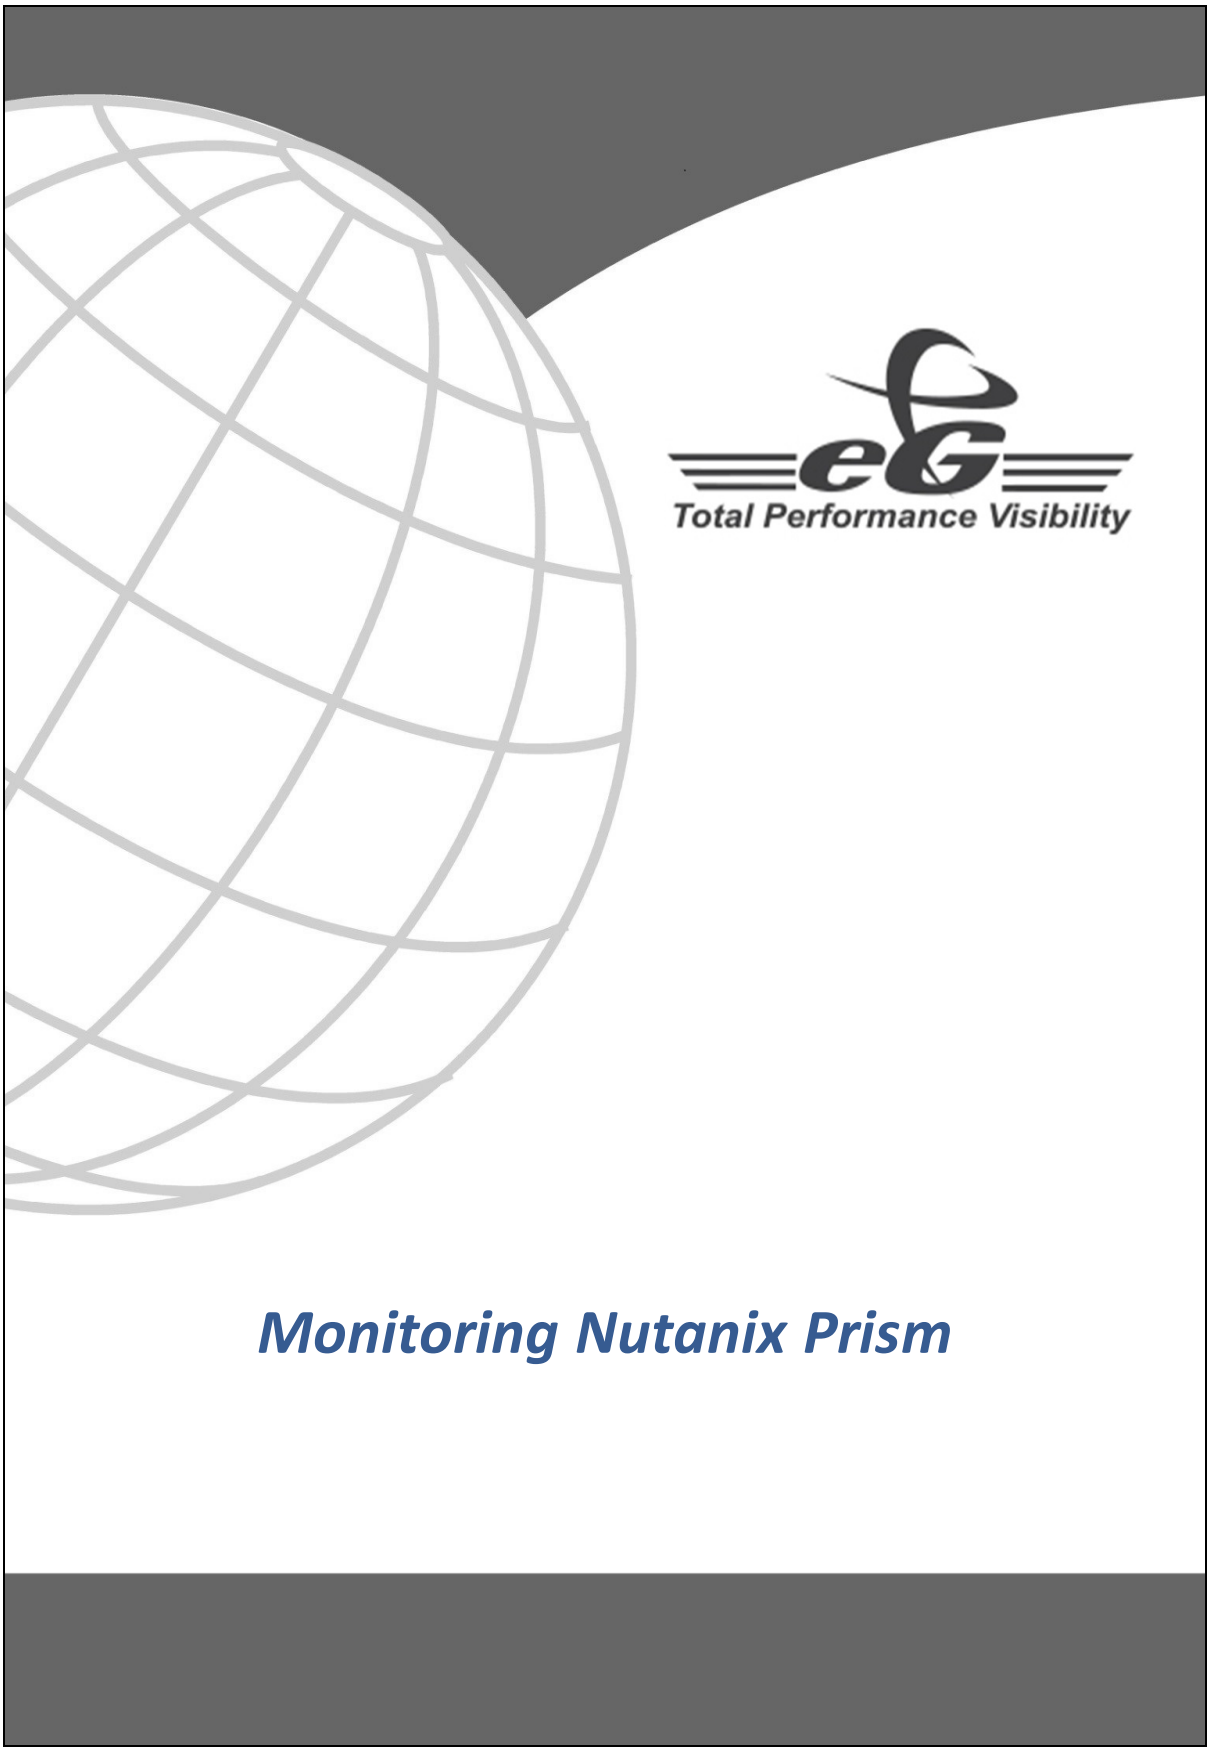

Monitoring Nutanix Prism

Table of contents

Table of Figures

1

Introduction

Prism is a distributed resource management platform which allows users to manage and monitor objects and services across their Nutanix environment.

These capabilities are broken down into two key categories:

- Interfaces

  HTML5 UI, REST API, CLI, PowerShell CMDlets, etc.

- Management

  Policy definition and compliance, service design and status, analytics and monitoring

Figure 1.1 highlights an image illustrating the conceptual nature of Prism as part of the Nutanix platform:

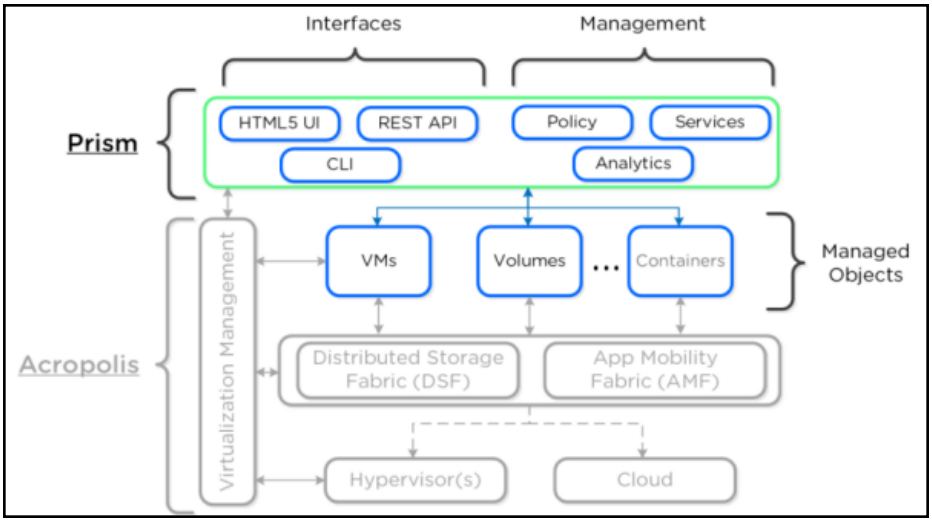

Figure 1.1: High-level Prism Architecture

Prism is broken down into two main components:

- **Prism Central (PC)**
  - Multi-cluster manager responsible for managing multiple Acropolis Clusters to provide a single, centralized management interface. Prism Central is an optional software appliance (VM) which can be deployed in addition to the Acropolis Cluster (can run on it).
  - 1-to-many cluster manager
- **Prism Element (PE)**
  - Localized cluster manager responsible for local cluster management and operations. Every Acropolis Cluster has Prism Element built-in.
    - 1-to-1 cluster manager

Figure 1.2 shows an image illustrating the conceptual relationship between Prism Central and Prism Element:

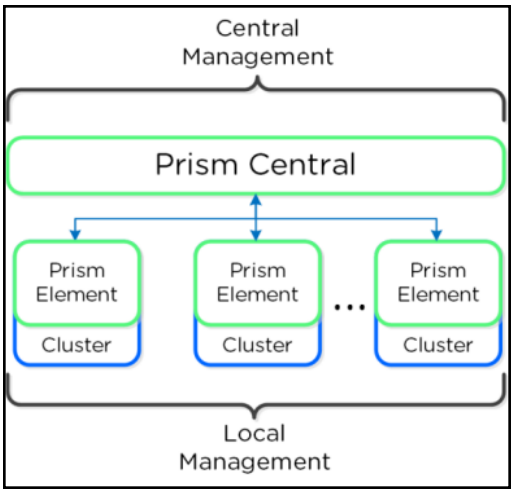

Figure 1.2: The Prism Architecture

A Prism service runs on every CVM with an elected Prism Leader which is responsible for handling HTTP requests. Similar to other components which have a Master, if the Prism Leader fails, a new one will be elected. When a CVM which is not the Prism Leader gets a HTTP request it will permanently redirect the request to the current Prism Leader using HTTP response status code 301.

Figure 1.3 shows a conceptual view of the Prism services and how HTTP request(s) are handled:

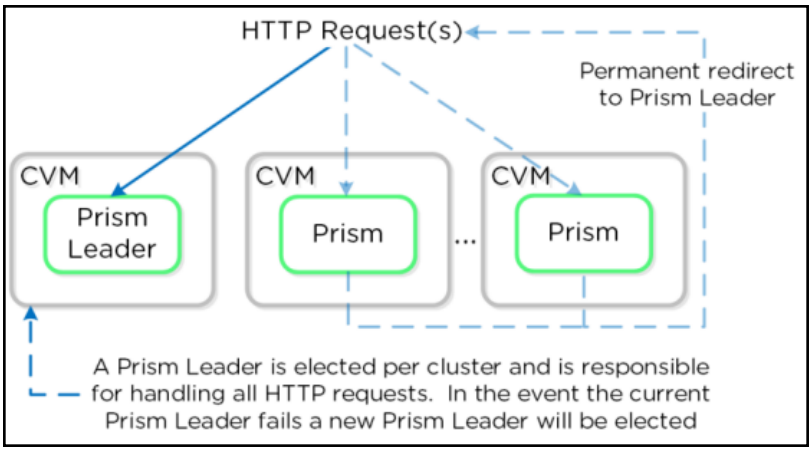

Figure 1.3: Prism Services - Request Handling

If the Prism is suddenly rendered inaccessible, then dependent applications will be denied access to the clustered resources, which in turn will result in serious degradations in application performance. Likewise, if the storage pools and containers managed by the Prism are not sized right, then VMs mapped to these storage resources will experience severe performance deficiencies. To avoid this, the availability and responsiveness of the Prism and the resource usage of the cluster(s) managed by Prism should be continuously tracked and administrators proactively alerted to the abnormalities. For this purpose, eG Enterprise provides a specialized monitoring model for the Nutanix Prism.

