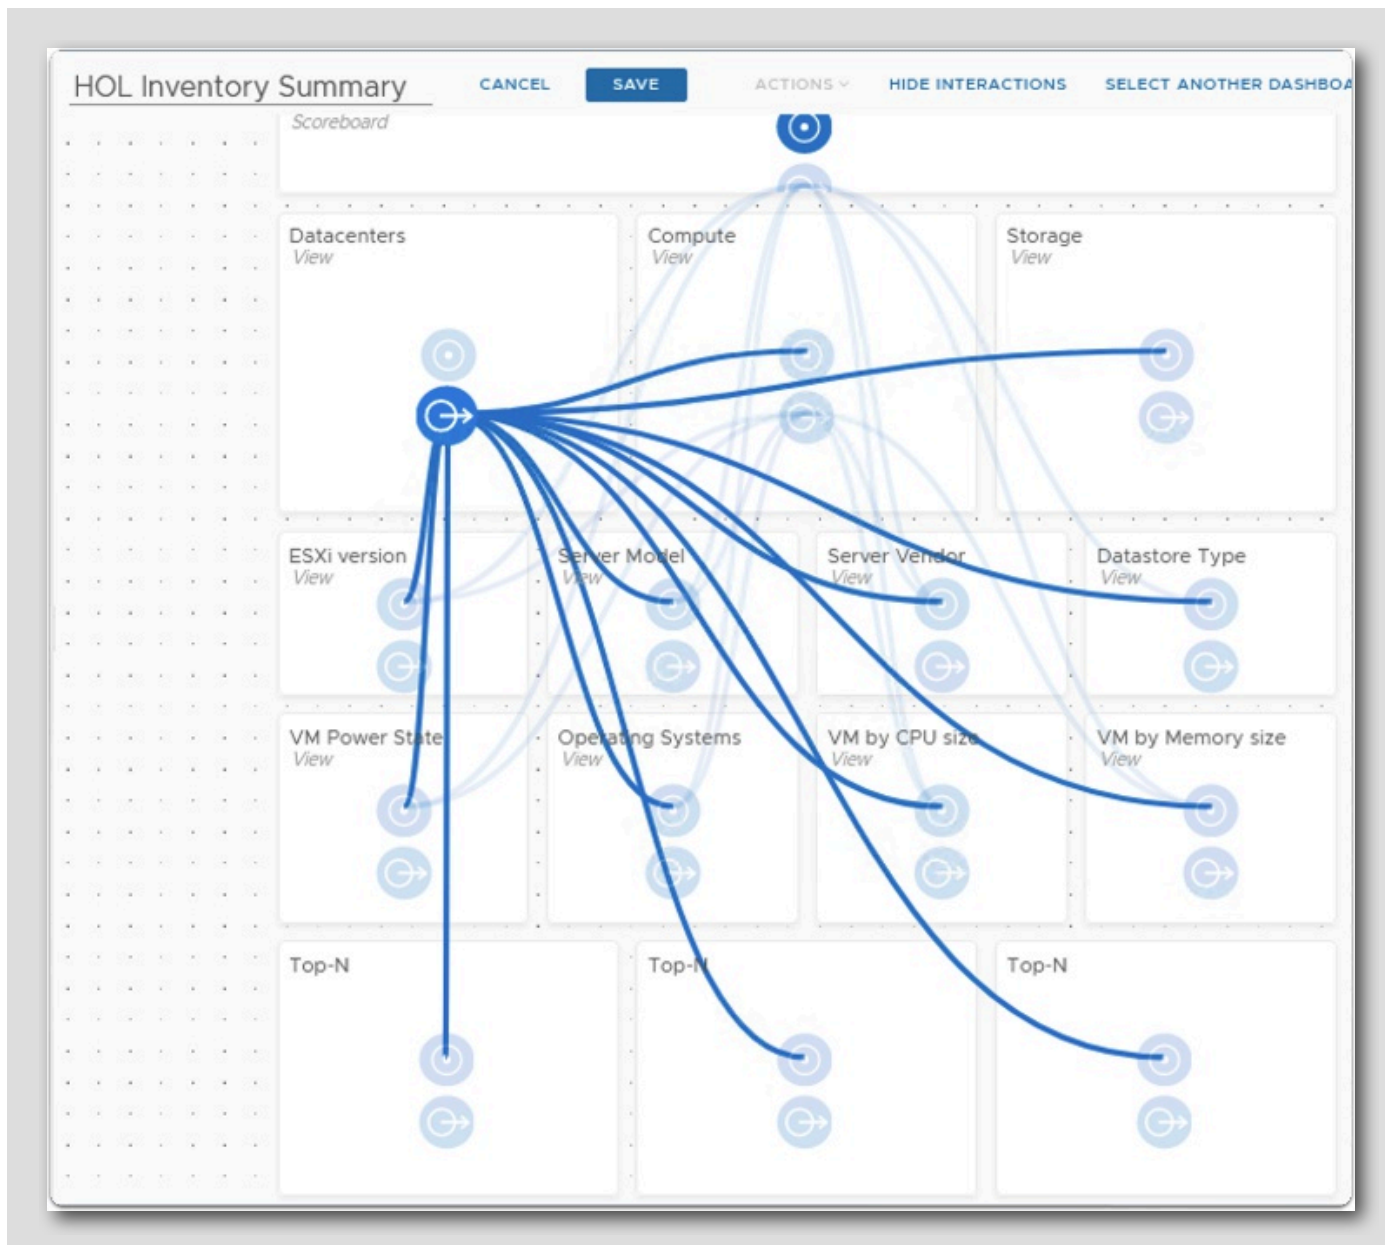
Here is where we need to connect and create the relationships between the "Datacenter View" widget and the three new Top-N widgets we have added. We will do this by dragging and dropping from the "Datacenter View" icon to each of the three icons in the new Top-N widgets we added.

1. Click on the circle with an arrow icon in the Datacenter View widget and drag it on top of the circle with the dot icon in the 1st Top-N widget.
2. Click on the circle with an arrow icon in the Datacenter View widget and drag it on top of the circle with the dot icon in the 2nd Top-N widget.
3. Click on the circle with an arrow icon in the Datacenter View widget and drag it on top of the circle with the dot icon in the 3rd Top-N widget.



## Edit Dashboard - Completed Relationships

[259]



Hover your mouse over the circle with an arrow in the Datacenters View and we should now see the lab environment match the screen capture.

As we see here, we have connecting relationship lines from the Datacenters View widget to each of the 3 new widgets we added.

## Hide Interactions

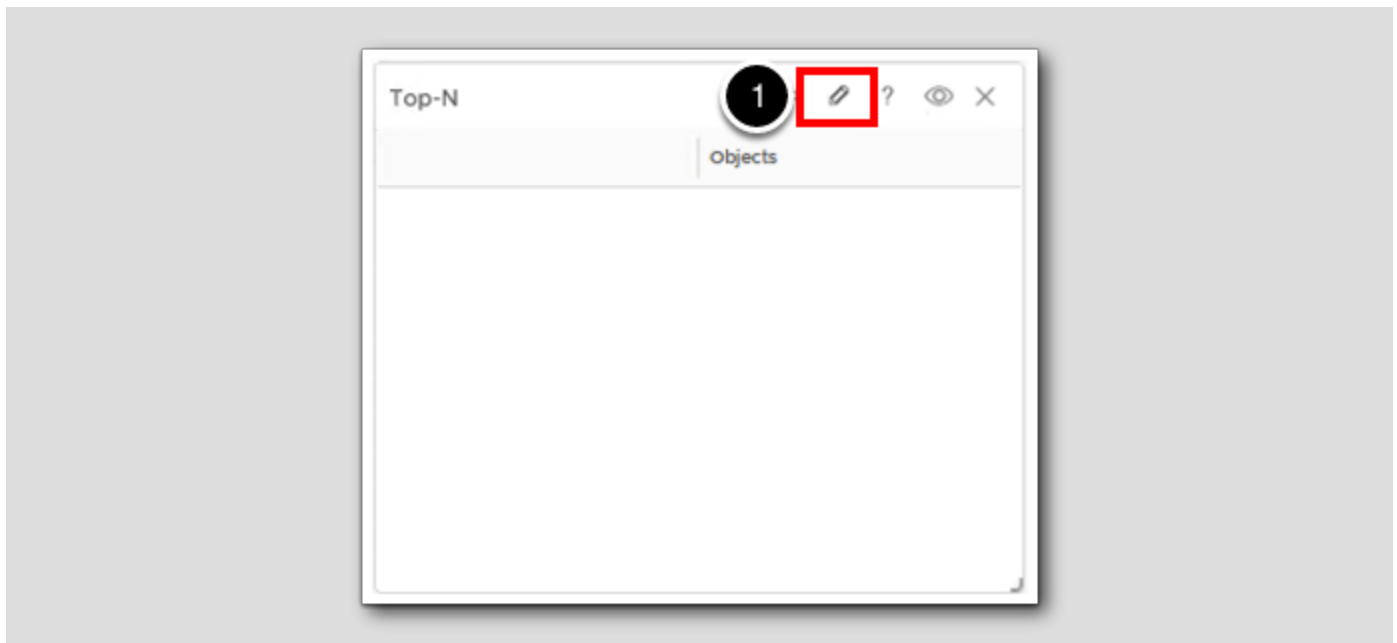
[260]



1. Click on HIDE INTERACTIONS to return to the Dashboard building page.

## Configure the Top-N widgets - CPU

[261]



1. Hover over the left most Top-N widget and click on the pencil icon when it appears.

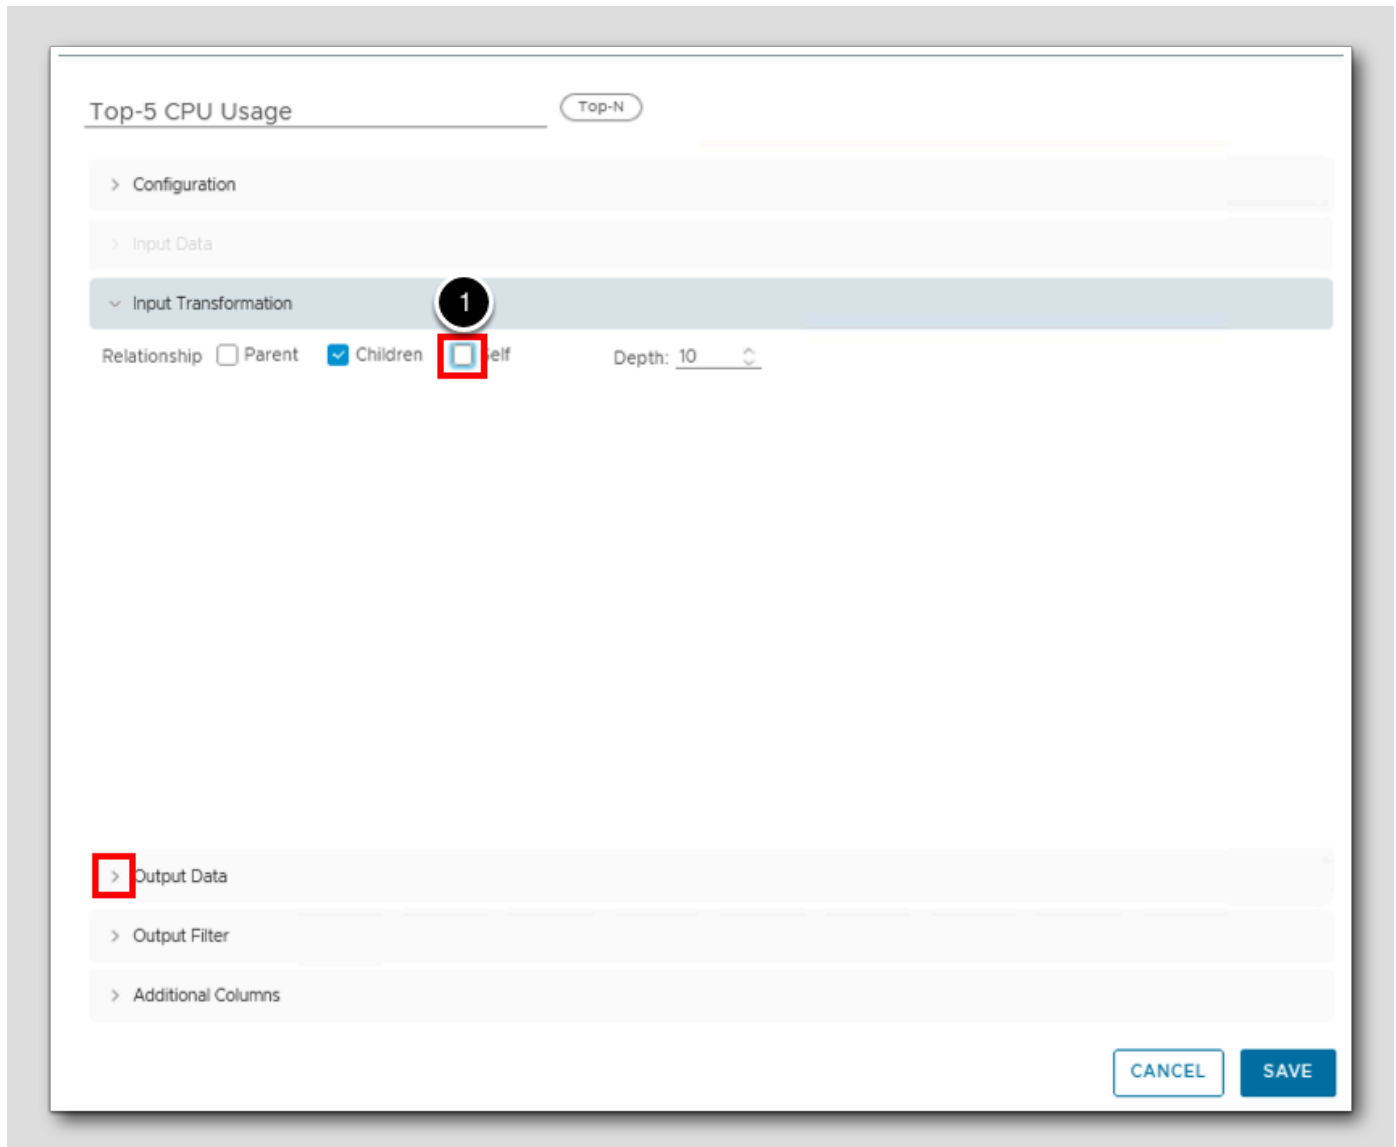
## Configure the Top-N widgets - CPU (Continued)

The screenshot shows the configuration interface for a 'Top-N' widget. The title 'Top-5 CPU Usage' is highlighted with a red box and labeled '1'. The 'Configuration' section includes settings for 'Refresh Content' (Off), 'Refresh Interval' (300 seconds), 'Self Provider' (Off), 'Redraw Rate' (15 minutes), 'Bars Count' (5), 'Round Decimals' (0), and 'Filter old metrics' (Off). Under 'Top-N Options', the 'Metric Analysis' radio button is selected and highlighted with a red box and labeled '2'. A dropdown menu is open, showing options: 'Top Highest Utilization' (selected), 'Top Lowest Utilization', 'Top Abnormal States', 'Top Highest Volatility', and 'Percentile'. The 'Input Data' section is expanded, with 'Input Transformation' highlighted by a red box and labeled '3'. At the bottom right, there are 'CANCEL' and 'SAVE' buttons.

1. Change the top name to Top-5 CPU Usage.
2. Check the **Metric Analysis** radio button because we want to display the CPU Usage (%) metric in this widget.
3. Expand **Input Transformation**.

## Input Transformation

Input Transformation will drive the data that is shown in the widget itself. Since we are driving data from the Datacenters View and want to show VM information, we need to specify that we want to show the Children Relationship as a datacenter is a parent of a VM.



1. Uncheck the Self Input Transformation.
2. Expand Output Data.

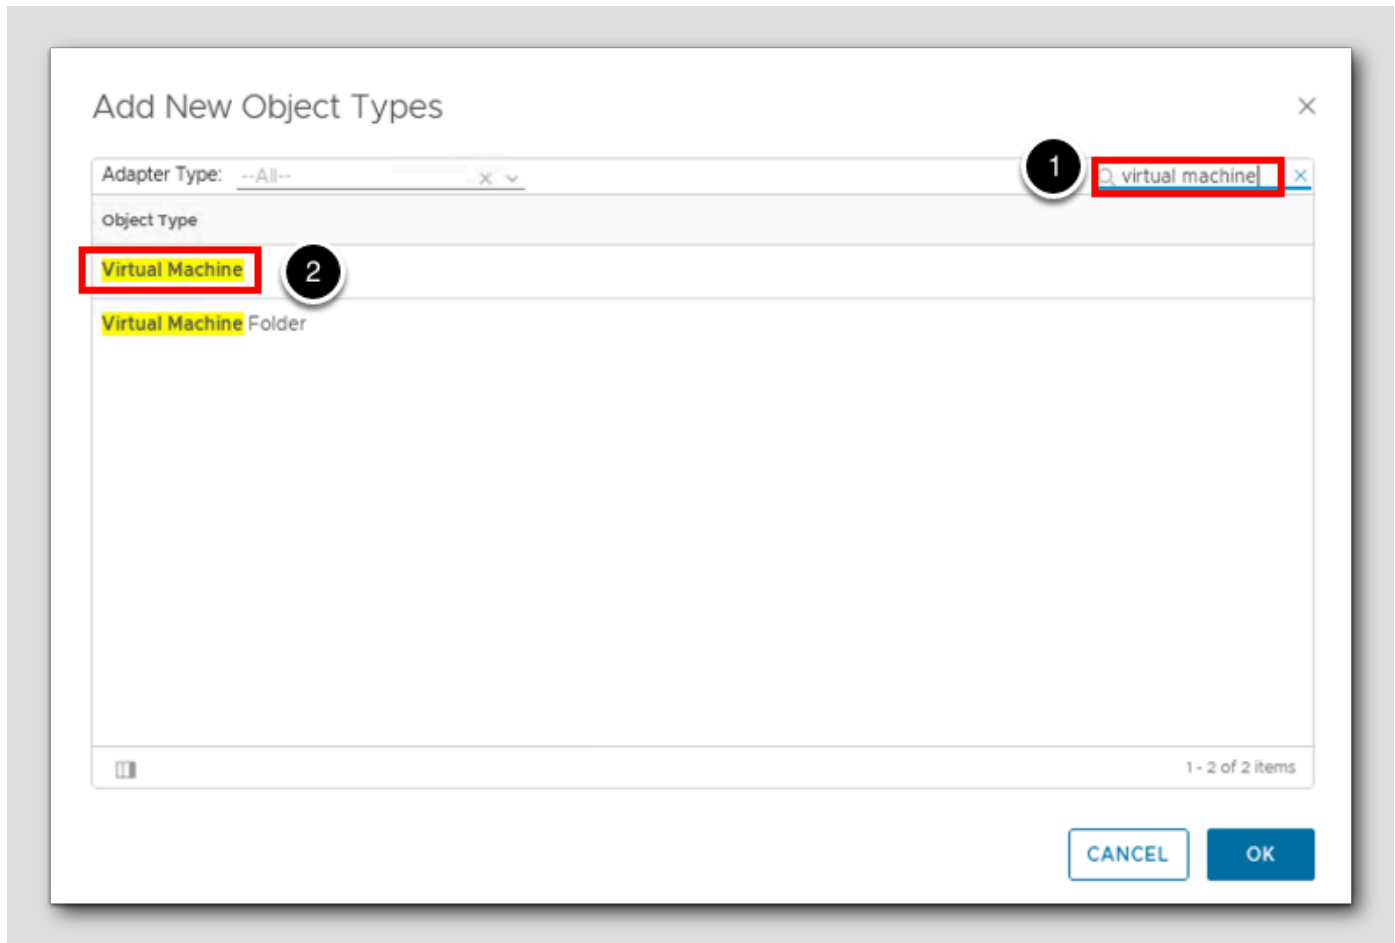
## Output Data

[264]

The screenshot shows the configuration interface for 'Top-5 CPU Usage'. The 'Output Data' section is the active tab, indicated by a '1' in a black circle. A red box highlights the '+' icon in the 'Object Type' table header. Below the table, there are fields for Metric, Label, Unit, Maximum, and Color Method. The 'Output Filter' and 'Additional Columns' sections are also visible. The 'CANCEL' and 'SAVE' buttons are at the bottom right.

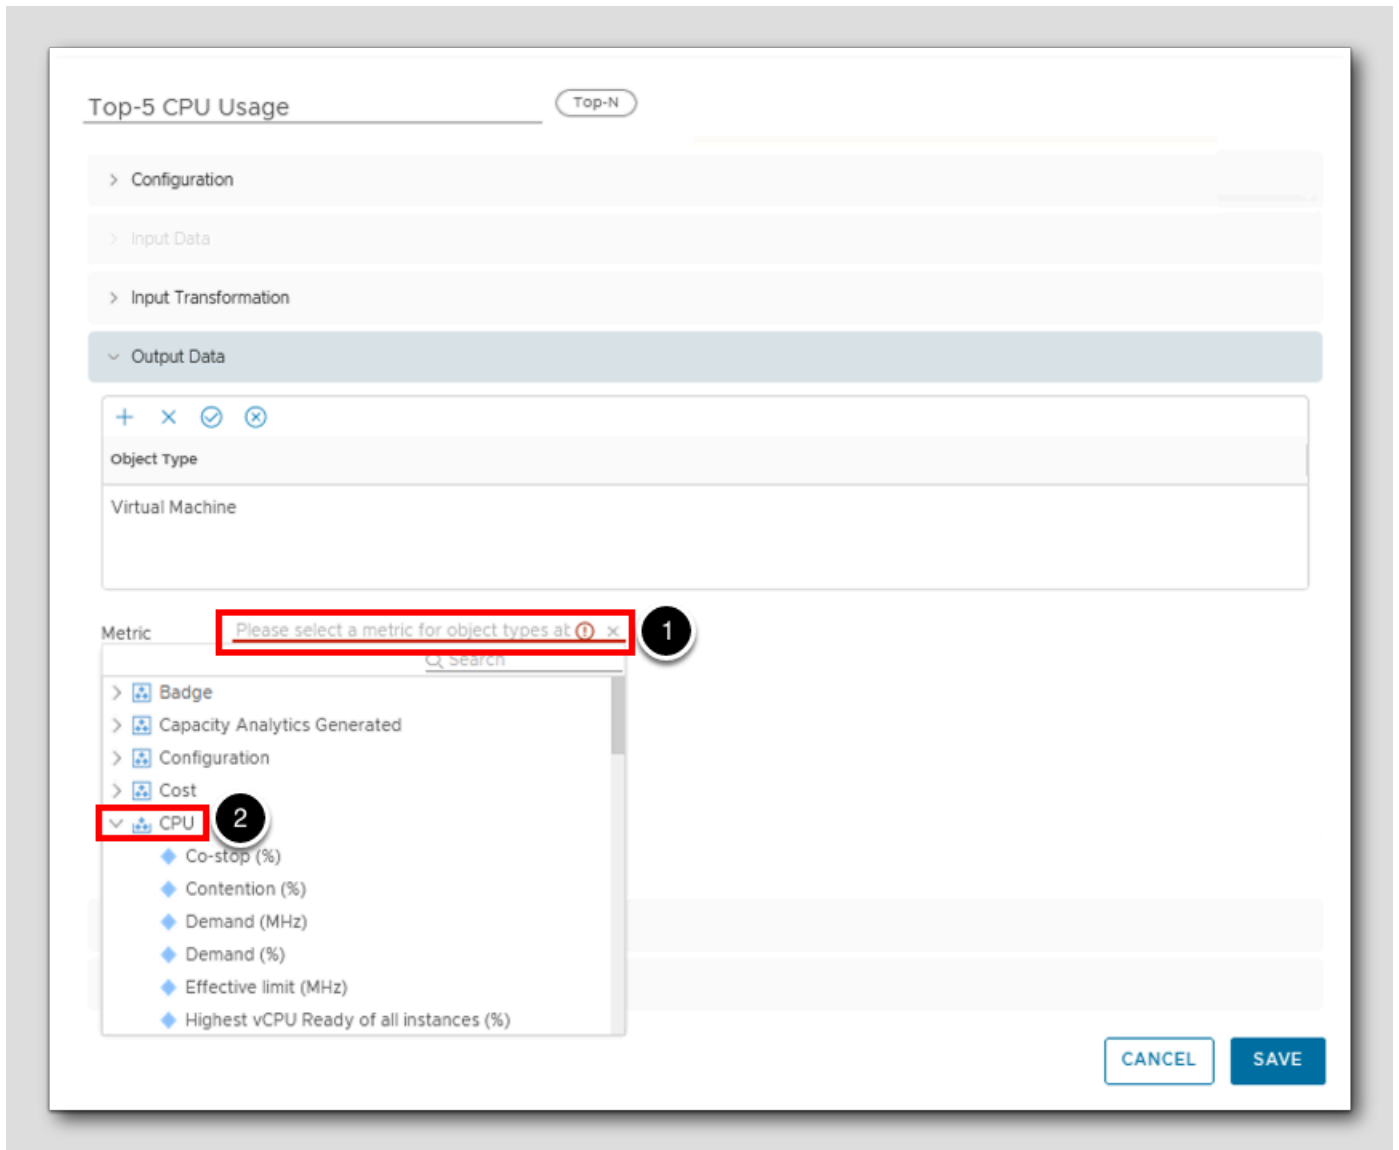
1. Click the + icon to add an Object Type.

## Add Virtual Machine Object Type



1. Type **virtual machine** in the top right search bar and hit **Enter**.
2. Double click on **Virtual Machine**.

## Add CPU Usage (%)



1. Click into the **Metric** search line.

2. Expand CPU.

## Add CPU Usage (%) (Continued)

The screenshot shows the 'Top-5 CPU Usage' configuration window. The 'Output Data' section is expanded, showing a list of metrics for 'Virtual Machine' objects. The 'Usage (%)' metric is highlighted with a red box and a circled '2'. A vertical scrollbar is highlighted with a red box and a circled '1', indicating the scroll action.

Top-5 CPU Usage Top-N

> Configuration

> Input Data

> Input Transformation

▼ Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric Please select a metric for object types at ⓘ x

Q Search

- ◆ Overlap (ms)
- ◆ Provisioned vCPU(s) (vCPUs)
- ◆ Ready (%)
- ◆ Run (ms)
- ◆ Swap wait (%)
- ◆ Total Capacity (MHz)
- ◆ Usage (%)
- ◆ Usage (MHz)
- ◆ Usage average Daily (MHz)
- ◆ vCPU Usage Disparity (%)
- ◆ Workload (%)

CANCEL SAVE

1. Scroll down until you see the Usage (%) metric.
2. Double click on the Usage (%) metric.



### Add CPU Usage (%) (Continued)

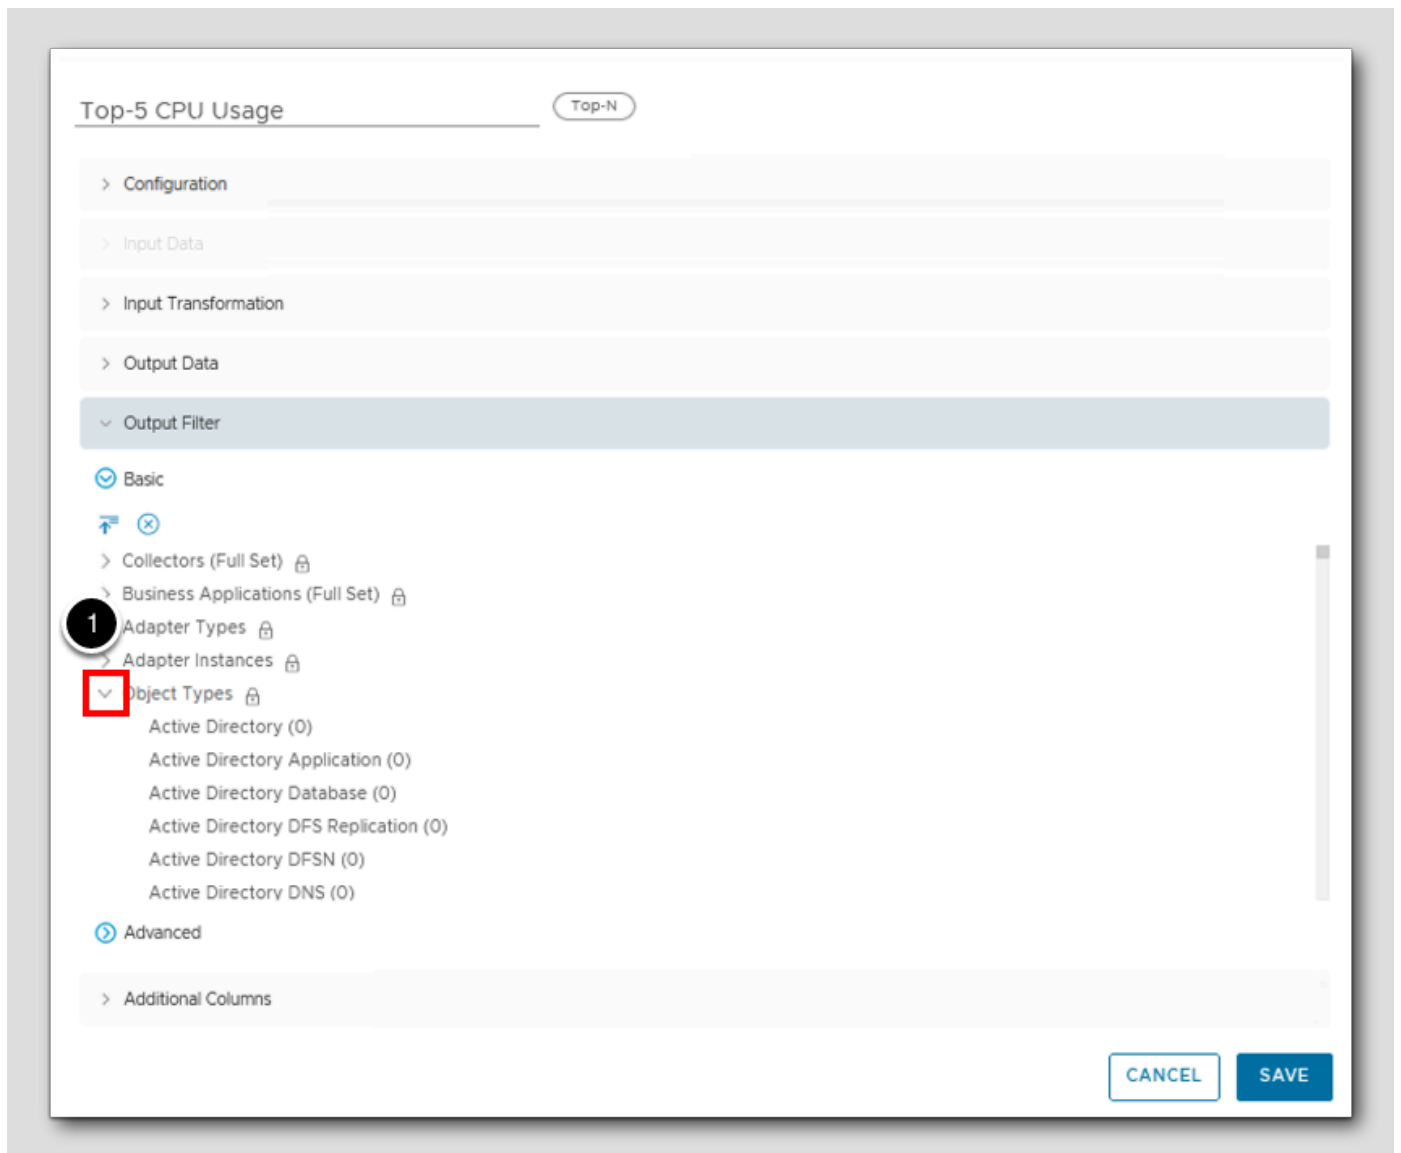
The screenshot shows the configuration interface for 'Top-5 CPU Usage'. The interface includes sections for Configuration, Input Data, Input Transformation, and Output Data. The 'Output Data' section is expanded, showing a table with one row for 'Virtual Machine'. Below the table, the 'Metric' is set to 'CPUUsage (%)', the 'Label' is 'CPU Usage', the 'Unit' is '%', and the 'Maximum' is 'Custom'. The 'Color Method' is set to '75', with '85' and '95' also visible as 'Yellow Bound', 'Orange Bound', and 'Red Bound' respectively. The 'Output Filter' section is also visible. Numbered callouts 1 through 8 highlight the following elements: 1. Metric field (CPUUsage (%)), 2. Label field (CPU Usage), 3. Unit field (%), 4. Maximum field (Custom), 5. Color Method field (75), 6. Yellow Bound field (85), 7. Red Bound field (95), and 8. Output Filter section.

1. Enter **CPU Usage** for the **Label**.
2. Change the unit to **%**.
3. Enter **100** for the **Maximum** value, this will set the graph bar to max out at a value of 100.
4. Change **Color Method** to **Custom**.
5. Enter **75** for **Yellow Bound**.
6. Enter **85** for **Orange Bound**.
7. Enter **95** for **Red Bound**.
8. Expand **Output Filter**.

## Output Filter

[269]

For consistency sake, we will add Virtual Machine as the Output Filter.



1. Expand Object Types.

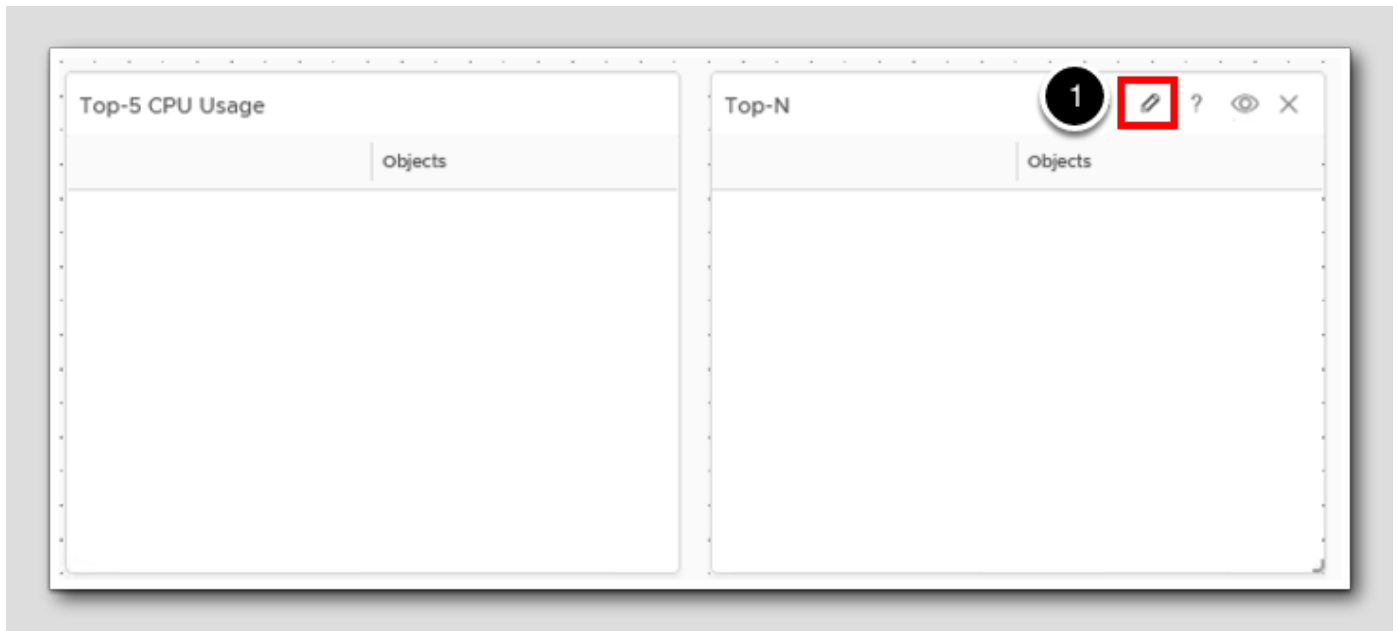
## Output Filter (Continued)

The screenshot displays the 'Top-5 CPU Usage' configuration page. Under the 'Output Filter' section, the 'Basic' tab is active. A list of object types is shown, including 'User (6)', 'vc.procmon process (0)', 'vCenter Server (1)', 'vcsyscmd process (0)', 'Velo Cloud Gateway (0)', 'Velo Cloud Gateway Application (0)', 'Velo Cloud Gateway Processes (0)', 'Velo Cloud Orchestrator (0)', 'Velo Cloud Orchestrator Application (0)', 'Virtual Machine (23)', and 'Virtual Machine Folder (7)'. The 'Virtual Machine (23)' item is selected and highlighted with a red box. A red vertical bar on the right side of the list indicates the scroll action. The 'SAVE' button at the bottom right is also highlighted with a red box. Numbered callouts 1, 2, and 3 point to the scroll bar, the selected 'Virtual Machine' item, and the 'SAVE' button respectively.

1. Scroll down until you see Virtual Machine in the Object Type list.
2. Single click on Virtual Machine Object Type.
3. Click SAVE.

## Add Memory Usage (%)

[27]



1. Hover over the middle **Top-N** widget and click on the pencil icon when it appears.

## Add Memory Usage (%) (Continued)

The screenshot shows the configuration interface for a widget. The widget title is "Top-5 Memory Usage" (highlighted with a red box and labeled 1). The "Configuration" section includes: Refresh Content (Off), Refresh Interval (300 seconds), Self Provider (Off), Redraw Rate (15 minutes), Bars Count (5), Round Decimals (0), and Filter old metrics (Off). The "Top-N Options" section has "Metric Analysis" selected (highlighted with a red box and labeled 2), with a dropdown menu showing options like "Top Highest Utilization". The "Input Data" section has "Input Transformation" expanded (highlighted with a red box and labeled 3). At the bottom right are "CANCEL" and "SAVE" buttons.

1. Change the top name to **Top-5 Memory Usage**.
2. Check the **Metric Analysis** radio button because we want to display the Memory Usage (%) metric in this widget.
3. Expand **Input Transformation**.

## Input Transformation

Top-5 Memory Usage Top-N

> Configuration

> Input Data

▼ Input Transformation 1

Relationship  Parent  Children  Self Depth: 10

2 > Output Data

> Output Filter

> Additional Columns

CANCEL SAVE

1. Uncheck the Self Input Transformation.
2. Expand Output Data.

## Output Data

Top-5 Memory Usage Top-N

> Configuration

> Input Data

> Input Transformation

**1** Output Data

+ × ✓ ⊗

Object Type

Metric	Please select a metric for object types above x
Label	
Unit	
Maximum	None
Color Method	

> Output Filter

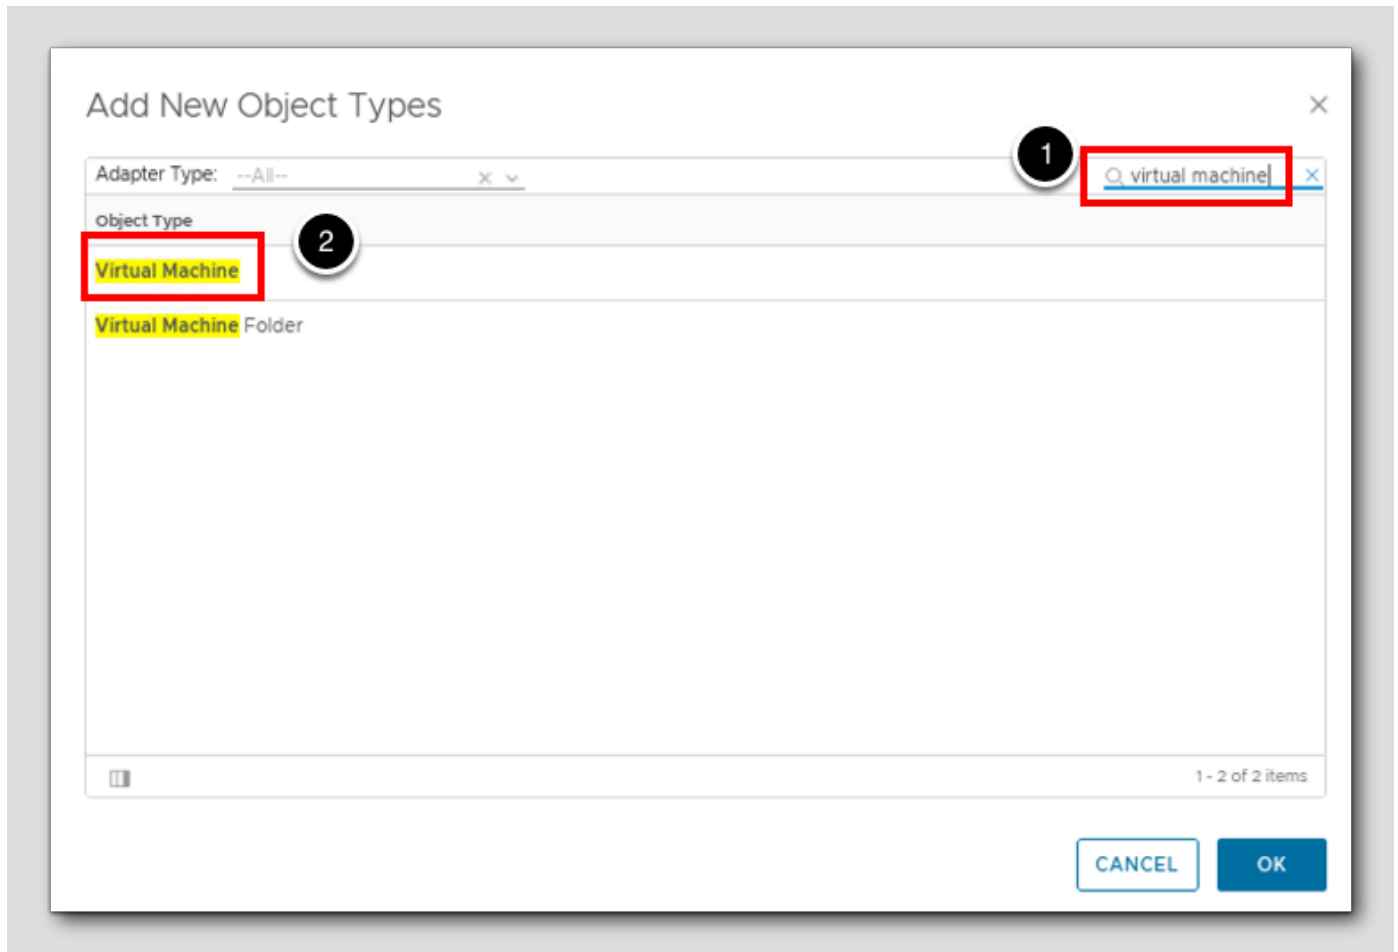
> Additional Columns

CANCEL SAVE

1. Click the + icon to add an Object Type.



## Add Virtual Machine Object Type



1. Type **virtual machine** in the top right search bar and hit **Enter**.
2. Double click on **Virtual Machine**.

## Add Memory Usage (%) (Continued)

The screenshot shows the 'Top-5 Memory Usage' configuration page. The 'Output Data' section is expanded, and a search bar for metrics is highlighted with a red box and labeled '1'. The search bar contains the text 'Please select a metric for object types at'. Below the search bar, a list of metrics is shown, with 'Memory' expanded and its sub-metrics listed. A red box highlights the 'Memory' category and is labeled '3'. A vertical scroll bar is highlighted with a red box and labeled '2'. The 'Memory' sub-metrics include: Balloon (%), Compressed (KB), Consumed (KB), Consumed (%), and Consumed average Daily (KB). At the bottom right, there are 'CANCEL' and 'SAVE' buttons.

1. Click into the **Metric** search line.
2. Scroll down until you see the **Memory** metric list.
3. Expand the **Memory** metric list.

## Add Memory Usage (%) (Continued)

The screenshot shows the 'Top-5 Memory Usage' configuration window. The 'Output Data' section is expanded, showing a list of metrics for 'Virtual Machine' objects. The 'Usage (%)' metric is highlighted with a red box and a circled '2'. A red vertical bar highlights the scroll area, with a circled '1' indicating the scroll action.

1. Scroll down until you see the Usage (%) metric.
2. Double click on the Usage (%) metric.

### Add Memory Usage (%) (Continued)

The screenshot shows the configuration interface for 'Top-5 Memory Usage'. The interface includes sections for Configuration, Input Data, Input Transformation, and Output Data. The 'Output Data' section is expanded, showing a table with 'Object Type' set to 'Virtual Machine'. Below the table, the 'Metric' is 'MemoryUsage (%)', 'Label' is 'Memory Usage', 'Unit' is '%', and 'Maximum' is 'Custom'. The 'Color Method' is set to '75' with a 'Bound' of '95'. The 'Output Filter' section is also visible, with an 'Additional Columns' section below it. The 'CANCEL' and 'SAVE' buttons are at the bottom right.

Top-5 Memory Usage Top-N

> Configuration

> Input Data

> Input Transformation

▼ Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric MemoryUsage (%) ×

Label Memory Usage 1

Unit % 2

Maximum 100 3

Color Method Custom 4

75 5 Bound 85 6 Bound 95 7 Bound

8

> Output Filter

> Additional Columns

CANCEL SAVE

1. Enter **Memory Usage** for the **Label**.
2. Change the unit to **%**.
3. Enter **100** for the **Maximum** value, this will set the graph bar to max out at a value of 100.
4. Change **Color Method** to **Custom**.
5. Enter **75** for **Yellow Bound**.
6. Enter **85** for **Orange Bound**.
7. Enter **95** for **Red Bound**.
8. Expand **Output Filter**.

## Output Filter

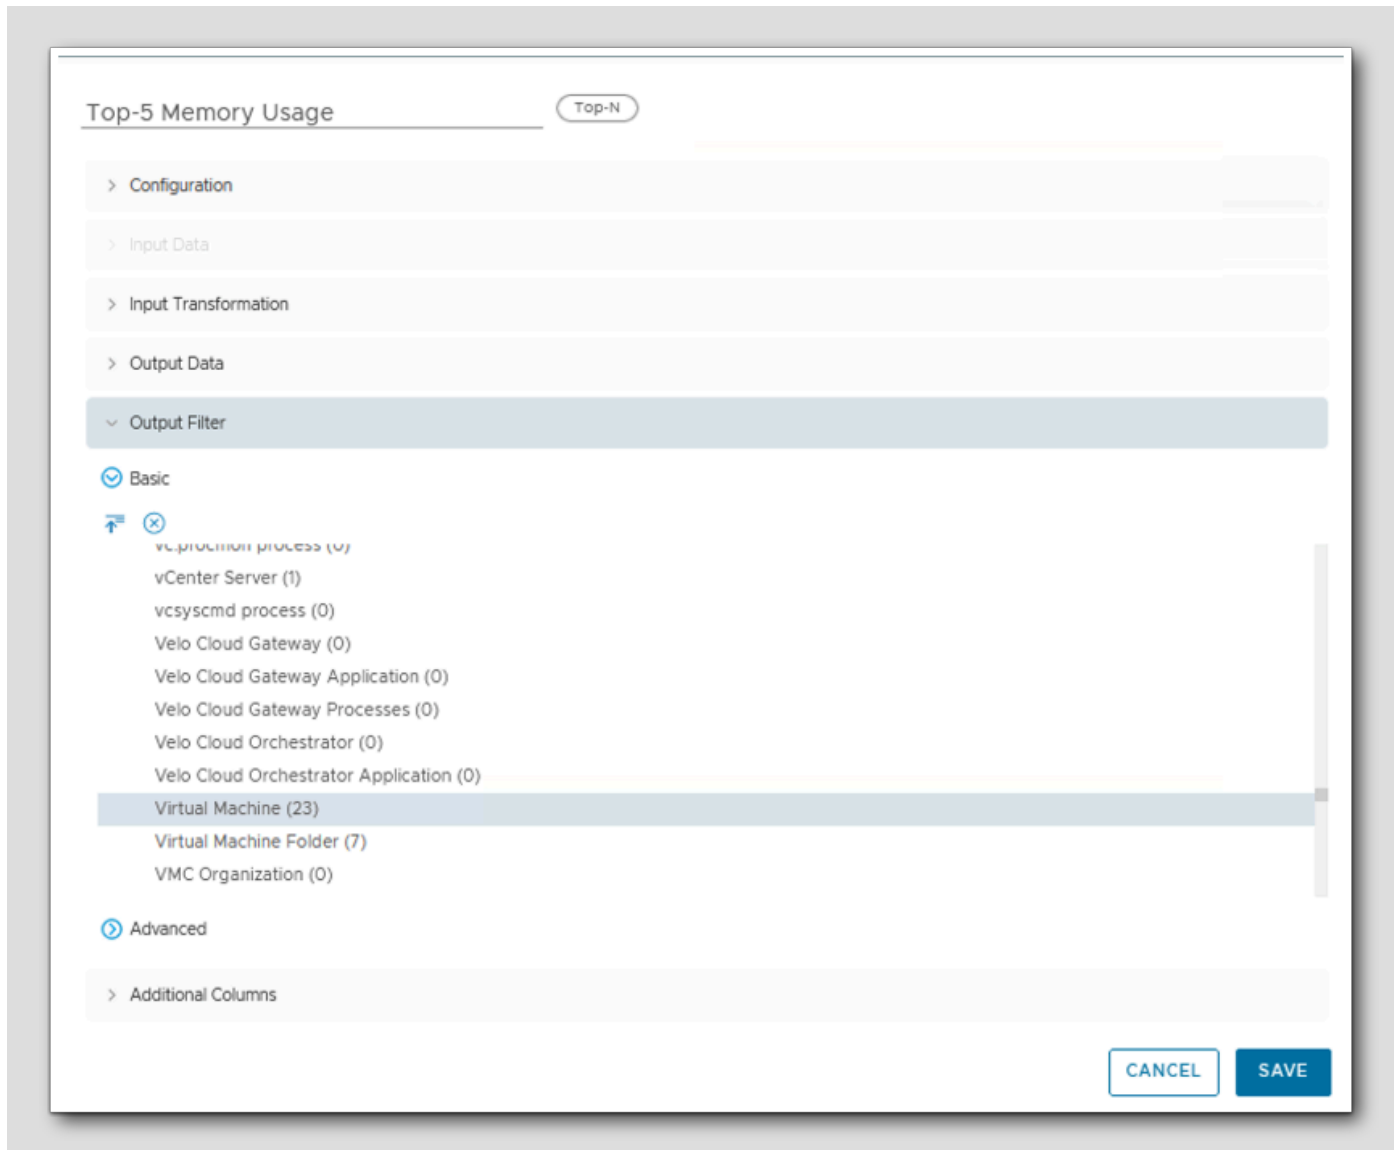
Top-5 Memory Usage Top-N

- > Configuration
- > Input Data
- > Input Transformation
- > Output Data
- ▼ Output Filter
  - ☑ Basic
    - ⌵ Collectors (Full Set) ⌵
    - ⌵ Business Applications (Full Set) ⌵
    - ⌵ Adapter Types ⌵
    - ⌵ Adapter Instances ⌵
    - ⌵ Object Types ⌵
      - Active Directory (0)
      - Active Directory Application (0)
      - Active Directory Database (0)
      - Active Directory DFS Replication (0)
      - Active Directory DFSN (0)
      - Active Directory DNS (0)
  - ⌵ Advanced
- > Additional Columns

CANCEL SAVE

1. Expand Object Types.

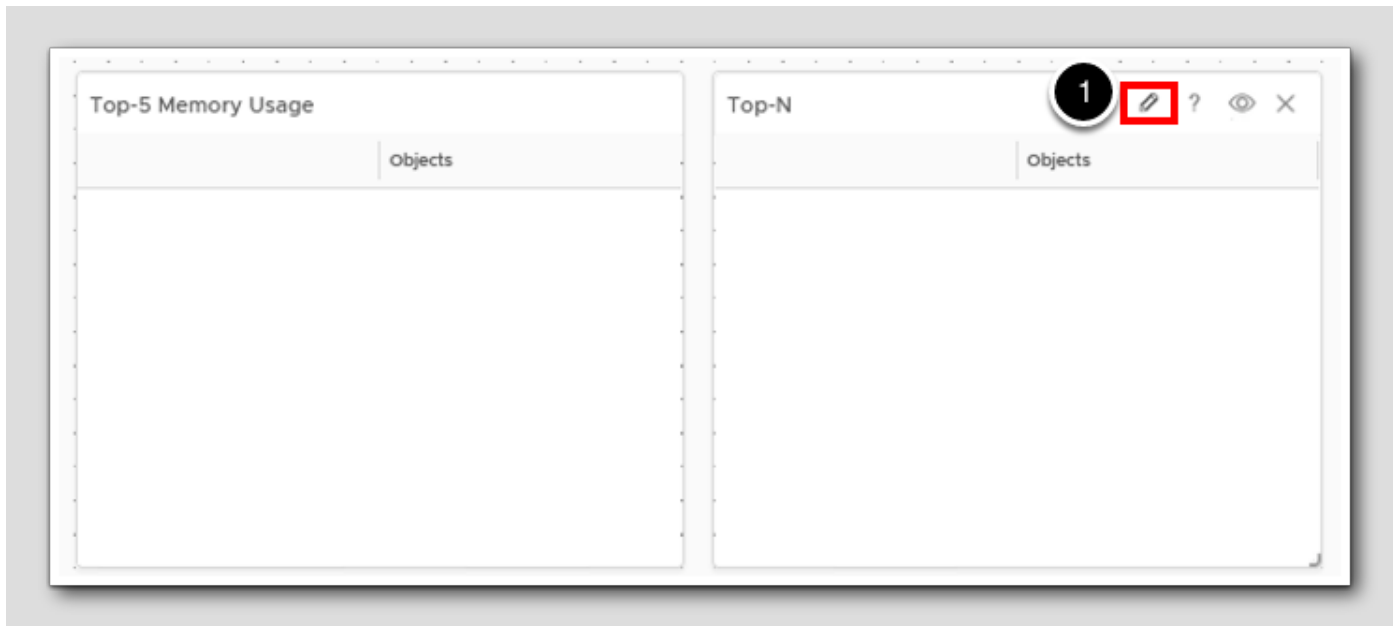
## Output Filter (Continued)



1. Scroll down until you see **Virtual Machine** in the Object Type list.
2. Single click on **Virtual Machine** Object Type.
3. Click **SAVE**.

## Configure Top-5 Storage Usage

[281]



1. Hover over the far right **Top-N** widget and click on the pencil icon when it appears.



## Add Storage IOPS

The screenshot shows the configuration interface for a widget titled "Top-5 IOPS Usage". The interface includes a "Configuration" section with various settings:

- Refresh Content:  On  Off
- Refresh Interval: 300 (seconds)
- Self Provider:  On  Off
- Redraw Rate: 15 minutes
- Bars Count: 5
- Round Decimals: 0
- Filter old metrics:  On  Off
- Top-N Options:  Application Health and Performance  Metric Analysis
  - Top Highest Utilization
  - Top Lowest Utilization
  - Top Abnormal States
  - Top Highest Volatility
  - Percentile %

Below the configuration section, there are expandable sections for "Input Data", "Input Transformation", "Output Data", and "Output Filter". The "Input Transformation" section is expanded, and the "Input Data" section is also expanded. The "CANCEL" and "SAVE" buttons are located at the bottom right of the configuration panel.

1. Change the top name to Top-5 IOPS Usage.
2. Check the **Metric Analysis** radio button because we want to display the **Highest IOPS of all instances** metric in this widget.
3. Expand **Input Transformation**.

## Input Transformation

Top-5 IOPS Usage Top-N

> Configuration

> Input Data

▼ Input Transformation 1

Relationship  Parent  Children  Self Depth: 10

2 > Output Data

> Output Filter

> Additional Columns

CANCEL SAVE

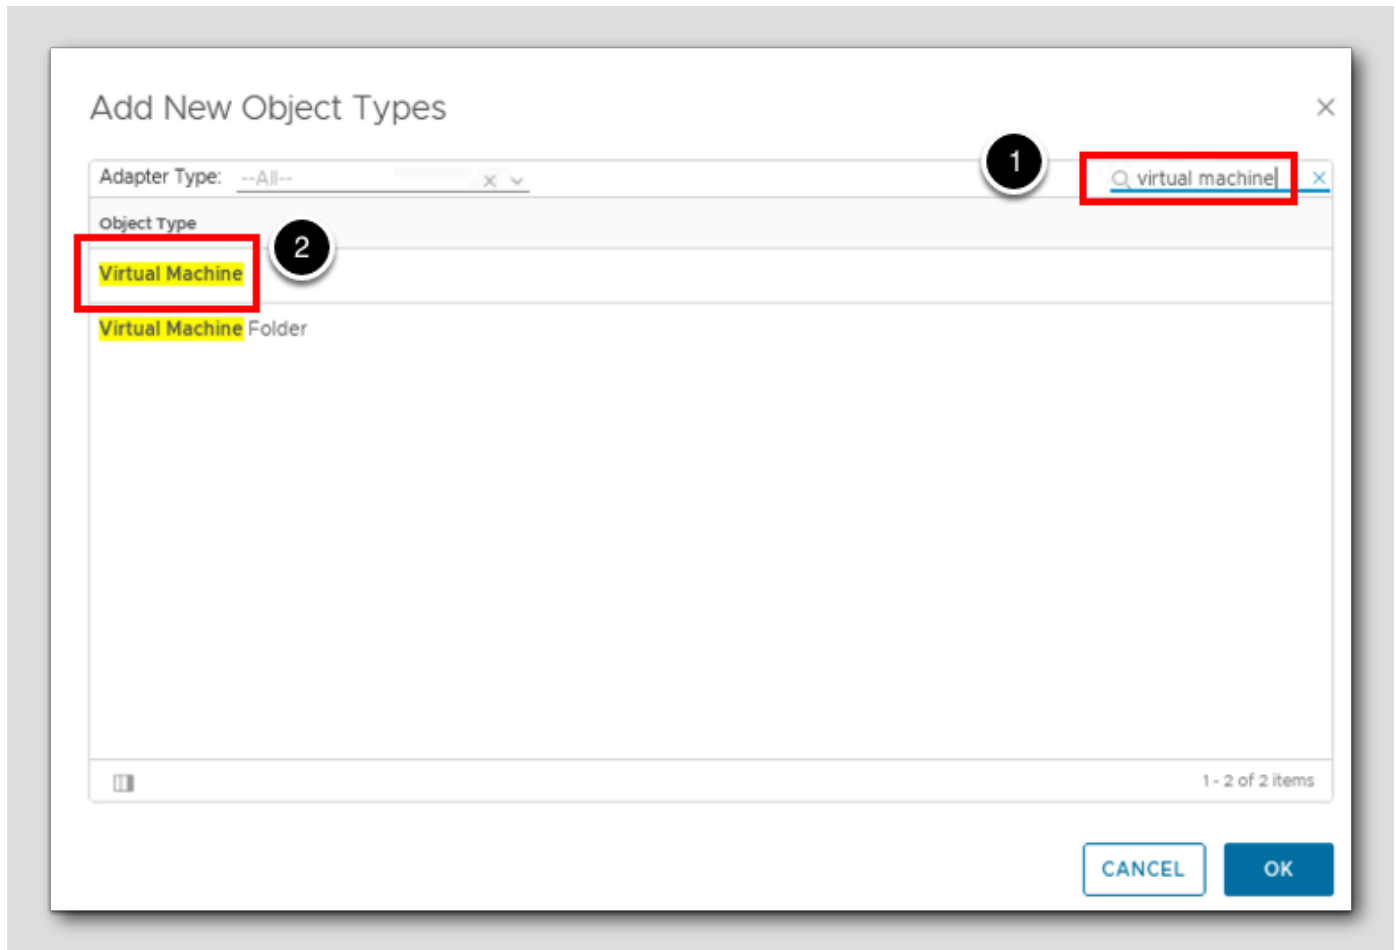
1. Uncheck the Self Input Transformation.
2. Expand Output Data.

## Output Data

The screenshot shows a configuration window titled "Top-5 IOPS Usage" with a "Top-N" button. The interface is divided into several sections: Configuration, Input Data, Input Transformation, Output Data, Output Filter, and Additional Columns. The "Output Data" section is currently active and highlighted in blue. A red circle with the number "1" is placed over a "+" icon in a toolbar within the "Output Data" section. Below the toolbar, there is a text input field for "Object Type". Underneath, there are several rows of configuration options: "Metric" with a red error message "Please select a metric for object types at 1 x", "Label", "Unit" (a dropdown menu), "Maximum" (a dropdown menu set to "None"), and "Color Method". At the bottom right of the window, there are "CANCEL" and "SAVE" buttons.

1. Click the + icon to add an Object Type.

## Add Virtual Machine Object Type



1. Type **virtual machine** in the top right search bar.
2. Double click on **Virtual Machine**.

## Add Storage IOPS (Continued)

The screenshot shows the 'Top-5 IOPS Usage' configuration page. The 'Output Data' section is expanded, and the 'Metric' search bar is highlighted with a red box and a '1' callout. Below the search bar, the 'Virtual Disk' category is expanded, and the 'Highest IOPS of all instances' metric is highlighted with a red box and a '4' callout. The scroll bar for the metric list is also highlighted with a red box and a '2' callout. The 'Virtual Disk' category is highlighted with a red box and a '3' callout. At the bottom right, there are 'CANCEL' and 'SAVE' buttons.

1. Click into the **Metric** search line.
2. Scroll down until you see the **Virtual Disk** metric list.
3. Expand the **Virtual Disk** metric list.
4. Double click on the **Highest IOPS of all instances** metric.

## Add Storage IOPS (Continued)

Top-5 IOPS Usage Top-N

> Configuration

> Input Data

> Input Transformation

▼ Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric

Label  1

Unit

Maximum  2

Color Method

3 Yellow Bound

4 Orange Bound

5 Red Bound

> Output Filter

6 Additional Columns

CANCEL SAVE

1. Enter Highest IOPS for the Label.
2. Change Color Method to Custom.
3. Enter 75 for Yellow Bound.
4. Enter 85 for Orange Bound.
5. Enter 95 for Red Bound.
6. Expand Output Filter.

## Output Filter

The screenshot shows the 'Top-5 IOPS Usage' configuration window. The 'Output Filter' section is expanded, showing a list of categories. 'Object Types' is highlighted with a red box and a circled '1', indicating it should be expanded. The list includes 'Active Directory' and its sub-items, all with a count of 0. 'Advanced' and 'Additional Columns' are also visible.

Top-5 IOPS Usage Top-N

- > Configuration
- > Input Data
- > Input Transformation
- > Output Data
- ▼ Output Filter
  - ☑ Basic
    - ⚙️
    - > Collectors (Full Set) 🔒
    - > Business Applications (Full Set) 🔒
    - 1. Object Types 🔒
      - Active Directory (0)
      - Active Directory Application (0)
      - Active Directory Database (0)
      - Active Directory DFS Replication (0)
      - Active Directory DFSN (0)
      - Active Directory DNS (0)
  - ⌄ Advanced
- > Additional Columns

CANCEL SAVE

1. Expand Object Types.

## Output Filter (Continued)


The screenshot displays the 'Top-5 IOPS Usage' configuration interface. The 'Output Filter' section is expanded, showing a list of object types. The 'Virtual Machine (23)' item is selected and highlighted with a red box, with a '2' callout. A vertical scrollbar on the right side of the list is highlighted with a red box and a '1' callout. At the bottom right, the 'SAVE' button is highlighted with a red box and a '3' callout. Other callouts include a 'Top-N' button at the top right and a 'Basic' tab at the bottom left.

1. Scroll down until you see **Virtual Machine** in the Object Type list.
2. Single click on **Virtual Machine** Object Type.
3. Click **SAVE**.



## Edit Dashboard - Save

[290]



1. Click on the **SAVE** button at the top of the dashboard to save our changes to the dashboard.

## HOL Inventory Summary - Modified Dashboard

HOL Inventory Summary ACTIONS ▾

1H 6H 24H 7D CUSTOM ☆ 🔗 ?

Summary

18 Total VM 16 Running VM 5 Hosts 2 Clusters 7 Datastores 1 Datacenters

Datacenters

Name	Clusters
vSphere World	2
RegionA01	2

2

1 - 2 of 2 items

Compute

Name	CPU Capacity
Management	41.9 GHz
Workload1	25.14 GHz

1 - 2 of 2 items

Storage

Name	Capacity
RegionA01-ISCSI01-...	0.24 TB
Total	0.24 TB

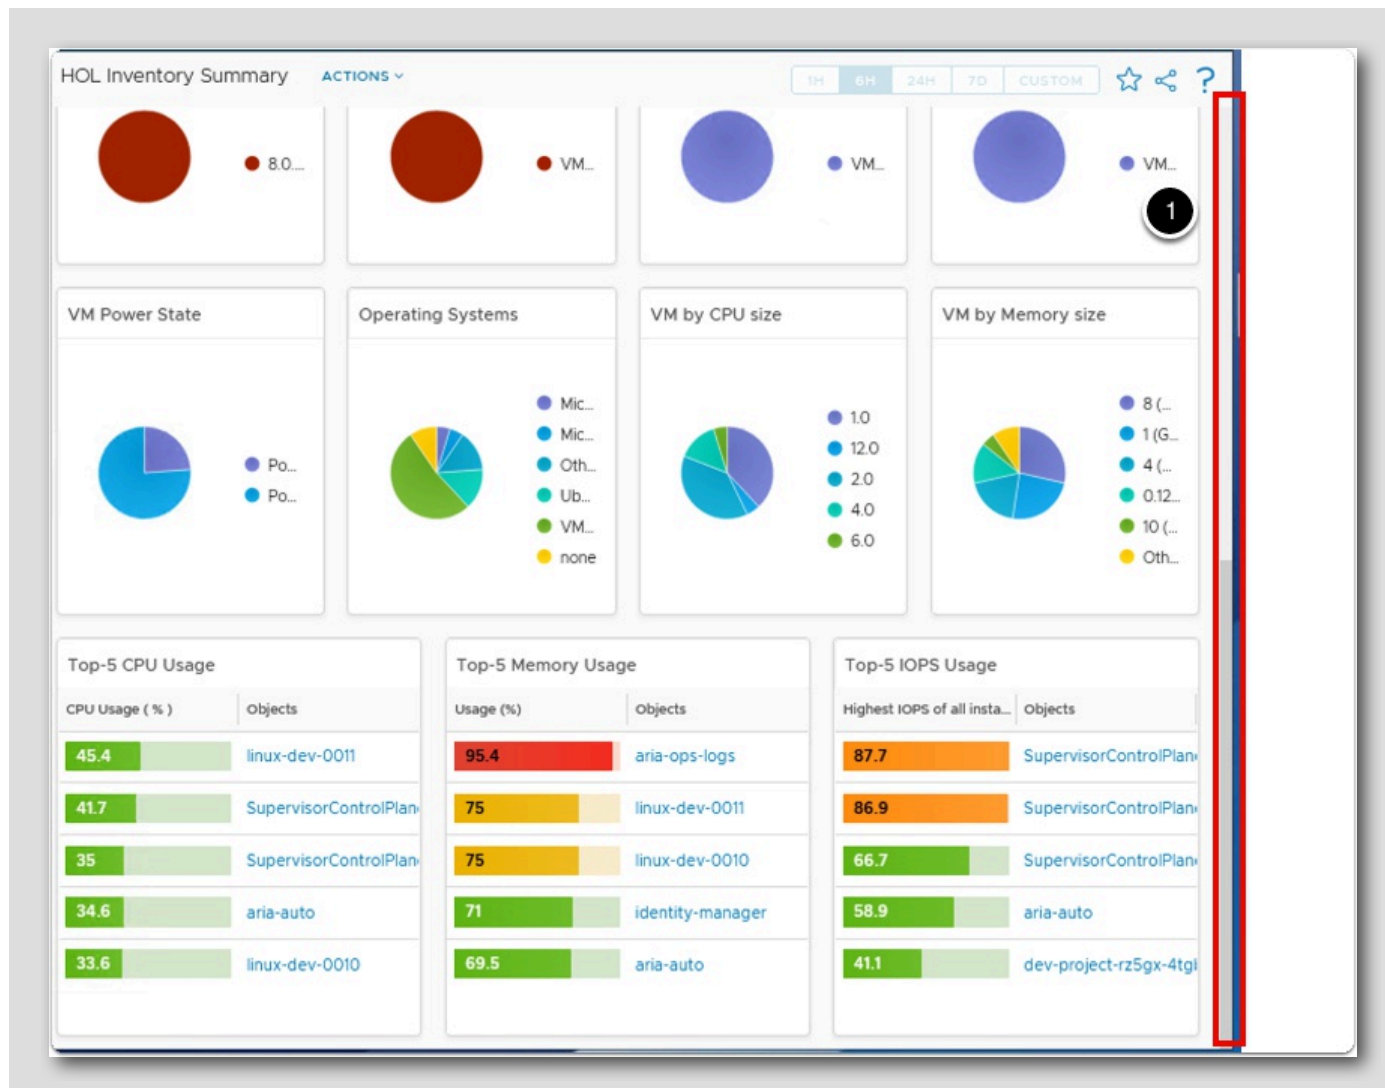
1 - 1 of 1 items

ESXi version 8.0... Server Model VM... Server Vendor VM... Datastore Type VM...

1. Scroll to the top of the HOL Inventory Dashboard.
2. Click on the RegionA01 row (not on the RegionA01 text).

## Completed Dashboard

[292]



1. Drag the scroll bar down to the very bottom of the dashboard.

We can now see the Top-N widgets we added each showing the top consumers of CPU, Memory and IOPS.

## Lesson End

[293]

Congratulations, we just completed the Clone and Modify Existing Dashboards lesson!

In this lesson, we started out by cloning the Inventory Summary dashboard and then customized the cloned dashboard. We added three Top-5 widgets to show the top CPU, Memory and IOPS troublemakers.

## Creating a New Custom Dashboard

[294]

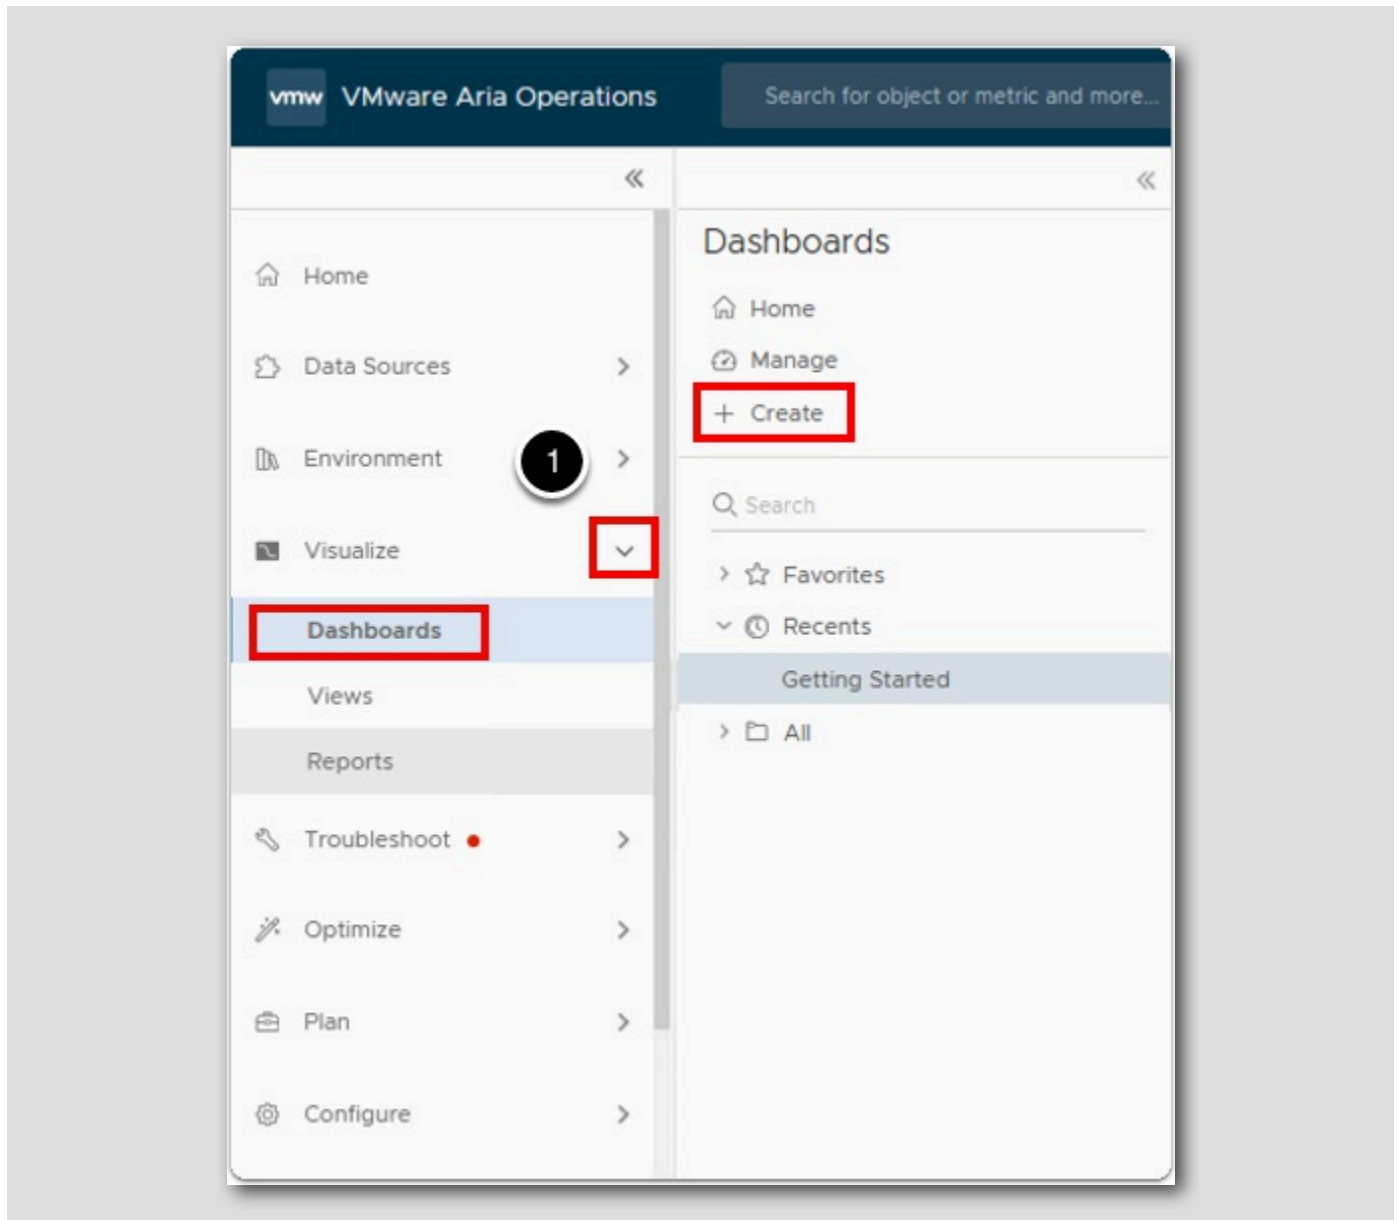
In this lesson, we will learn how to create a new dashboard from scratch.

We will create a brand new dashboard from scratch that will contain an Object List for a list of virtual machines. We will then add the following widgets to the dashboard as well:

- Object Relationship Topology
- Top Alerts
- Health Heat Map
- Top-N for CPU
- Top-N for Memory
- Top-N for Disk Space

## Dashboards

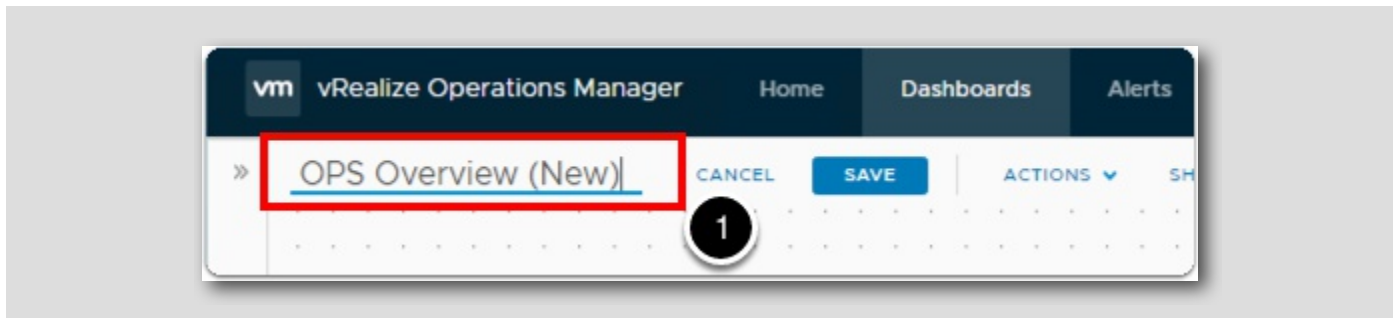
[295]



1. Expand **Visualize** from the left hand toolbar.
2. Click on **Dashboards**.
3. Click **+ Create**.

Create Dashboard - Name

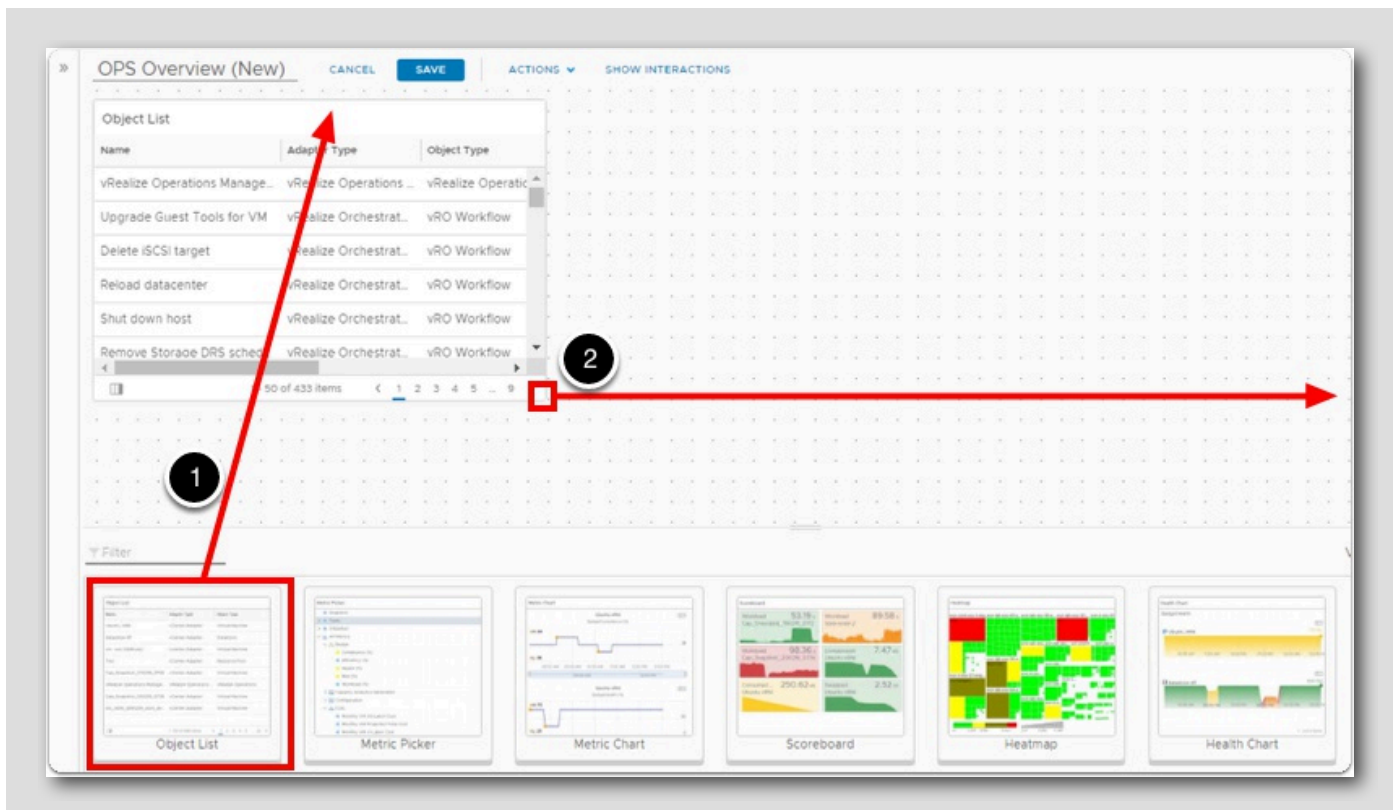
[296]



1. Replace the New Dashboard text with OPS Overview (New) in the name text field.

Create Dashboard - Object List

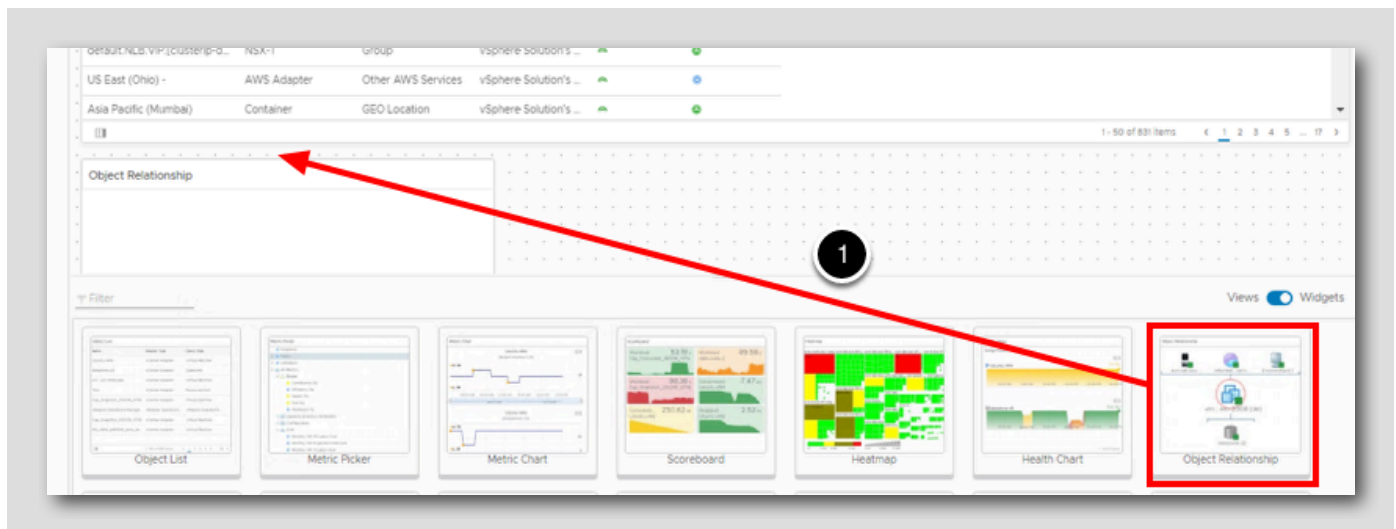
[297]



1. Click on and drag the Object List widget to the left side of the open space.
2. Click on the lower right-hand corner of the widget and drag it all the way to the right of the dashboard interface.

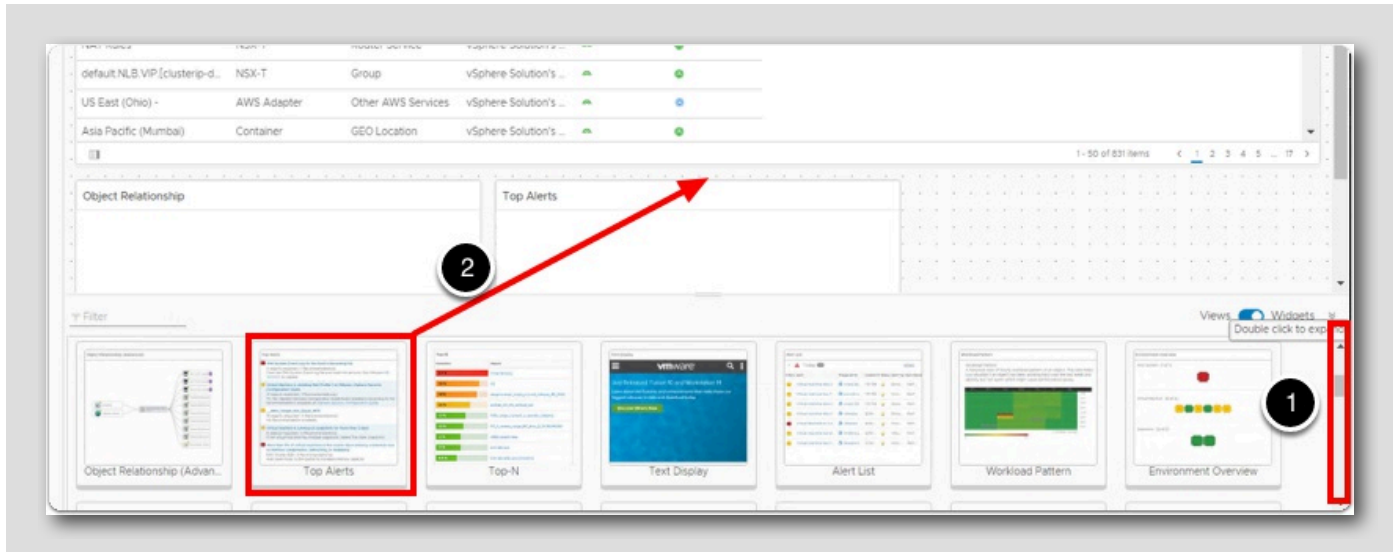
## Create Dashboard - Object Relationship

[298]



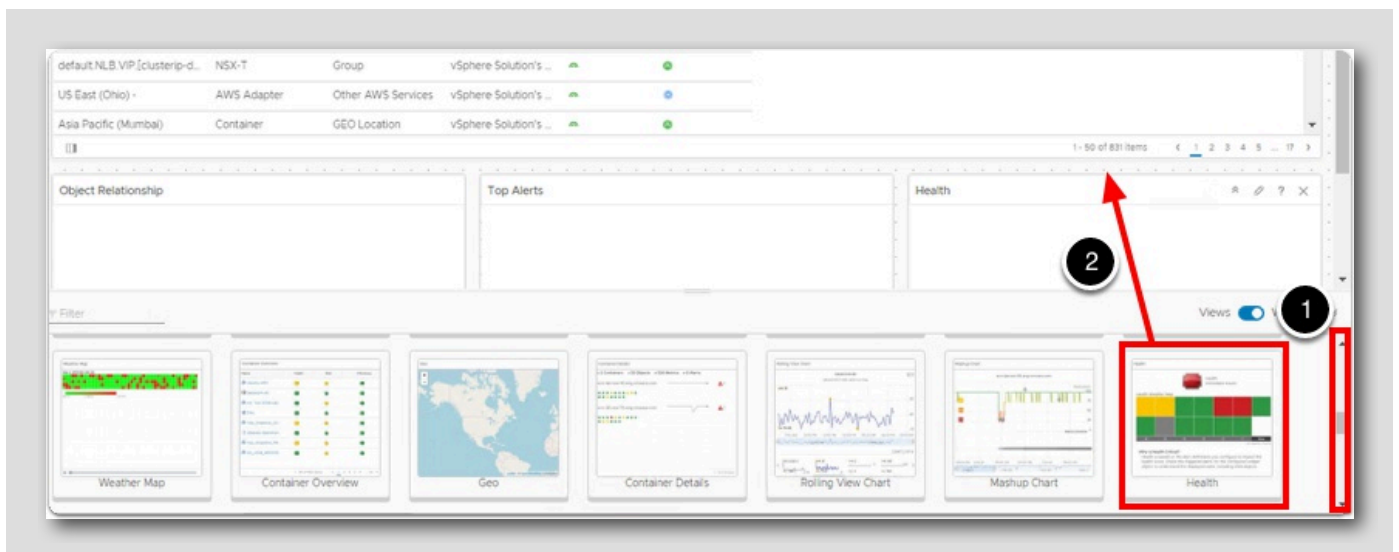
1. Click on and drag the Object Relationship widget to the left most column in the open space below the Object List widget. (You may need to scroll down a level to see it.)

