PSO TECHNICAL WHITE PAPER – **FEB 2017**

AUTHOR – **PETER TYMBEL**

ALIGNING vREALIZE OPERATIONS WITH BUSINESS OUTCOMES

Create New Value with Business-Oriented Reporting in vRealize Operations

vmware®

Table of Contents

Executive Summary 5

Introduction 6

Solution 8

What Are Custom Groups? 8

What Are Super Metrics? 9

Defining Group Types 10

Creating Custom Groups 12

Creating Super Metrics 21

- Quantity-Based Super Metrics 22
- Utilization-Based Super Metrics 23
- Performance-Based Super Metrics 24
- Application Pool Super Metrics 25
- Binding Custom Groups and Super Metrics 26
- Enabling Super Metrics 27

Building Custom Dashboards and Reports 28

Enhancing vRealize Automation Reporting 31

Authorizing Users and Providing Role-Based Access 32

- Configuring AD Authentication 33
- Creating a New Role 33
- Managing User Groups 35
- Sharing Custom Dashboards 37

Automating Business-Oriented Infrastructure Reporting 39

Conclusion 41

References 42

About the Author 42

Acknowledgments 42

Table of Figures

Figure 1: Typical vSphere Cluster-Level Reporting Structure Diagram 6
Figure 2: Desired Business Reporting Structure Diagram....7
Figure 3: Custom Group Business-Level Reporting Diagram 8
Figure 4: Example Super Metric Function Diagram....9
Figure 5: Group Types Diagram 10
Figure 6: Create Group Types 11
Figure 7: Create Custom Group.... 12
Figure 8: Custom Group Settings 12
Figure 9: Custom Group Object Type.... 13
Figure 10: Custom Group Criteria Set 13
Figure 11: Object Name–Based Custom Groups....14
Figure 12: vSphere Cluster Relationship–Based Custom Groups 15
Figure 13: vCenter VM Folders in vROps 16
Figure 14: vCenter VM Folder Relationship–based Custom Group 16
Figure 15: vSphere Tag–Based Custom Groups.... 17
Figure 16: vRealize Infrastructure Navigator–Based Custom Group 18
Figure 17: EPOps Agent–Based Custom Groups – Example 1: VMs....19
Figure 18: EPOps Agent–Based Custom Groups – Example 2: Services....20
Figure 19: Create Super Metric Wizard.... 21
Figure 20: vSphere Cluster Summary Metrics Example.... 22
Figure 21: Binding Super Metrics to Custom Groups via Group Type....26
Figure 22: Enabling Super Metrics in the Policy 27
Figure 23: Business Unit Summary Dashboard Example.... 28
Figure 24: Business Unit Compute Analysis Dashboard Example....29
Figure 25: Business Unit Storage Analysis Dashboard Example.... 29
Figure 26: Business Unit Distributions Dashboard Example....30
Figure 27: Business Unit Trends Dashboard Example 30
Figure 28: vRA MP Object Types and Metrics.... 31
Figure 29: Privileged Access Diagram.... 32
Figure 30: Active Directory Authentication Configuration 33
Figure 31: Create User Role....34
Figure 32: Assign Permissions To Role....34
Figure 33: Import Active Directory Groups.... 35

Figure 34: Import Users and Groups.. 35
Figure 35: User Roles and Object Permissions...36
Figure 36: Share Dashboards.. 37
Figure 37: Share Custom Dashboard with User Group... 37
Figure 38: Stop Sharing Dashboards with Everyone Group...38
Figure 39: Example of Limited Access...38
Figure 40: vROps Alert Definition and Notification Construct..40
Figure 41: vROps Custom Groups and Policies Diagram... 41

Executive Summary

vRealize Operations Manager (vROps) helps organizations monitor and analyze their virtual infrastructure, providing valuable performance and capacity planning tools. While vROps is great at mapping the hierarchical infrastructure relationships out-of-the-box (OOTB), there is a need for more in-depth business-oriented reporting based around logical business constructs such as services, applications, departments, groups, or any other type of logical business structure. Quantifying performance, utilization, and consumption at the logical business unit level is more important to business leaders than at the vSphere Cluster level. This is because in today's world, services, applications, and business units can span multiple infrastructures from private, through hybrid, to public. Not only is business-oriented reporting more intuitive for a business to consume, it also makes the IT more transparent and aligned with actual business outcomes.

vRealize Operations Manager is ready to answer this challenge. The software ensures managers and application owners can make better decisions, and it optimizes their workloads with minimal forethought and planning. vROps provides the tools needed to map out how the business structure relates to the underlying IT infrastructure and enables IT to create a more comprehensive business-oriented reporting solution that delivers new business value.

Introduction

Frequently, service, application, and business unit stakeholders will ask IT to provide reporting around their workload's performance, utilization, and consumption as a whole. For example, a business owner may ask the following questions:

1. What is the VM Growth for Service X?
2. How many CPUs have been allocated to all VMs in Application Y and how much of the allocated CPU power are those VMs utilizing?
3. How much storage space are all VMs in Department Z consuming?
4. What are the poorly performing VMs in an Application Stack and why?

These questions are mostly about the underlying virtual infrastructure, but instead of being posed at a vSphere Cluster level, they are positioned around the logical business constructs. Traditionally, it has been very difficult for IT to answer these types of questions because typical Element Managers, including vCenter, used by IT, represent various objects in a very rigid hierarchical fashion that does not translate well to the elastic business structure. Many monitoring tools, including vROps by default, rely on this hierarchical inventory structure and base their reporting around it. For example, vSphere Cluster is a natural capacity container for compute and storage resources with which virtualization teams are intimately familiar. Naturally, then, vROps uses the vSphere Cluster construct to roll-up and report performance, utilization, and consumption at this level. This model works well for virtualization engineers as it aids in effective capacity planning, but it provides negligible value to the business stakeholders.

Figure 1: Typical vSphere Cluster-Level Reporting Structure Diagram

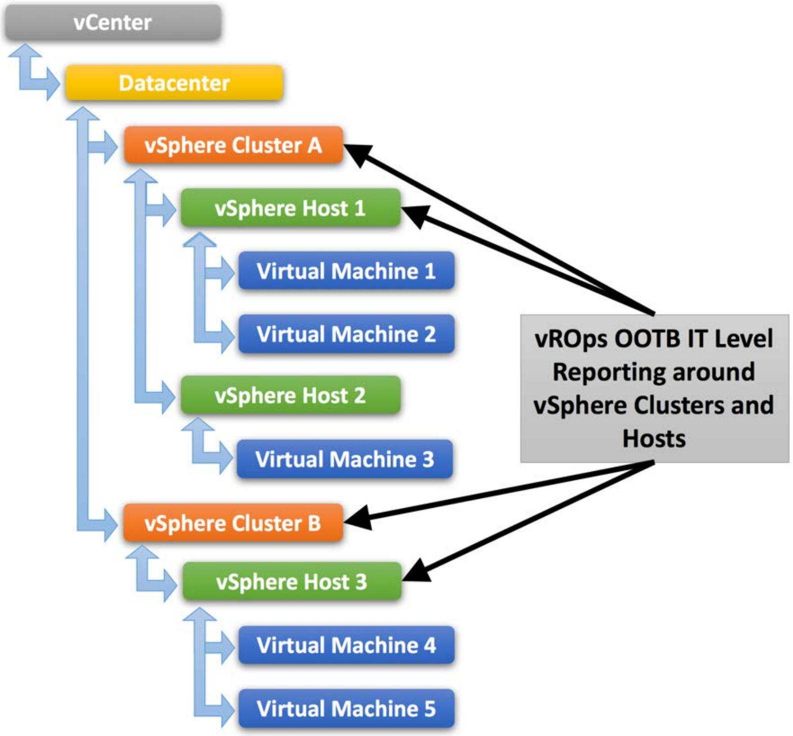

Another factor contributing to poor business reporting in IT is lack of metadata that can tie the various IT elements to their related business objects. The problem gets compounded with abstraction technologies and virtualization, which rely on complete separation for workload portability, resilience, and security.

vROps OOTB reporting around the hierarchical resource container model fails if VMs to be reported on are spread across multiple vSphere Clusters, or even worse if they are located in different physical datacenters, cloud providers, and are managed by separate vCenter instances.

A slightly different approach is necessary to address these challenges and empower business stakeholders with better information, and vROps is just the tool to do this. vROps enables IT to analyze the underlying infrastructure and present the information in a context consumable by business.

Figure 2: Desired Business Reporting Structure Diagram

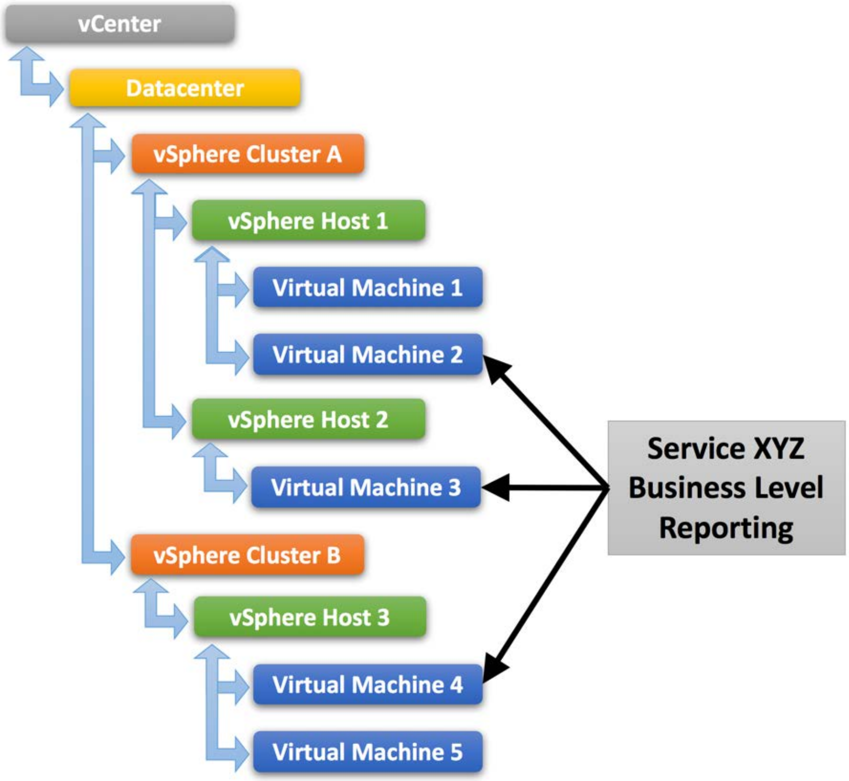

Solution

vROps allows for the organization of VMs into logical groupings that represent the business structure needed to easily answer the questions posed earlier. Taking advantage of vROps features such as Custom Groups and Super Metrics, necessary logic can be built to quantify and answer the business questions being asked. These constructs abstract and decouple the performance, utilization, and capacity data from the underlying infrastructure, and massage it into business-relevant information, thereby creating new business value.

What Are Custom Groups?

Custom Groups are a special type of container in vROps that can be used to organize VMs into logical buckets, which mimic the business units, services, or applications. Wrapping these soft boundaries around the VMs is a crucial step in enabling vROps to understand the business structure of the organization.

Custom Groups can contain other Object Types (Resource Kinds) such as vSphere Hosts, Clusters, Datastores, etc. as well as a mix of different objects. However, this example will only focus on virtual machine-based Custom Groups as they are most relevant to the business reporting use case.

Figure 3: Custom Group Business-Level Reporting Diagram

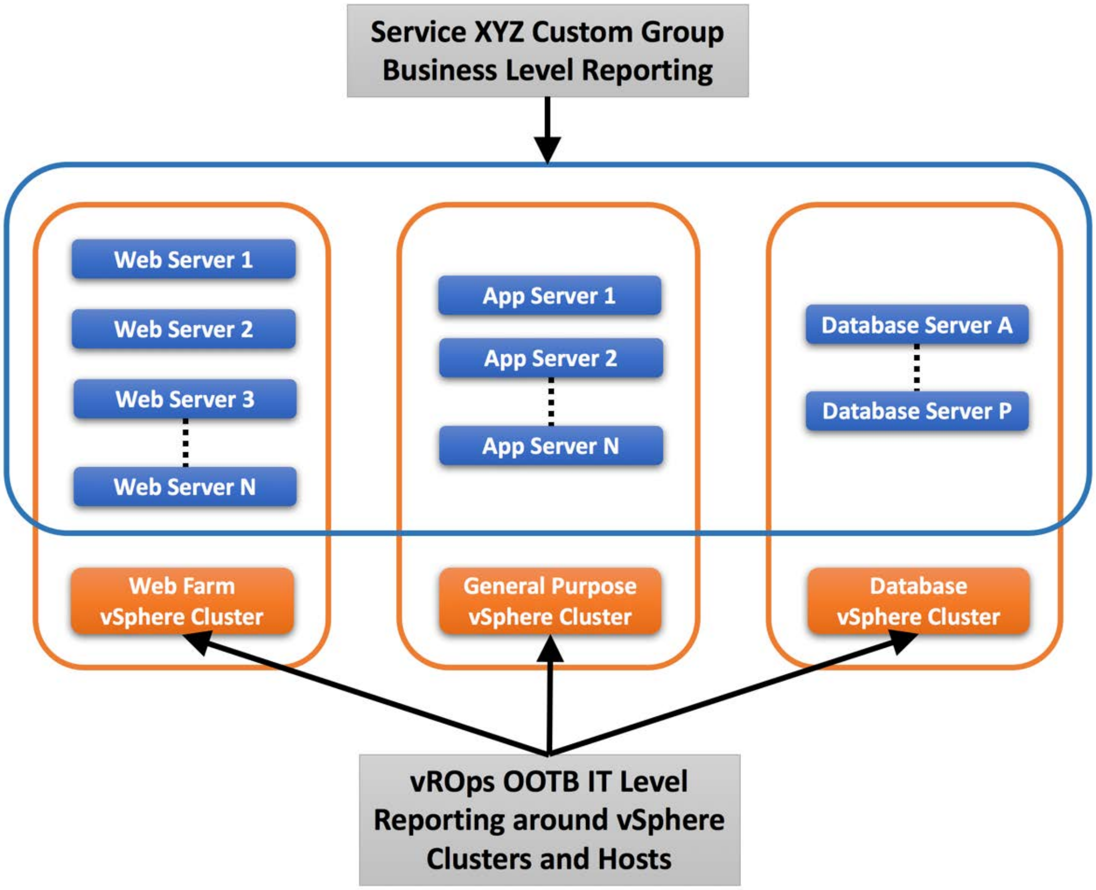

What Are Super Metrics?

Super Metrics are derived metrics that can be used to measure performance, utilization, and consumption by different business units and applications, or in this case, the Custom Group into which the VMs are placed. Super Metrics contain simple algebraic formulas that allow vROps to measure various aspects of the Custom Groups, which contain business unit–related VMs.

At the most basic level, Super Metrics enable the ability to extend built-in metrics by adding new metrics that provide additional insight. This is especially important for Custom Groups as they only come with basic badge metrics. Having a set of more comprehensive predefined metrics is not optimal since Custom Groups can contain a mix of vastly different objects such Hosts and Datastores. For example: CPU- and Memory-related metrics would not be applicable to the Datastores.

This is also where the not-so-obvious power of Super Metrics becomes clearly apparent. The Super Metrics employed to quantify Custom Groups use the concept of Relative Reference, or Depth, which allows vROps to look inside of a container/bucket, measure what is in it, and return the value at the container level rather than the object being measured. This ability sets vROps apart from other monitoring tools on the market as it empowers the user to create almost an infinite number of algorithms leveraging the built-in functions and collected time-series data, thus opening the door for unparalleled opportunity to create new business value.

Figure 4: Example Super Metric Function Diagram

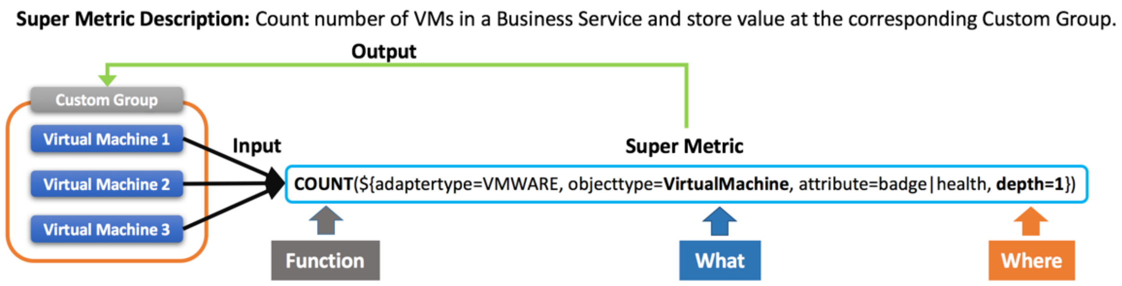
Super Metric Description: Count number of VMs in a Business Service and store value at the corresponding Custom Group.
Output
Custom Group
Virtual Machine 1
Virtual Machine 2
Virtual Machine 3
Input
Super Metric
COUNT(${adaptertype=VMWARE, objecttype=**VirtualMachine**, attribute=badge|health, **depth=1**})
Function
What
Where

After defining the logical structures and measuring them, various Custom Dashboards, Views, and Reports can be built that answer the business questions posed earlier.

Defining Group Types

Before creating Custom Groups and Super Metrics, Group Types must be defined. Group Types are the glue that bind Super Metrics to Custom Groups and will help organize categories of groups in vROps. This step must be done prior to creating Custom Groups, because a Group Type is required during the Custom Group creation process, and it is difficult to change this later unless the group is cloned.

Think of Group Types in terms of categories similar to Objects Types (Resource Kinds), for example: Virtual Machines, Datastores, vSphere Hosts and Clusters. Just like built-in objects have metrics associated with them via the Object Type, Group Types do the same thing for Custom Groups.