This document discusses how to monitor Nutanix Prism using eG Enterprise.

2

Configuring eG Enterprise to Monitor a Nutanix Prism

To monitor a Nutanix Prism using eG Enterprise, you need to manage the Prism using the eG admin interface and configure its tests. However, prior to that, you need to make sure that the pre-requisites for Prism monitoring are fulfilled. This topic elaborates on the pre-requisites and the procedures involved in configuring eG Enterprise to monitor a Nutanix Prism.

2.1 Pre-requisites for Monitoring Nutanix Prism

- The Nutanix Prism should be monitored in an agentless manner.
- The remote agent monitoring the Prism can be deployed on a Windows or a Linux host. If installed on a Windows host, the remote agent should run using *domain administrator* privileges.
- Enable the remote agent to communicate with the eG manager port (default: 7077).
- If a Prism Central is to be monitored, then make sure that there is IP connectivity between the remote agent and the IP address of Prism Central configured at the time of installing it.
- If a Prism Element is to be monitored, then make sure that there is IP connectivity between the remote agent and the external cluster IP address configured for the cluster that is managed by the Prism Element.
- Ensure that the remote agent has web access to the **WEBPORT** (port 9440, by default) configured for the Nutanix Prism.
- To connect to the Nutanix Prism and collect metrics from it, the eG agent should be configured with the credentials of a Prism user with the *Viewer* role. The steps for creating such a user are as follows:
    - Open a browser and connect to the Prism console using the URL: *http://<Prism_console_IP>:<Prism_Port>/console*. If the Prism is SSL-enabled, then the URL will be: *https://<Prism_console_IP>:<Prism_Port>/console*
    - Login to the console as a Cluster administrator or User administrator.
    - Figure 2.1 will then appear.

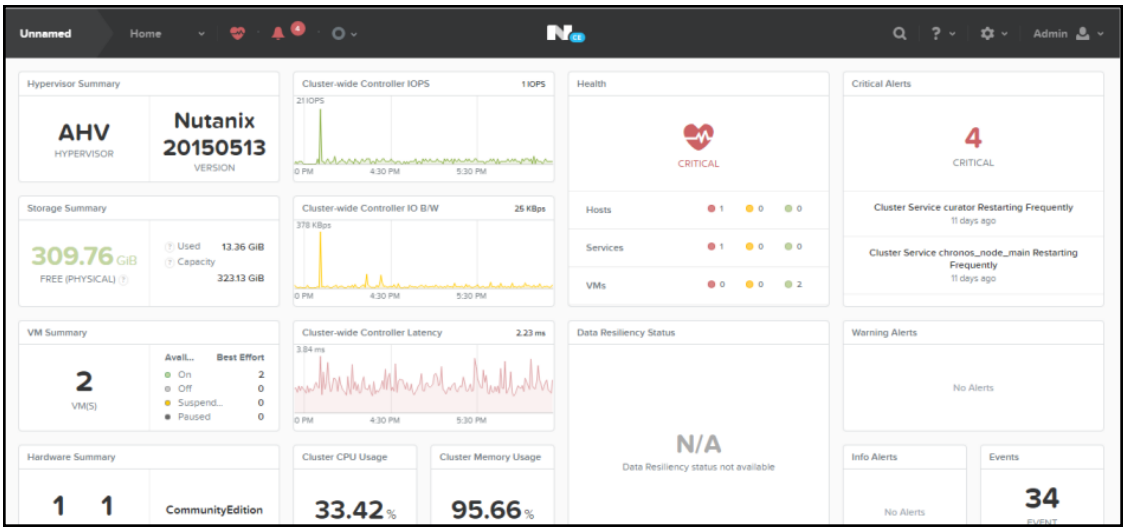

Figure 2.1: The Prism console

- In the console tool bar, click on the down-arrow next to the ⚙ tool to view the menu depicted by Figure 2.2. From the menu, choose the **User Management** option.

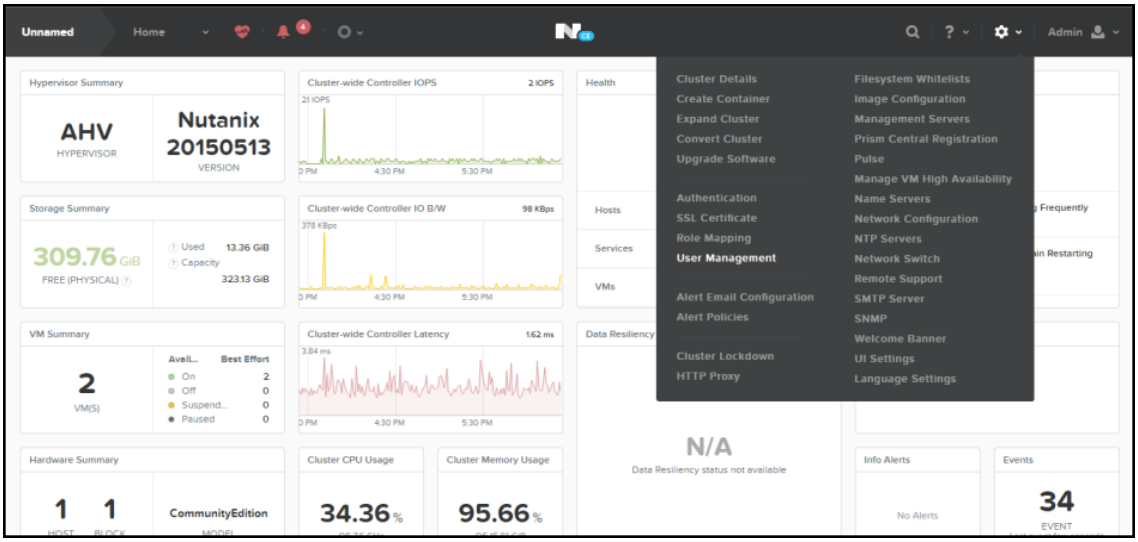

Figure 2.2: Choosing the User Management option

- Figure 2.3 will then appear. Click the **New User** button in Figure 2.3 to create a new user.