### Create Dashboard - Top Alerts



1. Drag the scroll bar down until we see the Top Alerts widget in the list. (should be the second row of new widgets)
2. Click on and drag the Top Alerts widget to the middle column in the open space below the Object List widget.

### Create Dashboard - Health

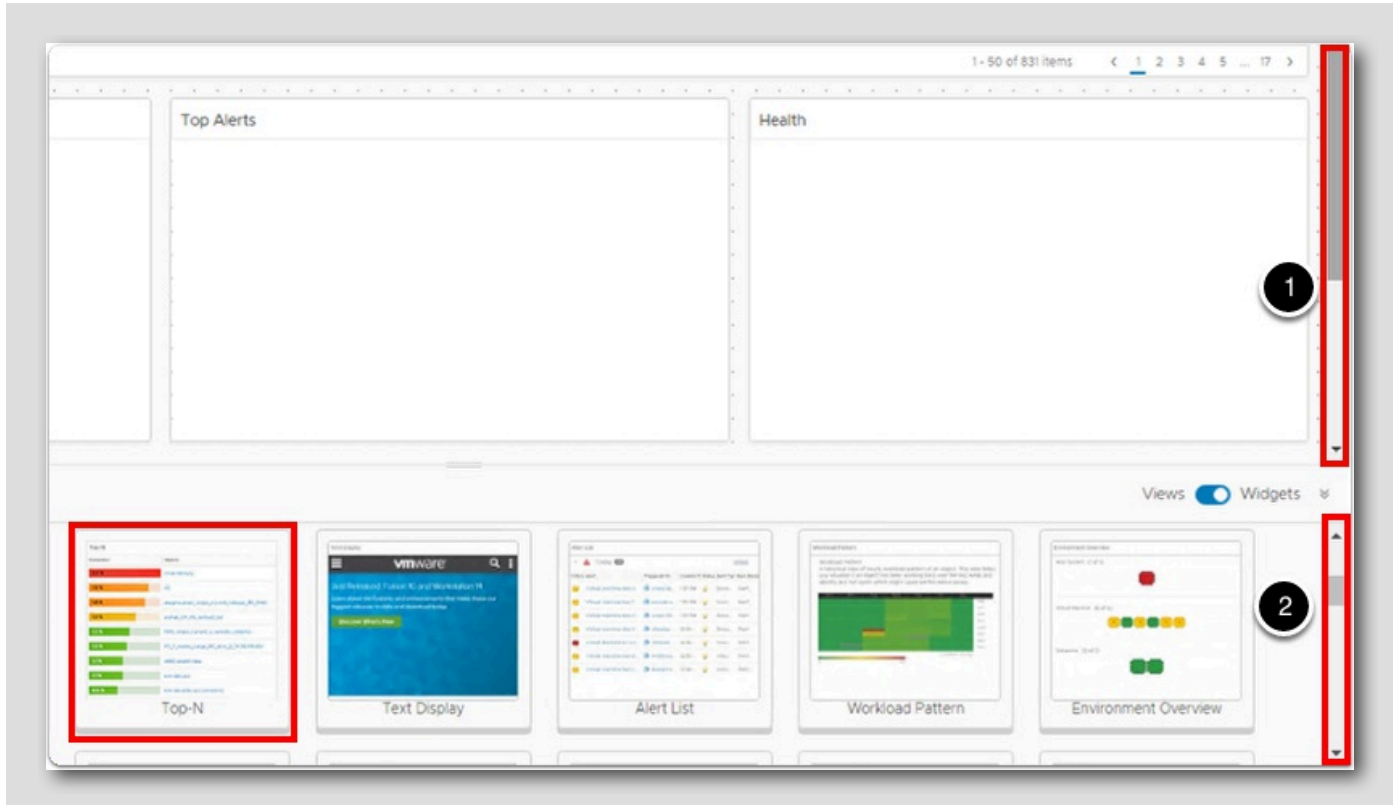




1. Drag the scroll bar down until we see the Health widget in the list. (should be on the fourth or fifth row of new widgets)
2. Click on and drag the Health widget (not the Scoreboard Health widget) to the far right column in the open space below the Object List widget.

## Create Dashboard - Top-N

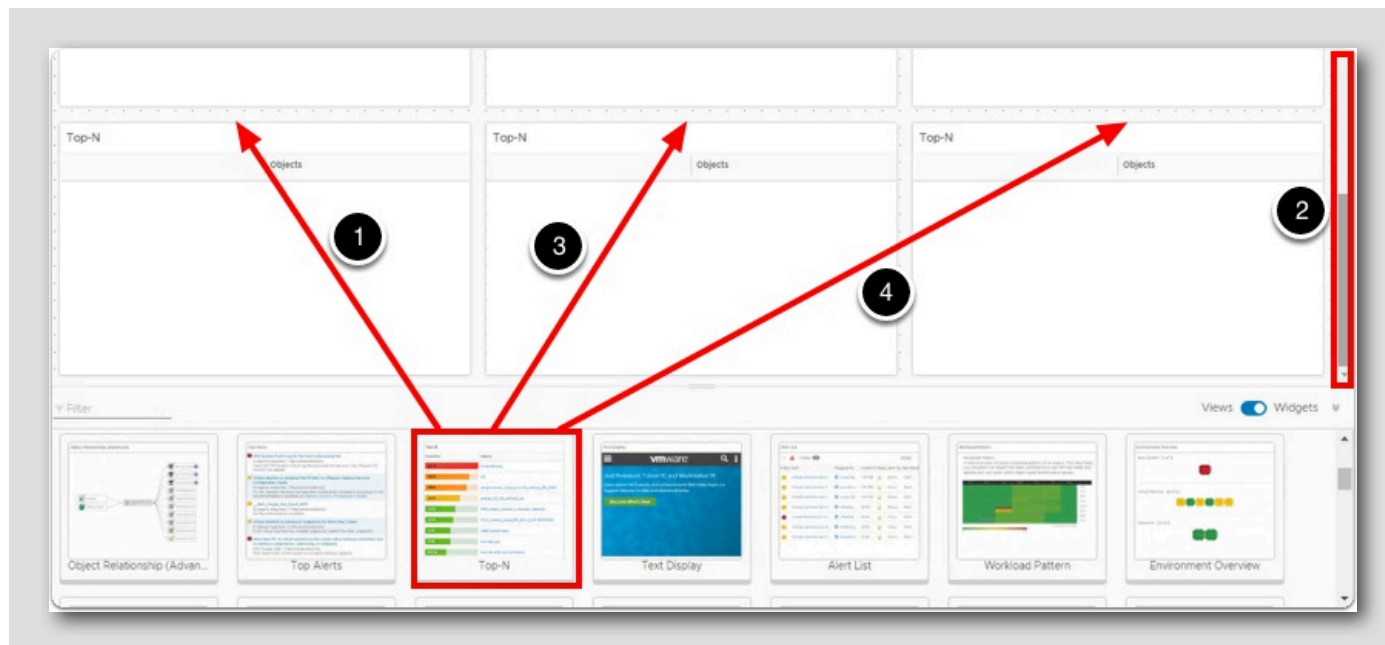
[301]



1. Drag the scroll bar all the way down to the bottom.
2. Drag the widgets scroll bar up until we can see the Top-N widget in the second row of widgets.

## Create Dashboard - Top-N

[302]



1. Click on the Top-N view and drag it to the left column.
2. Drag the scroll bar all the way down to the bottom so we can see the new Top-N view.
3. Click on the Top-N view and drag it to the middle column.
4. Click on the Top-N view and drag it to the right column.

## Create Dashboard - Show Interactions

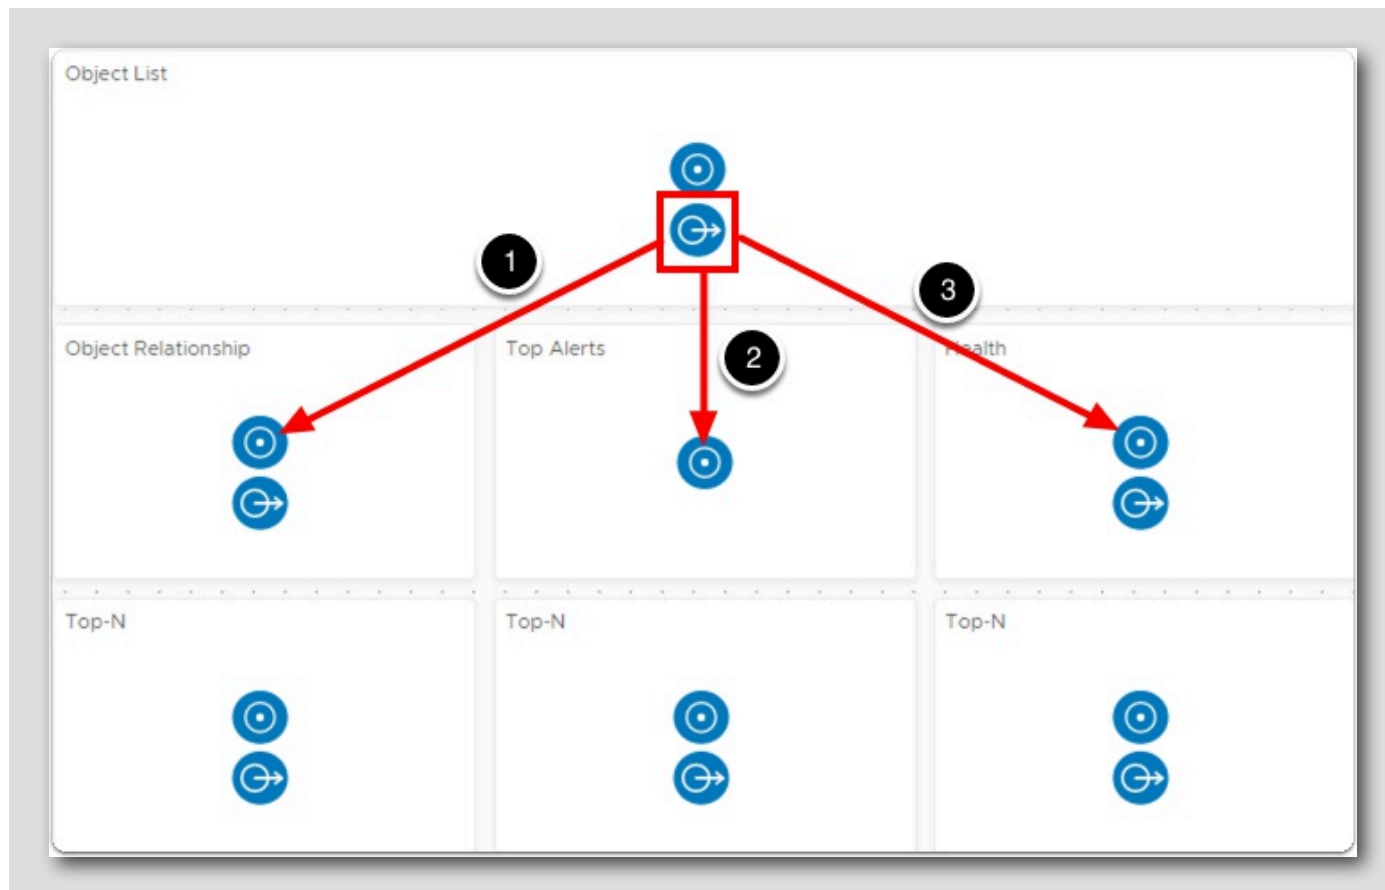
[303]



1. Click on SHOW INTERACTIONS text link at the top of the user interface.

## Create Dashboard - Create Relationships

[304]

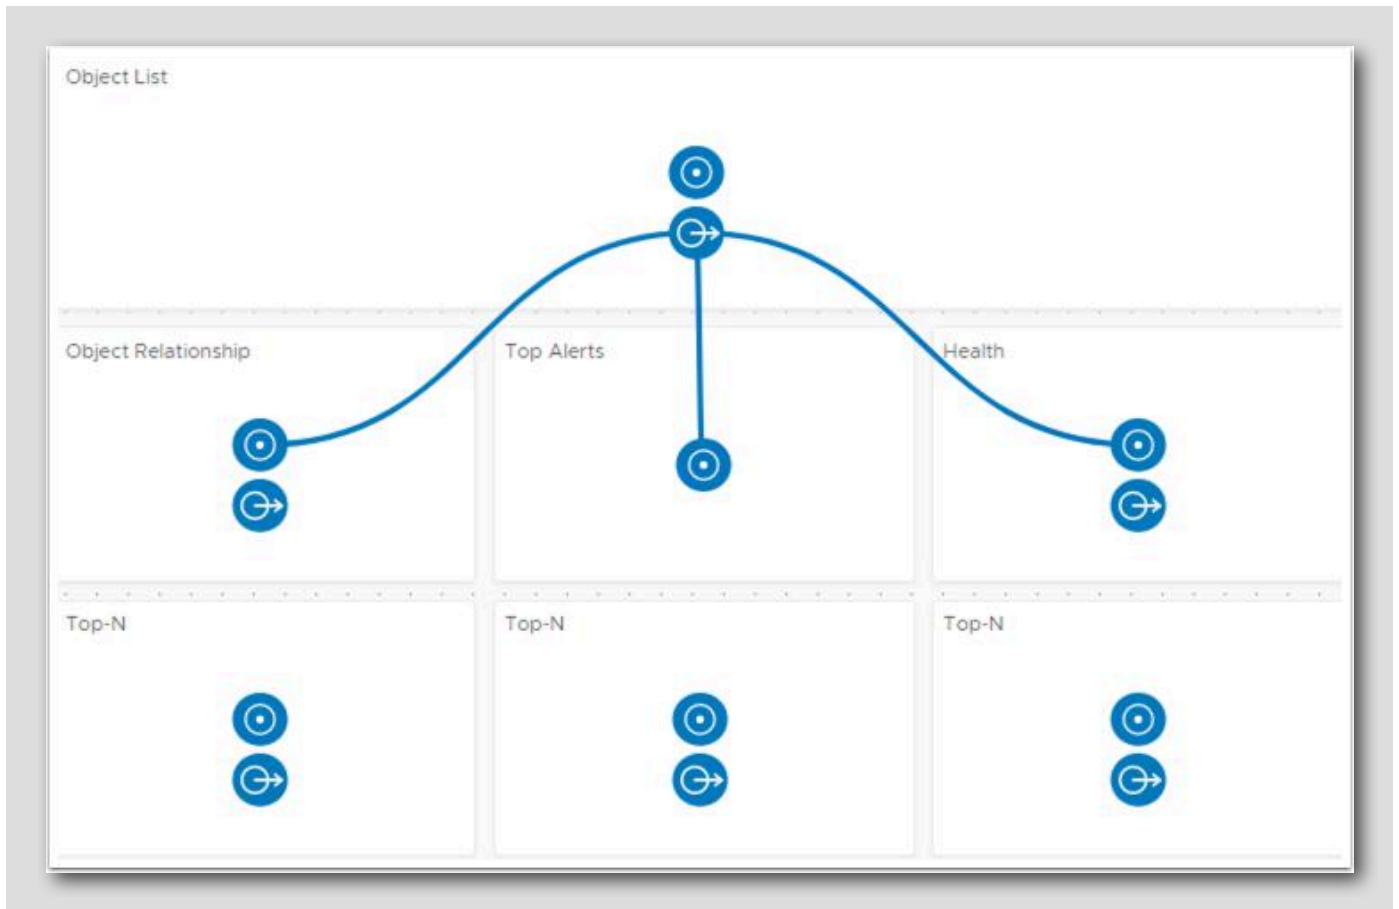


We now have to create the relationships between the widgets. We want to be able to click on a virtual machine in the Object List widget and have the rest of the widgets present the data associated with what we selected in the Object List.

1. Click on the circle and arrow icon in the Object List widget and drag it to the circle with a dot icon in the Object Relationship widget.
2. Click on the circle and arrow icon in the Object List widget and drag it to the circle with a dot icon in the Top Alerts widget.
3. Click on the circle and arrow icon in the Object List widget and drag it to the circle with a dot icon in the Health widget.

## Create Dashboard - Create Relationships (continued)

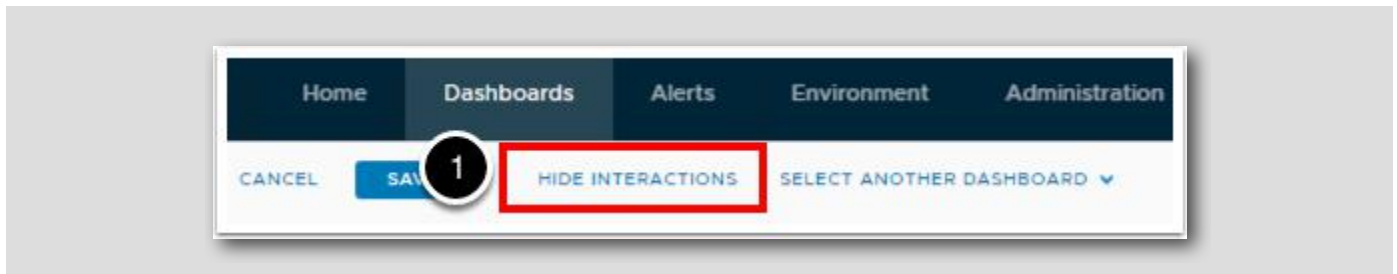
[305]



After completing the previous steps, we should now see the connecting line from the Object List to the Object Relationship, Metric Chart and Health widgets. We will not be connecting the Object List to the (3) Top-N widgets since we want them to show the Top 10 virtual machines with contention for CPU, Memory and Disk Space. We will see this later once we are done configuring the entire dashboard.

## Create Dashboard - Show Interactions

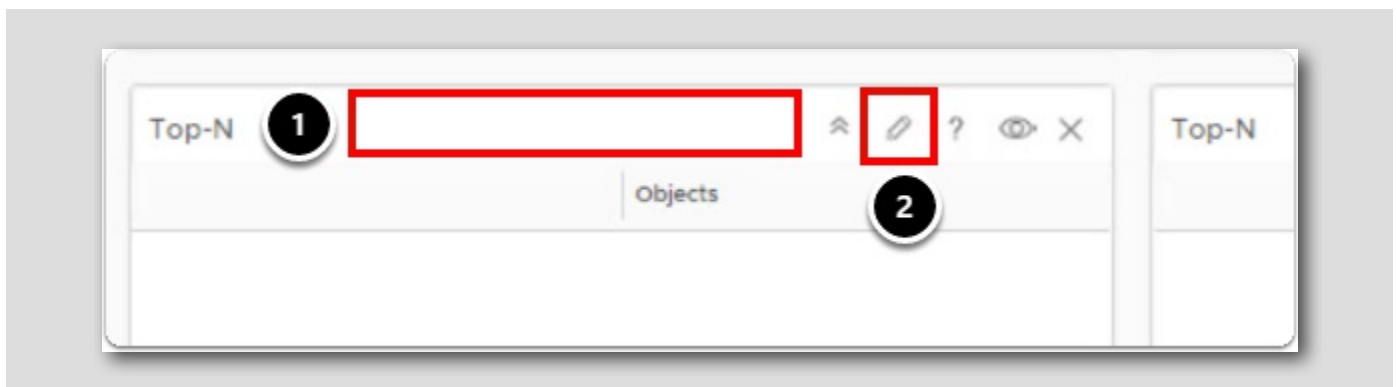
[306]



1. Click on **HIDE INTERACTIONS** text link at the top of the user interface.

## Configure Top-N Widget - Menu

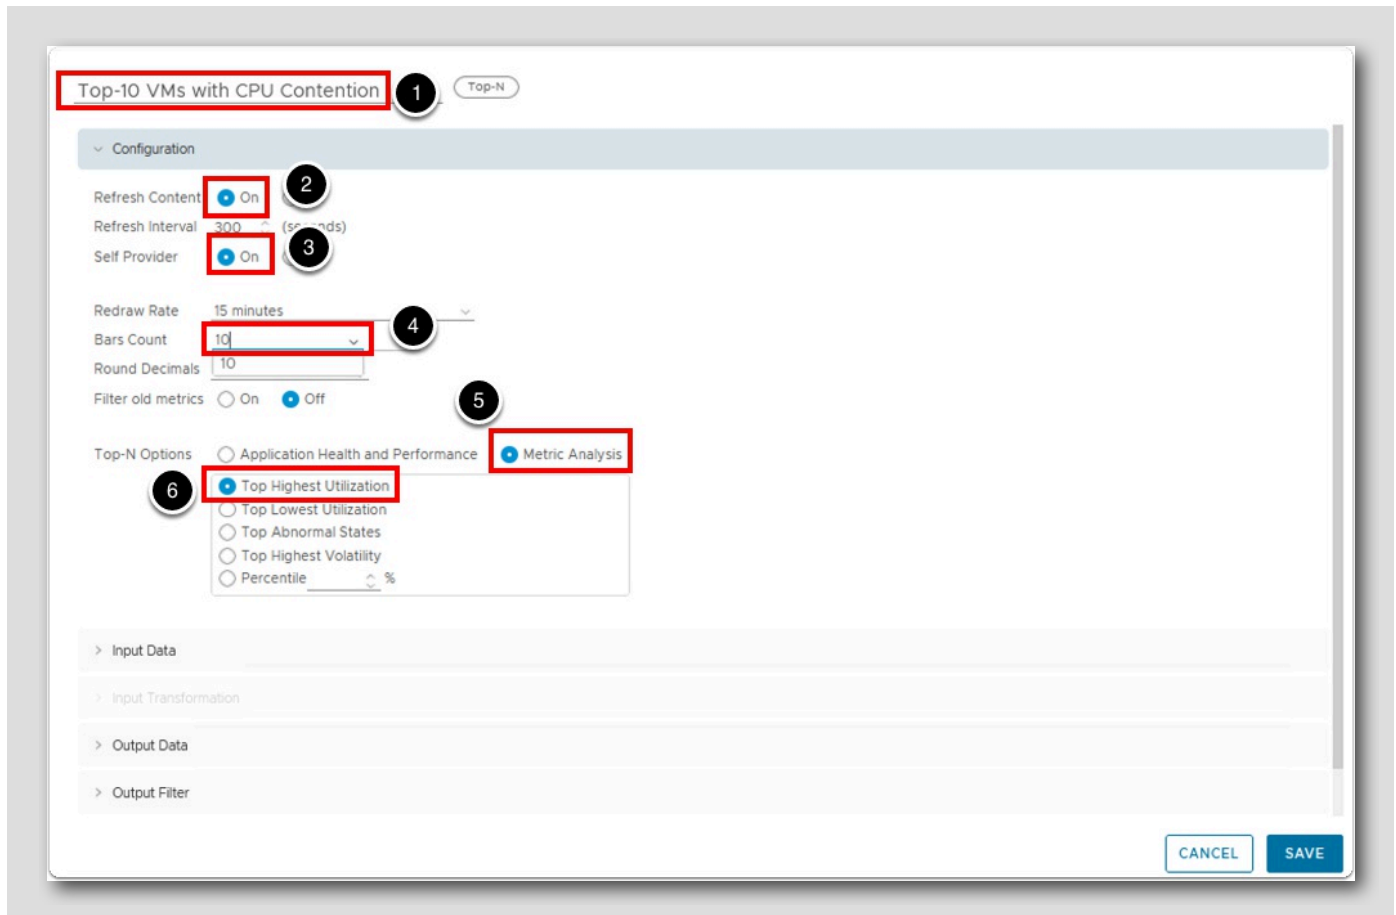
[307]



We now need to go into the settings of the widgets to make some configuration changes so that they will present the appropriate data in each of the widgets.

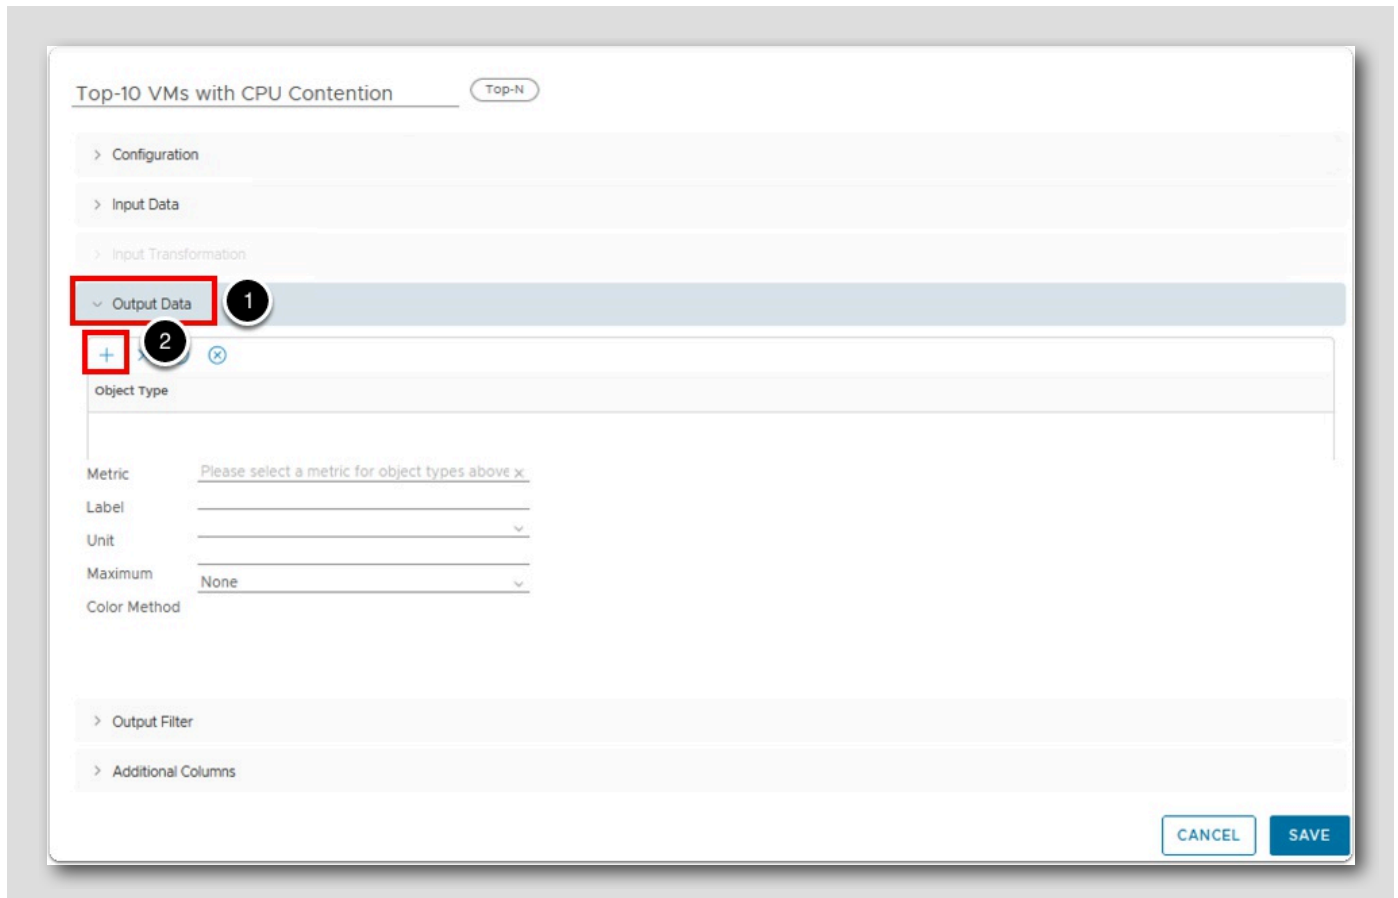
1. Hover the mouse over the top of the 1st Top-N widget in the widget to expose the hidden menu.
2. Click on the **Edit Widget** (*pencil*) icon to edit the widget.

## Configure Top-N Widget - Change Name



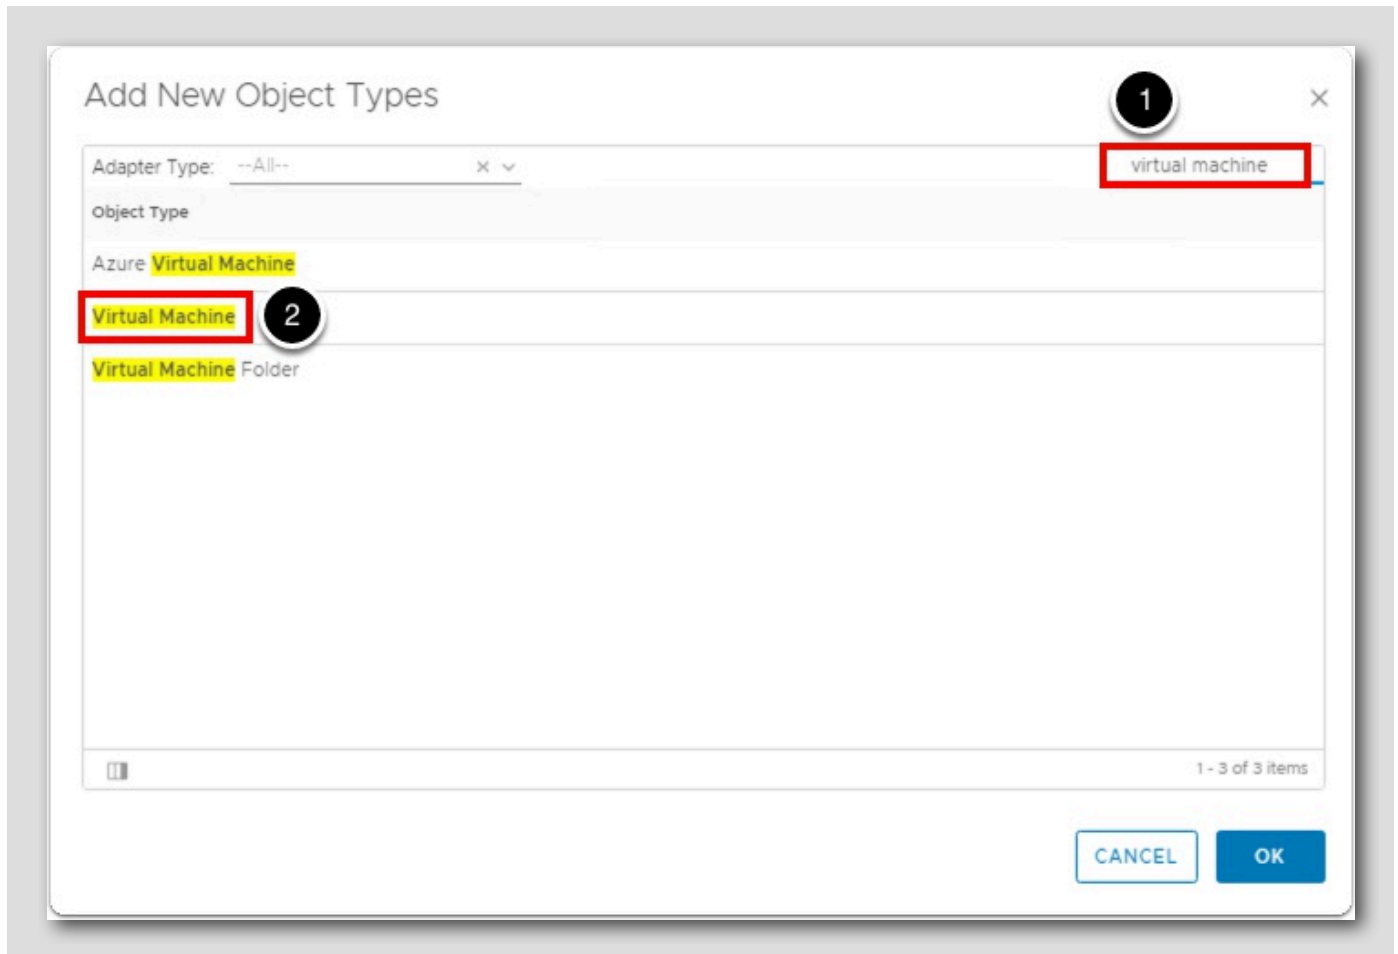
1. Replace the text Top-N with Top-10 VMs with CPU Contention.
2. Click on the On radio button next to Refresh Content.
3. Click on the On radio button next to Self Provider.
4. Change Bars Count to 10.
5. Click on the Metric Analysis radio button next to Top-N Options.
6. We see that it automatically selected the Top Highest Utilization radio button for us.

## Configure Top-N Widget - Object Types



1. Click on the **Output Data** selection to expand it.
2. Click on the **Add Object Type** (*plus sign*) to add an object type.

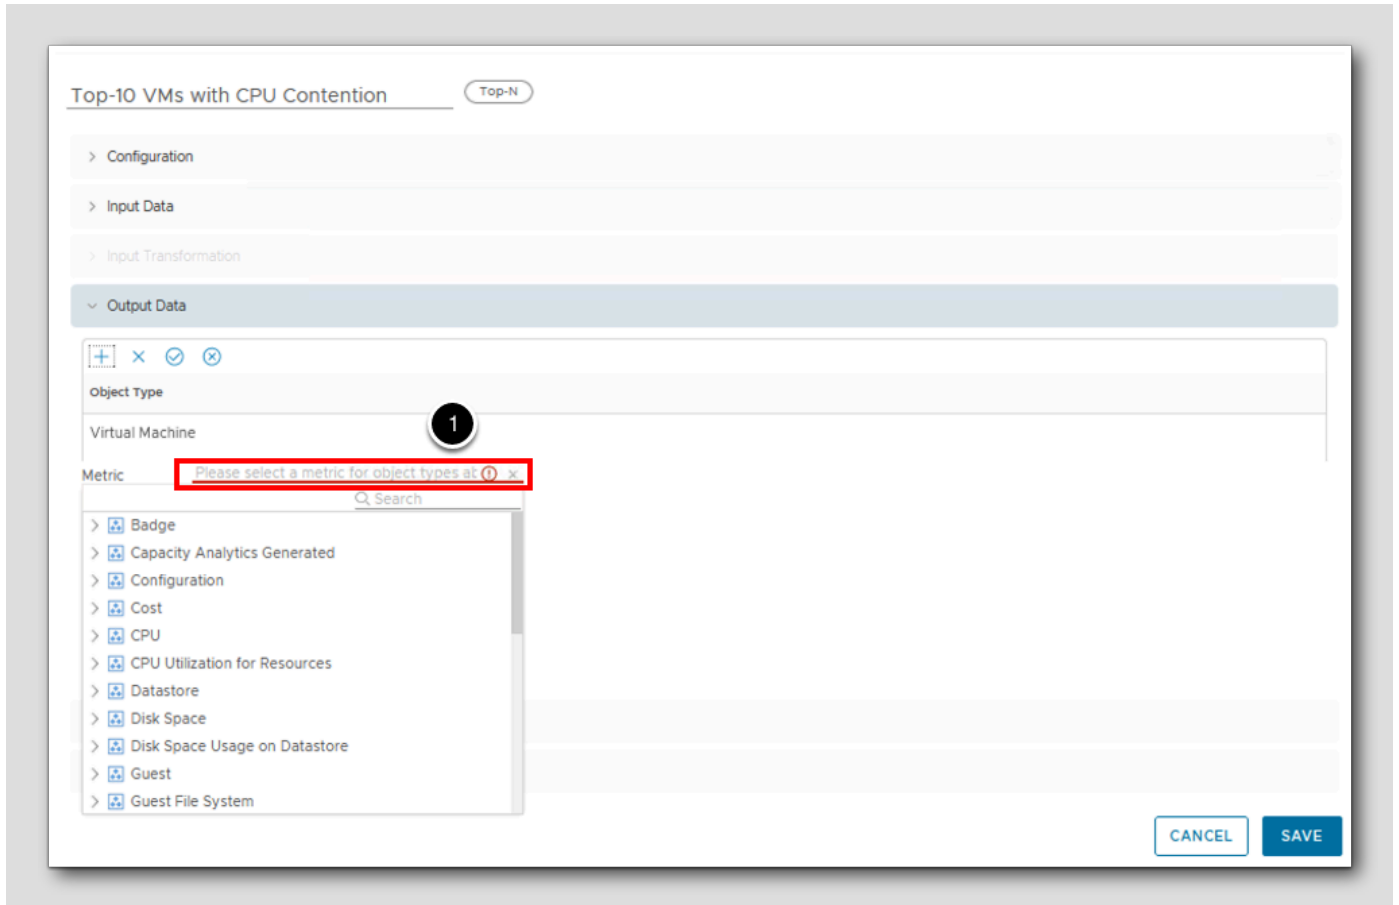
## Configure Top-N Widget - Virtual Machine



1. Type **virtual machine** into the Filter text field and hit the **ENTER** key on the keyboard.
2. Double click on **Virtual Machine**.

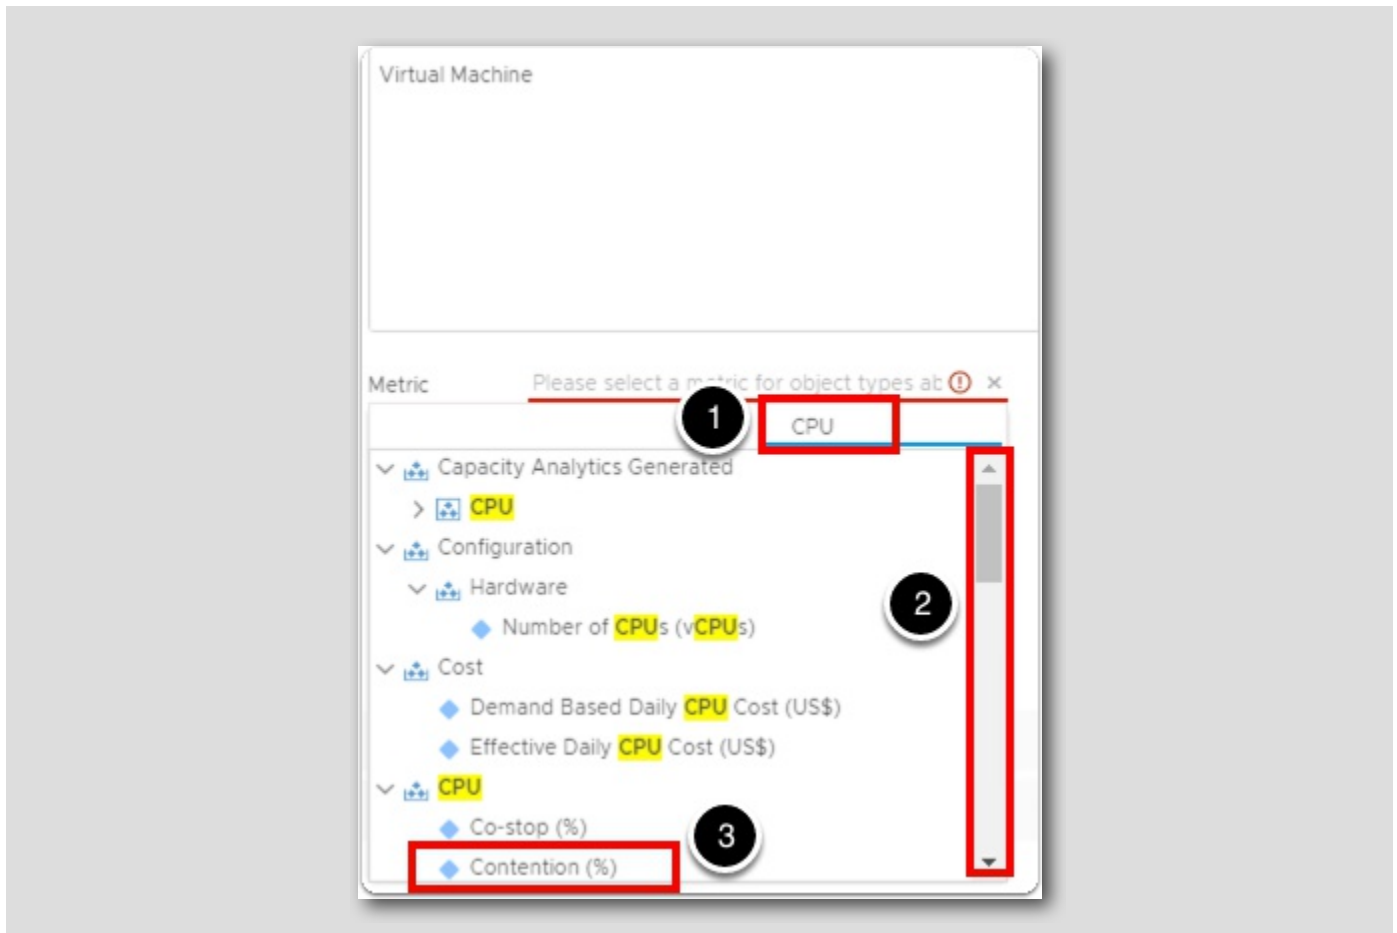


## Configure Top-N Widget - Widget



1. Click anywhere inside the **Metric** text field in order to expose the filter option.

## Configure Top-N Widget - CPU Contention (%)



1. Type CPU into the Metric text field and hit ENTER on the keyboard to filter for it.
2. Drag the scroll bar down (if needed) until we can see Contention (%) (CPU > Contention) in the drop-down list.
3. Double-Click on Contention (%).

## Output Data

Top-10 VMs with CPU Contention Top-N

> Configuration

> Input Data

> Input Transformation

▼ Output Data

Object Type

Virtual Machine

Metric CPUContention (%) ×

Label CPU Contention 1

Unit % 2

Maximum 100 3

Color Method Custom 4

75 5 85 6 95 7

Yellow Bound Orange Bound Red Bound

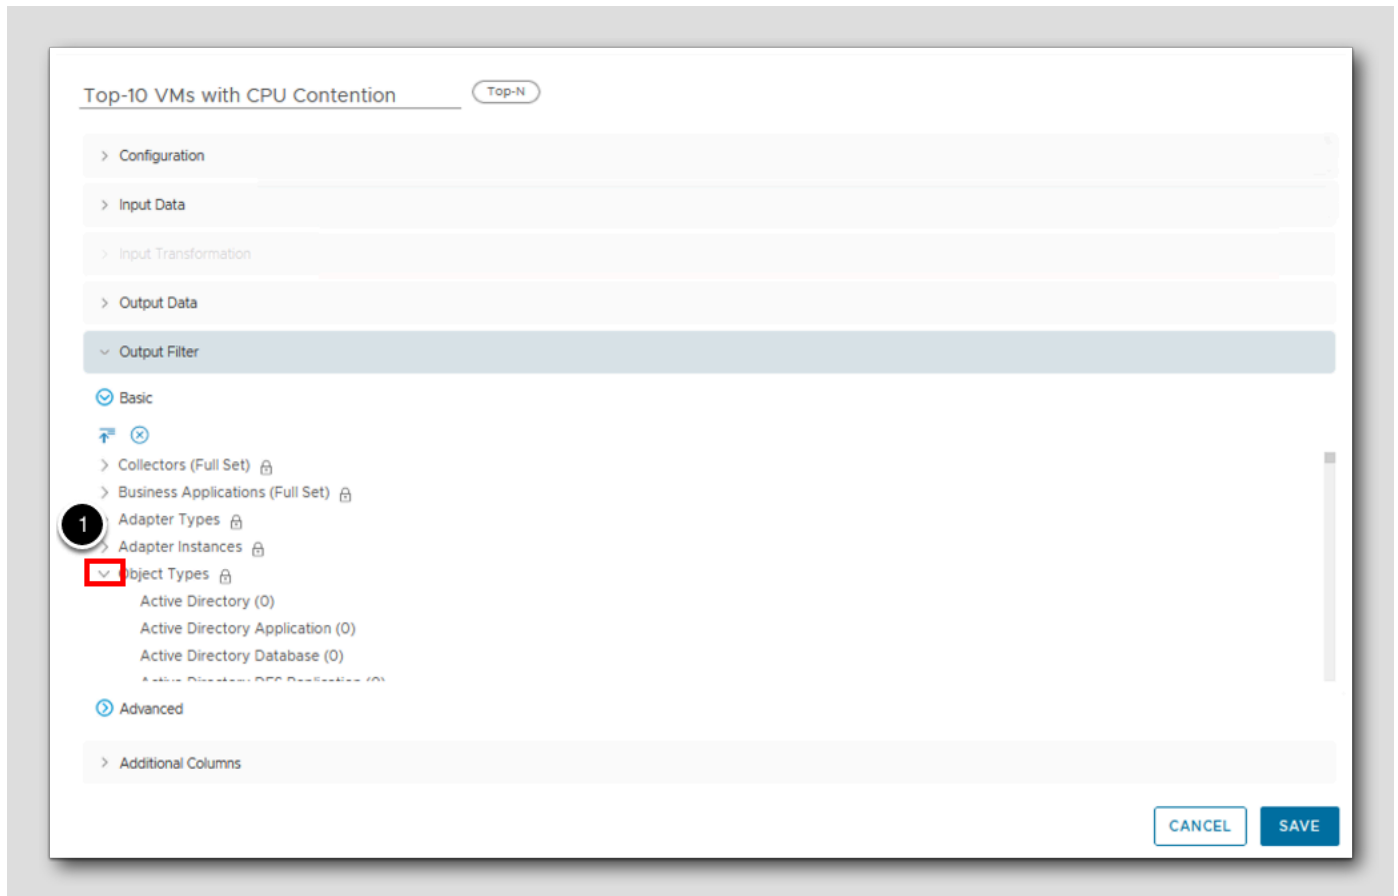
> Output Filter 8

> Additional Columns

CANCEL SAVE

1. Enter CPU Contention for the Label.
2. Change the unit to %.
3. Enter 100 for the Maximum value, this will set the graph bar to max out at a value of 100.
4. Change Color Method to Custom.
5. Enter 75 for Yellow Bound.
6. Enter 85 for Orange Bound.
7. Enter 95 for Red Bound.
8. Expand Output Filter.

## Output Filter



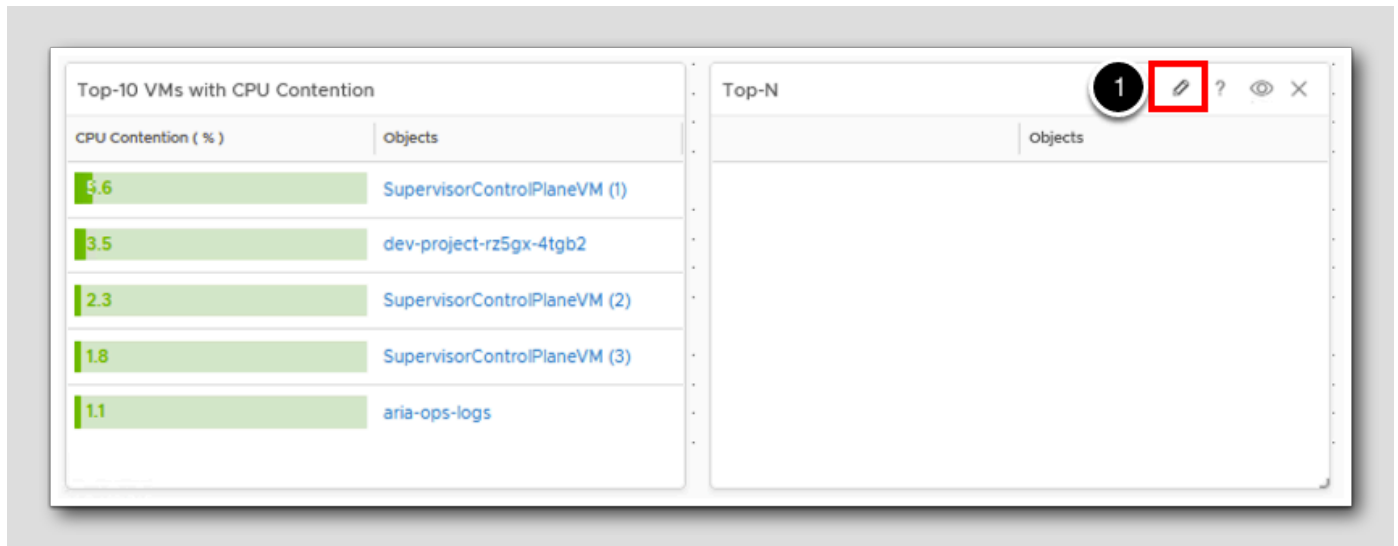
1. Expand Object Types.

## Output Filter (Continued)

The screenshot displays the 'Output Filter' configuration for 'Top-10 VMs with CPU Contention'. The interface is divided into 'Basic' and 'Advanced' sections. Under the 'Basic' section, a list of object types is shown, including 'Velo Cloud Gateway Processes (0)', 'Velo Cloud Orchestrator (0)', 'Velo Cloud Orchestrator Application (0)', 'Virtual Machine (23)', 'Virtual Machine Folder (7)', 'VMC Organization (0)', 'VMC Region (22)', and 'VMC Region Per Account (0)'. The 'Virtual Machine (23)' item is selected and highlighted with a red box. A red vertical bar on the right side of the list is labeled with a circled '1'. The 'Virtual Machine (23)' item is labeled with a circled '2'. At the bottom right, the 'SAVE' button is highlighted with a red box and labeled with a circled '3'. The 'CANCEL' button is also visible.

1. Scroll down until you see **Virtual Machine** in the Object Type list.
2. Single click on **Virtual Machine** Object Type.
3. Click **SAVE**.

## Configure Top-N Widget - Top-N



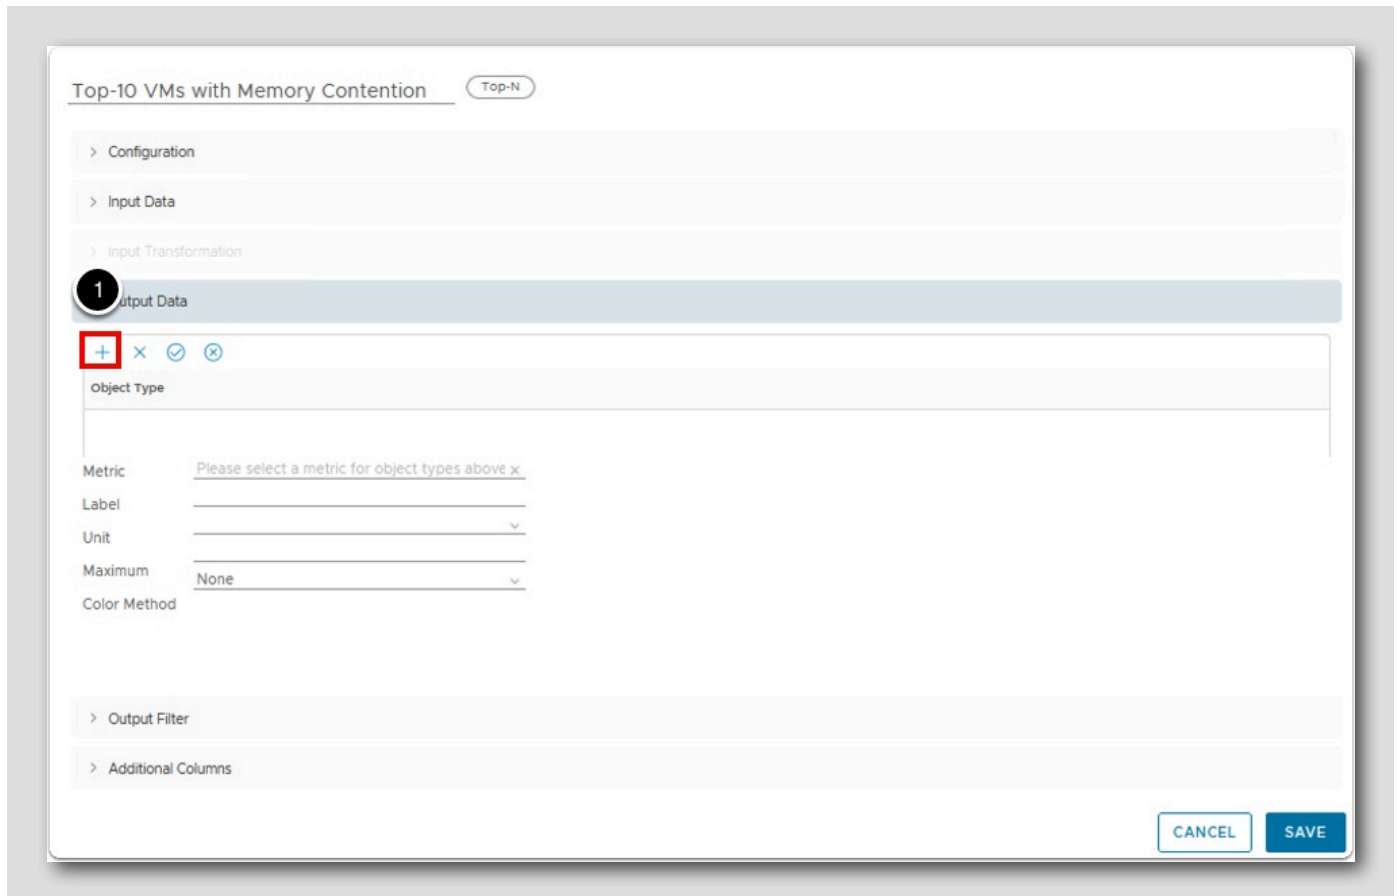
1. Hover over the second Top-N widget and click on the Edit Widget (*pencil*) icon when it appears.

## Configure Top-N Widget - Change Name

The screenshot shows the configuration for a Top-N widget. The title is 'Top-10 VMs with Memory Contention'. The configuration section includes: Refresh Content (On), Refresh Interval (300 seconds), Self Provider (On), Redraw Rate (15 minutes), Bars Count (10), Round Decimals (0), and Filter old metrics (Off). The Top-N Options section shows 'Metric Analysis' selected, with 'Top Highest Utilization' automatically selected. The Output Data section is expanded.

1. Replace the text **Top-N** with **Top-10 VMs with Memory Contention**.
2. Click on the **On** radio button next to **Refresh Content**.
3. Click on the **On** radio button next to **Self Provider**.
4. Change **Bars Count** to **10**.
5. Click on the **Metric Analysis** radio button next to **Top-N Options**.
6. We see that it automatically selected the **Top Highest Utilization** radio button for us.
7. Click on the **Output Data** selection to expand it.

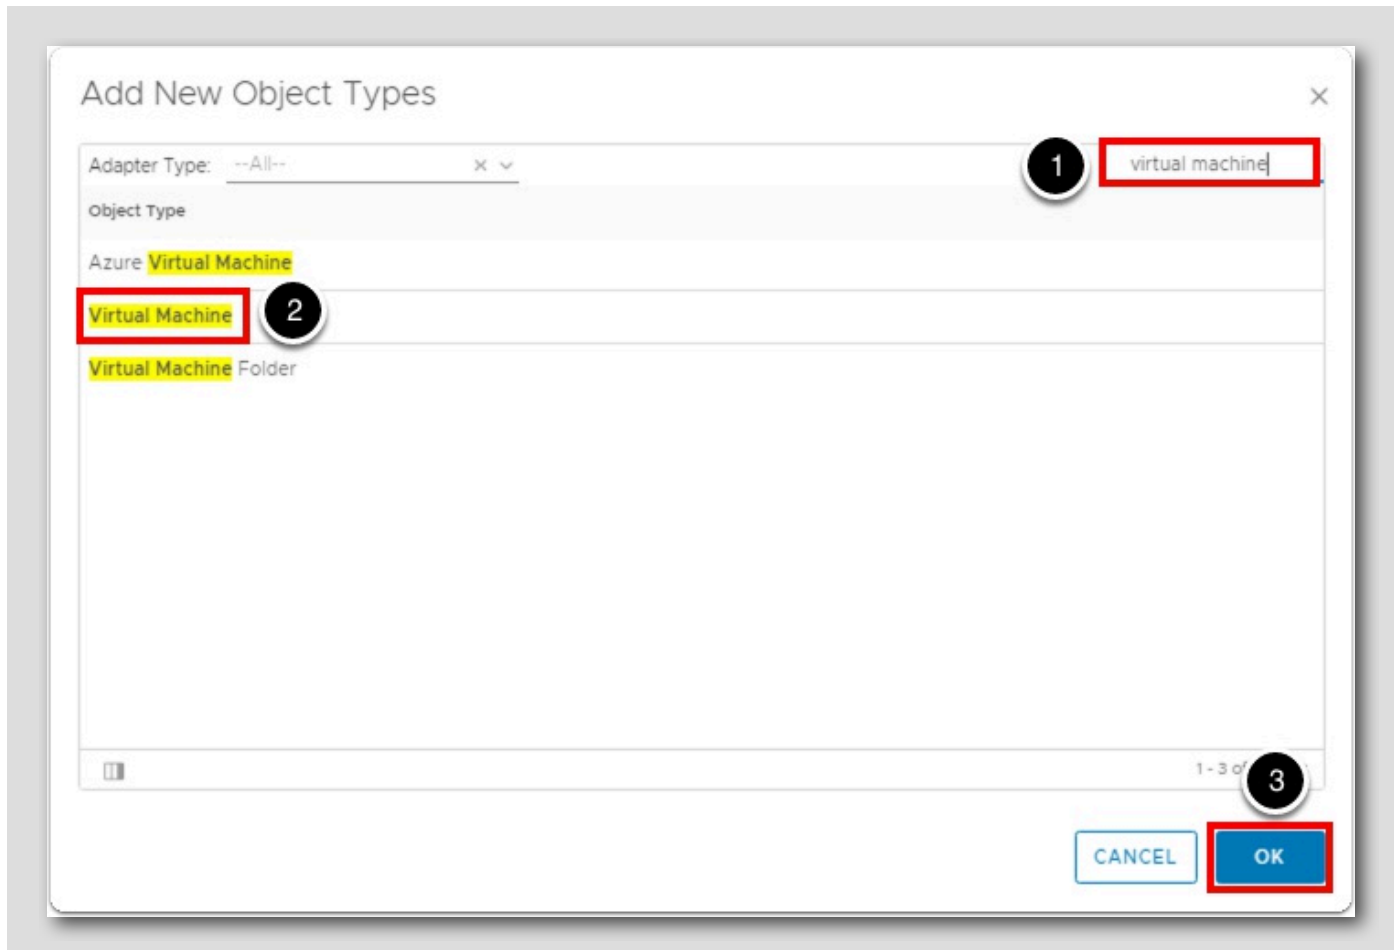
## Configure Top-N Widget - Object Types



1. Click on the Add Object Type (*plus sign*) to add an object type.



## Configure Top-N Widget - Virtual Machine



1. Type **virtual machine** into the Filter text field and hit the ENTER key on the keyboard.
2. Click on **Virtual Machine** in the list to select it.
3. Then click on the OK button.

## Configure Top-N Widget - Metric

Top-10 VMs with Memory Contention Top-N

> Configuration

> Input Data

> Input Transformation

▼ Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric 1 Please select a metric for object types above x

Label \_\_\_\_\_

Unit \_\_\_\_\_

Maximum None

Color Method \_\_\_\_\_

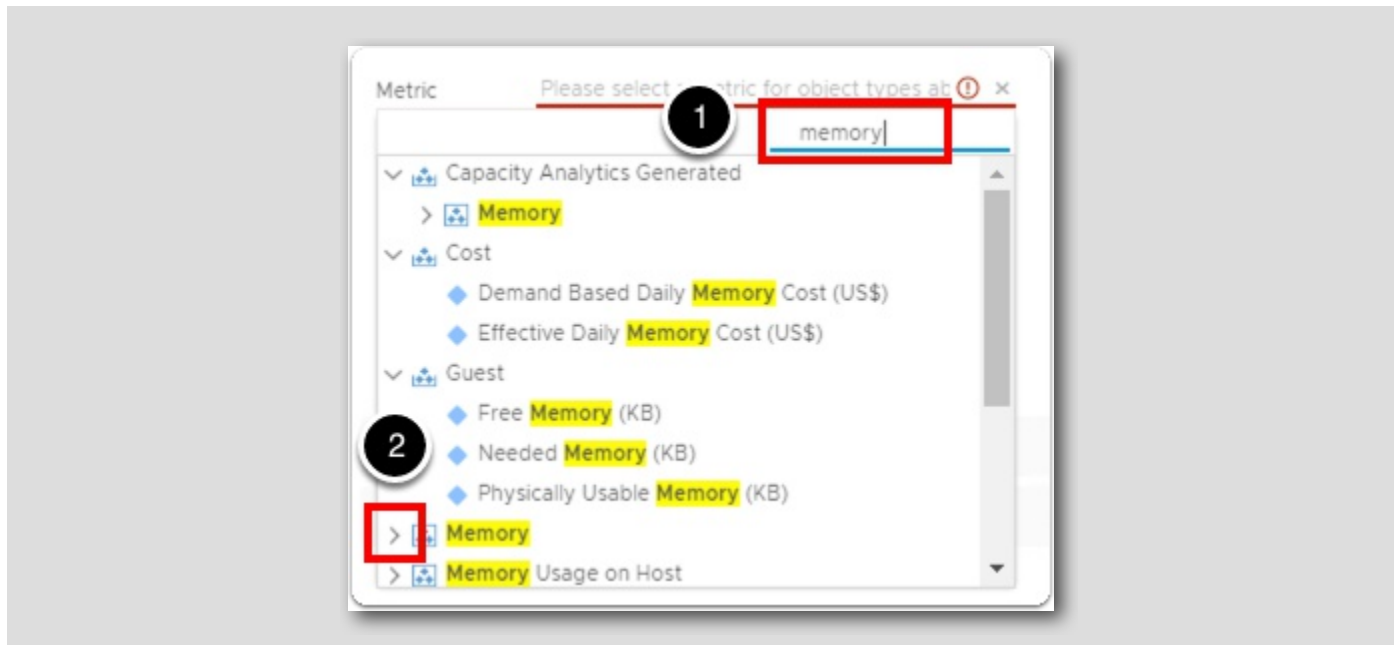
> Output Filter

> Additional Columns

CANCEL SAVE

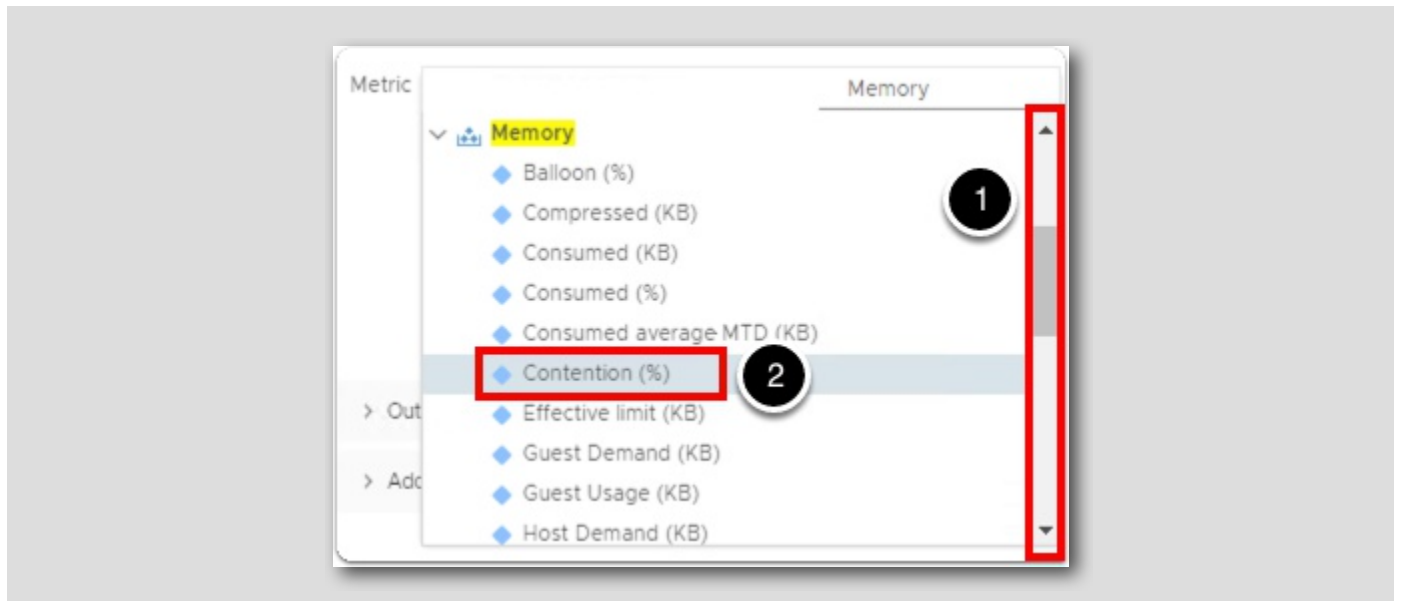
1. Click anywhere inside the **Metric** text field in order to expose the filter option.

## Configure Top-N Widget - Memory



1. Type **Memory** into the filter text field and hit **ENTER** on the keyboard to filter for it.
2. Then click on the arrow next to **Memory** to expand its drop-down menu.

## Configure Top-N Widget - Memory Contention (%)



1. Drag the scroll bar down until we can see **Contention (%)** in the drop-down list.
2. Double-Click on **Contention (%)**.

## Configure Top-N Widget - Memory Contention (%) (Continued)

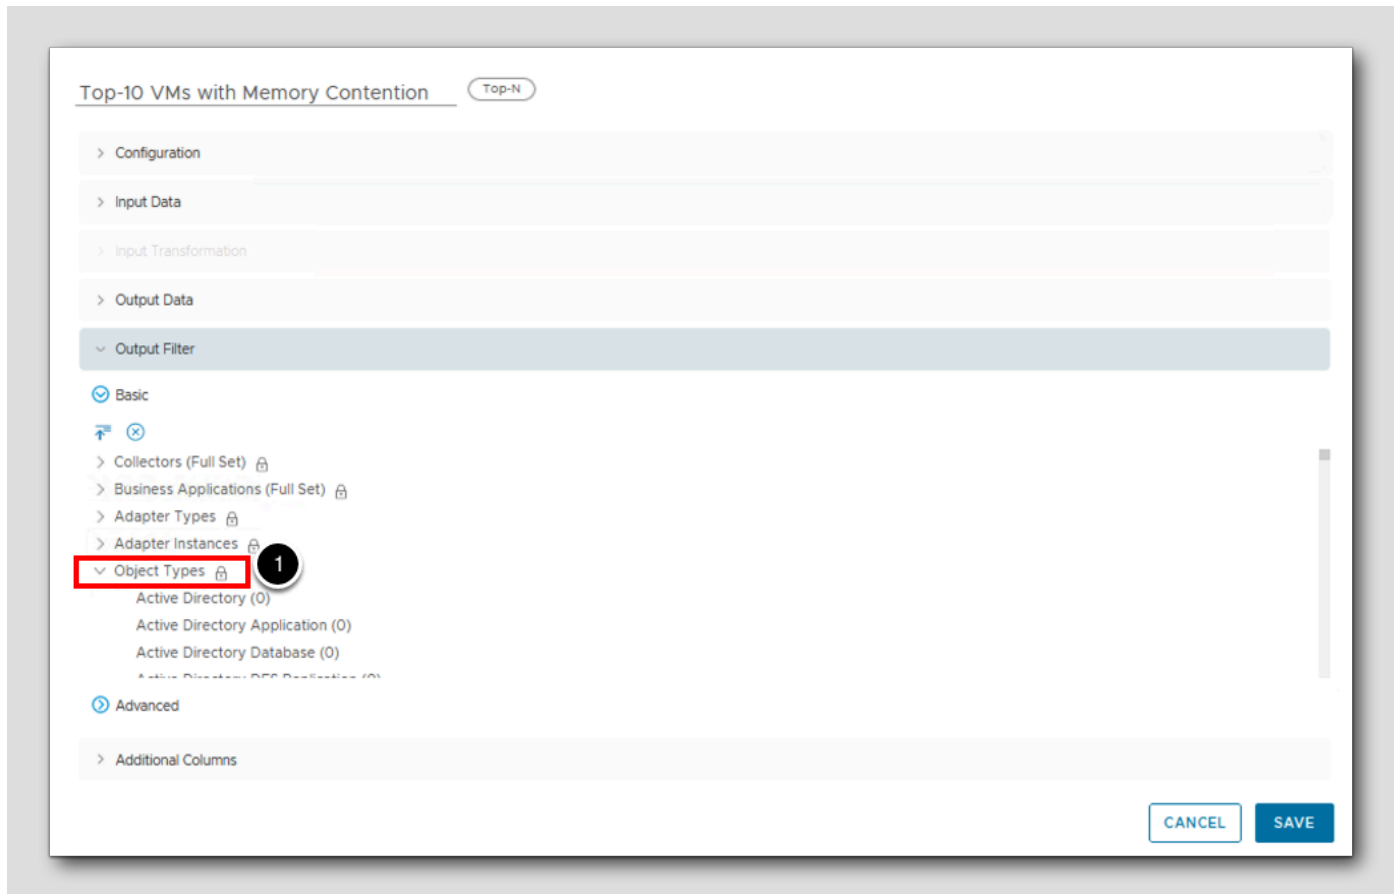
The screenshot shows the configuration interface for a Top-N widget. The widget is titled "Top-10 VMs with Memory Contention". The configuration is as follows:

- Metric: Memory Contention (%)
- Label: Memory Contention
- Unit: %
- Maximum: 100
- Color Method: Custom
- Yellow Bound: 75
- Orange Bound: 85
- Red Bound: 95

The "Output Filter" section is also expanded. The configuration is saved.

1. Enter **Memory Contention** for the Label.
2. Change the unit to **%**.
3. Enter **100** for the **Maximum** value, this will set the graph bar to max out at a value of 100.
4. Change **Color Method** to **Custom**.
5. Enter **75** for **Yellow Bound**.
6. Enter **85** for **Orange Bound**.
7. Enter **95** for **Red Bound**.
8. Expand **Output Filter**.

## Output Filter



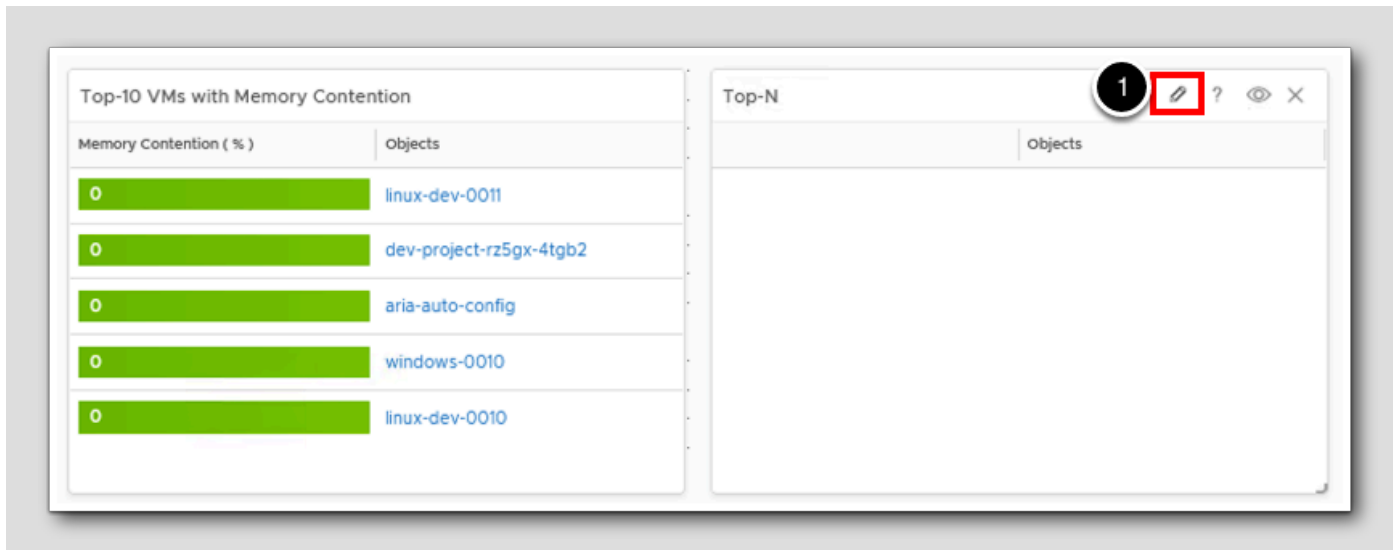
1. Expand Object Types.

## Output Filter (Continued)

The screenshot shows the 'Output Filter' configuration page for 'Top-10 VMs with Memory Contention'. The page is divided into sections: Configuration, Input Data, Input Transformation, Output Data, and Output Filter. The 'Output Filter' section is expanded, showing a list of object types under the 'Basic' tab. The 'Virtual Machine (23)' option is highlighted with a red box and a circled '2'. A vertical scrollbar on the right is also highlighted with a red box and a circled '1'. At the bottom right, the 'SAVE' button is highlighted with a red box and a circled '3'. The 'CANCEL' button is also visible.

1. Scroll down until you see Virtual Machine in the Object Type list.
2. Single click on Virtual Machine Object Type.
3. Click SAVE

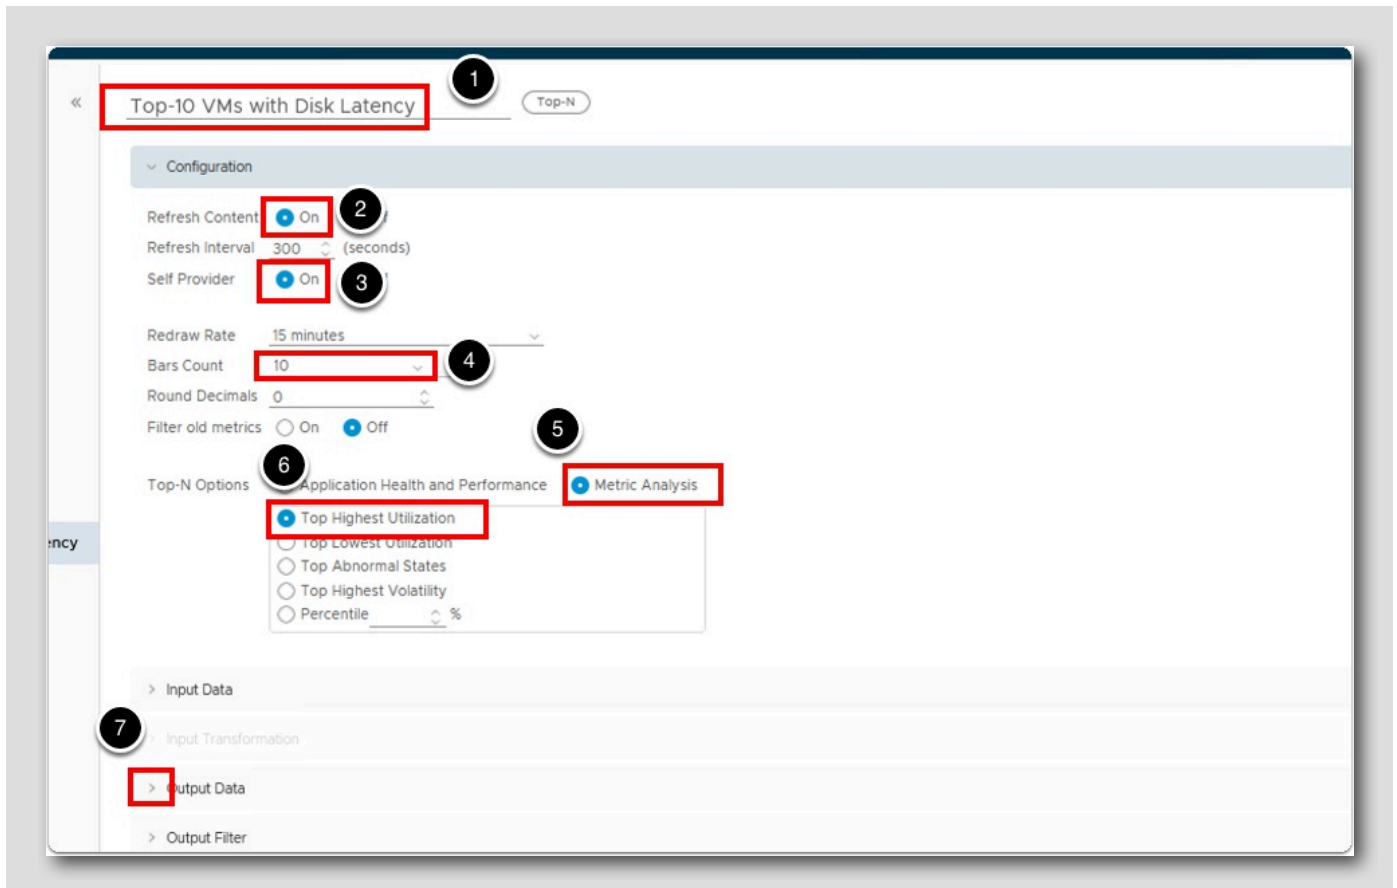
## Configure Top-N Widget - Top-N



1. Hover over the right most Top-N widget and click on the Edit Widget (*pencil*) icon when it appears.



## Configure Top-N Widget - Change Name



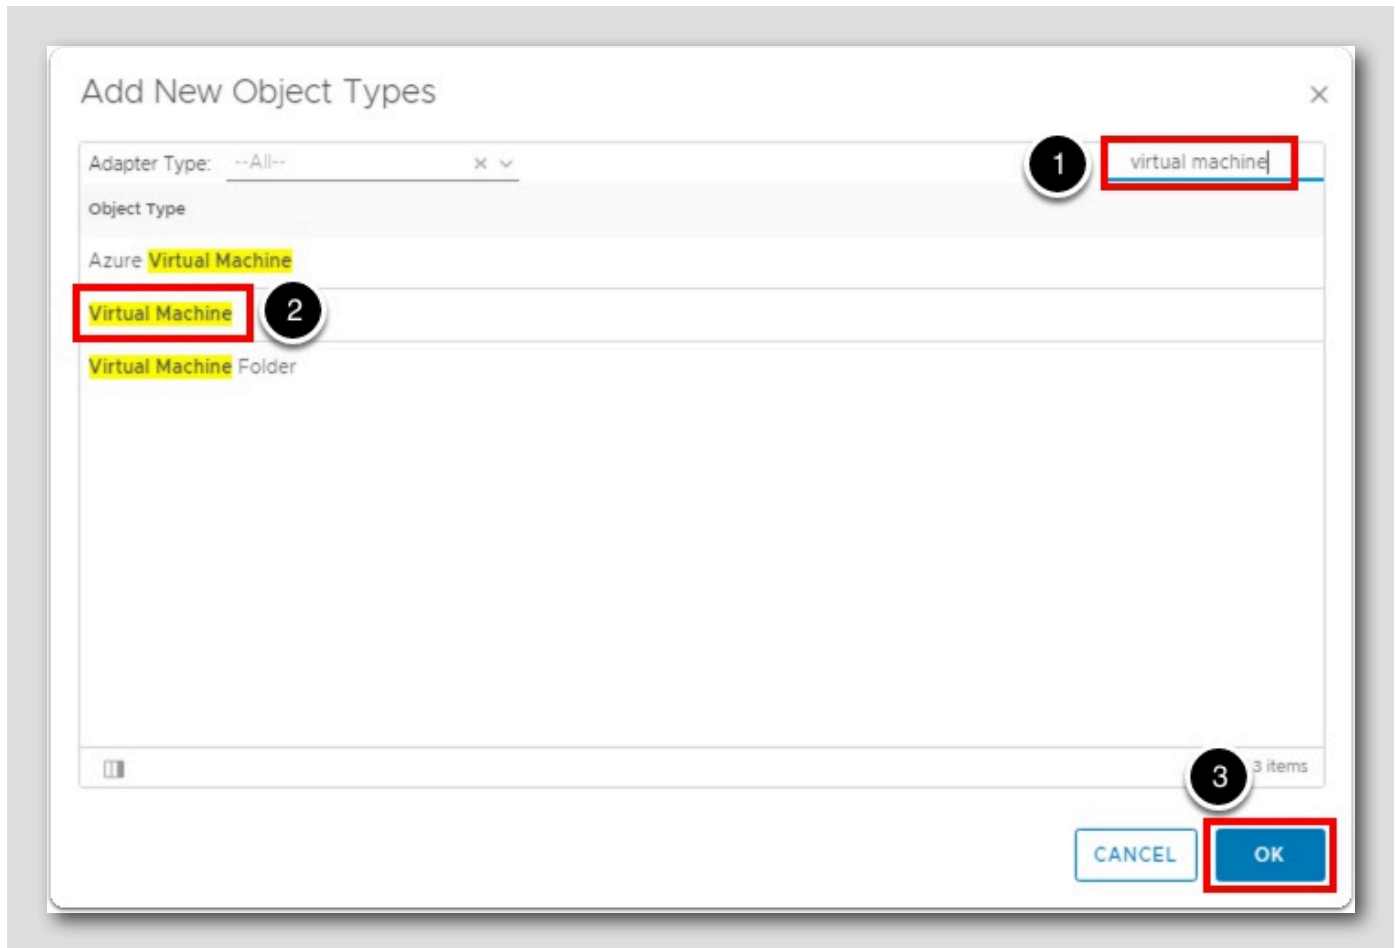
1. Replace the text **Top-N** with **Top-10 VMs with Disk Latency**.
2. Click on the **On** radio button next to **Refresh Content**.
3. Click on the **On** radio button next to **Self Provider**.
4. Change **Bars Count** to **10**.
5. Click on the **Metric Analysis** radio button next to **Top-N Options**.
6. We see that it automatically selected the **Top Highest Utilization** radio button for us.
7. Click on the **Output Data** selection to expand it.

## Configure Top-N Widget - Object Types

The screenshot shows the configuration interface for a Top-N widget. The widget title is "Top-10 VMs with Disk Latency" and it has a "Top-N" button. The configuration is organized into sections: Configuration, Input Data, Input Transformation, Output Data, Output Filter, and Additional Columns. The "Output Data" section is highlighted with a blue bar and a red circle containing the number "1". Below this section, there is a list of object types with a red square highlighting a plus sign icon. Below the list, there are fields for Metric, Label, Unit, Maximum, and Color Method.

1. Click on the Add Object Type (*plus sign*) to add an object type.

## Configure Top-N Widget - Virtual Machine



1. Type **virtual machine** into the Filter text field and hit the **ENTER** key on the keyboard.
2. Click on **Virtual Machine** in the list to select it.
3. Then click on the **OK** button.

## Configure Top-N Widget - Metric

Top-10 VMs with Disk Latency Top-N

> Configuration

> Input Data

> Input Transformation

< Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric Please select a metric for object types above x

Label

Unit

Maximum None

Color Method

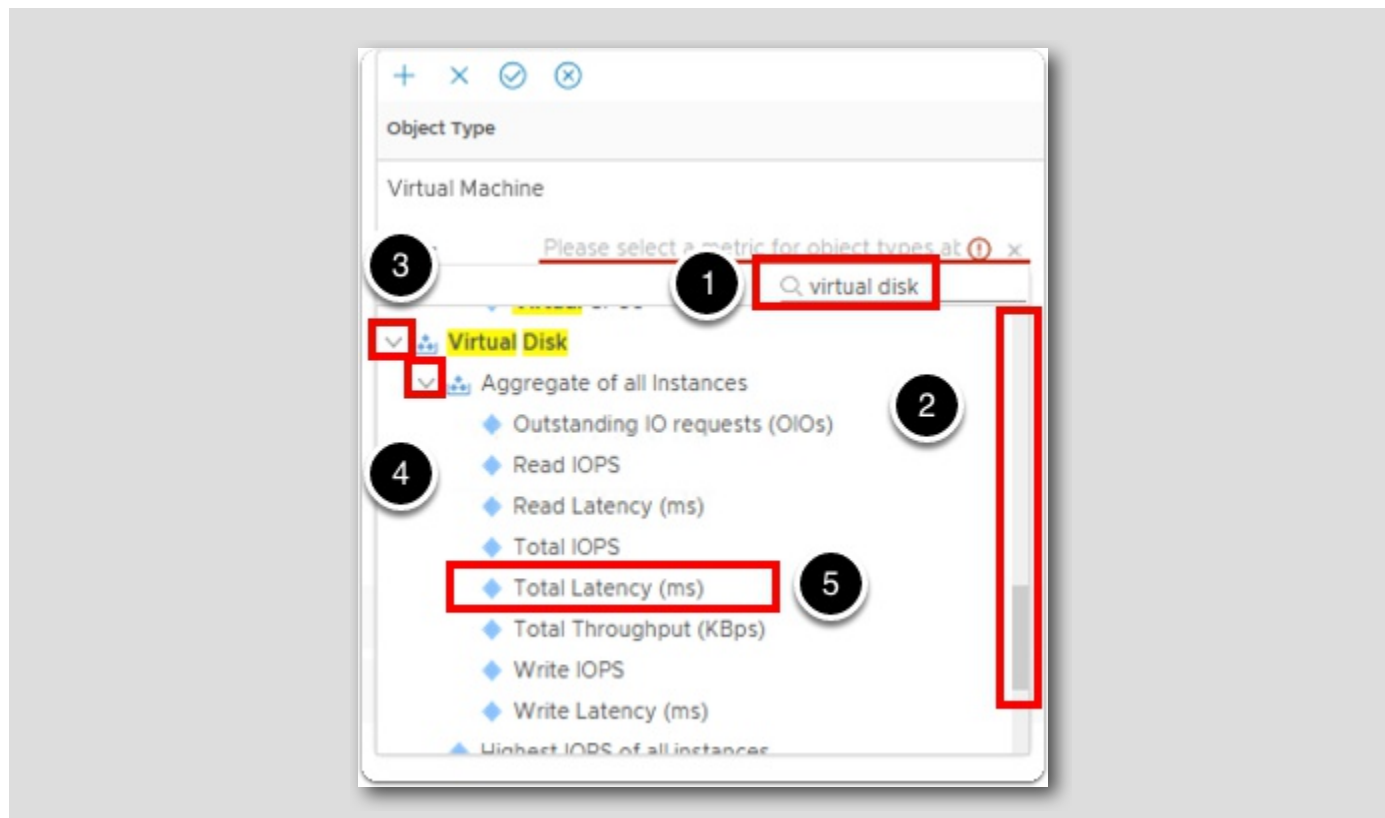
> Output Filter

> Additional Columns

CANCEL SAVE

1. Click anywhere inside the **Metric** text field in order to expose the filter option.

## Configure Top-N Widget - Virtual Disk Total Latency (ms)



1. Type **virtual disk**: into the Metric text field and hit **ENTER** to filter for it.
2. **Scroll Down** to the bottom of the list to see **Virtual Disk**.
3. Click on the Chevron beside **Virtual Disk** to expand it.
4. Click on the Chevron beside **Aggregate of all Instances** to expand it (You may need to scroll down more in the list to see it).
5. **Double-Click** on **Total Latency (ms)**.