Figure 5: Group Types Diagram

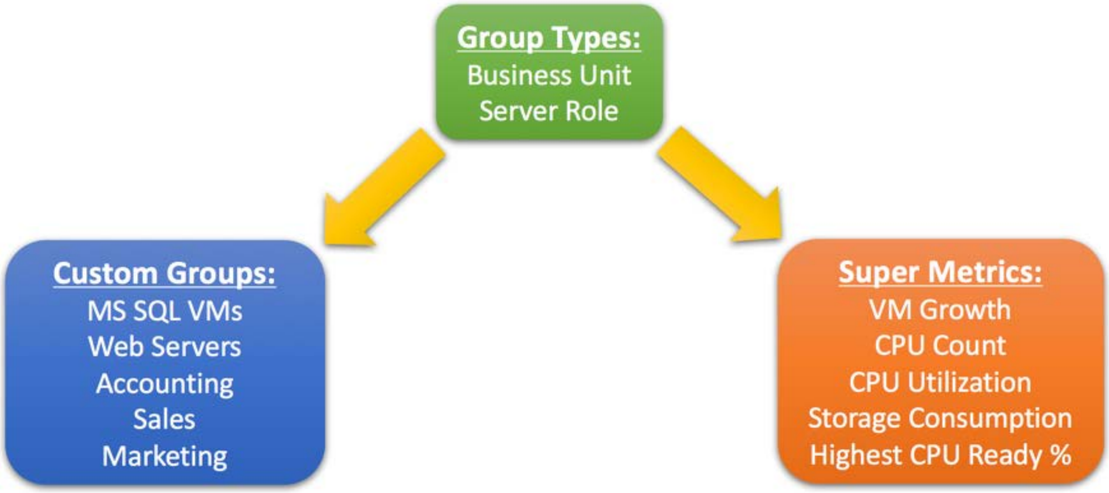

Typically, two or more Group Types are needed. One is needed to hold all of the business unit–related structures such as department, group, team, etc. Another one is needed to hold IT-related constructs such as applications, tiers, server roles, etc. In large and more complex environments, more Group Types can be defined if needed. However, while creating more Group Types will enable greater flexibility by providing more granularity, they come at a higher management overhead cost. Sometimes, keeping it simple can provide the right balance.

Another important factor that will come into play later during Custom Dashboard creation is the ability to filter by Group Type within the Widgets, similar to filtering by Object Type. This will be especially handy when different dashboards need to be constructed displaying different Group/Object categories depending on the dashboard audience.

New Group Types can be created by navigating to Content > Group Types and clicking Add Group Type.

Figure 6: Create Group Types

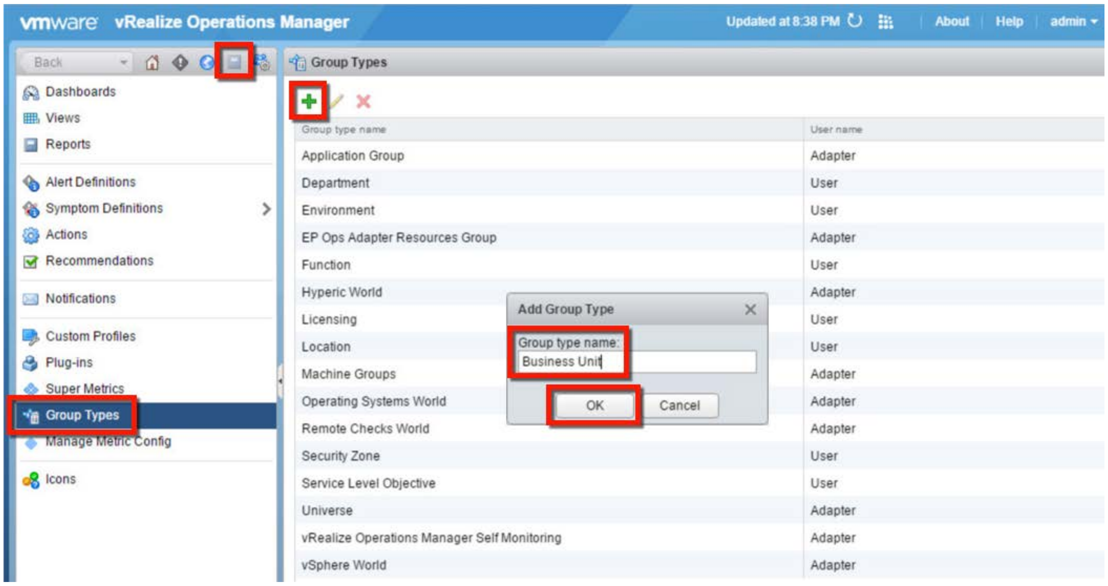

Group type name: Business Unit

Group type name	User name
Application Group	Adapter
Department	User
Environment	User
EP Ops Adapter Resources Group	Adapter
Function	User
Hyperic World	Adapter
Licensing	User
Location	User
Machine Groups	Adapter
Operating Systems World	Adapter
Remote Checks World	Adapter
Security Zone	User
Service Level Objective	User
Universe	Adapter
vRealize Operations Manager Self Monitoring	Adapter
vSphere World	Adapter

Creating Custom Groups

Once the Group Types are defined, Custom Groups based on different criteria sets can be created. The Custom Groups represent the logical business constructs such as services, applications, departments, teams, or any other type of logical business unit structure.

At their most basic level, Custom Groups are similar to SQL query SELECT statements with a WHERE clause, but created using a simple GUI. For example:

SELECT *column_name*
FROM *table_name*
WHERE *column_name operator value*;

New Custom Groups can be created by navigating to Environment and clicking New Group.

Figure 7: Create Custom Group

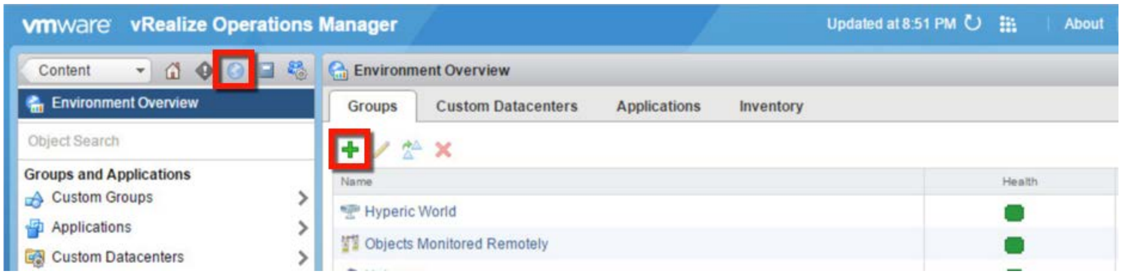

On the “New group” wizard, provide a descriptive Name, select one of the previously created Group Types, select a Policy that will apply to objects in this group, and click “Keep group membership up to date”. This last option ensures the group will dynamically include any objects (VMs) added after it has been created, providing a maintenance-free operation.

Figure 8: Custom Group Settings

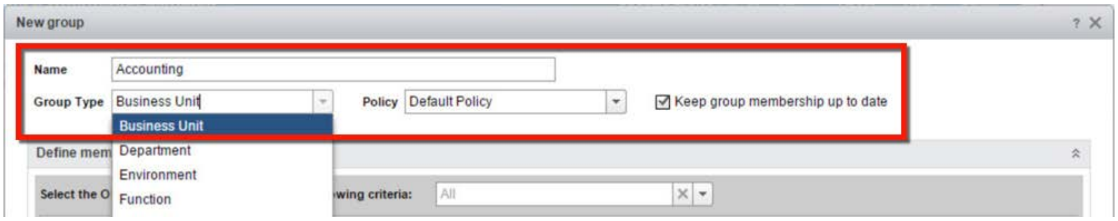
New group

Name: Accounting

Group Type: Business Unit

Policy: Default Policy

☑ Keep group membership up to date

Under Define membership criteria, in the Select the Object Type that matches all of the following criteria, find and select Virtual Machine under vCenter Adapter.

Figure 9: Custom Group Object Type

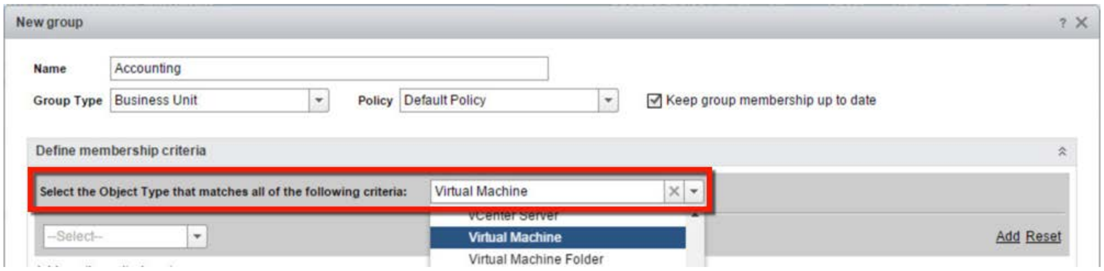

New group

Name: Accounting

Group Type: Business Unit

Policy: Default Policy

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Below the Object Type, select one of the criteria set types that the Custom Group will be based on. Most groups will be based on an Object name matching some naming convention, a Property such as a vSphere Tag, or Relationship to a vSphere Cluster.

Figure 10: Custom Group Criteria Set

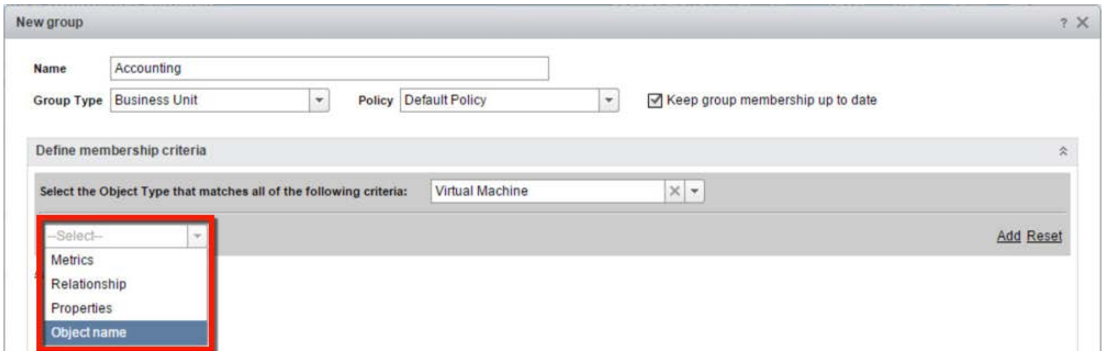

New group

Name: Accounting

Group Type: Business Unit

Policy: Default Policy

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

From here, the steps will vary depending on which option was selected for the criteria set type. Next is an examination of how to leverage each criteria set type.

Figure 11: Object Name-Based Custom Groups

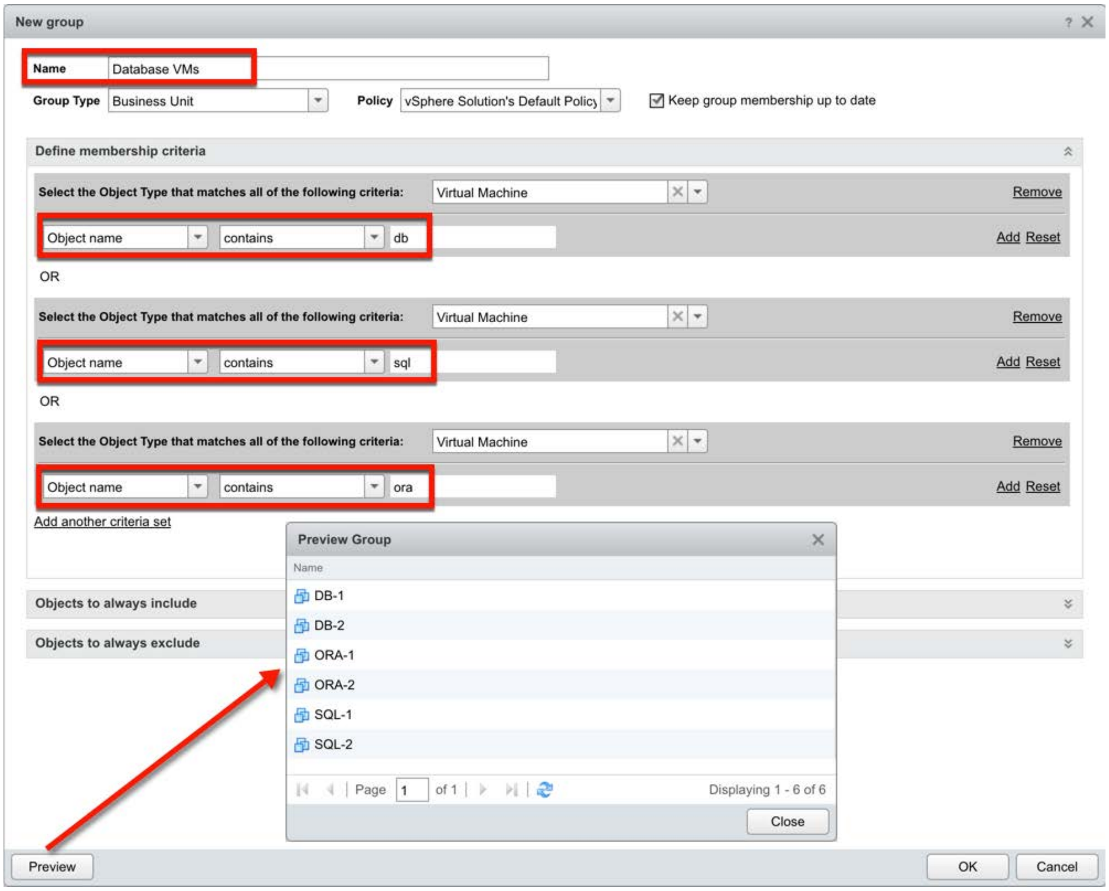
New group

Name: Database VMs

Group Type: Business Unit

Policy: vSphere Solution's Default Policy

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Object name contains db

OR

Select the Object Type that matches all of the following criteria: Virtual Machine

Object name contains sql

OR

Select the Object Type that matches all of the following criteria: Virtual Machine

Object name contains ora

Preview Group

Name
DB-1
DB-2
ORA-1
ORA-2
SQL-1
SQL-2

Page 1 of 1

Displaying 1 - 6 of 6

One of the easiest methods of grouping a number of related VMs together involves leveraging the VM naming convention for the object name criteria set. For instance, suppose all of the database VMs in the environment have SQL, ORA, or DB in their names. Create a Custom Group and use those keyword patterns in the criteria set. This will find all VMs with those keywords and include them in the group. Check the "Keep group membership up to date" checkbox to keep the group populated with new VMs introduced to the environment after the group has been created. Similarly, in an environment with a VM naming convention standard, other keyword patterns emerge that naturally fit other applications, server roles, or business units.

Figure 12: vSphere Cluster Relationship-Based Custom Groups

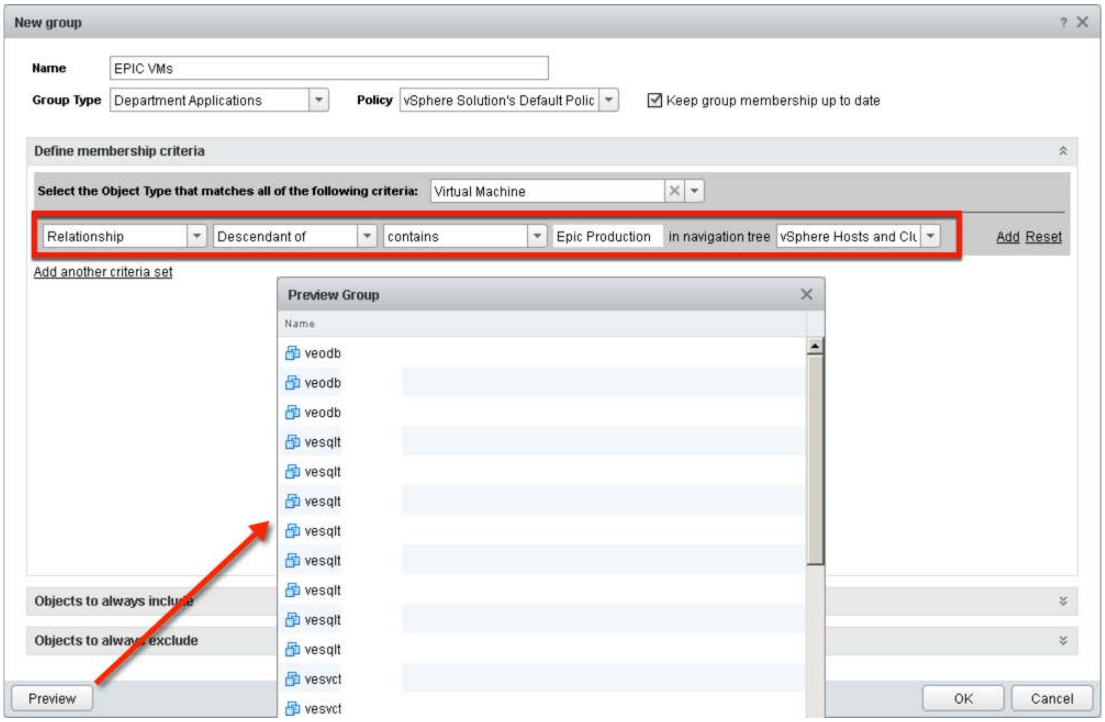

New group

Name: EPIC VMs

Group Type: Department Applications

Policy: vSphere Solution's Default Polic

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Relationship | Descendant of | contains | Epic Production | in navigation tree | vSphere Hosts and Clu

Preview Group

Name
- veodb
- veodb
- veodb
- vesqlt
- vesqlt
- vesqlt
- vesqlt
- vesqlt
- vesqlt
- vesqlt
- vesqlt
- vesvct
- vesvct

Often, solutions like SQL, Oracle, SAP, and/or Exchange, etc. may be walled off in their own dedicated vSphere Clusters for licensing or performance reasons. This is a perfect criteria set for creating a Custom Group of VMs based on the relationship to their vSphere Cluster.

Another example of leveraging a Relationship criteria set involves the vCenter VM Folder structure now supported in vROps.

Figure 13: vCenter VM Folders in vROps

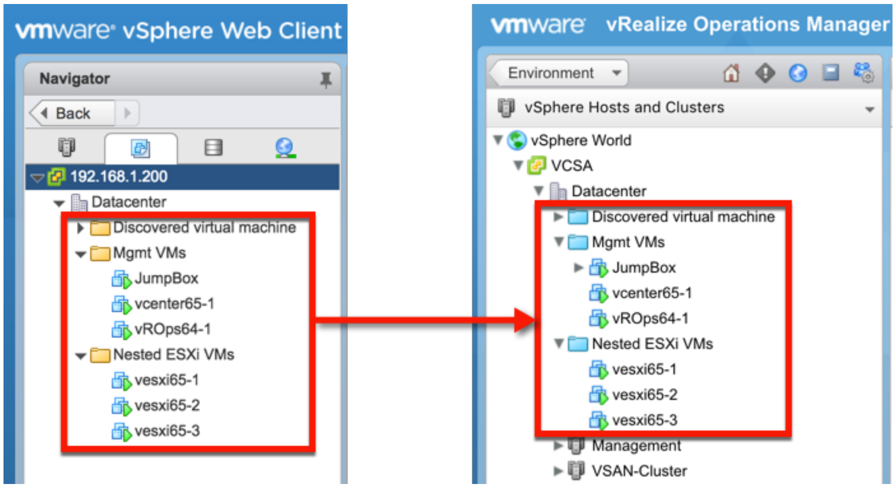

vmware® vSphere Web Client

Navigator

Back

- 192.168.1.200
  - Datacenter
    - Discovered virtual machine
    - Mgmt VMs
      - JumpBox
      - vcenter65-1
      - vROps64-1
    - Nested ESXi VMs
      - vesxi65-1
      - vesxi65-2
      - vesxi65-3

vmware vRealize Operations Manager

Environment

vSphere Hosts and Clusters

- vSphere World
  - VCSA
    - Datacenter
      - Discovered virtual machine
      - Mgmt VMs
        - JumpBox
        - vcenter65-1
        - vROps64-1
      - Nested ESXi VMs
        - vesxi65-1
        - vesxi65-2
        - vesxi65-3
      - Management
      - VSAN-Cluster

Often virtualization admins create folders in vCenter to organize different workloads by application or department to make management easier. This structure is exposed to vROps and can be used as a relationship criteria set for Custom Groups.