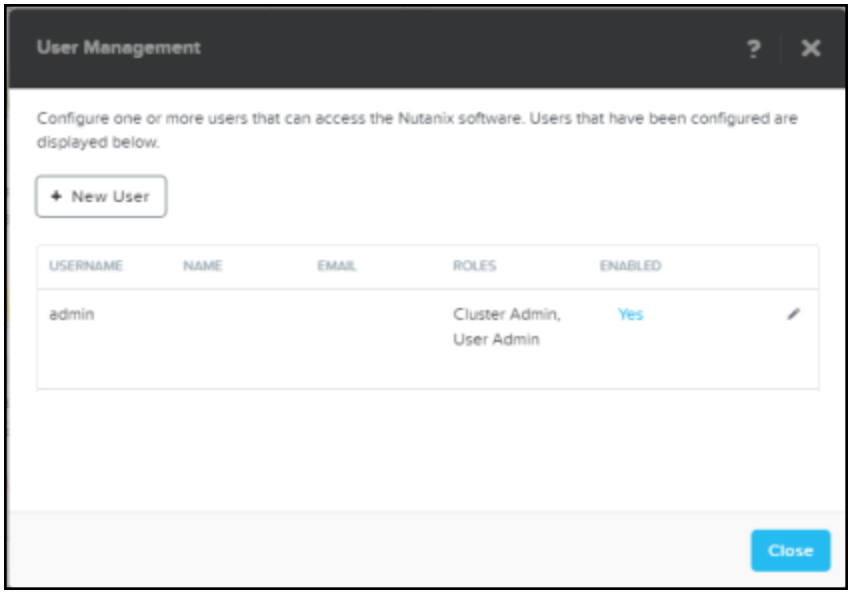

Figure 2.3: List of users who are already registered with Prism

- When Figure 2.4 appears, provide the details of the new user.

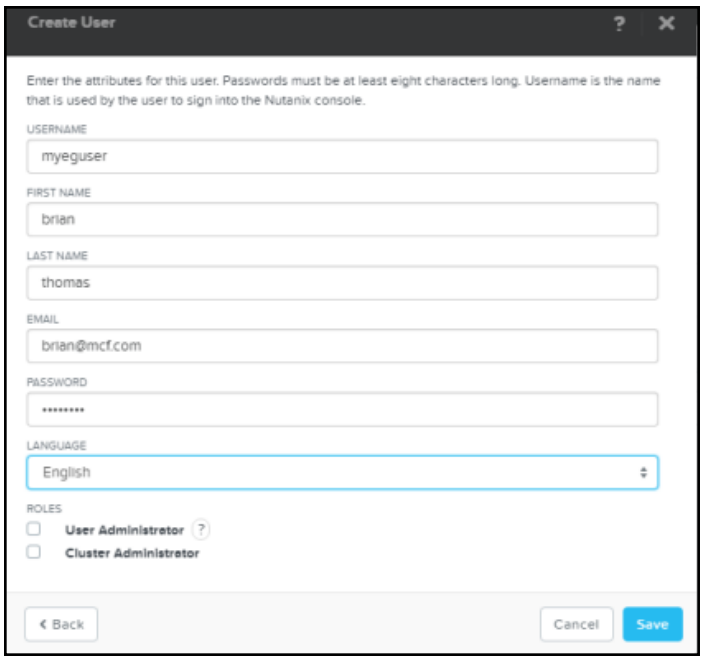

Figure 2.4: Providing the details of the new user

- When configuring the new user, make sure that the **User Administrator** and **Cluster Administrator** check boxes in Figure 2.4 are deselected. This is because, any user who is not assigned either of these roles will automatically assume the *Viewer* role.
- Finally, click the **Save** button in Figure 2.4 to save the changes.

- Figure 2.5 will then appear, wherein you can see that the new user has been successfully added, and that the Viewer role has been automatically assigned to this user.

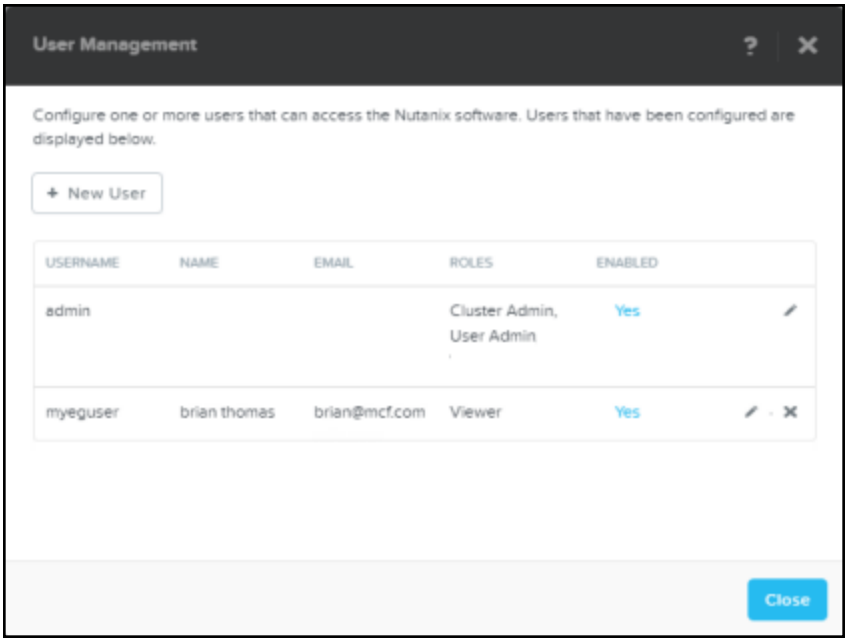

Figure 2.5: The new user's listing

2.2 Managing a Nutanix Prism

eG Enterprise cannot automatically discover a Nutanix Prism. This means that you have to manually add a Prism using the eG admin interface. The steps to achieve this are as follows:

1. Login to the eG admin interface.
2. From the Infrastructure tile of the Admin tile menu, select the Components -> Add/Modify option.