## Configure Top-N Widget - Virtual Disk Total Latency (ms) (Continued)

Top-10 VMs with Disk Latency Top-N

> Configuration

> Input Data

> Input Transformation

< Output Data

+ × ✓ ✕

Object Type

Virtual Machine

Metric Virtual Disk:Aggregate of all instances|Total I x

Label Total Latency 1

Unit ms 2

Maximum None ms

Color Method

3

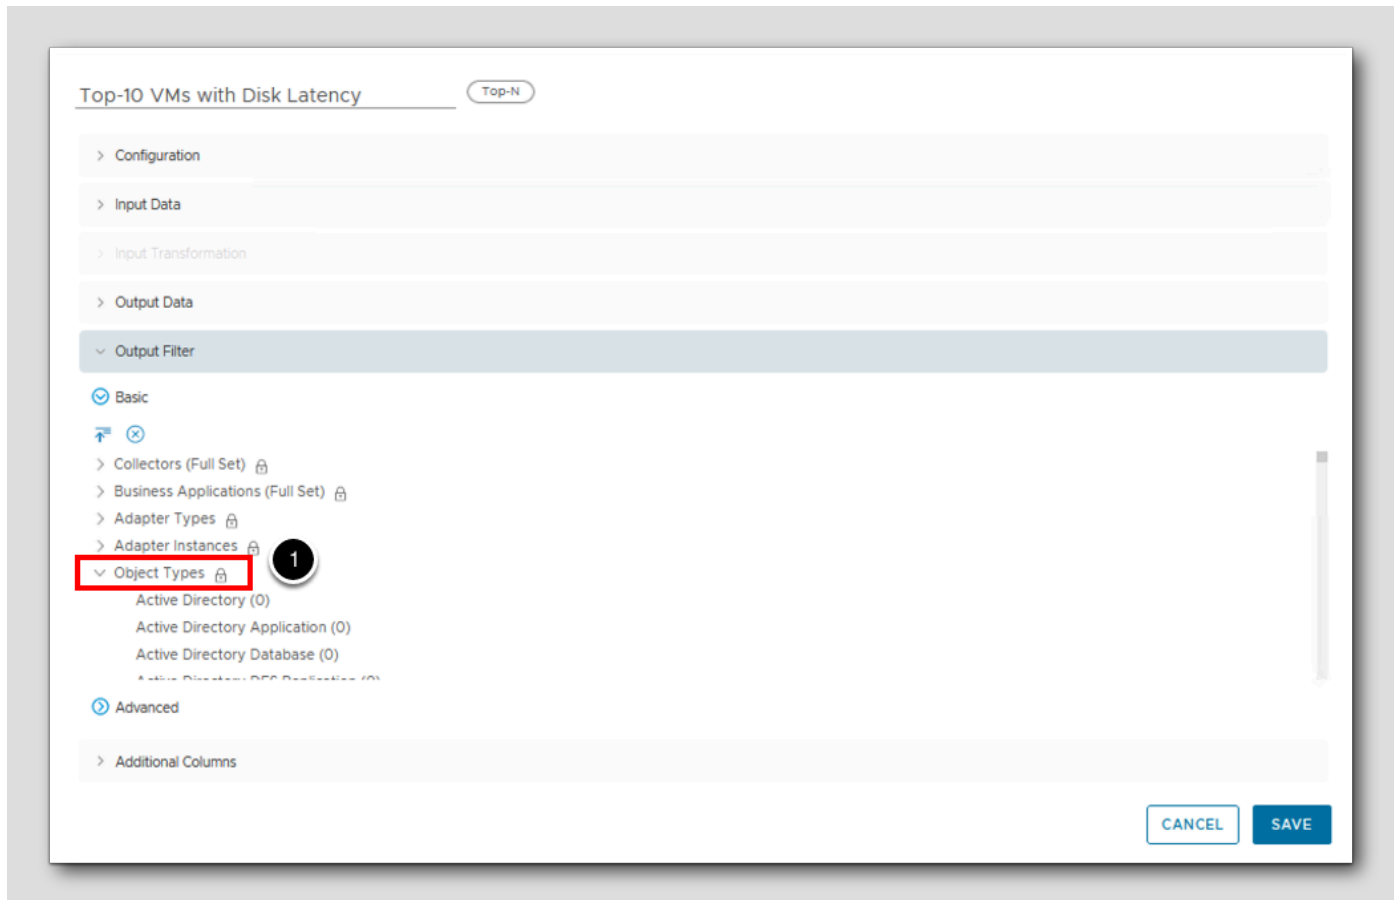
> Output Filter

> Additional Columns

CANCEL SAVE

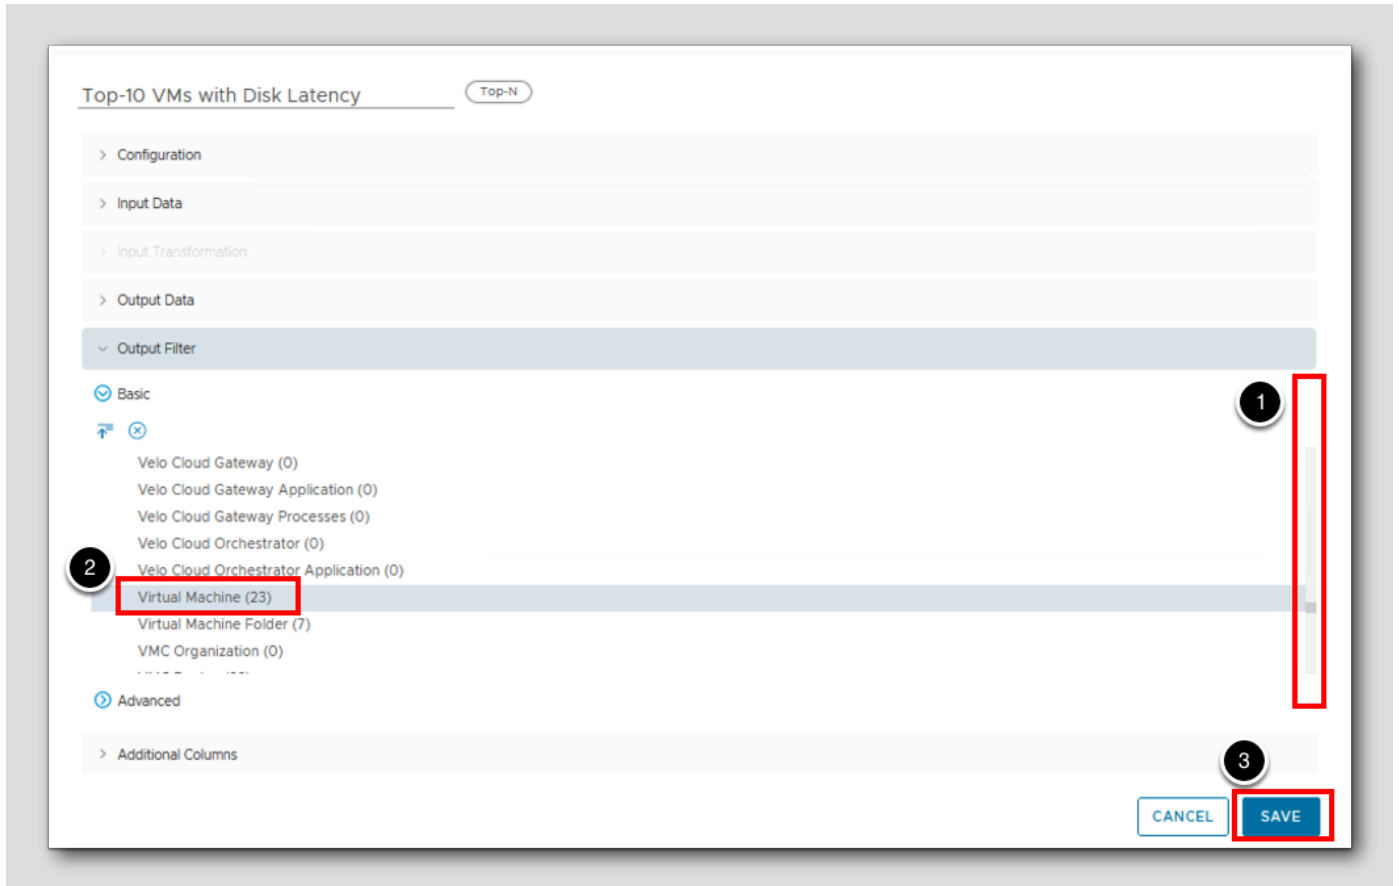
1. Type **Total Latency** for the Label.
2. Change the Unit to **ms**.
3. Expand **Output Filter**.

## Output Filter



1. Expand Object Types.

## Output Filter (Continued)

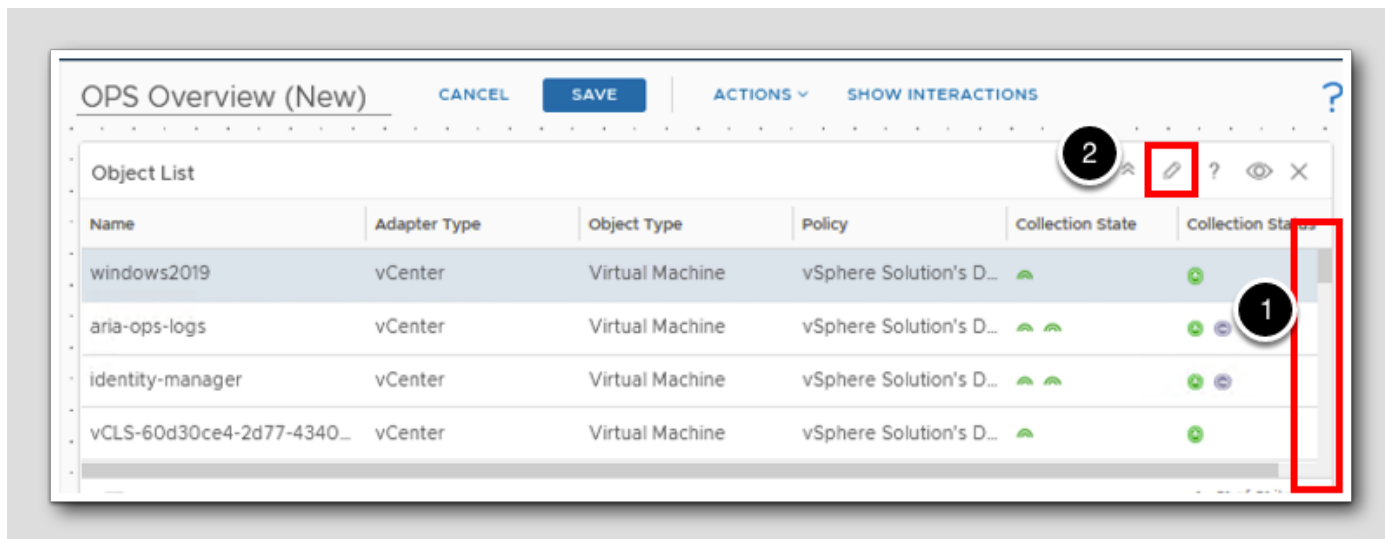


1. Scroll down until you see Virtual Machine in the Object Type list.
2. Single click on Virtual Machine Object Type.
3. Click SAVE



## Configure Object List

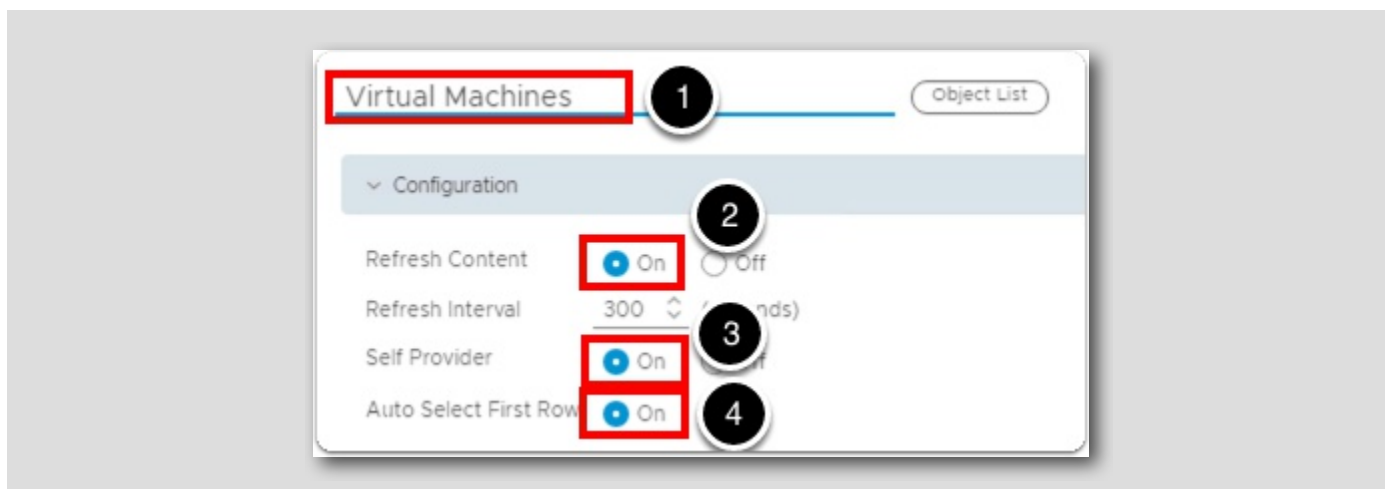
[335]



1. Scroll back up to the top of the dashboard.
2. Hover over the Object List widget and click on the Edit Widget (*pencil*) icon when it appears.

## Configure Object List - Change Name

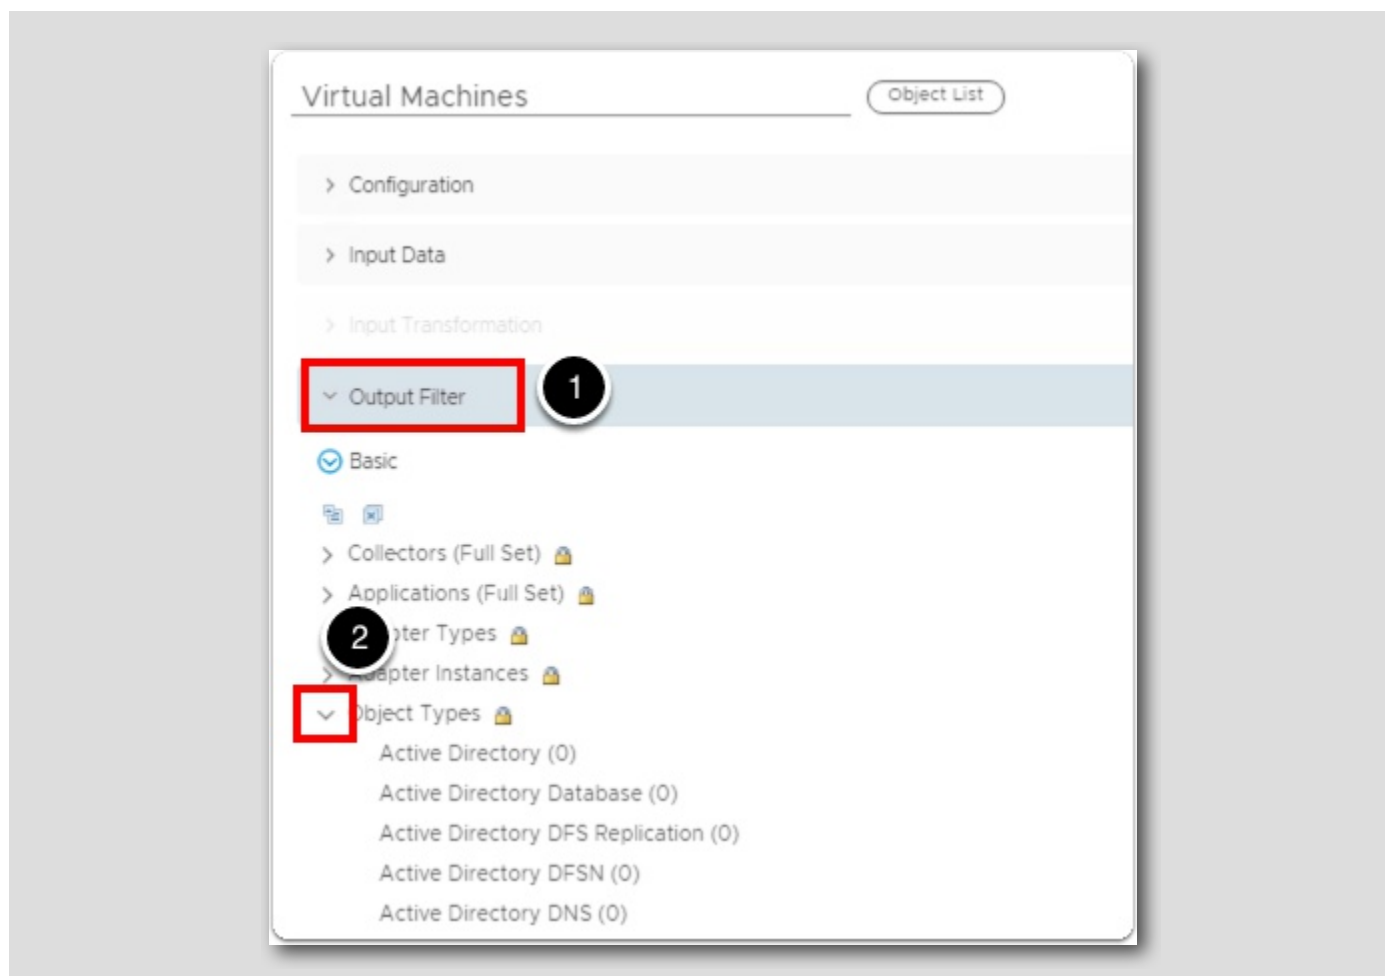
[336]



1. Type **Virtual Machines** into the Name text field.
2. Click on the **On** radio button to the right of Refresh Content.
3. Click on the **On** radio button to the right of Self Provider.
4. Click on the **On** radio button to the right of Auto Select First Row.

## Configure Object List - Output Filter

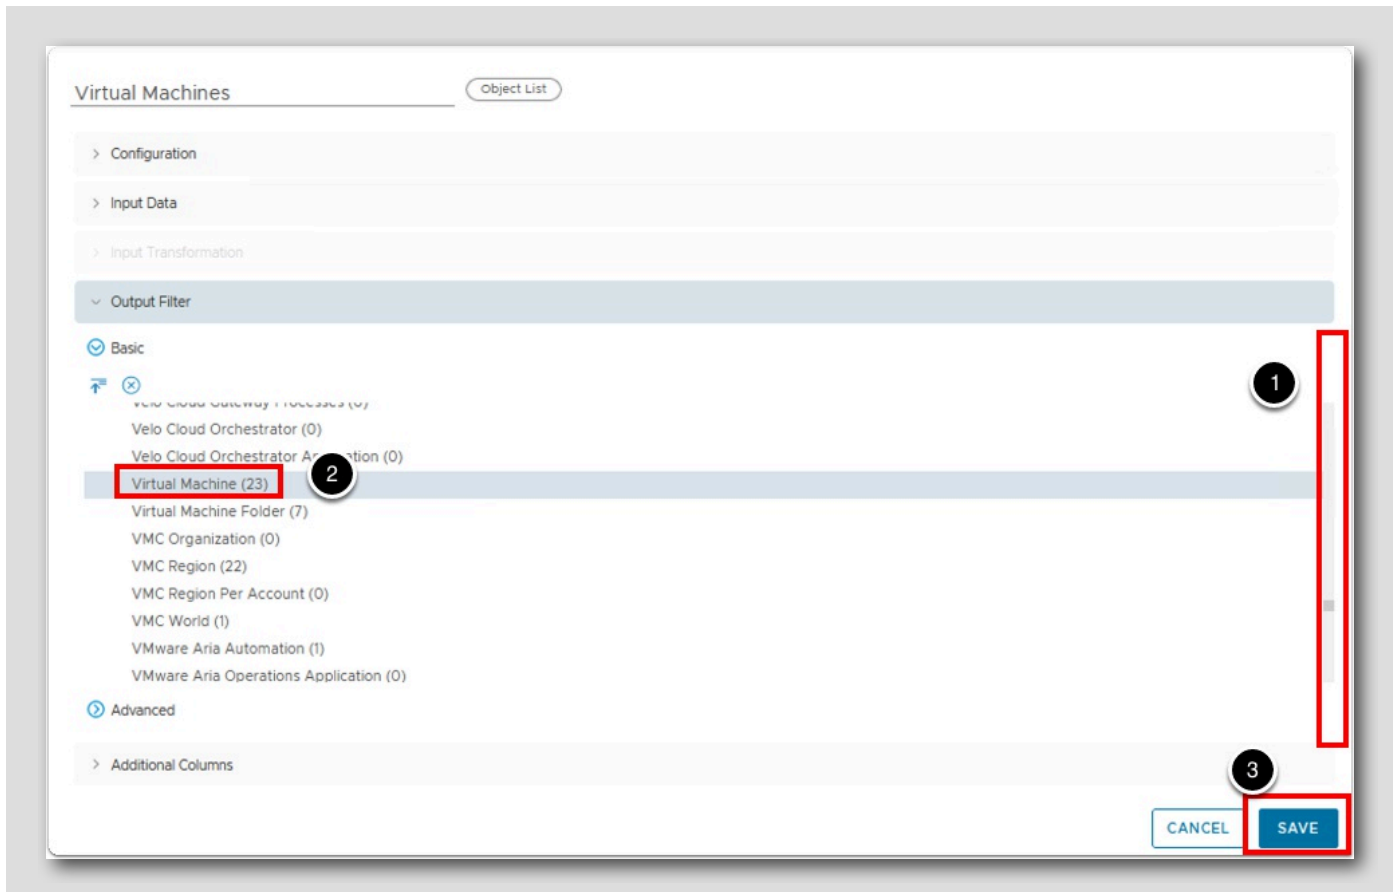
[337]



1. Click on **Output Filter**.
2. Then click on the **arrow** next to **Object Types** to expand its menu.

## Configure Object List - Virtual Machine


[338]



1. Drag the scroll bar down until we can see **Virtual Machine** in the drop-down list.
2. Then click on **Virtual Machine**.
3. Click **SAVE**.

## Save Dashboard

[339]



1. Click on the **SAVE** button to save our Dashboard.

## Review Dashboard

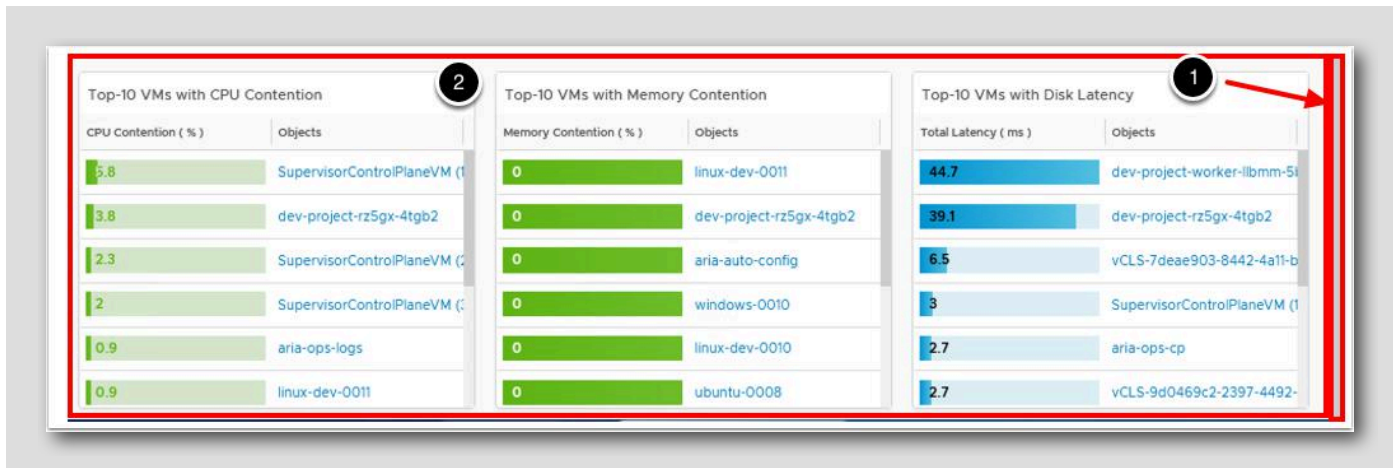
The screenshot displays the VMware Ops Overview (New) dashboard. At the top, there's a header with 'Ops Overview (New)' and 'ACTIONS'. Below this is a table of Virtual Machines. The table has columns for Name, Adapter Type, Object Type, Policy, Collection State, and Collection Status. The VM 'ubuntu-0008' is highlighted. Below the table are three main widgets: Object Relationship (showing a hierarchy with 'ubuntu-0008' at the center), Top Alerts (showing 'No Issues'), and Health (showing 'Health' with 'Immediate issues'). At the bottom, there are three more widgets: Top-10 VMs with CPU Contention, Top-10 VMs with Memory Contention, and Top-10 VMs with Disk Latency. Red boxes and numbered callouts (1-4) indicate the steps to be followed.

Name	Adapter Type	Object Type	Policy	Collection State	Collection Status
dev-project-rz5gx-4tgb2	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢
aria-auto-config	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢
windows-0010	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢
linux-dev-0010	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢
ubuntu-0008	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢
aria-auto	vCenter	Virtual Machine	vSphere Solution's D...	🟢	🟢

1. Drag the scroll bar all the way to the top of the new dashboard.
2. Drag the scroll bar in the Virtual Machines widget until you see ubuntu-0008.
3. Click on the the VM ubuntu-0008 to view the stats for that object (You may need to scroll down a little in the list of VMs to see it).
4. We now see that the three widgets (Object Relationship, Top Alerts and Health) now are populated with metrics related to the ubuntu-0008 virtual machine.

## Review Dashboard (continued)

[341]



1. Drag the scroll bar all the way to the bottom of the dashboard.
2. We see that we have the (3) Top-10 VMs widgets based on CPU Contention, Memory Contention and Disk Latency.

## Lesson End

[342]

Congratulations, we have completed the lesson on **Creating a New Custom Dashboard!**

In this lesson, we created a brand new custom dashboard that contained an Object List of virtual machines that had relationships to all the other widgets. However, we did not create the relationship from the virtual machine in the Object List widget to the Top-N widgets. This ensures that no matter which virtual machine we selected from the Object List widget, the Top-N widgets will always show the Top-10 VMs with CPU contention, Memory contention and Disk latency.

## Importing New Dashboards

[343]

In this lesson, we will learn how to import new dashboards into Aria Operations.

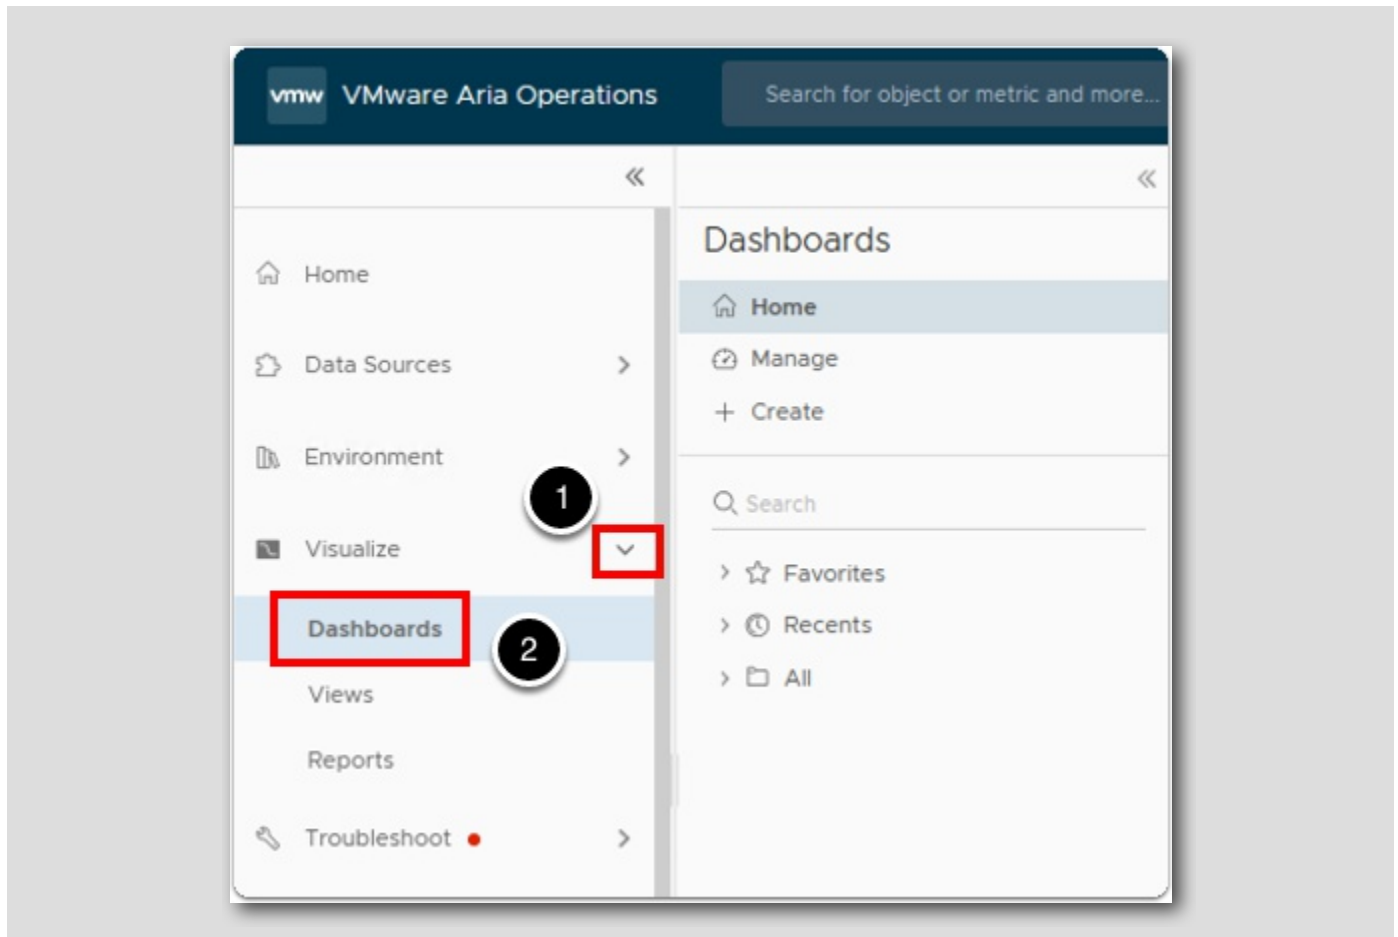
Perhaps you have multiple Aria instances and you need to import a dashboard that one of your teammates has created on a different Aria instance. This could be very useful and is very common when we have multiple Aria Operations instances or we have a Development instance that we use to develop and test our custom content.

We also have a great website to download and contribute cool dashboards we have made called the Dashboard Exchange. We can get to the VMware Aria Operations Sample Exchange quickly by visiting this site at <https://aria.vmware.com/sample-exchange/vmware-aria-operations-sample-exchange>

In this lesson, we'll highlight some downloadable community content that is available for use to use in vRealize Operations and we'll show how to access and import that content.

## Dashboards

[344]



1. Expand Visualize.
2. Click on the Dashboards.

## VMware Aria Operations Sample Exchange

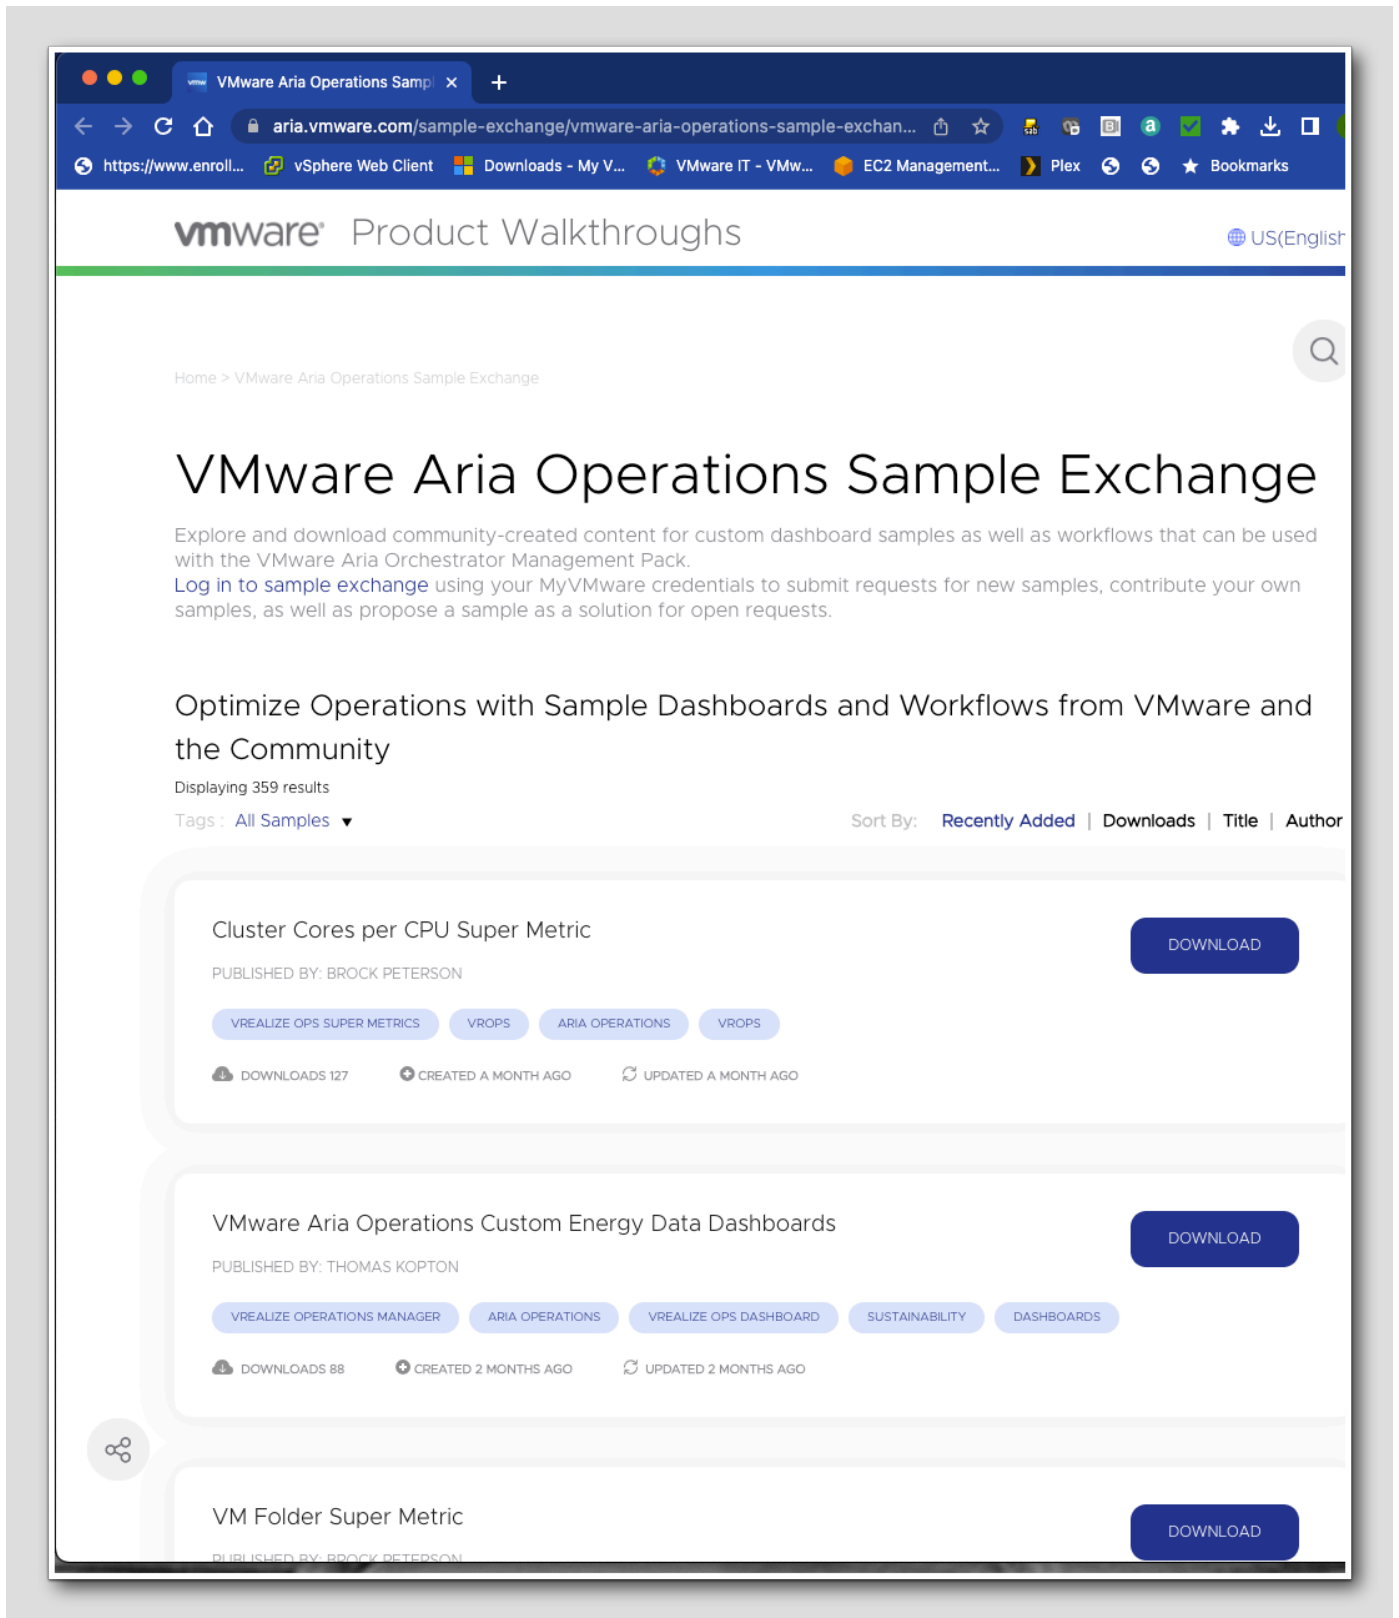
[345]

As we mentioned, we have a Dashboard Exchange that we can use to download community content. However, for this lab environment, we will use your browser to access that page instead of going through the Lab Environment UI due to firewall/proxy set up in the lab pod.

[Click here](#) to open the Aria Operations Sample Exchange in your browser.

## Aria Operations Sample Exchange

[346]



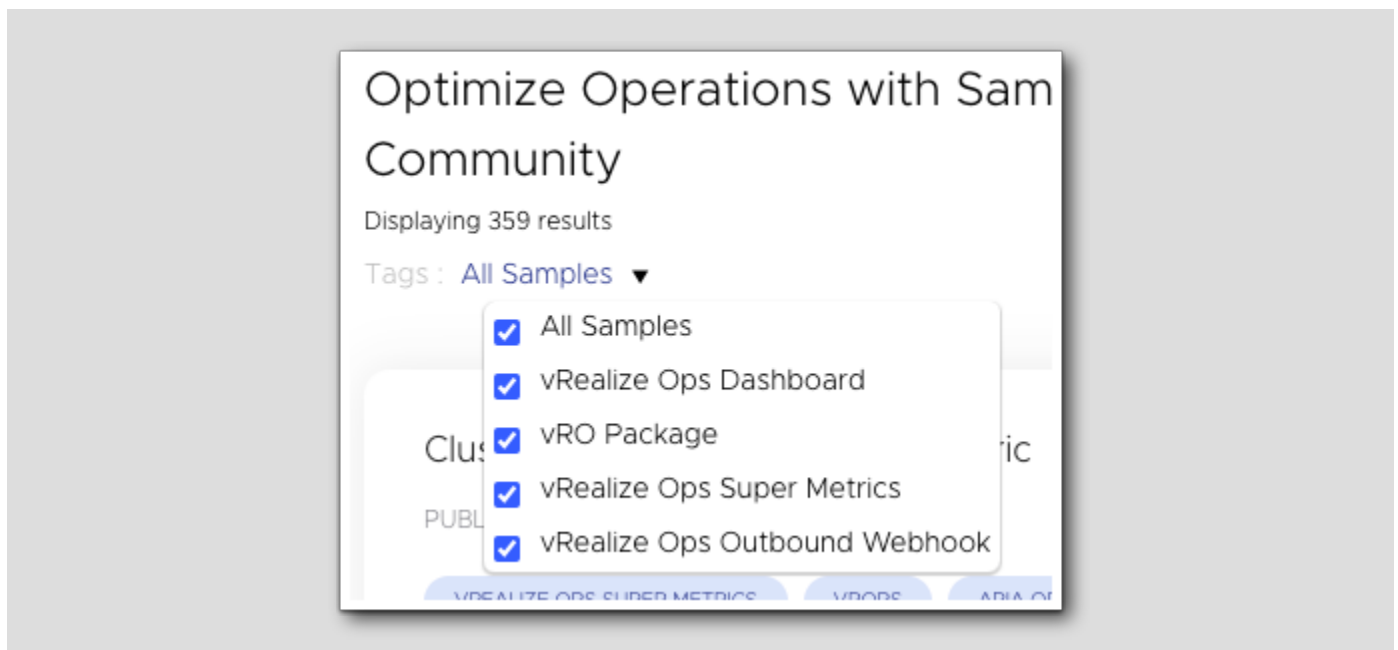


Everything in the Sample Exchange is community provided content and can be downloaded free of charge. If you login with your MyVMware account, you can also submit content that you've created to share with other users as well.

1. Hover over **All Samples** to show the different types of content that are available.

## Sample Types

[347]

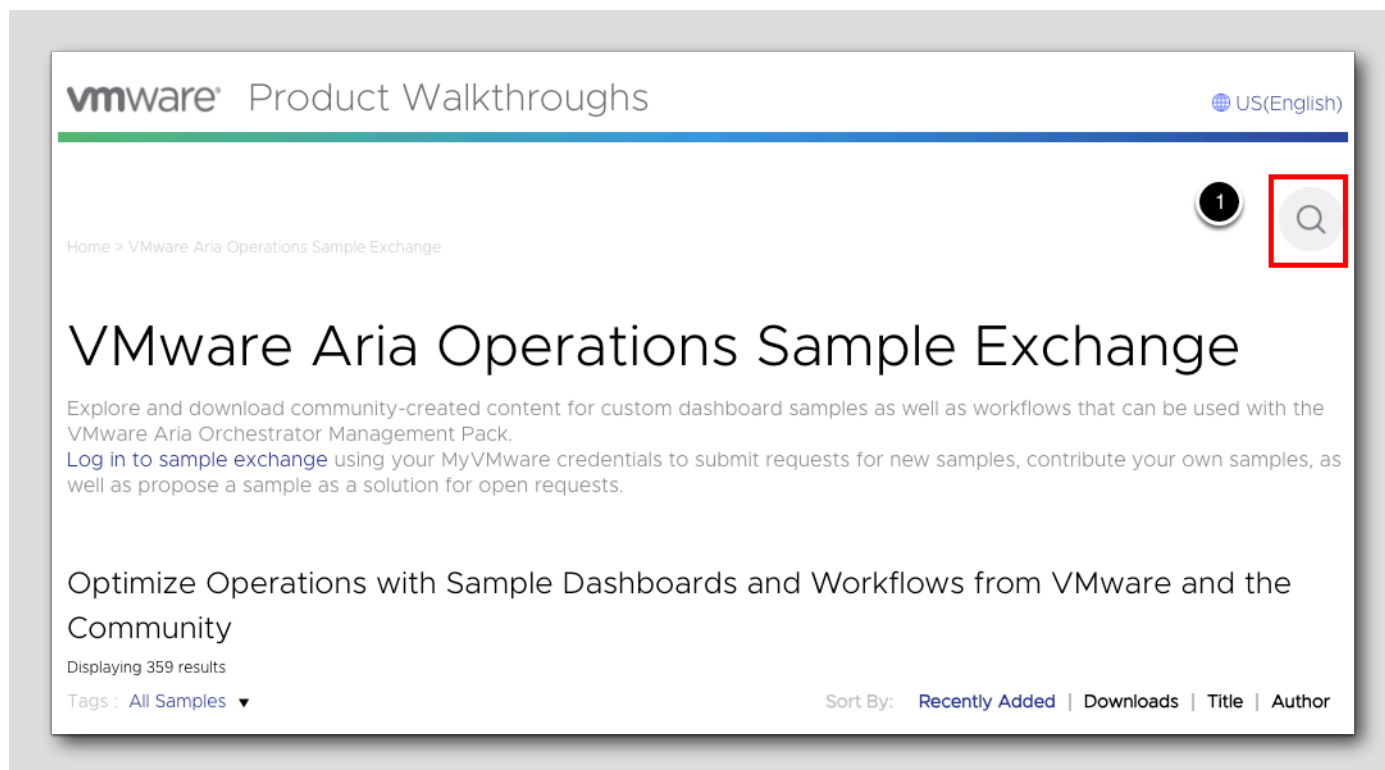


Here we can see the different types of content we can download from the sample exchange.

- vRealize Ops Dashboard - Dashboards, as we have covered in previous lessons, are great ways to quickly see a summary of various parts of your environment. There are many great dashboards built-in to Aria Operations, but the customization options are endless and they can be easily modified to suit your needs.
- vRO Package - Aria Orchestrator (vRO) packages can be used to further integrate Aria Operations into other products into the Aria Suite. For example, we can use a Aria Orchestrator package to allow us to assign a vRO workflow as a recommended action inside of a Aria Operations alert which greatly extends the capabilities of recommendations and/or automated actions inside of Aria Operations.
- vRealize Ops Super Metrics - Super Metrics are a way to create custom metrics inside of Aria Operations to discover metrics about your environment that the built in metrics won't cover. Note, HOL-2201-09-CMP will cover these super metrics in more detail.
- vRealize Ops Outbound Webhook - A webhook, which can also be called a HTTP push API, is a way for Aria Operations to provide other applications with data or information. The webhooks in the sample exchange provide examples of webhooks to some popular applications that you may want to receive data from Aria Operations.

## Search for our Sample Dashboard

[348]



vmware Product Walkthroughs US(English)

Home > VMware Aria Operations Sample Exchange

# VMware Aria Operations Sample Exchange

Explore and download community-created content for custom dashboard samples as well as workflows that can be used with the VMware Aria Orchestrator Management Pack.

[Log in to sample exchange](#) using your MyVMware credentials to submit requests for new samples, contribute your own samples, as well as propose a sample as a solution for open requests.

## Optimize Operations with Sample Dashboards and Workflows from VMware and the Community

Displaying 359 results

Tags: [All Samples](#) ▼

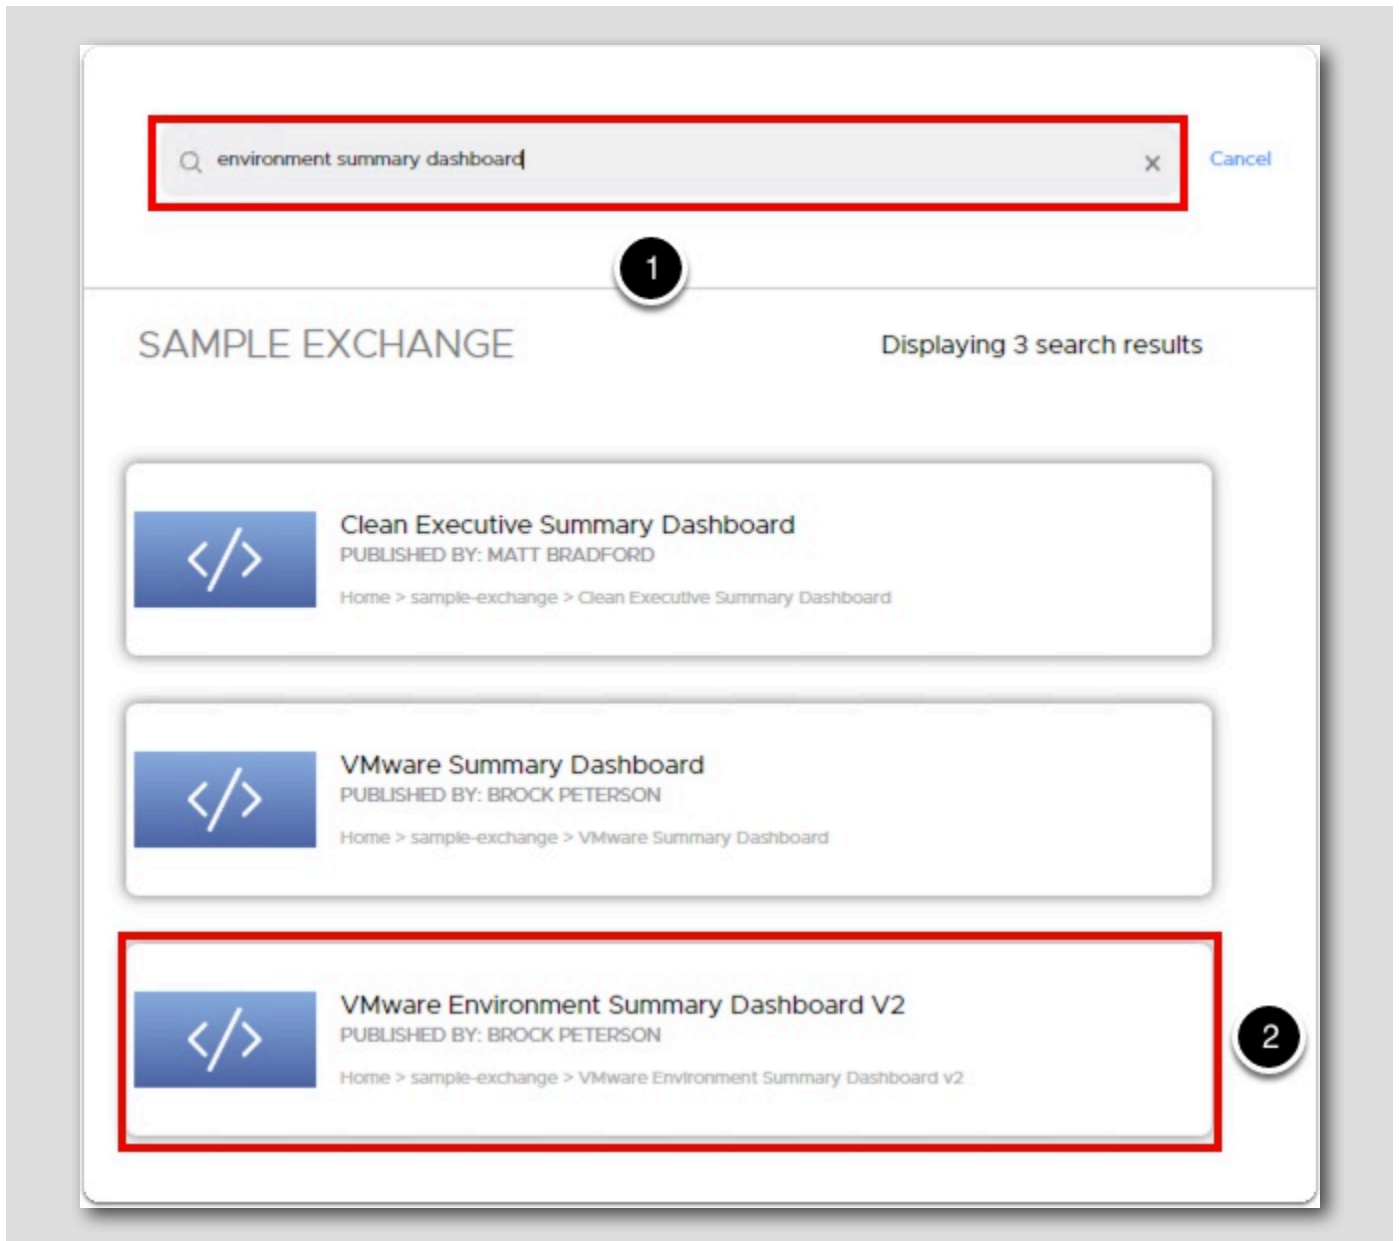
Sort By: [Recently Added](#) | [Downloads](#) | [Title](#) | [Author](#)

Let's search for a dashboard to use for this lesson.

1. Click the Search Icon on the sample exchange.

## Search for Environment Summary Dashboard

[349]



In this example, we will be selecting a dashboard that we can use in a Network Operations Center (NOC) to give us a good high level visual overview about the health of our environment. This search will return several results - a count this is continually increasing as people add additional content to this community sample exchange.

1. In the search box, enter **environment summary dashboard**, and hit **Enter**.
2. Scroll down or use your browser's search function (Ctrl-F) to search for **environment summary dashboard** and select the VMware Environment Summary Dashboard v2 once you find it.

Alternately, you can use this [direct link](#) to find the dashboard.

## Download Dashboard

[350]

# VMware Environment Summary Dashboard v2

PUBLISHED BY: BROCK PETERSON  DOWNLOADS 744  UPDATED 56 MONTHS AGO

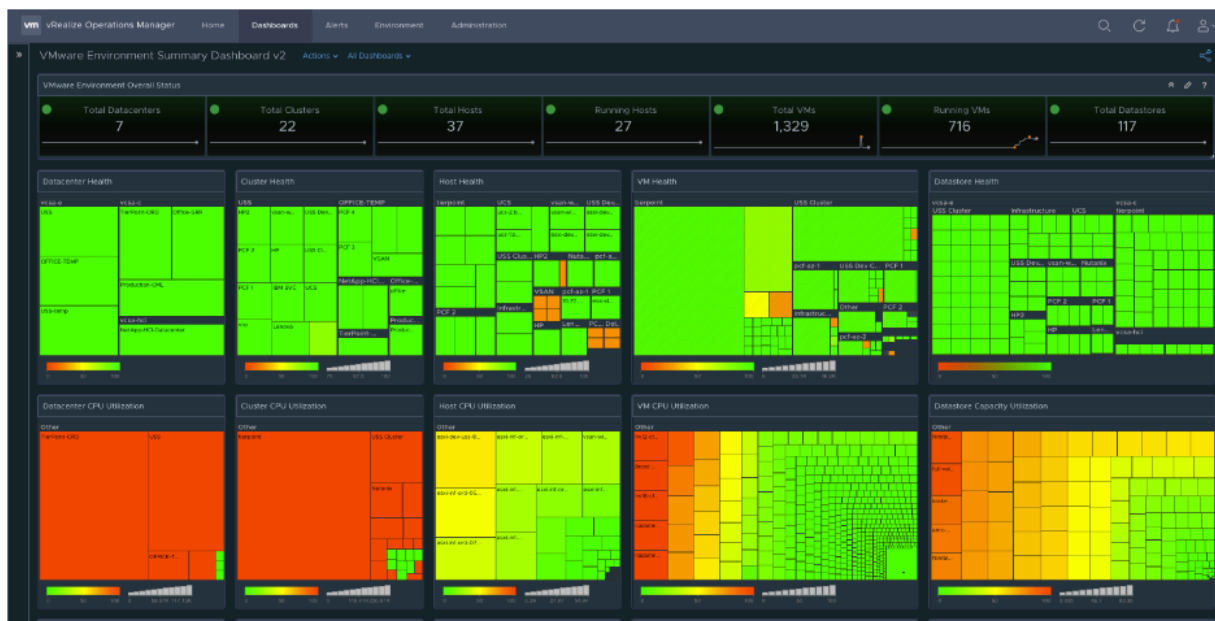
VREALIZE OPS DASHBOARD

VMWARE

VROPS

VROPS

This dashboard was developed in vROps 7.0 and is a riff on the original NOC dashboard from Luciano Da Silveira Gomes (<https://code.vmware.com/samples/4722/noc-dashboard?h=NOC>). It is designed to show important vSphere objects (metrics): Datacenters (CPU and Memory), Clusters (CPU and Memory), ESXi Hosts (CPU and Memory), VMs (CPU and Memory), and Datastores (Capacity and Latency).

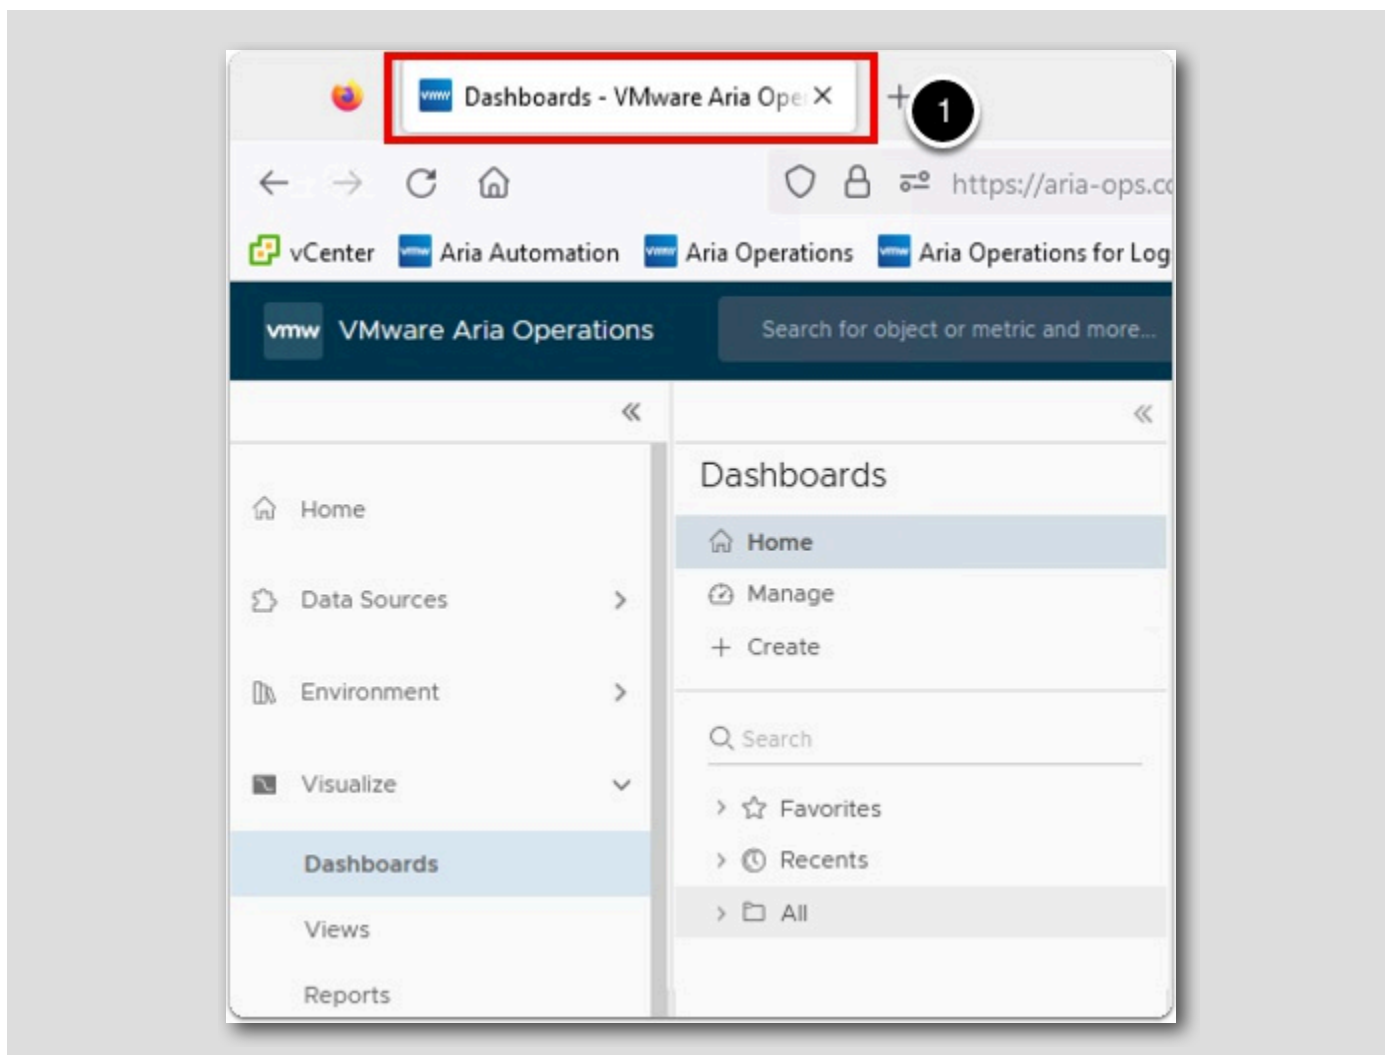
On this page we can see details about the dashboard. Some dashboards samples will also include screenshots of the dashboard showing what data it will include.

1. If you wanted to download the dashboard file, you would click on the **DOWNLOAD** button. However, we already have a copy of the file in our lab pod.

You can download the file and view the json if you want to see how a dashboard is stored as code for portability.

## Return to vRealize Operations

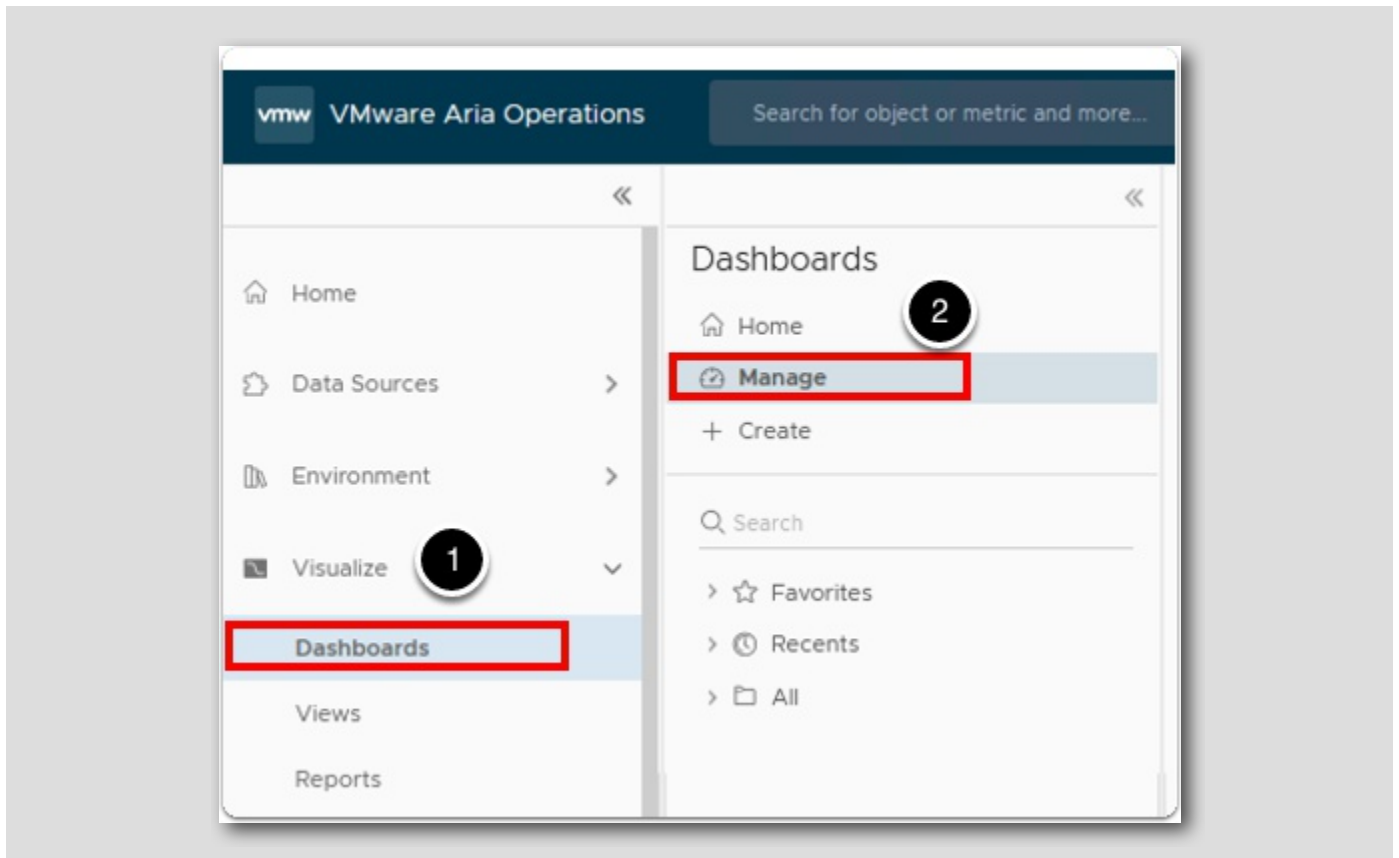
[351]



1. Return to the vRealize Operations Manager tab in your lab console.

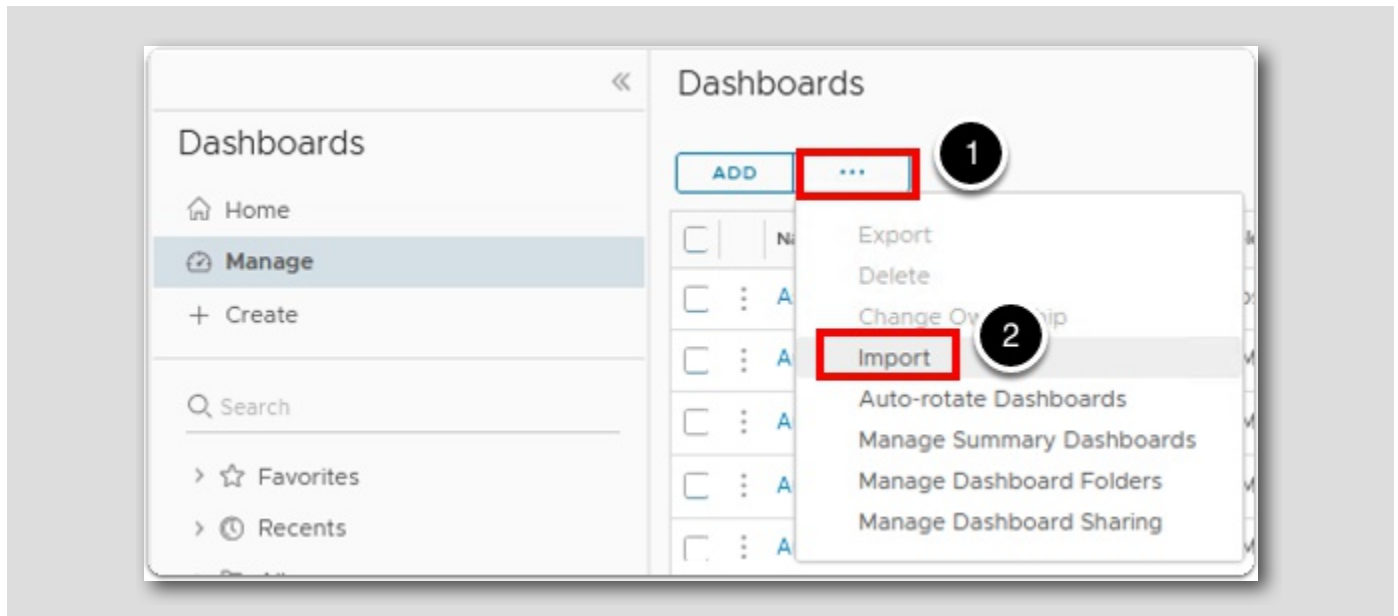
## Manage Dashboards

[352]



1. Click on Dashboards.
2. Click on Manage.

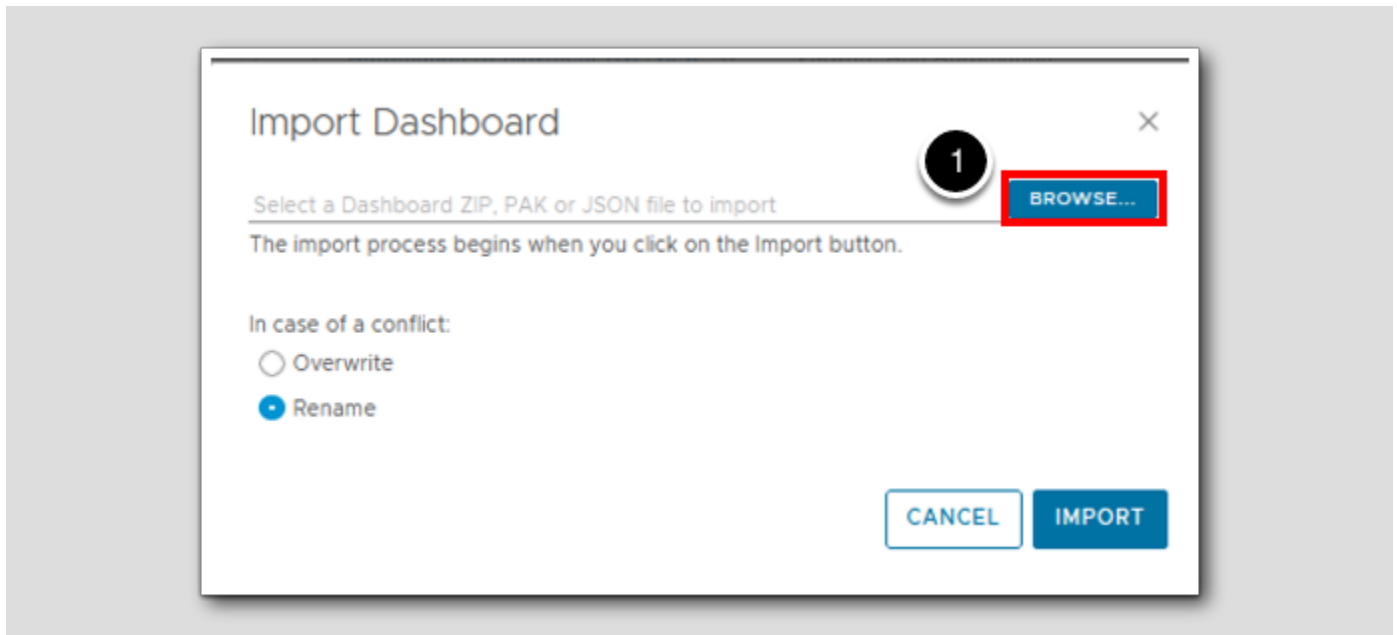
## Import Dashboard



1. On the Dashboards page, click the bottom with 3 dots to open the action menu.
2. Then click **Import** to start the import process.

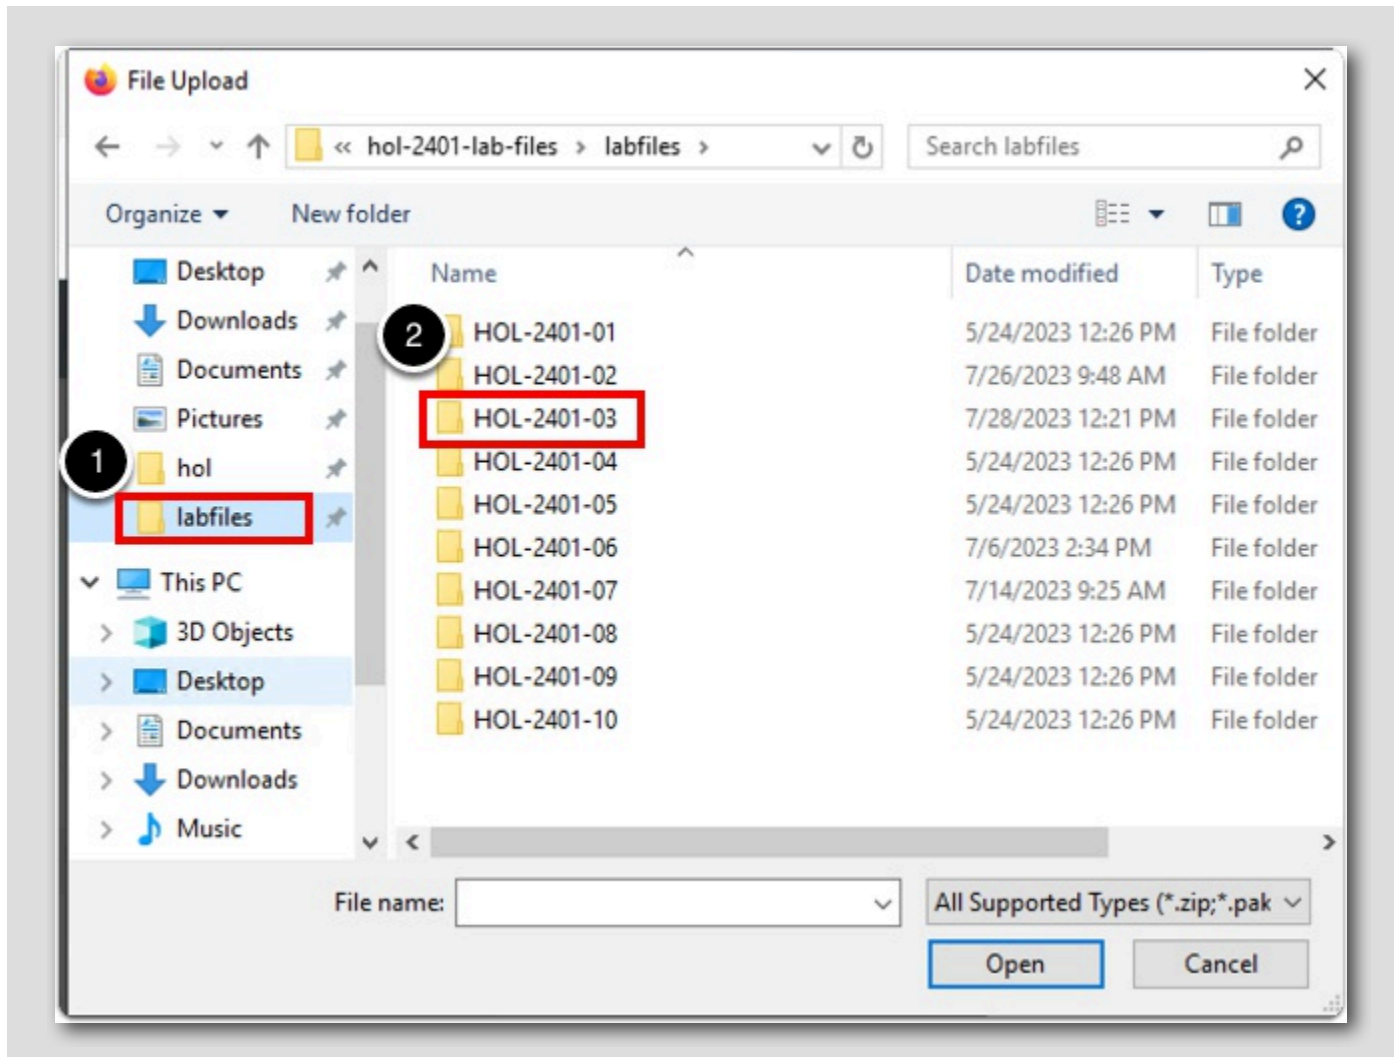
Browse to the Dashboard to import

[354]






## Import Dashboard File



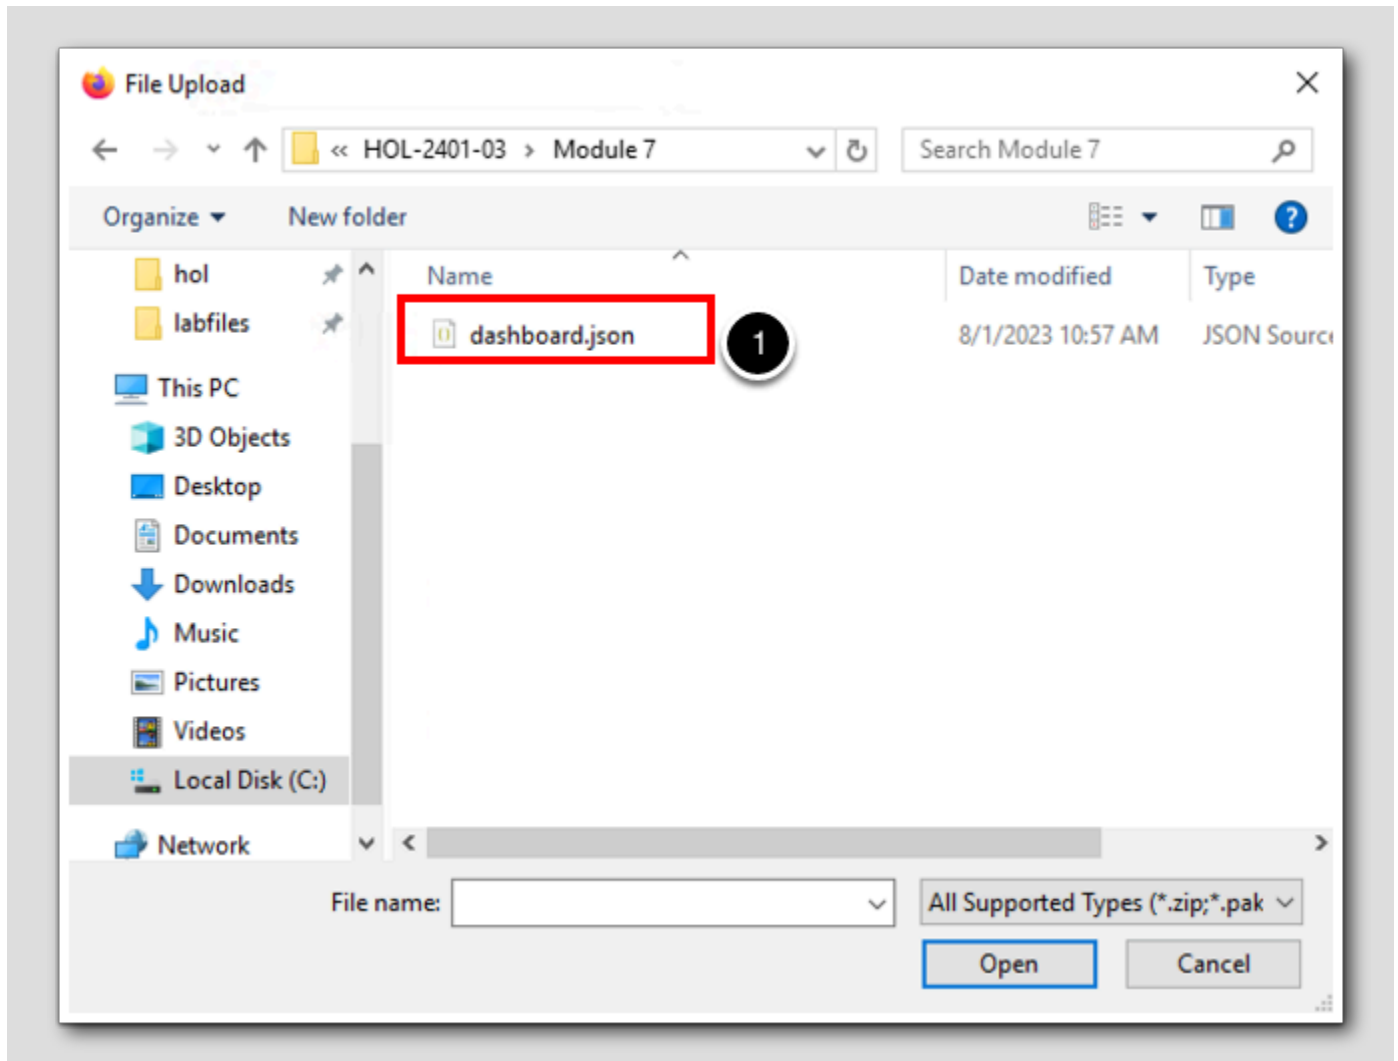
Note: We have already downloaded the dashboard file for you. It is in the Lab Files --> HOL-2401-03 --> Module 7 directory on your Main Console VM in the lab environment.

1. In the next window, click Lab Files
2. Double Click HOL-2401-03.



1. Double click on Module 7.

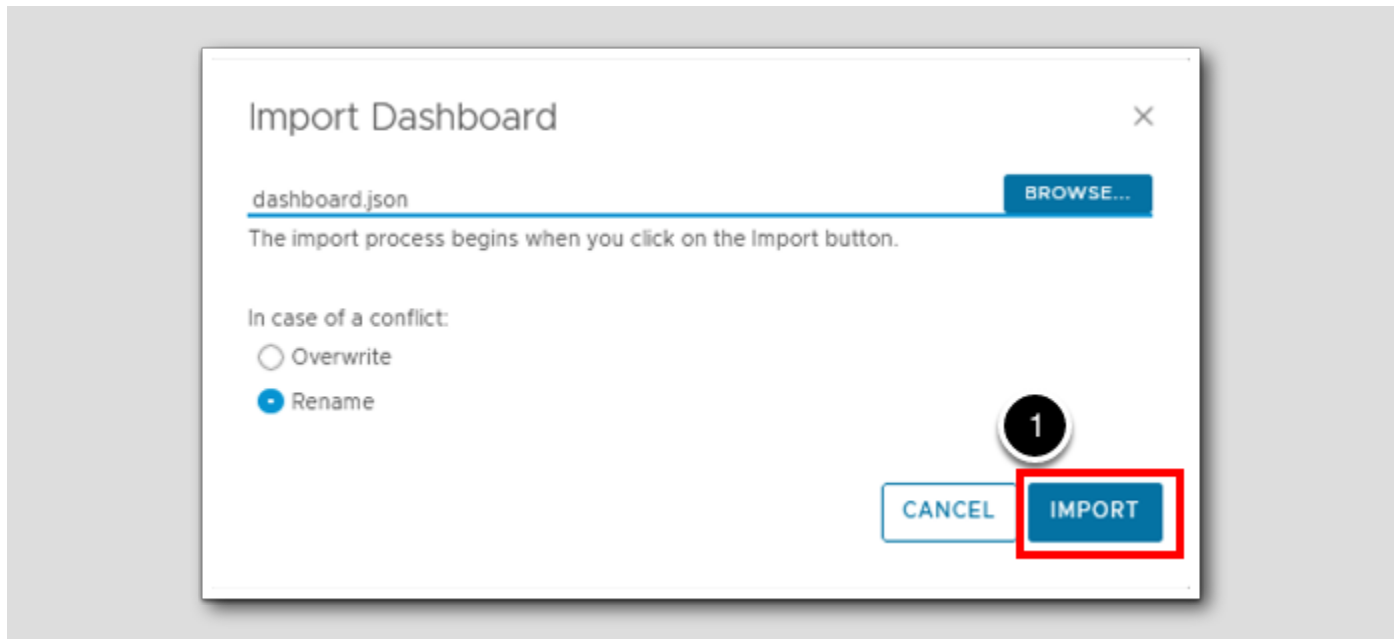
## Import Dashboard.json



1. Double click on the file Dashboard.json.

## Import

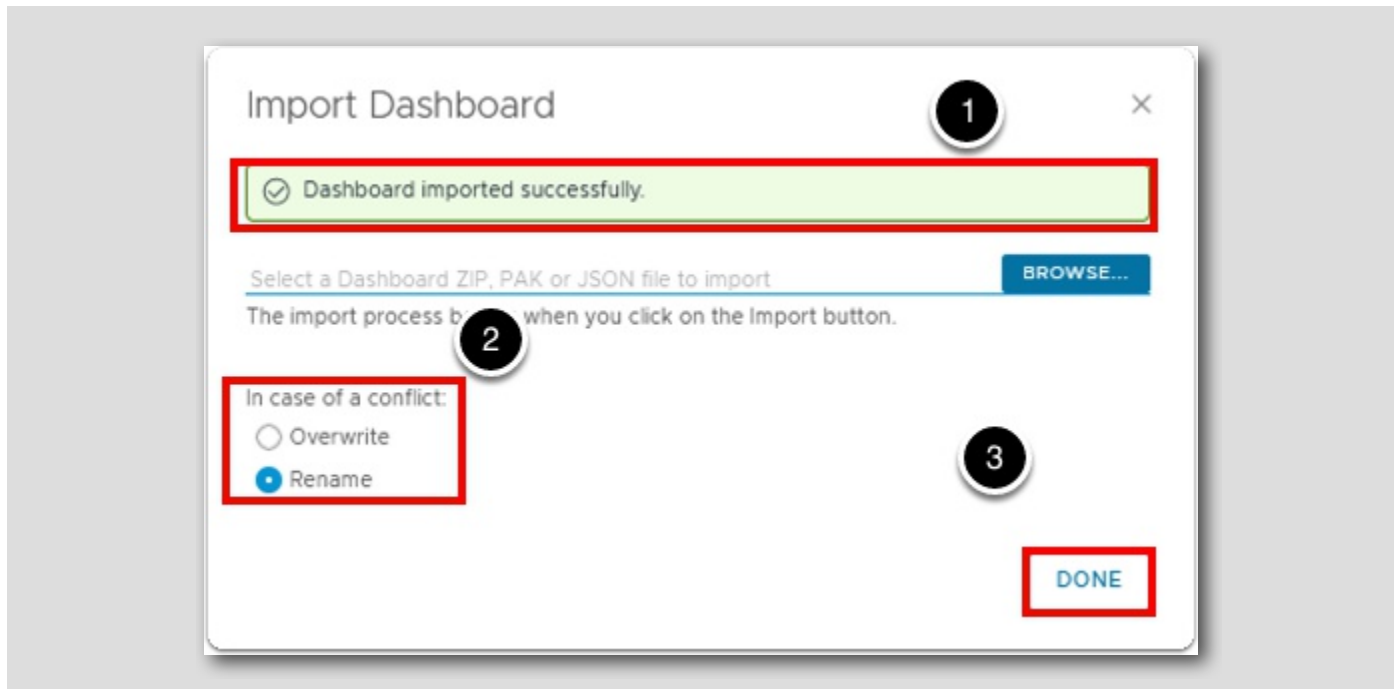
[358]



1. Click **IMPORT**.

## Dashboard Imported Successfully

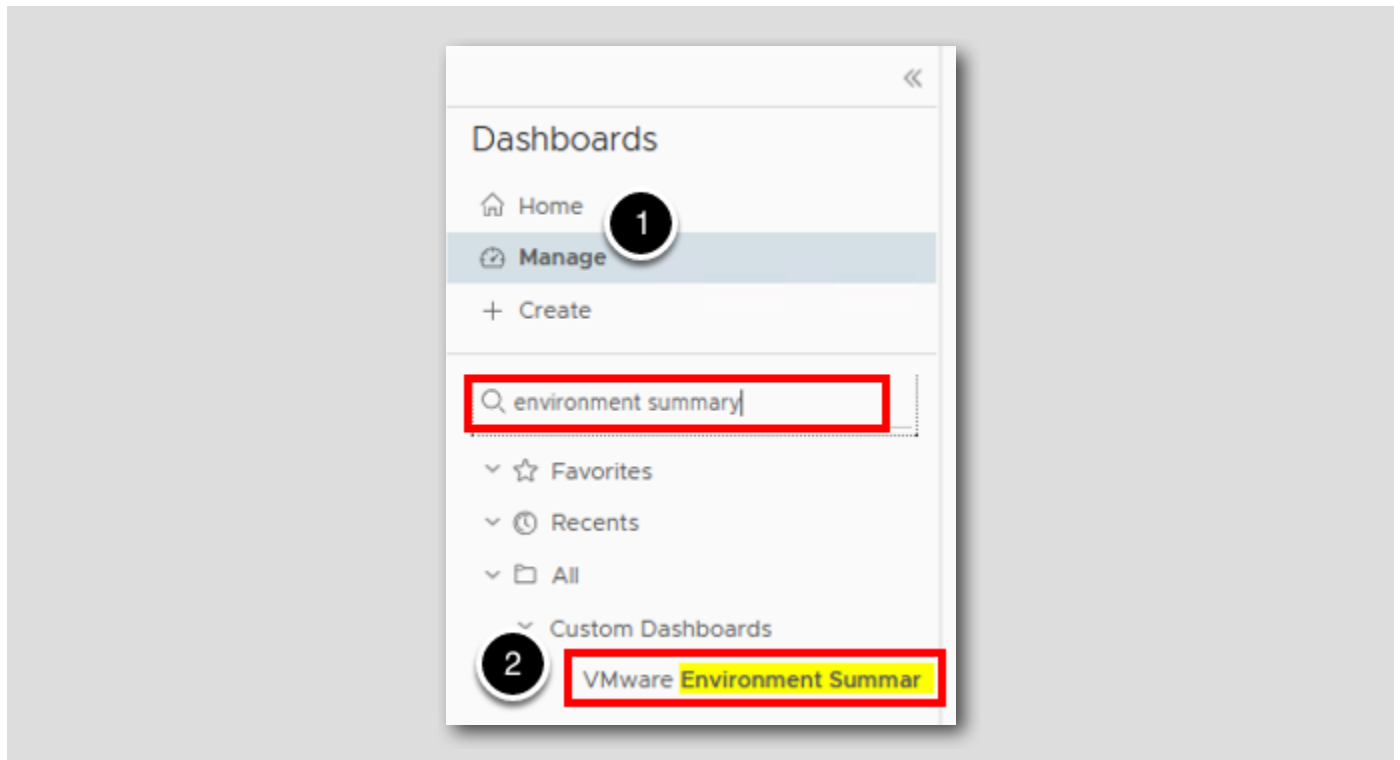
[359]



1. Here we can see the dashboard was successfully imported.
2. Notice here also that we have the option to either overwrite any existing dashboards that may have the same name as what you're importing. Or we can choose to rename the newly imported dashboard instead of overwriting existing content.
3. Click **DONE** to close this window.