Figure 14: vCenter VM Folder Relationship-based Custom Group

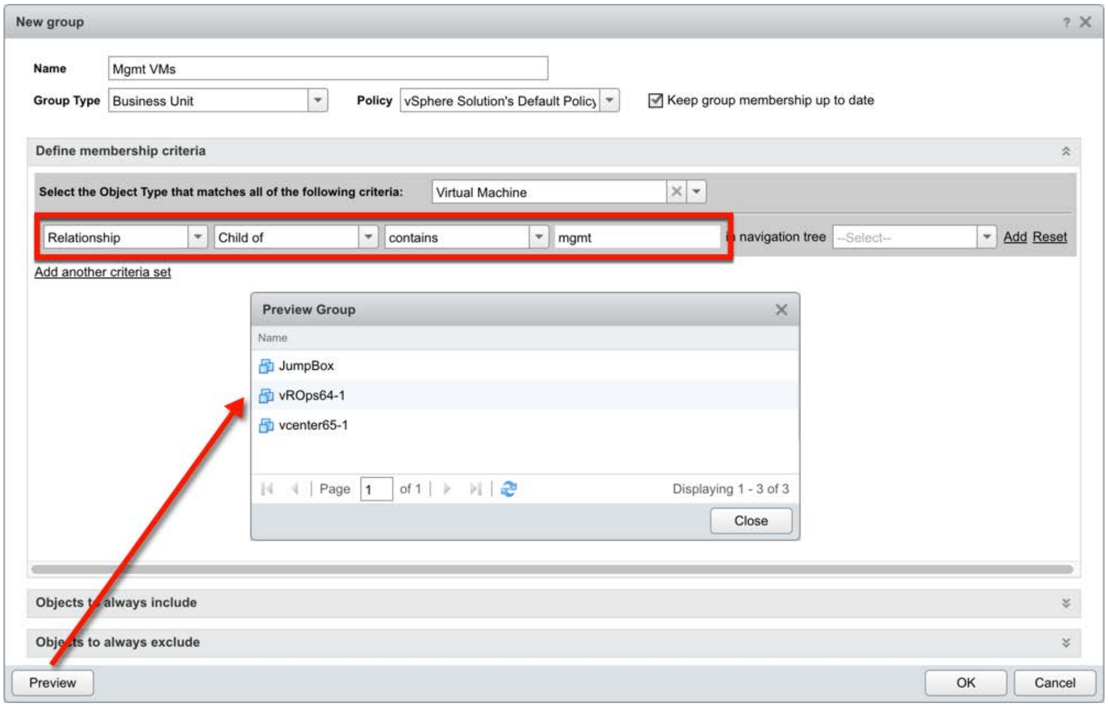

New group

Name: Mgmt VMs

Group Type: Business Unit

Policy: vSphere Solution's Default Policy

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Relationship | Child of | contains | mgmt | in navigation tree | --Select-- | Add | Reset

Add another criteria set

Preview Group

Name

- JumpBox
- vROps64-1
- vcenter65-1

Page 1 of 1

Displaying 1 - 3 of 3

Close

Objects to always include

Objects to always exclude

Preview | OK | Cancel

Figure 15: vSphere Tag-Based Custom Groups

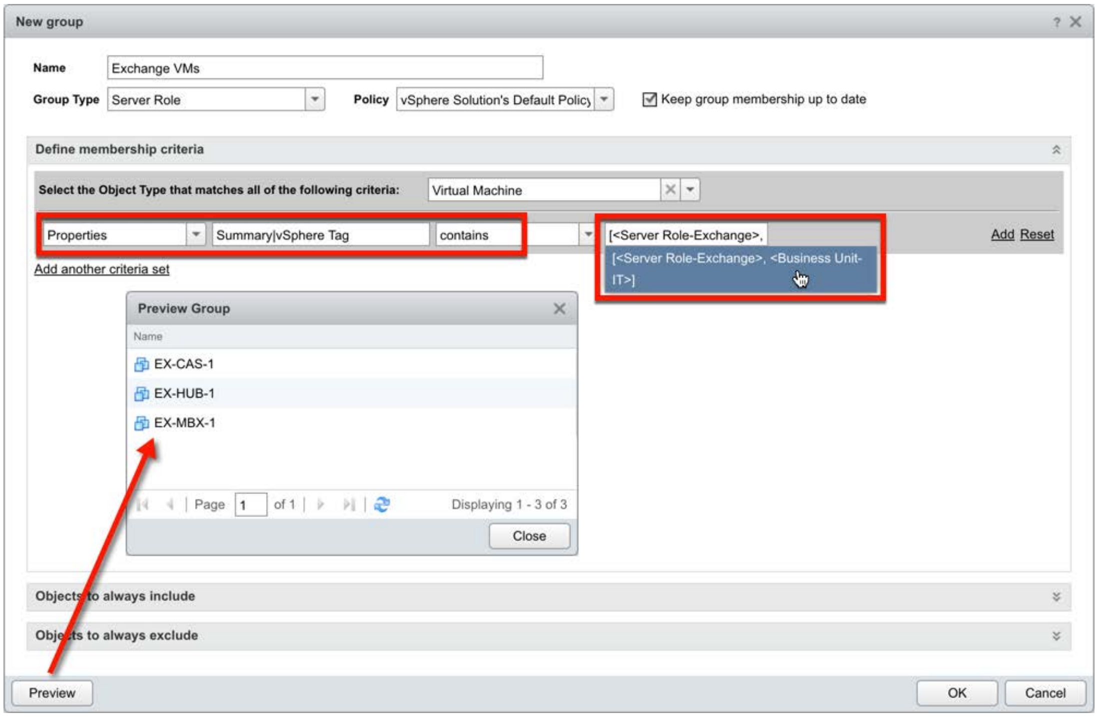

New group

Name: Exchange VMs

Group Type: Server Role

Policy: vSphere Solution's Default Policy

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Properties

Summary|vSphere Tag

contains

[<Server Role-Exchange>,

[<Server Role-Exchange>, <Business Unit-IT>]

Preview Group

Name
EX-CAS-1
EX-HUB-1
EX-MBX-1

Page 1 of 1

Displaying 1 - 3 of 3

Using vSphere Tags to organize VMs in a vCenter provides another mechanism that can be used to create criteria sets based on VM vSphere Tag properties. vSphere Tags in vCenter can be manually defined and assigned to VMs, or programmatically populated from an external source such as a CSV file, CMDB, or ITSM system using PowerCLI or vRealize Orchestrator.

Another important fact about vSphere Tags is that in the age of almost ubiquitous automation, tags should be assigned during the build process by the automation and orchestration engines, such as VMware's vRealize Automation and Orchestrator. This solves the missing link between IT infrastructure and business constructs that are so frequently overlooked in many environments.

Defining vSphere Tags in vCenter serves a dual purpose in this case as they can be leveraged by both vROps and vCenter alike when searching for VMs, not to mention various automation tools.

NOTE: While newer versions of vROps support vCenter Annotations, they can be easily converted to vSphere Tags in bulk.

Figure 16: vRealize Infrastructure Navigator–Based Custom Group

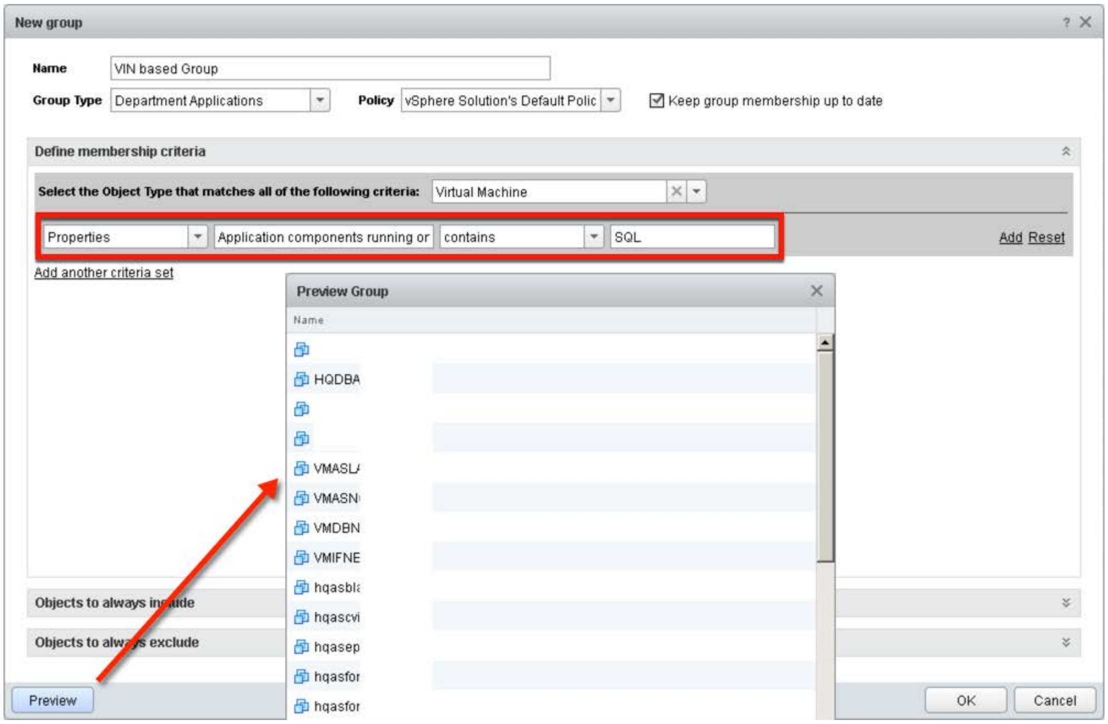

New group

Name: VIN based Group

Group Type: Department Applications

Policy: vSphere Solution's Default Polic

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Properties | Application components running or | contains | SQL

Preview Group

Name

HQDBA

VMASL/

VMASN

VMDBN

VMIFNE

hqasbl

hqascvi

hqasep

hqasfor

hqasfor

Another tool that can be utilized to create Custom Groups for VMs is vRealize Infrastructure Navigator (vRIN), formerly known as VIN. vRIN is an Application Dependency Mapping (ADM) tool that scans VMs for services and ports being used by those VMs. Relationships are then created among VMs communicating with each other. This relationship data about VMs in turn can be fed into vROps via a Management Pack (MP) and mapped to the VM objects. Once in vROps, this new data about ports and/or services can be used as a Custom Group criteria set to populate it with VMs.

Besides leveraging vRIN data for Custom Groups, it can also be used in vROps to simply report on what types of applications are running in the virtualized environment if this information is not available in the organization’s CMDB.

Figure 17: EPOps Agent-Based Custom Groups – Example 1: VMs

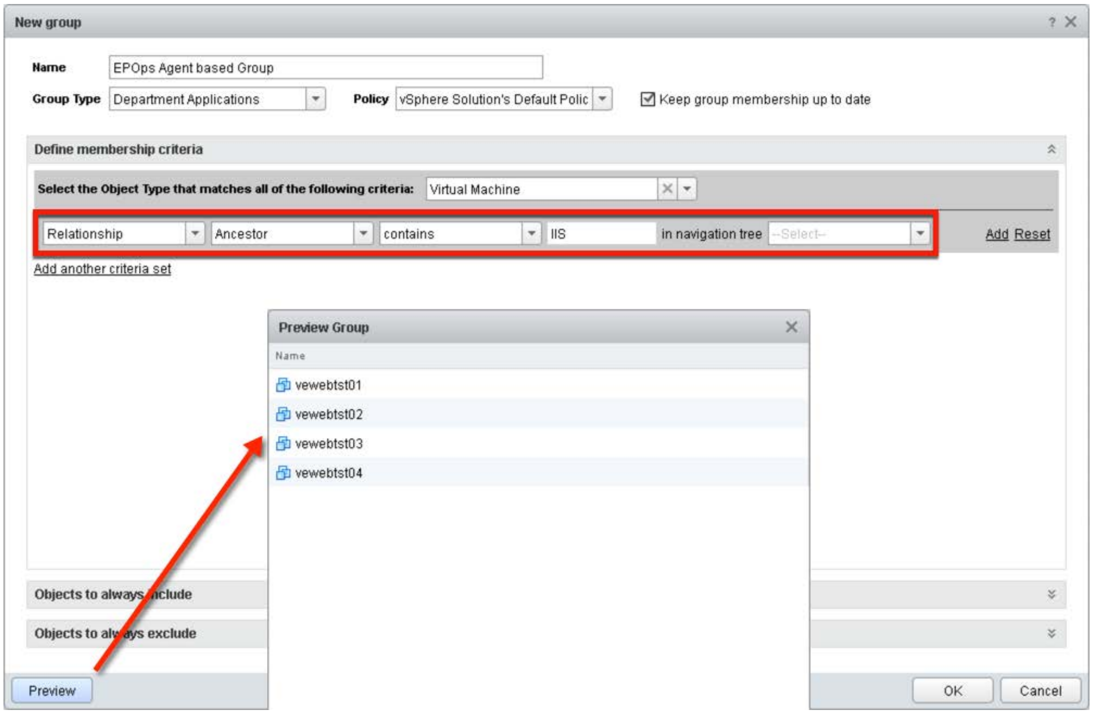

New group

Name: EPOps Agent based Group

Group Type: Department Applications

Policy: vSphere Solution's Default Polic

☑ Keep group membership up to date

Define membership criteria

Select the Object Type that matches all of the following criteria: Virtual Machine

Relationship | Ancestor | contains | IIS | in navigation tree | --Select--

Add another criteria set

Preview Group

Name

vewebtst01

vewebtst02

vewebtst03

vewebtst04

Objects to always include

Objects to always exclude

End-Point Operations (EPOps) Agent data and the relationships established between the VM, OS, and Services/Applications running on a VM can also be used as the criteria set for Custom Groups. In addition to using EPOps to determine VM group membership, the added application metric data can be used to create Super Metrics that will help monitor pools of actual application servers, not just VMs. That is a very powerful concept in itself as it allows for aggregating metrics from multiple application nodes and provides a summary.

Figure 18: EPOps Agent–Based Custom Groups – Example 2: Services

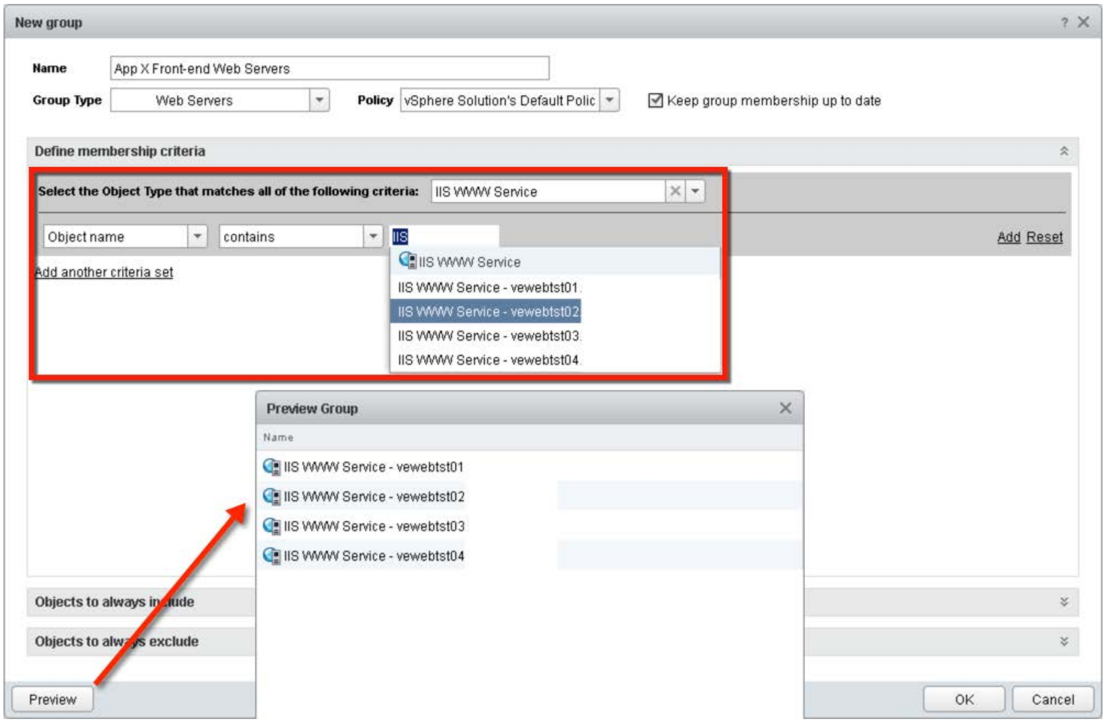
New group
Name
Group Type
Policy
Define membership criteria
Select the Object Type that matches all of the following criteria:
Add Reset
Add another criteria set
Preview Group
Name
Objects to always include
Objects to always exclude
Preview
OK
Cancel

Name: App X Front-end Web Servers

Group Type: Web Servers

Policy: vSphere Solution's Default Polic

☑ Keep group membership up to date

Select the Object Type that matches all of the following criteria: IIS WWW Service

Object name | contains | IIS

IIS WWW Service
IIS WWW Service - vewebtst01
IIS WWW Service - vewebtst02
IIS WWW Service - vewebtst03
IIS WWW Service - vewebtst04

Preview Group

IIS WWW Service - vewebtst01
IIS WWW Service - vewebtst02
IIS WWW Service - vewebtst03
IIS WWW Service - vewebtst04

Above is an example of what an application pool–based Custom Group would look like. Suppose there is an application stack with a pool of front-end web servers and the application owner would like to know how many total sessions are being handled by all the web servers, or which server has the worst response time. This may be very useful information when trying to keep an eye on the web farm utilization before making a scale-out decision. This concept can be extended to any application that is supported by the EPOps Agent, allowing the organization to go beyond the basic infrastructure and monitor various services as a single construct comprised of multiple servers and applications.

With a few examples of VM and application-based Custom Groups, the following section explains which Super Metrics are needed to answer the business-related questions mentioned earlier.