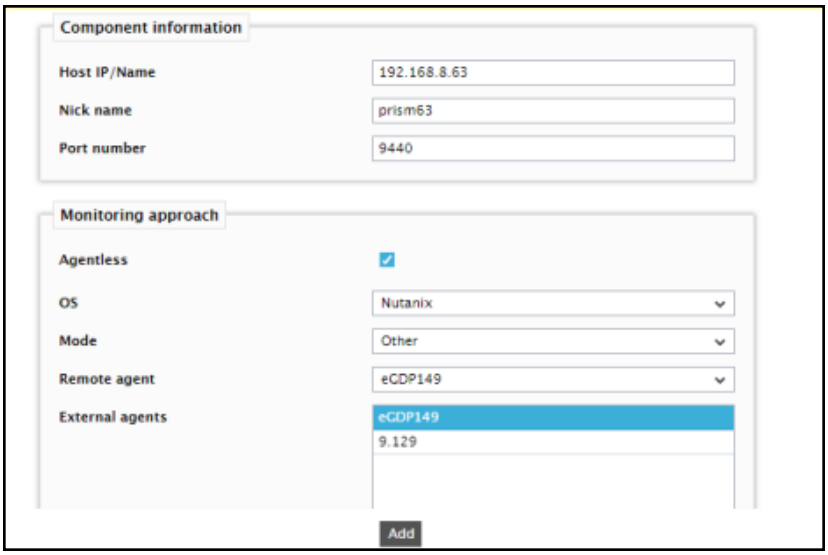

Figure 2.6: Adding a Nutanix Prism for monitoring

3. If the Prism that you want to monitor is a Prism Central, then, against **Host IP/Name**, specify the IP address using which the Prism was installed and configured.

   On the other hand, if the Prism that you want to monitor is a Prism Element, then, **against Host IP/Name**, specify the cluster IP address that is configured for communicating with the cluster that is managed by that Prism Element.

4. Provide a **Nick name** for the Prism and enter the **Port number** at which the Prism listens. By default, the port will be 9440.
5. By default, eG Enterprise monitors the Prism in an **Agentless** manner only.
6. Set the **OS** as *Nutanix* and the **Mode** as *Other*.
7. Pick a **Remote Agent** and an **External Agent** for the Nutanix Prism, and click on **Add** button to add the Prism for monitoring.

2.3 Configuring Tests for Nutanix Prism

After adding a Nutanix Prism for monitoring, try to sign out of the eG administrative interface. You will now be prompted to configure the tests depicted by Figure 2.7.

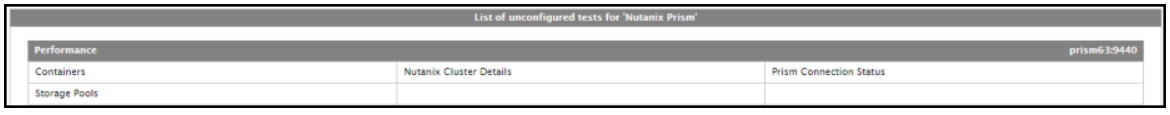

Figure 2.7: List of unconfigured tests for a Nutanix Prism

Click on any test in Figure 2.7 to configure it. For the purpose of this example, let us choose the **Nutanix Cluster Details** test for configuration. Figure 2.8 will then appear.

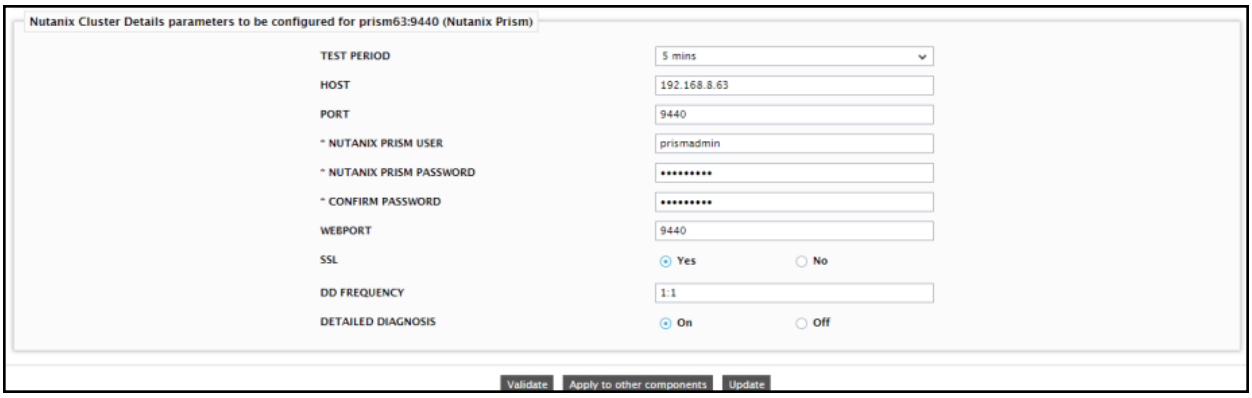

Figure 2.8: Configuring the Nutanix Cluster Details Test

To know how to configure this test, refer to the **Monitoring Nutanix Prism** topic.

Once the test is configured, click the **Update** button to save the changes. Configuring a single test automatically configures all other tests mapped to the managed Nutanix Prism. Therefore, proceed to sign out of the eG admin interface.

Once the configured tests run, they pull a variety of metrics from the Nutanix Prism and report the same to the eG manager. The eG manager then presents the metrics received in real-time in the eG monitoring console using a specialized monitoring model for the Nutanix Prism. Refer to the **Monitoring Nutanix Prism** topic for more details about this model.