## Find our Imported Dashboard

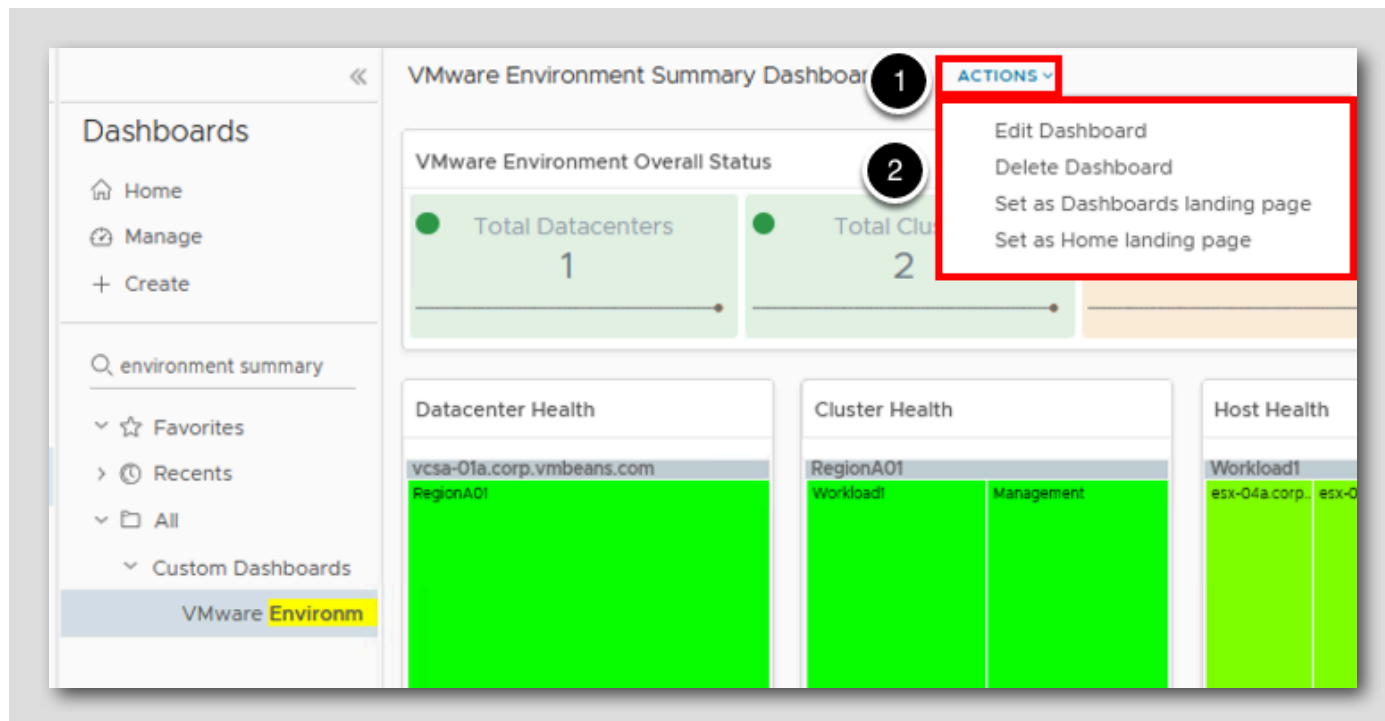
[360]



1. To find our imported dashboard, type `environment summary` in the search box and hit `Enter`.
2. Click on `VMware Environment Summary Dashboard V2`.

## Dashboard Actions

[361]



Now we can see our imported dashboard! As a final step, let's look at some of the available options.

1. Select **ACTIONS** beside the dashboard name.
2. Here we can see the available options for this newly imported dashboard to further integrate it into our environment.

## Lesson End

[362]

Congratulations, we have just completed the **Importing Dashboards** lesson.

In this lesson, we learned how to import dashboards into Aria Operations. We also learned about the Sample Exchange, which is a website where we can find some great ready-to-use content that we can use to easily extend the power of Aria Operations.

## Sharing Dashboards

[363]

In this lesson, we will learn how to share the numerous dashboards available in Aria Operations.

There are several very useful options for administrators to share dashboards to other personnel in their company. Now we can share a dashboard using a URL that can be given to ANYONE in our organization and they don't even need to be able to access our Aria Operations environment. This is a super useful feature when we need to share performance or capacity information to others in the organization, but don't want them logging into our Aria Operations instance.

We will see that we can also set an expiration time frame for the shared dashboard to be available. This is also really useful when you just want to give someone a view into a specific portion of the infrastructure for a limited period of time.

We can also share a dashboard through an email just by selecting the correct SMTP instance we have already set up in Aria Operations and entering the email of the recipient you want to have your new dashboard. Like with the other sharing options, we can also put an expiration time frame for the email as well.

We can even embed the dashboard into any other web page by simply copying the HTML code provided and pasting it into any system like Confluence or our own internal intranet portal.

Group sharing is simply giving dashboard access to any group that currently is set up through the authentication source we already have configured in Aria Operations.

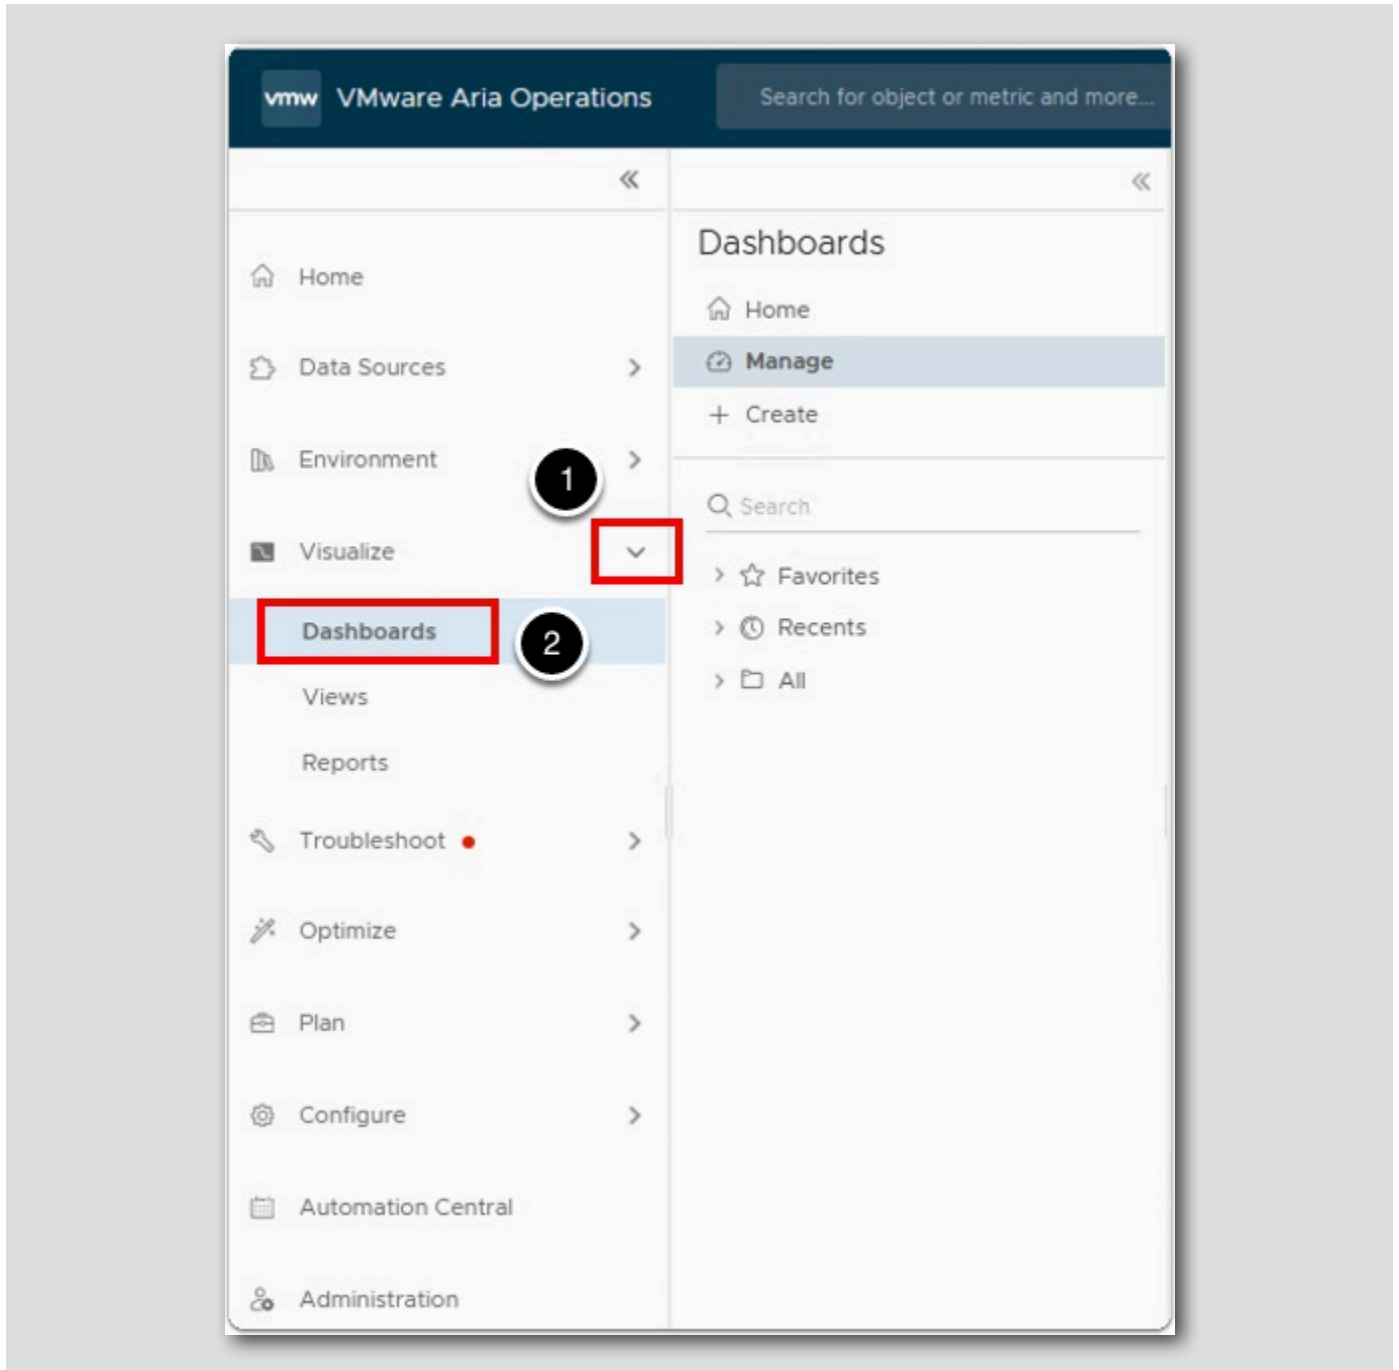
The final option gives us the ability to export the dashboard and move it to any other Aria Operations environment. This is very useful when we have multiple Aria Operations instances or we have a Development instance that we use to develop and test our custom content.

We have commonly seen the (NOC) Network Operations Center of an IT organization share dashboards on their large monitors in their NOC. They have created web pages that contain various bits of information from various monitoring systems in order to minimize the amount of monitors they have to have in the NOC. We can easily give them what they need by providing them an embedded link to the dashboard in which they can embed into their existing web portal. That way they don't have to add an additional monitor to house the Aria Operations dashboard. We will use this scenario in this lesson to learn how to share out the VMware Environment Summary Dashboard to them.



## Dashboards

[364]

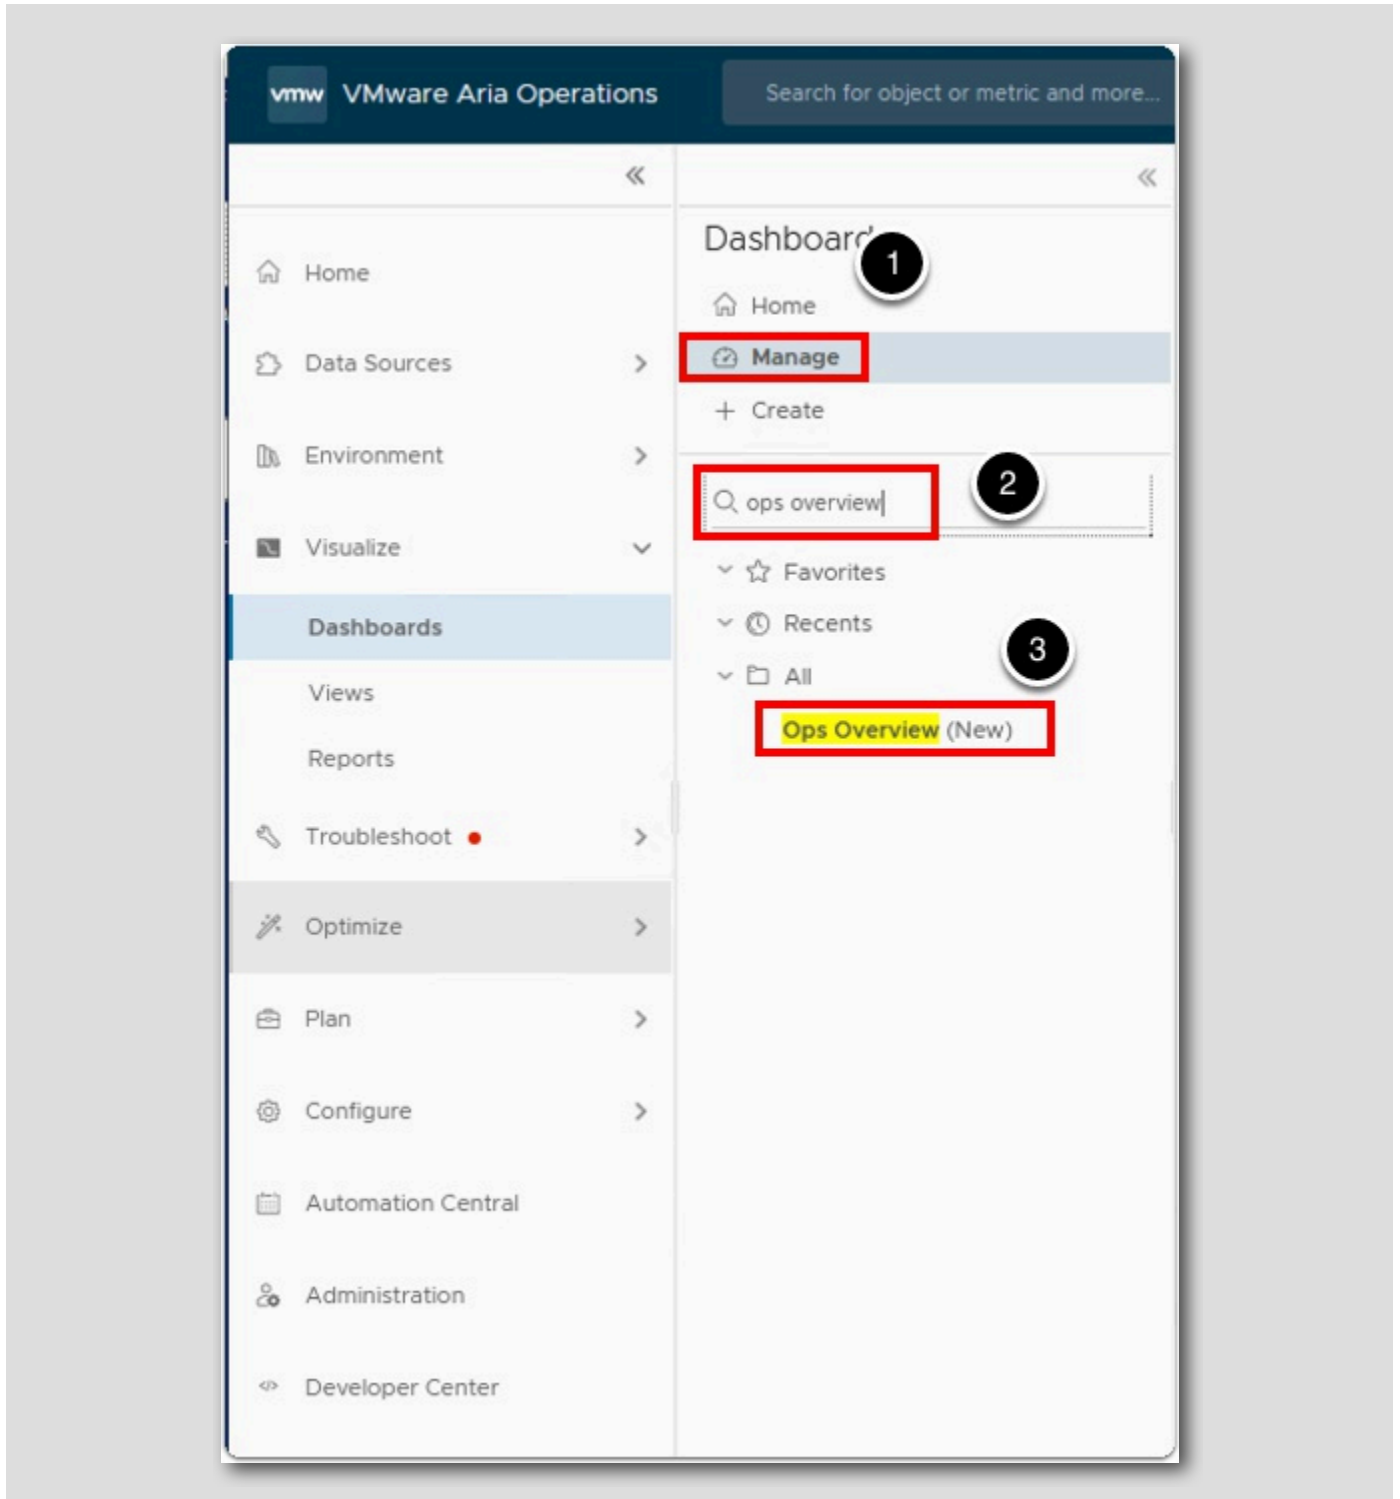


1. Expand Visualize.
2. Click on Dashboards.



## VMware Environment Summary Dashboard

[365]



In our example, we want to share a dashboard with the (NOC) Network Operations Center, so let's go to the Environment Summary dashboard.

1. Click on **Manage**.
2. In the search bar type **ops overview**.
3. Then click on **Ops Overview (New)**.

## Ops Overview (New) Dashboard Review

[366]

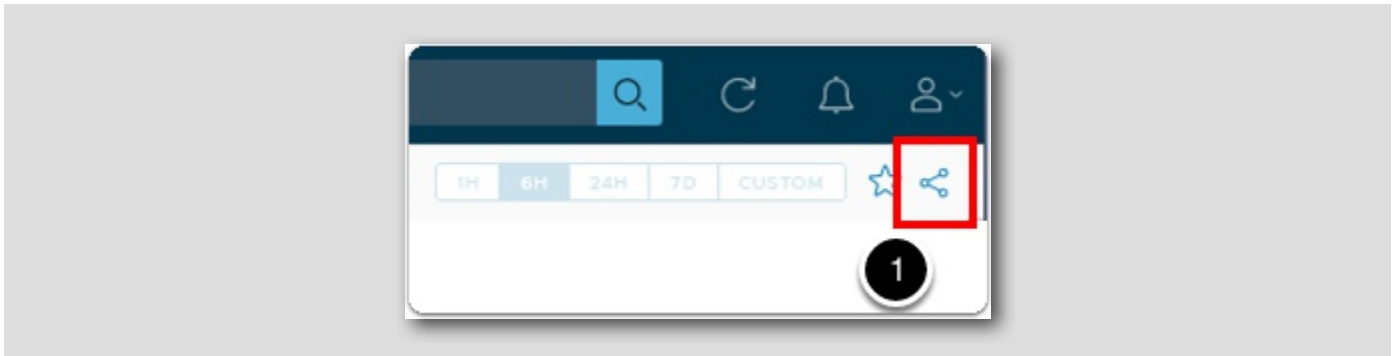
The screenshot displays the 'Ops Overview (New)' dashboard. At the top, there are filters for time ranges (1H, 6H, 24H, 7D, CUSTOM) and icons for favorites and sharing. The main content area is divided into several sections:

- Virtual Machines Table:** A table listing VMs with columns for Name, Adapter Type, Object Type, Policy, Collection State, and Collection Status. The first row is 'windows2019' with a green status icon.
- Object Relationship Diagram:** A hierarchical diagram showing 'Templates' and 'VM-RegionA...' at the top, with 'windows2019' in the middle, and 'RegionA01-I...' at the bottom. A red circle highlights the 'windows2019' node.
- Top Alerts:** A section titled 'No Issues'.
- Health:** A section with a green status icon, titled 'Health' with 'Immediate issues' below it. It includes a 'Health Trend' chart and a 'Why is health Good?' section.
- Top-10 VMs with CPU Contention:** A section with a table header 'CPU Contention (%)' and 'Objects'.
- Top-10 VMs with Memory Contention:** A section with a table header 'Memory Contention (%)' and 'Objects'.
- Top-10 VMs with Disk Latency:** A section with a table header 'Total Latency (ms)' and 'Objects'.

To recap this scenario, the NOC personnel want to have the new Ops Overview (New) Dashboard in the NOC at all times so they can monitor the troublemaking virtual machines after hours. We need to share this dashboard with them, but remember they have a web portal that they use. Therefore, we will need to provide them the embedded link that they can simply add to their existing web portal.

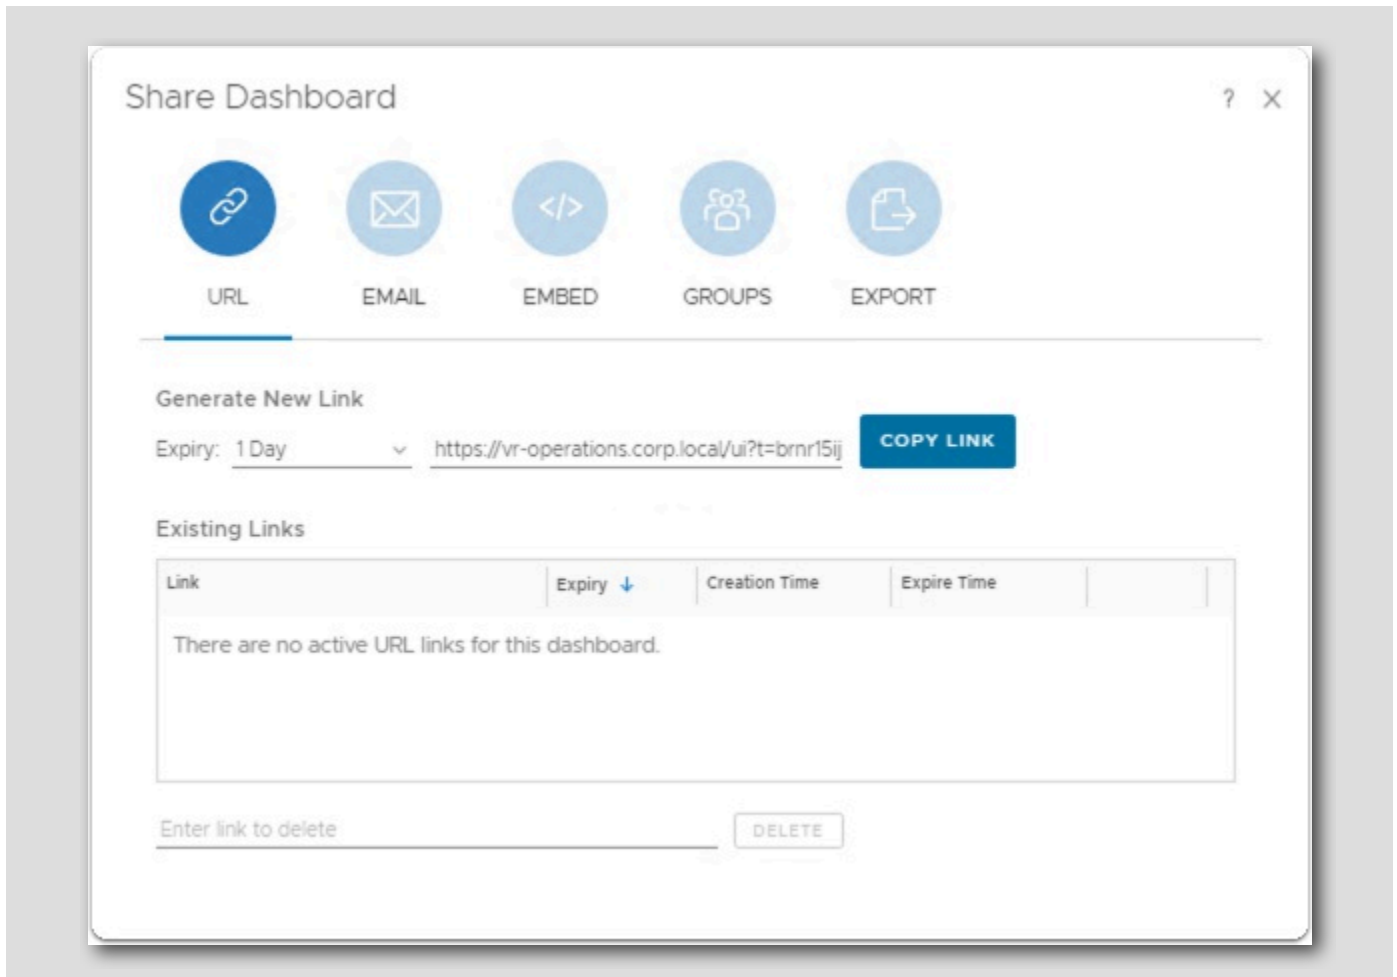
## Share Dashboard

[367]



1. Click on the **Share Dashboard** (three connected circles) icon in the upper right-hand side of the VMware Environment Summary Dashboard.

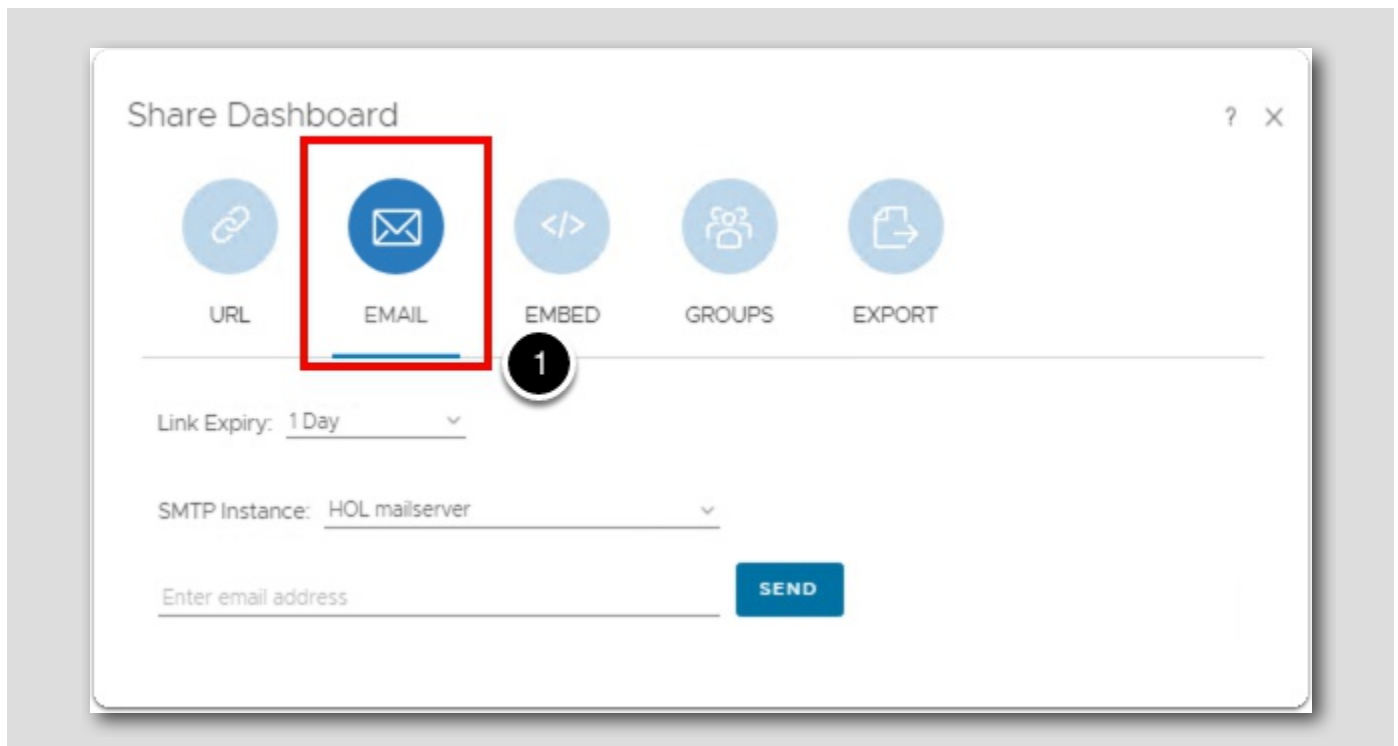
## Share Dashboard - URL



In this example we can simply create a URL to provide to anyone so they can view the dashboard. For Link Expiry, we have the options to select (1) Day, (1) Week, (1) Month, (3) Months and Never Expire. We see that the link to the dashboard is already filled in. We would then click on the COPY LINK button to copy it to the computer's clipboard allowing us to copy it into a file, email, etc.

## Share Dashboard - Email

[369]



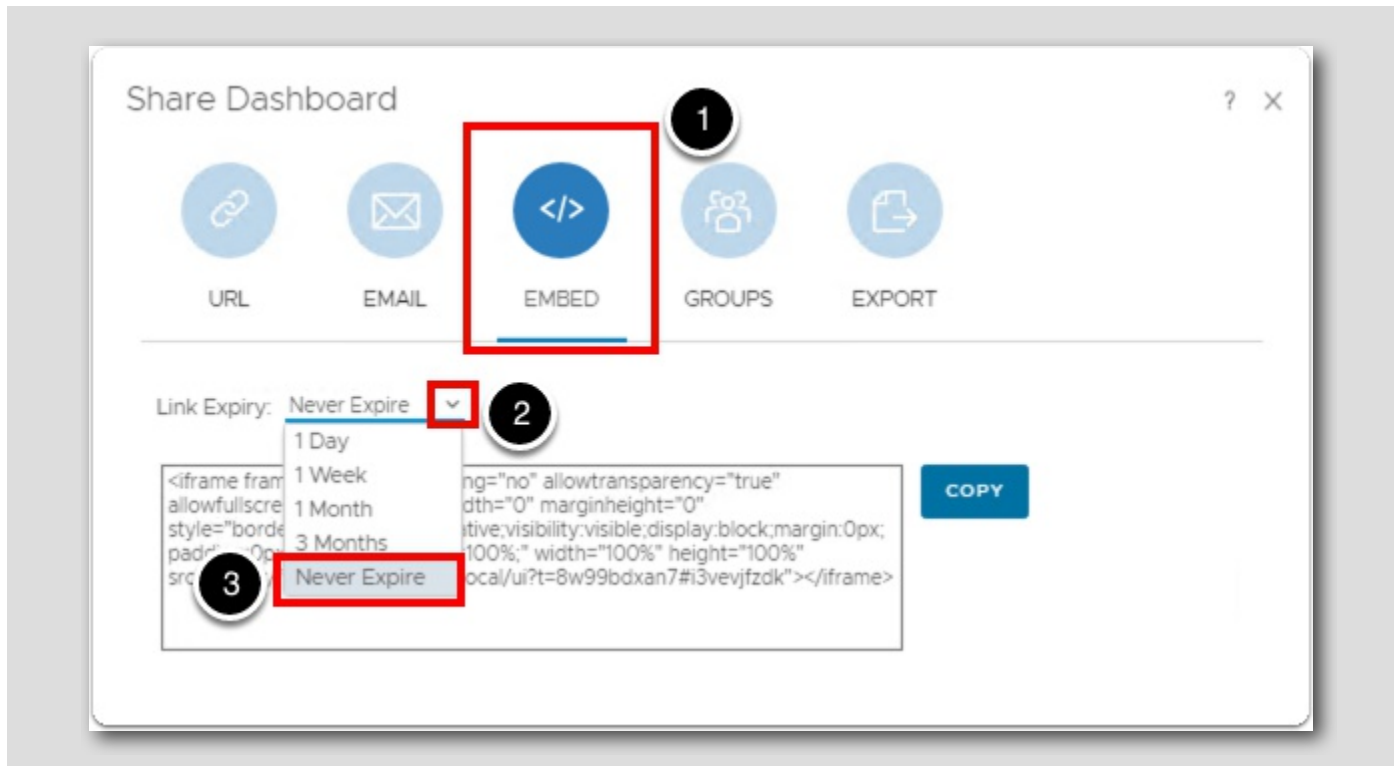
In this example, we want to send the dashboard link to someone via an email address directly from the Aria Operations interface. As a note, we won't actually be sending the link to the dashboard to an email address. We will just run through the steps as though we are going to.

1. Click on the **EMAIL** icon to select email as an option to send the link.

Again, we have the options to select (1) Day, (1) Week, (1) Month, (3) Months and Never Expire. In this lab environment, we do not have an SMTP instance configured. In a production environment, we would configure this with the company email server information by clicking on the CONFIGURE button if it wasn't already configured within Aria Operations. Then we would type the email address of the individual we are sending the link to. Finally, we would click the SEND button to send the email with the link to the dashboard to the receiver.

## Share Dashboard - Embed

[370]

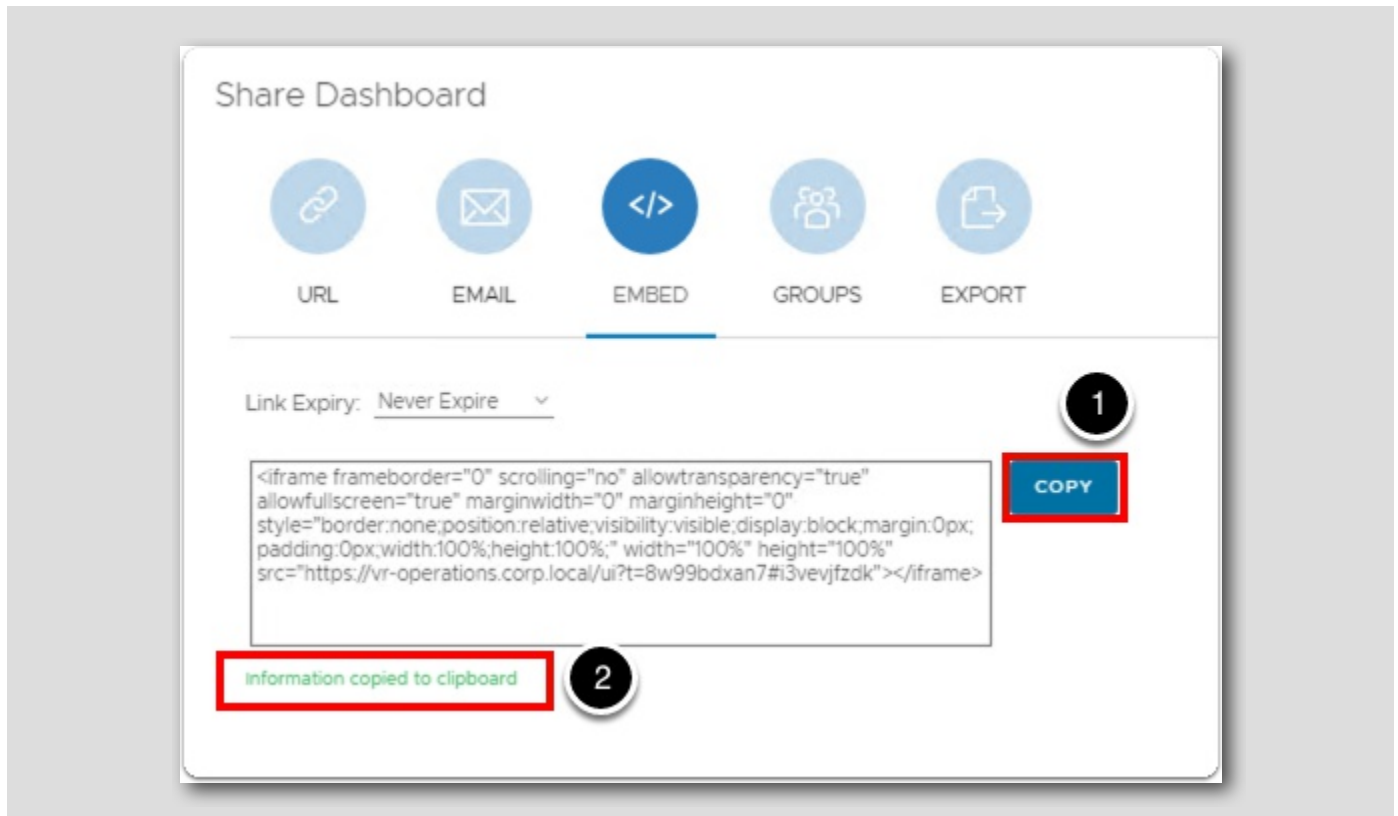


In the introduction of this lesson, we discussed the scenario of the (NOC) Network Operations Center having a web page that they wanted to embed the Ops Overview (New) Dashboard in. We will now go through the steps associated to accomplish providing them the embedded dashboard.

1. Click on the **EMBED** icon in the Share Dashboard pop-up window. We want this selection since they want the dashboard included in their existing web portal in the NOC.
2. Next click on the **arrow** next to **Link Expiry**: to expand its drop-down menu.
3. Click on **Never Expire** since we want this for the NOC and do not want their access to expire.

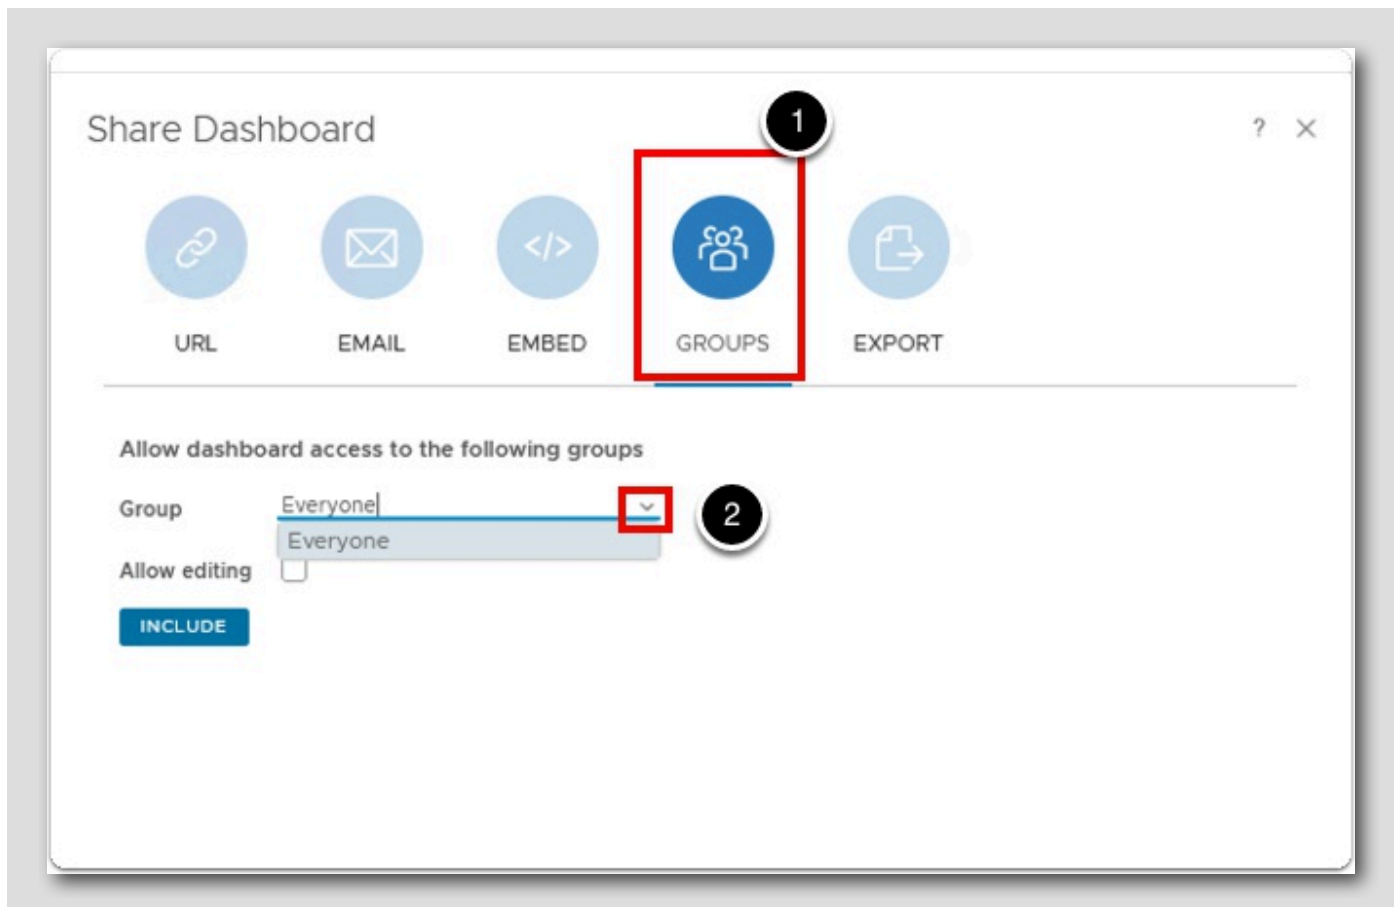


## Share Dashboard - Copy to Clipboard



1. Click on the COPY button.
2. We see that by clicking the COPY button, that the embedded link has been successfully copied to the clipboard. It is now ready to be copied into an email or some other method in which we can provide the NOC the link to use in their web portal.

## Share Dashboard - Groups



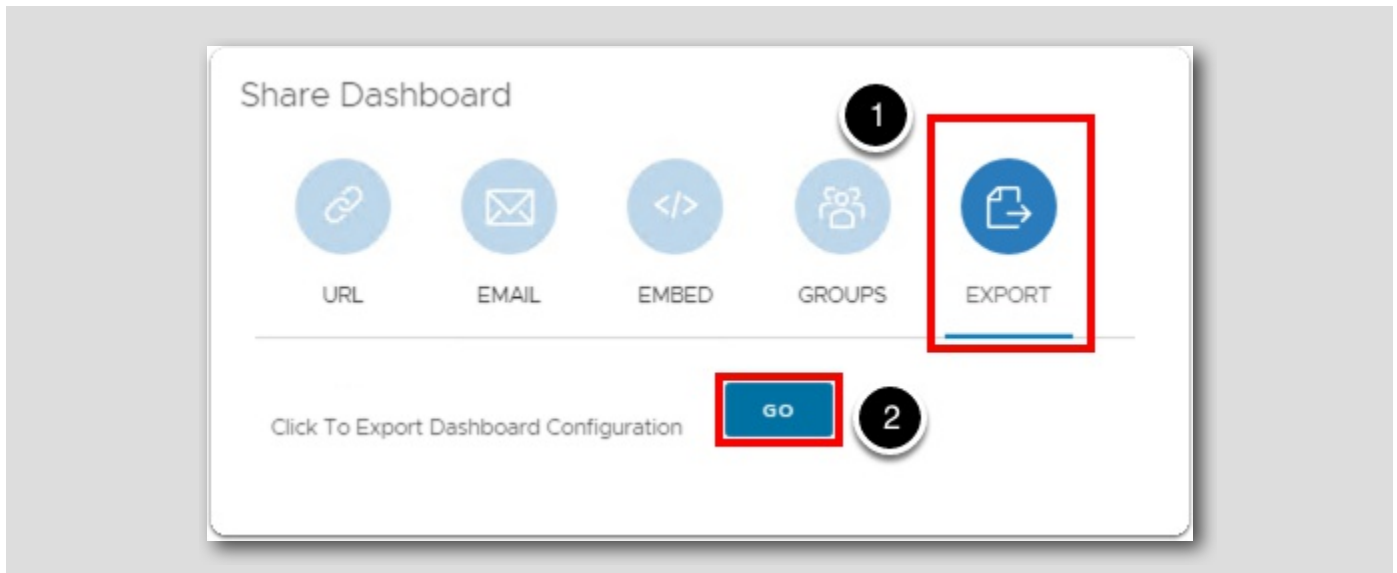
In this example, we need to authorize only a previously established security group in Aria Operations access to this dashboard. Currently the Everyone group has access to this dashboard.

1. Click on the **GROUPS** icon in the pop-up window.
2. Click on the **arrow** to the left of the **INCLUDE** button. We see that we only have two options configured in this environment. This list will vary from one environment to another based on what groups have been configured within Aria Operations.

We would then click on the **INCLUDE** button to give this group(s) access to the dashboard.

## Share Dashboard - Export

[373]




Remember, this Ops Overview (New) dashboard we are currently in is a custom dashboard is not a default out-of-the-box dashboard. We want to export this dashboard because we have another instance of Aria Operations in a (DR) Disaster Recovery datacenter and want to have the same dashboard in that instance as well. So we need to export the dashboard and then import it into the instance in the DR datacenter.

1. Click on the **EXPORT** button in the pop-up window.
2. Then click on the **GO** button to export it into a ZIP file format.

## Share Dashboard - Download File

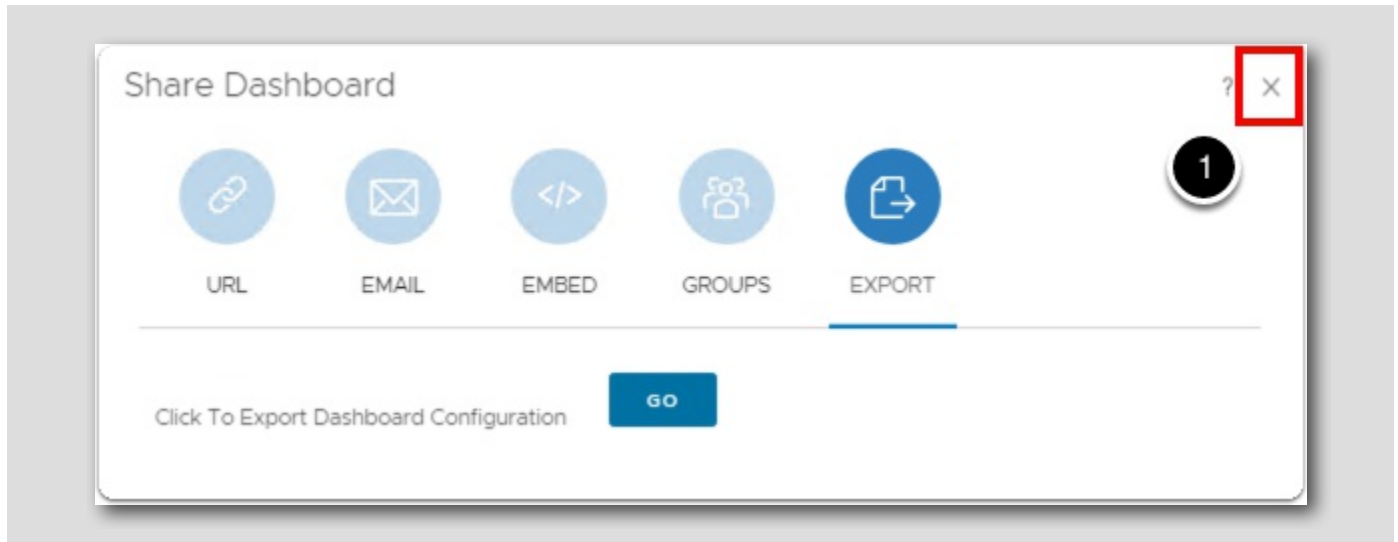
[374]



We see that it will download the dashboard as a ZIP file. We could then copy this ZIP file over the DR site and then import it into that Aria Operations instance.

## Share Dashboard - Exit

[375]



That's it, we have gone through all the options for sharing dashboards in Aria Operations!

1. Click on the X in the upper right-hand corner of the Share Dashboard pop-up window to close it.

## Lesson End

[376]

Congratulations, we have just completed the **Sharing Dashboards** lesson which is the last lesson of **Creating and Managing Dashboards!**

In this lesson, we learned how to share Aria Operations dashboards through various methods. We can share them via a URL, Email, Embedded file, Groups or Export the dashboard to import into another instance of Aria Operations.

## Conclusion

[377]

In this module, cloned an existing Dashboard, built a Dashboard from scratch and learned how to share Dashboards with other groups.

## You've finished the module

[378]

Congratulations on completing the lab module.

For more information on getting started with Aria Operations, see the [VMware Aria Operations: Journey to Success](#) guide at the [VMware Apps & Cloud Management Tech Zone](#).

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Module 8 - Enhancing depth of VMware Aria Operations with Super Metrics (15 minutes) Basic

### Introduction

[380]

The Troubleshooting Workbench is where you perform advanced troubleshooting tasks on an alert that triggered on an object. You can investigate both known and unknown issues in VMware Aria Operations. It was specifically designed to focus in and out of an object to quickly identify if there is an issue with a specific object or, by providing the ability to zoom out the scope, to see if there is a systemic issue within the infrastructure.

### Log in to Aria Operations

[381]

We will log in to a live instance of Aria Operations running in this lab.

### Open the Firefox Browser from the Windows Task Bar

[382]

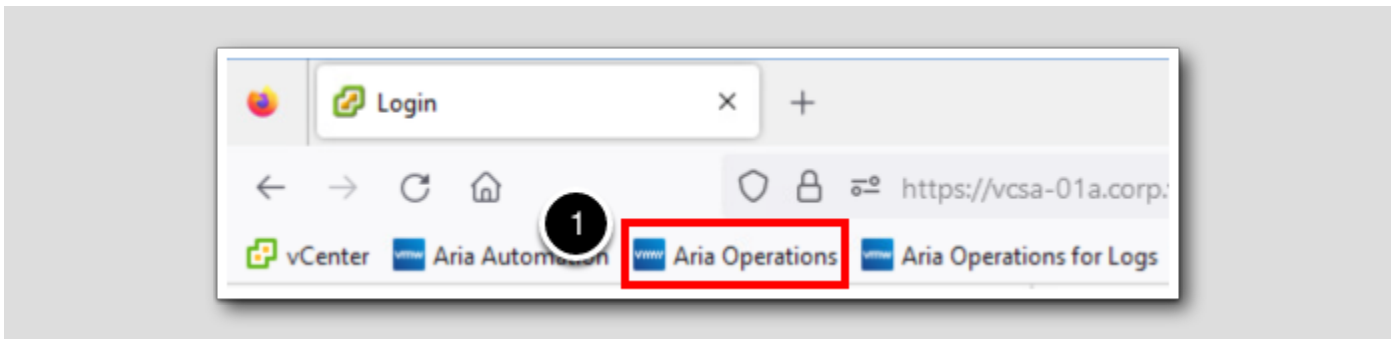


If the browser is not already open, launch Firefox.

1. Click the **Firefox** icon in the Windows Quick Launch Task Bar at the bottom of the screen.

### Navigate to Aria Operations


[383]



1. Click the **Aria Operations** bookmark in the bookmarks toolbar.

## Log in to Aria Operations

[384]



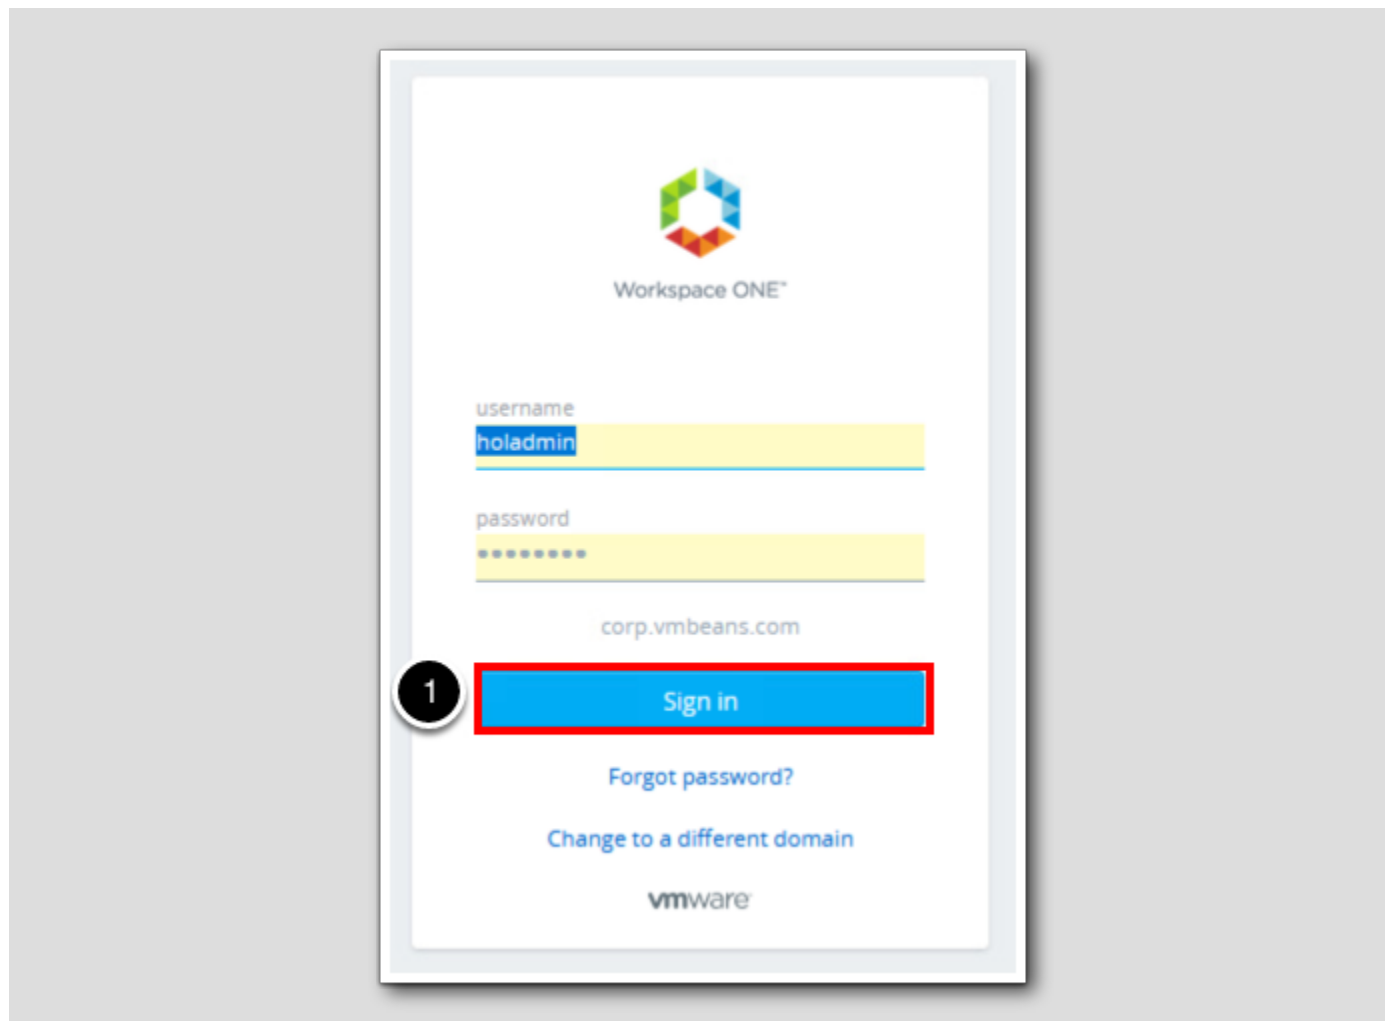
Aria Operations is integrated with VMware Workspace ONE Assist (also known as VMware Identity Manager) in this lab. This integration is listed as vIDMAuthSource in our live lab environment.

vIDMAuthSource may be pre-selected as the default identity source. If it is not, then you will need to select it.

1. Click the **drop-down arrow** and select vIDMAuthSource if it is not already selected.
2. Click **REDIRECT** to be taken to the authentication page.

## VMware Identity Manager Login

[385]



VMware Identity Manager acts as the identity provider for the Active Directory authentication source in this lab.

Credentials for the default user, holadmin, have already been provided.

1. Click **Sign in**

## Understanding Hierarchical Relationships in Aria Operations

[386]

Before we jump into creating super metrics, it is first important to understand that Aria Operations maintains several hierarchical relationship trees. And whenever you install additional management packs for extensibility, each management pack will add at least one additional hierarchy in Aria Operations.

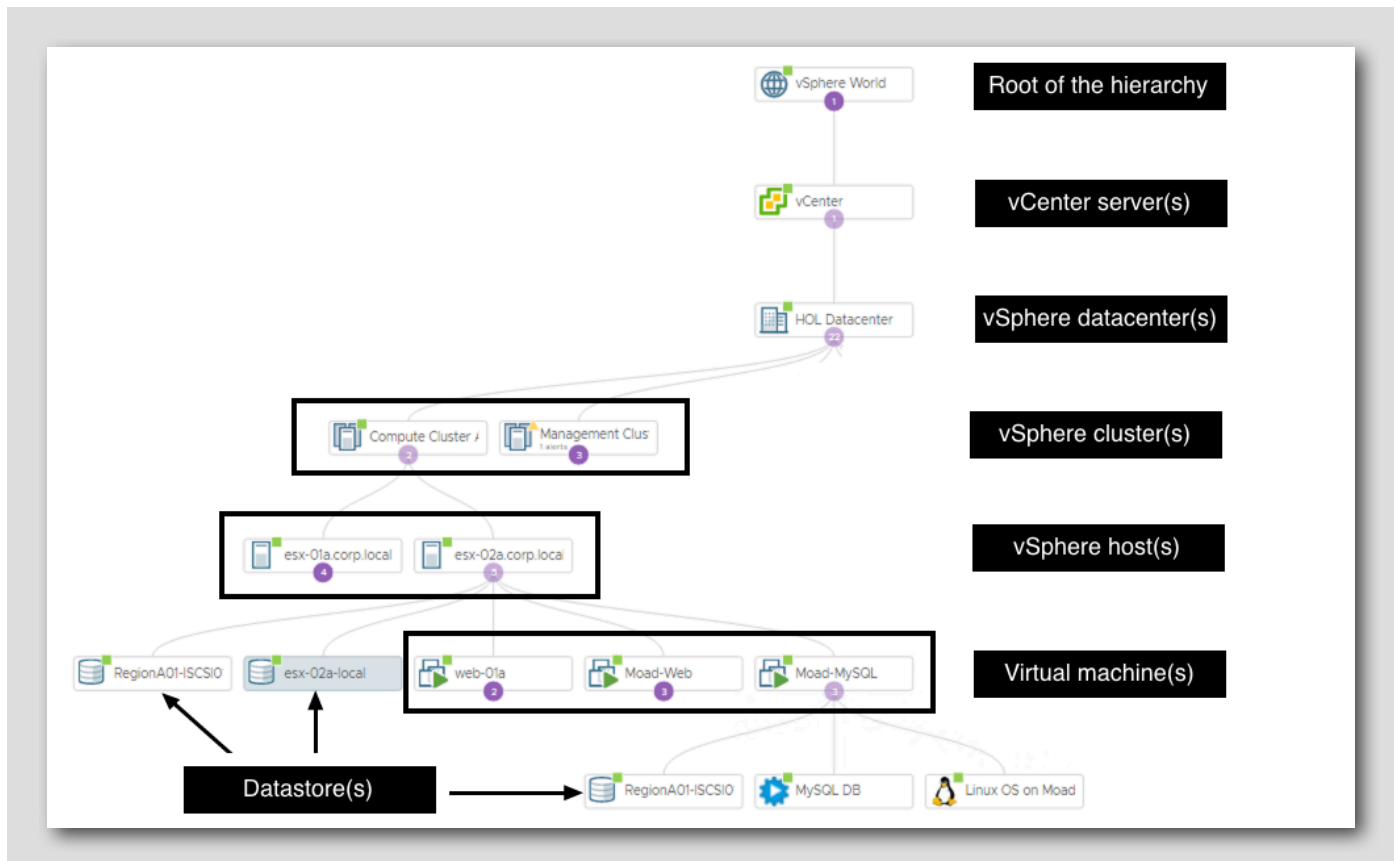
This is important to understand in the context of super metrics because unless you are creating a new metric on an object or object type



that is based only on metrics from that same object/object type you will need to know where in the hierarchy the related object types are. For example, in the vSphere Hosts and Clusters hierarchy, a virtual machine is a child of a host. If you want to create a super metric for hosts that shows the average CPU usage across all virtual machines that are running on a given host, you need to write your super metric formula with the proper syntax to look one level down from the host to the virtual machines for the metric inputs to the super metric.

## vSphere Hosts and Clusters Hierarchy

[387]

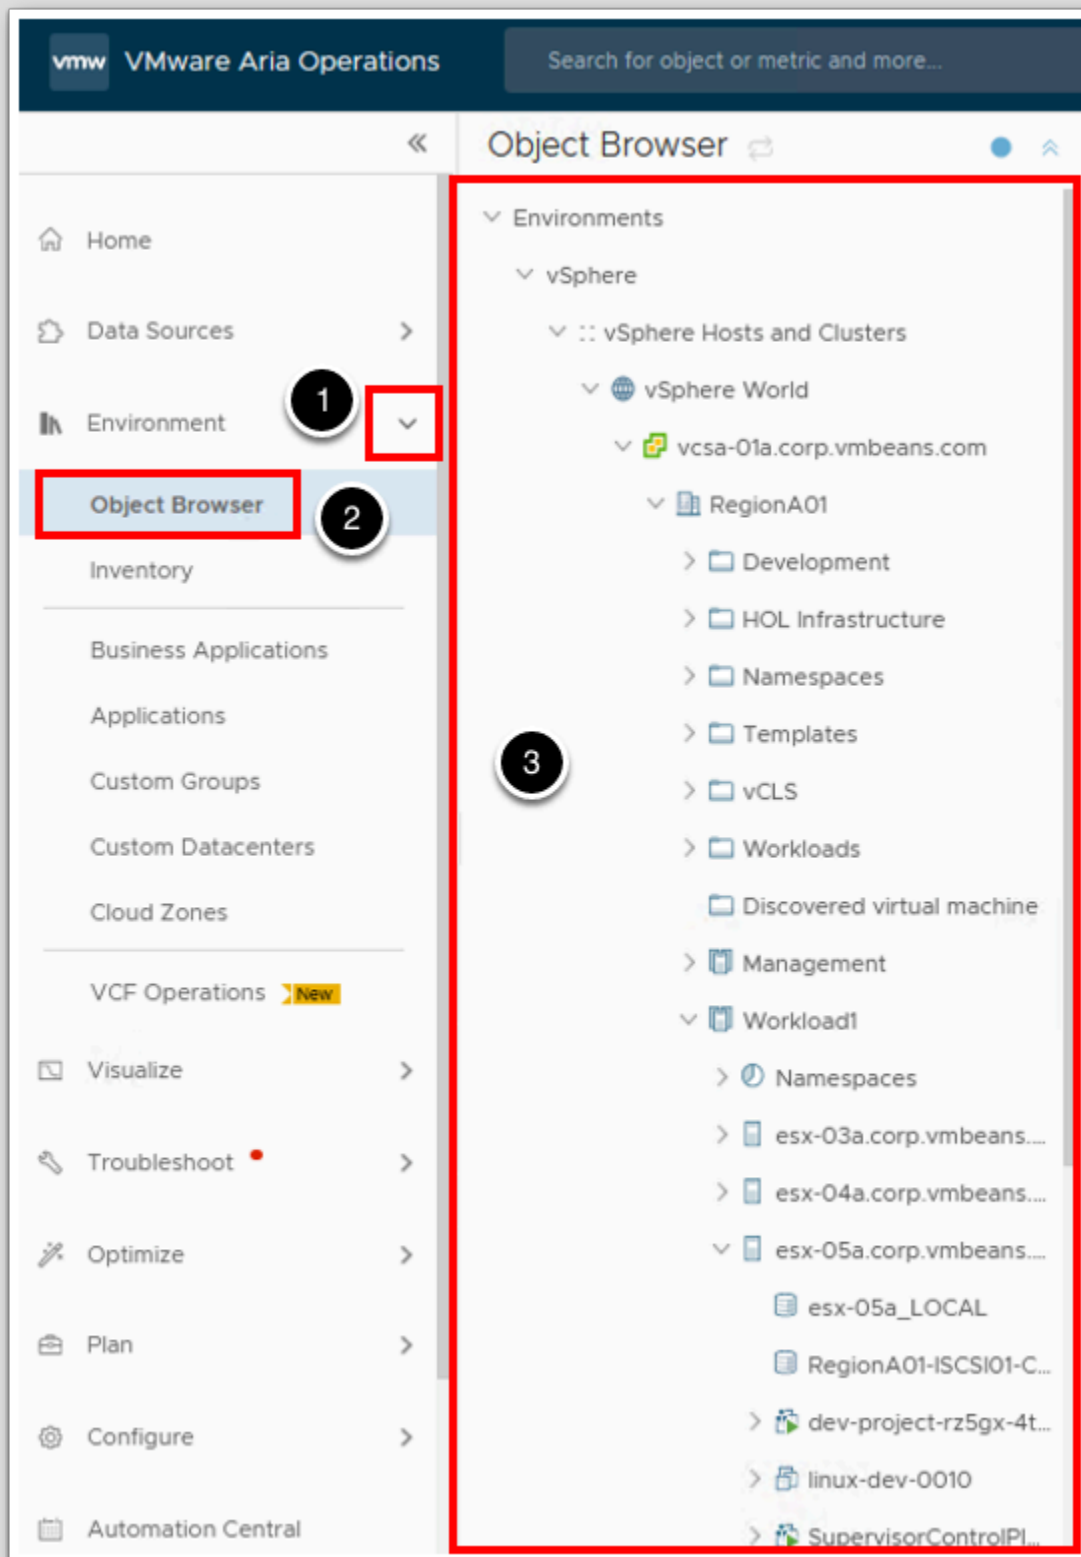


We will focus here on the vSphere Hosts and Clusters hierarchy because that's the one we will be using for the examples in this lab module. The hierarchy is shown in the graphic. There would also be other object types in the hierarchy if they existed in our lab vCenter server (for example, resource pools).

For this hierarchy you can see that virtual machines are two levels below clusters. And that vSphere hosts are one or two levels above datastores (this dual relationship can be found in other places as well). In the super metric formulas, the relationship distance (number of hops) is represented by the **depth** parameter and we will use that parameter in some examples later in this module.

## Object Browser

[388]



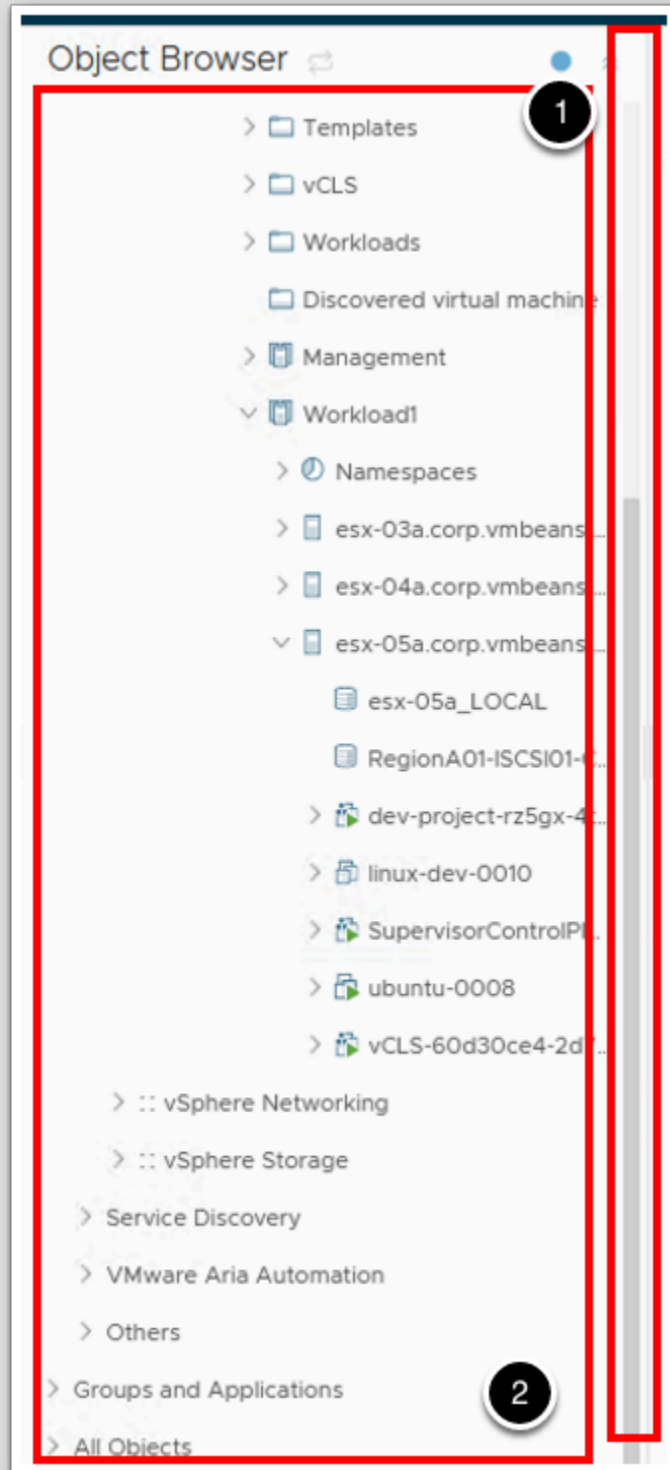
To see another way of looking at the vSphere Hosts and Clusters hierarchy within Aria Operations:

1. Expand **Environment** on the left menu bar.
2. Click **Object Browser**.
3. Expand the vSphere Hosts and Clusters hierarchy by clicking the > arrows

The levels of indentation in this view indicate the relative depth of each object type.

## Object Hierarchies

[389]

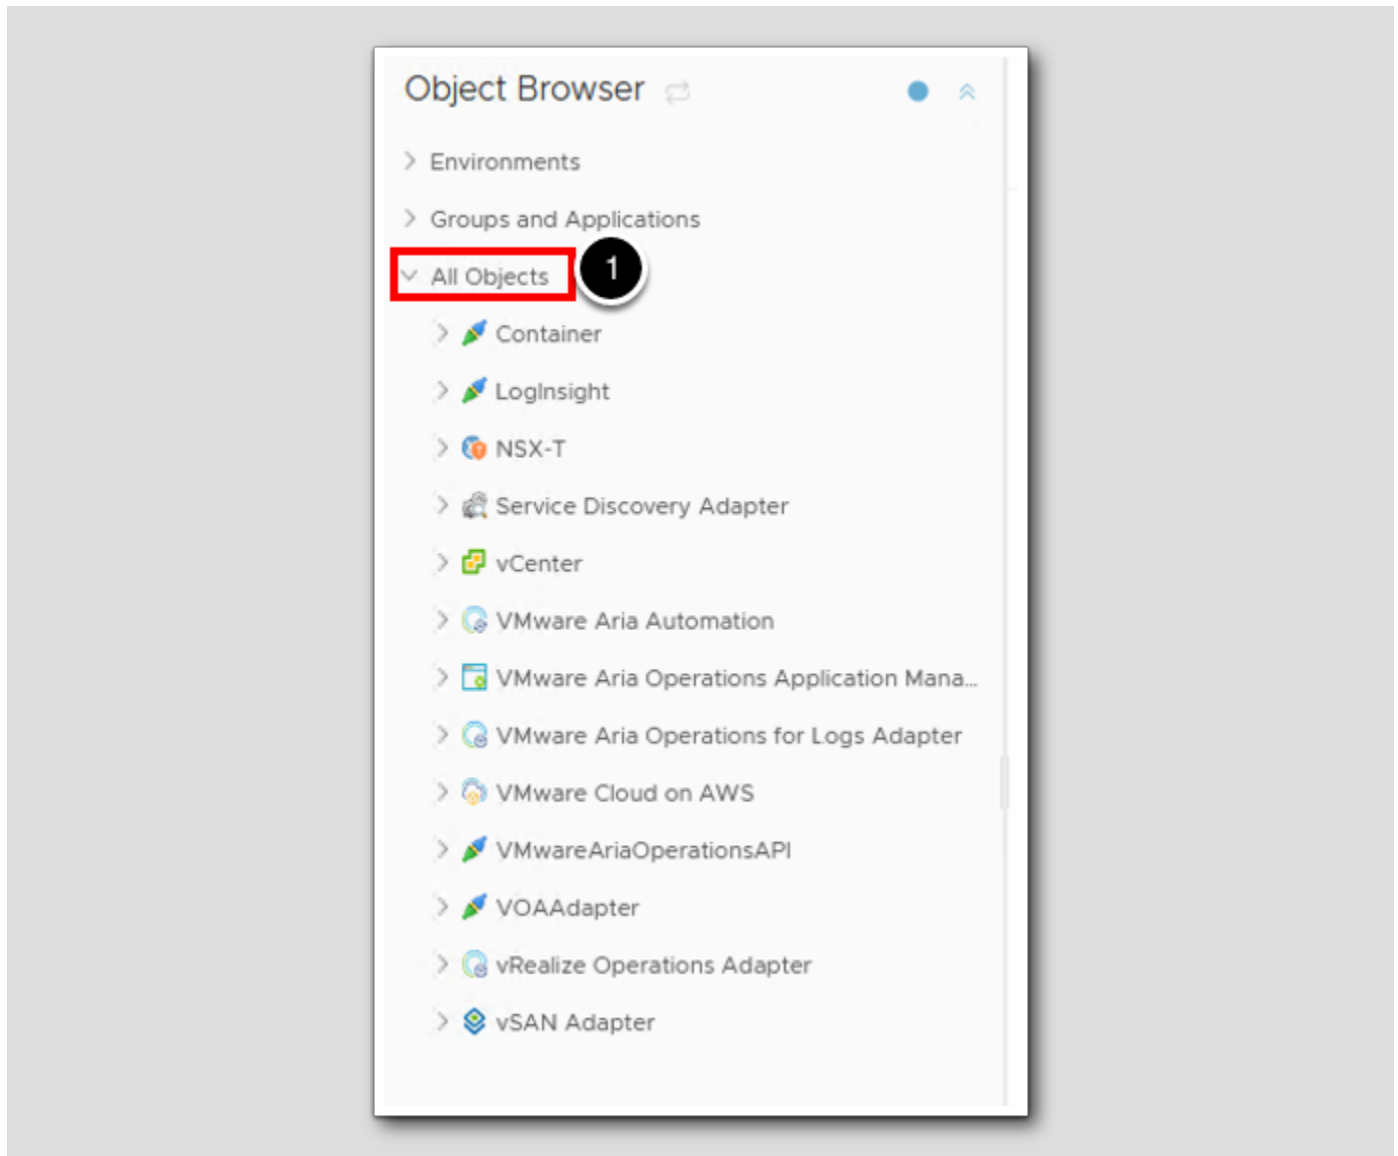


To see the available hierarchies within Aria Operations:

1. Scroll **down** in the Object Browser.
2. Note all of the available hierarchies in this Aria Operations instance.

## All Objects

[390]



1. Expand All Objects.

As stated earlier, if additional management packs were installed for extensibility (for example, NetApp or Dell EMC storage) hierarchies for those objects would also be here.

## Lesson End

[391]

You should have an understanding of vSphere hierarchy and structure of objects.

## Create Your First Super Metric

[392]

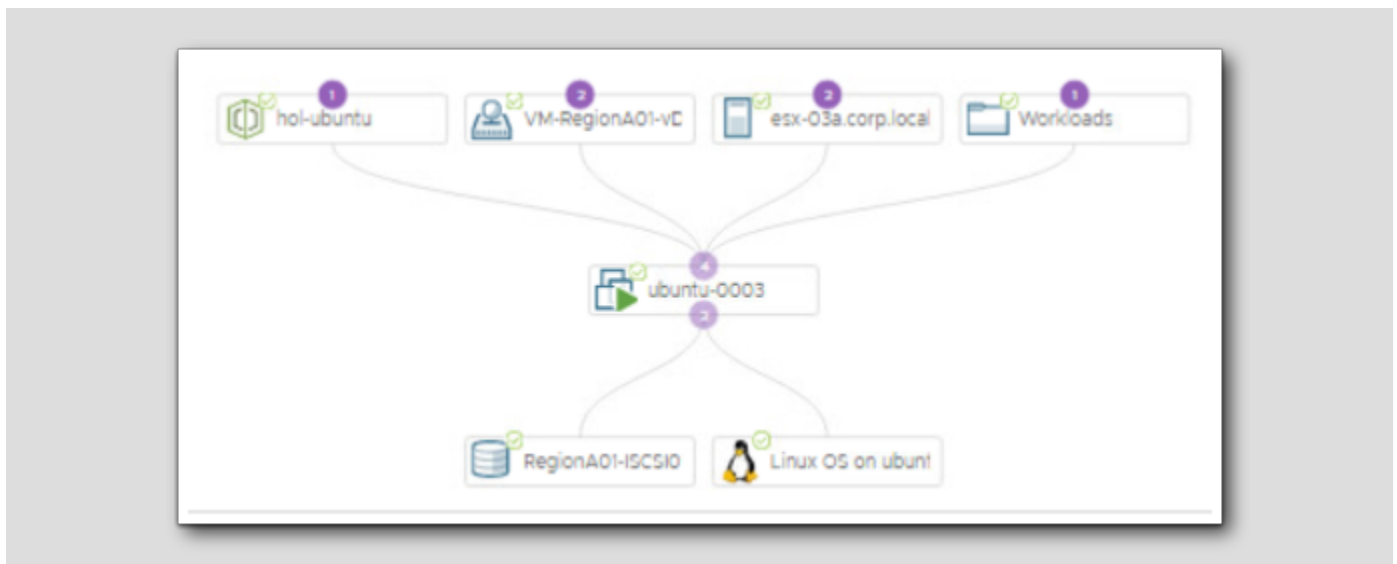
When you think "Super Metric", think metric math that will use some kind of a formula to get the desired outcome. In this first example, we will create a simple super metric and explore the depth parameter in a super metric formula.

Your first assignment is to create a super metric that will calculate the average memory utilization across all virtual machines running on a vSphere host or in a vSphere cluster. This is an example of creating a metric on an object (host or cluster) that is based on metrics from related objects (virtual machines).

## The Hierarchical Relationship

[393]

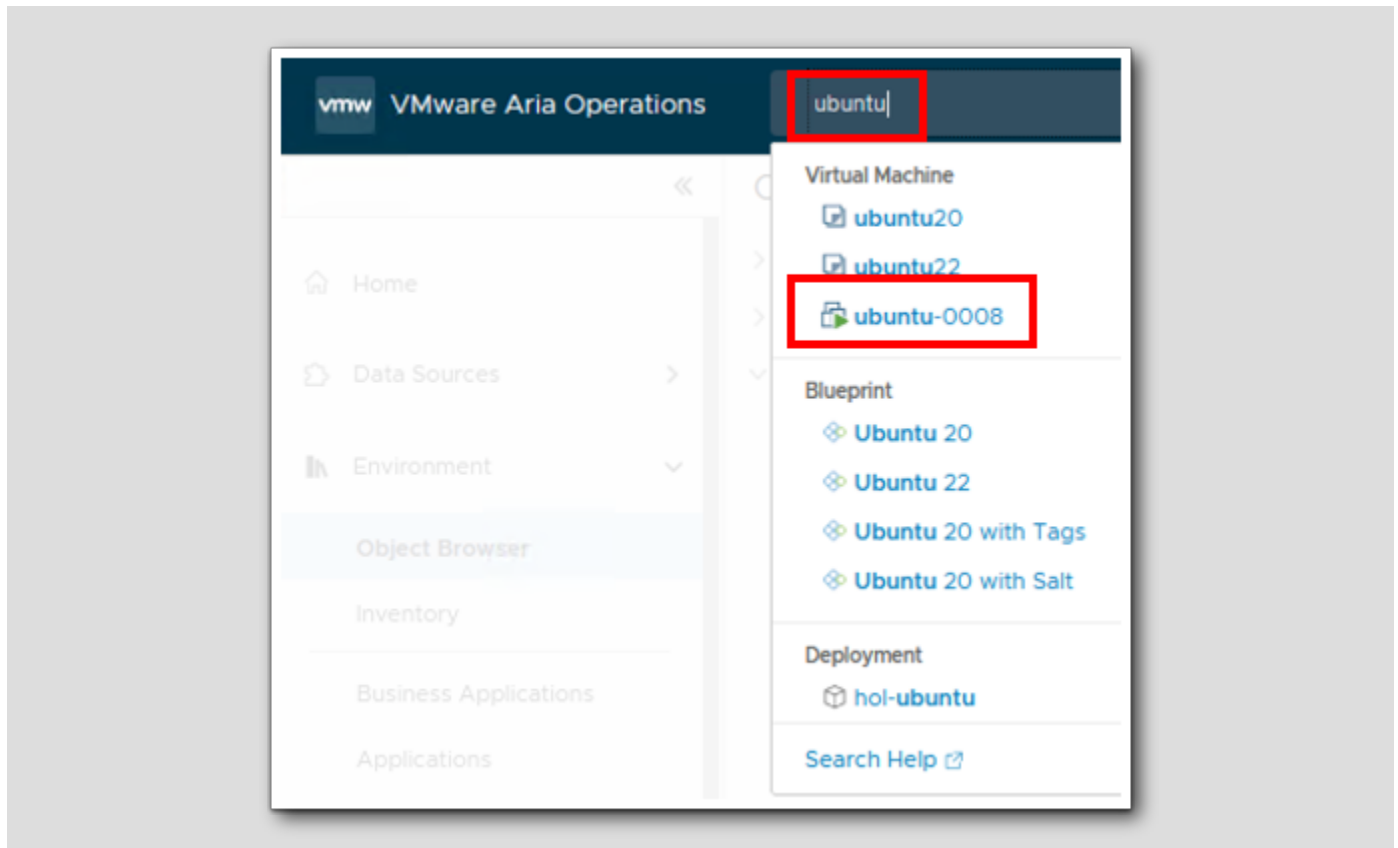
The below screenshot is not applicable to the current environment you are using but is being used because it is a good illustration of an object's hierarchy (parent and child objects).