Creating Super Metrics

Super Metrics used for Custom Groups generally fall into three categories and align with the business-oriented goals to quantify performance, utilization, and consumption of resources by related VMs in Custom Groups.

To create a Super Metric, go to Content > Super Metrics > Add New Super Metric.

Figure 19: Create Super Metric Wizard

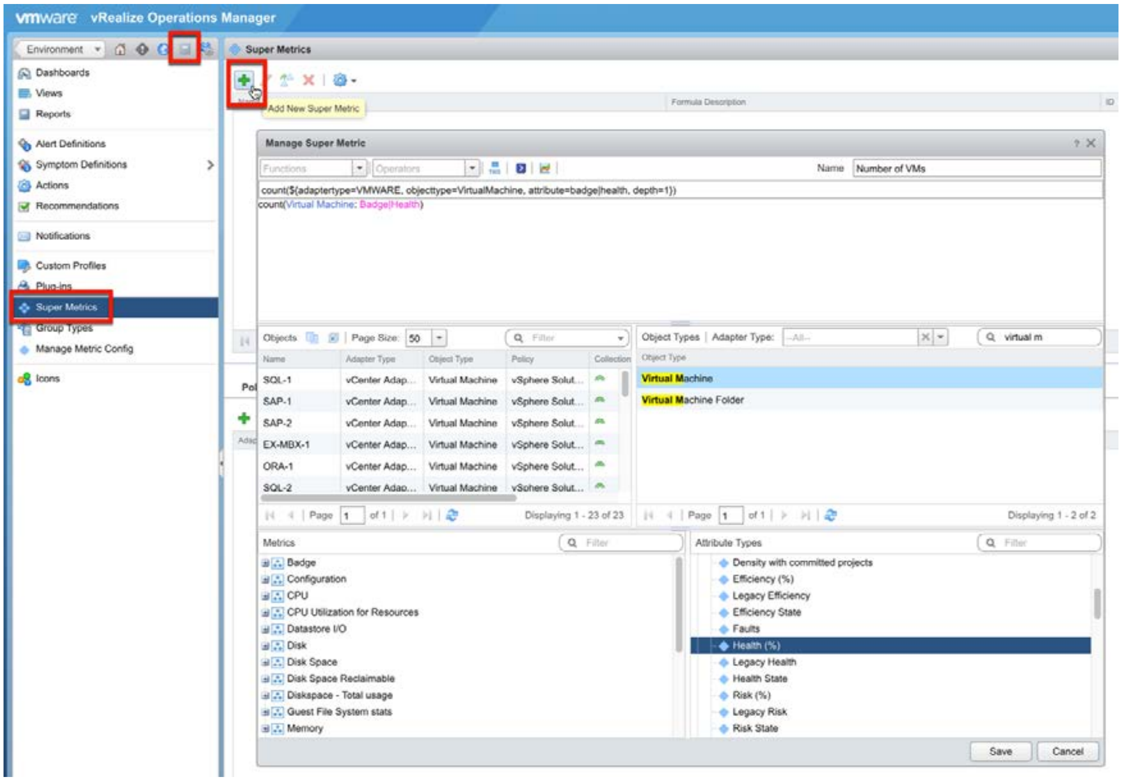

Name: Number of VMs

```
count(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=badge|health, depth=1})
count(Virtual Machine: Badge|Health)
```

Objects | Page Size: 50

Name	Adapter Type	Object Type	Policy	Collection
SQL-1	vCenter Adap...	Virtual Machine	vSphere Solut...	
SAP-1	vCenter Adap...	Virtual Machine	vSphere Solut...	
SAP-2	vCenter Adap...	Virtual Machine	vSphere Solut...	
EX-MBX-1	vCenter Adap...	Virtual Machine	vSphere Solut...	
ORA-1	vCenter Adap...	Virtual Machine	vSphere Solut...	
SQL-2	vCenter Adap...	Virtual Machine	vSphere Solut...	

Page 1 of 1 Displaying 1 - 23 of 23

Object Types | Adapter Type: --All-- | virtual m

Object Type
Virtual Machine
Virtual Machine Folder

Page 1 of 1 Displaying 1 - 2 of 2

Next are examples of Super Metrics that provide business-related answers to the team and application managers.

Quantity-Based Super Metrics

Following is a discussion about which Super Metrics may be useful from a Custom Group perspective to a cloud/virtualization administrator or business unit owner. Please note that many Super Metrics in this category are very similar to the vSphere Cluster-level metrics found under the Summary section.

Figure 20: vSphere Cluster Summary Metrics Example

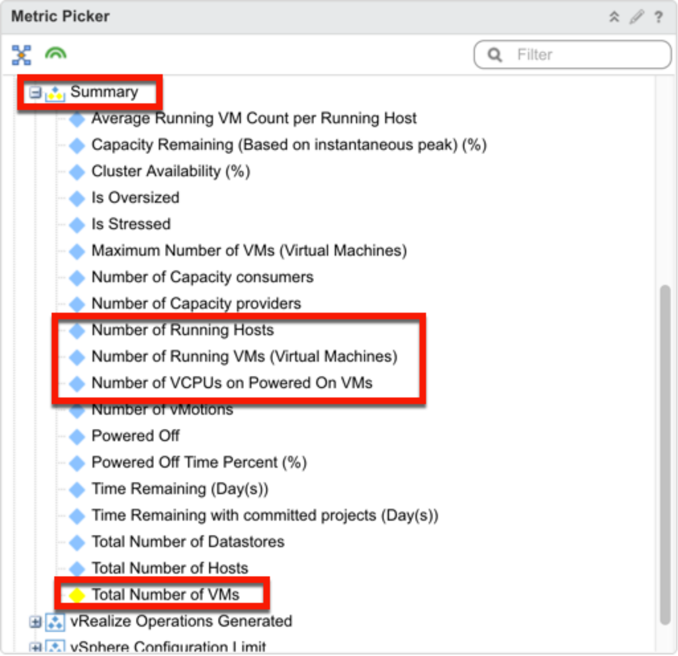

Metric Picker

Filter

- Summary
  - Average Running VM Count per Running Host
  - Capacity Remaining (Based on instantaneous peak) (%)
  - Cluster Availability (%)
  - Is Oversized
  - Is Stressed
  - Maximum Number of VMs (Virtual Machines)
  - Number of Capacity consumers
  - Number of Capacity providers
  - Number of Running Hosts
  - Number of Running VMs (Virtual Machines)
  - Number of VCPUs on Powered On VMs
  - Number of vMotions
  - Powered Off
  - Powered Off Time Percent (%)
  - Time Remaining (Day(s))
  - Time Remaining with committed projects (Day(s))
  - Total Number of Datastores
  - Total Number of Hosts
  - Total Number of VMs
- vRealize Operations Generated
- vSphere Configuration Limit

As mentioned earlier, a Custom Group is a container akin to a vSphere Cluster, therefore counting what is in the bucket is very natural. This is especially useful when management wants to know, for example, how many SQL VMs there are across all virtualized environments. Related questions arise, such as how many VMs are powered on vs. powered off, how many VMs are idle and/or oversized, etc. Following is a list of a few basic Super Metrics created for the previously defined Custom Groups:

1. Number of VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=badge|health, depth=1})
2. Number of Powered ON VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, metric=sys|poweredOn, depth=1, where="==1"})
3. Number of Powered OFF VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, metric=sys|poweredOn, depth=1, where="==0"})
4. Number of Idle VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=summary|idle, depth=1, where="==1"})

5. Number of Oversized VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, metric=summary|oversized, depth=1, where="==1"})

6. Number of Stressed VMs:
   count(${adaptertype=VMWARE, objecttype=VirtualMachine, metric=summary|isStress, depth=1, where="==1"})

Utilization-Based Super Metrics

The utilization-based Super Metrics measure how many resources have been allocated and are being used by the workloads. Thus, the percentage of the allocated resources being utilized is easily visible, providing insight into what percentage of the allocated capacity is actually being used.

1. CPU GHz Allocated to BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|actual.capacity.normalized, depth=1})/1000

2. CPU GHz Demanded by BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|demandmhz, depth=1})/1000

3. CPU GHz Used by BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|usagemhz_average, depth=1})/1000

4. CPU % Demanded in BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|demandmhz, depth=1})/1000

5. CPU % Used in BU:
   (${this, metric=Super Metric|sm_CPU_GHz_Demanded>}/${this, metric=Super Metric|sm_<CPU_GHz_Used>})*100

6. CPU Count in BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|corecount_provisioned, depth=1})

7. CPU Count Recommended to BU
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|size.recommendation, depth=1})

8. CPU Count Reclaimable from BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=cpu|wasteValue, depth=1})

9. Memory GB Allocated to BU:
   sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=mem|guest_provisioned, depth=1})/1048576

10. Memory GB Consumed by BU:
    sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=mem|consumed_average, depth=1})/1048576

11. Memory % Used by BU:
    sum(${this, metric=Super Metric|sm_Memory_GB_Consumed>}/${this, metric=Super Metric|sm_<Memeory_GB_ Allocated>})*100

12. Disk Space Provisioned to BU:
    sum(${adaptertype=VMWARE, objecttype=VirtualMachine,
    attribute=mem|consumed_average, depth=1})/1048576

13. Disk Space Used by BU:
    sum(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=diskspace|used,
    depth=1})

14. VM Snapshots GB in BU:
    sum(${adaptertype=VMWARE, objecttype=VirtualMachine,
    attribute=diskspace|snapshot, depth=1, where=">0"})

15. Disk Space % Used by BU:
    (${this, metric=Super Metric|sm_<Disk_Space_Used>}/${this, metric=Super
    Metric|sm_<Disk_Space_Provisioned>})*100

16. Total BU IOPS:
    sum(${adaptertype=VMWARE, objecttype=VirtualMachine,
    attribute=datastore|commandsAveraged_average, depth=1})

Performance-Based Super Metrics

The performance-based Super Metrics are intended to track and report any SLA breaches. For example, are there any VMs whose CPU Ready % is above a certain threshold? This can minimize the potential impact from poorly performing workloads, as well as track the SLA.

1. Highest VM CPU Contention %:
   max(${adaptertype=VMWARE, objecttype=VirtualMachine,
   attribute=cpu|capacity_contentionPct, depth=2})

2. Highest VM CPU Ready % per Core:
   max(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=Super
   Metric|sm_<VM_CPU_Ready_%_per_Core_Super_Metric>, depth=1})

3. Highest VM CPU Co-Stop %:
   max(${adapterkind=VMWARE, resourcekind=VirtualMachine,
   attribute=cpu|costopPct, depth=1})

4. Highest VM CPU Swap Wait %:
   max(${adapterkind=VMWARE, resourcekind=VirtualMachine,
   attribute=cpu|swapwaitPct, depth=1})

5. Highest VM CPU IO Wait %:
   max(${adapterkind=VMWARE, resourcekind=VirtualMachine,
   attribute=cpu|iowaitPct, depth=1})

6. Highest VM Mem Swap In Rate (KBps):
   max(${adaptertype=VMWARE, objecttype=VirtualMachine,
   attribute=mem|swapinRate_average, depth=1})

7. Highest VM Mem Decompression Rate (KBps):
   max(${adaptertype=VMWARE, objecttype=VirtualMachine,
   attribute=mem|decompressionRate_average, depth=1})

8. Highest VM Mem Ballooned (KB):
   max(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=mem|balloonPct, depth=1})
9. Highest VM Disk IOPS:
   max(${adapterkind=VMWARE, resourcekind=VirtualMachine, attribute=virtualDisk|commandsAveraged_average, depth=1})
10. Highest VM Disk Latency (ms):
    max(${adapterkind=VMWARE, resourcekind=VirtualMachine, attribute=datastore|totalLatency_average, depth=1})
11. Highest VM Network Packets Dropped:
    max(${adapterkind=VMWARE, resourcekind=VirtualMachine, attribute=net|dropped, depth=1})
12. Highest VM Network Usage Rate (KBps):
    max(${adaptertype=VMWARE, objecttype=VirtualMachine, attribute=net|usage_average, depth=1})

Application Pool Super Metrics

Providing a list of application-specific Super Metrics is more difficult as each application has a different set of metrics used as the base input. Following are suggestions based on MS IIS. This logic can be extrapolated and applied to other applications. The key point is that working with the specific service or application subject matter experts will help identify the Key Performance Indicators (KPIs) and thresholds important to assessing service health by the stakeholders.

1. Number of Sessions
2. Average Response Time
3. Highest Response time
4. Etc.

NOTE: This is by no means an exhaustive list of all possible Super Metrics. These are merely suggestions of the most common Super Metrics encountered in a typical environment. Many more Super Metrics can be created to answer various business questions and provide valuable insight.

NOTE: Please see the reference section of this white paper on how to create Super Metrics. This white paper only provides general Super Metric examples that can be created for Custom Groups.

NOTE: Super Metric naming convention is an individual preference, but the recommendation is to establish a convention and follow it. This makes it much easier to find, group, and filter Super Metrics later while using them.

Binding Custom Groups and Super Metrics

After defining the Super Metrics and Custom Groups, the next step involves binding them together. This process tells vROps for which objects to track Super Metrics, and where all the previous work of creating Group Types, Custom Groups, and Super Metrics comes together into a cohesive definition.

To bind the Super Metrics to their Custom Groups via their associated Group Type, select each Super Metric in the list and add each Group Type defined at the beginning of this process.

Figure 21: Binding Super Metrics to Custom Groups via Group Type

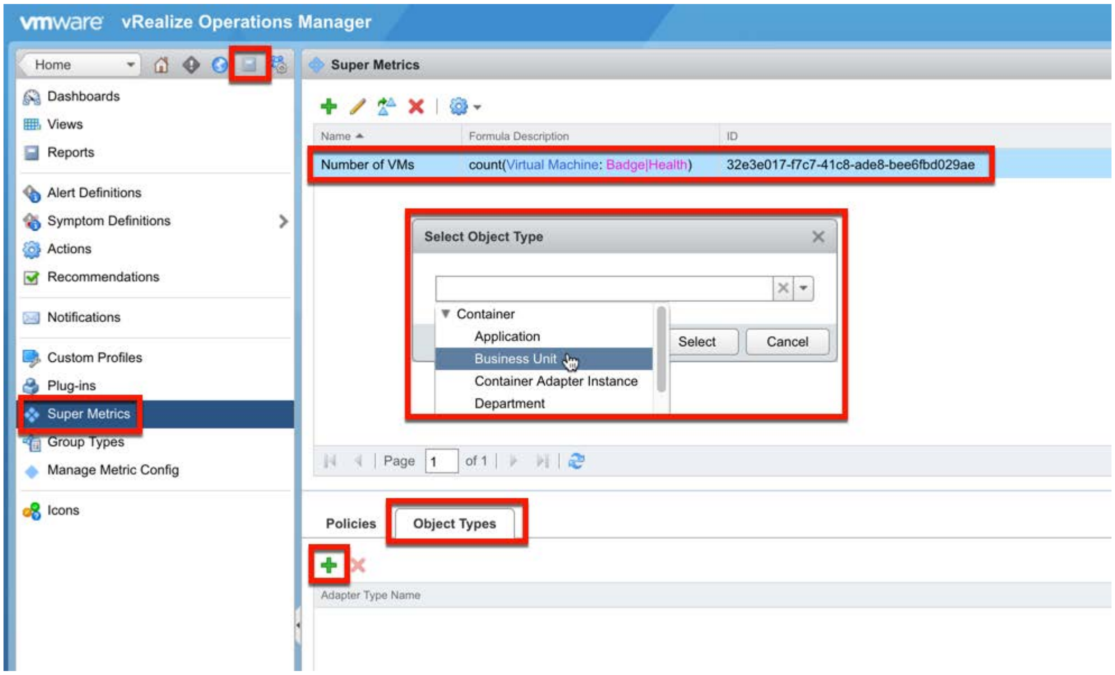

Name	Formula Description	ID
Number of VMs	count(Virtual Machine: Badge|Health)	32e3e017-f7c7-41c8-ade8-bee6fbd029ae

Enabling Super Metrics

Before Super Metrics can start collecting any data about the Custom Groups, they must be enabled in the policy.

Go to Administration > Policy > Policy Library, and then select and edit the policy that will track the new Super Metrics. On the Edit Monitoring Policy pop-up window, go to Collect Metrics and Properties, find Super Metrics created previously, and enable them.

Warning: Only enable Super Metrics for the respective Object Type. Enabling Super Metrics for All Object Types will collect extraneous data not relevant to most objects, use more disk space, clutter the user interface, and have a potential performance impact. For example, there is no need to capture CPU-related Super Metrics for Datastores.