3

Monitoring Nutanix Prism

Figure 3.1 depicts the specialized web-based monitoring model that eG Enterprise provides for the Nutanix Prism.

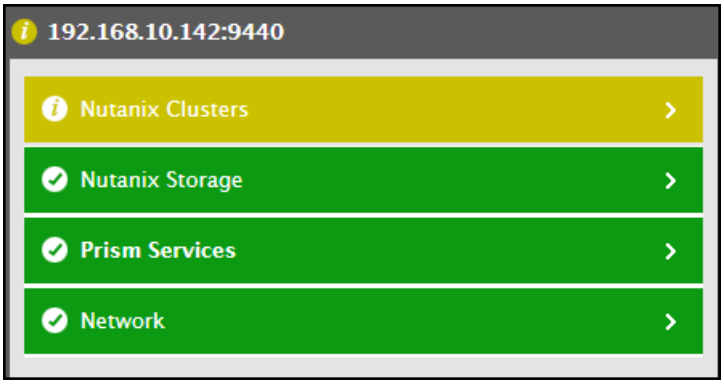

Figure 3.1: Layer model of the Nutanix Prism

Each layer of Figure 1 is mapped to tests that invoke REST API commands to report a wide variety of metrics related to Prism health and the health of the cluster(s) managed by the Prism. Using these metrics, administrators can find quick and accurate answers to the following performance queries:

- Is the Prism available? If so, how quickly is it responding to HTTP requests?
- Which clusters are being managed by the Prism?
- Is any cluster over-utilizing its compute and storage resources?
- How is the I/O load on containers and/or storage pools? Is any container/storage pool overloaded with I/O requests?
- Is any container/storage pool utilizing bandwidth excessively when serving I/O requests?
- Is any container/storage pool latent?
- Is any container/storage pool running out of storage space?

This topic elaborates on each of the top 3 layers of Figure 3.1.

3.1 The Prism Services Layer

The tests mapped to this layer check the availability of the Prism and its responsiveness.

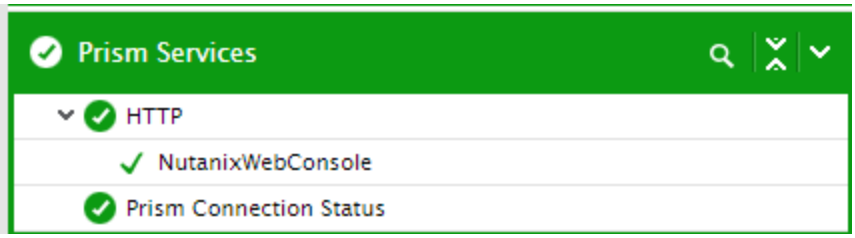

Figure 3.2: The tests mapped to the Prism Services layer

Since the HTTP test is already discussed in the **Monitoring Web Servers** document, the Prism Connection Status test alone will be discussed in the Section **3.1.1** topic.

3.1.1 Prism Connection Status Test

The Nutanix Prism typically manages one/more Nutanix Acropolis clusters. Each node in a cluster runs a Nutanix Controller VM (CVM), which serves all of the I/O operations for the hypervisor and all VMs running on that node. A Prism service runs on every CVM with an elected Prism Leader which is responsible for handling HTTP requests. The cluster external IP will always be hosted by the Prism Leader. If the Prism Leader fails, then a new leader will be elected, which will assume the cluster external IP. This implies that the Prism service will be rendered unavailable, only if all the CVMs in a cluster fail. Such a failure can cause the dependent applications to suffer prolonged outages. Likewise, if the Prism leader takes too long to respond to HTTP requests, it will once again degrade application performance.

If this is to be avoided, then the availability and responsiveness of the Prism service should be periodically checked and abnormalities should be promptly reported. This can be achieved using the **Prism Connection Status** test.

This test emulates an HTTP request to a Nutanix Prism, reports whether the Prism service is available or not, and if available, also reports how quickly it responds to the request. Sudden breaks in Prism availability and poor responsiveness of the Prism service can be promptly detected in the process. The test also reports the response code returned by the Prism service, so that the nature of the response - whether it is an error response or not - can be determined.

Target of the test : A Nutanix Acropolis Prism

Agent deploying the test : A remote agent

Outputs of the test : One set of results for the Nutanix Acropolis Prism

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **NUTANIX PRISM USER** and **NUTANIX PRISM PASSWORD** - To connect to the Nutanix Prism and collect metrics from it, the eG agent should be configured with the credentials of a Prism user with the

Viewer role. The steps for creating such a user are detailed in the Pre-requisites for Monitoring Nutanix Prism topic.

5. **CONFIRM PASSWORD** - Confirm the **NUTANIX PRISM PASSWORD** by retyping it here.
6. **SSL** - By default, the Nutanix Prism server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the Prism server via HTTPS by default.
7. **WEBPORT** - By default, the Nutanix Prism server listens on port 9440. This implies that to monitor a Nutanix Prism server, the eG agent connects to the server via port 9440.
8. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For a *Nutanix Acropolis Prism* server, this is set to *1:1* by default. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Prism availability	Indicates whether or not the Prism service is available.	Percent	If the value of this measure is 100, it indicates that the Prism service is available. The value 0 on the other hand denotes that the Prism service is unavailable.
Response time	Indicate the time taken by the Prism service to respond to HTTP requests.	Secs	If the value of this measure consistently increases, it indicates that the performance of the Nutanix prism is gradually deteriorating. Ideally, the value of this measure should be low.
Response code	Indicates the response code returned by the emulated HTTP request.	Number	The possible values for this measure and what they represent are detailed in the table below: **Measure Value** / **Description** 200 / The API request was successful and received a response. 201 / The API request was successful and created an object.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
<table>
<tr><th>Measure Value</th><th>Description</th></tr>
<tr><td>400</td><td>The API request was malformed and could not be processed.</td></tr>
<tr><td>401</td><td>You have no access and/or are not authorized.</td></tr>
<tr><td>403</td><td>You are authorized but do not have the privileges for this API.</td></tr>
<tr><td>404</td><td>The URL was not found</td></tr>
<tr><td>405</td><td>The called method is not allowed or is not supported</td></tr>
<tr><td>408</td><td>The request timed out (20 seconds maximum).</td></tr>
<tr><td>500</td><td>The API request was received but there was a server error.</td></tr>
<tr><td>503</td><td>Service unavailable at this time or too early to process.</td></tr>
<tr><td>508</td><td>HTTP other than 1.1 not supported.</td></tr>
</table>
</td></tr>
</table>

3.2 The Nutanix Storage Layer

The tests mapped to this layer monitor the storage pools and containers managed by the Prism, track the IOPS performed on them, measure their space usage, and report abnormal usage patterns and I/O performance.