If you recall from a previous lesson, we learned that virtual machines are children of hosts and "grandchildren" of clusters in the vSphere Hosts and Clusters hierarchy. So if we create a super metric on the cluster object type and on the host object type and have it look one or two levels down the hierarchy to create the sum of the metric representing memory usage on virtual machines, we will have completed the assignment for this lesson.



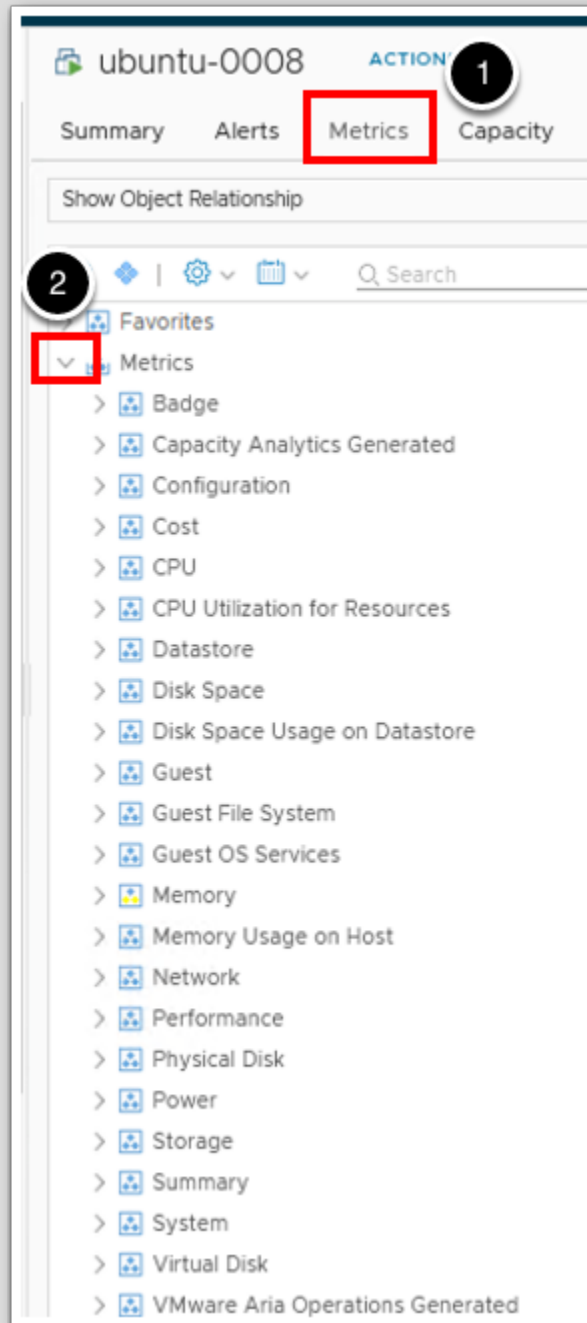
## Which Metric Will We Be Using?



Before we get started with the super metric, let's understand which virtual machine metric we will be using for this lesson. Since we want to average a vm metric (memory utilization), let's go find a vm to see what metrics are available. We will take a look at the ubuntu-0008 virtual machine.

1. In the search box, type **ubuntu**
2. Click the **ubuntu-0008** link under the Virtual Machine object type

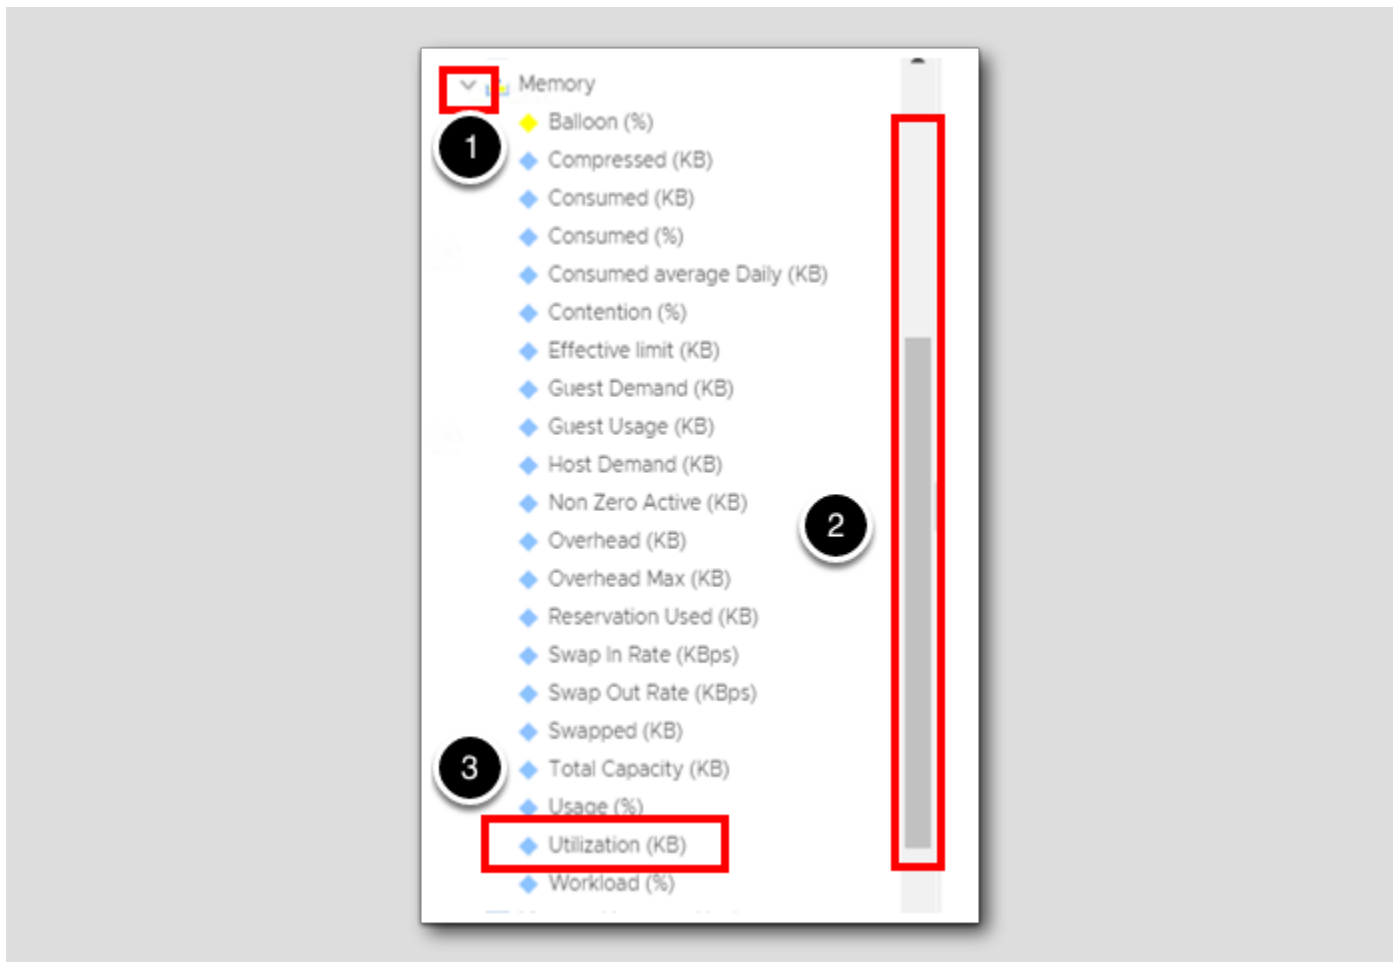
### Expand the All Metrics Tree



1. On the ubuntu-0008 object page, click the **Metrics** tab.
2. Click the > to expand the **Metrics** section.

## Expand the Memory Metric

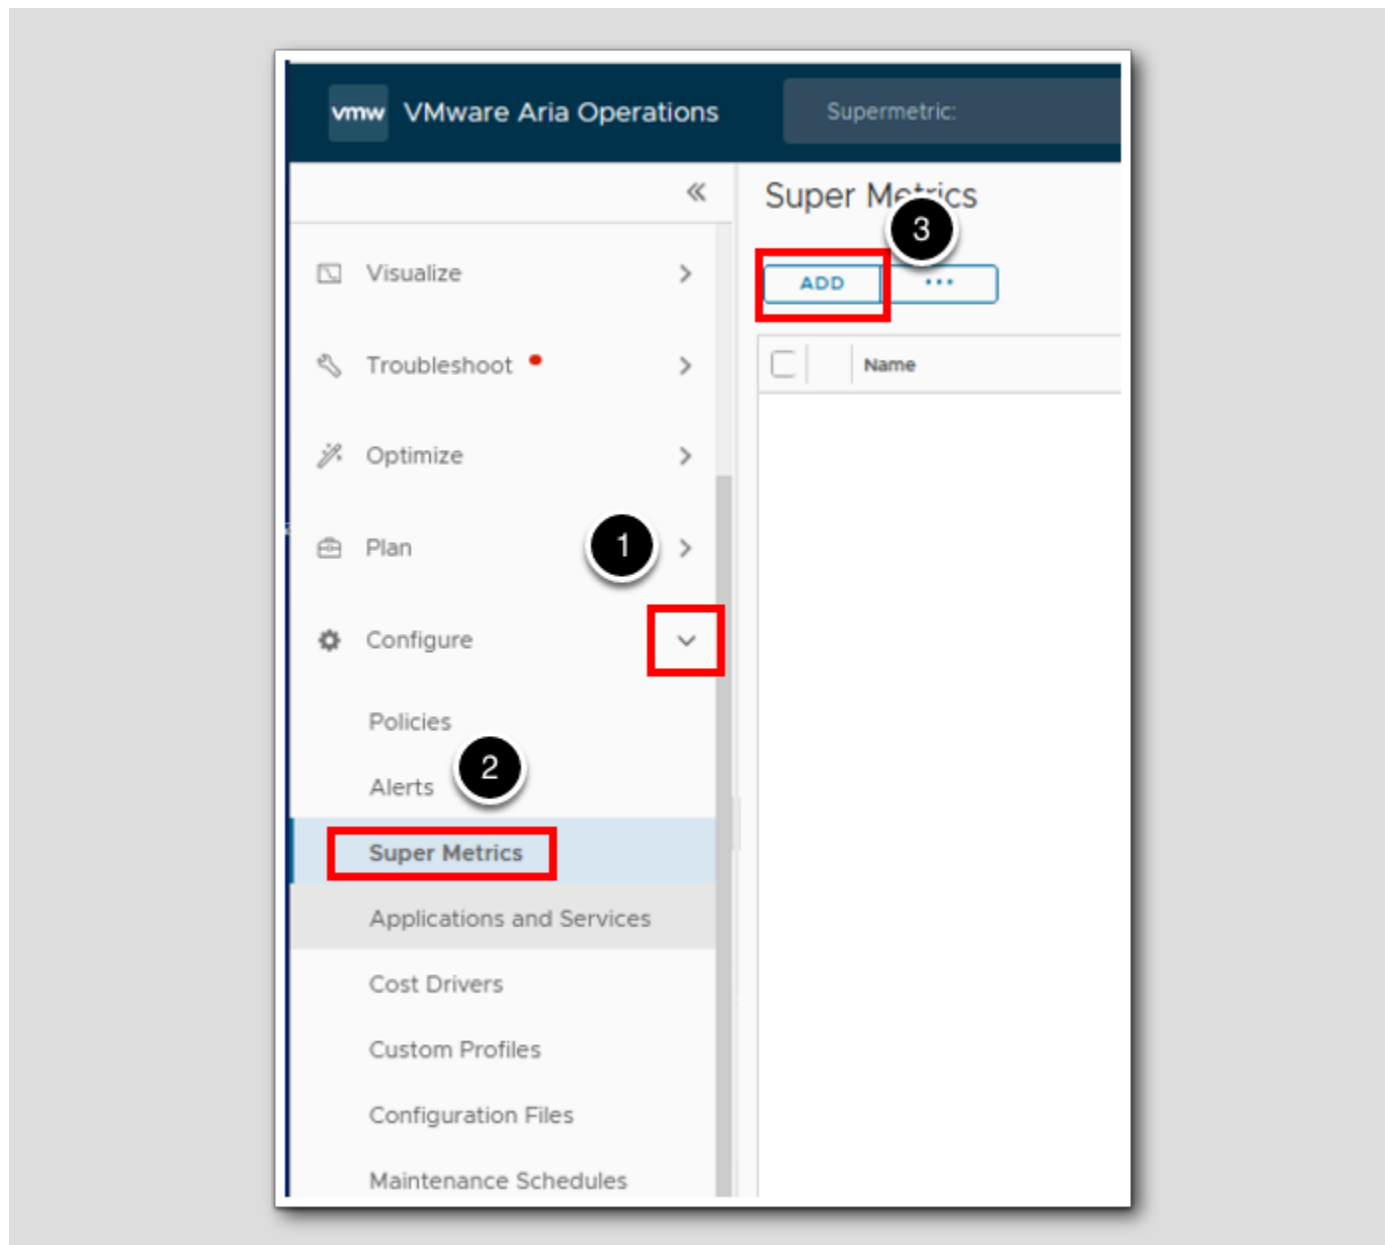
[396]



1. Click the > to expand the **Memory** section.
2. Scroll down to see the list of available memory metrics.
3. Note the **Utilization (KB)** metric - this is the metric we will be using to create our super metric.

## Create The Super Metric

[397]



Now that we know which virtual machine metric we will be using, let's navigate to the new super metric editor window. The new super metric workspace can be found in the Configure section of Aria Operations.

1. Expand the Configure.
2. Click Super Metrics.
3. Click ADD to create a new super metric

Super Metric tab

[398]

**Create Super Metric**

1 - Super Metric 2 - Object Types

**Name:** Average Mem Usage Across All Descendant VMs (GB)

**Description (Optional):** Calculates the average memory usage of all VMs that are descendants of this object]

PREVIOUS **NEXT** CREATE CANCEL

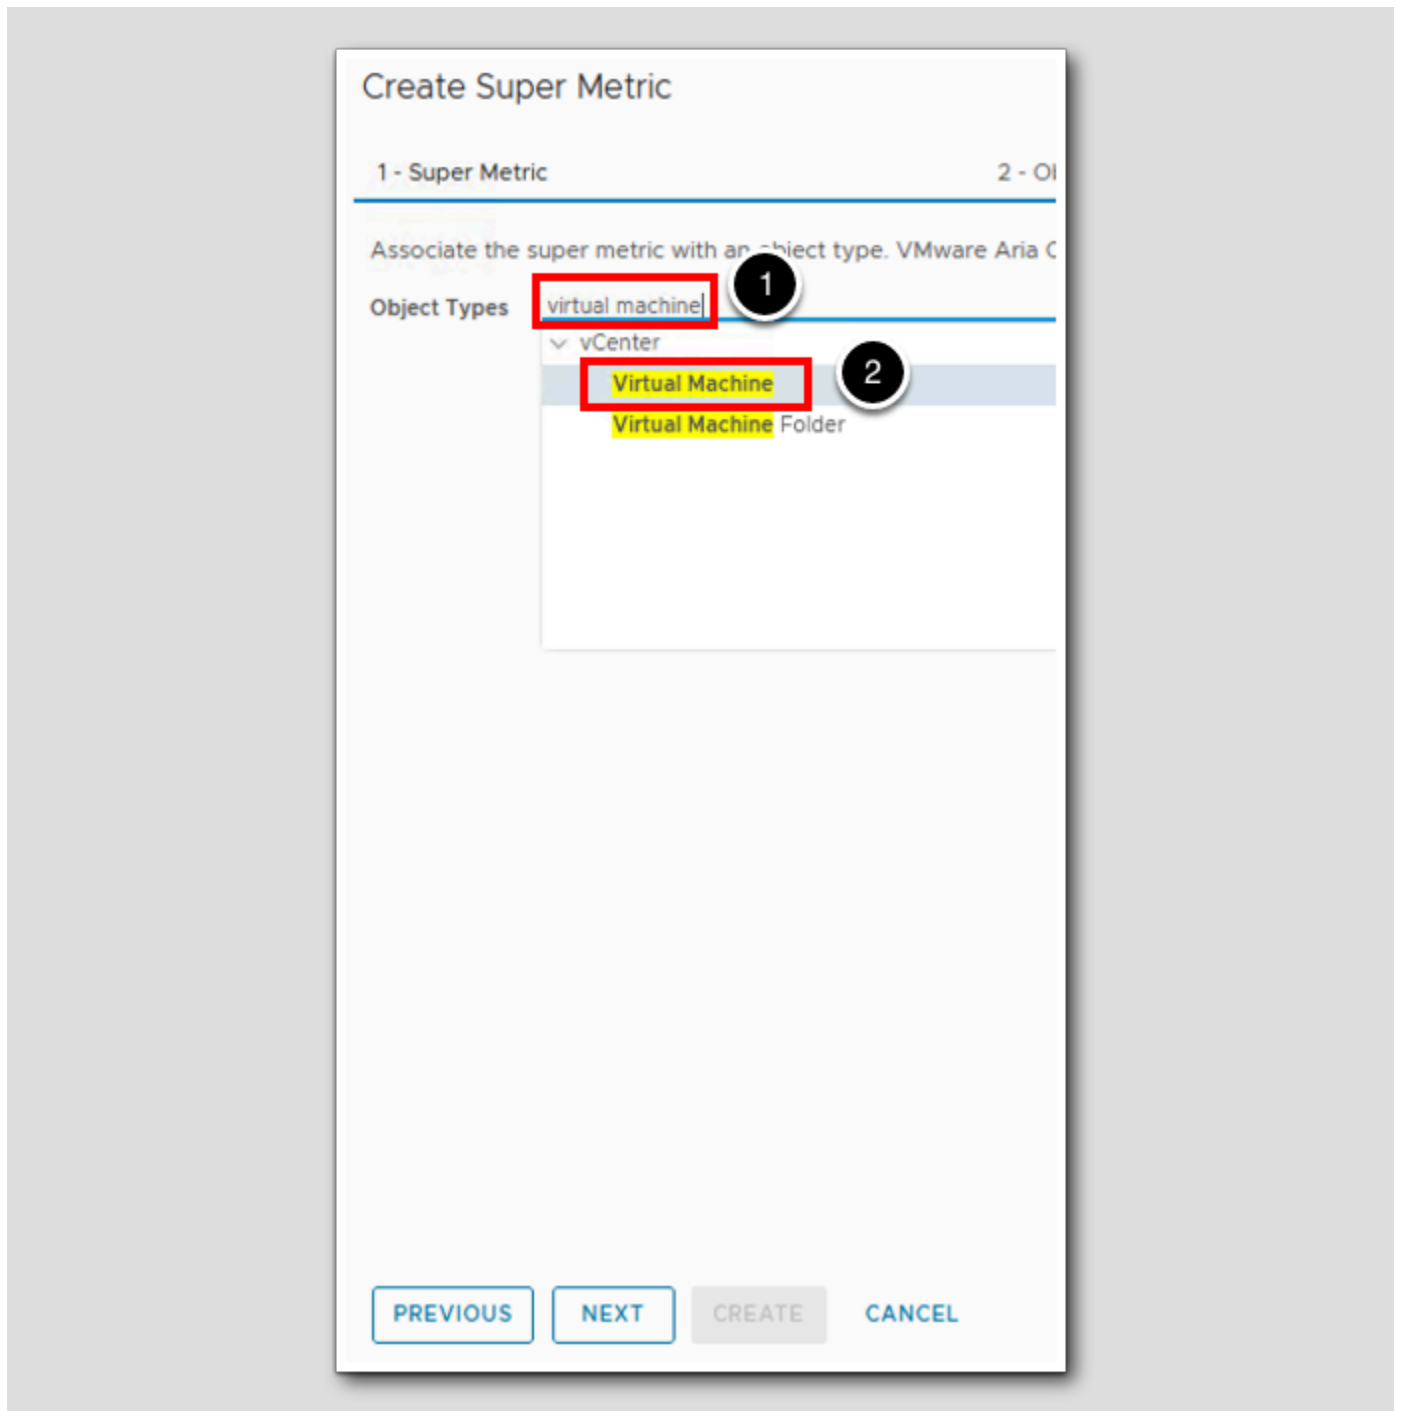
Let's enter some basic information about the super metric. You want to create a name that is descriptive enough so you or others will understand what it is calculating when you use it later in dashboards or reports or alert definitions. It is also a good idea to include the unit of measure in the metric name - in this case we will calculate the value in gigabytes (GB).

1. In the Name field, type **Average Mem Usage Across All Descendant VMs (GB)**
2. In the Description field, type **Calculates the average memory usage of all VMs that are descendants of this object.**
3. Click the **NEXT** button.

## Object Types

[399]

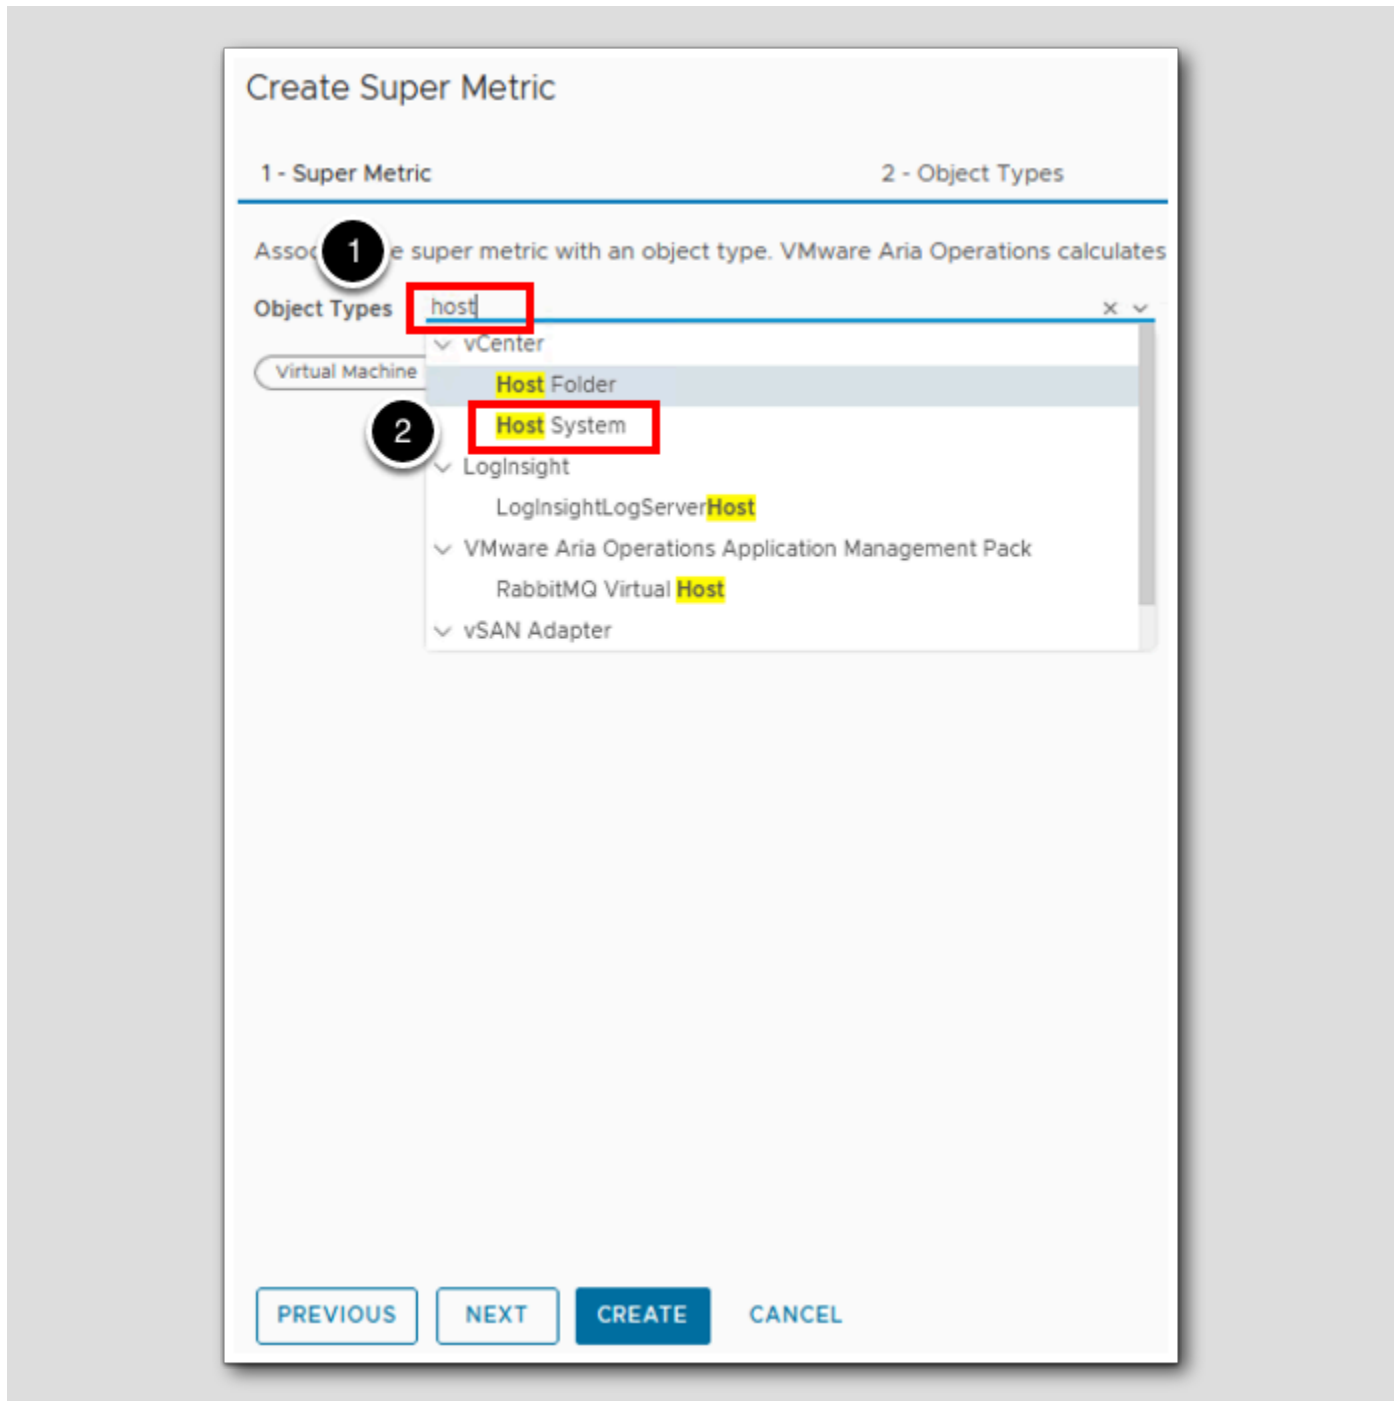
Since we will be getting the average VM memory usage across multiple objects we will add 3 Object Types.



1. In the Object Types search line, type Virtual Machine.
2. Single click on Virtual Machine to add Virtual Machine as an Object Type.

## Object Types (Continued)

[400]

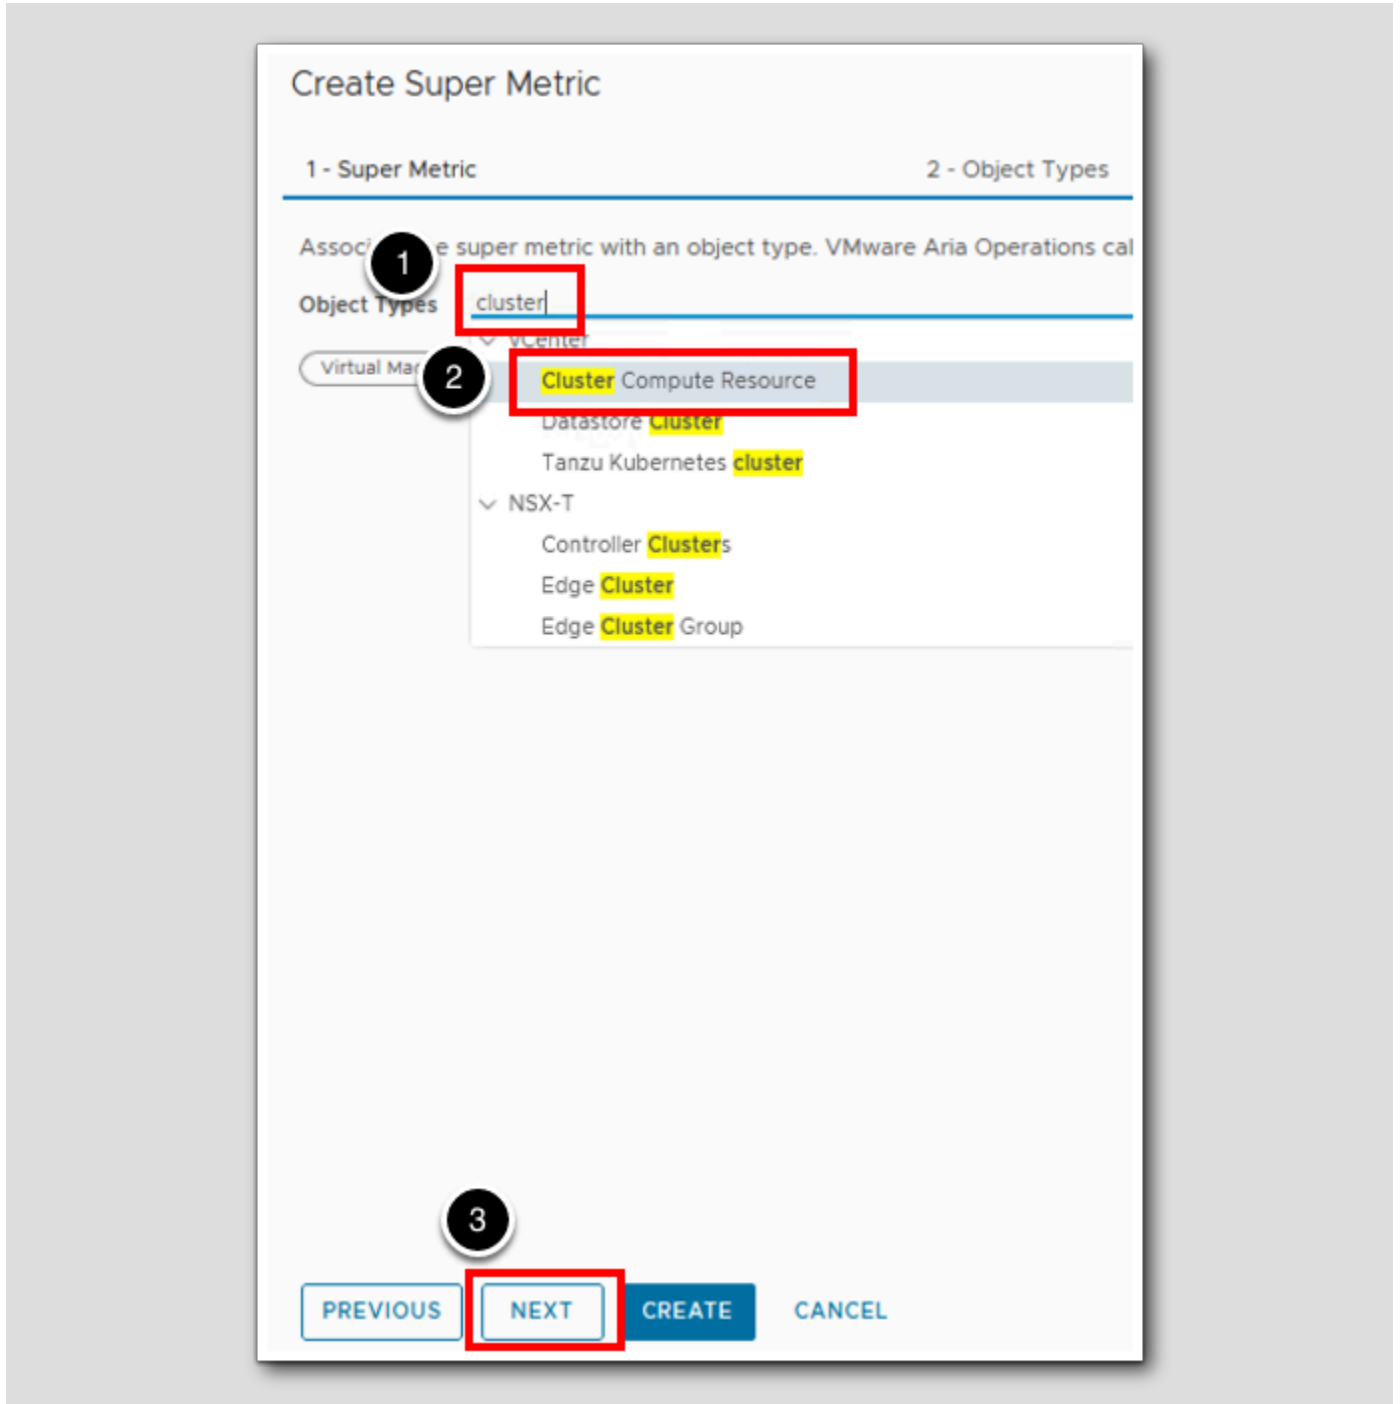


1. In the Object Types search line, type host.
2. Single click on Host System to add Hosts as an Object Type.



Object Types (Continued)

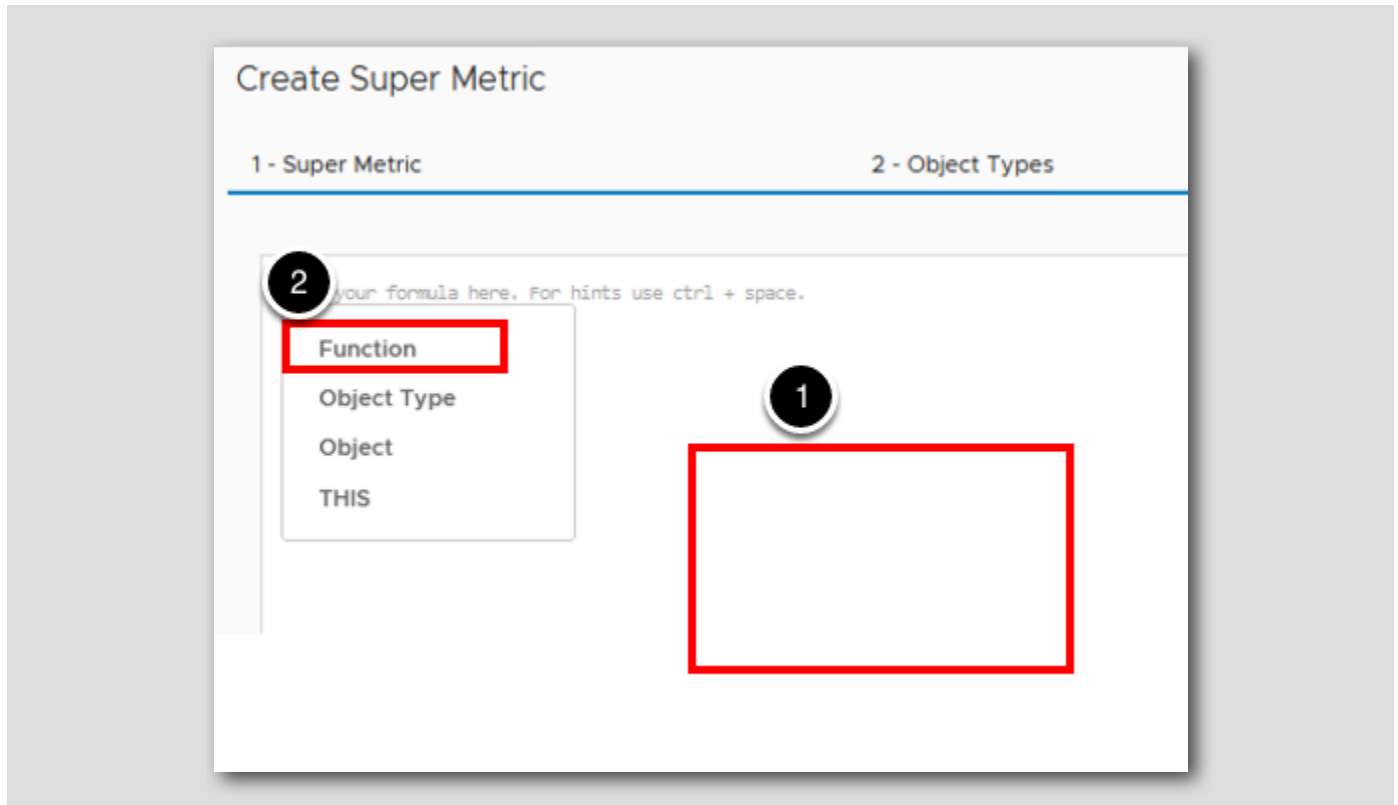
[401]



1. In the **Object Types** search line, type cluster.
2. Single click on **Cluster Compute Resource** to add Clusters as an Object Type.
3. Click **NEXT**.

## Formula Functions

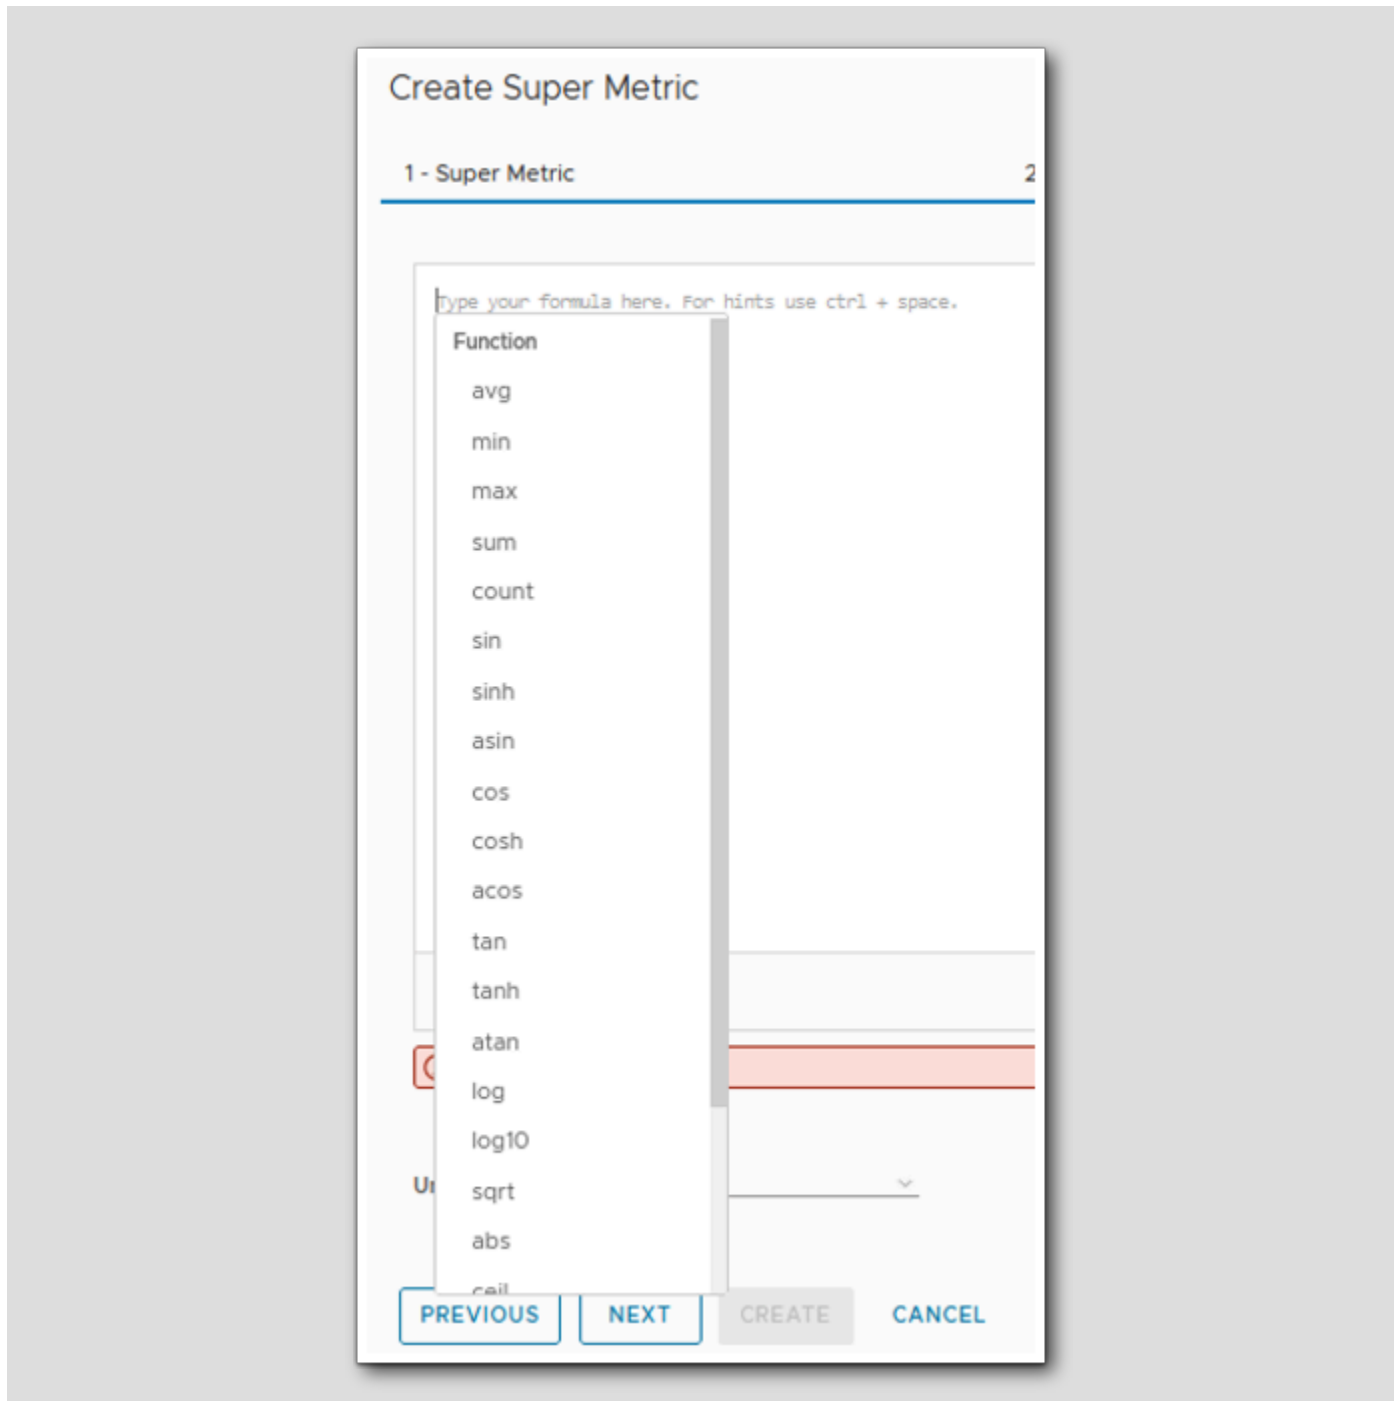
[402]



1. Click anywhere in the empty formula box.
2. Click the **Functions** drop-down to see a list of all available functions.

## Formula Functions (Continued)

[403]



The list includes looping functions (avg, combine, count, max, min and sum) that work on more than one input value and can return either a single value or a set of values depending on the formula syntax. The remainder of the functions are single functions. They work on only a single value or a single pair of values.

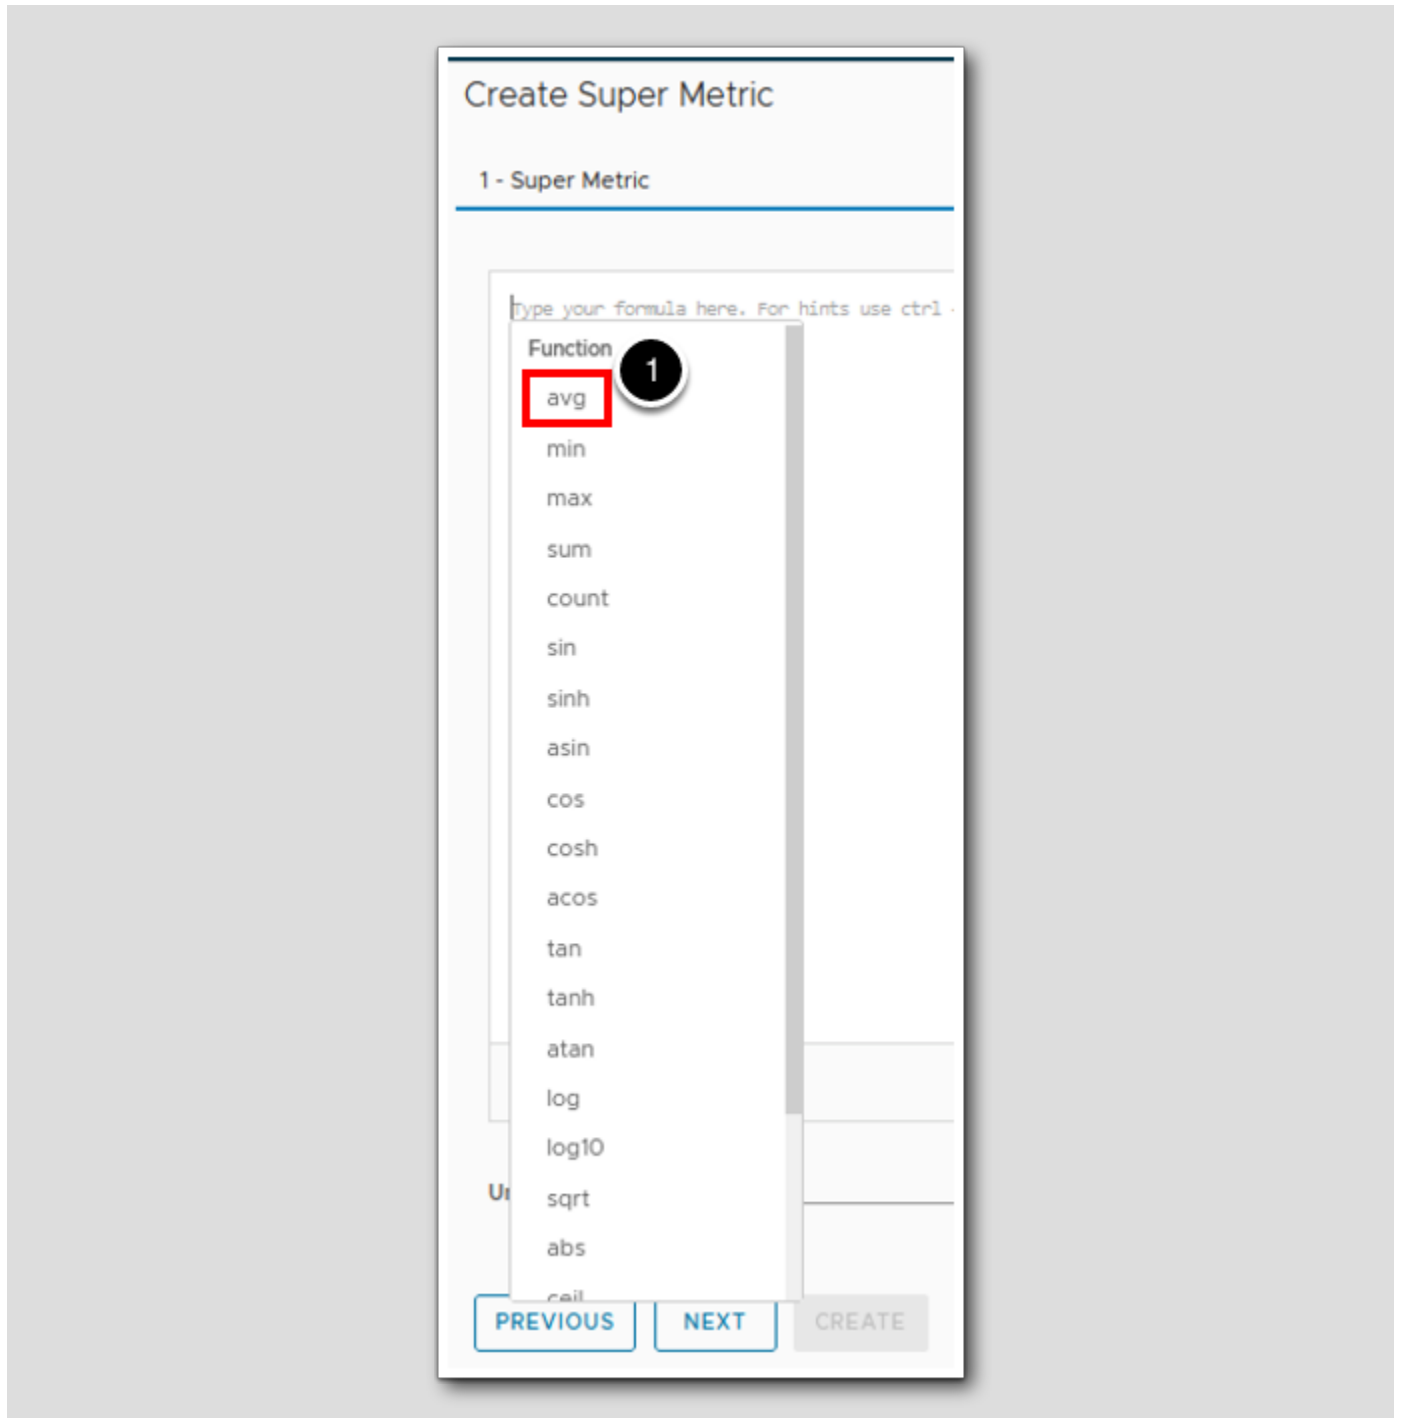
To better understand the concept of looping functions, think about the example metric we are going to create in this lesson. We want to look for all descendant virtual machines (could be one or could be many), get the value for memory usage for each of those virtual machines, and then calculate an average of those values which we will then store a single super metric on our object (in this case a vSphere host or cluster). In the process, we will use a looping function to "loop through" all of the descendant virtual machines to get the memory usage metric value for each one.

Note: The product documentation for super metric functions and operators can be found [here](#).

## Create a formula

[404]

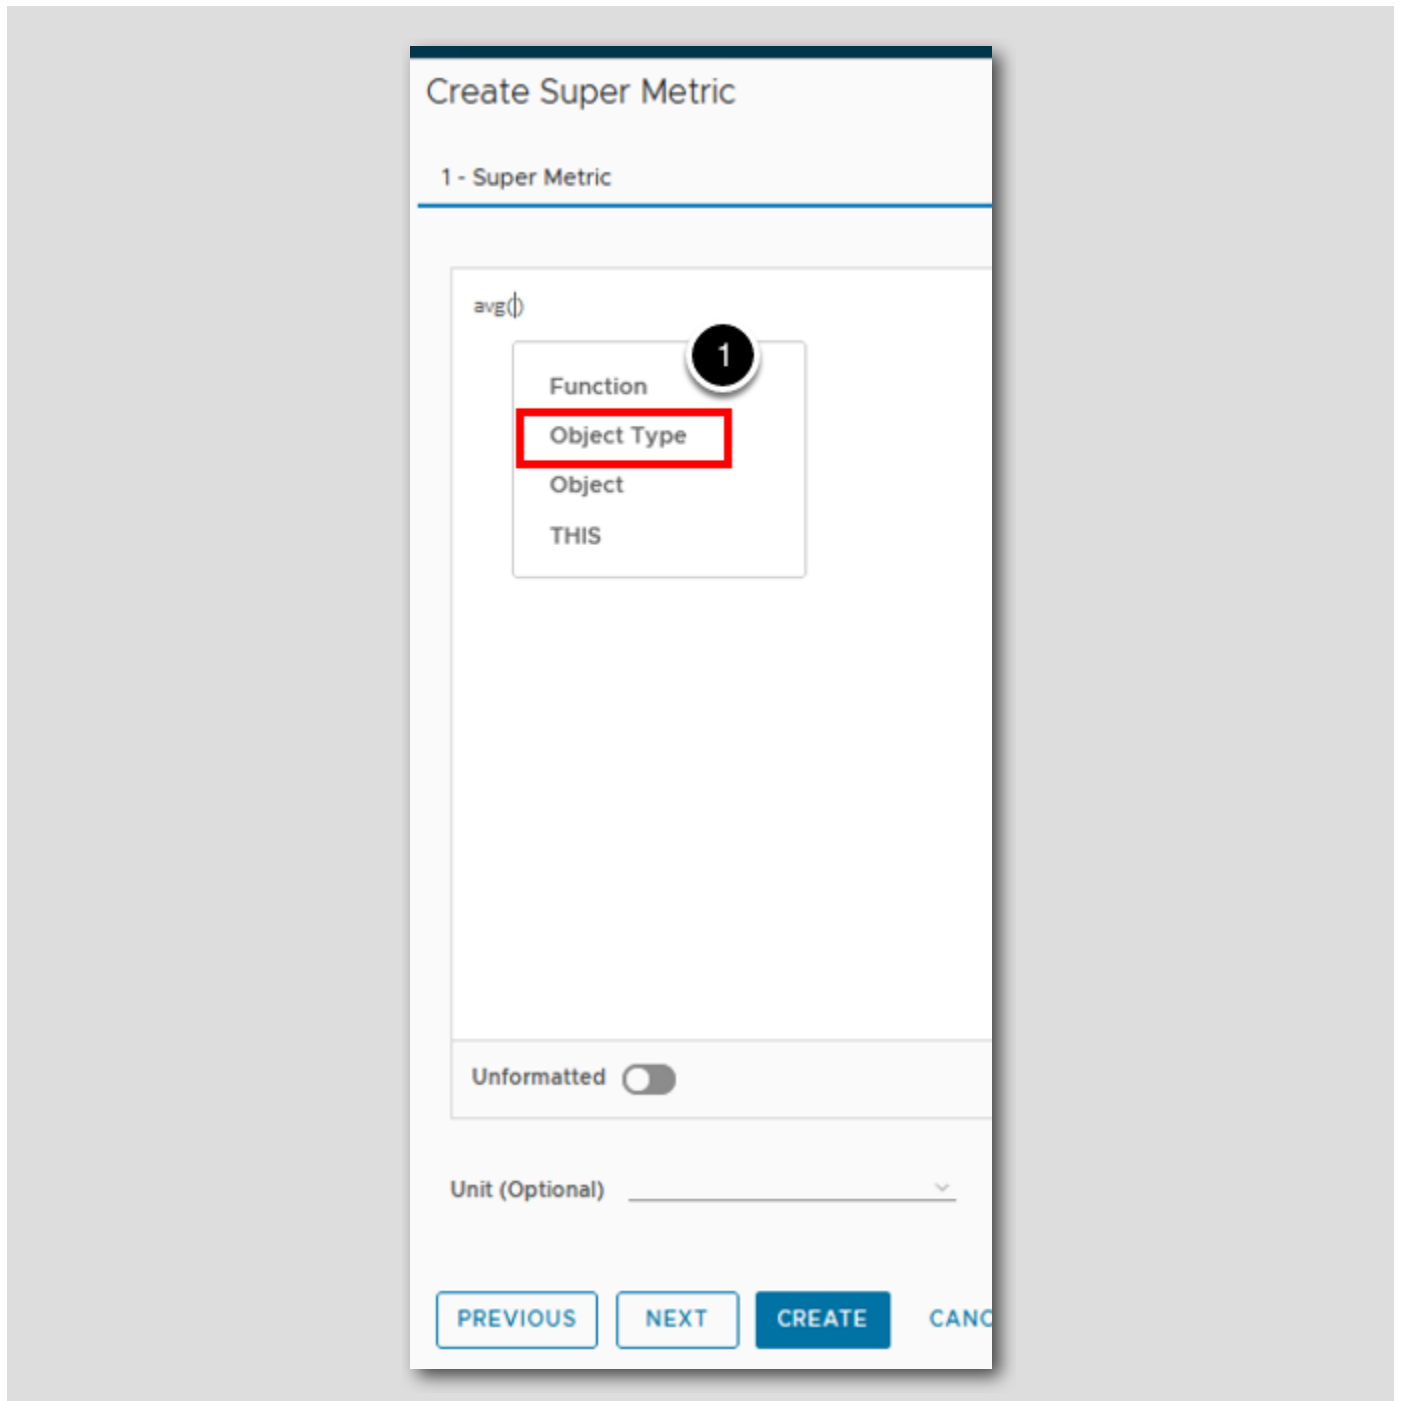
Recall that we want to create an average of the memory usage across all virtual machines on our host or in our datacenter so let's start by adding the **avg** function to our formula.



1. Click on avg in the Function list.

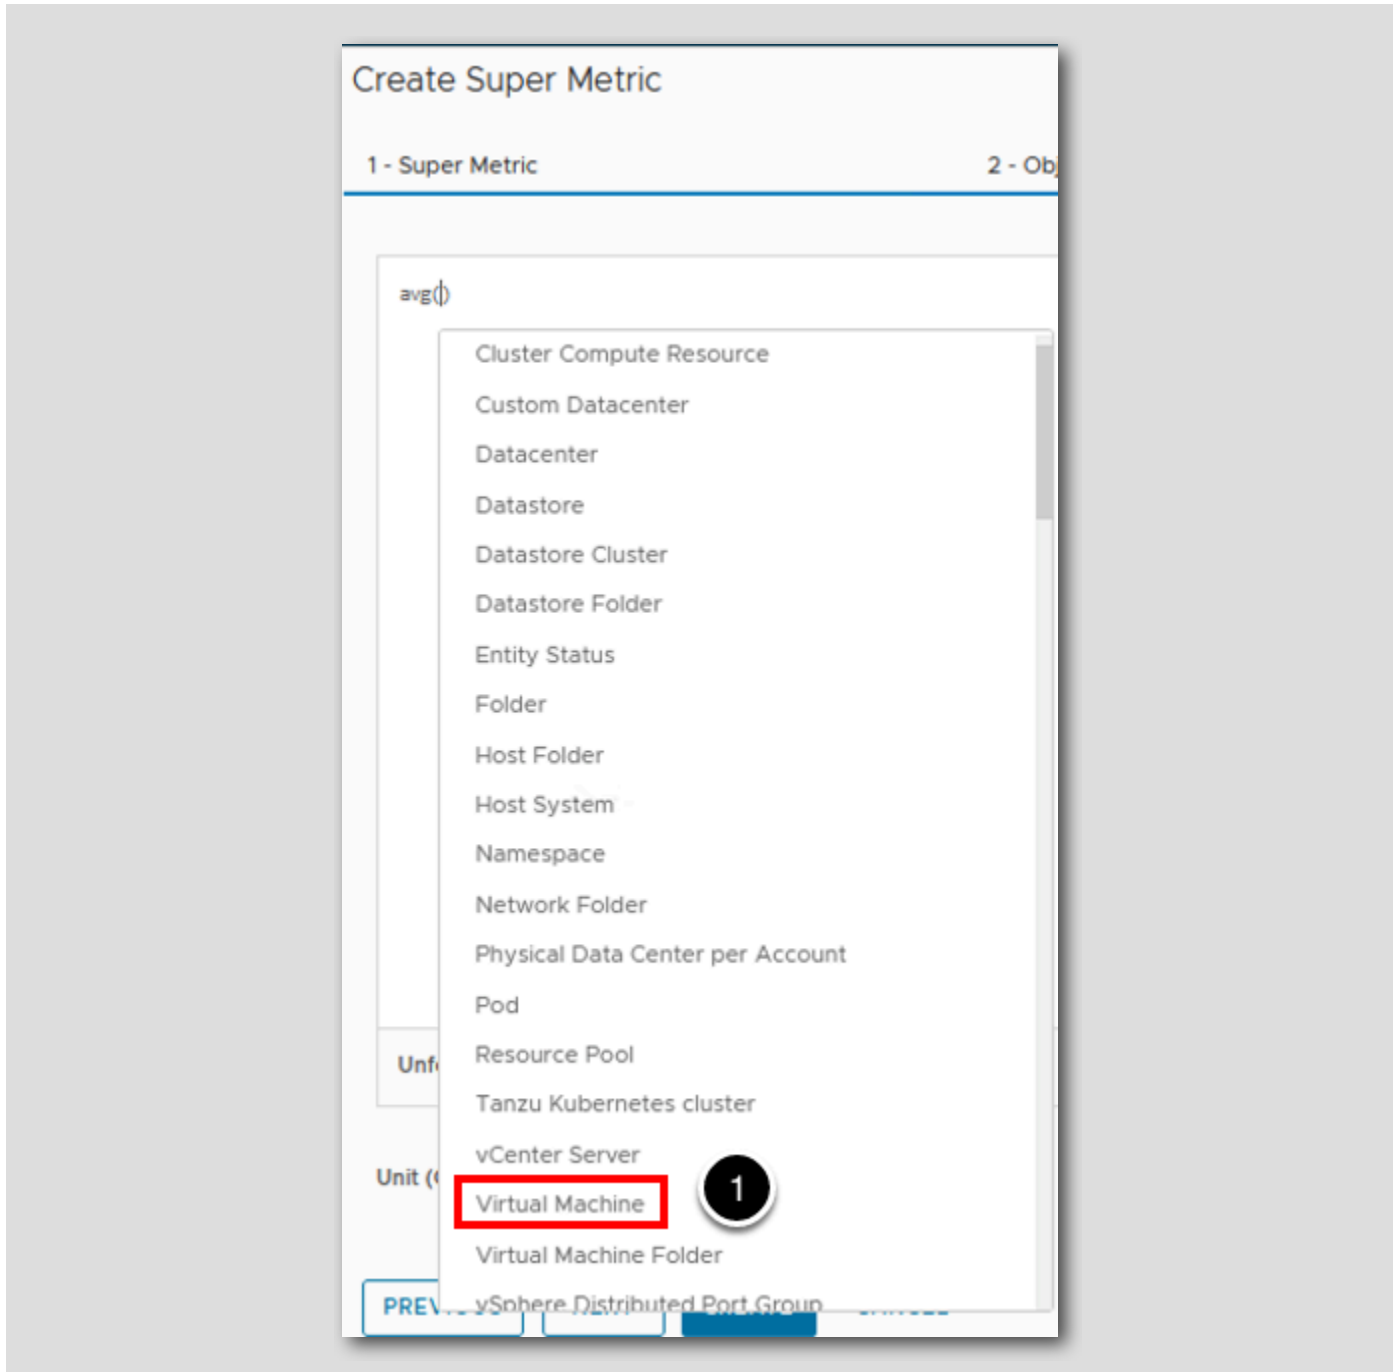
## Start Creating the Super Metric Formula

[405]



1. Click on **Object Type** to select what we will average on.

### Select the Virtual Machine Object Type

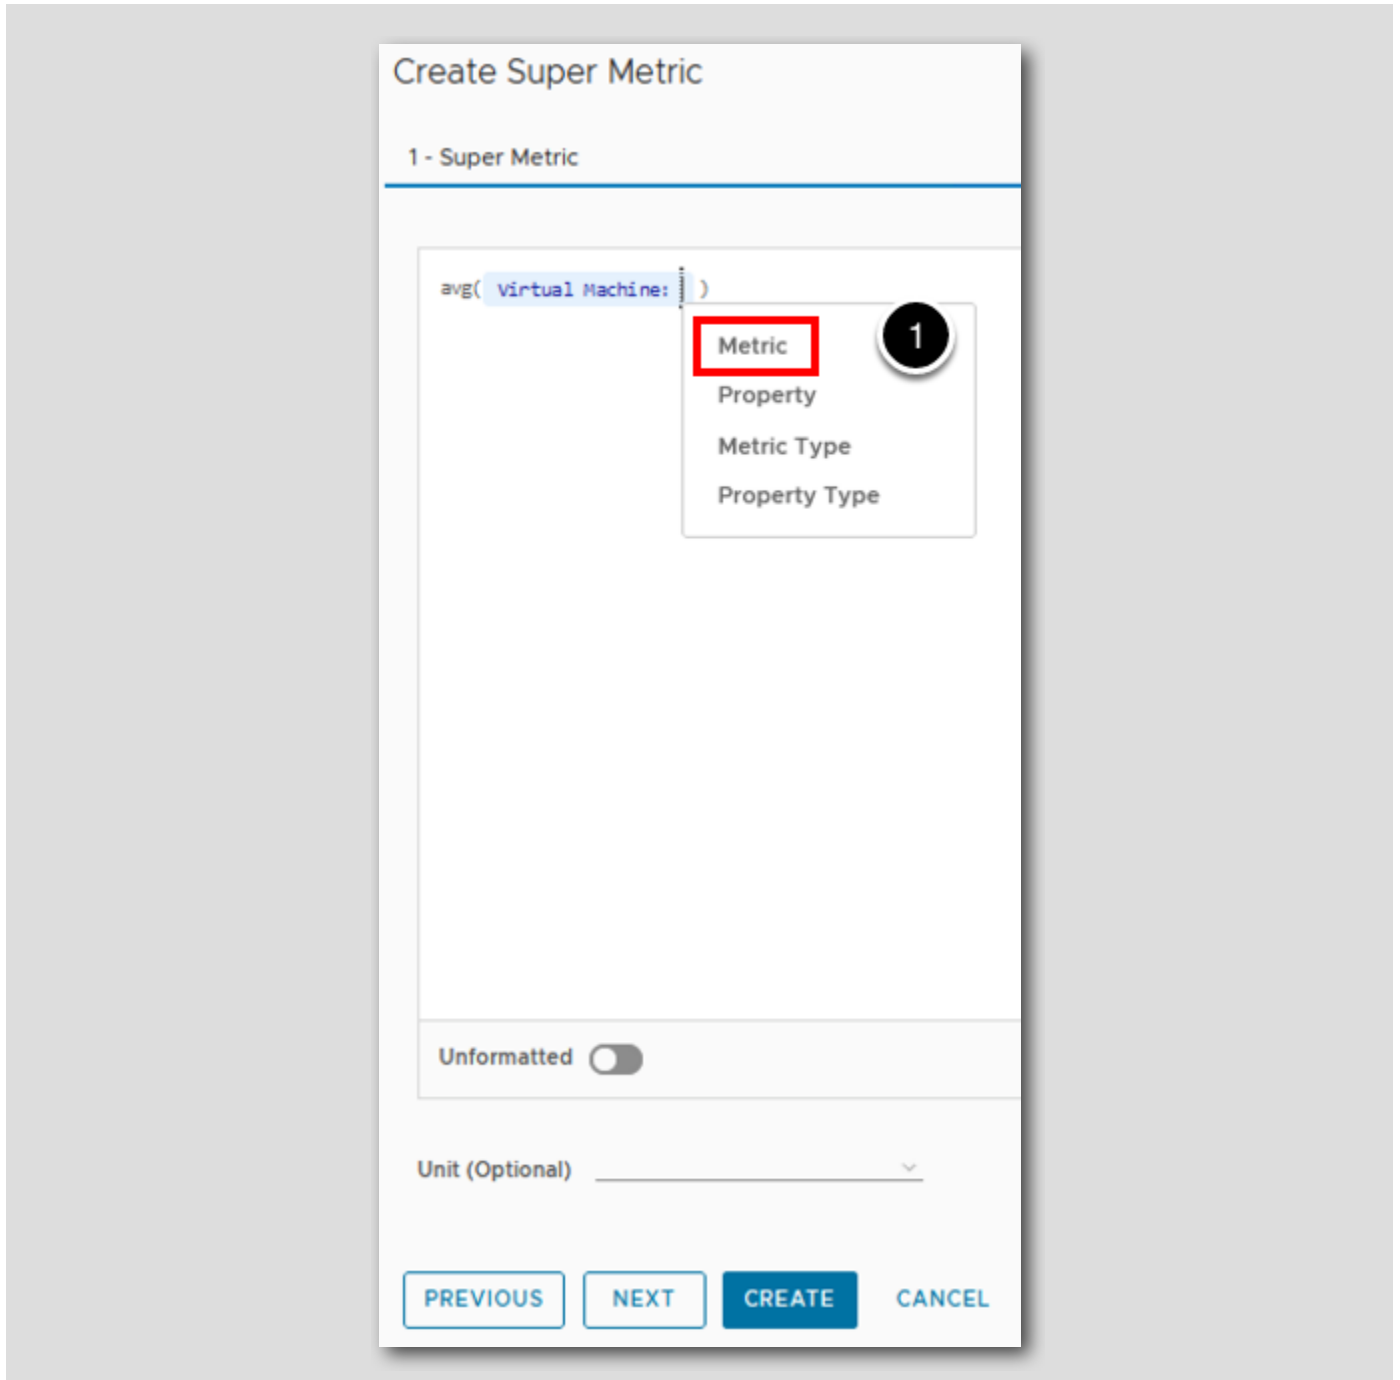


2. Select Virtual Machine as our Object Type.



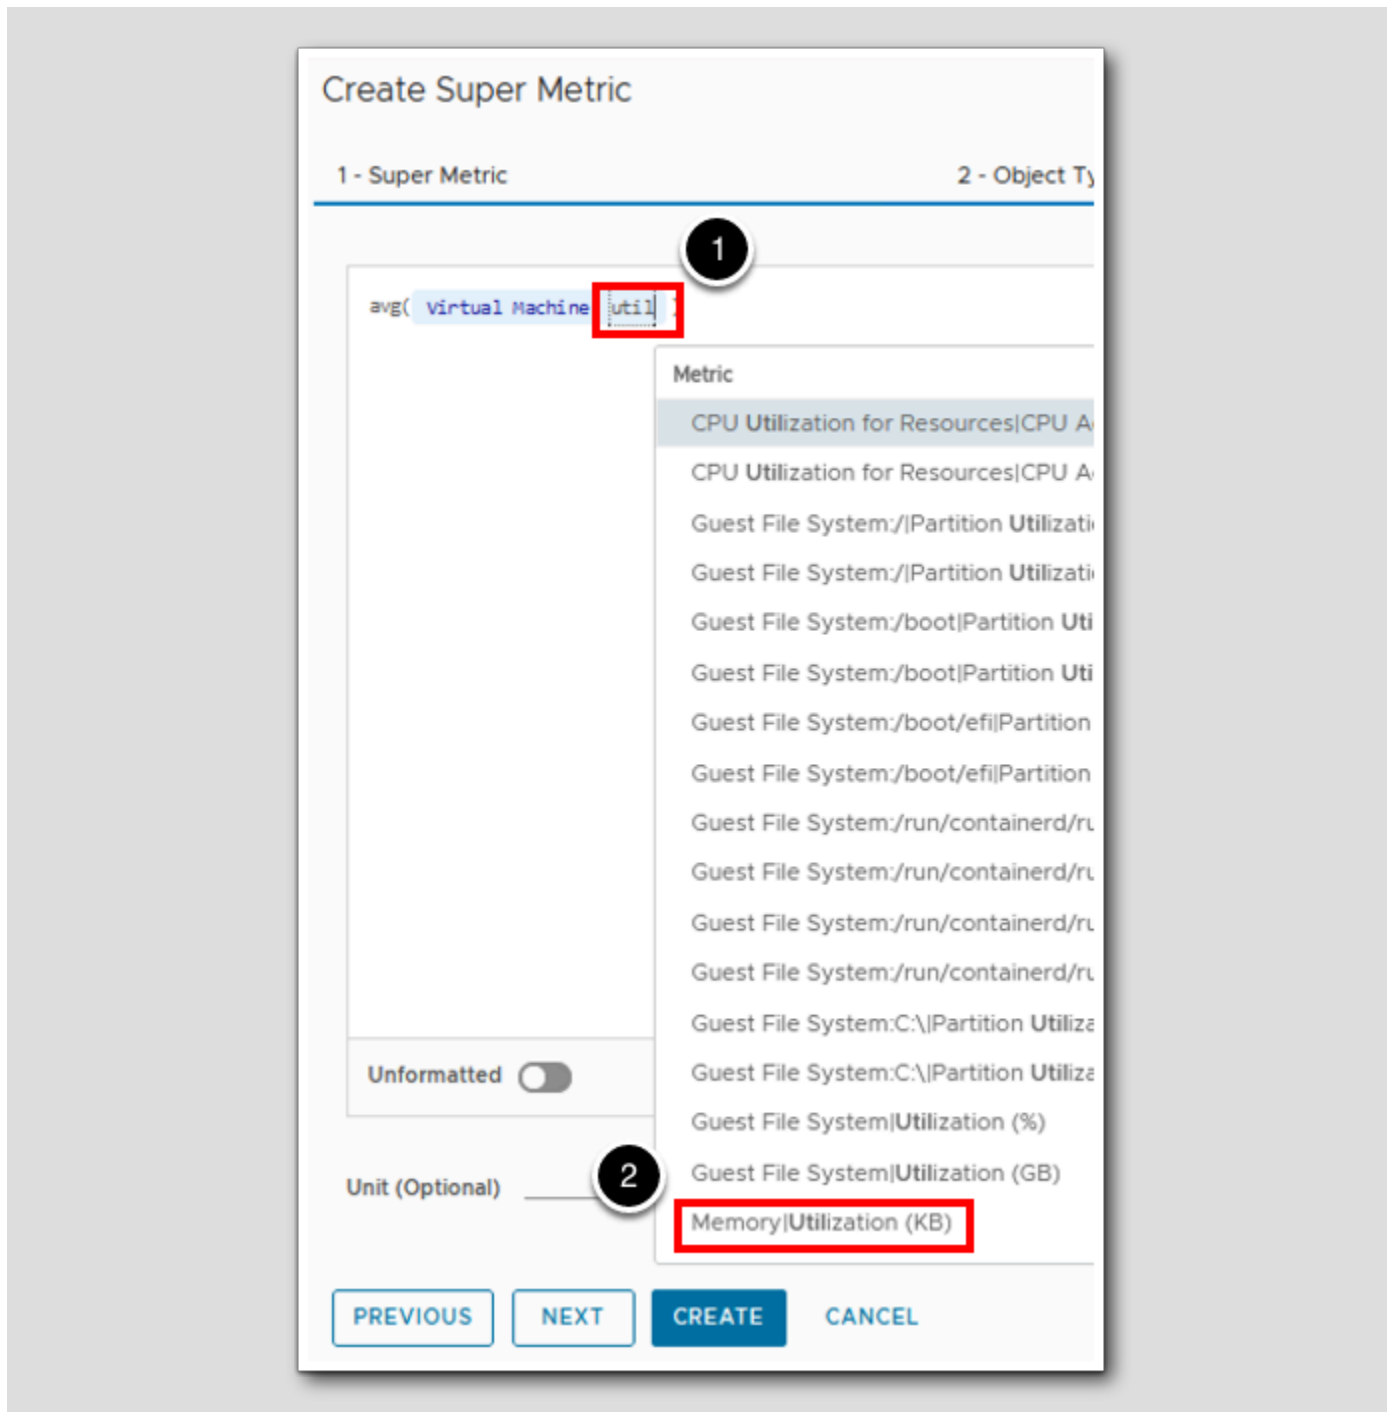
## Add the Attribute

[407]



1. Since we want to average a specific metric from all of the virtual machines, click **Metric**

Select the Metric



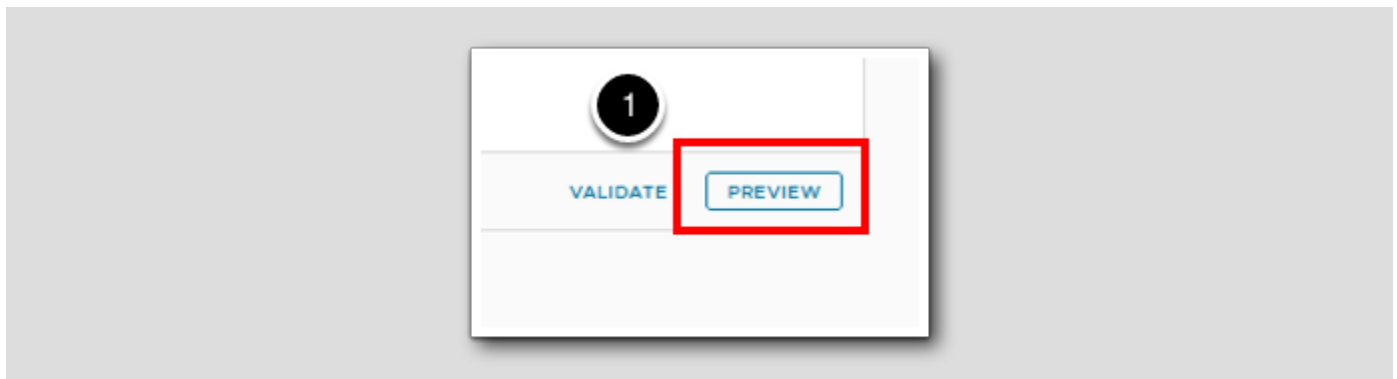
1. Start typing some text from the metric name if you know it. Type **util**
2. Click **Memory|Utilization (KB)** (note the highlighted portion of the metric.) You should remember this metric from the beginning of the lesson.

Note that the the units of memory utilization are in KB but we want our super metric value to be in GB. That's OK because we can just add the additional math to the formula to do the conversion from KB to GB.

## Preview

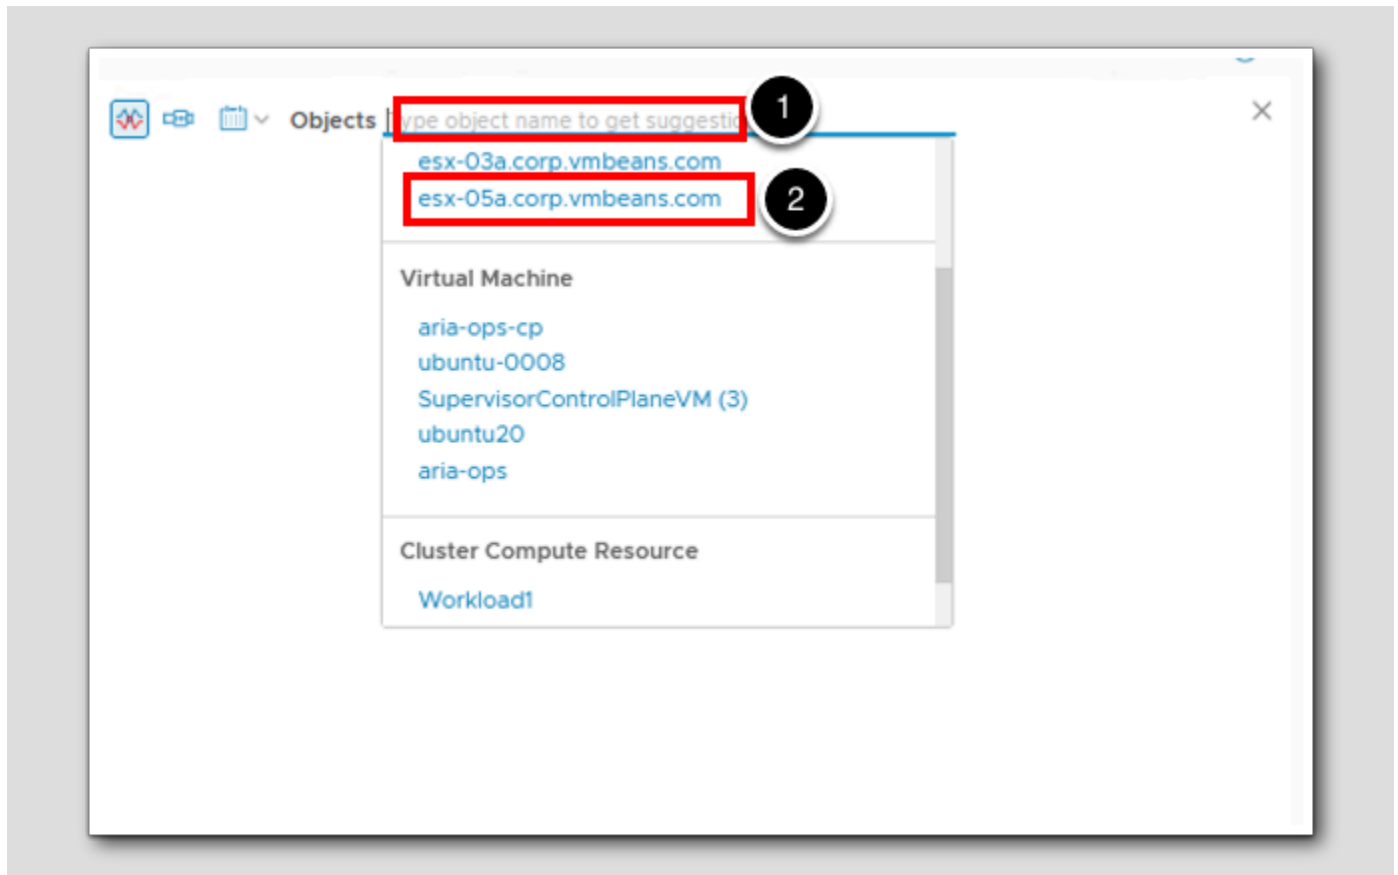
[409]

Lets take a look at our formula now that it's got something to calculate



1. Click PREVIEW in the lower right hand corner of the screen.

## Object List

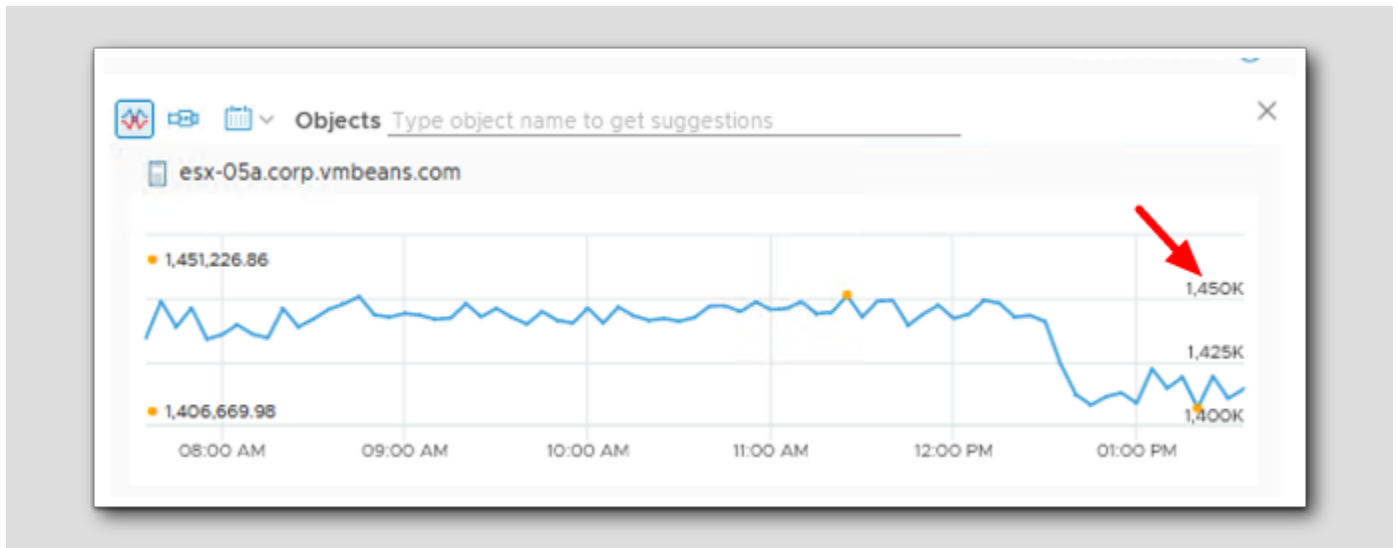


1. Click in the Objects search bar to bring up a list of Objects.
2. Select esx-05a.corp.vmbeans.com.

Notice that the list has a combination of Hosts, Virtual Machines and Cluster Compute Resource based on the Object Types we added earlier.

## Average Memory usage on the host

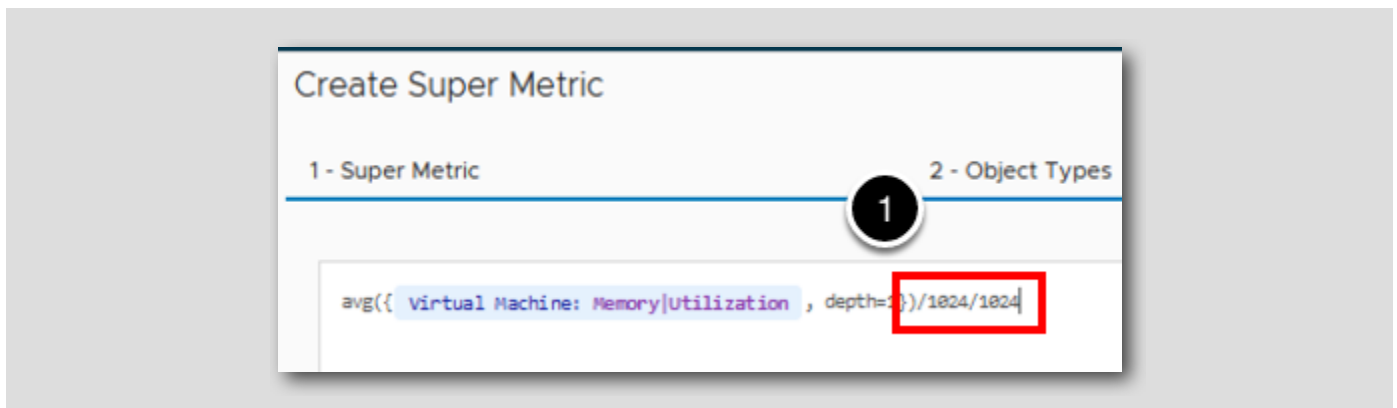
[411]



Notice that the chart is displaying memory usage in KB and not GB. Lets adjust our formula to accomodate that.