Figure 22: Enabling Super Metrics in the Policy

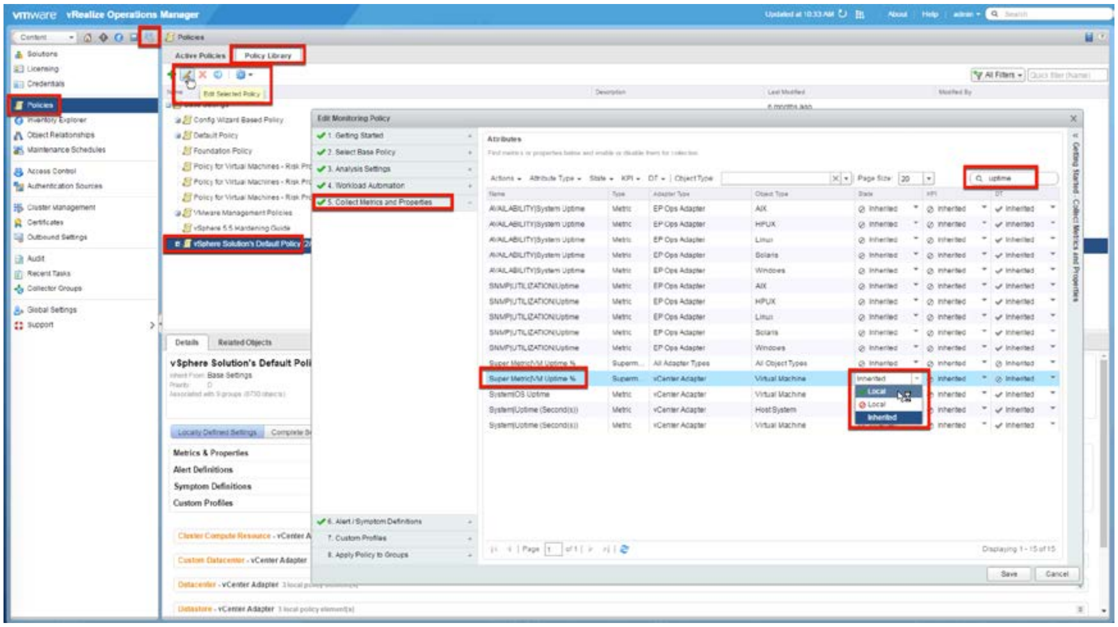

Attributes

Actions ▾ Attribute Type ▾ State ▾ KPI ▾ DT ▾ | Object Type | Page Size: 20 | uptime

Name	Type	Adapter Type	Object Type	State	KPI	DT
AVAILABILITY|System Uptime	Metric	EP Ops Adapter	AIX	Inherited	Inherited	✓ Inherited
AVAILABILITY|System Uptime	Metric	EP Ops Adapter	HPUX	Inherited	Inherited	✓ Inherited
AVAILABILITY|System Uptime	Metric	EP Ops Adapter	Linux	Inherited	Inherited	✓ Inherited
AVAILABILITY|System Uptime	Metric	EP Ops Adapter	Solaris	Inherited	Inherited	✓ Inherited
AVAILABILITY|System Uptime	Metric	EP Ops Adapter	Windows	Inherited	Inherited	✓ Inherited
SNMP|UTILIZATION|Uptime	Metric	EP Ops Adapter	AIX	Inherited	Inherited	✓ Inherited
SNMP|UTILIZATION|Uptime	Metric	EP Ops Adapter	HPUX	Inherited	Inherited	✓ Inherited
SNMP|UTILIZATION|Uptime	Metric	EP Ops Adapter	Linux	Inherited	Inherited	✓ Inherited
SNMP|UTILIZATION|Uptime	Metric	EP Ops Adapter	Solaris	Inherited	Inherited	✓ Inherited
SNMP|UTILIZATION|Uptime	Metric	EP Ops Adapter	Windows	Inherited	Inherited	✓ Inherited
Super Metric|VM Uptime %	Superm...	All Adapter Types	All Object Types	Inherited	Inherited	Inherited
Super Metric|VM Uptime %	Superm...	vCenter Adapter	Virtual Machine	Inherited ✓ Local Local Inherited	Inherited	Inherited
System|OS Uptime	Metric	vCenter Adapter	Virtual Machine		Inherited	✓ Inherited
System|Uptime (Second(s))	Metric	vCenter Adapter	Host System		Inherited	✓ Inherited
System|Uptime (Second(s))	Metric	vCenter Adapter	Virtual Machine		Inherited	✓ Inherited

With these useful business-oriented Super Metrics defined, enabled, and bound to a business unit–based Group Type, the next step involves using them in custom dashboards and reports.

Building Custom Dashboards and Reports

Well-designed Custom Dashboards should strive to provide all relevant information in a single pane-of-glass, be intuitive and easy to use, and offer all relevant information to satisfy a use case without having to navigate further. The goal should be to empower the user to make better business decisions without having to look for information in multiple places or master the intricacies of a product.

Well-designed Custom Dashboards are the culmination of forethought and planning, and combine all previously touched-on elements into a cohesive story that paints a picture and informs the user. These dashboard examples show how to put the newly defined Business Unit–based groups and Super Metrics to good use.

Figure 23: Business Unit Summary Dashboard Example

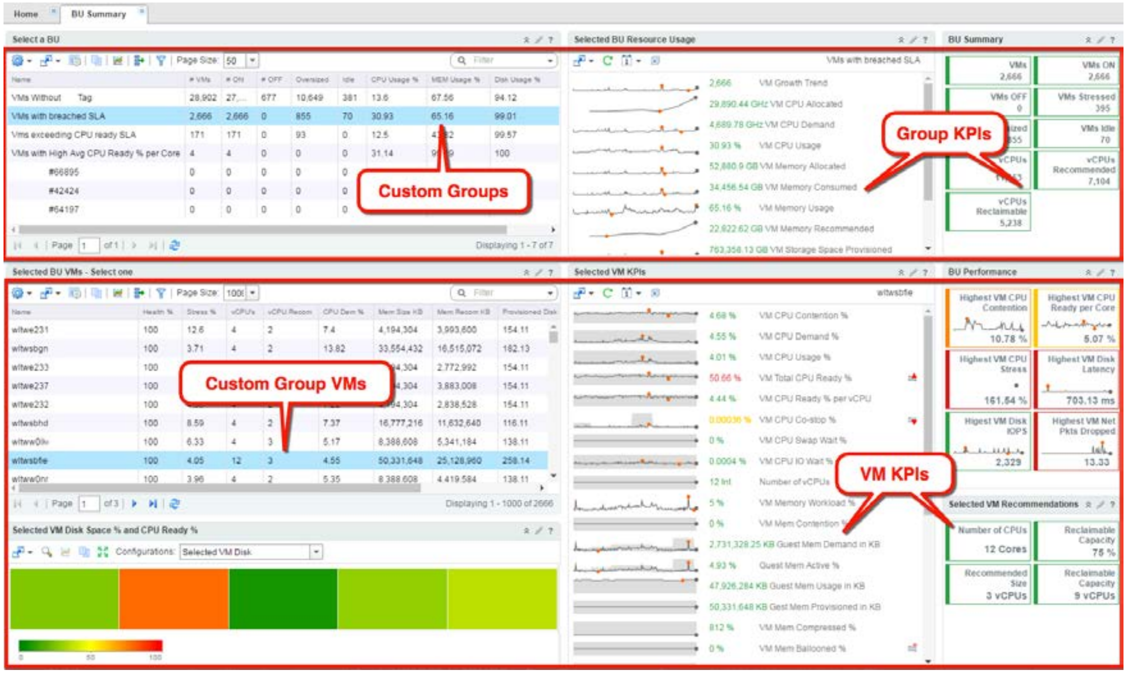

Select a BU

Page Size: 50

Name	# VMs	# ON	# OFF	Oversized	Idle	CPU Usage %	MEM Usage %	Disk Usage %
VMs Without Tag	28,902	27,...	677	10,649	381	13.6	67.56	94.12
VMs with breached SLA	2,666	2,666	0	855	70	30.93	65.16	99.01
Vms exceeding CPU ready SLA	171	171	0	93	0	12.5	[illegible]	99.57
VMs with High Avg CPU Ready % per Core	4	4	0	0	0	31.14	[illegible]	100
#66895	0	0	0	0	0			
#42424	0	0	0	0	0			
#64197	0	0	0	0	0			

Page 1 of 1 Displaying 1 - 7 of 7

Selected BU Resource Usage

VMs with breached SLA

2,666 VM Growth Trend
29,890.44 GHz VM CPU Allocated
4,689.78 GHz VM CPU Demand
30.93 % VM CPU Usage
52,880.9 GB VM Memory Allocated
34,456.54 GB VM Memory Consumed
65.16 % VM Memory Usage
22,822.62 GB VM Memory Recommended
763,358.13 GB VM Storage Space Provisioned

BU Summary

VMs 2,666
VMs ON 2,666
VMs OFF 0
VMs Stressed 395
[illegible]ized [illegible]55
VMs Idle 70
vCPUs [illegible]
vCPUs Recommended 7,104
vCPUs Reclaimable 5,238

Selected BU VMs - Select one

Page Size: 1000

Name	Health %	Stress %	vCPU's	vCPU Recom	CPU Dem %	Mem Size KB	Mem Recom KB	Provisioned Disk
witwe231	100	12.6	4	2	7.4	4,194,304	3,993,600	154.11
witwsbgn	100	3.71	4	2	13.82	33,554,432	16,515,072	182.13
witwe233	100	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	2,772,992	154.11
witwe237	100	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	3,883,008	154.11
witwe232	100	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	2,838,528	154.11
witwsbhd	100	8.59	4	2	7.37	16,777,216	11,632,640	116.11
witww0liv	100	6.33	4	3	5.17	8,388,608	5,341,184	138.11
witwsbfe	100	4.05	12	3	4.55	50,331,648	25,128,960	258.14
witww0nr	100	3.96	4	2	5.35	8,388,608	4,419,584	138.11

Page 1 of 3 Displaying 1 - 1000 of 2666

Selected VM Disk Space % and CPU Ready %

Configurations: Selected VM Disk

0 50 100

Selected VM KPIs

witwsbfe

4.68 % VM CPU Contention %
4.55 % VM CPU Demand %
4.01 % VM CPU Usage %
50.66 % VM Total CPU Ready %
4.44 % VM CPU Ready % per vCPU
0.00036 % VM CPU Co-stop %
0 % VM CPU Swap Wait %
0.0004 % VM CPU IO Wait %
12 Int Number of vCPUs
5 % VM Memory Workload %
0 % VM Mem Contention %
2,731,328.25 KB Guest Mem Demand in KB
4.93 % Guest Mem Active %
47,926,284 KB Guest Mem Usage in KB
50,331,648 KB Gest Mem Provisioned in KB
812 % VM Mem Compressed %
0 % VM Mem Ballooned %

BU Performance

Highest VM CPU Contention 10.78 %
Highest VM CPU Ready per Core 5.07 %
Highest VM CPU Stress 161.54 %
Highest VM Disk Latency 703.13 ms
Higest VM Disk IOPS 2,329
Highest VM Net Pkts Dropped 13.33

Selected VM Recommendations

Number of CPUs 12 Cores
Reclaimable Capacity 75 %
Recommended Size 3 vCPUs
Reclaimable Capacity 9 vCPUs

The Business Unit Summary dashboard can be used by business unit and application owners, as well as by operation teams. It provides summarized reporting at the container-level along with a breakdown of individual VMs and their infrastructure components. This powerful dashboard allows the user to go top-down, from a high-level Business Unit container summary at the top, to low-level virtual machine details below, all in a single pane-of-glass.

Figure 24: Business Unit Compute Analysis Dashboard Example

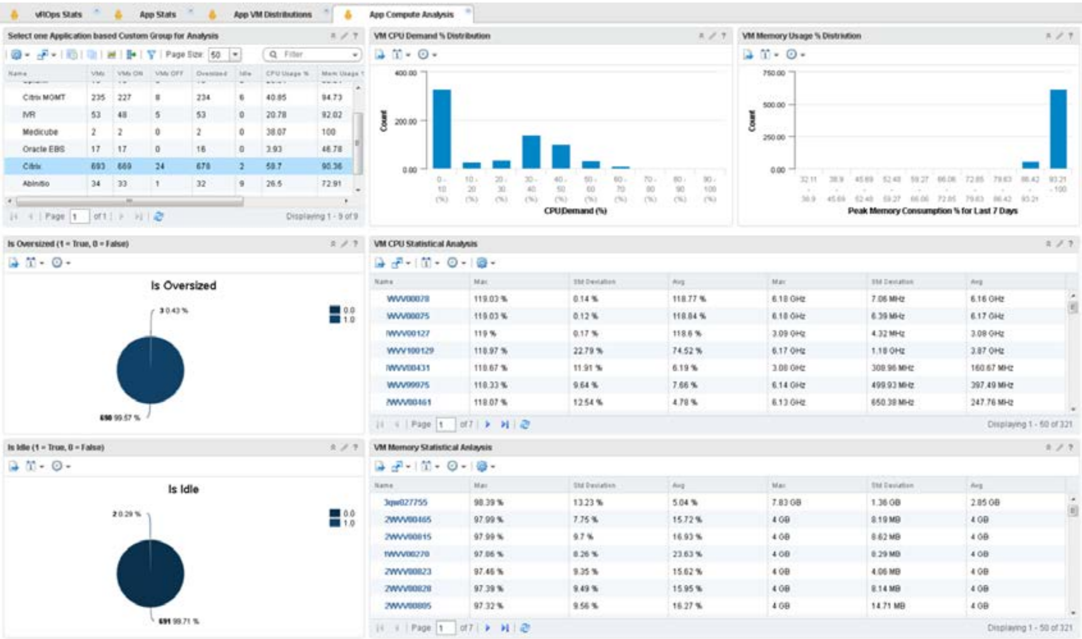

Select one Application based Custom Group for Analysis

Name	VMs	VMs ON	VMs OFF	Oversized	Idle	CPU Usage %	Mem Usage %
Citrix MOMT	235	227	8	234	6	40.85	94.73
IVR	53	48	5	53	0	20.78	92.02
Medicube	2	2	0	2	0	38.07	100
Oracle EBS	17	17	0	16	0	3.93	46.78
Citrix	693	669	24	678	2	58.7	90.36
AbInitio	34	33	1	32	9	26.5	72.91

Page 1 of 1 Displaying 1 - 9 of 9

VM CPU Statistical Analysis

Name	Max	Std Deviation	Avg	Max	Std Deviation	Avg
WVV00078	119.03 %	0.14 %	118.77 %	6.18 GHz	7.06 MHz	6.16 GHz
WVV00075	119.03 %	0.12 %	118.84 %	6.18 GHz	6.39 MHz	6.17 GHz
IWVV00127	119 %	0.17 %	118.6 %	3.09 GHz	4.32 MHz	3.08 GHz
WVV100129	118.97 %	22.79 %	74.52 %	6.17 GHz	1.18 GHz	3.87 GHz
IWVV00431	118.67 %	11.91 %	6.19 %	3.08 GHz	308.96 MHz	160.67 MHz
WVV99975	118.33 %	9.64 %	7.66 %	6.14 GHz	499.93 MHz	397.49 MHz
2WVV00461	118.07 %	12.54 %	4.78 %	6.13 GHz	650.38 MHz	247.76 MHz

Page 1 of 7 Displaying 1 - 50 of 321

VM Memory Statistical Anlaysis

Name	Max	Std Deviation	Avg	Max	Std Deviation	Avg
3qw027755	98.39 %	13.23 %	5.04 %	7.83 GB	1.36 GB	2.85 GB
2WVV00465	97.99 %	7.75 %	15.72 %	4 GB	8.19 MB	4 GB
2WVV00815	97.99 %	9.7 %	16.93 %	4 GB	8.62 MB	4 GB
1WVV00270	97.86 %	8.26 %	23.63 %	4 GB	8.29 MB	4 GB
2WVV00823	97.46 %	9.35 %	15.62 %	4 GB	4.06 MB	4 GB
2WVV00828	97.39 %	9.49 %	15.95 %	4 GB	8.14 MB	4 GB
2WVV00805	97.32 %	9.56 %	16.27 %	4 GB	14.71 MB	4 GB

Page 1 of 7 Displaying 1 - 50 of 321

The BU Compute Analysis Dashboards provides a statistical analysis of Virtual Machine CPU and Memory Demand. The dashboard brakes down CPU and Memory Demand by Percentile as well as calculates CPU and Memory Demand Standard Deviation for each virtual machine.

Figure 25: Business Unit Storage Analysis Dashboard Example

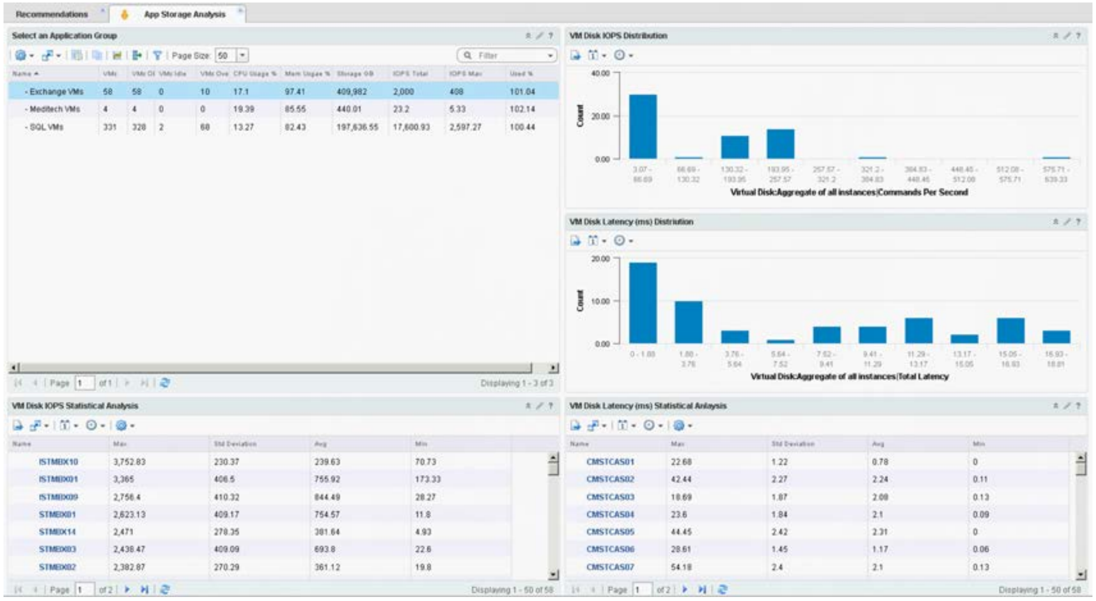

Select an Application Group

Name	VMs	VMs Ol	VMs Idle	VMs Ove	CPU Usage %	Mem Usage %	Storage GB	IOPS Total	IOPS Max	Used %
- Exchange VMs	58	58	0	10	17.1	97.41	409,982	2,000	408	101.04
- Meditech VMs	4	4	0	0	19.39	85.55	440.01	23.2	5.33	102.14
- SQL VMs	331	328	2	68	13.27	82.43	197,636.55	17,600.93	2,597.27	100.44

Page 1 of 1 Displaying 1 - 3 of 3

VM Disk IOPS Statistical Analysis

Name	Max	Std Deviation	Avg	Min
ISTMBX10	3,752.83	230.37	239.63	70.73
ISTMBX01	3,365	406.5	755.92	173.33
ISTMBX09	2,756.4	410.32	844.49	28.27
STMBX01	2,623.13	409.17	754.57	11.8
STMBX14	2,471	278.35	381.64	4.93
STMBX03	2,438.47	409.09	693.8	22.6
STMBX02	2,382.87	270.29	361.12	19.8

Page 1 of 2 Displaying 1 - 50 of 58

VM Disk Latency (ms) Statistical Anlaysis

Name	Max	Std Deviation	Avg	Min
CMSTCAS01	22.68	1.22	0.78	0
CMSTCAS02	42.44	2.27	2.24	0.11
CMSTCAS03	18.69	1.87	2.08	0.13
CMSTCAS04	23.6	1.84	2.1	0.09
CMSTCAS05	44.45	2.42	2.31	0
CMSTCAS06	28.61	1.45	1.17	0.06
CMSTCAS07	54.18	2.4	2.1	0.13

Page 1 of 2 Displaying 1 - 50 of 58

The BU Storage Analysis Dashboard provides statistical analysis of virtual machine Disk IOPS and Latency. The dashboard breaks down IOPS and Latency by Percentile, plus it calculates IOPS and Latency Standard Deviation for each virtual machine.