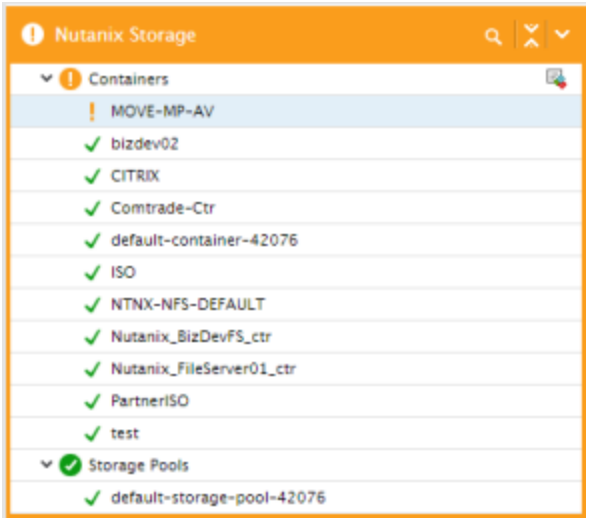

Figure 3.3: The tests mapped to the Nutanix Storage layer

3.2.1 Storage Pools Test

A storage pool is a group of physical storage devices including PCIe SSD, SSD, and HDD devices for the cluster. The storage pool can span multiple Nutanix nodes and is expanded as the cluster scales. In most configurations, only a single storage pool is leveraged.

Since the VMs and nodes in a cluster rely heavily on the storage pools for their availability and overall performance, it is imperative that the storage pools be sized and tuned right. If not, the dependent VMs and nodes will experience serious performance setbacks ranging from a slowness to a standstill!

To determine whether/not a storage pool needs to be resized, an administrator must first know how much storage space is available to that pool, how this space has been utilized, what is the typical I/O load on the pool, and how well it processes this load. The **Storage Pools** test reports these statistics for each storage pool that is managed by the Nutanix Prism. With the help of this information, administrators can proactively detect a potential space contention, an I/O overload, and even processing latencies that may impact storage performance, and can initiate measures to avert them. Additionally, the test also measures and reports the effectiveness of the storage optimization methodologies that are applied currently, t

Target of the test : A Nutanix Acropolis Prism

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every storage pool managed by the Nutanix Acropolis Prism

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.

4. **NUTANIX PRISM USER** and **NUTANIX PRISM PASSWORD** - To connect to the Nutanix Prism and collect metrics from it, the eG agent should be configured with the credentials of a Prism user with the *Viewer* role. The steps for creating such a user are detailed in the Pre-requisites for Monitoring Nutanix Prism topic.
5. **CONFIRM PASSWORD** - Confirm the **NUTANIX PRISM PASSWORD** by retyping it here.
6. **SSL** - By default, the Nutanix Prism server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the Prism server via HTTPS by default.
7. **WEBPORT** - By default, the Nutanix Prism server listens on port 9440. This implies that to monitor a Nutanix Prism server, the eG agent connects to the server via port 9440.
8. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For a Nutanix Acropolis Prism server, this is set to *1:1* by default. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. It is recommended that you do not change the default setting of this parameter.
9. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

   The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Physical disks	Indicates the number of disks pooled in this storage pool.	Number	The detailed diagnosis of this measure lists the ID of the disk, status of the disk and host name.
Disk physical usage	Indicates the total amount of physical storage space used by the disks in this storage pool.	GB	
Storage capacity	Indicates the amount of space in the cluster that is available to this storage	GB	Where there are multiple storage pools, you can compare the value of this

Measurement	Description	Measurement Unit	Interpretation
	pool.		measure across the pools to know which pool has been sized with the maximum storage space.
Storage used space	Indicates the total amount of physical storage space used in this storage pool.	GB	A consistent increase in the value of this measure is indicative of rapid usage of space in the pool, which could lead to a storage space contention.
Storage free space	Indicates the total amount of physical storage space that is unused in this pool.	GB	Ideally, the value of this measure should be high. A very low value for this measure could indicate that the pool is running short of storage resources and may require expansion.
Storage space usage	Indicates the percentage of physical storage space used in this storage pool.	Percent	A value close to 100% is a cause for concern as it indicates a probable contention for storage space on the pool. You may want to consider resizing the pool to make sure that VM operations continue uninterrupted.
Storage free space	Indicates the percentage of physical storage space that is unused in this storage pool.	Percent	A value less than 50% is a cause for concern as it indicates a probable contention for storage space on the pool. You may want to consider resizing the pool to make sure that VM operations continue uninterrupted.
Storage logical usage	Indicates the total amount of logical storage space used in this storage pool.	GB	
Total I/O latency	Indicates the average I/O latency for physical disk requests in this storage pool.	Secs	Ideally, the value of this measure should be very low. A high value or a steady increase in this value could indicate an I/O processing bottleneck on the pool. In such a case, compare the value of the Read IO latency and