## Convert from KB to GB

[412]

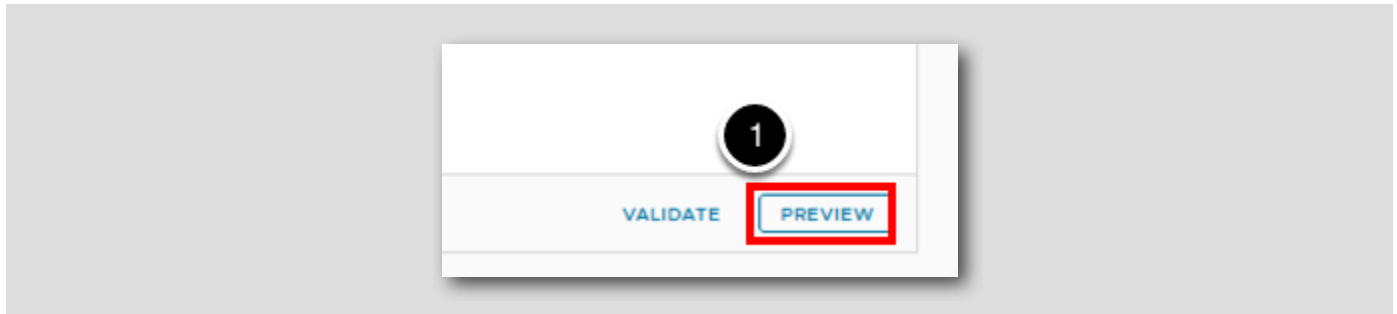


In order to convert our resulting value in kilobytes (KB) to gigabytes (GB), we need to divide the resultant average by 1024 to get to megabytes (MB) and then divide again by 1024 to get to GB.

1. At the end of the formula add /1024/1024 so the formula looks like - `avg({Virtual Machine: Memory|Utilization, depth=1})/1024/1024`

Do we see GB now

[413]




1. Click Preview again to expand the preview section.

Get rid of the old chart

[414]

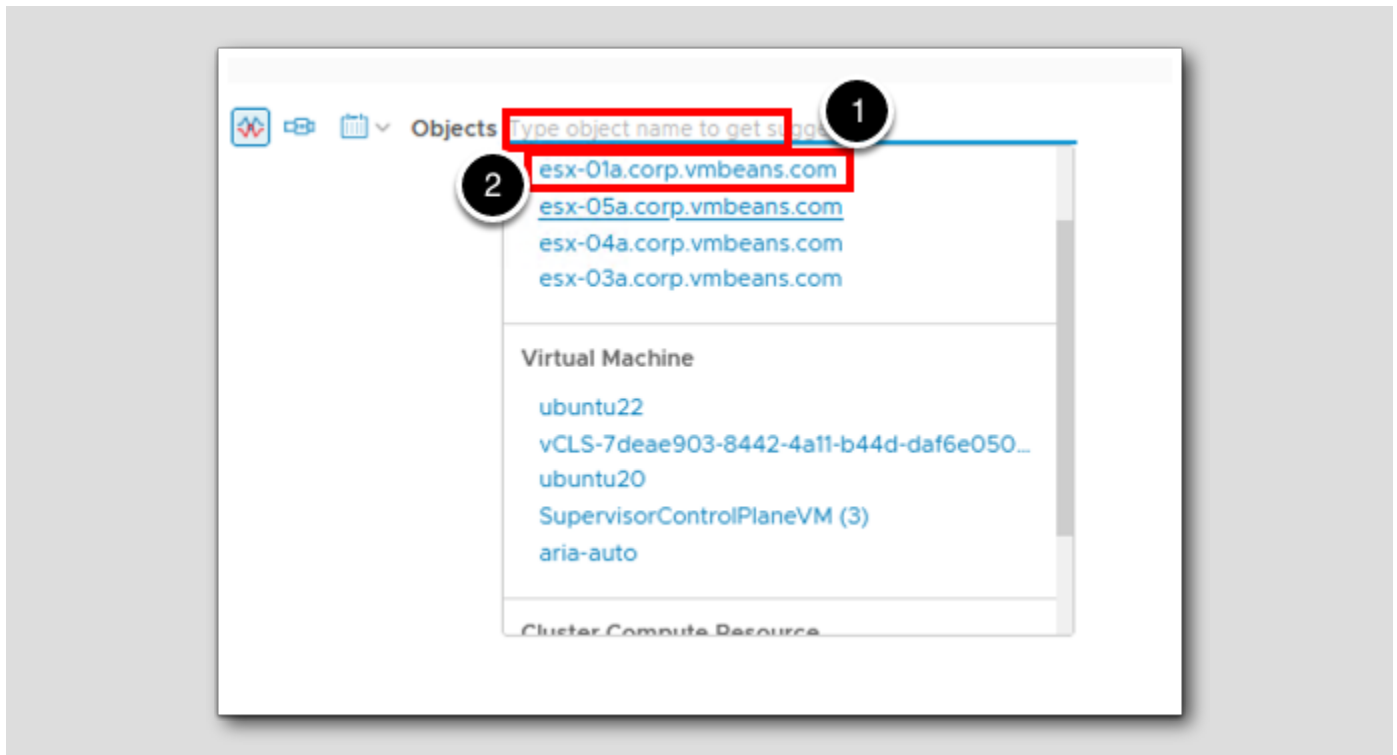
You will notice the old chart is still in preview and is showing in KB still. Lets close that chart and re-open it.



1. Hover your mouse over the chart and the blue X will appear, click on the blue X.

You will close out Preview if you click on the upper right X instead of the blue X. Just re-click on PREVIEW to get back in.

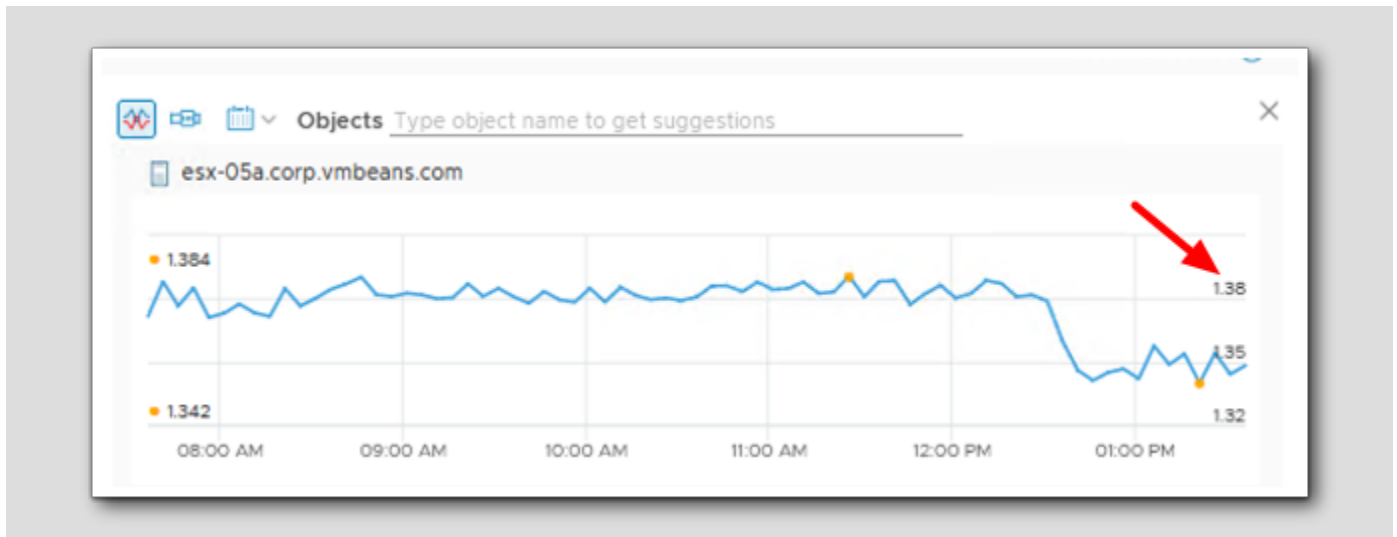
## Re-Add Esx-05



1. Click in the Objects search bar to bring up a list of Objects.
2. Select esx-05a.corp.vmbeans.com.

Shows GB now doesn't it

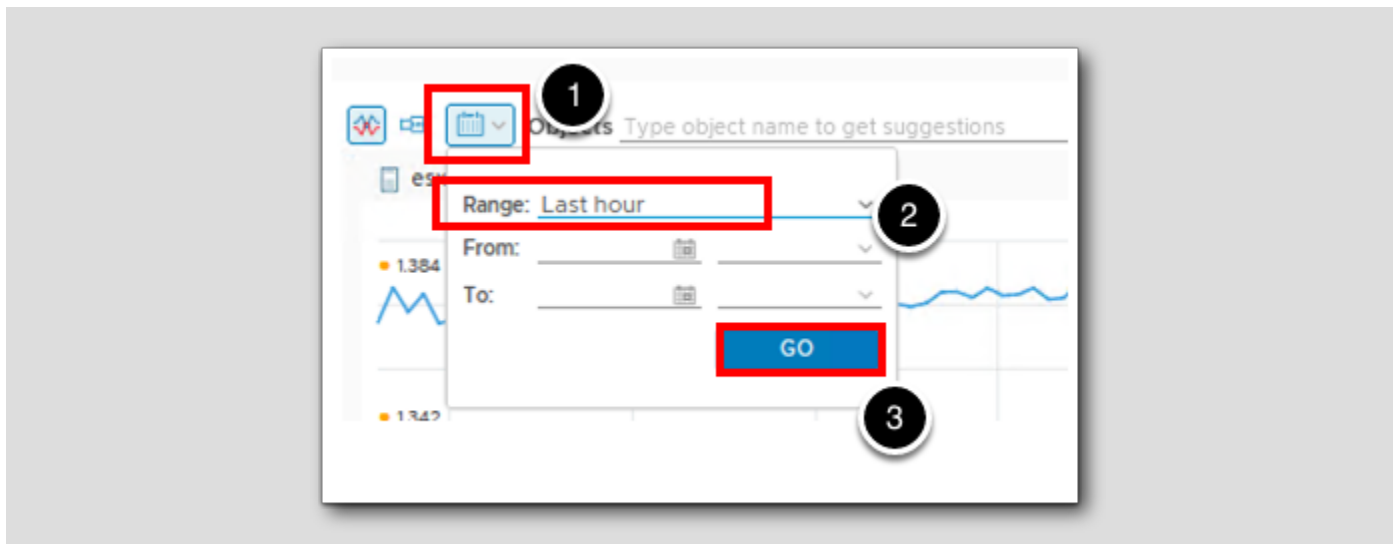
[416]



You should see that the chart is now displaying in GB vs KB with the extra /1024/1024 in there.

View the Host Preview

[417]



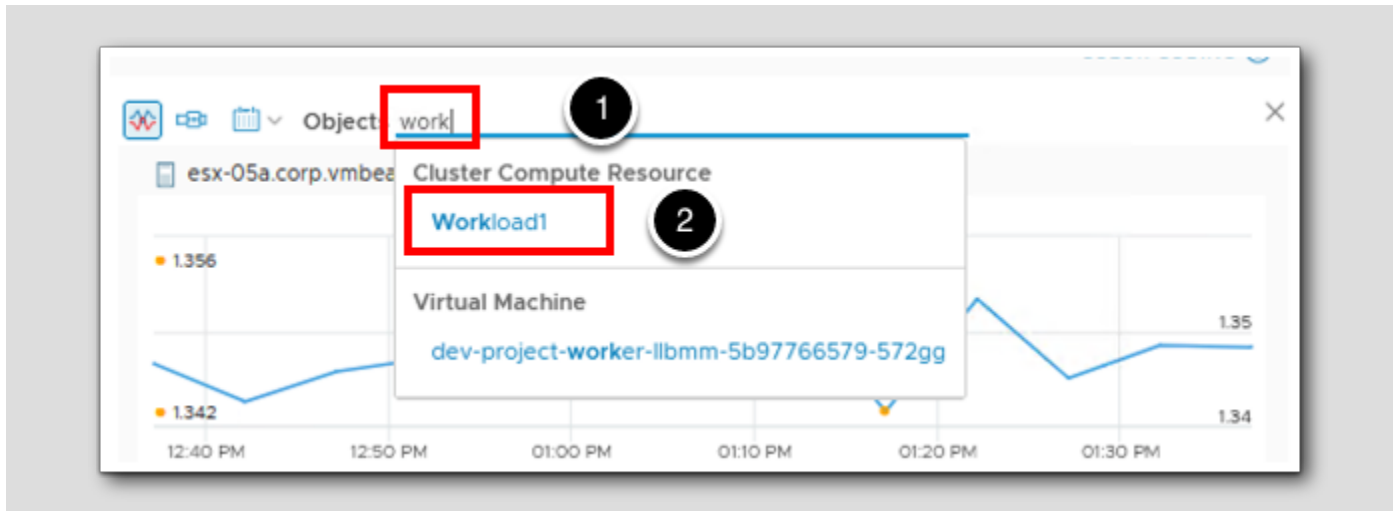
You should see a preview of your super metric on the esx-05a.corp.vmbeans.com host. Note that your values will likely be different and you may or may not see the graph cover the entire time period depending on how long your lab environment has been running before you started this lesson.



1. You can change the time period of the preview if desired. Click the **calendar icon**.
2. Click the **Range** drop-down and change the option from Last 6 hours to **Last hour**.
3. Click **GO** to apply the changes.

## Select the Workload 1 Cluster For Preview

[418]



Since we wanted our super metric to show the average vm memory utilization for hosts and clusters, let's test our metric on the Workload 1 vSphere cluster. There are four VMs running in that cluster so we should see the average memory utilization across the four VMs, right?

1. Delete the esx-03a entry, and type the first few letters of your cluster: **workl**
2. Select the **Workload 1** cluster

View the Cluster Preview

The screenshot shows the 'Create Super Metric' wizard in vSphere, specifically the 'Object Types' step. The formula entered is `avg({ Virtual Machine: Memory|Utilization depth=1 } 1824/1824)`. A red box highlights the `depth=1` parameter in the formula. To the right, there are two preview charts. The top chart, for `esx-05a.corp.vmbeans.com`, shows a line graph of memory utilization over time with values ranging from 1.34 to 1.356. The bottom chart, for `Workload1`, shows a flat line at zero utilization. A red box highlights this zero-value chart, and a '1' in a circle points to it. At the bottom, there are buttons for 'PREVIOUS', 'NEXT', 'CREATE', and 'CANCEL'. A 'Unit (Optional)' dropdown is also visible.

1. Now our preview shows zero. What's going on? That cluster has four VMs running and certainly the average memory utilization is not 0 GB per vm.
2. It's time to discuss the **depth** parameter.

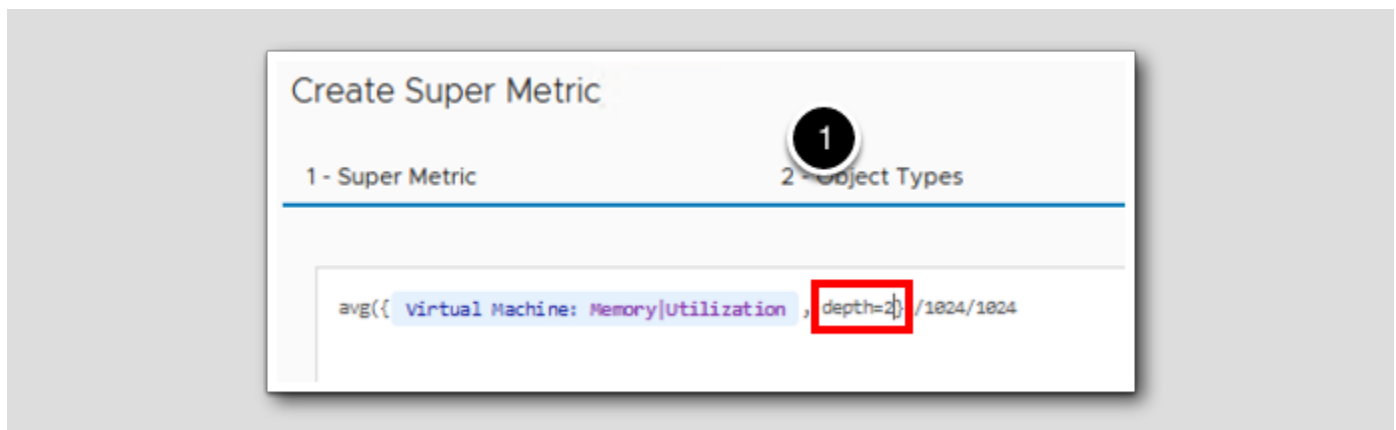
The depth parameter in a super metric formula is used to tell Aria Operations how far down (or up) the object hierarchy to look for the objects and their metrics when performing the calculation. As mentioned earlier, within Aria Operations there are multiple hierarchies (or traversal specs). Each adapter type will usually have at least one hierarchy. For example, the vCenter adapter creates vSphere Hosts and Clusters, vSphere Networking and vSphere Storage hierarchies.

If we look at the vSphere Hosts and Clusters hierarchy, it goes (from top to bottom): vSphere World --> vCenter Server(s) --> vSphere Datacenter(s) --> vSphere Cluster(s) --> vSphere Host(s) --> Virtual Machines --> Datastores. So in our case we want to calculate our super metric based on one (host --> vm) or two (cluster --> host --> vm) levels down the hierarchy. If you look at our super metric formula, you see that **depth=1** was added automatically which is why the preview worked on the esx-05a host (the vms were one level below the host) but not for the Workload 1 cluster (the vms were two levels below the cluster).

Something else you might notice about the **depth** parameter is that a positive value (1 in this case) will look down the hierarchy. If we wanted to look up the hierarchy, we would need to use a negative value for the depth parameter. That might seem opposite from what you would expect but you just need to remember that rule: positive depth = look down, negative depth = look up.

## Fix the Super Metric Formula

[420]



So let's update our formula to get it to look two levels down the hierarchy.

1. Delete the '1' and replace it with a 2 for the depth parameter.

Now the formula is calculating the average VM memory utilization for our cluster. But does that mean it won't work for hosts any longer? Since it is looking down two levels down in the hierarchy for vms will it look past the vms when applied to a host? The good news is that it will still work for hosts. In fact, a depth of 2 means it will look down one level and two levels. A depth of 5 would look down one, two, three, four and five levels for vms (or whatever object type is in the formula).

Back to Preview

[421]



1. Click the PREVIEW tab again.

Workload1 shows correctly

The screenshot shows the 'Create Super Metric' wizard in vSphere, specifically the '3 - Formula' step. The formula entered is `avg({ Virtual Machine: Memory[Utilization], depth=2})/1024/1024`. The 'Objects' list on the right shows two objects: 'esx-05a.corp.vmbeans.com' and 'Workload1'. Each object has a corresponding line chart showing memory utilization over time. A red arrow points to the 'Workload1' chart. At the bottom, the 'NEXT' button is highlighted with a red box and a circled '1', indicating the next step in the process.

Notice this time the chart preview updated without deleting and re-adding

1. Click NEXT.

## Assign The Super Metric To One or More Policies

**Create Super Metric**

1 - Super Metric      2 - Object Types      3 - Formula      4 - Policies

Select which policies you would like to enable this super metric in. You may also customize thresholds per policy. After one collection cycle, the super metric begins collecting and processing data, and it appears on each instance of the specified object type on the All Metrics tab.

Policy	Virtual Machine	Host System	Cluster Compute Resource
vSphere Solution's Default Policy (May 12, 2023 10:...	<input type="checkbox"/>	<input checked="" type="checkbox"/> 1	<input checked="" type="checkbox"/> 2

PREVIOUS    NEXT    CREATE    CANCEL

The final (optional) step is to enable the super metric for the object types in one or more policies. If you don't enable the metric calculation in a policy here, you will have to go edit the policy(ies) where you want to enable the calculation later in the policy editor.

In our lab we only have one policy that is being used. In a production environment you might have several or more policies active in Aria Operations. If you have multiple active policies you will see all of them listed on this screen and you can select which policies you want to activate the super metric calculation in for each object type.

1. Check the **box** to enable the metric on hosts in the policy.
2. Check the **box** to enable the metric on clusters in the policy.
3. Click **CREATE** to save your super metric.

## Lesson end

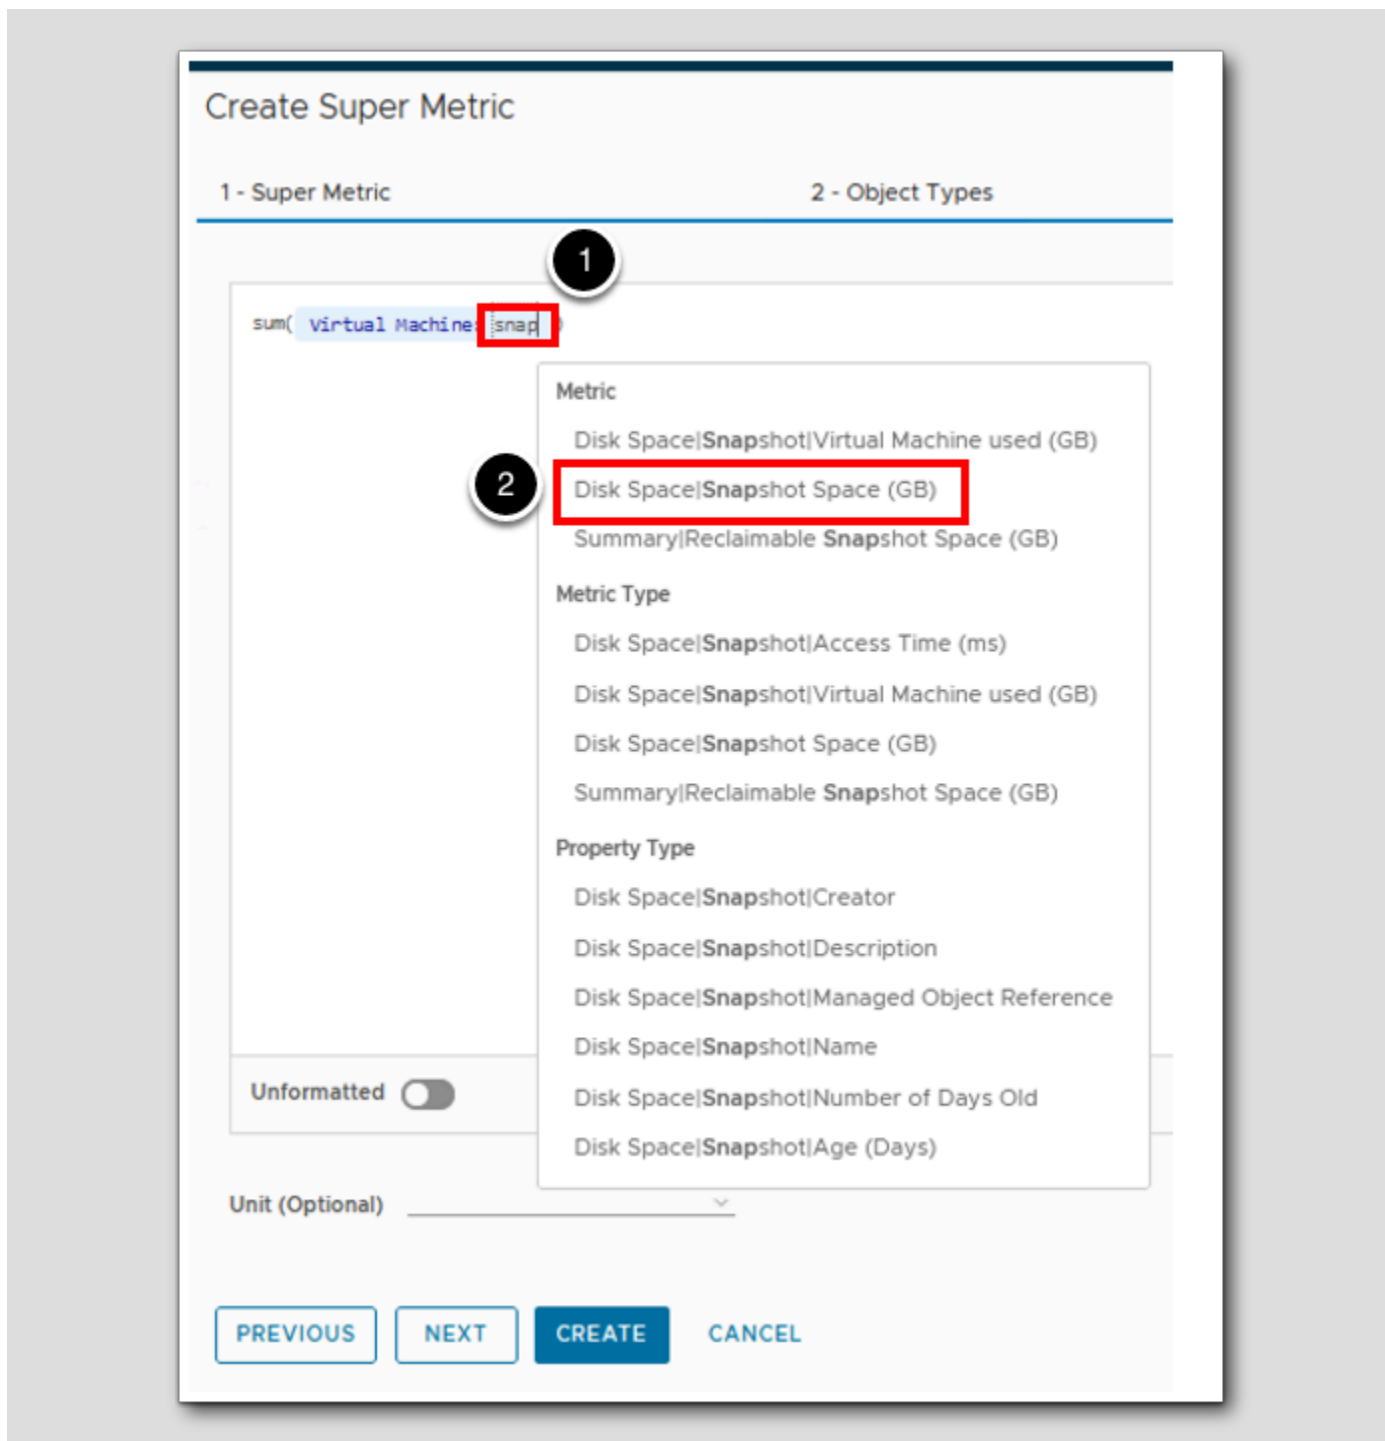
[424]

Congratulations! You have created your first super metric and applied to two object types in the active policy in your lab environment. There are a few more lessons ahead where we will explore creating other super metrics to learn about some additional super metric features. If you want to skip ahead and see the results of your work, use the Table of Contents at the top of the lab manual to jump past the other super metric creation lessons.

## Using 'This' and Negative Depth Parameters in Super Metric Formulas

[425]

Let's create another super metric. For this example, the assignment is to use a super metric to calculate the percentage of a datastore's capacity that is being used to store virtual machine snapshots.

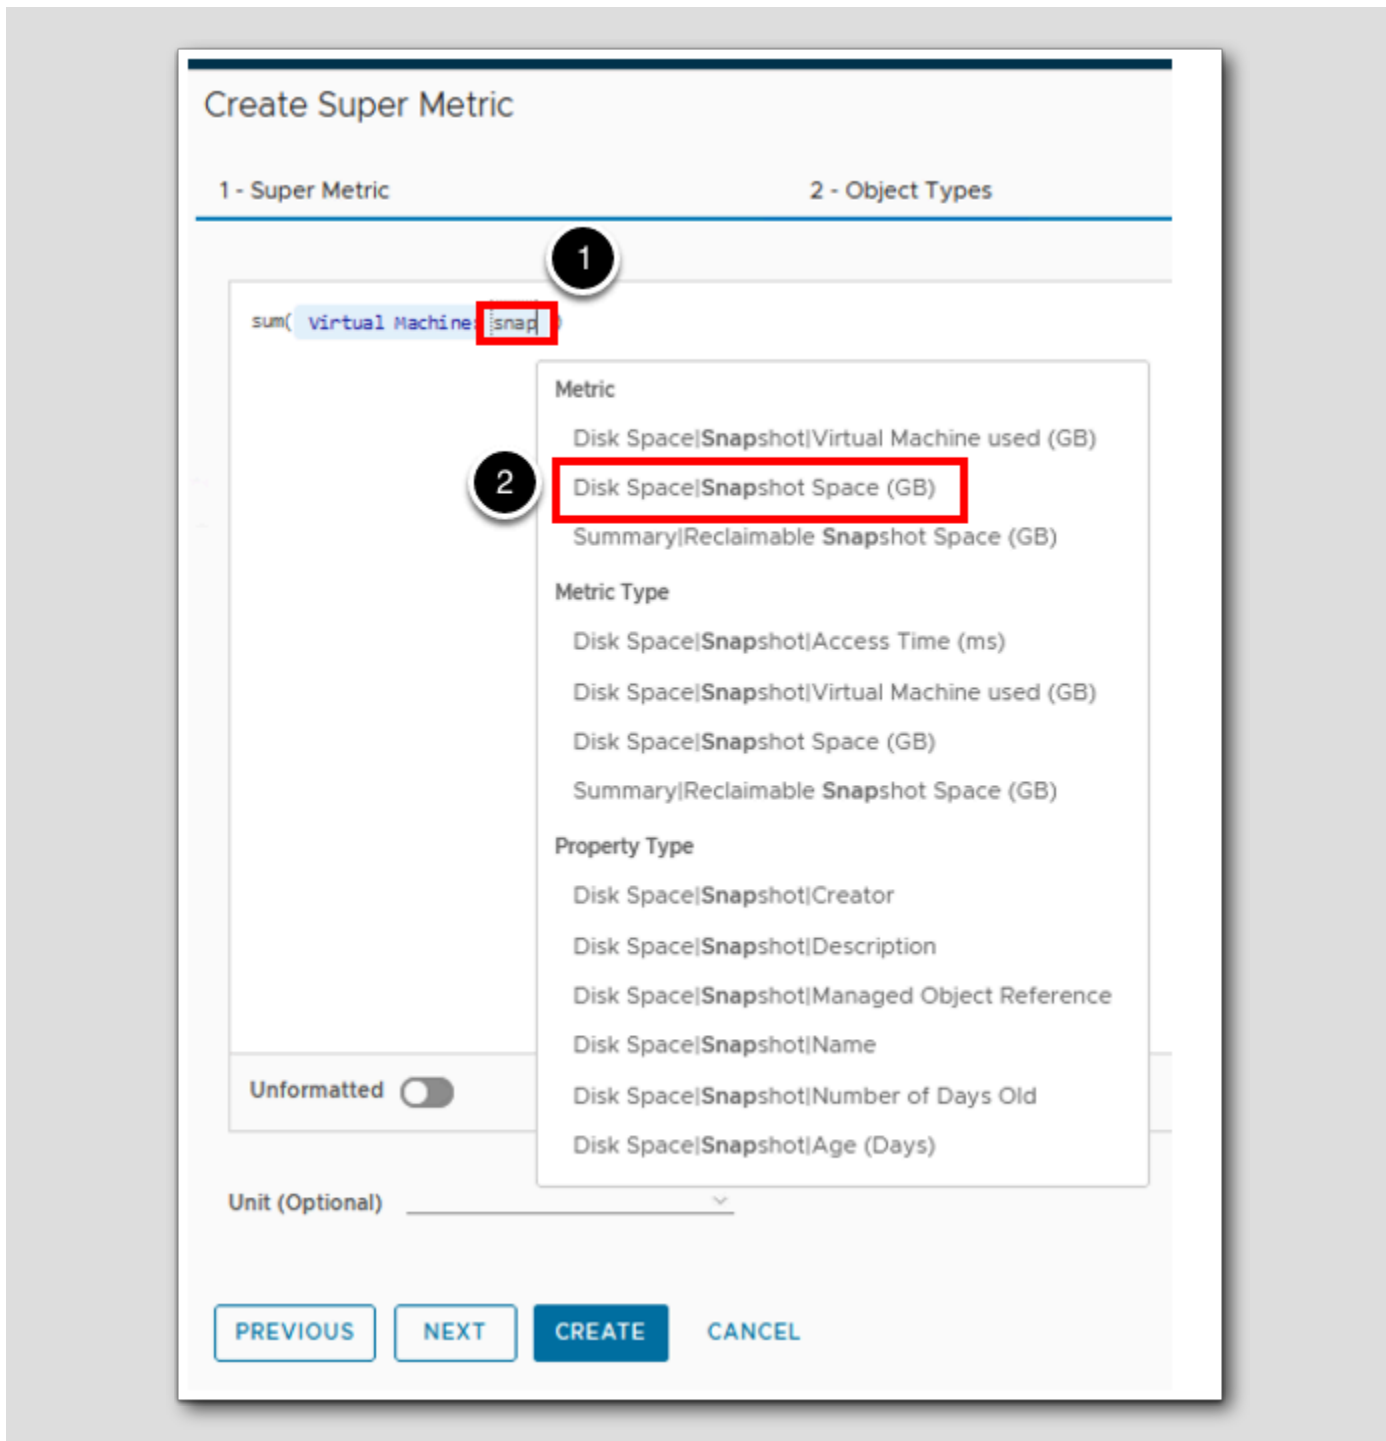


1. Type snap
2. Click Disk Space|Snapshot Space (GB) under the Metric section

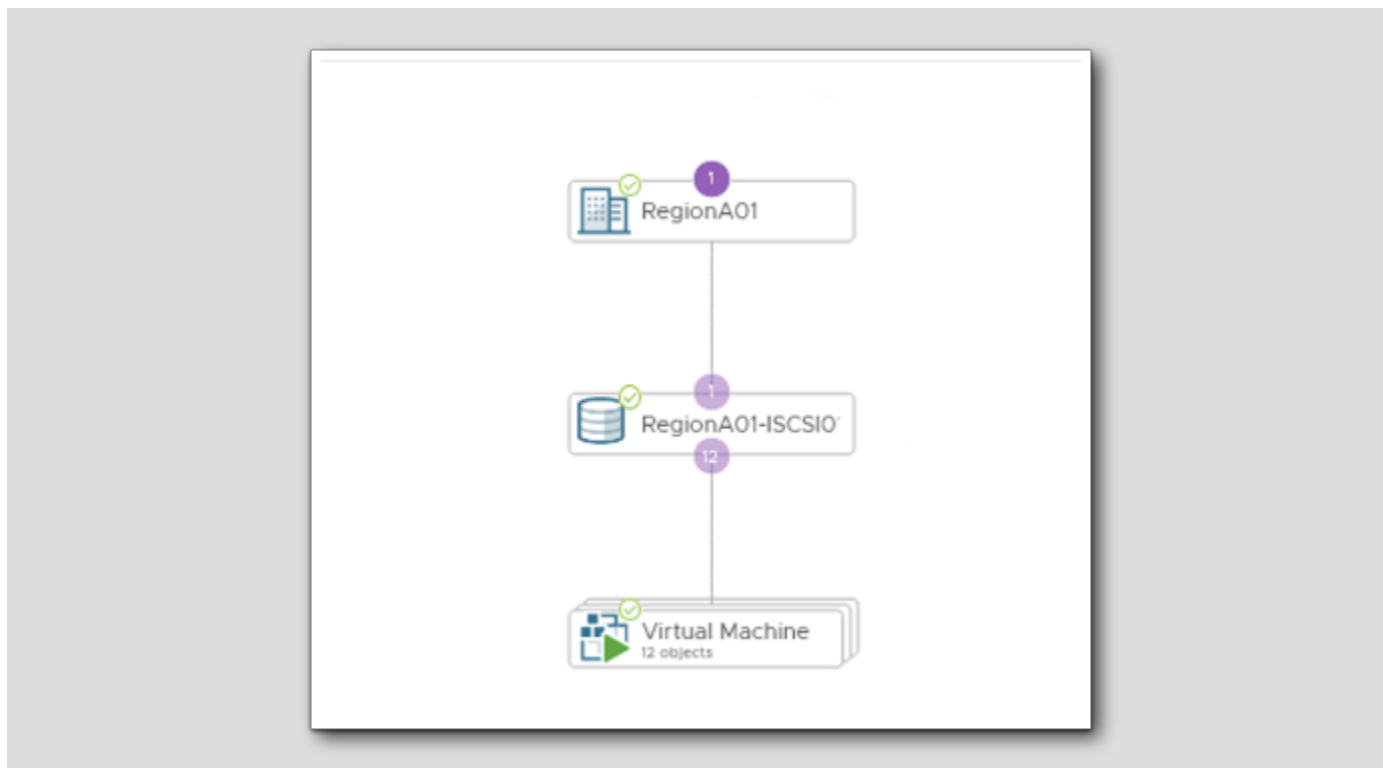
Note that the metric we want to use shows up both in the Metric Type category and the Metric category. Metric Type is a general attribute and should be used any time there might be more than one instance of the metric on an object (for example a CPU core's



usage where there are multiple cores in the host. Or the space used by individual snapshots when there are multiple snapshots on the virtual machine). In this case, the Disk Space|Snapshot Space is just a single metric that represents the total snapshot space used by the VM across all snapshots (if there are more than one).



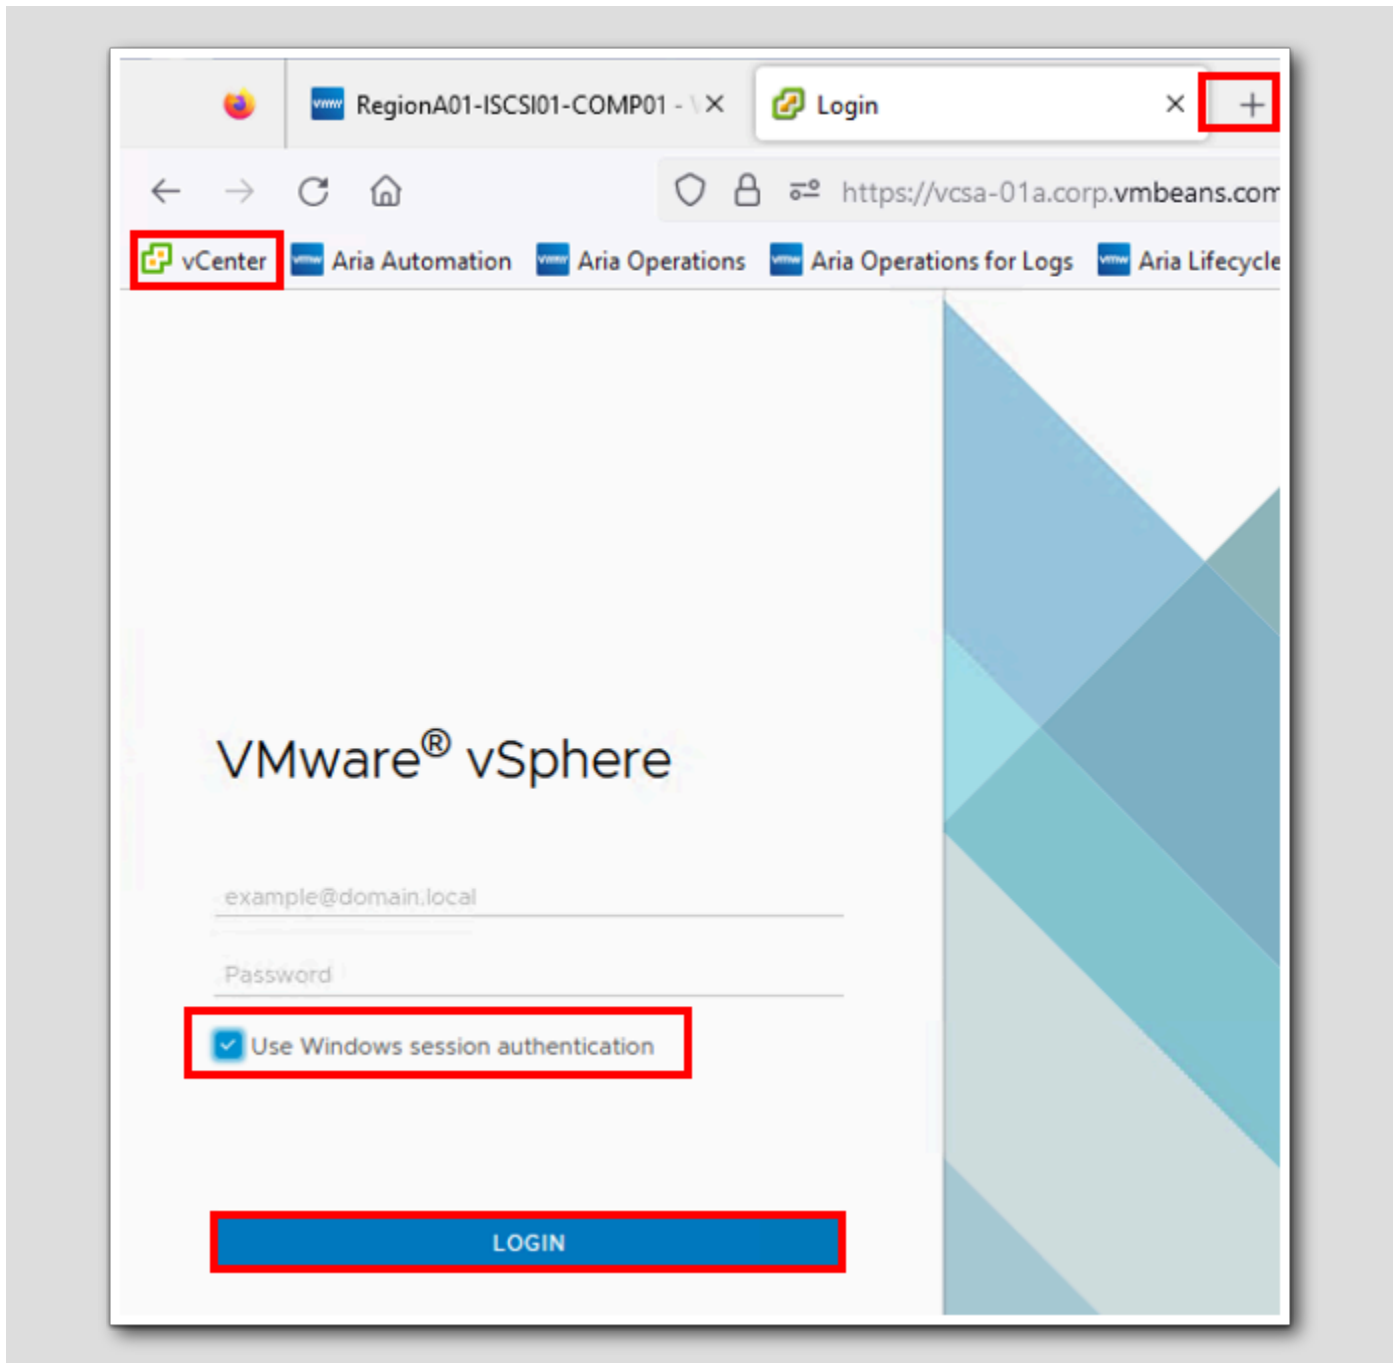
## The Hierarchical Relationship



If you recall from a previous lesson, we learned that a datastore is a child of hosts and of virtual machines in the vSphere Hosts and Clusters hierarchy. In this case, we will be using the VM <--> datastore relationship. Note in the graphic (and in our lab environment) that the RegionA01-ISCSI datastore supports Twelve virtual machines. So if we create a super metric on the datastore object type and have it look one level up the hierarchy to create the sum of the metric representing snapshot space on virtual machines, we will have completed the assignment for this lesson.

## Preparing the Lab Environment

[427]

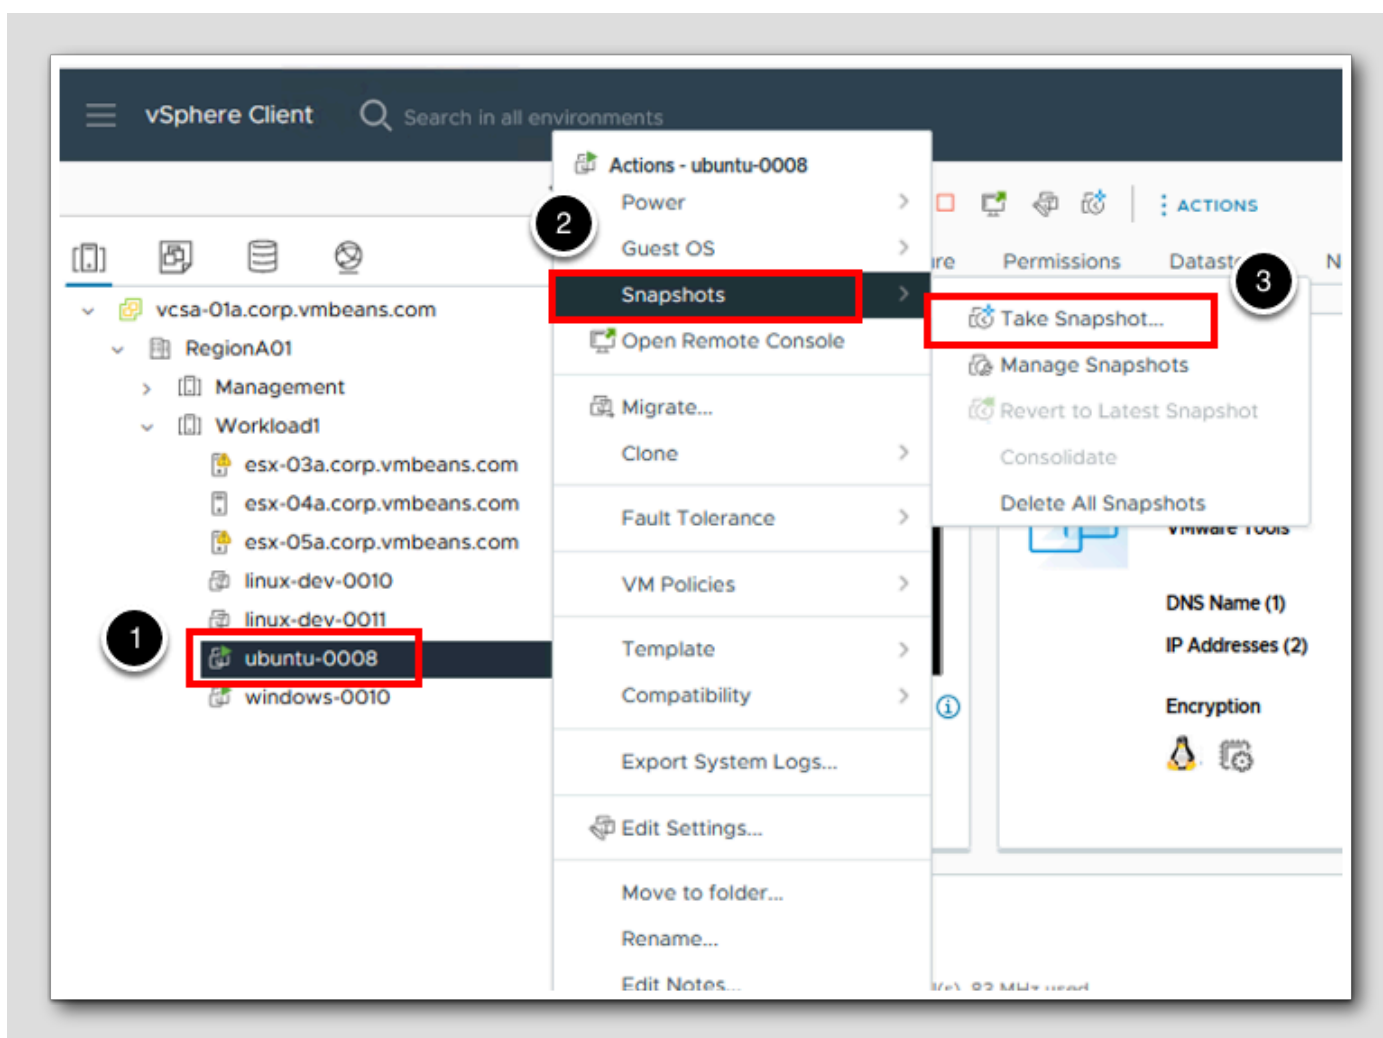


Before we begin this exercise, we must prepare the lab environment. None of the VMs in this lab currently have snapshots, so we will take one quickly.

1. Click the + sign in the Firefox title bar to open a new browser tab
2. Click the vCenter bookmark to navigate to the vSphere Client login
3. Click the checkbox next to Use Windows session authentication
4. Click LOGIN

## Take a Snapshot of the ubuntu-0008 Virtual Machine

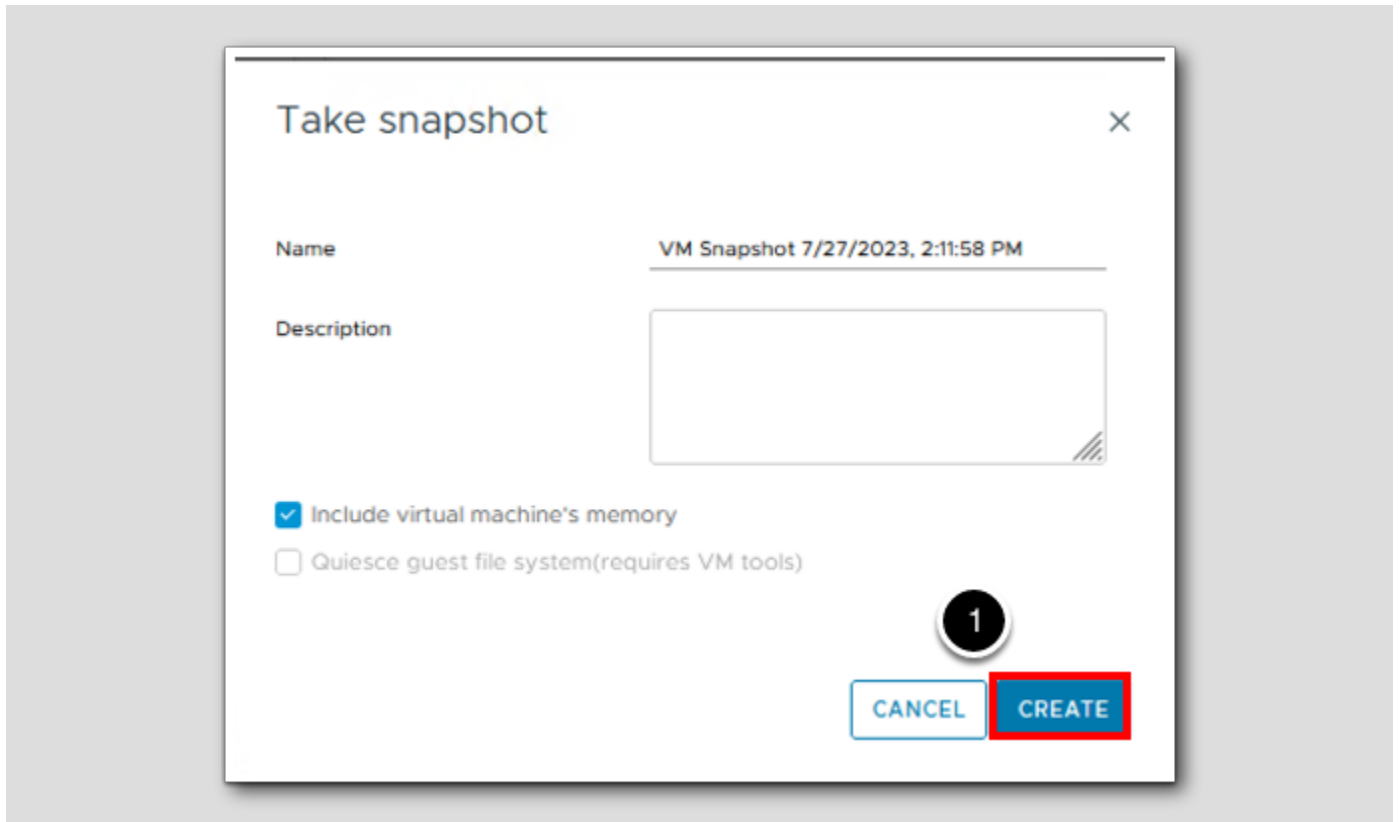
[428]



1. Right click on the ubuntu-0008 VM in the vSphere inventory to open the Actions - ubuntu-0008 menu.
2. Mouse over Snapshots to open the sub-menu
3. Click on Take Snapshot...

## Take a Snapshot of the ubuntu-0008 Virtual Machine

[429]



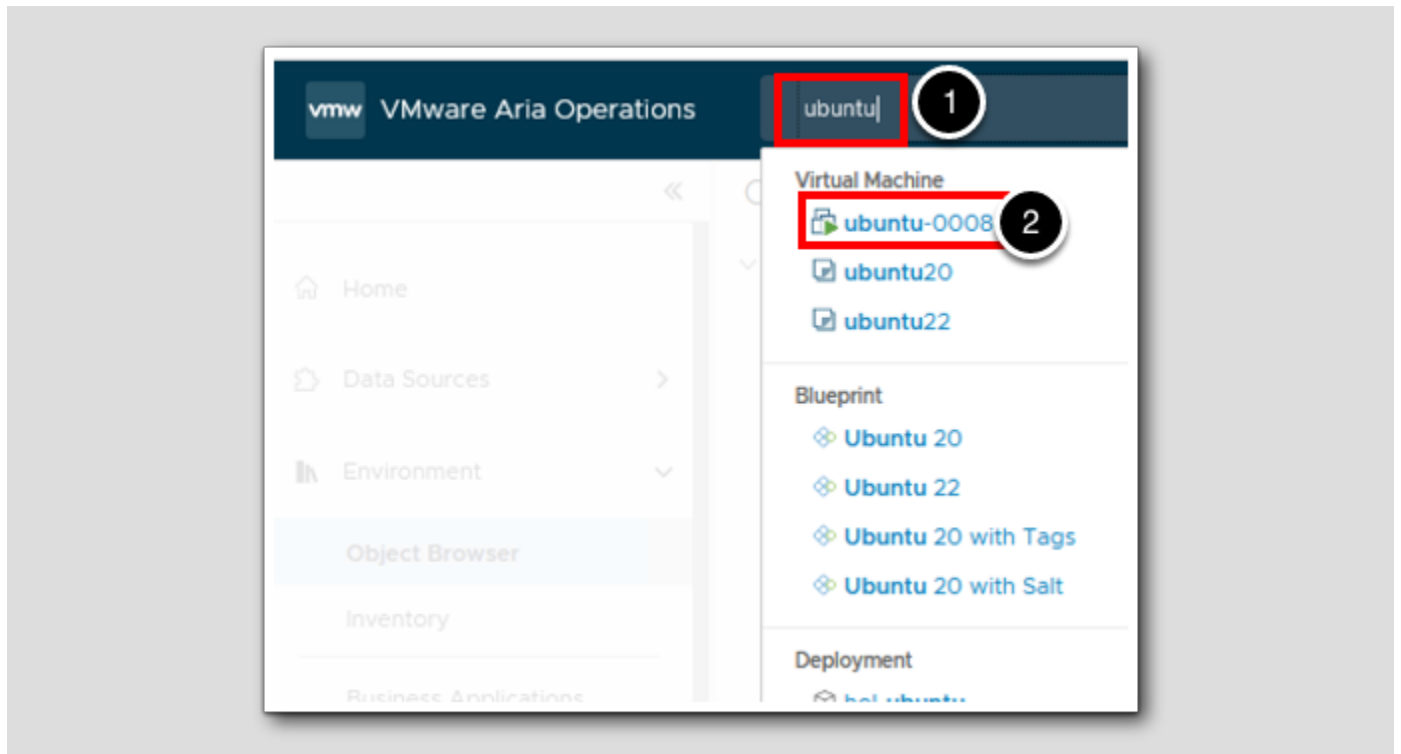
The default snapshot name and settings will suffice for this exercise.

1. Click **CREATE** to create the snapshot
2. Click the **Aria Operations Manager** browser tab (not shown) to return to Aria Operations

## Which Metric Will We Be Using?

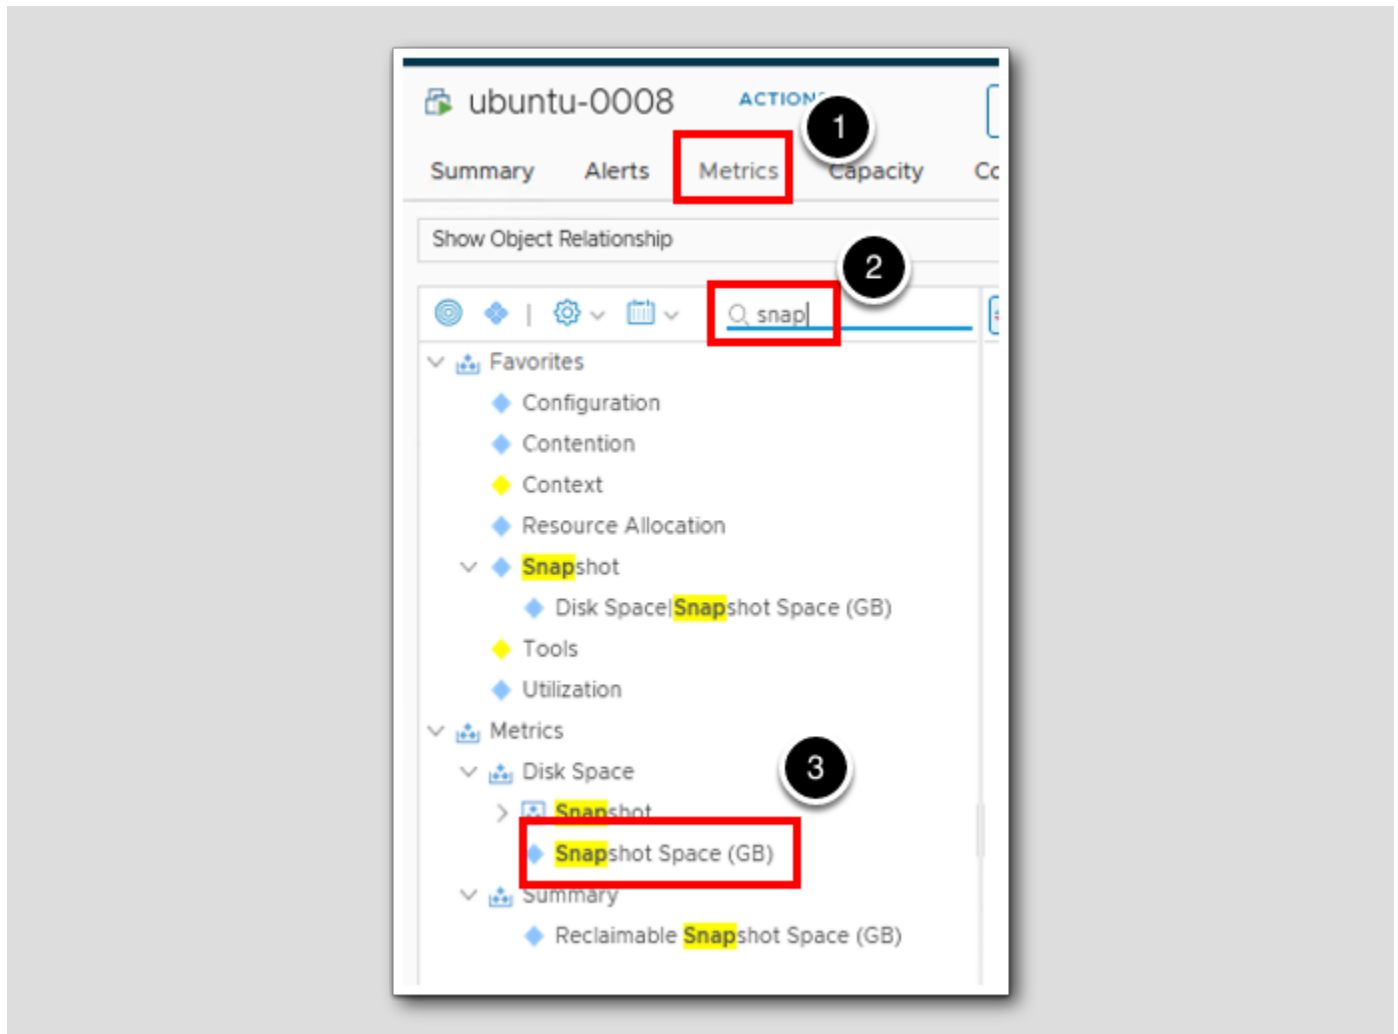
[430]

Before we get started with the super metric, let's understand which virtual machine metric we will be using for this lesson. Since we want to average a vm metric (disk snapshot space), let's go find a vm to see what metrics are available. We will again take a look at the ubuntu-0008 virtual machine.



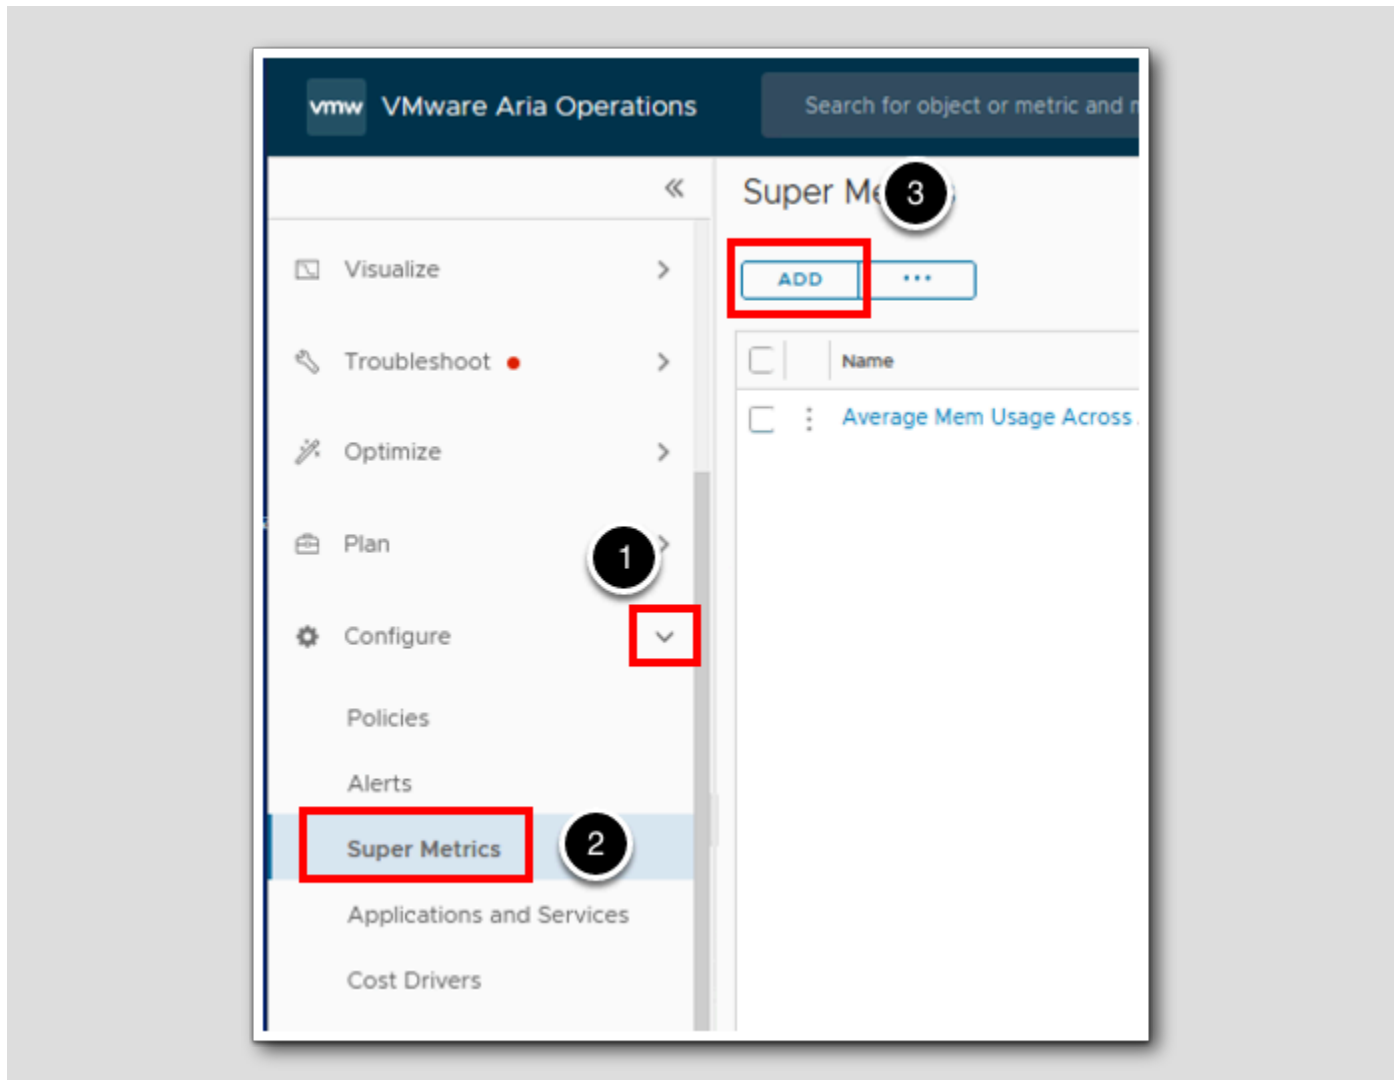
1. In the search box, type **ubuntu**
2. Click the **ubuntu-0008** link under the Virtual Machine object type

## Expand the All Metrics Tree



1. On the `ubuntu-0008` object page, click the **Metrics** tab.
2. In the filter field, type `snap` and press the **Enter** key to filter the metric results.
3. Metrics --> Disk Space --> **Snapshot Space (GB)** is the metric that represents the total space on the disk consumed by snapshots on this virtual machine.

## Create the Super Metric



1. Expand Configure.
2. Click on Super Metrics.
3. Click ADD.



### Create Super Metric

The screenshot displays the 'Create Super Metric' wizard interface. At the top, there are two steps: '1 - Super Metric' (indicated by a circled '1') and '2 - Object Types'. The 'Name' field contains the text 'Percentage of Datastore Capacity Used by Snapshots (%)'. The 'Description (Optional)' field contains the text 'On a datastore object, find the percentage of the total datastore capacity that is used by all VM snapshots on the datastore.' At the bottom, there are four buttons: 'PREVIOUS', 'NEXT' (highlighted with a red box and a circled '3'), 'CREATE', and 'CANCEL'.

**Create Super Metric**

1 - Super Metric **1** 2 - Object Types

Name: Percentage of Datastore Capacity Used by Snapshots (%)

Description (Optional): On a datastore object, find the percentage of the total datastore capacity that is used by all VM snapshots on the datastore.

**2**

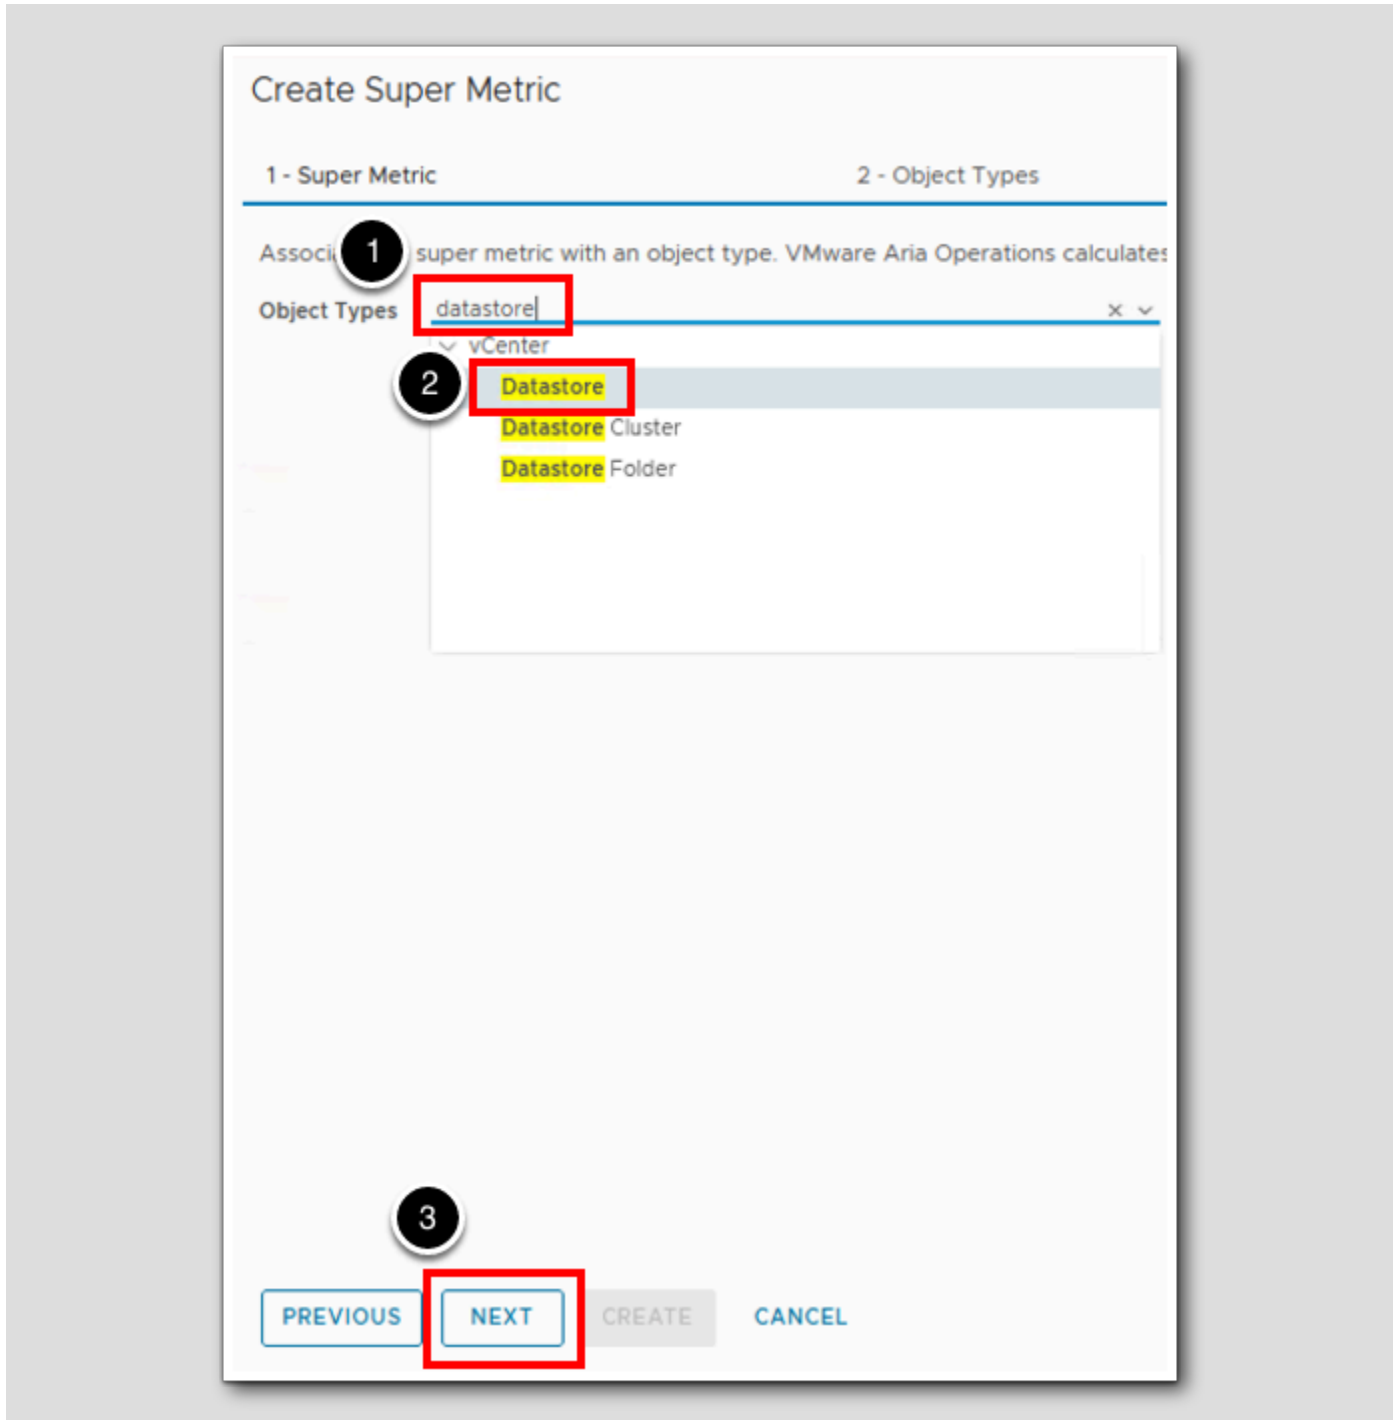
**3**

PREVIOUS **NEXT** CREATE CANCEL

1. Type the super metric Name: **Percentage of Datastore Capacity Used by Snapshots (%)**
2. Type a description for the metric (optional).
3. Click the **NEXT** button (not shown) to advance the wizard.

### Select the Object Types

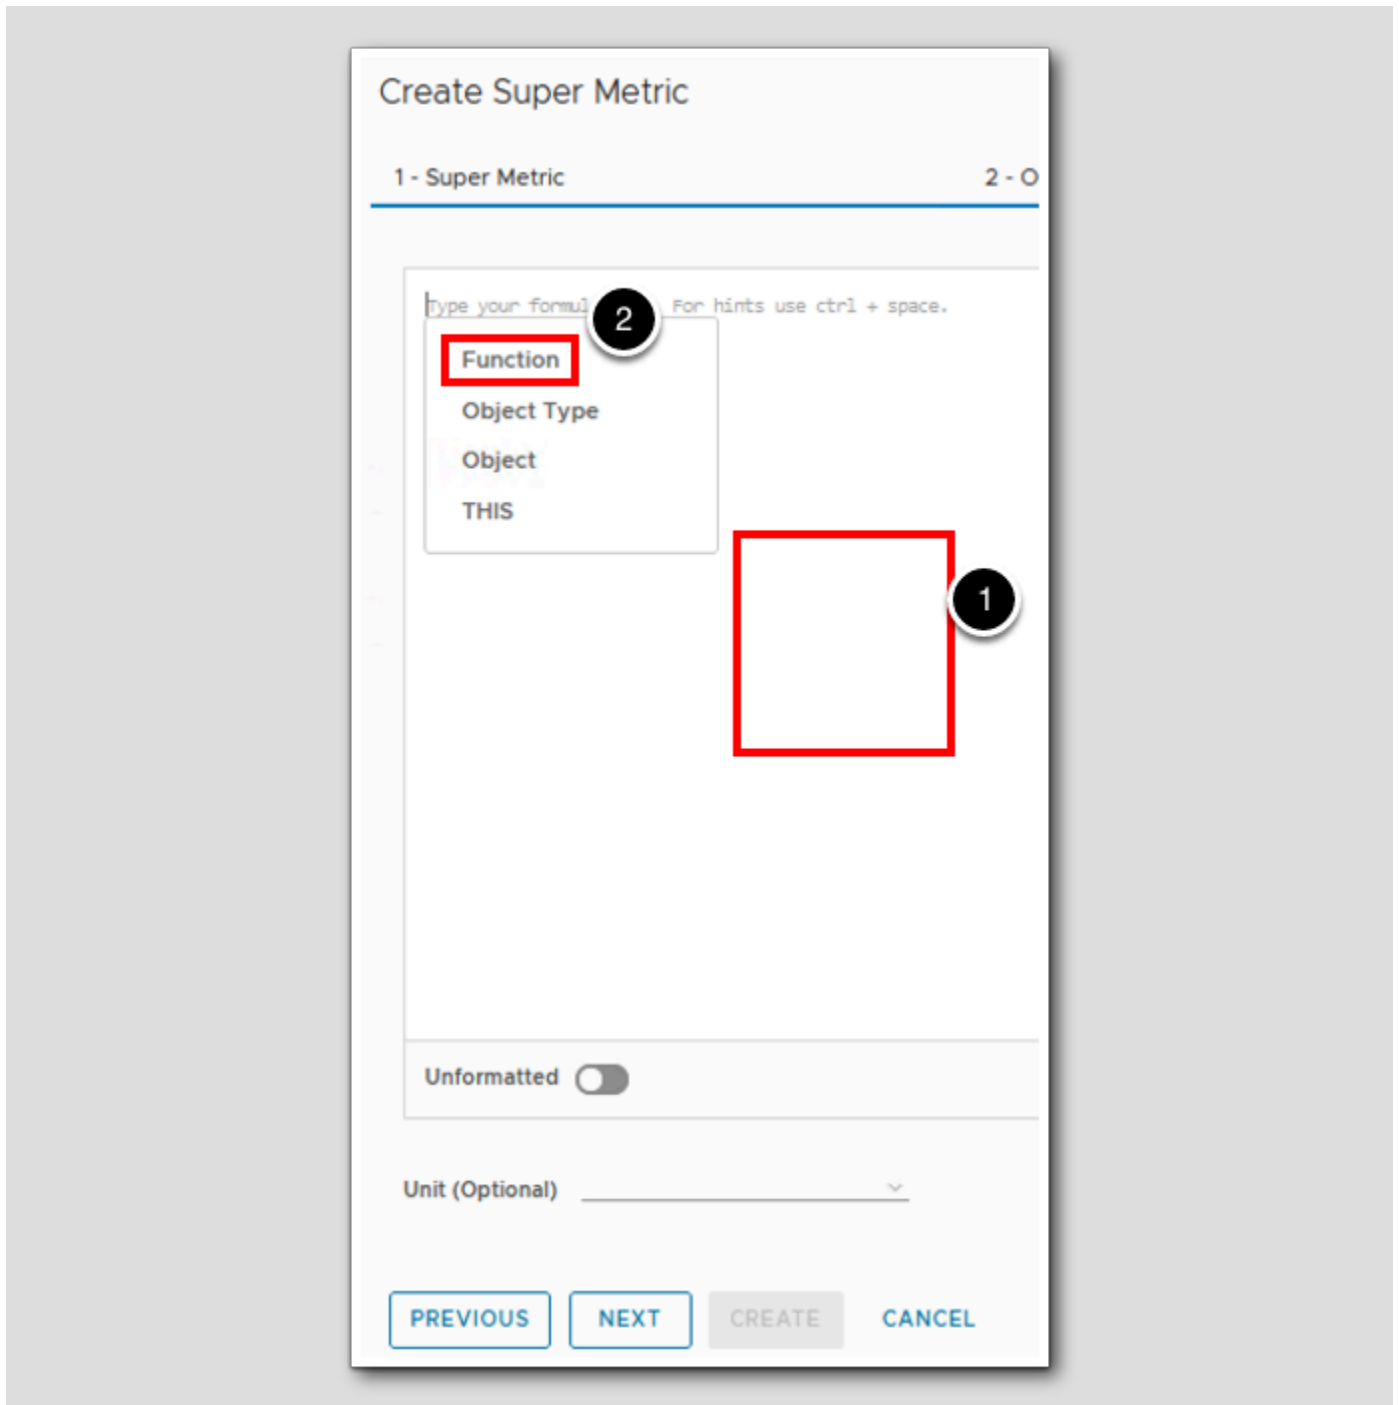
[434]



1. In the **Object Types** line type **datastore**.
2. Click **Datastore** under vCenter Adapter to select the object type.
3. Click **NEXT**.

## Start the Formula

[435]

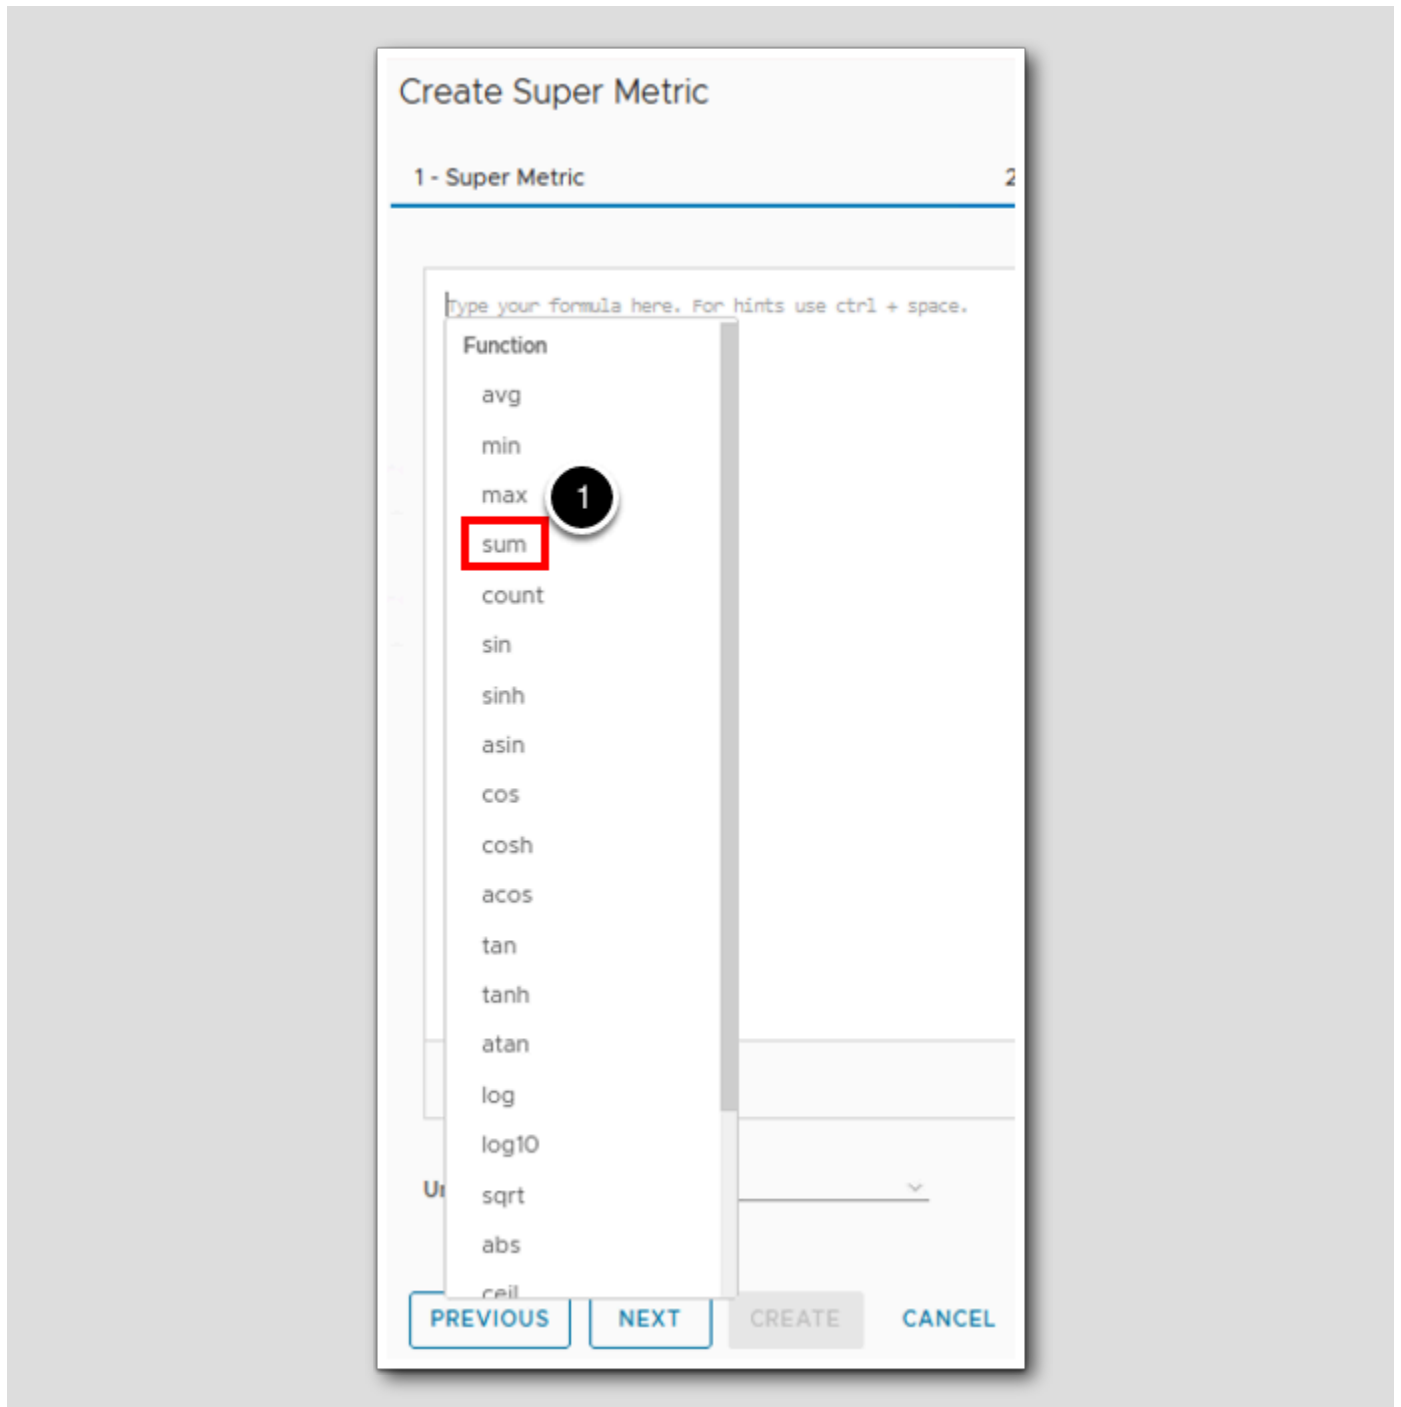


The formula will be: The sum of the snapshot space from all VMs on the datastore divided by the total capacity of the datastore.