Figure 26: Business Unit Distributions Dashboard Example

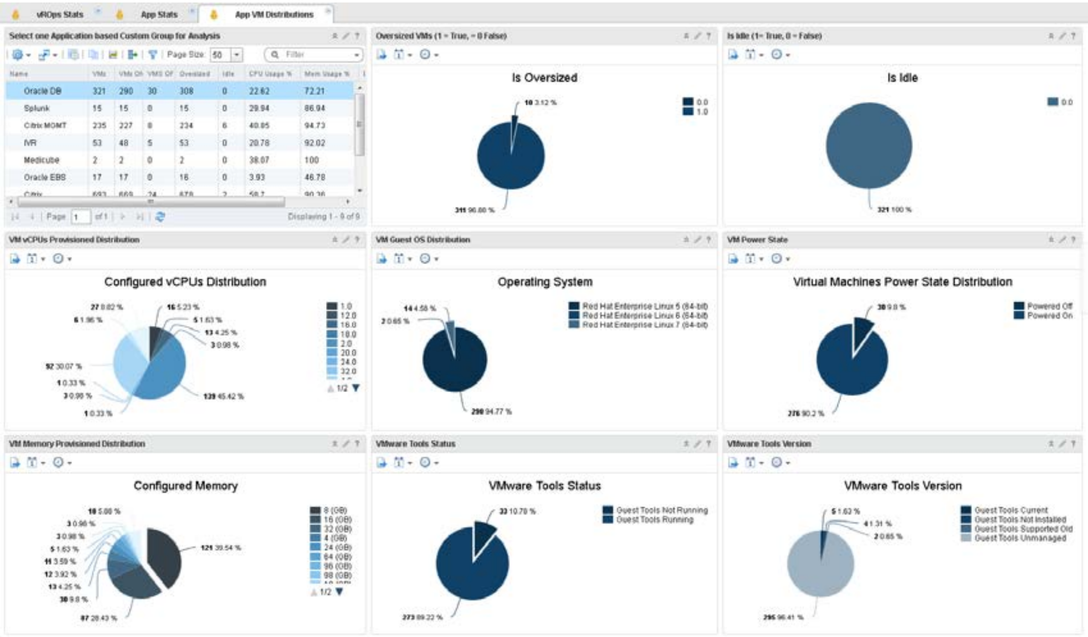

Select one Application based Custom Group for Analysis

Name	VMs	VMs On	VMS Off	Oversized	Idle	CPU Usage %	Mem Usage %
Oracle DB	321	290	30	308	0	22.62	72.21
Splunk	15	15	0	15	0	29.94	86.94
Citrix MGMT	235	227	8	234	6	40.85	94.73
IVR	53	48	5	53	0	20.78	92.02
Medicube	2	2	0	2	0	38.07	100
Oracle EBS	17	17	0	16	0	3.93	46.78
Citrix	693	669	24	678	2	58.7	90.36

Page 1 of 1 Displaying 1 - 9 of 9

The BU Distributions Dashboard provides the ability to see the breakdown of various types of systems in a Business Unit.

Figure 27: Business Unit Trends Dashboard Example

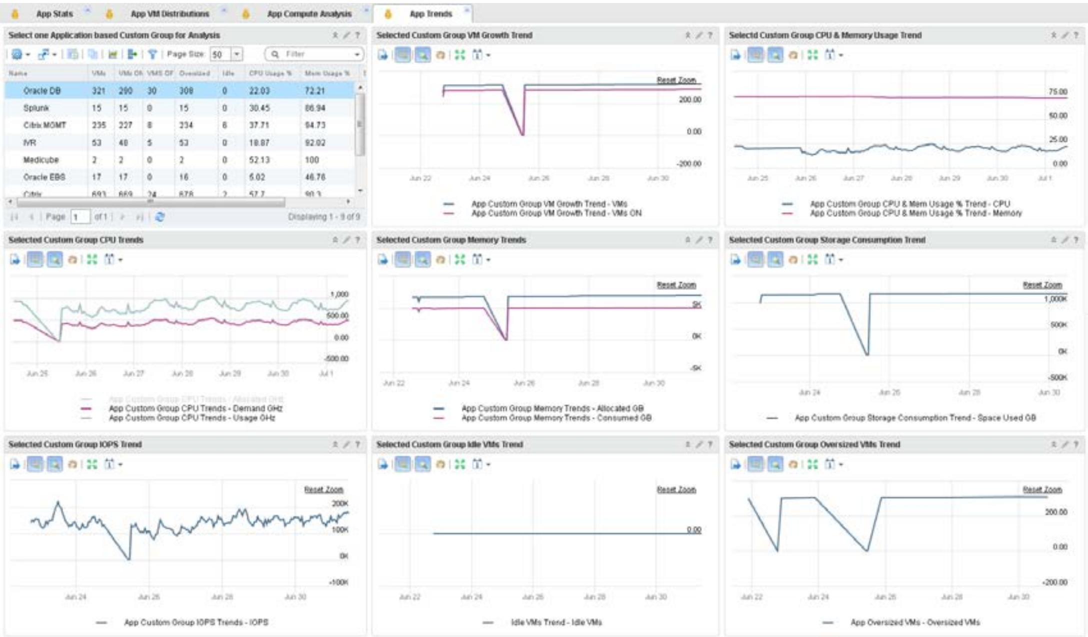

Select one Application based Custom Group for Analysis

Name	VMs	VMs On	VMS Off	Oversized	Idle	CPU Usage %	Mem Usage %
Oracle DB	321	290	30	308	0	22.03	72.21
Splunk	15	15	0	15	0	30.45	86.94
Citrix MGMT	235	227	8	234	6	37.71	94.73
IVR	53	48	5	53	0	18.87	92.02
Medicube	2	2	0	2	0	52.13	100
Oracle EBS	17	17	0	16	0	5.02	46.76
Citrix	693	669	24	678	2	57.7	90.3

Page 1 of 1 Displaying 1 - 9 of 9

App Stats
App VM Distributions
App Compute Analysis
App Trends
Page Size: 50
Filter
Selected Custom Group VM Growth Trend
Reset Zoom
App Custom Group VM Growth Trend - VMs
App Custom Group VM Growth Trend - VMs ON
Selectd Custom Group CPU & Memory Usage Trend
App Custom Group CPU & Mem Usage % Trend - CPU
App Custom Group CPU & Mem Usage % Trend - Memory
Selected Custom Group CPU Trends
App Custom Group CPU Trends - Demand GHz
App Custom Group CPU Trends - Usage GHz
Selected Custom Group Memory Trends
App Custom Group Memory Trends - Allocated GB
App Custom Group Memory Trends - Consumed GB
Selected Custom Group Storage Consumption Trend
App Custom Group Storage Consumption Trend - Space Used GB
Selected Custom Group IOPS Trend
App Custom Group IOPS Trends - IOPS
Selected Custom Group Idle VMs Trend
Idle VMs Trend - Idle VMs
Selected Custom Group Oversized VMs Trend
App Oversized VMs - Oversized VMs

The BU Trendlines Dashboard provides the ability to see the Workload growth as well as the CPU, Memory, Storage, and Network resources consumption trending over time.

Enhancing vRealize Automation Reporting

The Management Pack for vRealize Automation (vRA) extends operational management capabilities of the vRealize Operations Manager platform to provide tenant-aware operational visibility of the infrastructure supporting private clouds to cloud provider administrators. In addition to the 3 OOTB dashboards, this MP extends operational management platform, analytics and custom dashboarding capabilities of vROps to vRA's cloud constructs like Fabric Groups, Tenants, Business Groups, Reservations, etc.

The management pack provides many metrics and properties associated with vRA objects OOTB; however, the concepts explained in this white paper can be used to enhance the management pack functionality. Super Metrics similar to those for Custom Groups can be created to track additional information about vRA-related objects. Additional dashboards can be designed to provide more detailed information around vRA objects to specific stakeholders, further extending vROps functionality and reach.

Figure 28: vRA MP Object Types and Metrics

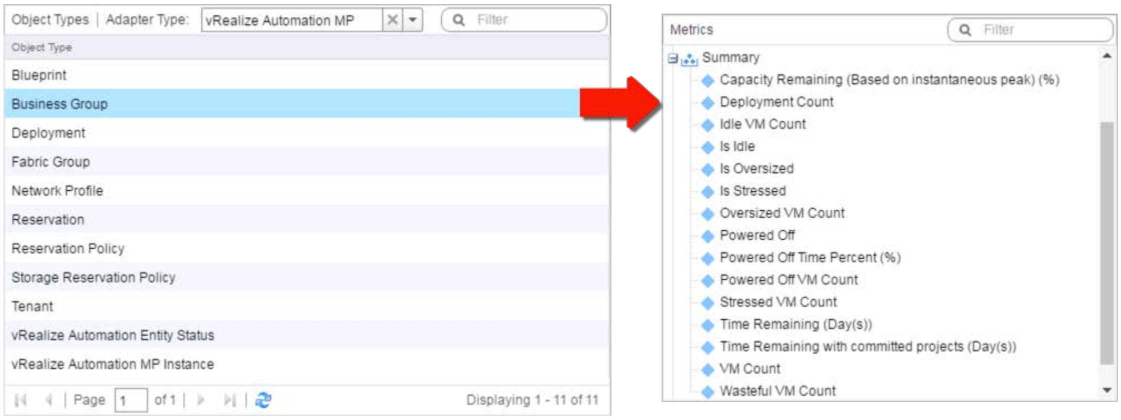

Object Type

Blueprint
Business Group
Deployment
Fabric Group
Network Profile
Reservation
Reservation Policy
Storage Reservation Policy
Tenant
vRealize Automation Entity Status
vRealize Automation MP Instance

Summary

- Capacity Remaining (Based on instantaneous peak) (%)
- Deployment Count
- Idle VM Count
- Is Idle
- Is Oversized
- Is Stressed
- Oversized VM Count
- Powered Off
- Powered Off Time Percent (%)
- Powered Off VM Count
- Stressed VM Count
- Time Remaining (Day(s))
- Time Remaining with committed projects (Day(s))
- VM Count
- Wasteful VM Count

See the reference section for additional vRA Management Pack documentation.

Authorizing Users and Providing Role-Based Access

With the Custom Groups, Super Metrics, and custom dashboards ready, it is time to provide secure access to the business unit members. The objective is to streamline the experience for the users by providing only the information they need, while minimizing the number of clicks and ensuring privileged access.

Figure 29: Privileged Access Diagram

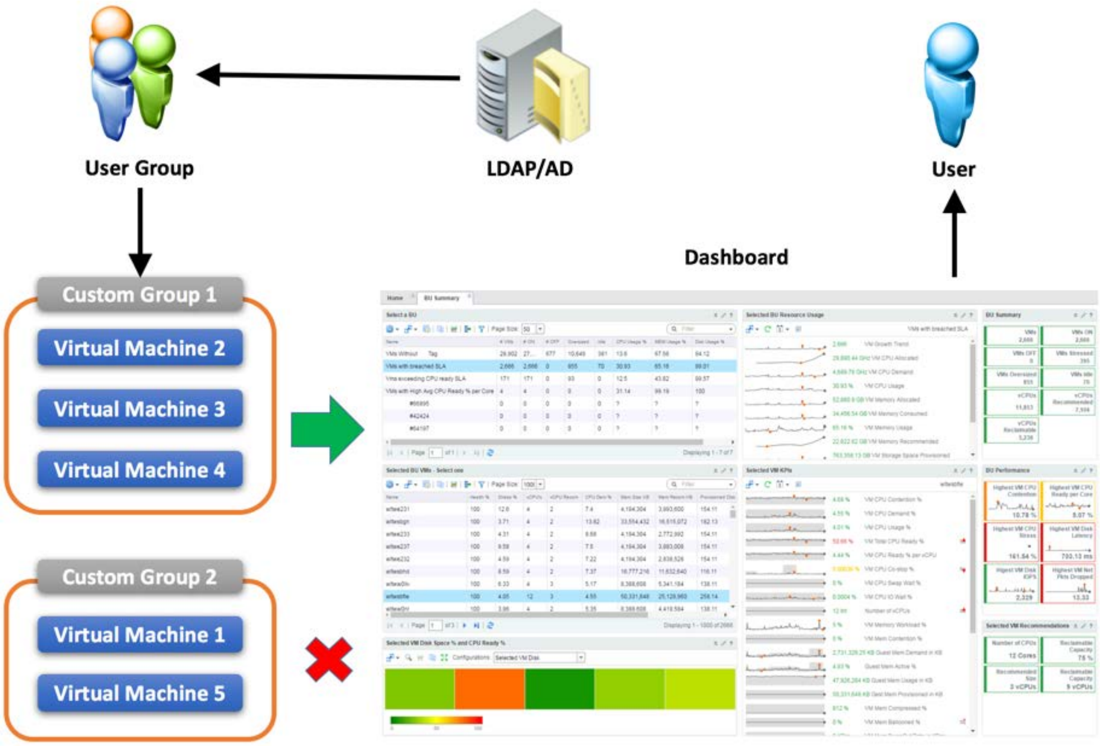

When this process is completed, the experience should be such that when the users log in, they are taken to the BU Summary dashboard and can only see the Custom Groups and Objects they have permissions to view. In effect, this will make the BU Summary dashboard just a reusable shell for any combination of Custom Groups and User Groups, thus simplifying ongoing support. With solely a single dashboard to maintain, it is much easier to make changes in the future that will be available to all users.

Configuring AD Authentication

vROps provides flexible options to authenticate users including Local, vCenter, SSO SAML, and Active Directory (AD). This example focuses on using Active Directory as it is the most prevalent directory service in today's enterprise.

Enabling AD integration is a relatively simple process in vROps and can be accomplished under Administration > Authentication Sources > and clicking Add. Fill in all required fields, click Test to validate and then OK to save.

Figure 30: Active Directory Authentication Configuration

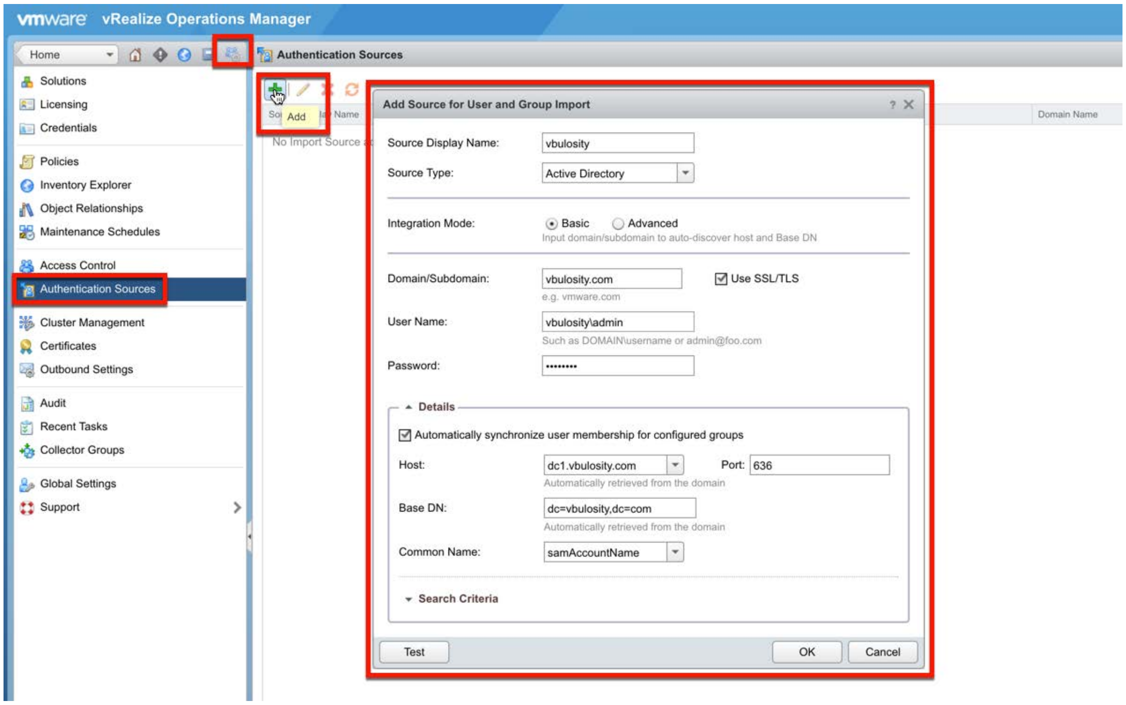

Add Source for User and Group Import

Source Display Name: vbulosity

Source Type: Active Directory

Integration Mode: ● Basic ○ Advanced
Input domain/subdomain to auto-discover host and Base DN

Domain/Subdomain: vbulosity.com ☑ Use SSL/TLS
e.g. vmware.com

User Name: vbulosity\admin
Such as DOMAIN\username or admin@foo.com

Password: ••••••••

Details

☑ Automatically synchronize user membership for configured groups

Host: dc1.vbulosity.com Port: 636
Automatically retrieved from the domain

Base DN: dc=vbulosity,dc=com
Automatically retrieved from the domain

Common Name: samAccountName

Search Criteria

Test OK Cancel

Creating a New Role

To simplify the user experience by providing only a single dashboard to view, all administrative and content creation functions must be disabled. To accomplish this, a special user role must be created with very limited privileges. This role will be even more restrictive than the Built-in Read Only one.

Go to Administration > Access Control > Roles and click add. Provide the name for the new role on the Create Role window and click OK.