1. Click anywhere in the empty formula box.
2. Select Function.

## Add Sum

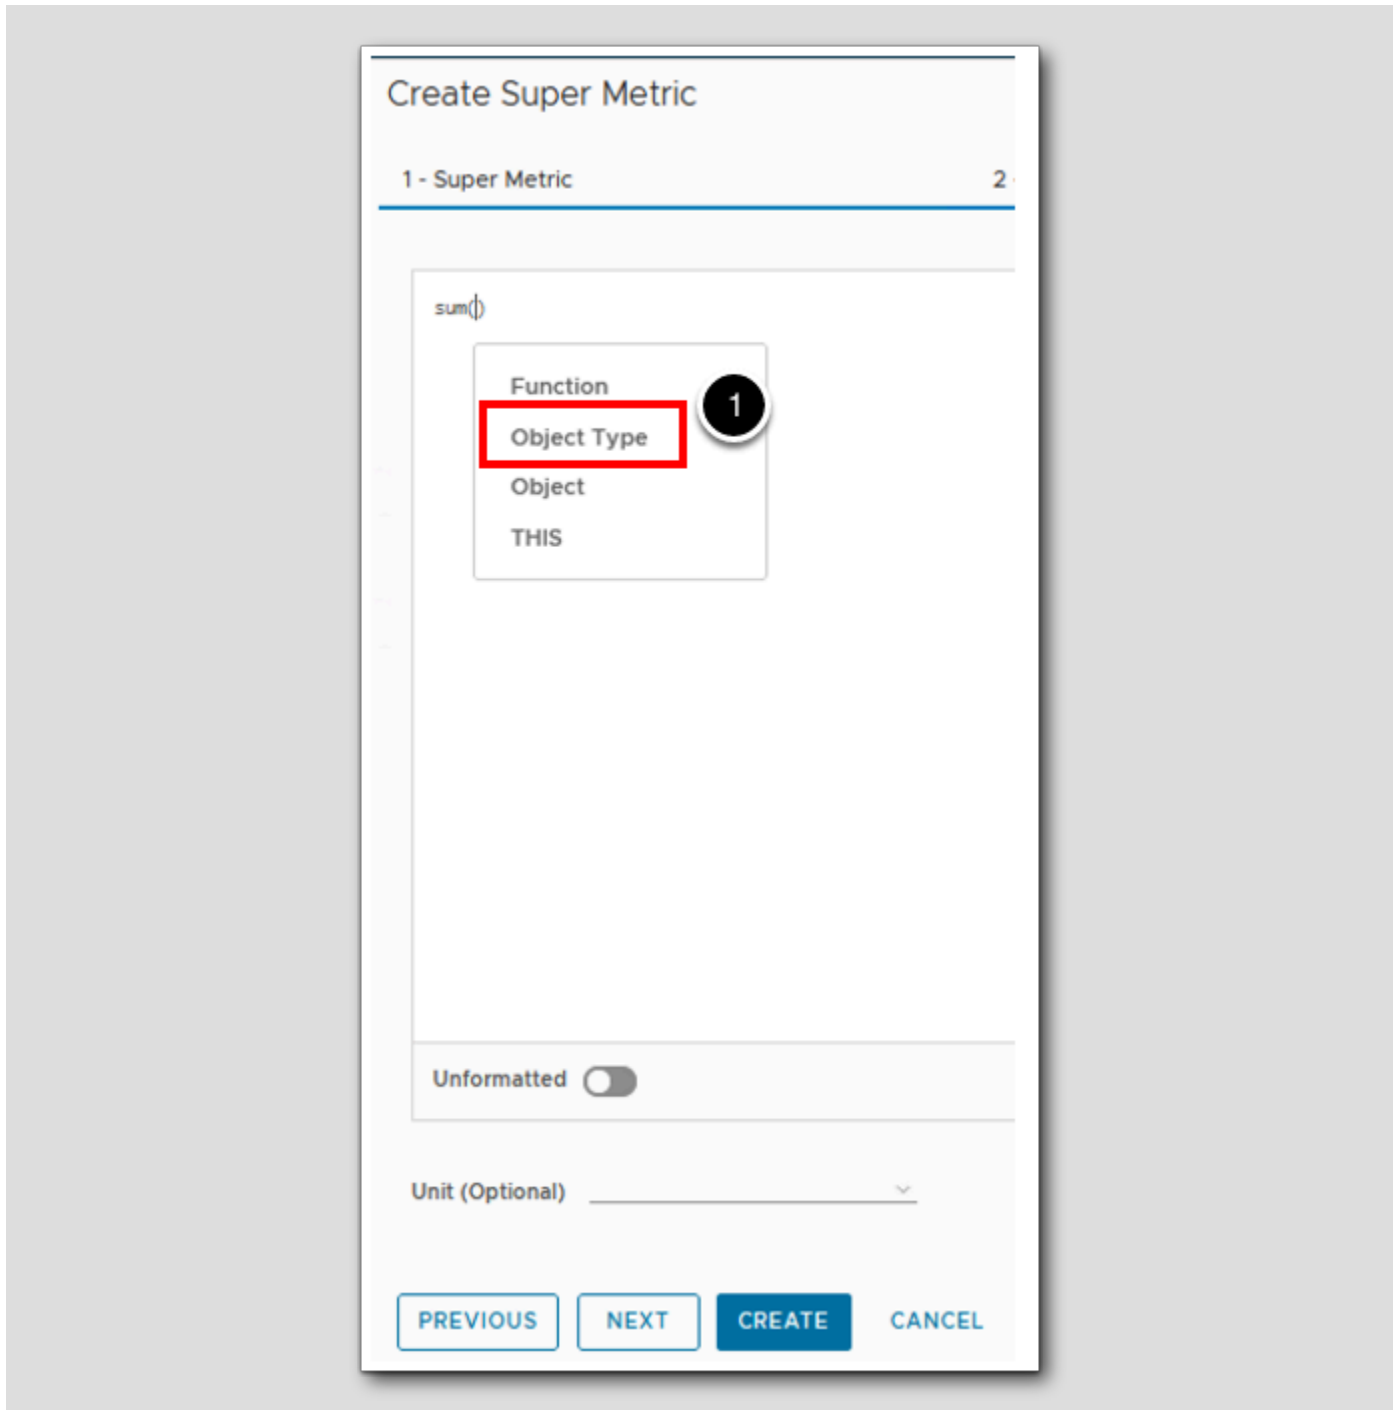
[436]



1. Select sum.

### Add the Virtual Machine Object Type

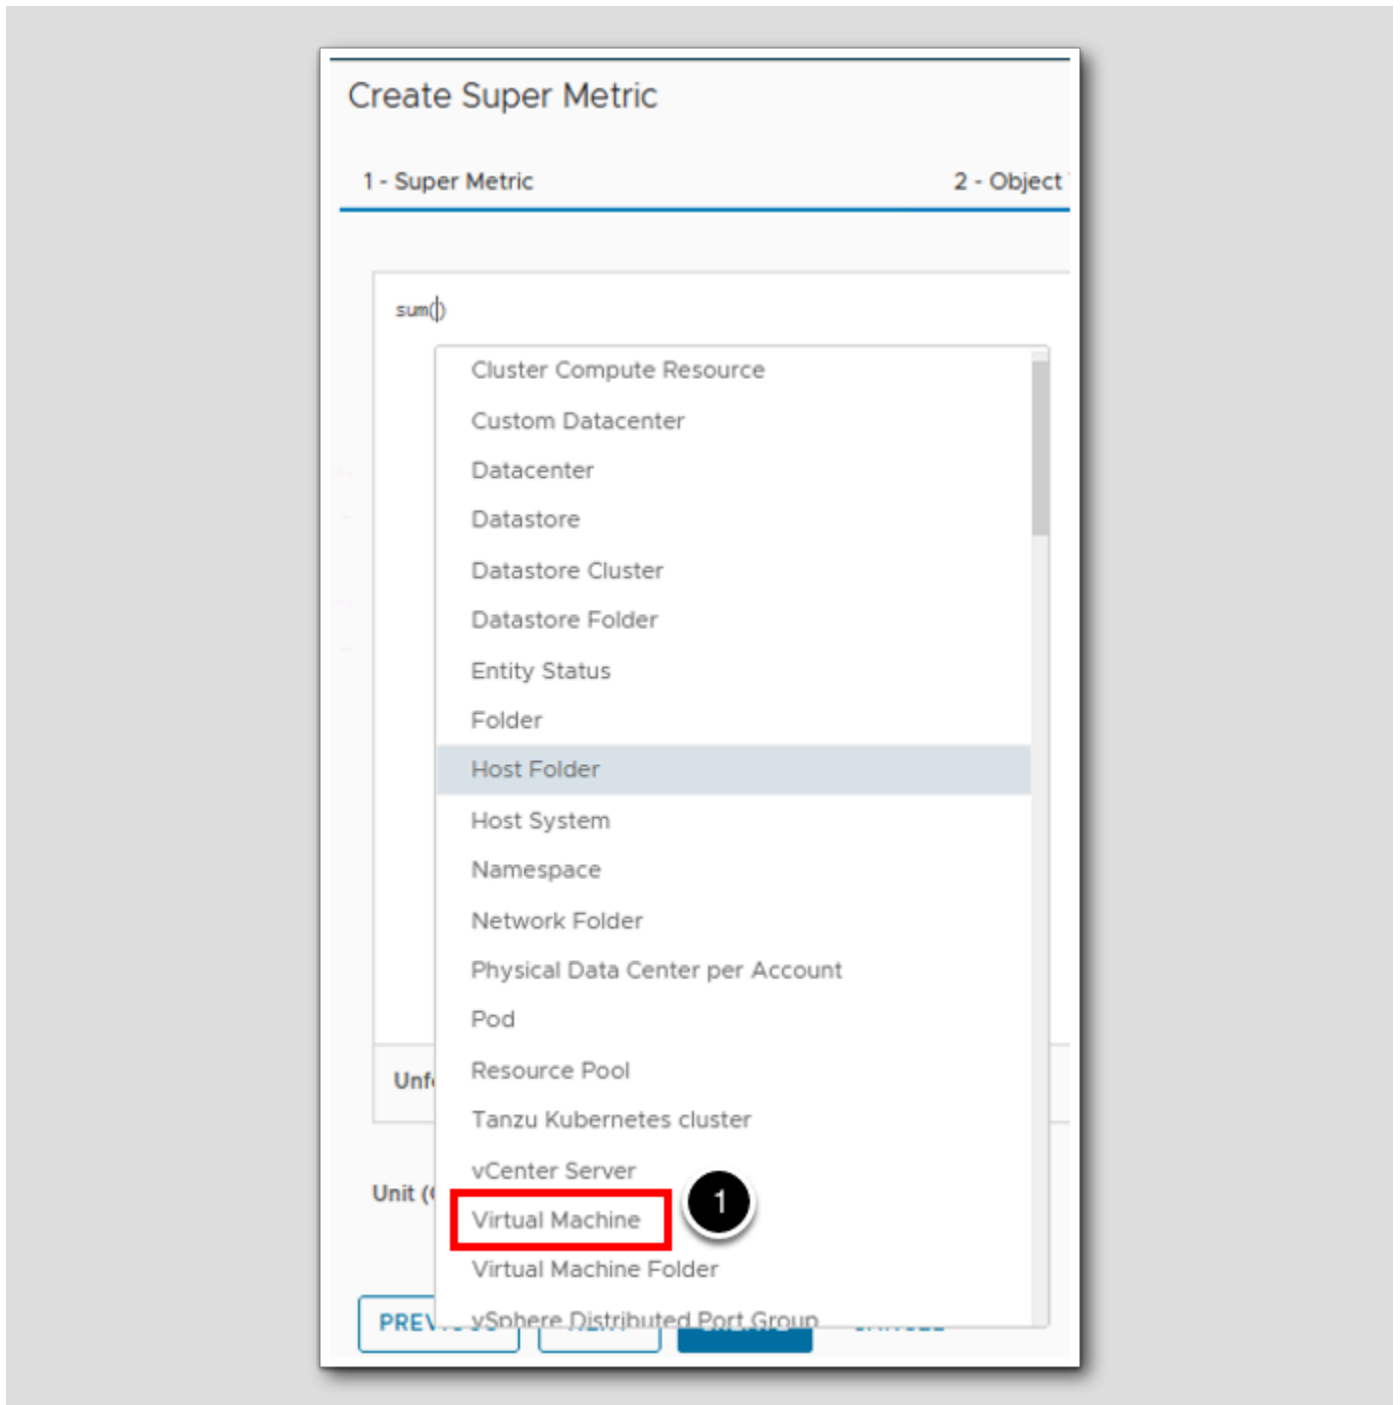
[437]





Select Object Type.

### Select Virtual Machine



1. Select Virtual Machine.

## Add the Metric

[439]

1. Select **Metric**.

Let's select the metric. Remember from earlier in this lesson that we will be using the **Disk Space|Snapshot Space (GB)** metric

## Continue Creating the Formula

[440]

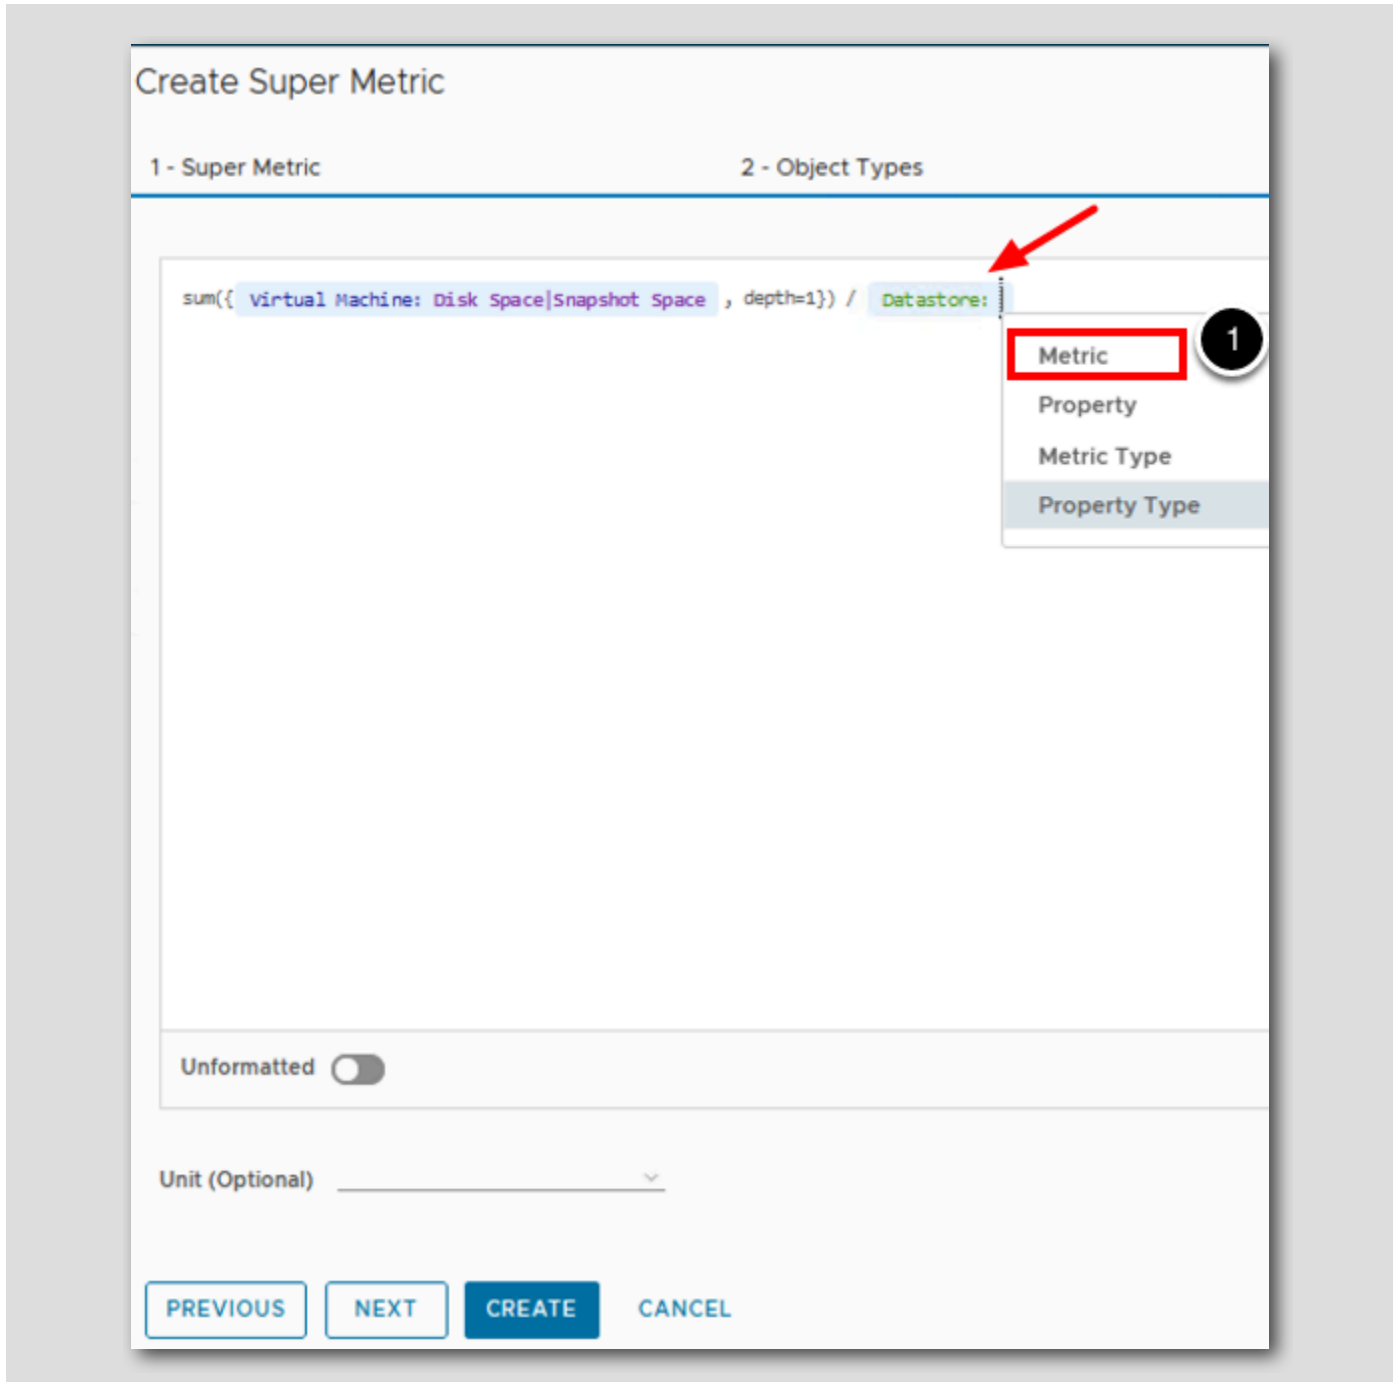
We have the numerator of our formula (the sum of the snapshot space from all VMs on the datastore). Let's add the division operator and get ready to add the denominator.



1. Move your cursor to the end of the formula and type a **space** followed by a **/** and then another **space** (note that the spaces are optional but they make the formula easier to read).
2. Select **THIS**.

Specifying 'This Object';

[441]



What happens when depth=0?

Let's take the example we are working on from the perspective of the datastore. The metric will be applied to datastore objects and we want to know for each datastore, what is the sum of the disk snapshot space from all of the VMs attached to that datastore (VMs are the parents) and then divide the sum by a metric on the datastore itself (the total capacity of the datastore). So if we are going to create a metric that will be attached to datastore objects and one of the calculation inputs is a metric from the datastore object itself, can we just say object type = datastore and depth = 0 in the super metric formula? Actually, there is special syntax for this type of situation ... instead of saying depth=0, it entails prefacing the metric or metric attribute with 'This Resource' and there is a special way of building that into the metric definition - the **THIS** button in the editor.

Clicking the THIS button has added a green Datastore: object in the formula.

1. Click on **Metric**.

## Select the Total Capacity Metric

**Create Super Metric**

1 - Super Metric      2 - Object Types      3 - Formula

sum({ Virtual Machine: Disk Space|Snapshot Space , depth=1}) / Datastore: total

Unformatted

Unit (Optional) \_\_\_\_\_

PREVIOUS    NEXT    CREATE    CANCEL

**Metric**

- Capacity|Total Capacity (GB)
- Capacity|Total Provisioned Consu
- Capacity Analytics Generated|Dis
- Cost|Monthly Total Cost (US\$/Mo
- Datastore|Total Latency (ms)
- Datastore|Total Throughput (KBp
- Devices:Aggregate of all Instance
- Devices:Aggregate of all Instance
- Devices:Aggregate of all Instance
- Disk Space|Total Capacity (GB)
- Disk Space|Total Provisioned Disk
- Summary|Total Number of Cluste
- Summary|Total Number of Hosts
- Summary|Total Number of VMs
- VMware Aria Operations Generat
- VMware Aria Operations Generat

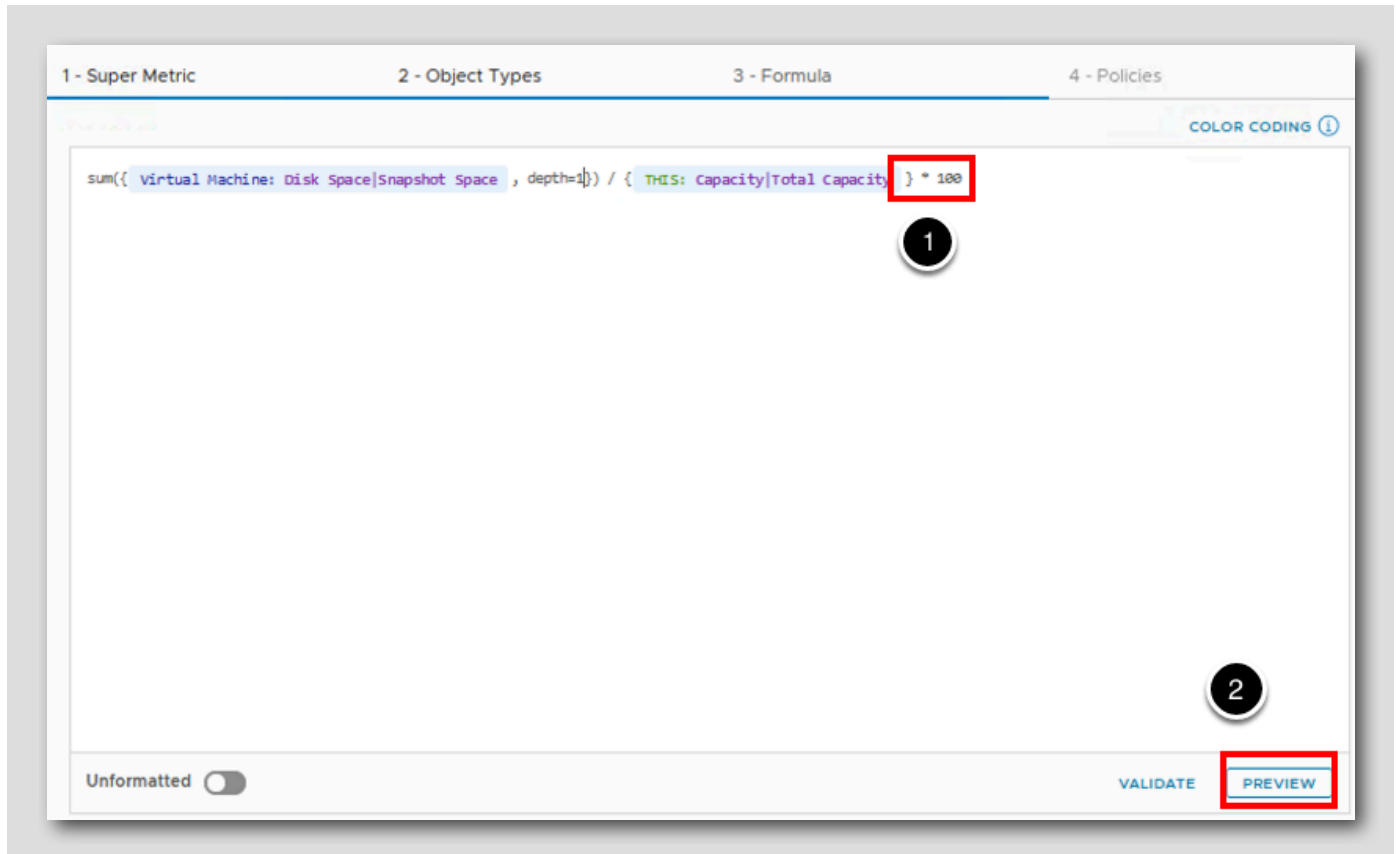
1. Type total to filter the list.

2. Click Capacity|Total Capacity (GB) from the Metric section to select the metric.



## Finish the Formula

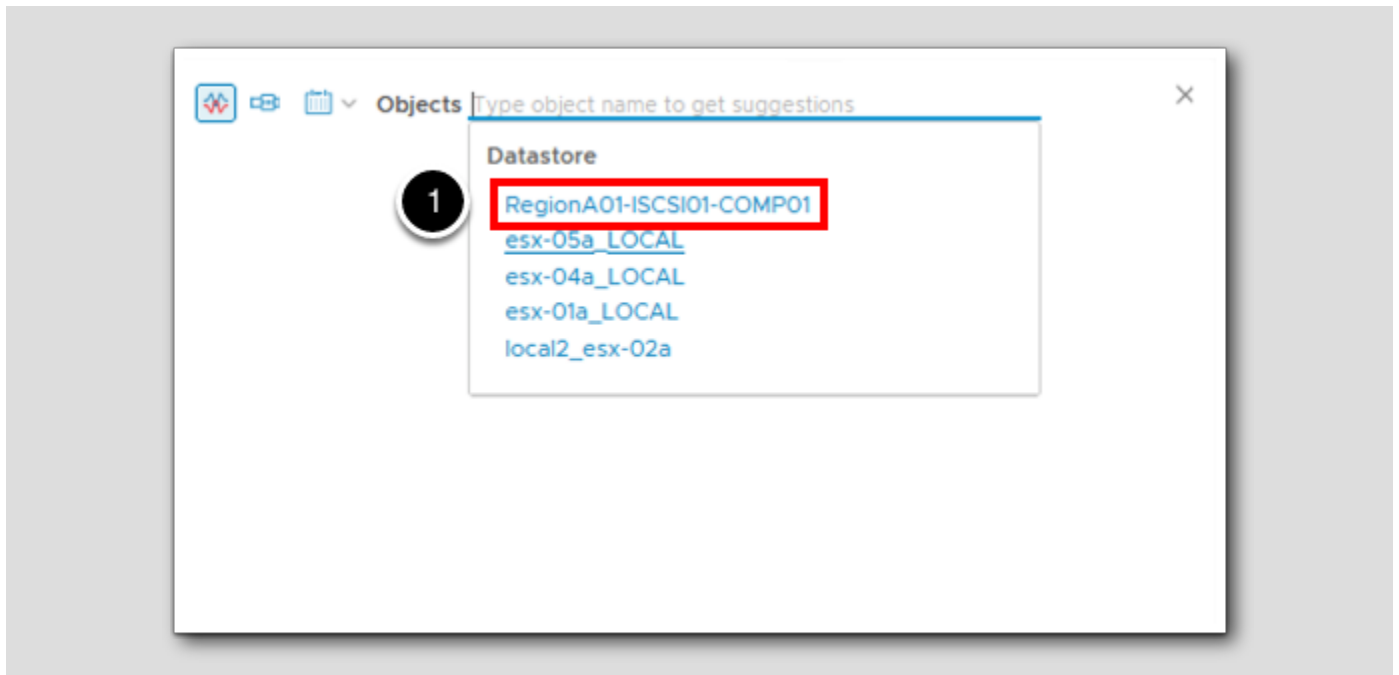
The result at this point will be a ratio of the sum of the snapshot space metric for all of the VMs on a datastore divided by the total capacity of that datastore. To convert it to a percentage, we just need to multiply by 100.



1. With your cursor at the end of the formula, type `space` then `*` then `space` then `100`
2. Click **PREVIEW** to open that section.

### Preview the Super Metric

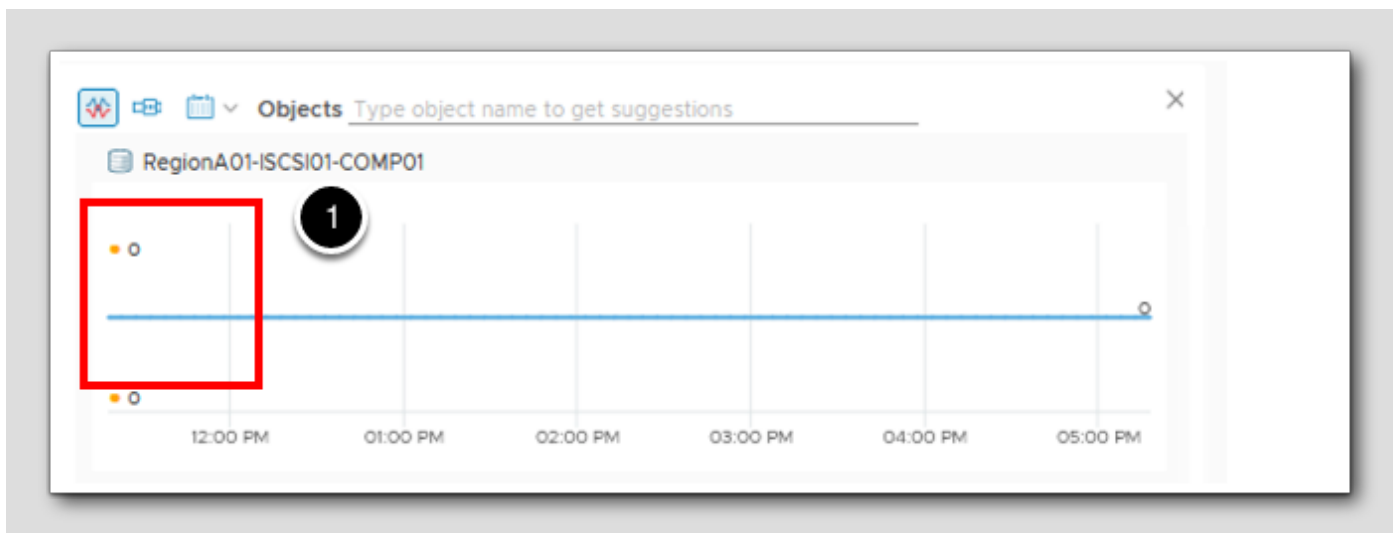
[444]



1. Click the RegionA01-ISCSI01-COMP-01 datastore object as the preview source.

### View the Preview

[445]



1. The preview shows zero percent of the datastore is used by snapshot storage. In this lab environment that is not the case. So why doesn't our formula work?

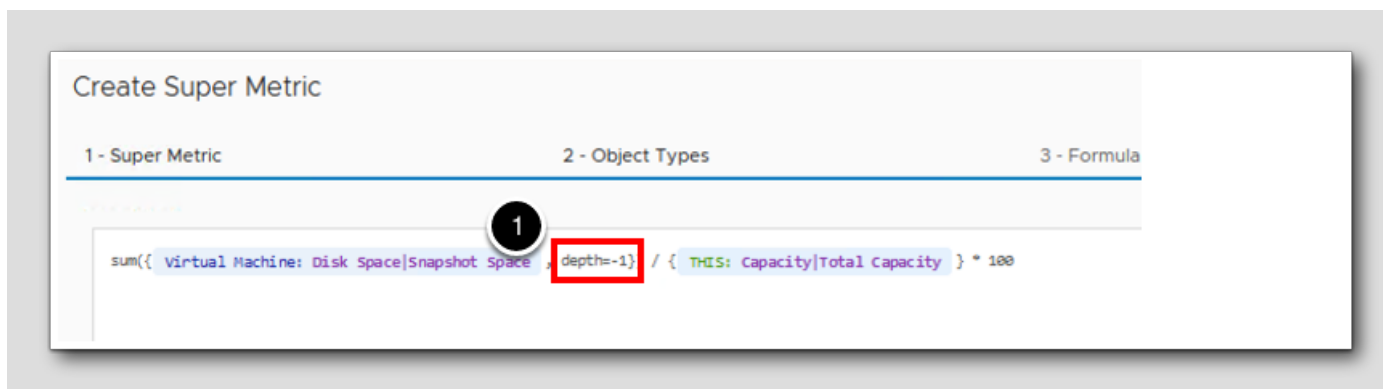
Do you remember the relationship hierarchy between datastores and VMs? Do you remember how the depth parameter works in a super metric formula?

In this case, virtual machines are parents of datastores. Our depth parameter on the datastore object is 1. Remember that a depth of 1 means one level down the hierarchy. But here we need to look up the hierarchy one level - from the datastore to the VM. So instead of depth=1, we need to have depth=-1.

Remember? Positive depth means look down the hierarchy. Negative depth means look up the hierarchy.

## Fix the Formula

[446]

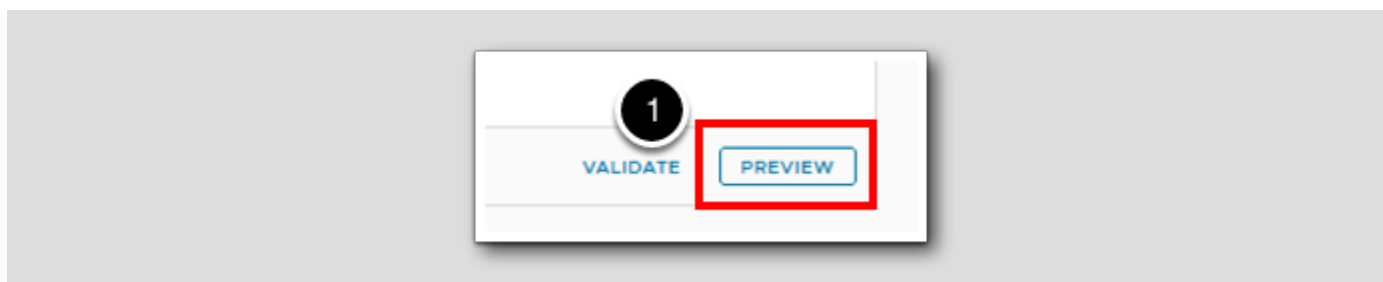


Let's fix the depth parameter and try the preview again.

1. Place your cursor just to the left of the 1 in the depth parameter and type a minus sign (-) to make the depth=-1

## Back into Preview

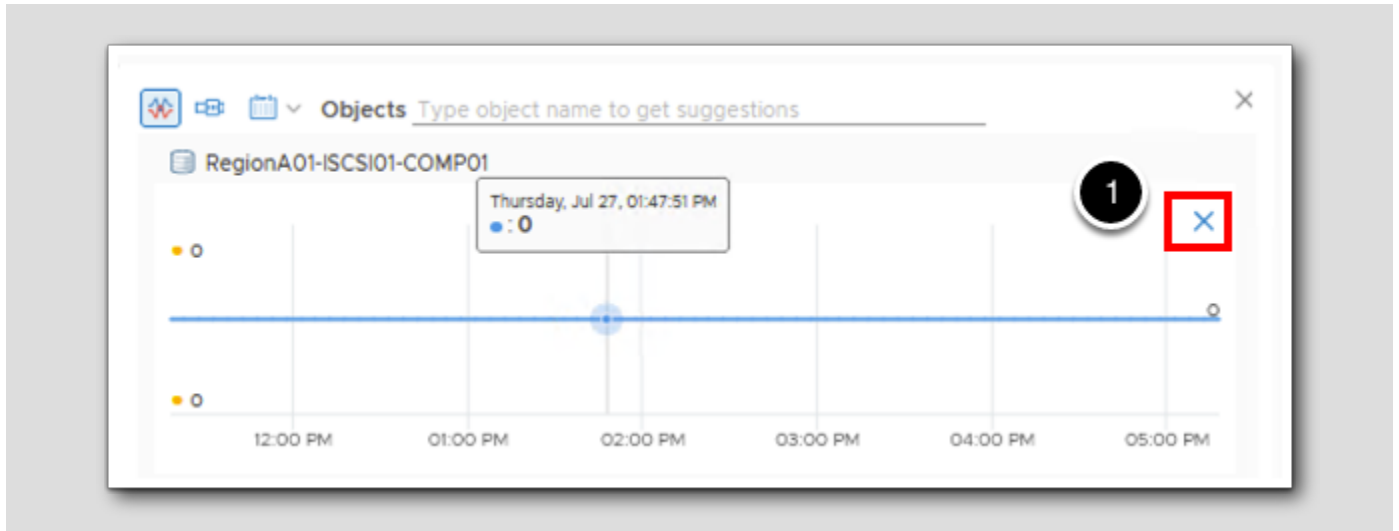
[447]



1. Click on PREVIEW again.

Close the old chart

[448]



1. Hover your mouse over the RegionA01-ISCSI01-COMP01 chart until the Blue X appears, click the blue X.

Re-select RegionA01

[449]

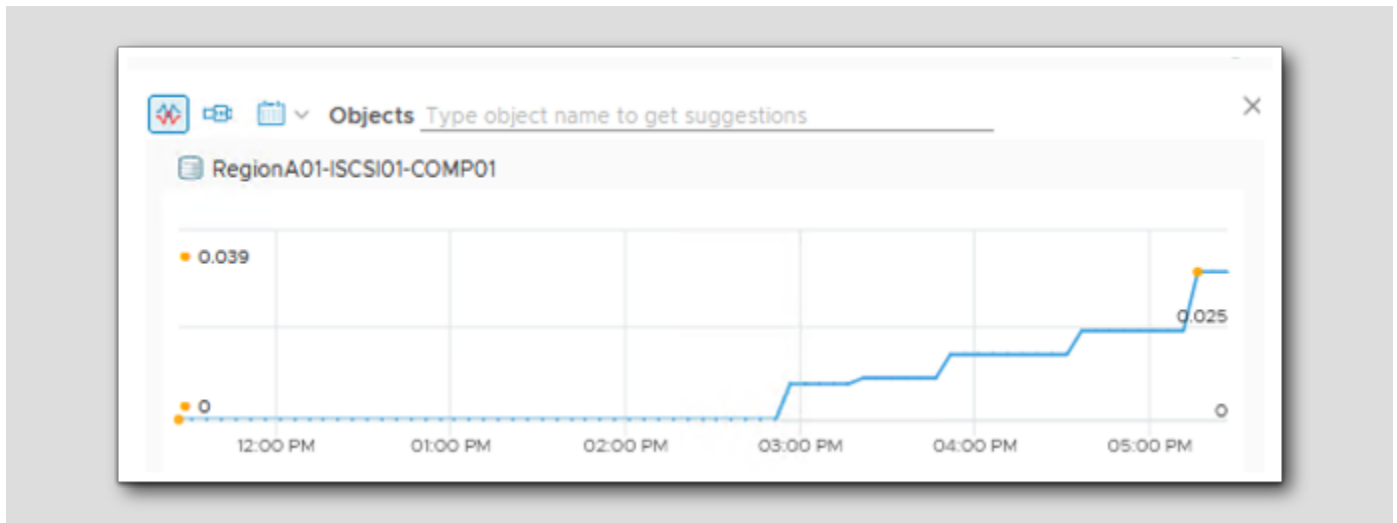


1. Click in the **Objects** line.
2. Click on **RegionA01-ISCSI01-COMP01**.

If you do not see RegionA01-ISCSI01-COMP01, close out of the preview window and re-open the preview window.

We should see the Super Metric working

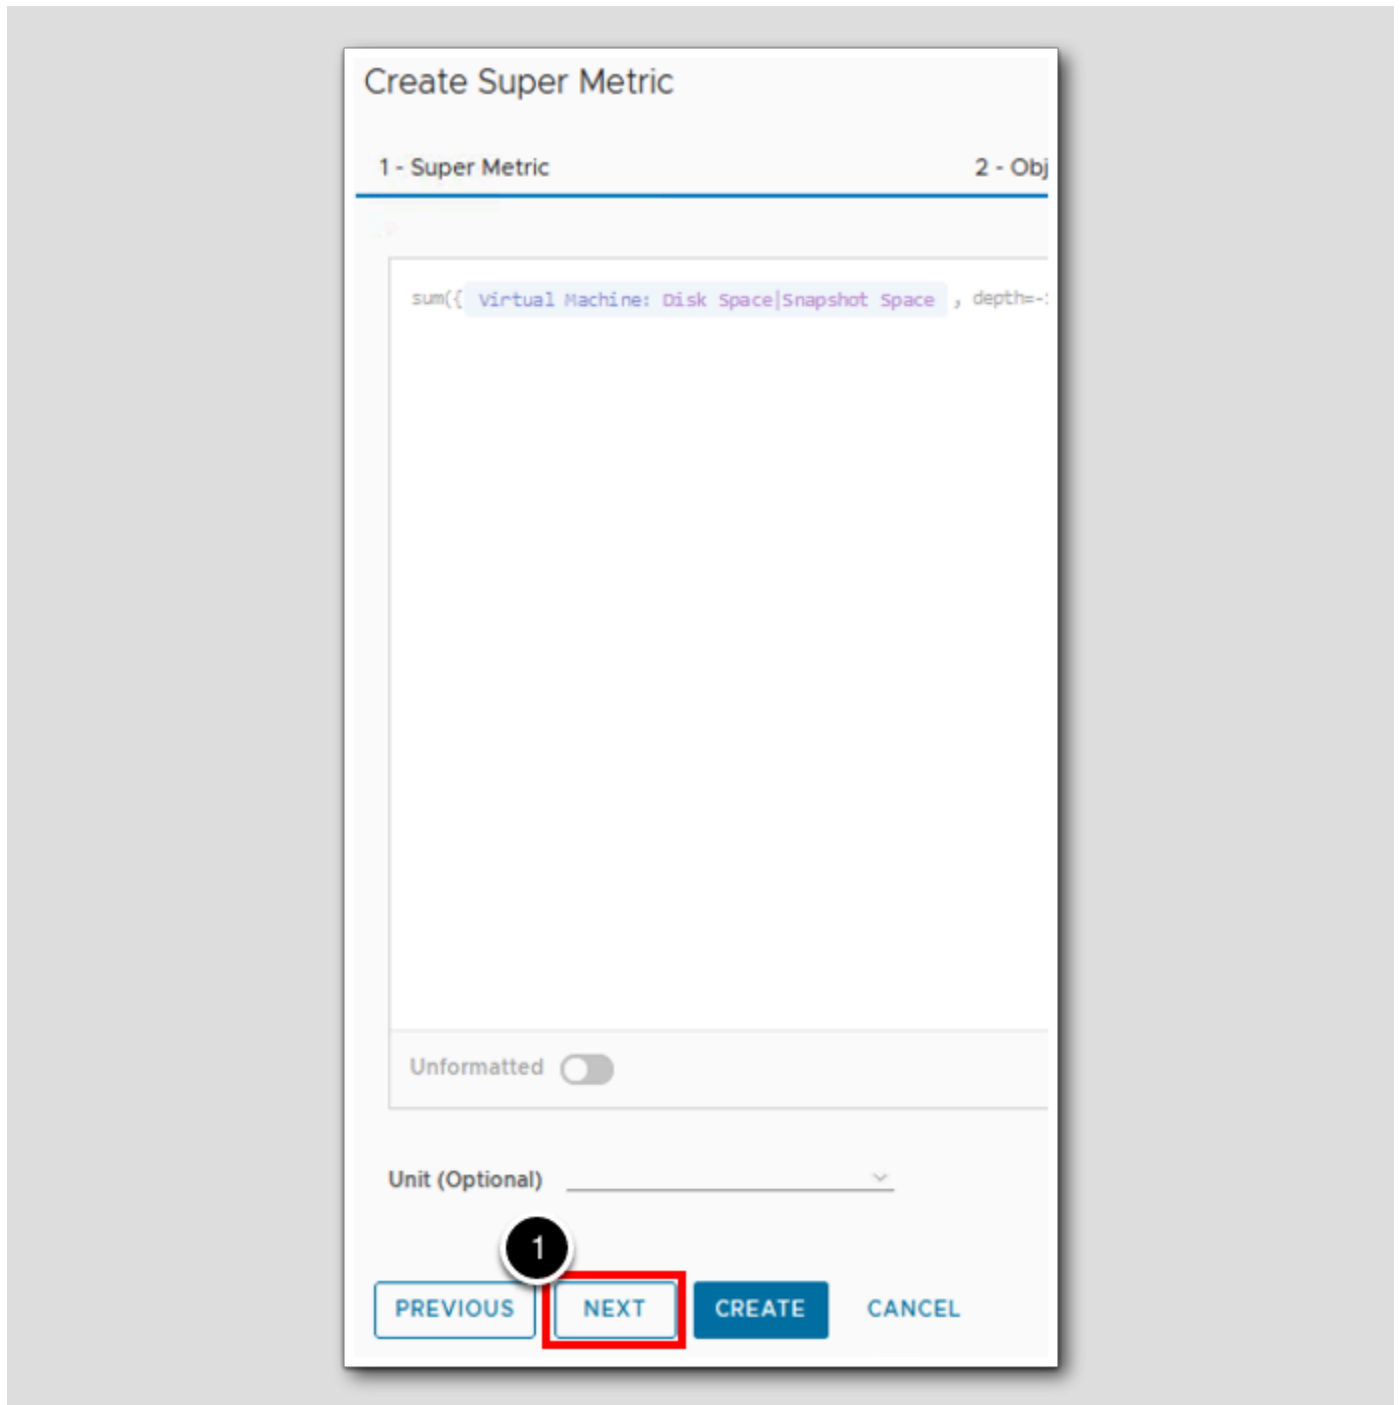
[450]



Note the chart's graph starts to grow when you added that snapshot

## On to Policies

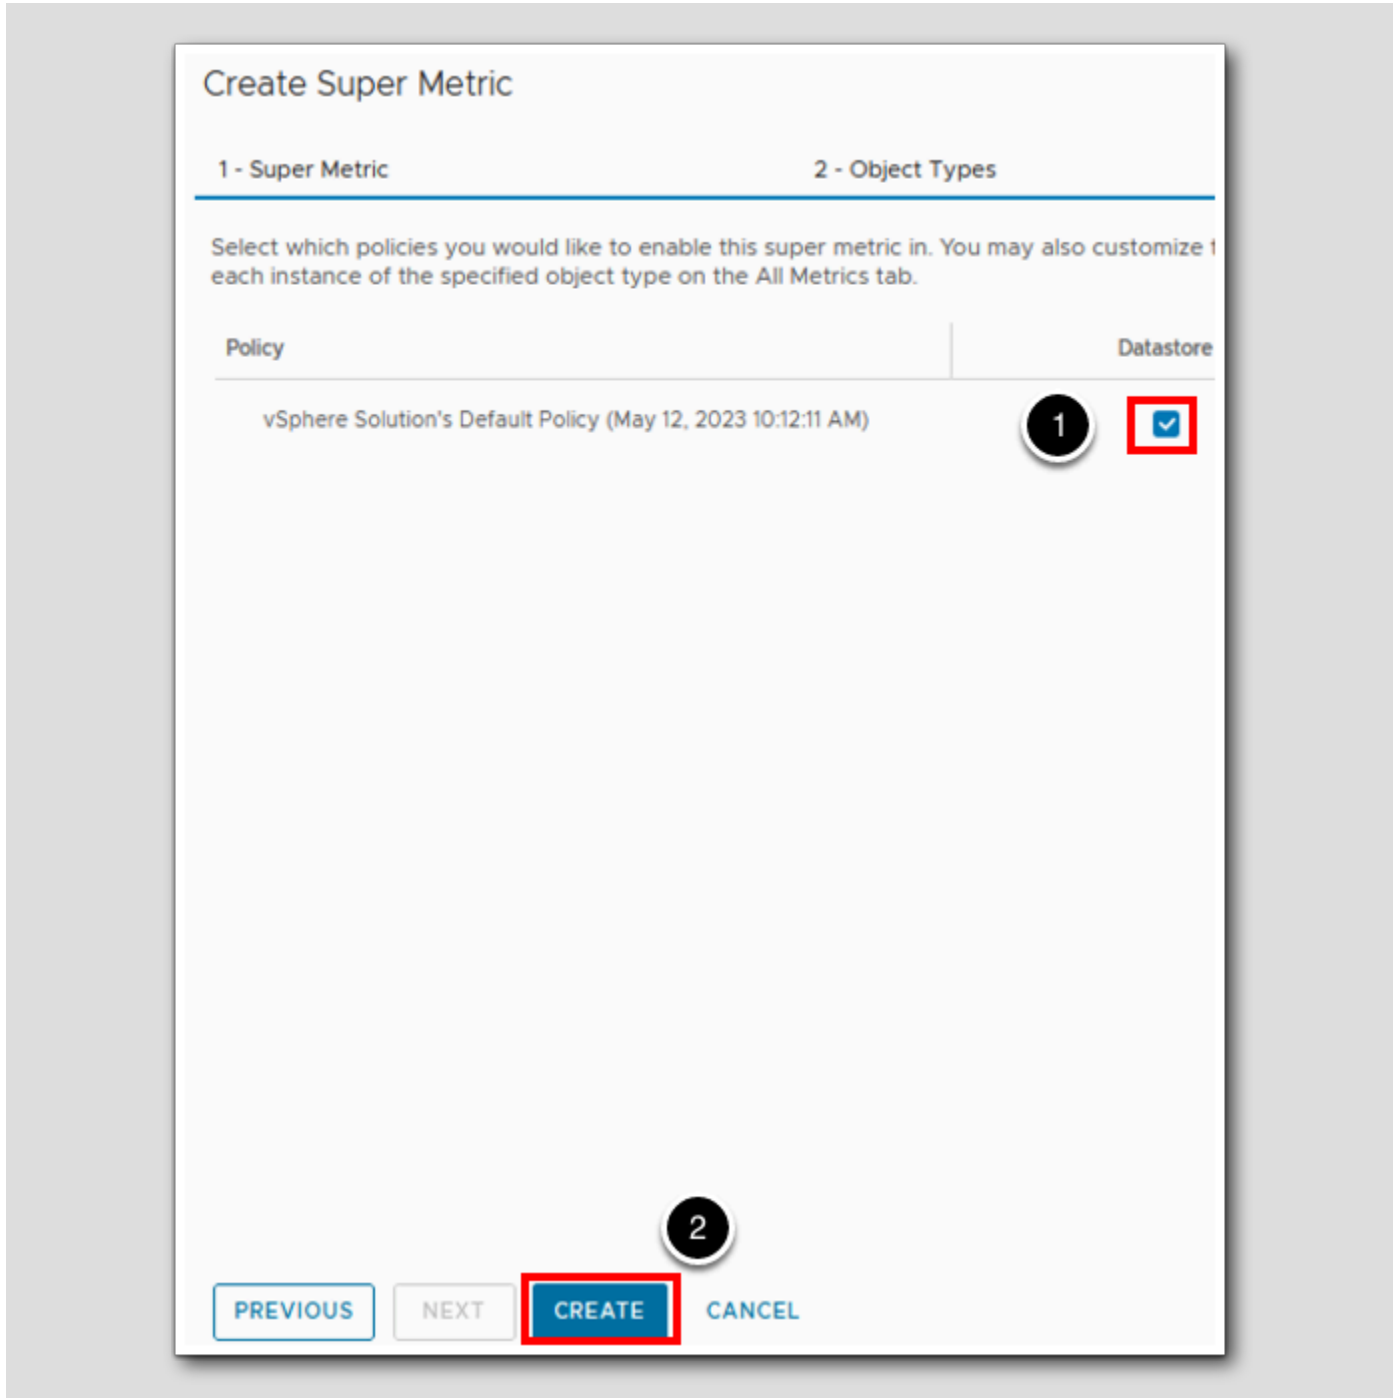
[451]



1. Click NEXT.

Select the Policy and Finish

[452]



Just like in the last lesson, we need to enable the super metric in one or more policies if we want it to actually be calculated and then we can finish the process.

1. Check the **box** to enable the metric on Datastore object types in the HOL Policy policy.
2. Click the **CREATE** button to complete the wizard.

## Lesson End

[453]

In this lesson we learned how to use the THIS function in a storage Super Metric.

## Handling Sets of Data Points in a Super Metric Formula

[454]

This topic confounds many people when they first start creating super metric formulas so it's worth spending some time to understand when you might run into this issue and how you can work around it. If you remember back in the lesson where we created our first super metric, there was a discussion about super metric functions and it was stated that the list of available functions includes looping functions (avg, combine, count, max, min and sum) that work on more than one input value and can return either a single value or a set of values depending on the formula syntax. The topic of this lesson centers on that notion of "either a single value or a set of values" depending on the syntax.

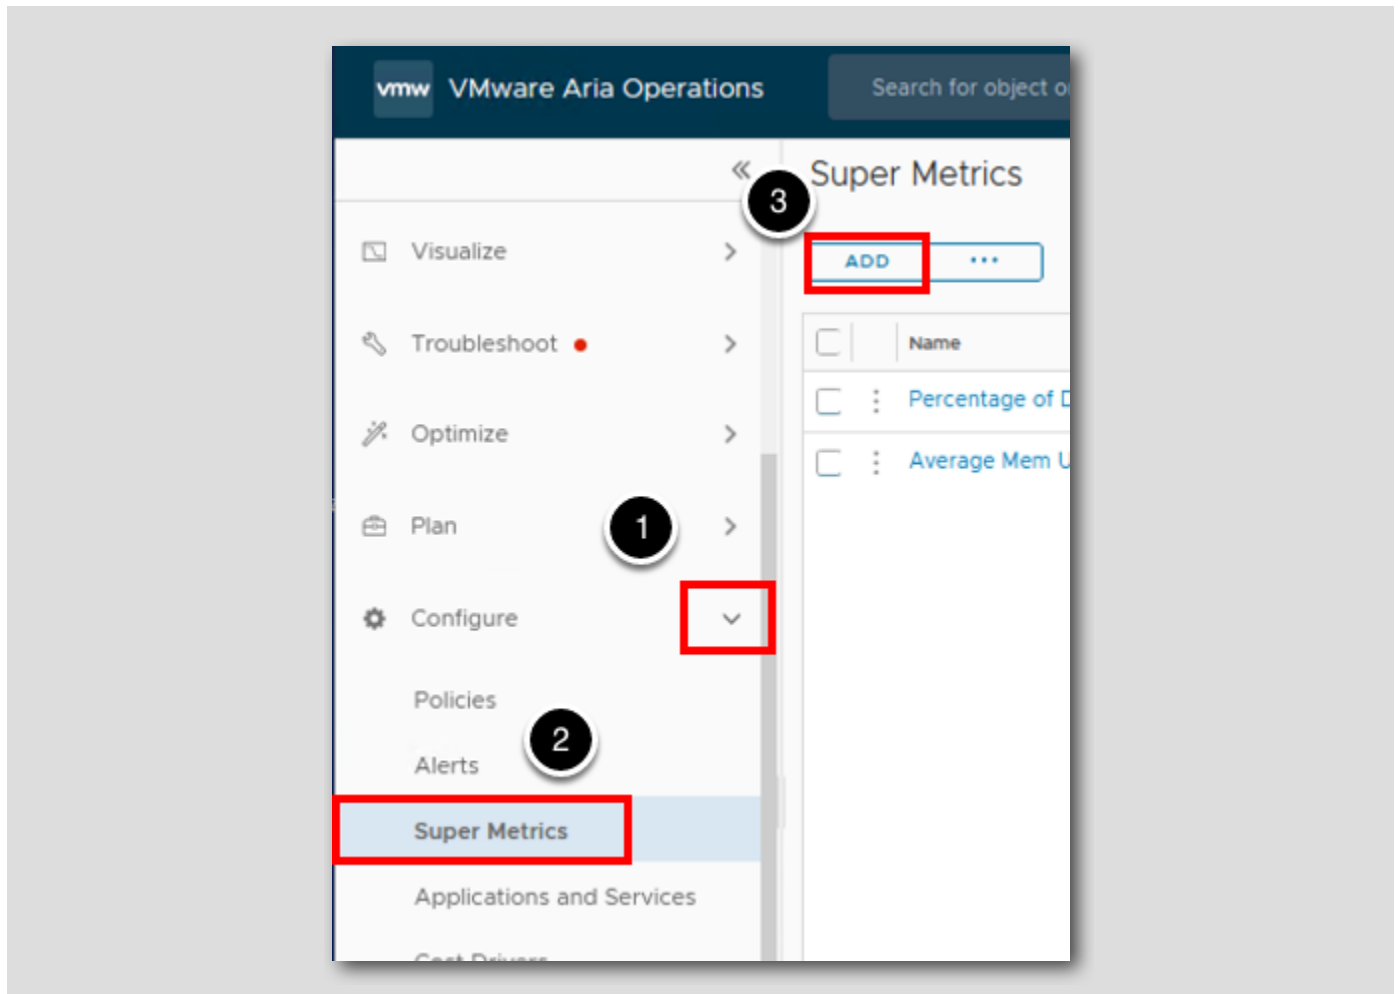
If you think back to the discussion about hierarchies in Aria Operations, you will recall for example that in the vSphere Hosts and Clusters hierarchy, virtual machines are children of hosts and that a virtual machine's parent is a host. We understand that a host can have one or more VMs as children but that a VM can only have a single host as its parent. But if we think about the relationship between hosts and datastores, we realize that a host can have one or more datastores as descendants and a datastore can have one or more hosts as ascendants. We know this because we understand vSphere enough to realize that. However, Aria Operations really has no way to know whether relationships between particular objects or object types are one-to-one or one-to-many. This is the thing that can cause confusion when creating a super metric formula until you understand the concept and how to work with it.

In this lesson we will explore this concept by creating a super metric that can be applied to virtual machine. It will calculate the percentage of a vSphere cluster's usable memory that the VM is using. For example, if a cluster has 200 GB of usable memory and a VM in that cluster was demanding 4 GB of memory, our value should be  $4/200*100$  (to make the ratio into a percentage). The assignment will require us to use some concepts that we covered in the previous lessons and will address the issue discussed above.



## Launch the Super Metric Wizard

[455]



1. Expand Configure.
2. Click on Super Metrics.
3. Click ADD.

### Name the Super Metric

**Create Super Metric**

1 - Super Metric      2 - Object Ty

Name: **1** VM Memory Usage As a Percentage of Cluster Memory

Description (Optional): **2** VM's percentage of usable cluster memory

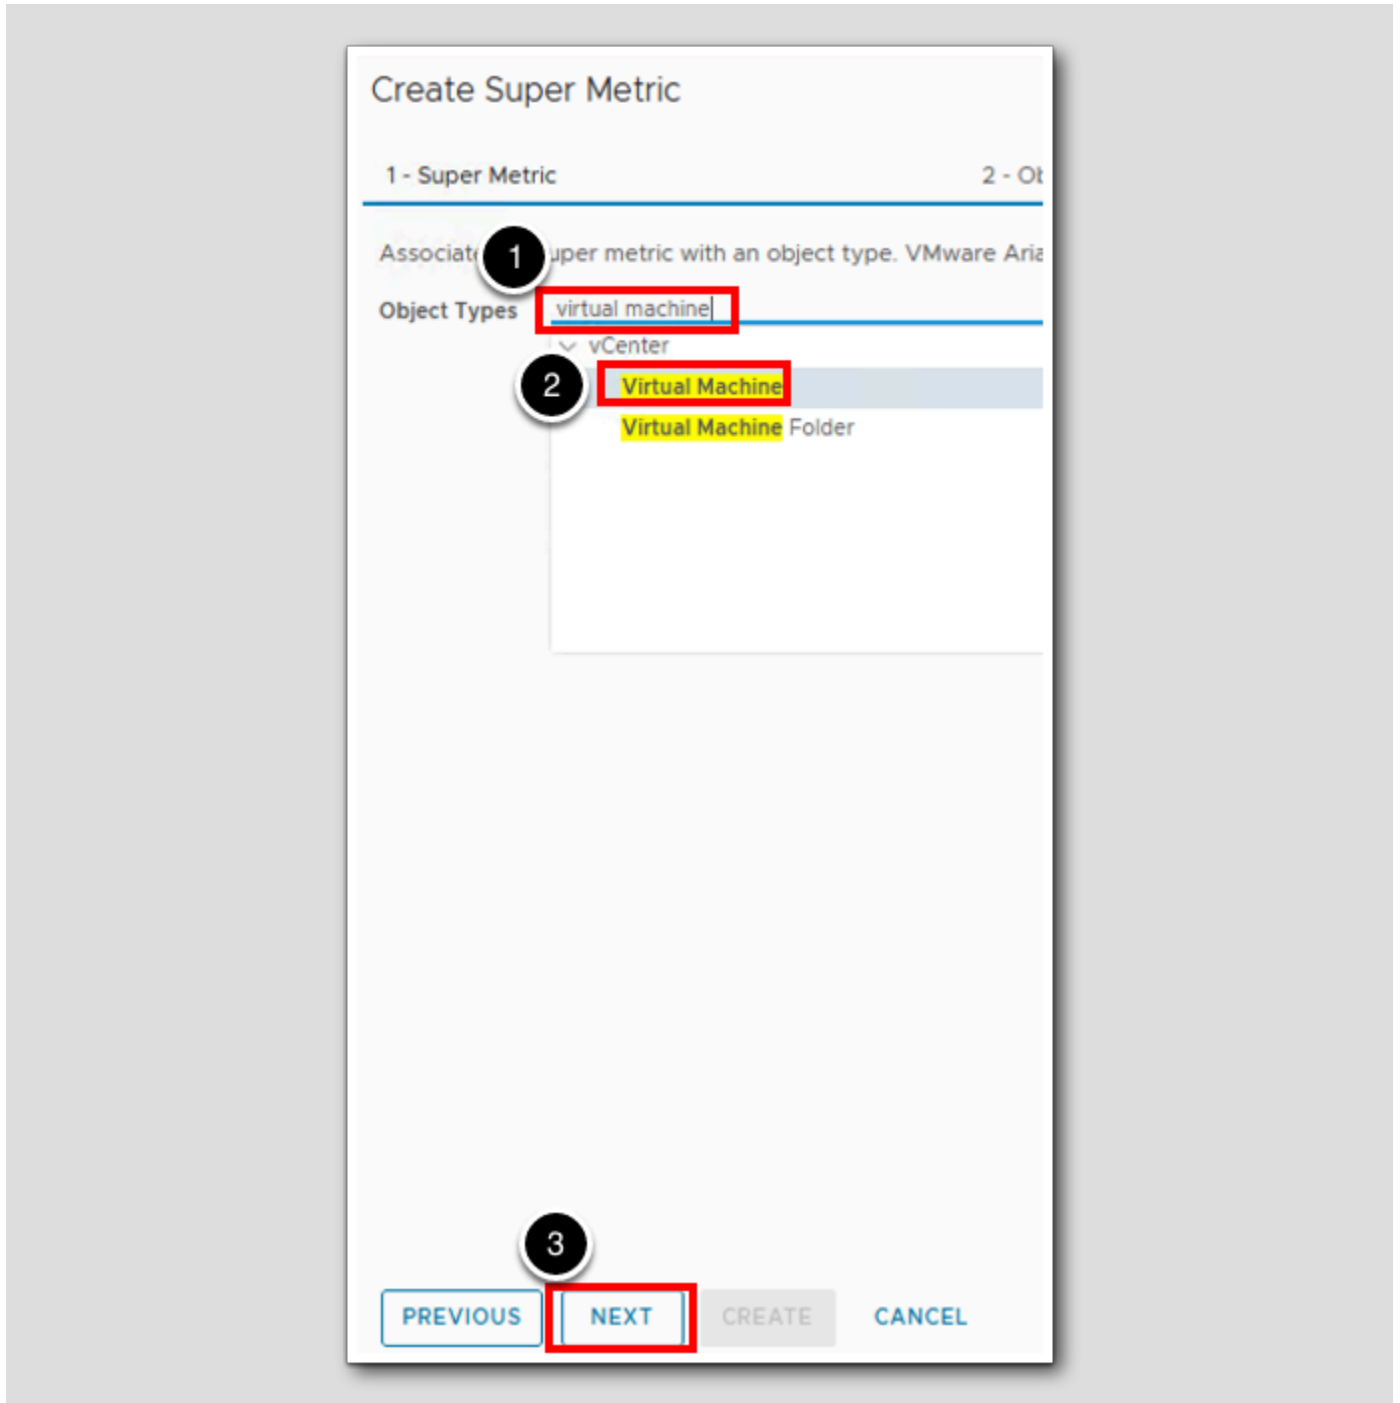
**3**

PREVIOUS   **NEXT**   CREATE   CANCEL

1. In the Name field, type VM Memory Usage As a Percentage of Cluster Memory (%).
2. In the Description field, type VM's percentage of usable cluster memory.
3. Click **NEXT**.

Assign to an Object Type

[457]

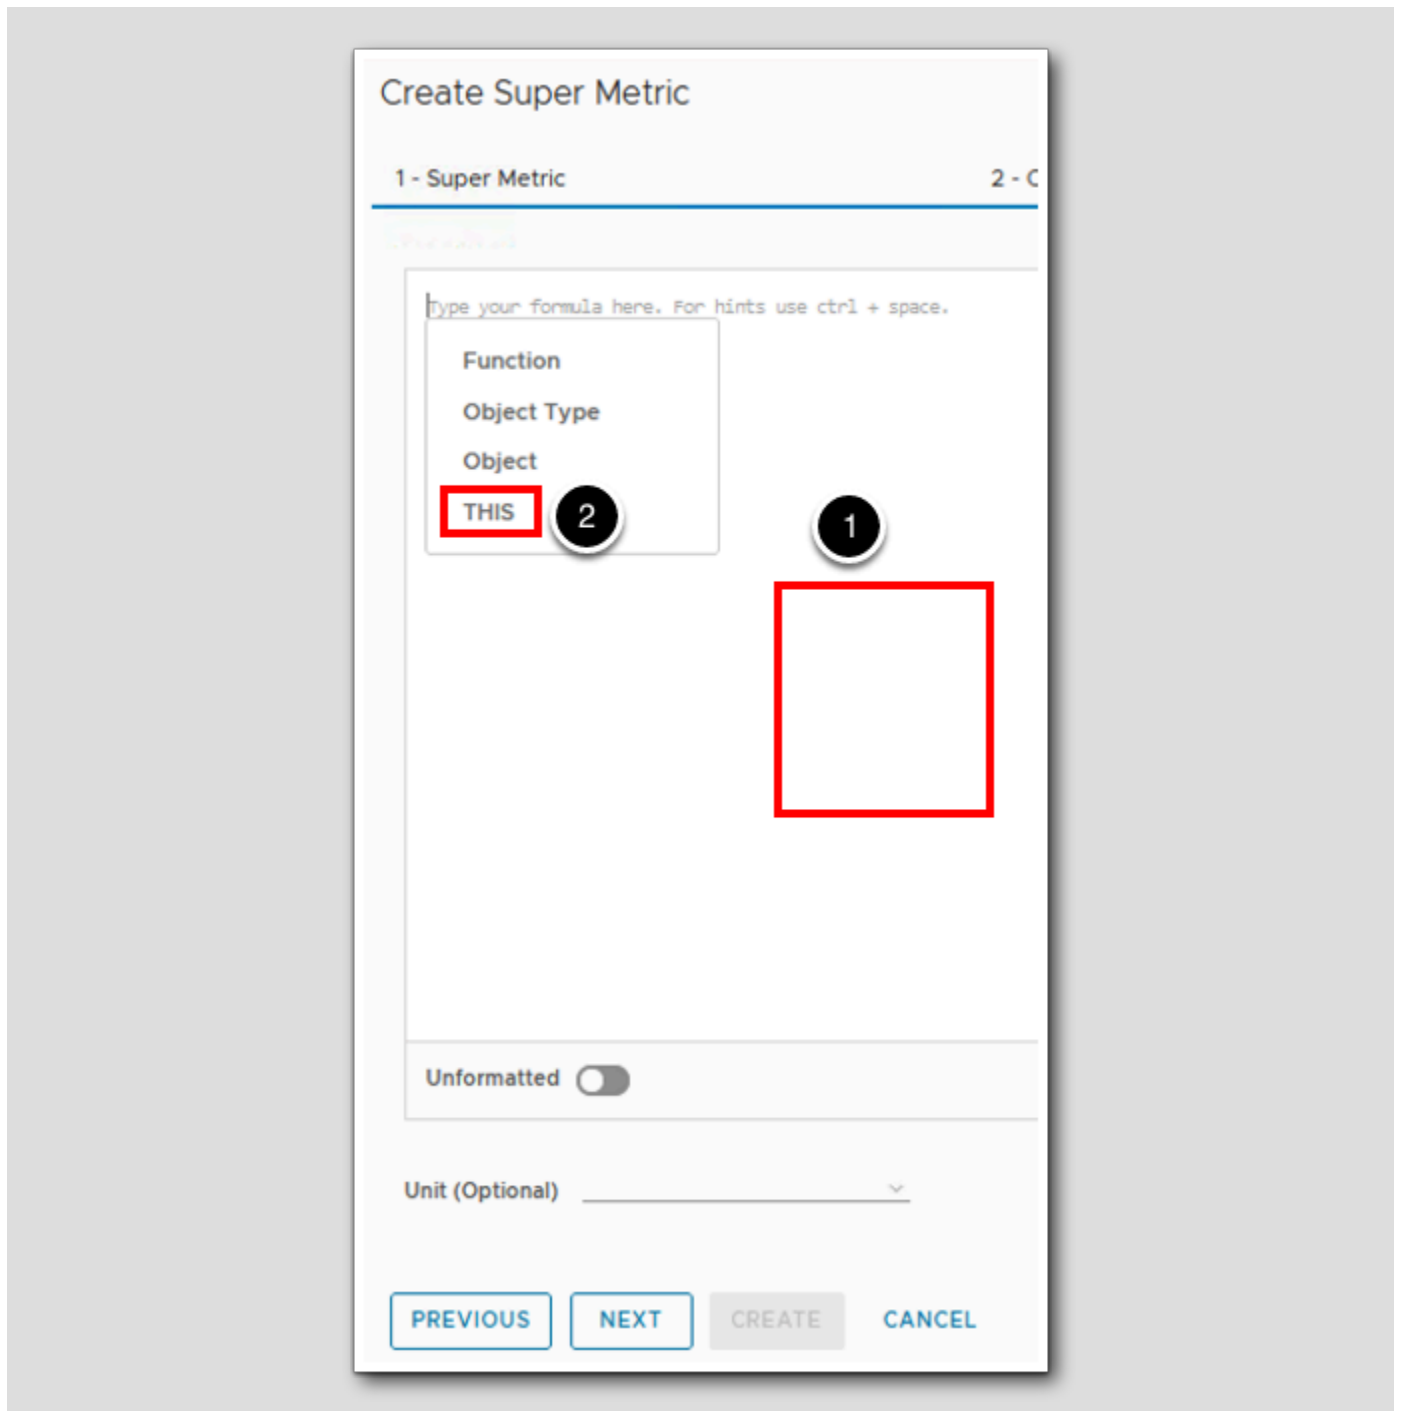


1. In the **Object Types** search field, type **virtual machine**.
2. Single click on **Virtual Machine**.
3. Click **NEXT**.

## Formula

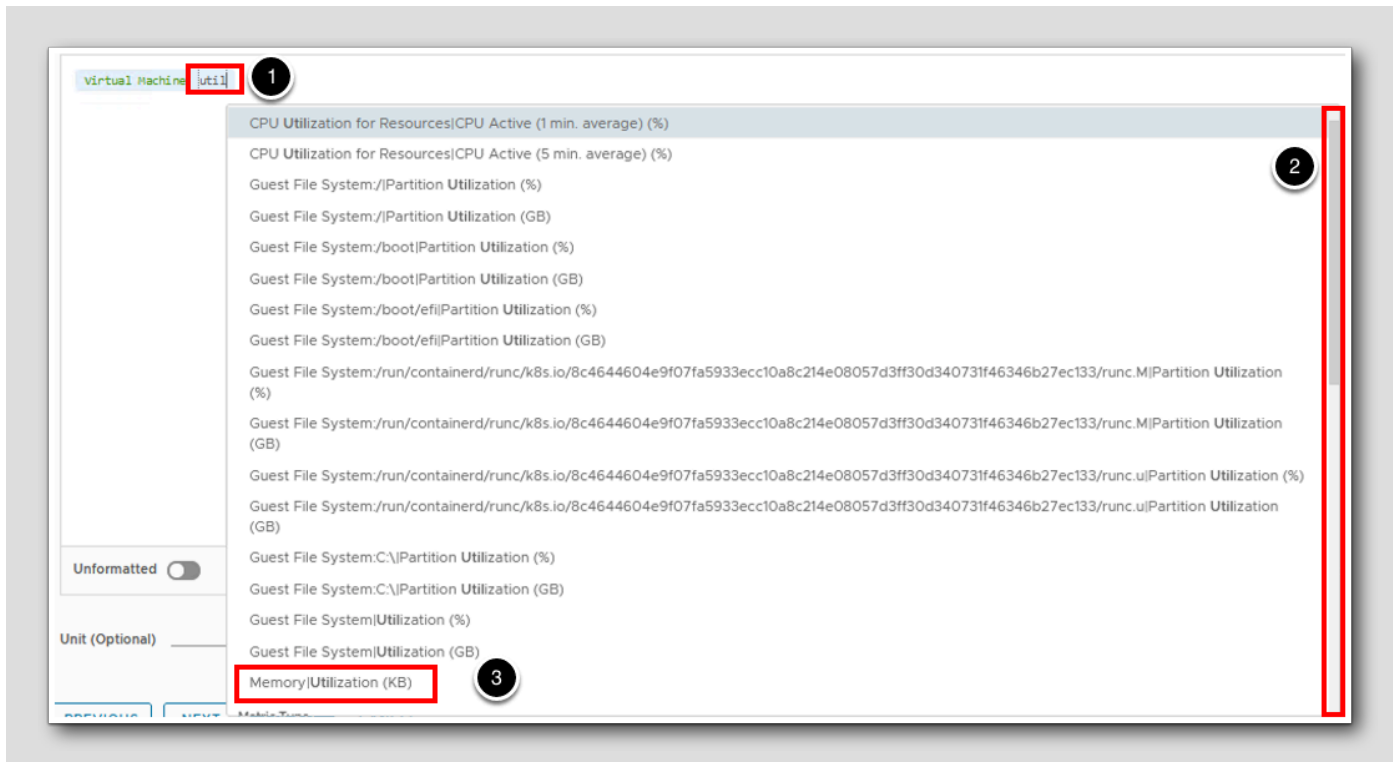
[458]

Since the super metric will be applied to virtual machines and the first metric (the numerator) in the formula is the vm's memory demand we will again use the **THIS** button here.



1. Click anywhere in the empty Formula box.
2. Click the THIS button.

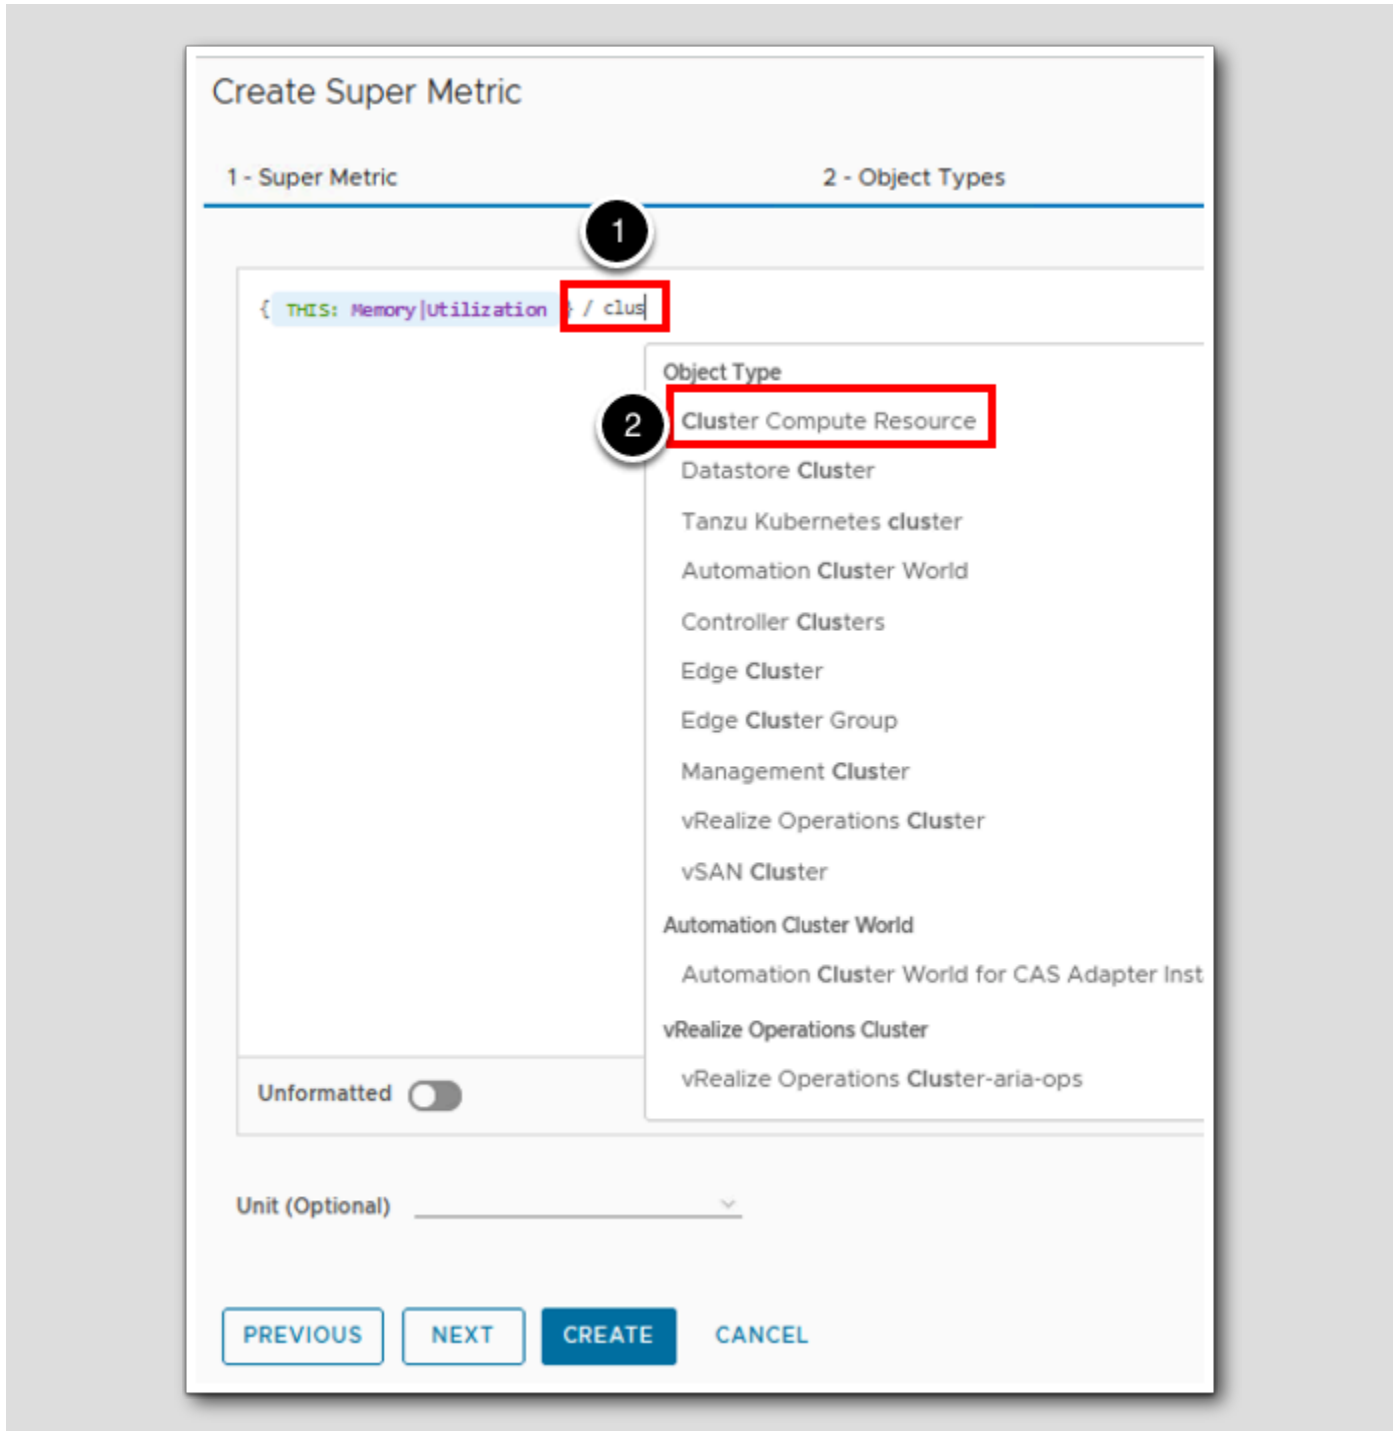
## Add the Metric



Since we want the VM's memory utilization metric,

1. Type util after the green Virtual Machine.
2. Use the scroll bar to scroll to the bottom of the list. Ensure that you are in the list of Metrics and not Metric Types
3. Click Memory|Utilization (KB) to add it to the formula.

### Choose the vSphere Cluster Object Type





1. At the end of the formula line, type a space then / then another space then **clus** (the spaces are not necessary but make the formula easier to read).
2. Click **Cluster Compute Resource** to select the object type.

## Select the Metric

[461]

The screenshot shows the 'Create Super Metric' wizard with three steps: 1 - Super Metric, 2 - Object Types, and 3 - Formula. The current step is '2 - Object Types'. The formula field contains the text: `{ THIS: Memory|utilization } / Cluster Compute Resource|usable`. A red box highlights the word 'usable' in the formula. A dropdown menu is open, showing a list of metrics. A red box highlights the metric 'Memory|Usable Memory (KB)'. A circled '1' is above the formula field, and a circled '2' is above the dropdown menu. At the bottom, there are buttons for 'PREVIOUS', 'NEXT', 'CREATE', and 'CANCEL'. There is also an 'Unformatted' toggle switch and a 'Unit (Optional)' dropdown menu.

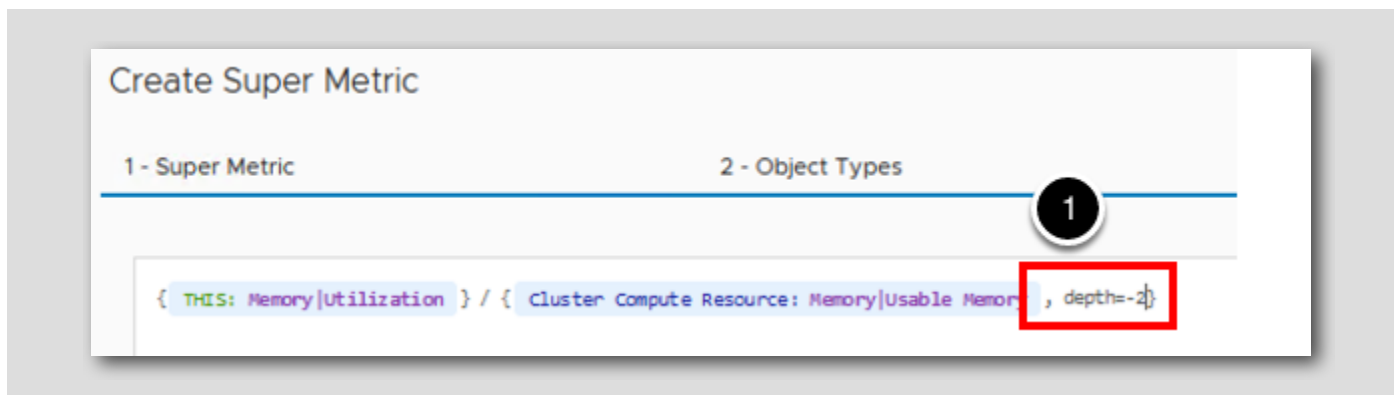
Be sure to select the correct metric here. There are a lot of similarly named that are returned by the filter.

1. On the formula line, type **usable** after **Cluster Compute Resource**.
2. Click **Memory|Usable Memory (KB)** in the **Metric** section to add it to the formula.

## Adjust the Formula Depth

[462]

Remembering what we learned earlier about the depth parameter and knowing that vSphere clusters are two levels above VMs in the hierarchy, we need to adjust the value. Remember for the depth parameter, a positive number means look down the hierarchy while a negative number means look up.