Figure 31: Create User Role

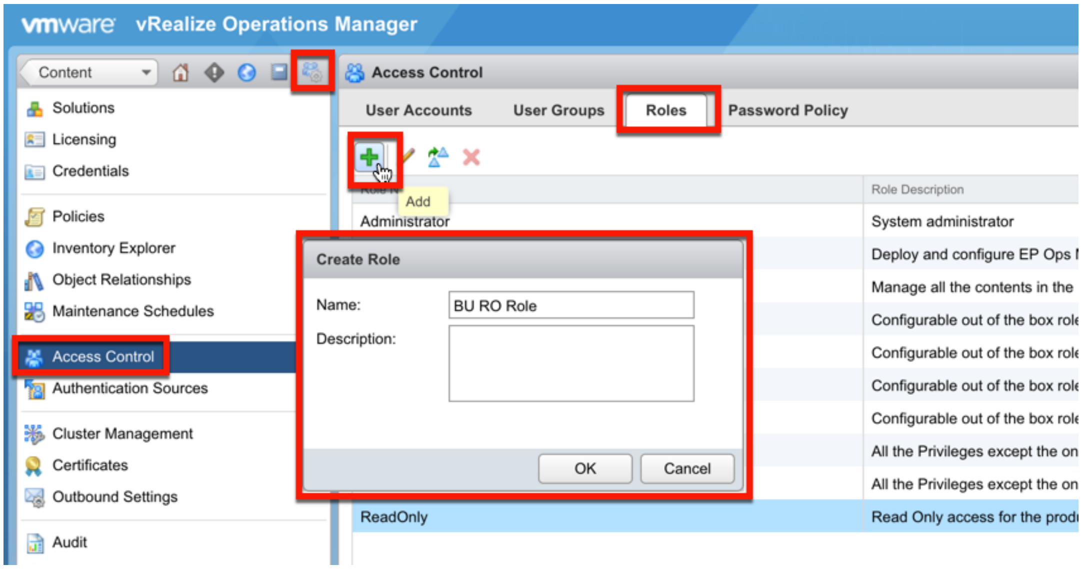

Create Role

Name: BU RO Role

Description:

Role Name	Role Description
Administrator	System administrator
	Deploy and configure EP Ops I
	Manage all the contents in the
	Configurable out of the box role
	Configurable out of the box role
	Configurable out of the box role
	Configurable out of the box role
	All the Privileges except the on
	All the Privileges except the on
ReadOnly	Read Only access for the prod

Once the role has been created, select it and click Permissions. On the “Assign Permissions To Role” window, select the following permissions and click Update:

1. Administration > Login Interactively
2. Content > Dashboard Management > Read and View Dashboards Page
3. Environment > View Dashboard Home Page

Figure 32: Assign Permissions To Role

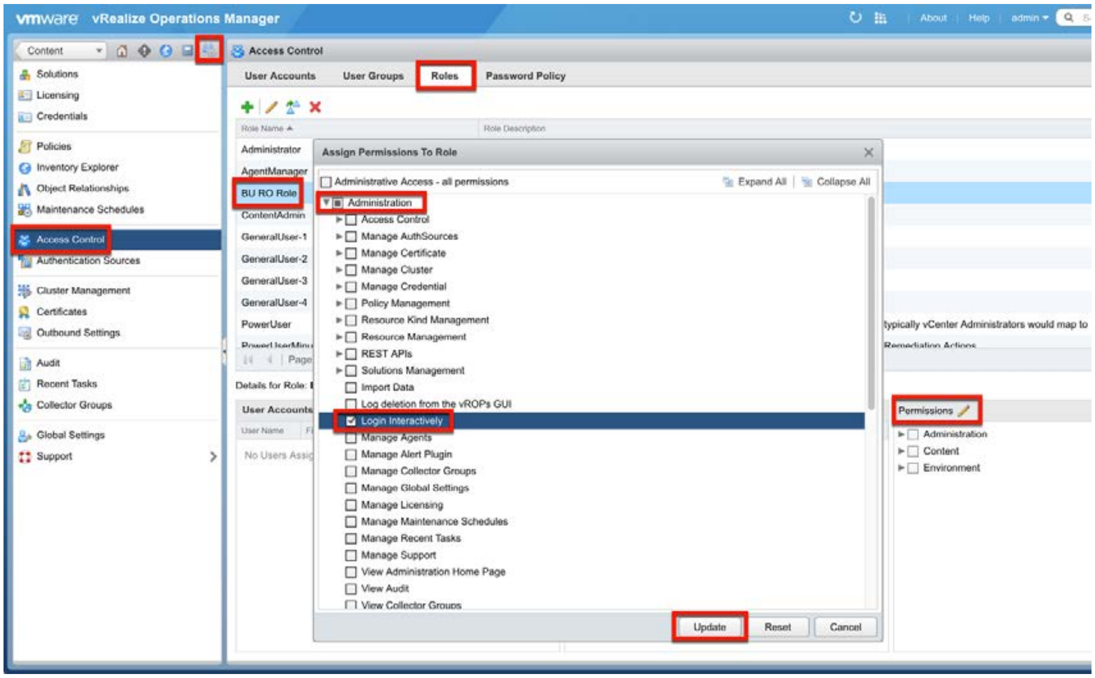

Assign Permissions To Role

- ☐ Administrative Access - all permissions
- ▼ ■ Administration
  - ▸ ☐ Access Control
  - ▸ ☐ Manage AuthSources
  - ▸ ☐ Manage Certificate
  - ▸ ☐ Manage Cluster
  - ▸ ☐ Manage Credential
  - ▸ ☐ Policy Management
  - ▸ ☐ Resource Kind Management
  - ▸ ☐ Resource Management
  - ▸ ☐ REST APIs
  - ▸ ☐ Solutions Management
  - ☐ Import Data
  - ☐ Log deletion from the vROPs GUI
  - ☑ Login Interactively
  - ☐ Manage Agents
  - ☐ Manage Alert Plugin
  - ☐ Manage Collector Groups
  - ☐ Manage Global Settings
  - ☐ Manage Licensing
  - ☐ Manage Maintenance Schedules
  - ☐ Manage Recent Tasks
  - ☐ Manage Support
  - ☐ View Administration Home Page
  - ☐ View Audit
  - ☐ View Collector Groups

NOTE: This screenshot does not show all permissions needed due to size.

Managing User Groups

With AD integration and User Roles in place, the next step involves setting up access for users and groups. vROps supports both local users and groups, and can import users and groups from corporate directories. This example showcases importing AD Security Groups since they reduce the management burden in vROps and are popular in enterprise settings by allowing users to leverage their existing AD accounts.

To import an AD Group, use the Import Group command in Administration > Access Control > User Groups.

Figure 33: Import Active Directory Groups

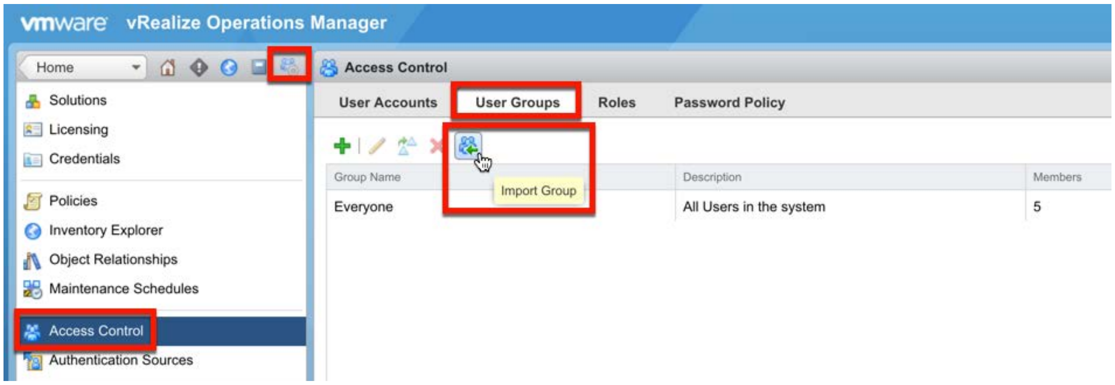

Group Name	Description	Members
Everyone	All Users in the system	5

On the Import User Groups wizard, select the directory configured previously in the Configure AD Authentication subsection and search for group(s) to be imported. Select groups to import and click Next.

Figure 34: Import Users and Groups

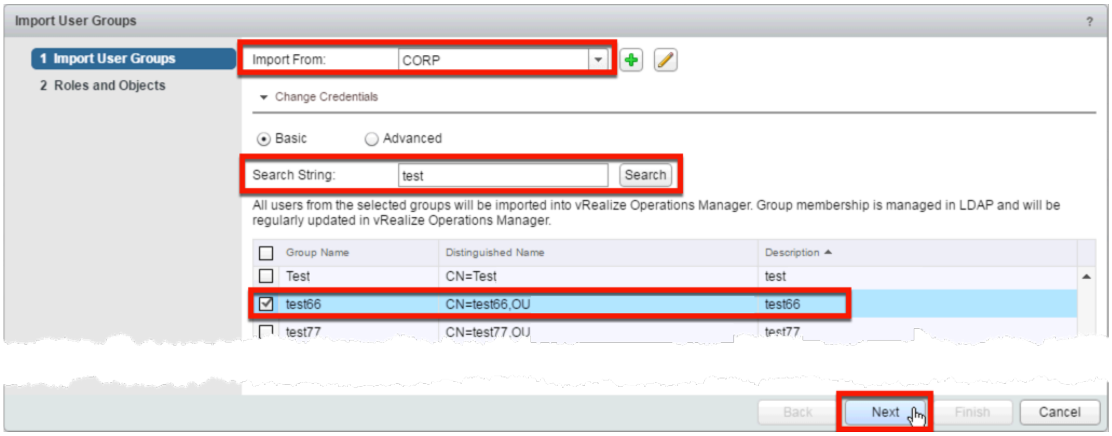

Import From: CORP

● Basic ○ Advanced

Search String: test

All users from the selected groups will be imported into vRealize Operations Manager. Group membership is managed in LDAP and will be regularly updated in vRealize Operations Manager.

□	Group Name	Distinguished Name	Description ▲
□	Test	CN=Test	test
☑	test66	CN=test66,OU	test66
□	test77	CN=test77,OU	test77

On the Roles and Objects step, select the role (BU RO Role) created earlier in the "Created New Role" subsections, and check the "Assign this role to group" checkbox. Click Custom Groups, under the Selected Objects expand the Group Type (Business Unit) created previously, and then select the Custom Group(s) the users in the imported group should have permissions to view. Select the Propagation checkmark next to the group name in order to allow users access to objects that are members of the custom group. Then click Finish.

Figure 35: User Roles and Object Permissions

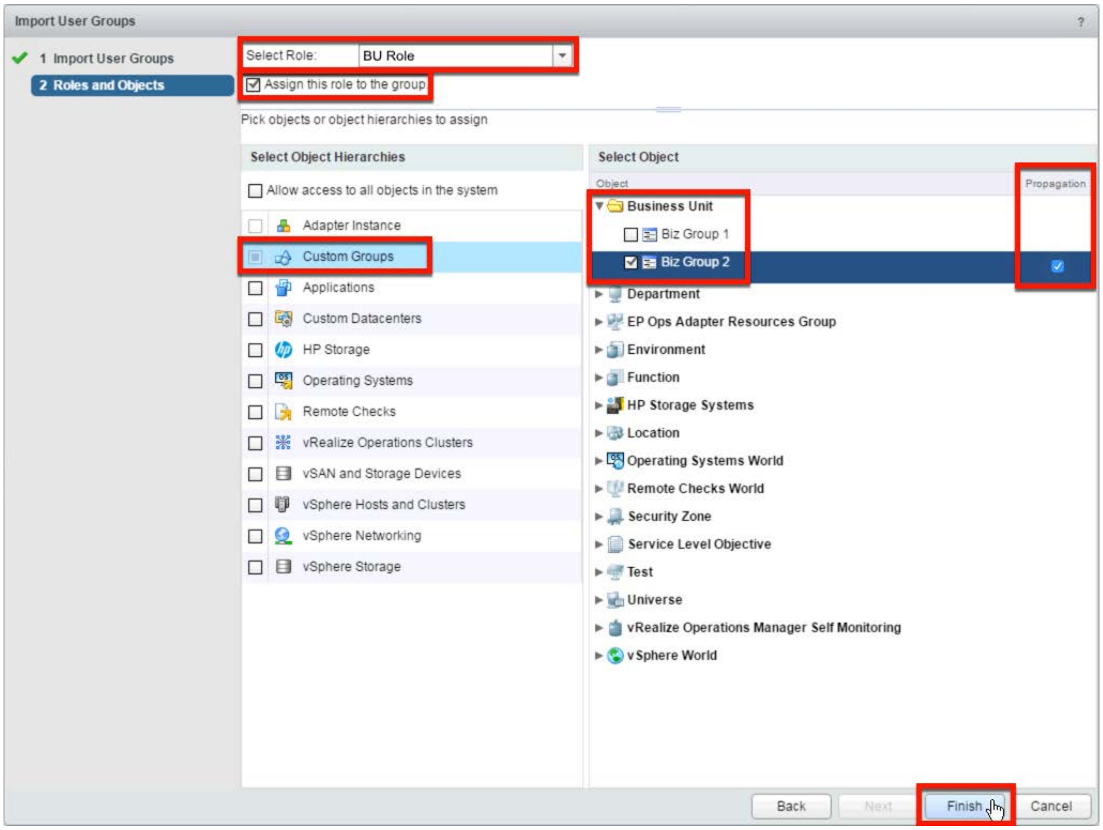

Import User Groups ?

✓ 1 Import User Groups

2 Roles and Objects

Select Role: BU Role

☑ Assign this role to the group

Pick objects or object hierarchies to assign

Select Object Hierarchies

□ Allow access to all objects in the system

- □ Adapter Instance
- ■ Custom Groups
- □ Applications
- □ Custom Datacenters
- □ HP Storage
- □ Operating Systems
- □ Remote Checks
- □ vRealize Operations Clusters
- □ vSAN and Storage Devices
- □ vSphere Hosts and Clusters
- □ vSphere Networking
- □ vSphere Storage

Select Object

Object	Propagation
▼ Business Unit	
□ Biz Group 1	
☑ Biz Group 2	☑
▶ Department	
▶ EP Ops Adapter Resources Group	
▶ Environment	
▶ Function	
▶ HP Storage Systems	
▶ Location	
▶ Operating Systems World	
▶ Remote Checks World	
▶ Security Zone	
▶ Service Level Objective	
▶ Test	
▶ Universe	
▶ vRealize Operations Manager Self Monitoring	
▶ vSphere World	

With AD Authentication configured, the user Role created, User Groups imported, and Custom Group permissions set, this completes the process of providing secure access to Business-Oriented Infrastructure Reporting with vROps.

Sharing Custom Dashboards

With everything in place, the last task that remains involves sharing custom dashboards with the imported AD Groups from the previous step. This is a rather trivial step that is necessary in order to expose content to the authorized users.

To share dashboards, go to Content > Dashboards > Actions > Share Dashboards.

Figure 36: Share Dashboards

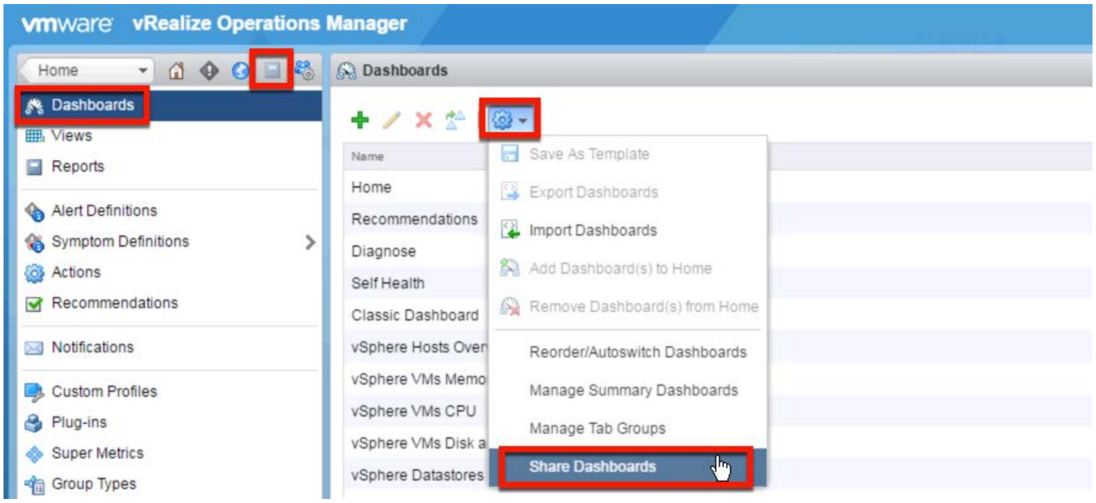

On the "Share Dashboard" wizard, locate the dashboard to be shared on the left side, select it, drag it onto one of the imported AD Groups, and let go. When successful, the Dashboards Assigned column will display the number of shared dashboards with that group, and the dashboard Shared icon will not be crossed out. Click Save to apply the new configuration.

Figure 37: Share Custom Dashboard with User Group

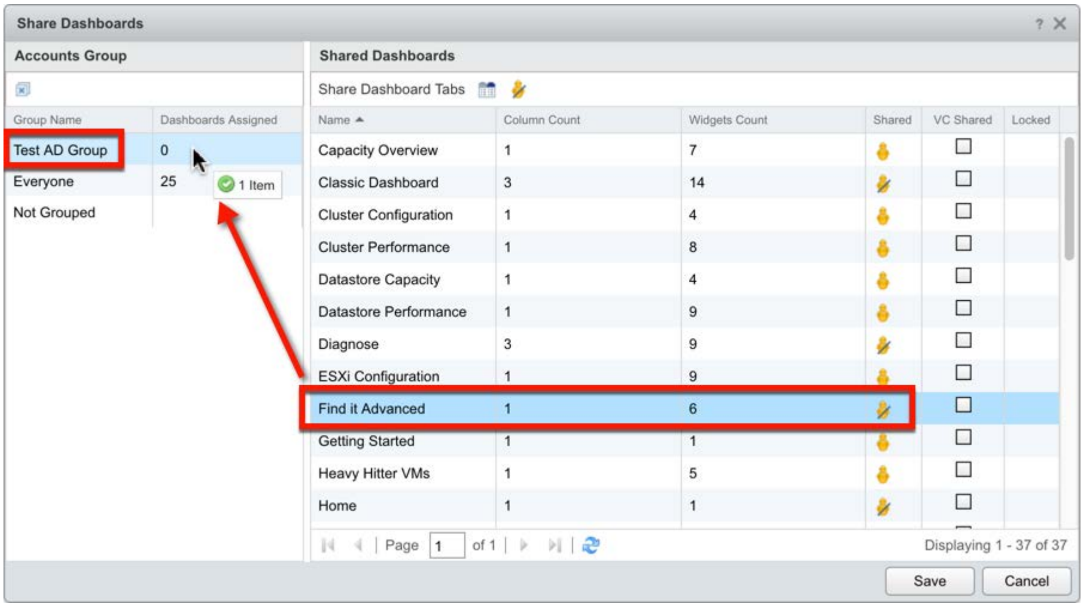

Group Name	Dashboards Assigned
Test AD Group	0
Everyone	25
Not Grouped	

Name	Column Count	Widgets Count	Shared	VC Shared	Locked
Capacity Overview	1	7	●	□	
Classic Dashboard	3	14	⊘	□	
Cluster Configuration	1	4	●	□	
Cluster Performance	1	8	●	□	
Datastore Capacity	1	4	●	□	
Datastore Performance	1	9	●	□	
Diagnose	3	9	⊘	□	
ESXi Configuration	1	9	●	□	
Find it Advanced	1	6	⊘	□	
Getting Started	1	1	●	□	
Heavy Hitter VMs	1	5	●	□	
Home	1	1	⊘	□	

In addition to sharing the BU Summary dashboard, all dashboards should be unshared with the Everyone group. This will further reduce the Dashboard List menu options.

Figure 38: Stop Sharing Dashboards with Everyone Group

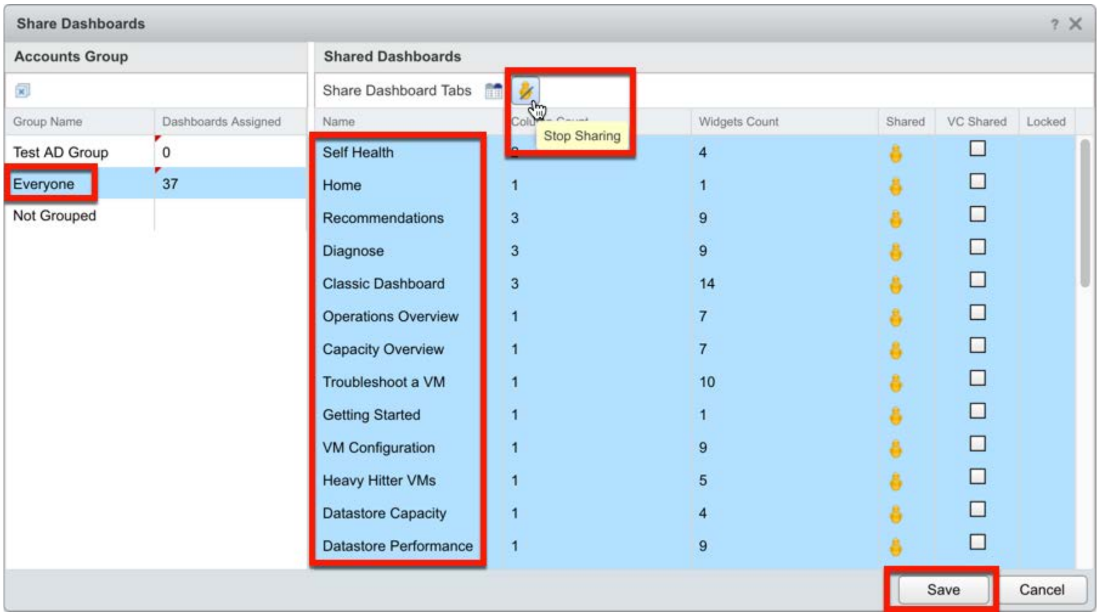

Share Dashboards

Accounts Group

Group Name	Dashboards Assigned
Test AD Group	0
Everyone	37
Not Grouped	

Shared Dashboards

Share Dashboard Tabs — Stop Sharing

Name	Columns Count	Widgets Count	Shared	VC Shared	Locked
Self Health	[illegible]	4	●	☐	
Home	1	1	●	☐	
Recommendations	3	9	●	☐	
Diagnose	3	9	●	☐	
Classic Dashboard	3	14	●	☐	
Operations Overview	1	7	●	☐	
Capacity Overview	1	7	●	☐	
Troubleshoot a VM	1	10	●	☐	
Getting Started	1	1	●	☐	
VM Configuration	1	9	●	☐	
Heavy Hitter VMs	1	5	●	☐	
Datastore Capacity	1	4	●	☐	
Datastore Performance	1	9	●	☐	

The end-result should be a clean, streamlined, and unobstructed access dashboards shared with the user groups. All other navigation options and dashboards are hidden providing simple and focused view for the end user.

Figure 39: Example of Limited Access

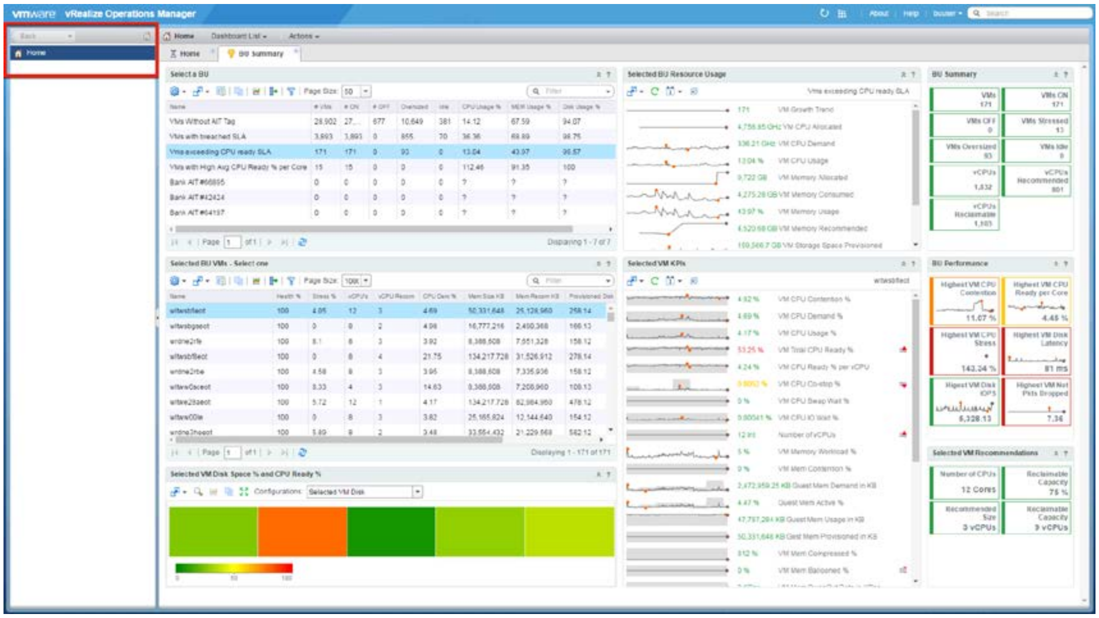

vmware vRealize Operations Manager
About | Help | bcuser ▾ | Search
Back
Home
Home | Dashboard List ▾ | Actions ▾
Home | BU Summary

Select a BU

Page Size: 50

Name	# VMs	# ON	# OFF	Oversized	Idle	CPU Usage %	MEM Usage %	Disk Usage %
VMs Without AIT Tag	28,902	27,...	677	10,649	381	14.12	67.59	94.07
VMs with breached SLA	3,893	3,893	0	855	70	36.36	69.89	96.75
Vms exceeding CPU ready SLA	171	171	0	93	0	13.04	43.97	96.57
VMs with High Avg CPU Ready % per Core	15	15	0	0	0	112.46	91.35	100
Bank AIT #66895	0	0	0	0	0	?	?	?
Bank AIT #42424	0	0	0	0	0	?	?	?
Bank AIT #64137	0	0	0	0	0	?	?	?

Page 1 of 1 — Displaying 1 - 7 of 7

Selected BU VMs - Select one

Page Size: 100K

Name	Health %	Stress %	vCPU's	vCPU Recom	CPU Dem %	Mem Size KB	Mem Recom KB	Provisioned Disk
wltwsbfleot	100	4.05	12	3	4.69	50,331,648	25,128,960	258.14
wltwsbgseot	100	0	8	2	4.98	16,777,216	2,450,368	166.13
wrdne2rfe	100	8.1	8	3	3.92	8,388,608	7,651,328	158.12
wltwsbf8eot	100	0	8	4	21.75	134,217,728	31,526,912	278.14
wrdne2rbe	100	4.58	8	3	3.95	8,388,608	7,335,936	158.12
wltwx0aceot	100	8.33	4	3	14.63	8,388,608	7,208,960	108.13
wltwe28aeot	100	5.72	12	1	4.17	134,217,728	82,984,960	478.12
wltwx00le	100	0	8	3	3.82	25,165,824	12,144,640	154.12
wrdne2hoeot	100	5.89	8	2	3.48	33,554,432	21,229,568	582.12

Page 1 of 1 — Displaying 1 - 171 of 171

Selected VM Disk Space % and CPU Ready %

Configurations: Selected VM Disk

Selected BU Resource Usage — Vms exceeding CPU ready SLA

171	VM Growth Trend
4,758.85 GHz	VM CPU Allocated
336.21 GHz	VM CPU Demand
13.04 %	VM CPU Usage
9,722 GB	VM Memory Allocated
4,275.28 GB	VM Memory Consumed
43.97 %	VM Memory Usage
4,520.69 GB	VM Memory Recommended
159,566.7 GB	VM Storage Space Provisioned

Selected VM KPIs — wltwsbfleot

4.82 %	VM CPU Contention %
4.69 %	VM CPU Demand %
4.17 %	VM CPU Usage %
53.25 %	VM Total CPU Ready %
4.24 %	VM CPU Ready % per vCPU
0.8092 %	VM CPU Co-stop %
0 %	VM CPU Swap Wait %
0.00041 %	VM CPU IO Wait %
12 ...	Number of vCPUs
5 %	VM Memory Workload %
0 %	VM Mem Contention %
2,472,959.25 KB	Guest Mem Demand in KB
4.47 %	Guest Mem Active %
47,797,284 KB	Guest Mem Usage in KB
50,331,648 KB	Gest Mem Provisioned in KB
812 %	VM Mem Compressed %
0 %	VM Mem Ballooned %

BU Summary

VMs 171	VMs ON 171
VMs OFF 0	VMs Stressed 13
VMs Oversized 93	VMs Idle 0
vCPUs 1,832	vCPUs Recommended 801
vCPUs Reclaimable 1,101	

BU Performance

Highest VM CPU Contention 11.07 %	Highest VM CPU Ready per Core 4.48 %
Highest VM CPU Stress 142.34 %	Highest VM Disk Latency 81 ms
Higest VM Disk IOPS 5,328.13	Highest VM Net Pkts Dropped 7.36

Selected VM Recommendations

Number of CPUs 12 Cores	Reclaimable Capacity 75 %
Recommended Size 3 vCPUs	Reclaimable Capacity 9 vCPUs

Automating Business-Oriented Infrastructure Reporting

What if the organization consists of hundreds or thousands of Business Units and needs to provide secure access only to specific workload owners? Creating all of these corresponding Custom Groups, importing AD Groups, and assigning permissions can be a time consuming process. To address this issue, various automation options exist, from simple one-off PowerShell scripts to fully automated vRealize Orchestrator (vRO) workflows. Both options rely on vROps Rest API calls with various degrees of integration and complexity, depending on a specific environment and needs.

Here are the high-level steps need to automate this process:

1. Automation solution prerequisites that can be completed manually as they only require initial setup (one-off scripts can also be used to automate some parts of this):
   a. Define Group Types
   b. Create or Import Super Metrics
   c. Bind Super Metrics to Group Types
   d. Enable Super Metrics
   e. Build the Custom Dashboards
2. The following steps should be fully automated with vRO workflows so that various systems can be kept in sync:
   a. Tag VMs in vCenter (optional)
   b. Create Custom Groups
   c. Create User Groups
   d. Import Users from AD Security Groups
   e. Assign User Roles to Custom Groups
   f. Assign Object Permissions to Custom Groups
   g. Share Custom Dashboards with User Groups

Automation is a very powerful concept in itself, but it is beyond the scope of this white paper.

NOTE: VMware offers Professional Services that can help organizations design, implement, and automate solutions described in this White Paper.

Alerts and Notifications

Any monitoring and reporting solution would be incomplete without a robust Alert and Notification subsystem. vROps provides a lot of flexibility and granularity in this area. Alert Definitions can be created using Symptoms based on collected metrics, properties, and events. Symptoms can use Dynamic or Hard Thresholds that can be preset in definitions and overwritten in Policies for specific workloads. Wait and Cancel Cycles can be used to smooth over peaks, reducing false positives. Logical conditions can be used to combine multiple Symptoms to create smart Alert Definitions that can consider a complex web of issues before triggering an alert. Alert Definitions can be categorized by their severity to all allow prioritization of issues and problems.

Recommendations can be customized to provide problem resolution guidance for the operations staff. Hyperlinks can be embedded in the Recommendations linking to Knowledge Base articles or internal company wikis.

Finally, Notifications can push the triggered Alerts to various external systems and notify the users when problems occur. vROps Notifications can be pushed via SNMP traps to an event monitor, SMTP to a team distribution email box, or Rest API to an IT Service Management (ITSM) system. The integration with ITSM via the vROps REST API opens a door for bidirectional communication between the two systems, allowing Alerts to become Tickets. Tickets can be tracked and when resolved, a call back to vROps can be made to cancel the Alert, completing the problem management lifecycle.

Figure 40: vROps Alert Definition and Notification Construct

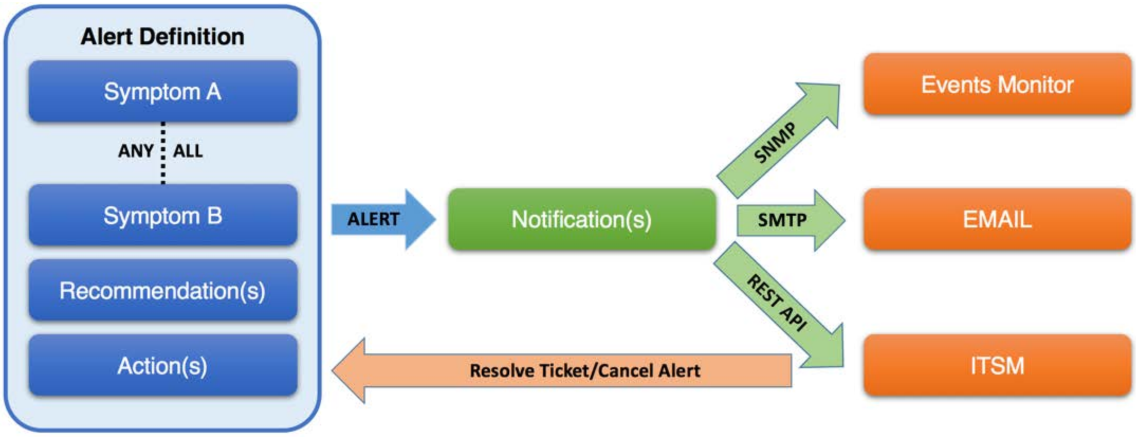

Policies

Policies provide much more than just the ability to enable select metrics as described in the Enable Super Metrics subsection. Different Service-Level Agreements (SLAs) can be configured to provide more targeted monitoring for various types of workloads such as production, test, and development. This enables an organization to have distinctive sets of Alerts and Thresholds, as well as use different Capacity Management models depending on the Workload tracked. In essence, policies provide numerous ways of fine-tuning almost all aspects of vROps, making it a very flexible monitoring and reporting solution for any enterprise.

Figure 41: vROps Custom Groups and Policies Diagram

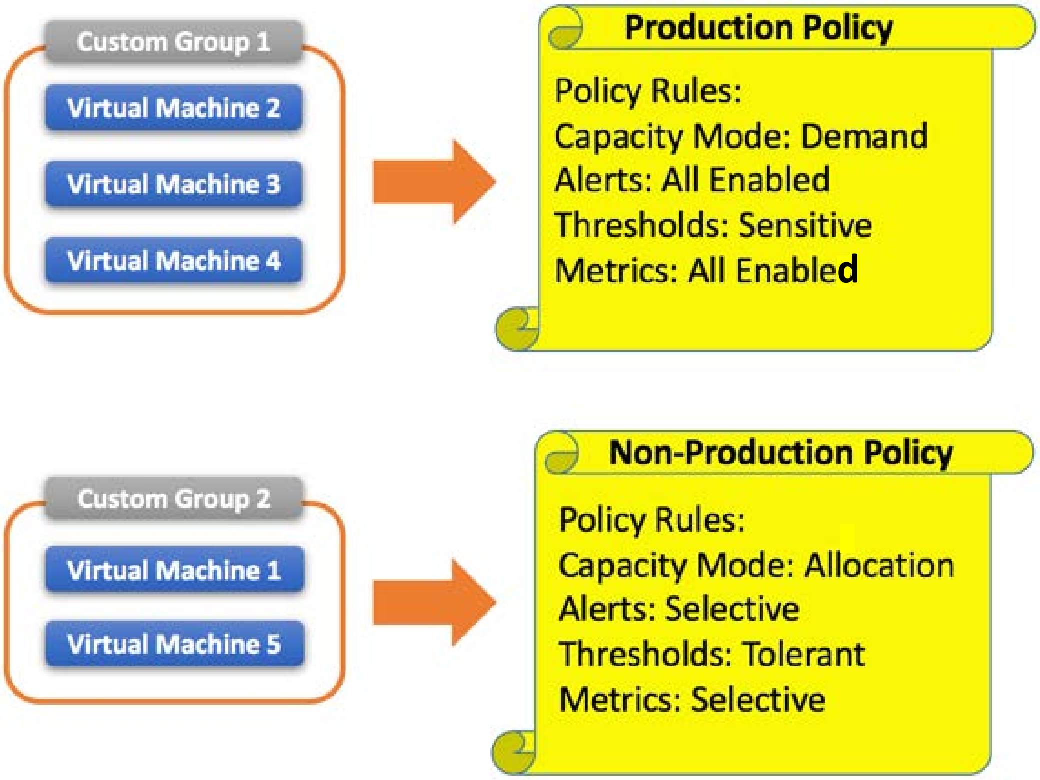

Conclusion

vRealize Operations Manager, straight out of the box, does a great job at providing infrastructure analytics; however, it does not automatically provide performance, utilization, and consumption around business units, applications, and services. IT can use vROps Custom Groups and Super Metrics to solve this problem. With these tools, IT can build a logical representation of a business. IT can use the same set of analytics they are already familiar with to answer real-world business questions and better align IT with business-oriented outcomes.

References

[1] VMware vRealize Operations Manager Documentation
https://www.vmware.com/support/pubs/vrealize-operations-manager-pubs.html

[2] Creating and Applying Super Metrics in a vSphere Environment – White Paper
https://vmwmanagement.files.wordpress.com/2012/01/vcops_super_metrics_in_vsphere_environment_v1_0.pdf

[3] How to Troubleshoot vSphere Performance Issues – White Paper
http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vsphere5x-perfts-vcops.pdf

[4] VMware vRealize Operations Manager – Blog
https://blogs.vmware.com/management/vrealize-operations

[5] VMware Management Pack for vRealize Automation
https://solutionexchange.vmware.com/store/products/management-pack-for-vrealize-automation

[6] VMware vRealize Infrastructure Navigator Documentation
https://www.vmware.com/support/pubs/vcenter-infrastructure-navigator-pubs.html

[7] Leveraging VMware vRealize Infrastructure Navigator – Blog
https://blogs.vmware.com/tam/2016/03/vmware-infrastructure-navigator.html

[8] VMware Professional Services Organization (PSO)
http://www.vmware.com/professional-services.html

About the Author

Peter Tymbel is a Senior Consultant with VMware Software Defined Datacenter (SDDC) Professional Services practice. He holds a degree in Computer Science from Webster University in St. Louis and various industry certifications. His work focuses on helping VMware customers operationalize their SDDC investment. He publishes a blog on various vROps related topics at vbulosity.com.

Acknowledgments

I would like to sincerely thank Mark Gandler for providing the customer use case and sparking the inception for many ideas that ultimately lead to the solution and this white paper. Special thanks to Julie Brodeur, Max Drury, Jeremiah Fowler, Gary Boy, Jason Gaudreau, and Neeraj Gandhi for reviews and contributions to the paper. I also would like extend thanks to my manager Christopher Dorencz for support in my writing endeavors.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com