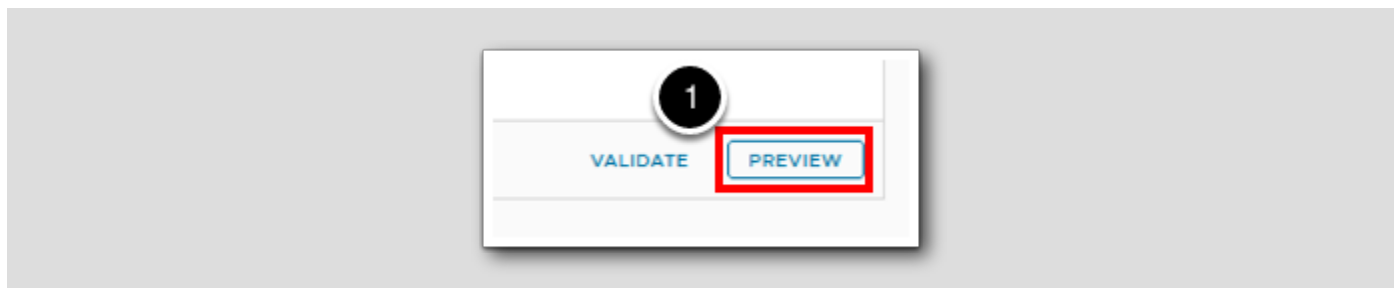


1. On the formula line, change the "1" and in its place, type -2

## Preview the Super Metric

[463]

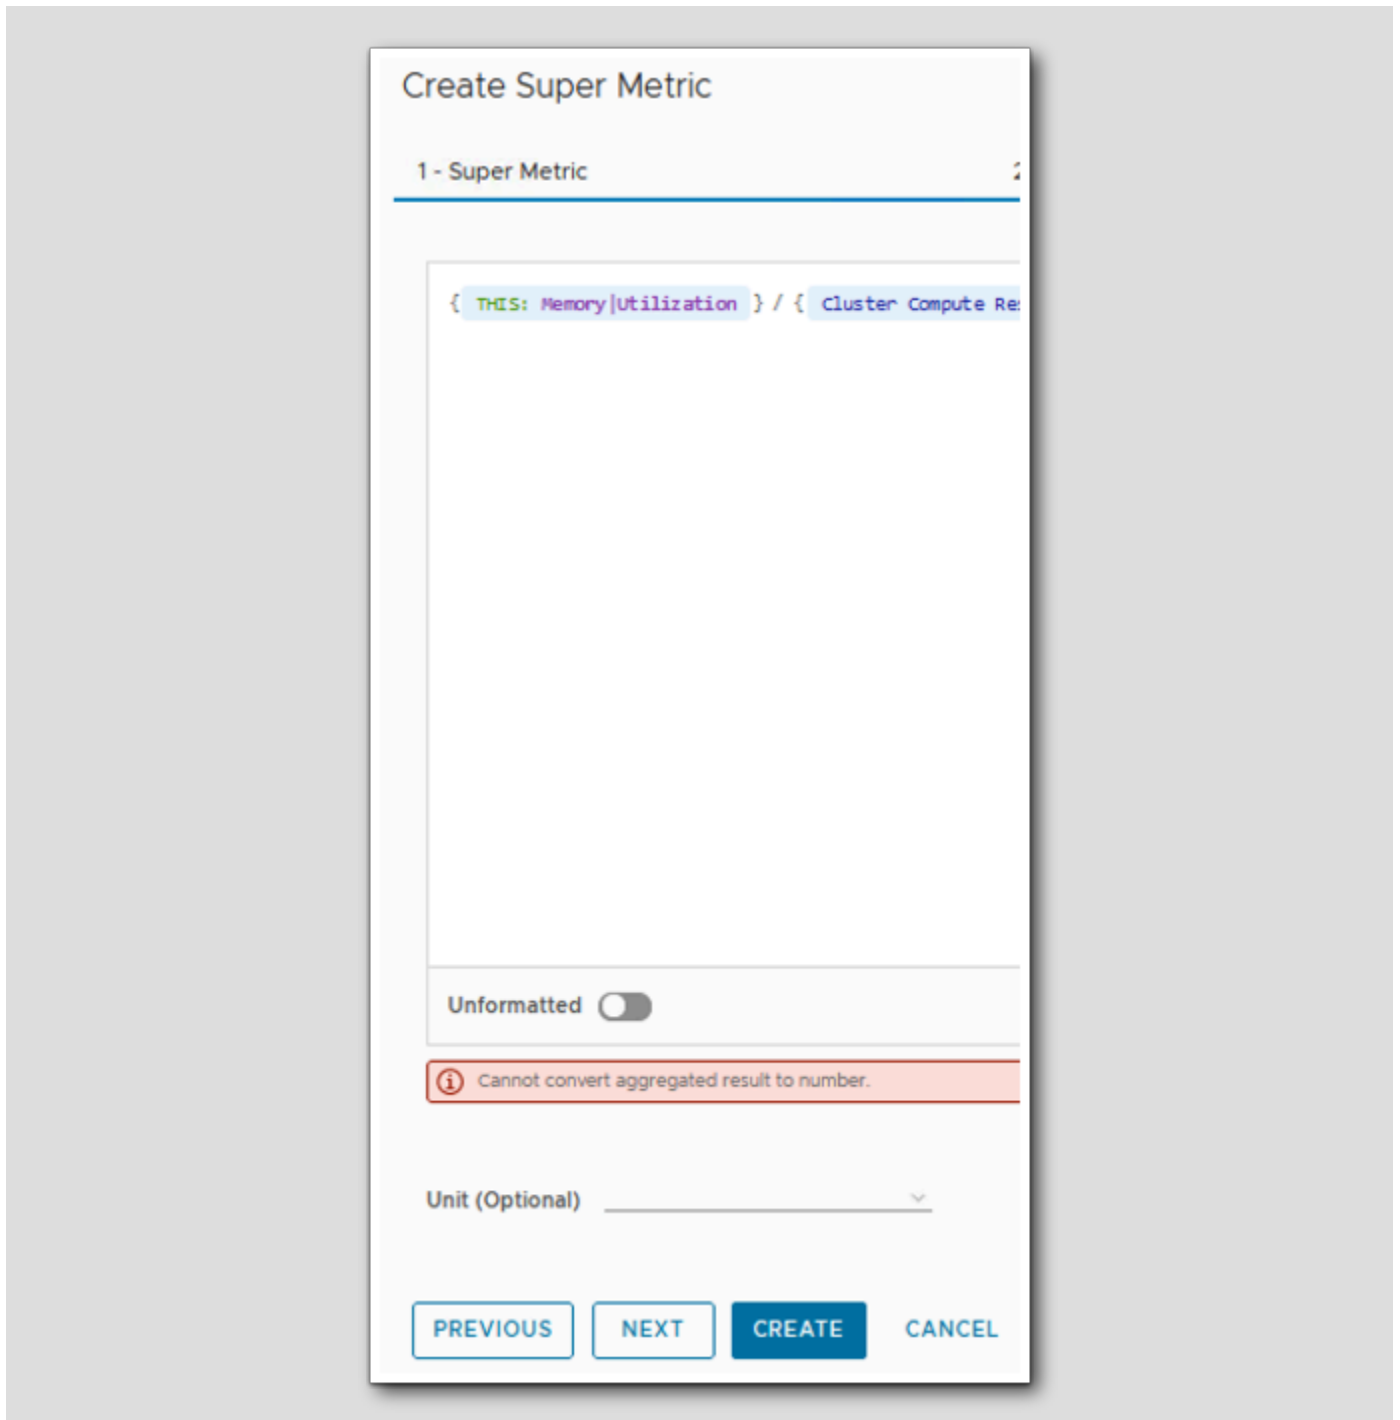
OK. We're done, right? Let's preview the super metric by selecting a virtual machine in our inventory.



1. Click Preview

Uh Oh!

[464]



Uh oh. We got an error - Cannot convert aggregated result to number. This is the issue that was discussed at the beginning of the lesson. Remember that while we know there can only be one cluster as an ascendant (2 levels up) from the VM, Aria Operations doesn't have any way of knowing that. As far as Aria Operations knows, there could be a set of cluster objects that are two levels above the VM.

So how do we handle this? We need to modify the formula using a looping function. If you recall from the beginning of the lesson, it was reiterated that looping functions (avg, combine, count, max, min and sum) work on more than one input value and can return either a single value or set of values depending on the formula syntax. What does that mean in this context? It means we can use many of those looping functions to convert the results of the cluster portion of the formula to a single value. Essentially we can tell Aria Operations to take the avg or min or max or sum of the values from all clusters above the VM and return a single number representing the calculation. What is the average or minimum or maximum or sum of a single number? It's that number.

In this case, we will use the max function (to find the maximum value from a set of one).

Revise the Formula

[465]

Create Super Metric

1 - Super Metric      2 - Object Types

1      2

```
{ THIS: Memory|Utilization } , max({ Cluster Compute Resource: Memory|Usable Memory , depth=2})
```

Unformatted

Unit (Optional) \_\_\_\_\_

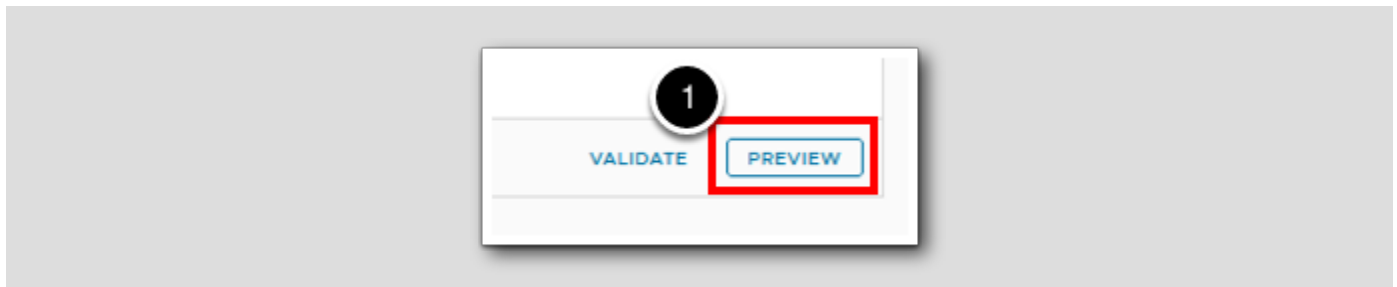
PREVIOUS    NEXT    CREATE    CANCEL

1. Place your cursor just before "{Cluster..." in the formula and type `max(` (don't click max in the list of suggestions or it will add both parenthesis there)
2. Press the **End** key on your keyboard to move to the end of the line and type a closing parentheses )

For reference, there is the completed formula so far: `{This Resource: Memory|Utilization} / max({Cluster Compute Resource: Memory|Usable Memory, depth=-2})`

Click Preview again

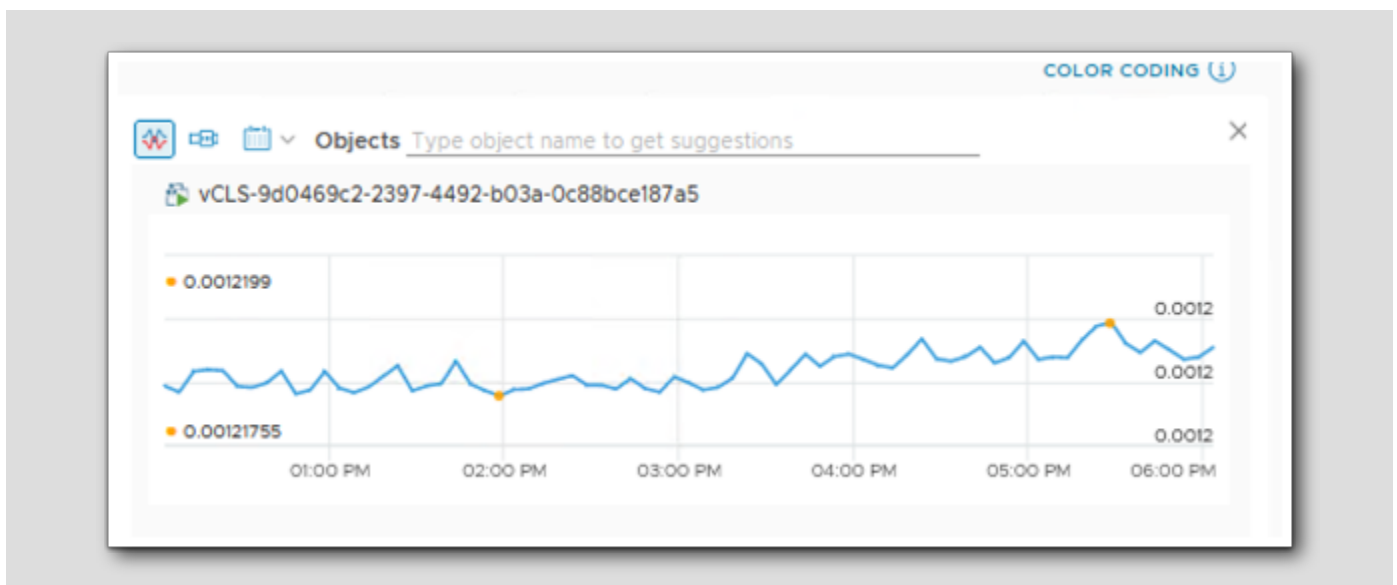
[466]



1. Click **PREVIEW**.

Hey it worked this time!

[467]

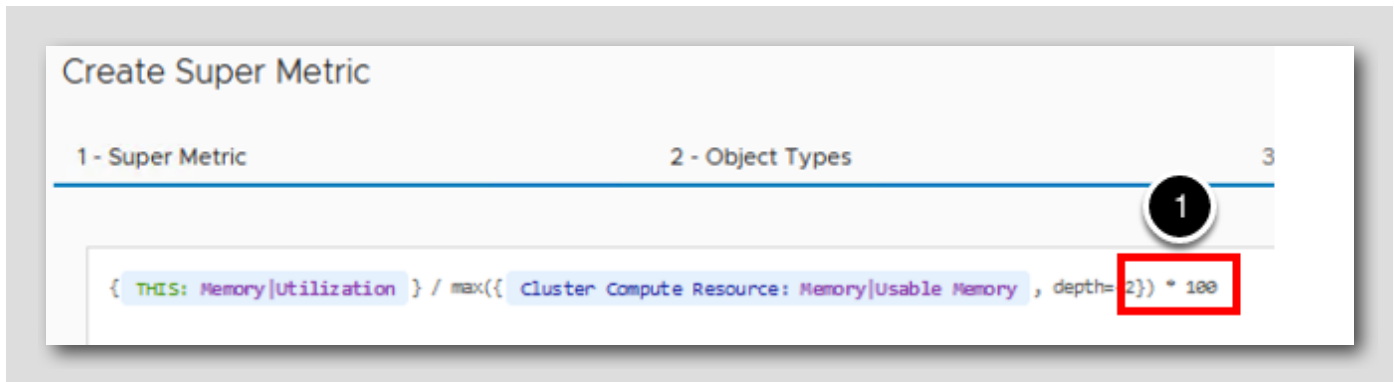


You should not get an error this time and get a populated chart.

## Convert the Ratio to a Percentage

[468]

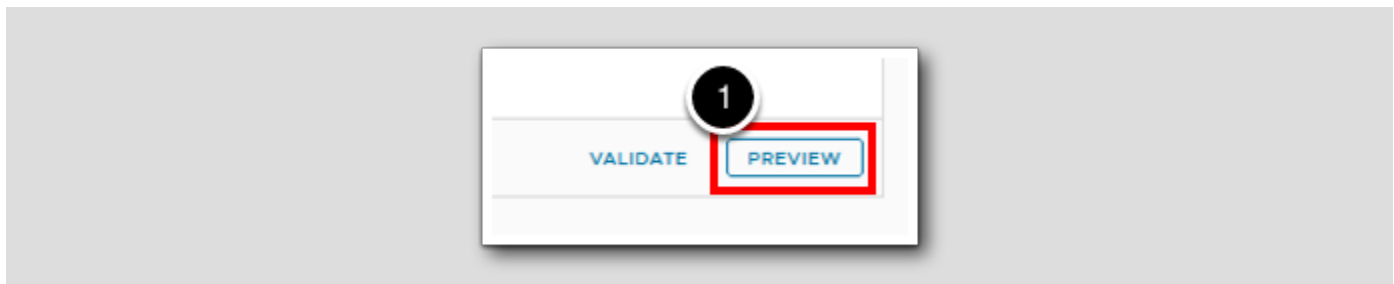
The formula is returning the ratio of vm memory utilization to cluster memory capacity. But the assignment was to calculate the value as a percentage.



1. At the end of the formula type space then \* then space then 100

## Preview one more time

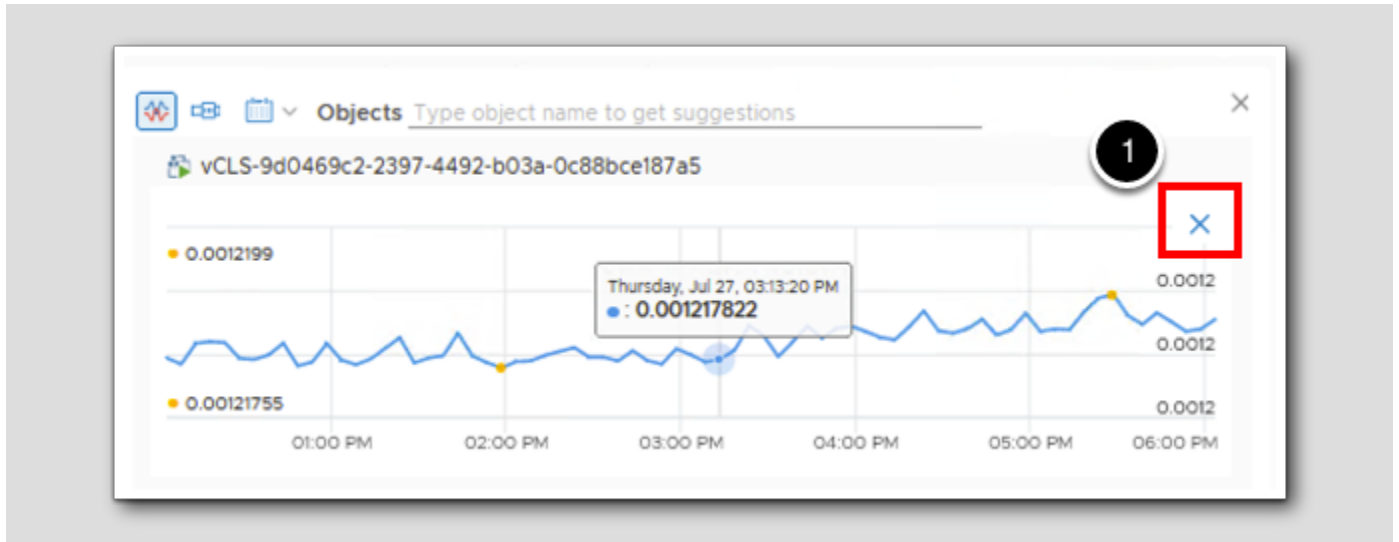
[469]



1. Click PREVIEW.

### Remove and re-add the chart

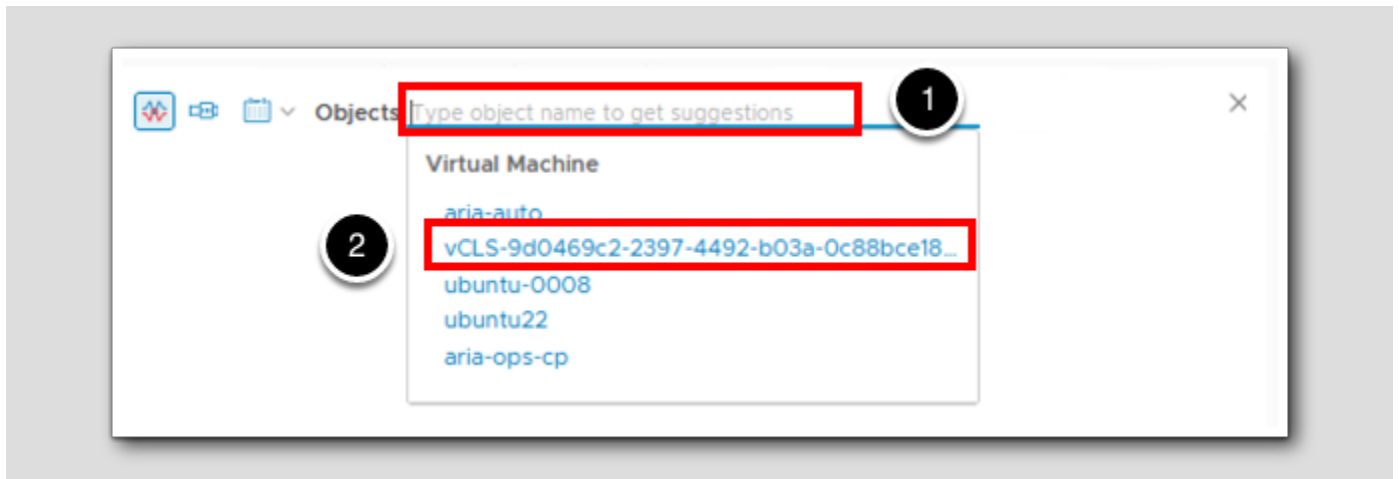
[470]



1. Hover your mouse over the chart until the Blue X appears, click the blue X.

### Add the object back in

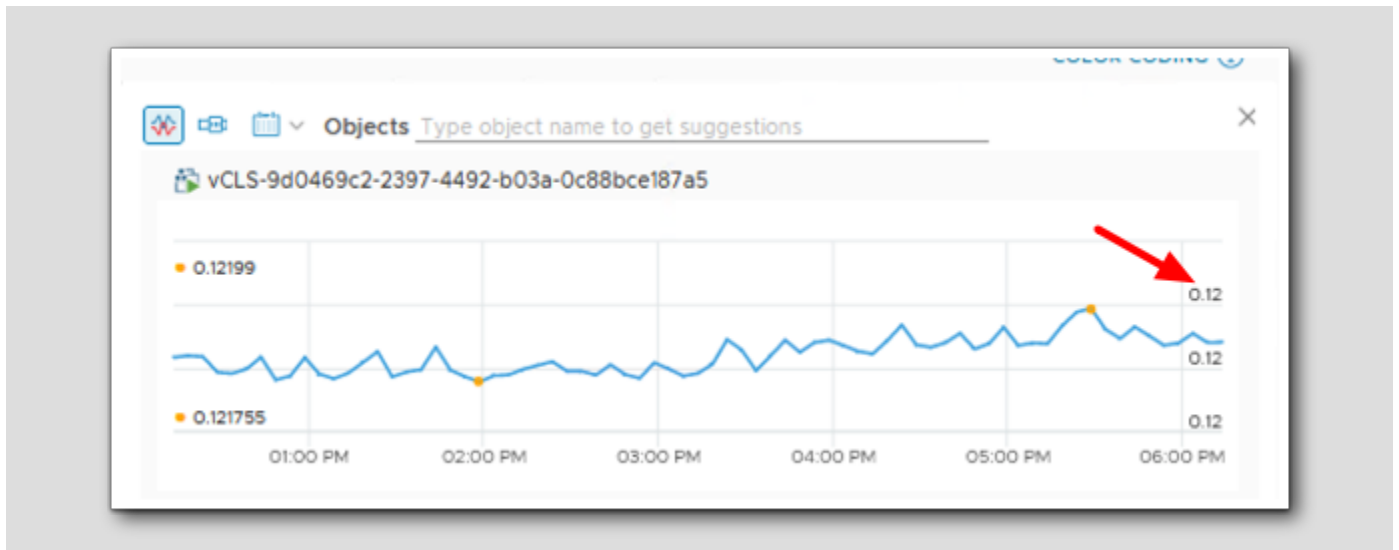
[471]



1. Click in the Objects line.
2. Select the vCLS-... virtual machine.



The Chart now shows %



Notice the Y axis values have changed to the desired scale.

## On to the Policy

[473]

Create Super Metric

1 - Super Metric 2 - Object

```
{ THIS: Memory|utilization } / max({ Cluster Compute Resource
```

Unformatted

Unit (Optional) \_\_\_\_\_

PREVIOUS NEXT CREATE CANCEL

1. Click NEXT.

## Select the Policy and Finish

[474]

Just like in the previous lessons, we need to enable the super metric in one or more policies if we want it to actually be calculated and then we can finish the process.



## Lesson End

[475]

In this lesson we learned how to handle sets of data points in a Super Metric formula and how to navigate the error Cannot convert aggregated result to number using the MAX function.

## Using String Operators and the "Where" Clause in a Super Metric Formula

[476]

Super metrics can also include some logic in the formula. In this lesson we will look at using the "where" clause and a string operator to evaluate a VM property (the guest OS).

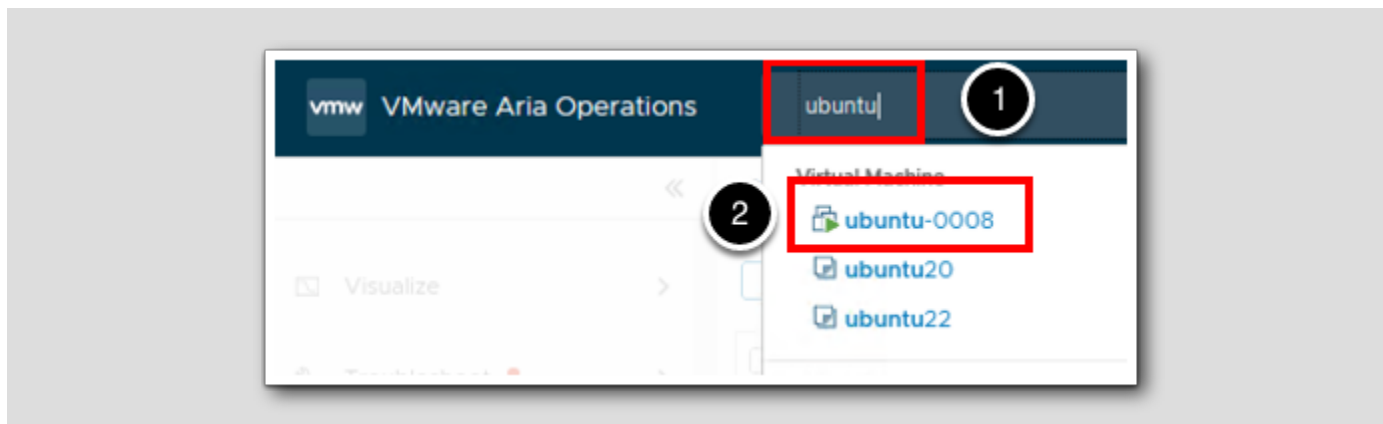
The task this time is to determine the total number of VMs in our datacenter that are running some variant of the ubuntu operating system.

The following string operators are available for use in a super metric formula. Note that string operators are valid only when used in a "where" clause to evaluate whether or not the specified text does or does not exist in the string.

## Examine the Guest OS Full Name Property

[477]

Let's first take a look at the VM property we are going to use in this super metric formula.

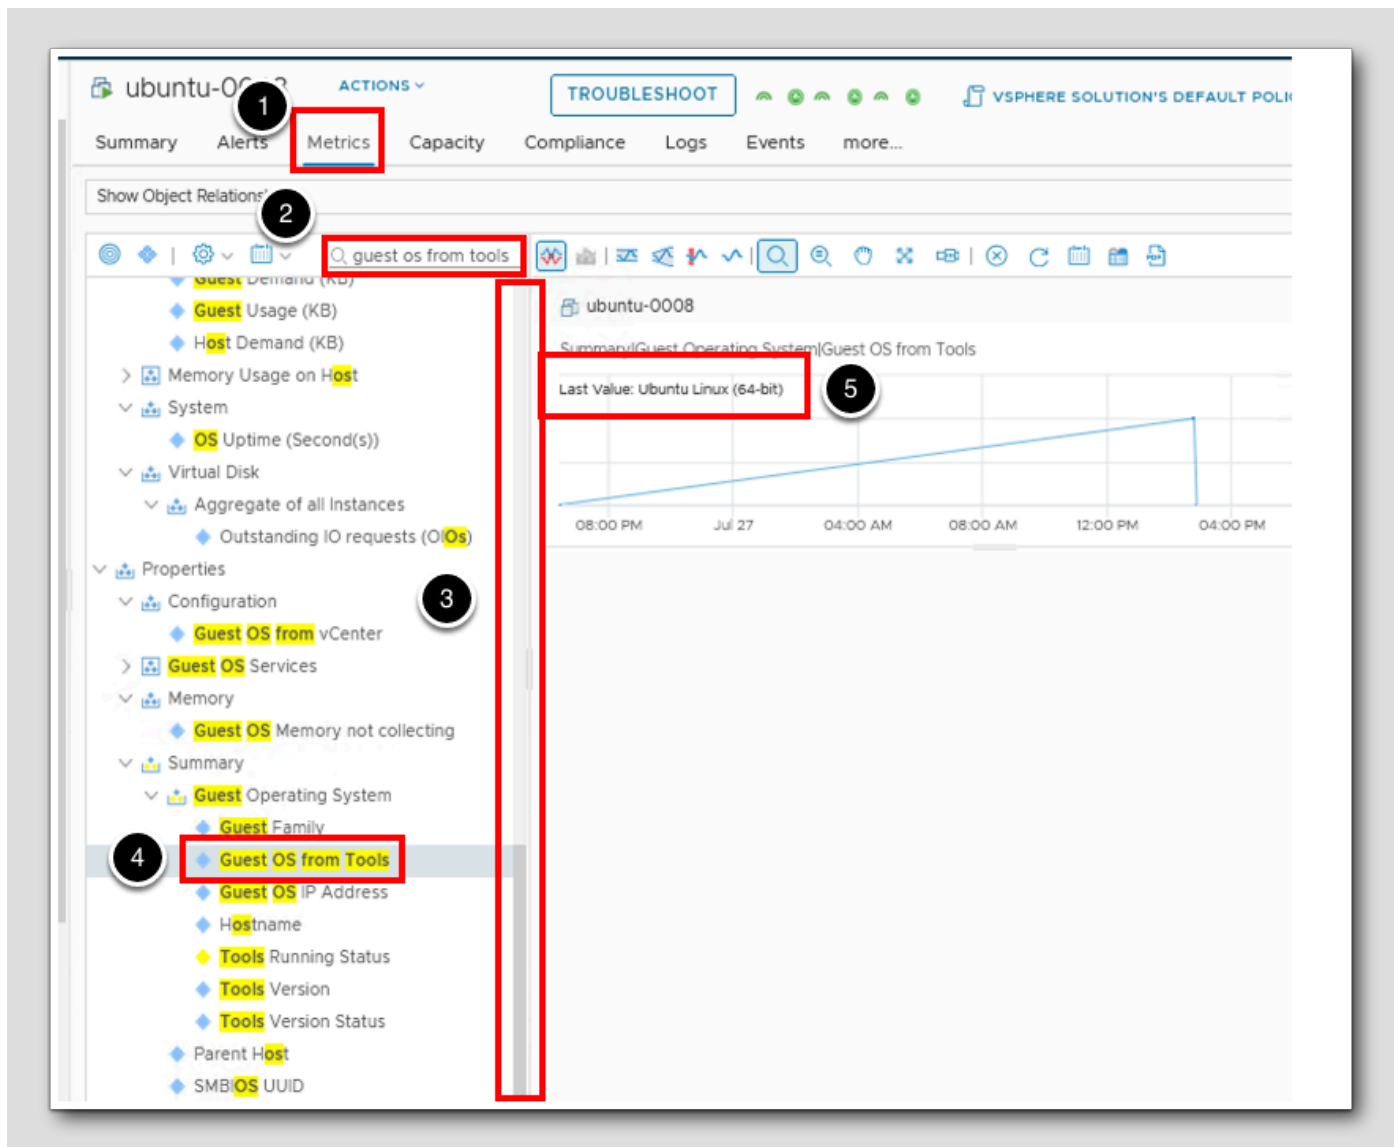


1. In the search box, type **ubuntu**.
2. Click to select the **ubuntu-0008** Virtual Machine.

## Select the OS Name Property

[478]

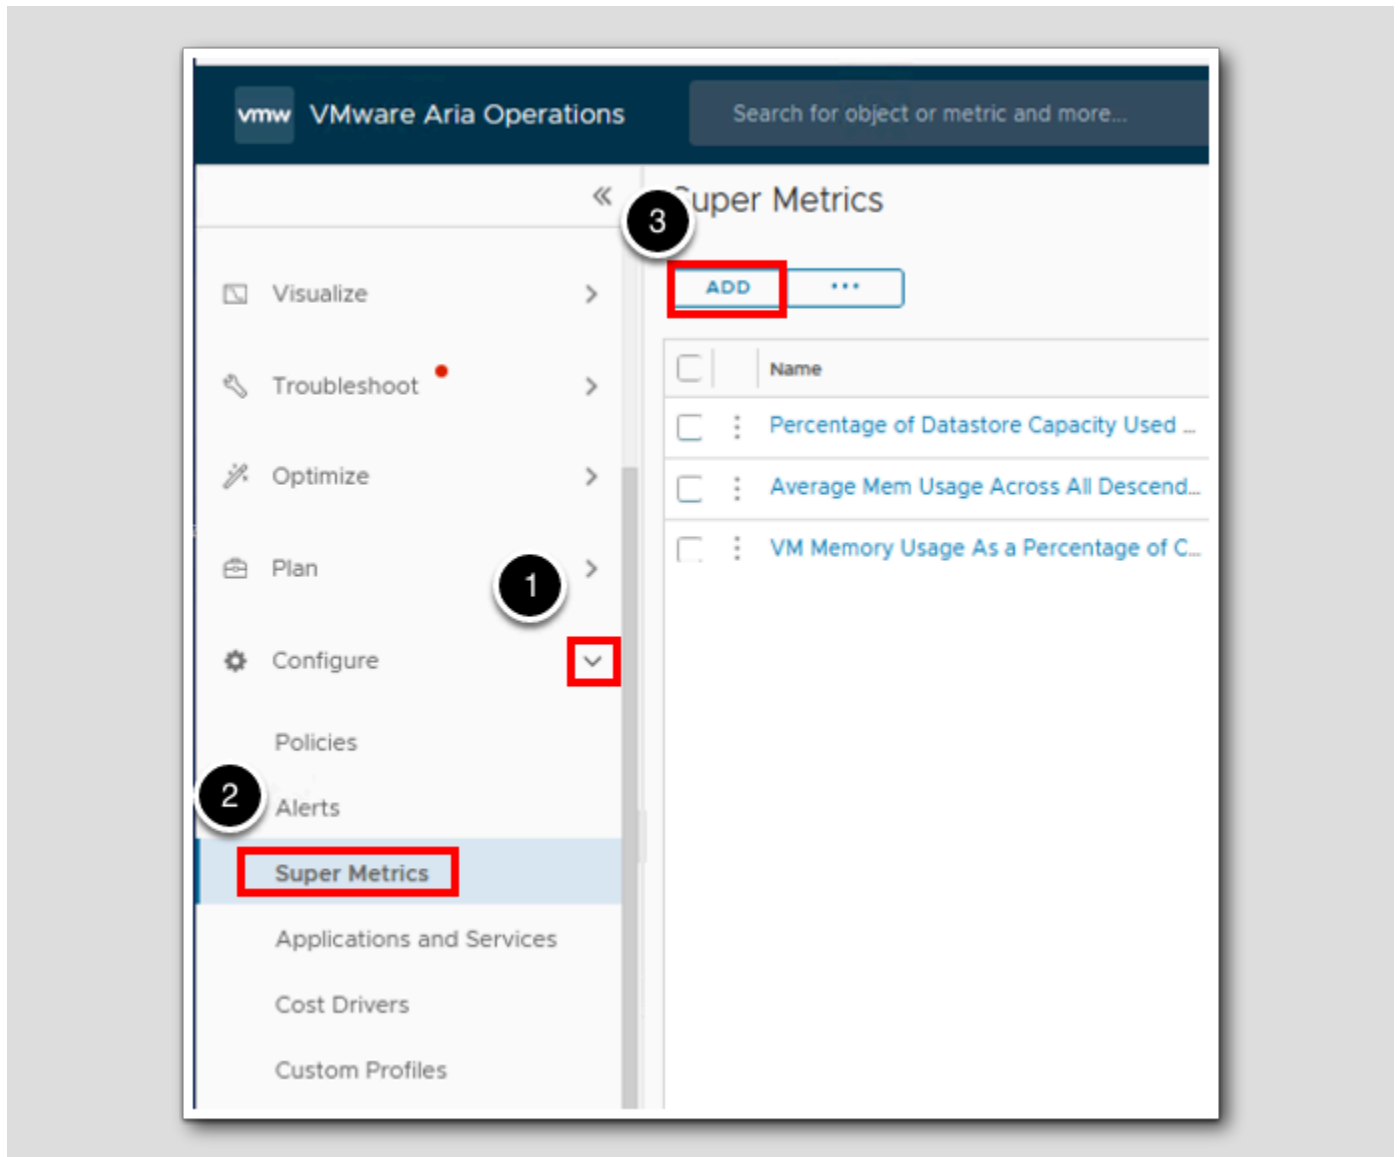
The guest operating system name is contained in the Guest OS Full Name property for a vm that is running VMtools.



1. Click the **Metrics** tab.
2. In the filter box, type **guest os from tools** and press the **Enter** key.
3. **Scroll down** to the Property section
4. Double-click the property **Guest OS from Tools**.
5. Note the OS name of this VM.

We will create a super metric that counts all of the VMs with the text "ubuntu" in that property field and then we can apply the super metric to our datacenter object type.

## Create the Super Metric



1. Expand the Configuration section.
2. Click Super Metrics.
3. Click ADD to create a new super metric.

### Name the Super Metric

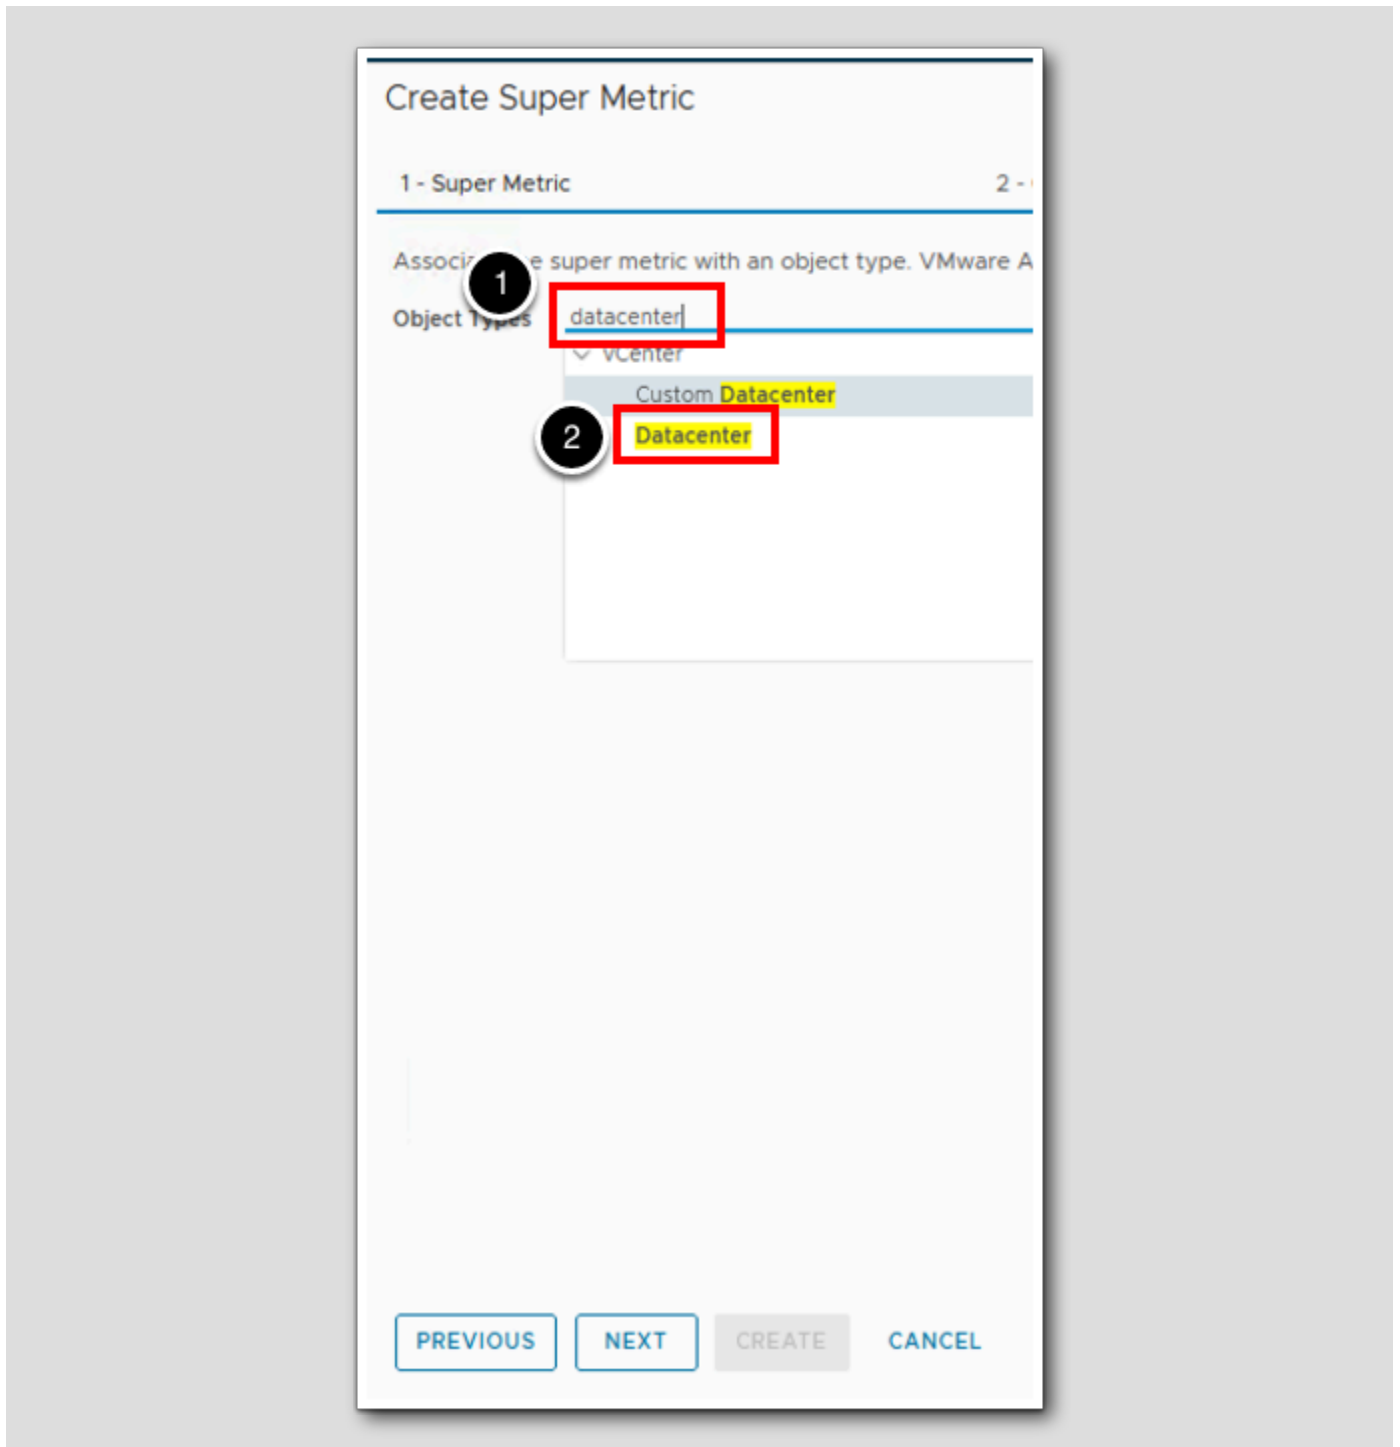
The screenshot shows the 'Create Super Metric' wizard with two tabs: '1 - Super Metric' (active) and '2 - Object Types'. The 'Name' field contains 'Count of Ubuntu VMs' and is marked with a red box and a circled '1'. The 'Description (Optional)' field contains 'Counts the number of VMs that are running Ubuntu as an operating system in the attached datacenter, cluster or host.' and is marked with a red box and a circled '2'. At the bottom, the 'NEXT' button is highlighted with a red box and a circled '3'. Other buttons include 'PREVIOUS', 'CREATE', and 'CANCEL'.



1. Type the name: Count of Ubuntu VMs.
2. Type a description.
3. Click **NEXT**.

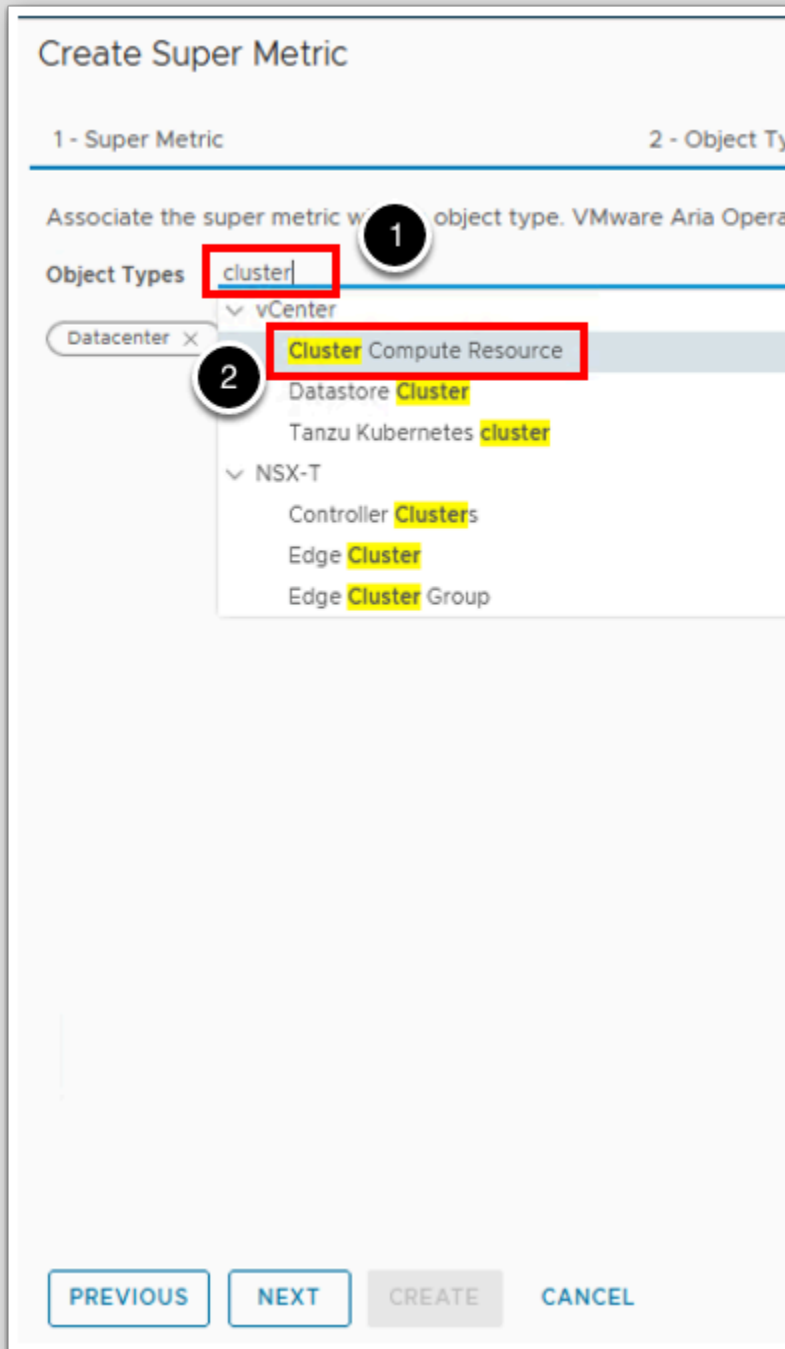
## Object Types

[481]



1. In the **Object Types** field type **datacenter**.
2. Single click on **Datacenter**.

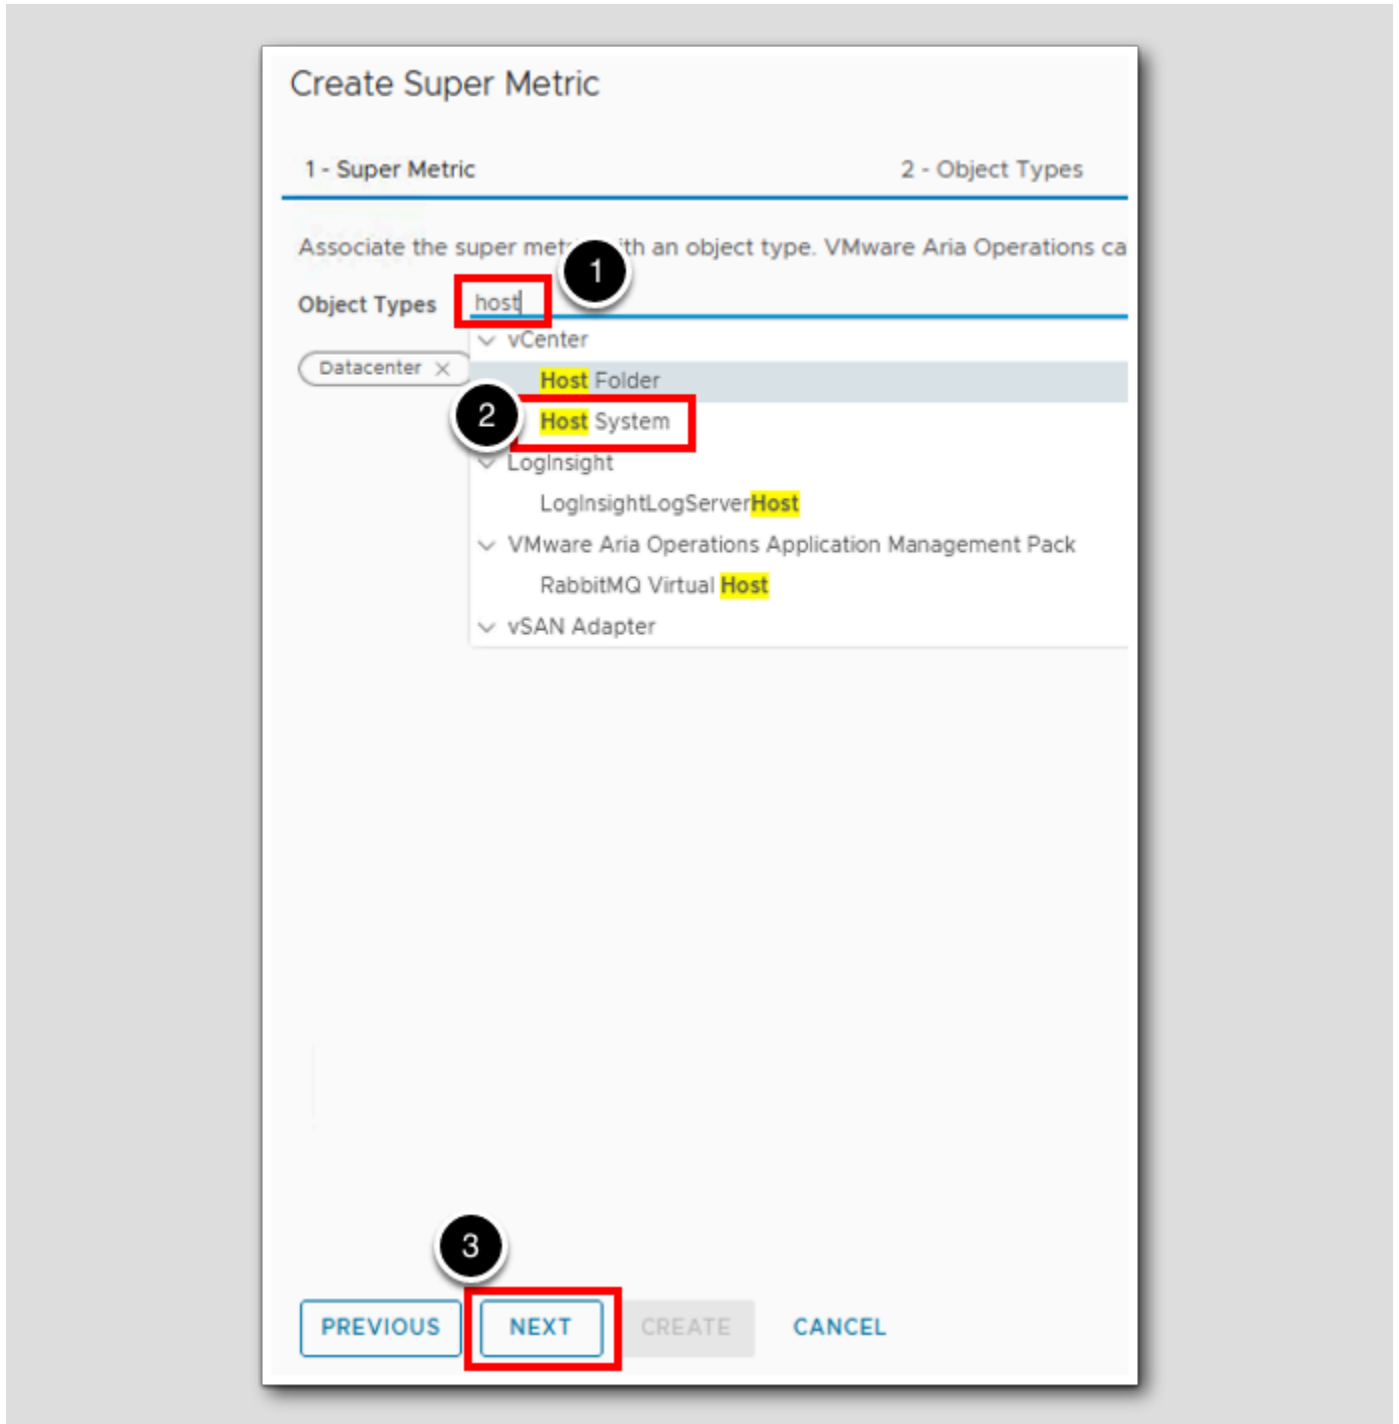
### Add Cluster as a type



1. In the **Object Types** field type cluster.
2. Single click on **Cluster Compute Resource**.

Now add Host System

[483]

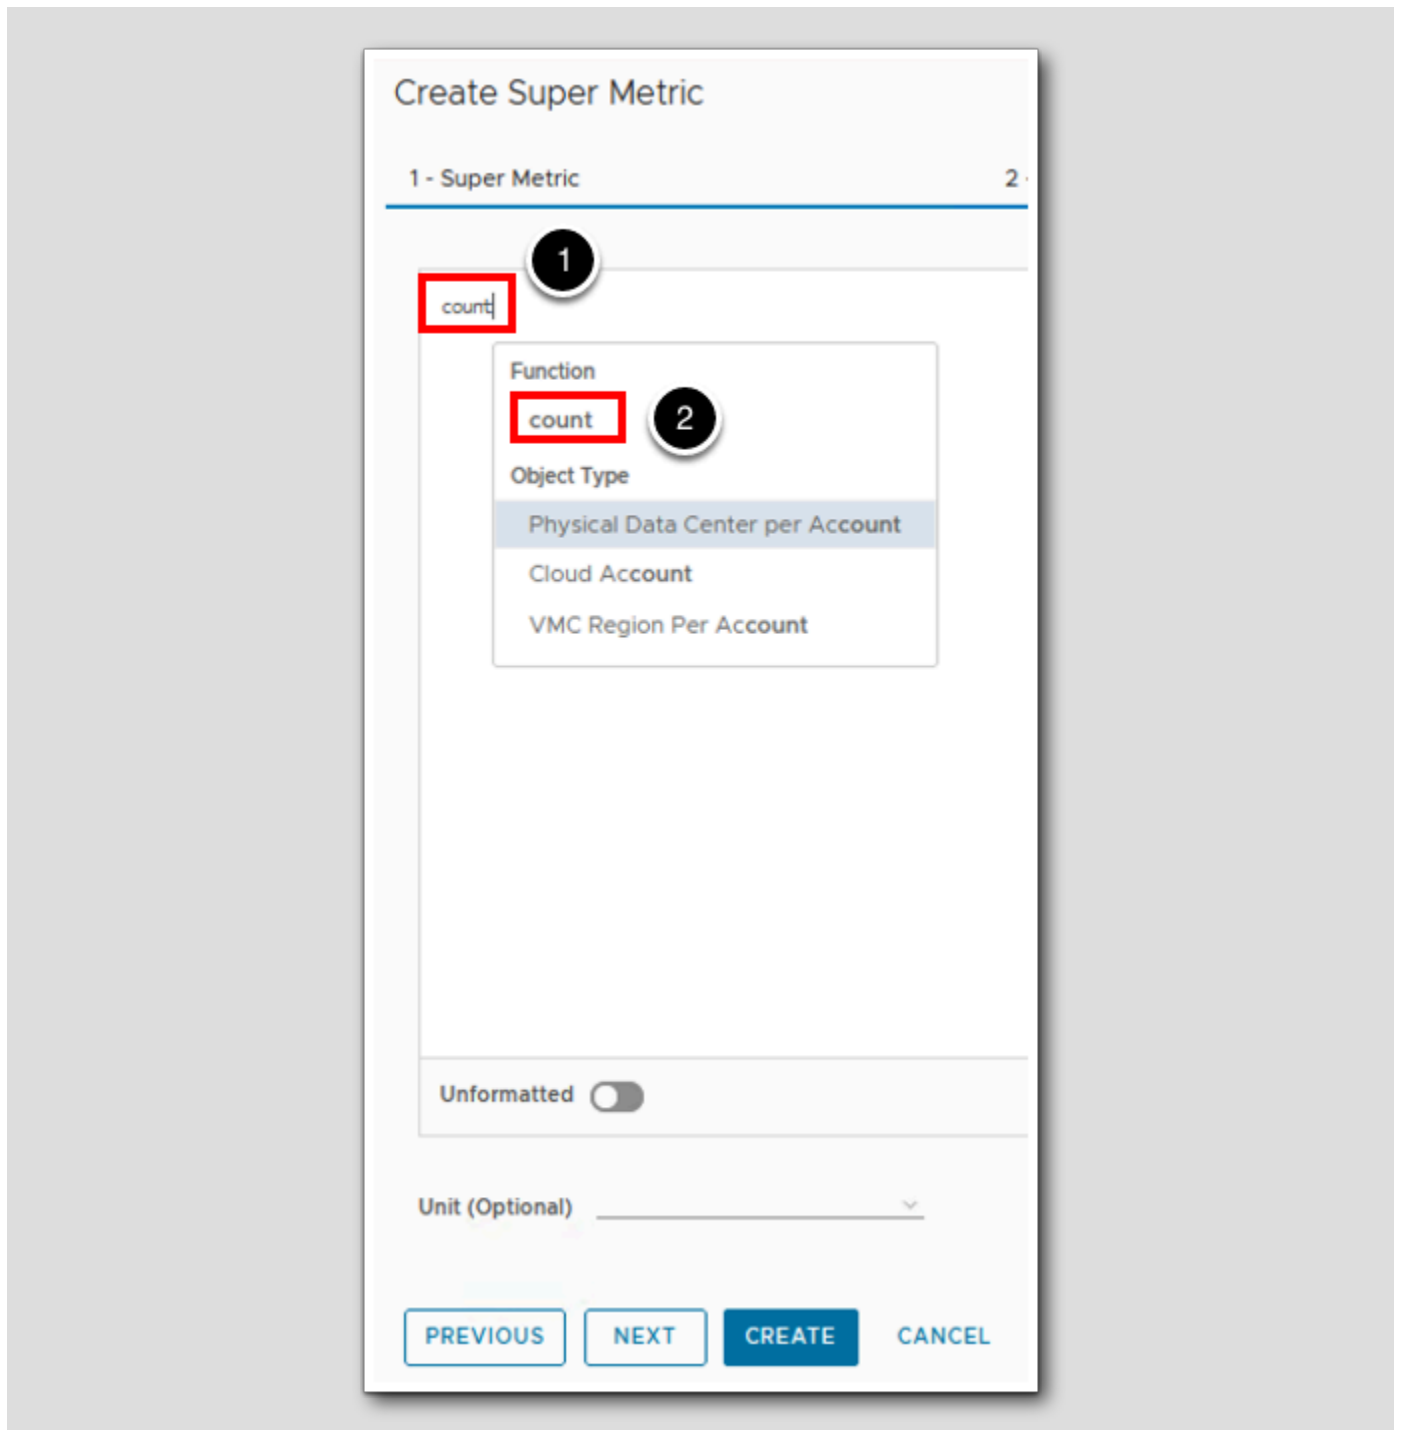


1. In the **Object Types** field type **host**.
2. Single click on **Host System**.
3. Click **NEXT**.

## Start the Formula

[484]

Remember that we want to count the number of VMs running the CentOS operating system so we will use the **count** looping function.

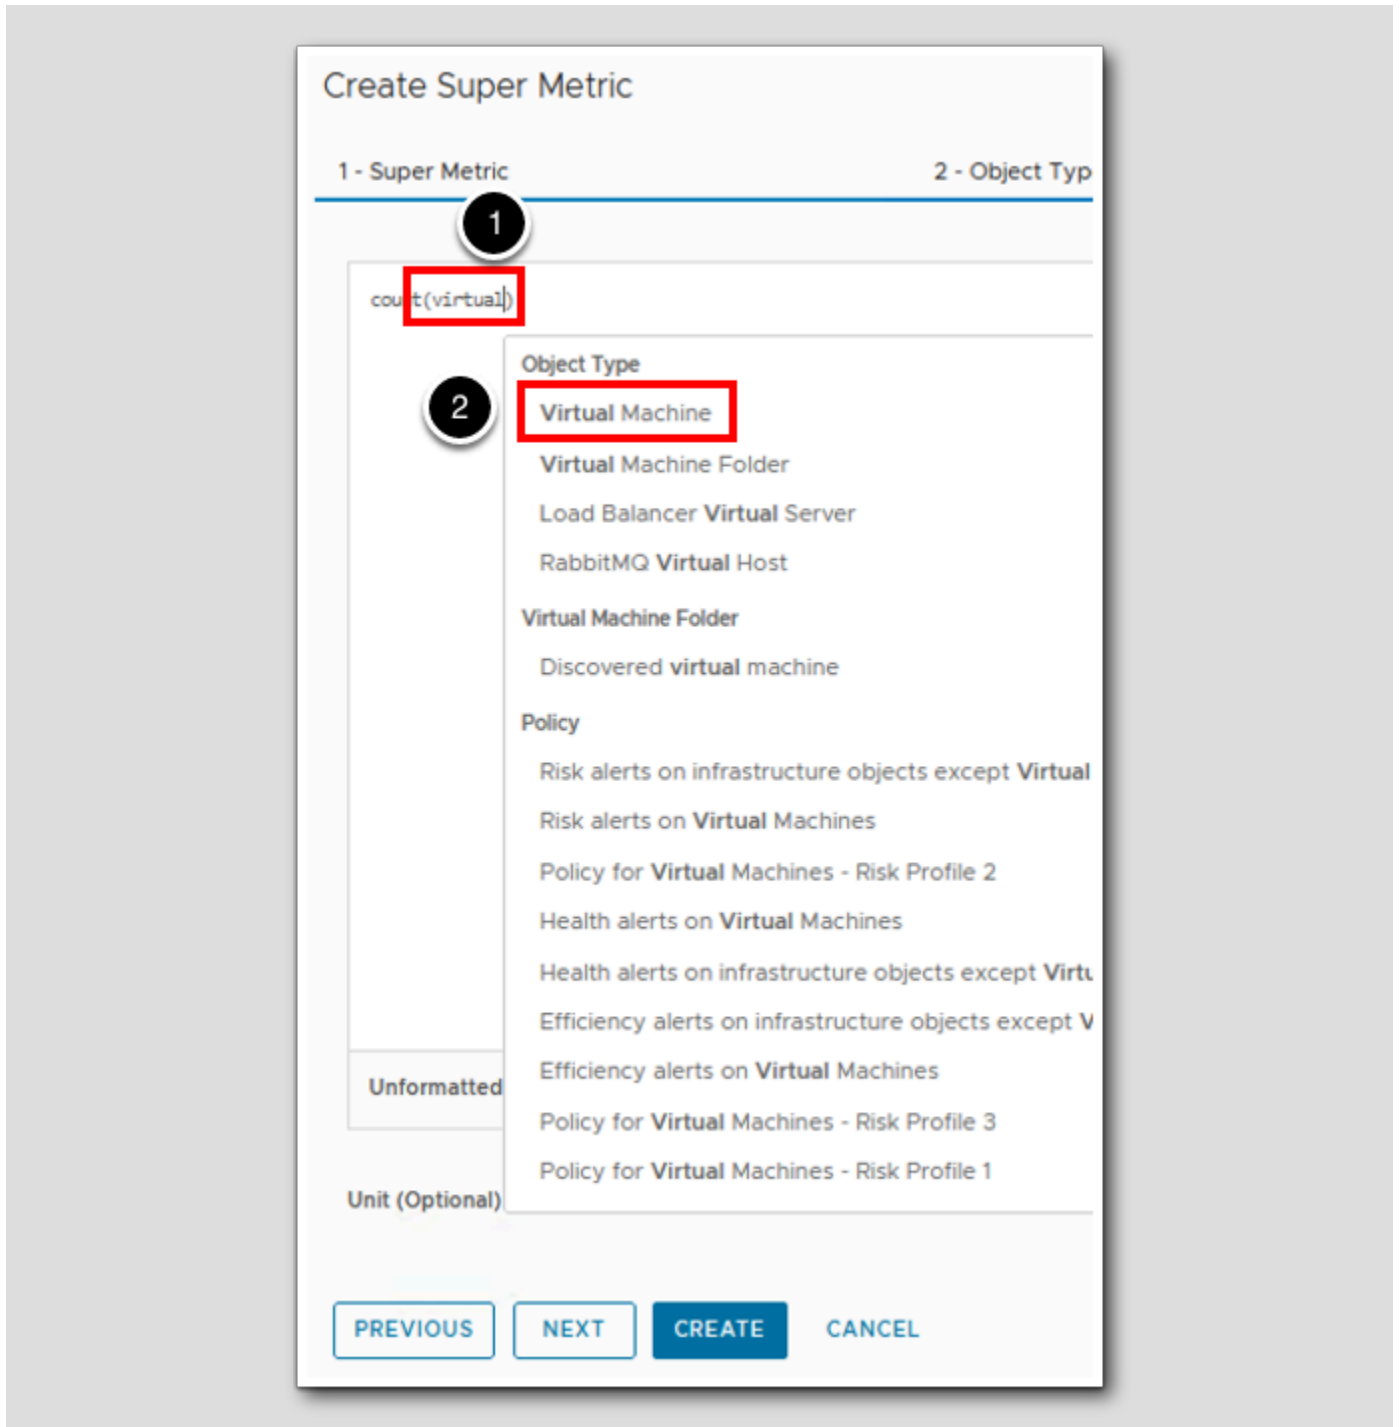


1. Type **count** to see a list of matching options.
2. Click the **count** function.



Select the Virtual Machine Object Type

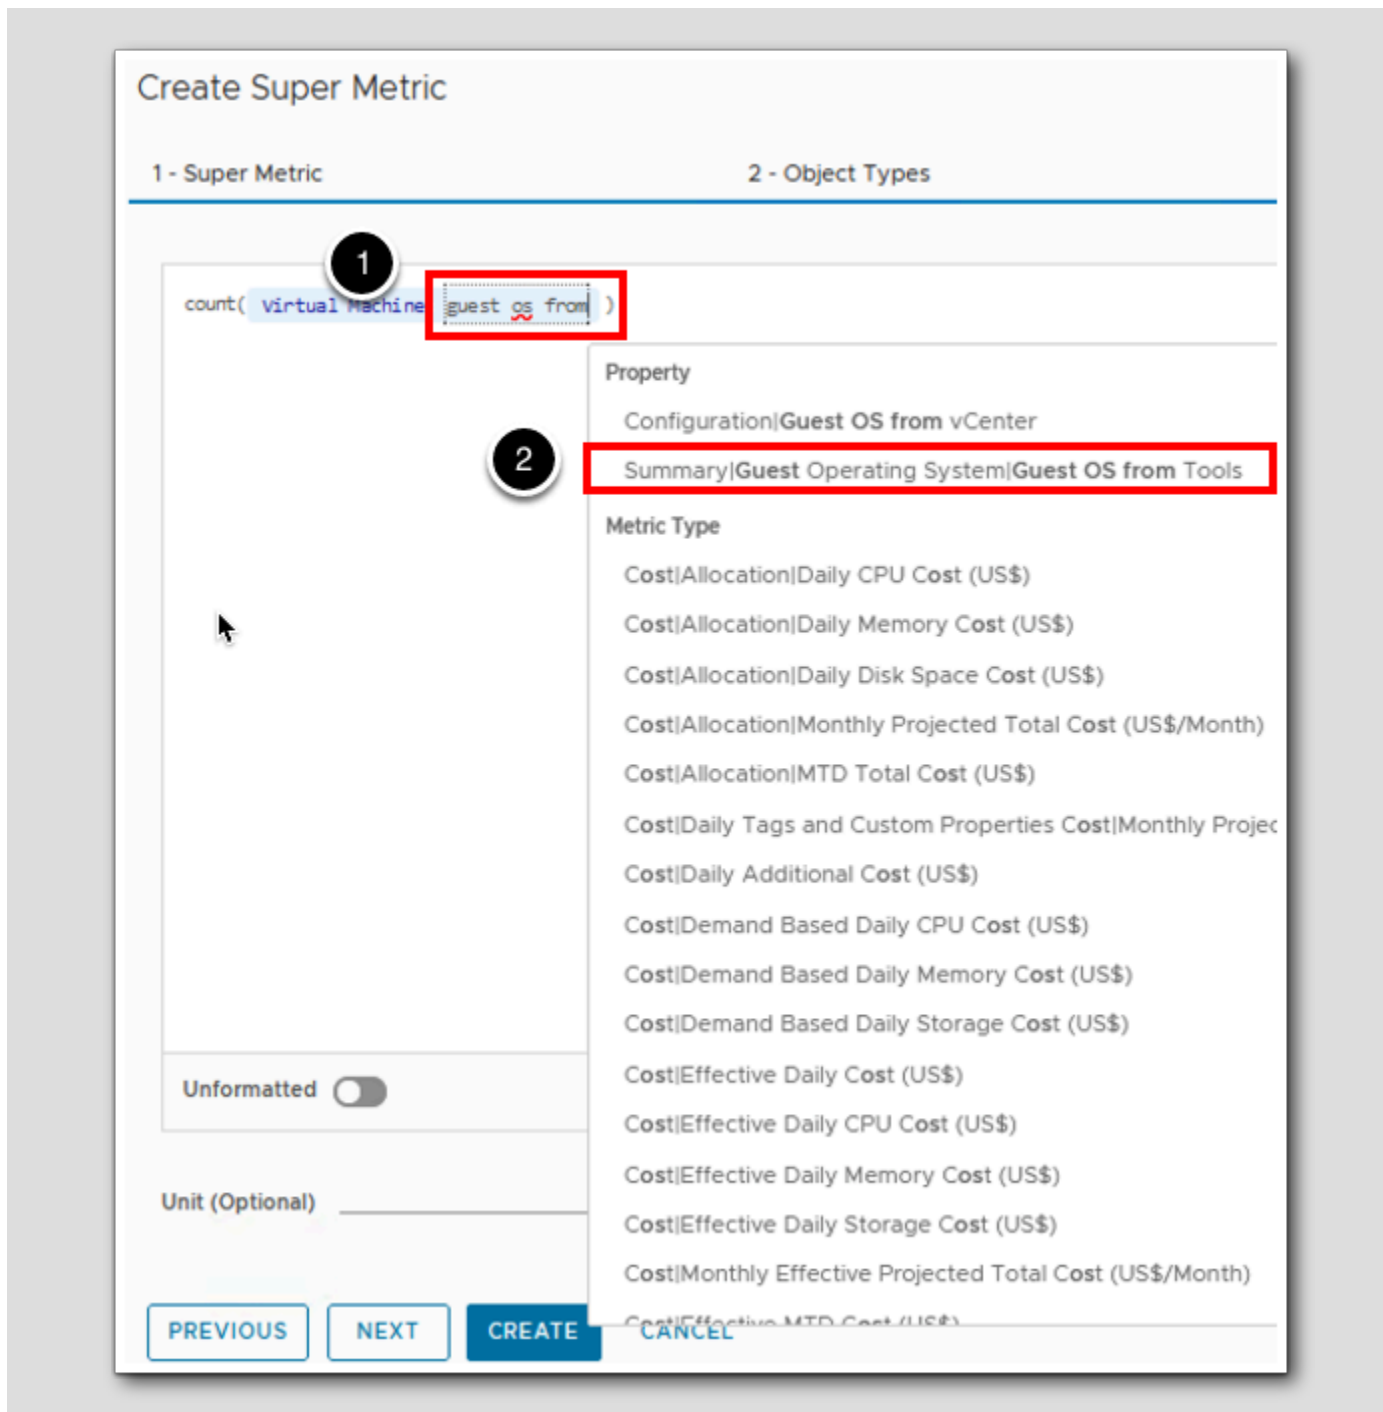
[485]



At the cursor position (between the parenthesis):

1. Type **virtual**.
2. From the match list, click the **Virtual Machine** object type to select it.

### Select the Guest OS From Tools Property



At the cursor position:

1. Type **guest os from**.
2. From the match list, click the **Summary|Guest Operating System|Guest OS from Tools** property.

## Adjust the Depth Parameter

[487]

Remember that we are going to want to apply this metric at the vSphere Datacenter object level. Going back to our discussion earlier about depth, we will need to set the depth to Datacenter --> Cluster --> Host --> VM or three levels down. Traversing down the hierarchy means a positive depth parameter so:

Create Super Metric

1 - Super Metric 2 - Object Types

1

```
count({ virtual Machine: Summary|Guest Operating System|Guest OS from Tools , depth=1})
```

Unformatted

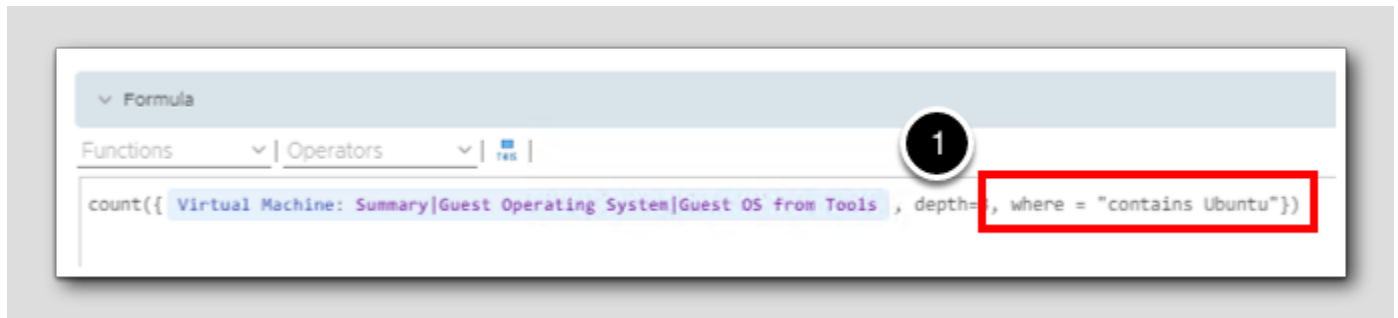
Unit (Optional) \_\_\_\_\_

PREVIOUS NEXT CREATE CANCEL

1. Delete the 1 for depth and then type 3 in its place.

## Add the Where Clause

[488]




At the cursor position (just to the right of the 3 you typed), type the following. Note the leading comma, the quotation marks and the exact case. The syntax may not seem intuitive but that is the way it needs to be written. It might be easiest to just highlight the text below and drag it to the HOL console.

1. After `depth=3` type `, where = "contains Ubuntu"`

## Let's Preview it

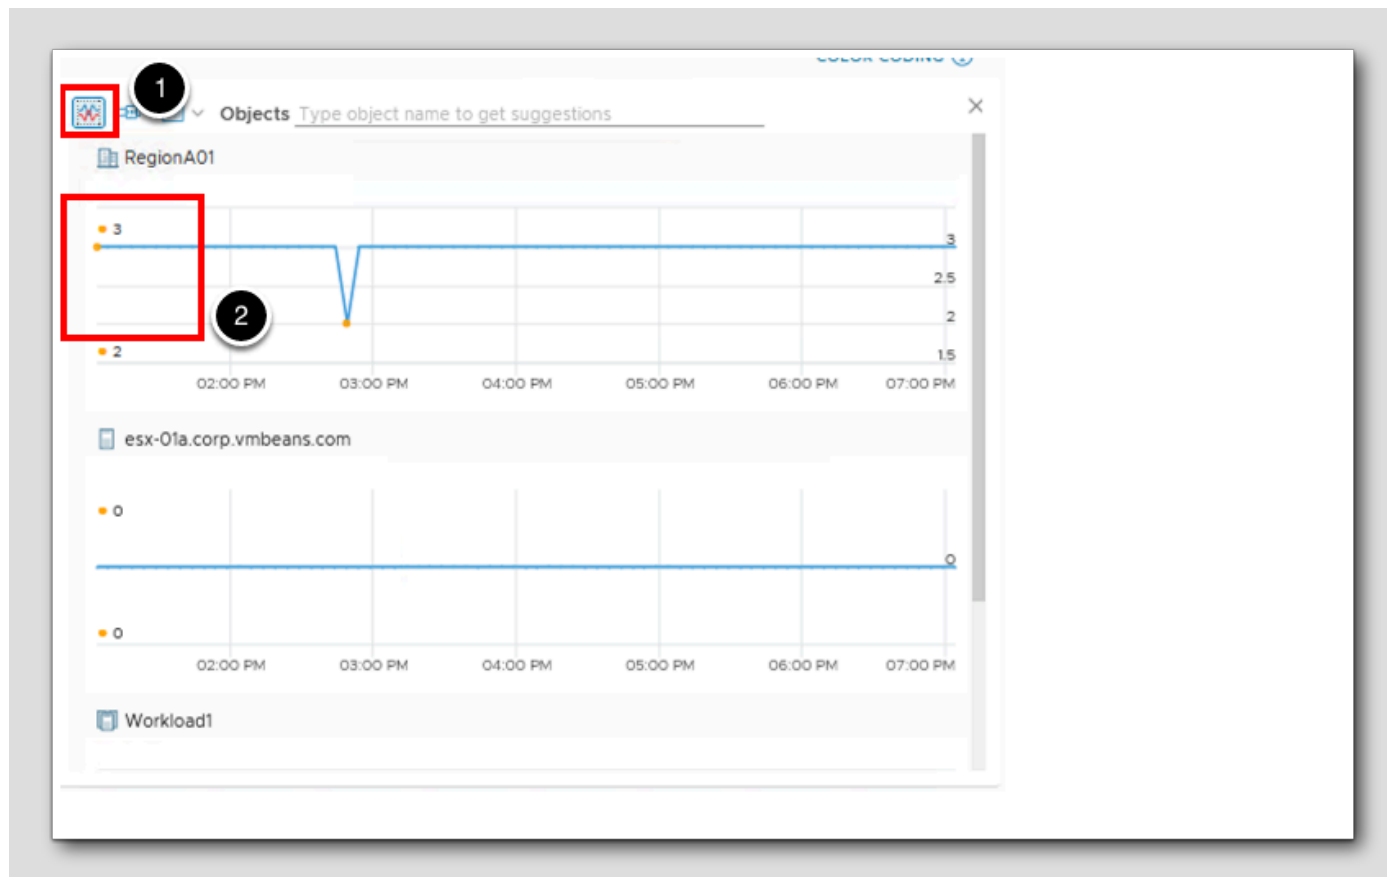
[489]



1. Click PREVIEW.



## View The Super Metric Preview



1. If the **Split Charts** button is not active, click on the **Split Charts** button.
2. Notice that **RegionA01** shows 3 Ubuntu VMs.

Note that you can also assign the super metric to **Host System** and **Cluster Compute Resource** object types with good results since this formula will look down 1, 2 and 3 levels to find Virtual Machine object types and check the operating system property for Ubuntu.



## On to Policy

[492]

Create Super Metric

1 - Super Metric 2 - Object

```
count({ virtual Machine: Summary|Guest Operating System|Guest
```

Unformatted

Unit (Optional)

1

PREVIOUS NEXT CREATE CANCEL

1. Click NEXT.

## Assign Policies for Object Types

[493]

**Create Super Metric**

1 - Super Metric      2 - Object Types      3 - Formula      4 - Policy

Select which policies you would like to enable this super metric in. You may also customize thresholds per policy. After one collection cycle, the super metric begins collecting each instance of the specified object type on the All Metrics tab.

Policy	Datacenter	Cluster Compute Resource	Host System
vSphere Solution's Default Policy (May 12, 2023 10:12:11 AM)	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

1      2      3

4

PREVIOUS    NEXT    **CREATE**    CANCEL

1. Check the Policy box for **Datacenter**.
2. Check the Policy box for **Cluster Compute Resource**.
3. Check the Policy box for **Host System**.
4. Click **CREATE**.

## Lesson End

[494]

In this lesson we created a Super Metric that used a Where clause to count Ubuntu systems across different Object Types.

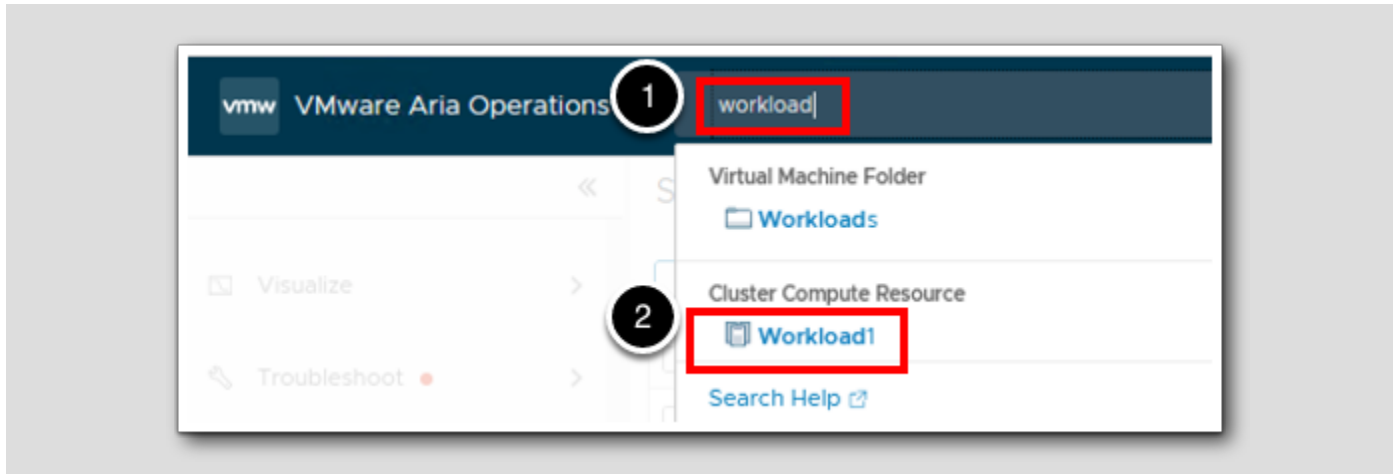
## Verifying Super Metric Calculation

[495]

We just created several super metrics. Let's check to make sure they are being calculated on the appropriate objects in our environment.

## Search for a Cluster

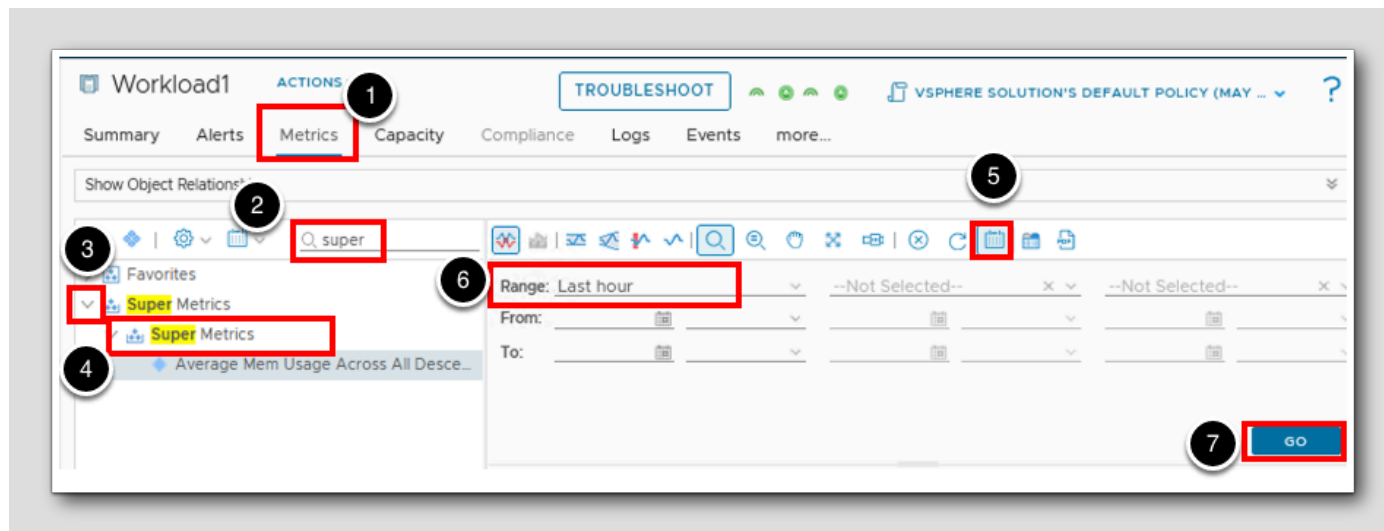
[496]



Let's first take a look at the Workload 1 vSphere cluster's metrics.

1. In the search box, type `workload`.
2. Click to select the `Workload 1` vSphere cluster.

## View the Workload 1 Super Metric




To see the calculated value of the cluster super metric:

1. Click the **Metrics** tab.
2. In the search box, enter **Super** and hit Enter.
3. Note that there is a metric category **Super Metrics**. This will only exist when there are one or more super metrics calculated for the object. Click to expand it.
4. Double-click **Average Mem Usage Across All Descendant VMs (GB)**.
5. Since the metric is new, let's change the time scale. Click the **calendar** icon.
6. Click **Last Hour** to change the time scale.
7. Click **GO**.

## Super Metric Graph

It is important to understand that super metric values will only be stored in the database from the time you create the metric and enable it in the appropriate policy.



1. Note that in this case the super metric has been calculated and stored in the database for several minutes. In your lab environment, the value and number of metric points will vary.

## Visualize Historical Super Metric Values

The screenshot shows the VMware vSphere Workload1 interface. The left sidebar has a search bar with 'super' entered. Under 'Super Metrics', 'Average Mem Usage Across All Descendant VMs (GB)' is selected. Under 'Super Metric Previews', a preview of the same metric is shown. The main area has a configuration panel for 'Workload1' with 'Range: Last 12 hours' selected. Below this are two line graphs. The top graph shows the current metric value over time, and the bottom graph shows a historical preview of the metric value over time. Red boxes and numbered callouts (1-4) highlight key UI elements: 1. Show previewable supermetrics button, 2. Super Metric name, 3. Historical preview graph, 4. Range selector.

We also have the ability to visualize what a super metric value would have been for time frames prior to when the metric was created.

1. Click the **Show previewable supermetrics** button. Note that there is now a **Super Metric Previews** category now, and it should be expanded. If it is not, then click to expand it.
2. **Double-click** the super metric name.
3. You can see that a historical view of the super metric is available.
4. It may help to look further back, so use the **calendar** icon to select a different time range. In this case, I've set my range to the **Last 12 hours**.
5. Click **GO**.

Note that the historical super metric calculation will be limited to the time range available for the metric(s) that are used in the super metric formula. In this lab environment, you may see large gaps in the data because of when the environment was created and the fact that the lab pod sits dormant (powered off) until shortly before you logged in and took this lab. Also note that while we have set a non-standard data collection interval of one minute in this lab pod (see frequency of data points in the top graph), the historical preview uses the standard 5-minute interval for calculations.

If you are interested, you can select VM, host, datacenter and datastore objects in this environment and confirm that the super metrics we created and enabled for each of those object types is also being calculated.

## Lesson End

[500]

In this lesson we checked our Super Metric creations from the Metrics tab.

## Conclusion

[501]

In this module, we created multiple Super Metrics to highlight the power of creating Super Metrics in Aria Operations.

## You've finished the module

[502]

Congratulations on completing the lab module.

For more information on getting started with Aria Operations, see the [VMware Aria Operations: Journey to Success](#) guide at the [VMware Apps & Cloud Management Tech Zone](#).

From here you can:

1. Click to advance to the next page and continue with the next lab module
2. Open the **TABLE OF CONTENTS** to jump to any module or lesson in this lab manual
3. End your lab and come back and start it again in the future

## Conclusion

### Learning Path Next Steps!

[504]

Learn More about Modern Apps and Cloud Management on Tech Zone



- Learn
- Try
- What's New

Visit <https://via.vmw.com/LearnMACM>







VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 [vmware.com](http://vmware.com).

Copyright © 2023 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at [vmware.com/go/patents](http://vmware.com/go/patents). VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Lab SKU: HOL-2401-03-CMP Version: 20230907-162130