Measurement	Description	Measurement Unit	Interpretation
			Write IO latency measures to figure out where the slowness is worst - when processing read requests? or write requests?
Read IO latency	Indicates the average time taken by this storage pool to process read I/O requests.	Secs	If the Total I/O latency measure reports an abnormally high value, then compare the value of these measures to figure out where the slowness is maximum - when processing read requests? or write requests?
Write IO latency	Indicates the average time taken by this storage pool to process write I/O requests.	Secs	
Total IO bandwidth	Indicates the bandwidth per second used by this storage pool when processing I/O requests.	KB/Sec	A high value for this measure denotes that the storage pool is processing bandwidth- intensive I/O. In such situations, you may want to compare the value of the Read IO bandwidth and Write IO bandwidth measures to know what type of I/O requests are truly contributing to the excessive bandwidth consumptions - read requests? or write requests?
Read IO bandwidth	Indicates the bandwidth per second used by this storage pool when processing read I/O requests.	KB/Sec	If the value of the Total IO bandwidth measure is high, then you may want to compare the value of the Read IO bandwidth and Write IO bandwidth measures to know what type of I/O requests are truly contributing to the excessive bandwidth consumption - read requests? or write requests?
Write IO bandwidth	Indicates the bandwidth per second used by this storage pool when processing write I/O requests.	GB	
Total IOPS	Indicates the number of I/O operations performed currently on this storage pool.	Number	This measure is a good indicator of the level of I/O activity on the storage pool. A steady and significant increase in the value of this measure could indicate a potential I/O overload. In such

Measurement	Description	Measurement Unit	Interpretation
			situations, you may want to compare the value of the Read IOPS and Write IOPS measures of the storage pool to know what type of IO operations are contributing to the overload.
Read IOPS	Indicates the number of read I/O operations performed currently on this storage pool.	Number	If the value of the Total IOPS measure is unusually high, then compare the value of these measures for that storage pool to know what is contributing to the unusual I/O activity levels - read requests? or write requests?
Write IOPS	Indicates the number of write I/O operations performed currently on this storage pool.	Number	
Total transformed usage	Indicates the amount of actual usage of storage (i.e., usage after compression and deduplication) in the storage pool.	GB	The Nutanix platform incorporates a wide range of storage optimization technologies that work in concert to make efficient use of available capacity for any workload. Compression and Deduplication are two such technologies. Compression can be inline or offline. Inline compression will compress sequential streams of data or large I/O sizes (>64K) in memory before it is written to the Extent Store. Offline compression will initially write the data as normal (in an un-compressed state) and then leverage the Curator framework to compress the data cluster wide. The Elastic Dedupe Engine in Nutanix allows for data deduplication in the capacity (Extent Store) and performance (Unified Cache) tiers. Streams of data are fingerprinted during ingest using a SHA-1 hash at a 16K

Measurement	Description	Measurement Unit	Interpretation
			granularity. This fingerprint is only done on data ingest and is then stored persistently as part of the written block's metadata. For duplicate data that can be deduplicated in the capacity tier, the data does not need to be scanned or re- read, essentially duplicate copies can be removed. The true effectiveness of these optimization methodologies can be measured by determining how much storage space in the pool these technologies helped save. By comparing the value of this measure with the value of the Storage usage measure of the pool, you should be able to make an accurate assessment of the effectiveness of these methodologies.

3.2.2 Containers Test

A container is a logical segmentation of the Storage Pool and contains a group of VMs or files (vDisks). Containers typically have a 1 to 1 mapping with a datastore (in the case of NFS/SMB).

The containers need to be adequately sized to handle the load imposed by the VMs they are attached to. Latent containers and the ones with insufficient storage space will adversely impact the performance of the dependent VMs and also affect the overall storage performance of the Nutanix environment. This is why, it is important that administrators know which VMs are mapped to which containers, determine how much space each container is configured with, and also keep track of the current demand for space and processing power on every container, so that containers that are not configured to meet this demand can be proactively detected and resized (if required). This is where the **Containers** test helps!

This test auto-discovers the containers managed by the Nutanix Prism, monitors the I/O load on and usage of each container , and precisely pinpoints those containers that are overloaded and under-sized. Additionally, the test also lists the VMs that are mapped to each container, so that you can quickly identify those VMs, the performance of which will be significantly affected owing to the problematic containers. Furthermore, the test also enables you to quickly review the capabilities that are turned on/off at the container-level, so that you can go back and make changes to the overall configuration of the container, if required.

Target of the test : A Nutanix Acropolis Prism

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every container managed by the Nutanix Acropolis Prism

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **NUTANIX PRISM USER** and **NUTANIX PRISM PASSWORD** - To connect to the Nutanix Prism and collect metrics from it, the eG agent should be configured with the credentials of a Prism user with the *Viewer* role. The steps for creating such a user are detailed in the Pre-requisites for Monitoring Nutanix Prism topic.
5. **CONFIRM PASSWORD** - Confirm the **NUTANIX PRISM PASSWORD** by retyping it here.
6. **SSL** - By default, the Nutanix Prism server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the Prism server via HTTPS by default.
7. **WEBPORT** - By default, the Nutanix Prism server listens on port 9440. This implies that to monitor a Nutanix Prism server, the eG agent connects to the server via port 9440.
8. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For a Nutanix Acropolis Prism server, this is set to *1:1* by default. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. It is recommended that you do not change the default setting of this parameter.
9. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

   The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Replication factor	Indicates the replication factor setting of this	Number	Each Nutanix container can be configured with a replication factor (RF)

Measurement	Description	Measurement Unit	Interpretation
	container.		of two or three. RF=2 ensures that two copies of data are maintained at all times, allowing the cluster to survive the failure of one node or disk. Similarly, when RF is set to 3 (RF=3), three copies of the data are maintained in the cluster, providing resilience from two simultaneous failures. This level of flexibility allows administrators to dynamically configure data redundancy based on application SLAs and the criticality of the data set. The value of this measure can be 2 or 3 depending upon the RF setting of the container.
Oplog replication factor	Indicates the replication factor setting for the Oplog of this container.	Number	The OpLog acts as a staging area to absorb incoming writes onto a low-latency SSD tier. Upon being written to the local OpLog, the data is synchronously replicated to another one or two Nutanix CVM's OpLog, depending upon the replication factor (RF) setting of Oplog, before being acknowledged (Ack) as a successful write to the host. For instance, if the RF setting of Oplog is 2, then data will be synchronously replicated to one other Oplog. If RF is 2, then data will be replicated to two other Oplogs. This ensures that the data exists in at least two or three independent locations and is fault tolerant. NOTE: For RF=3, a minimum of 5 nodes is required since metadata will be RF5. The value of this measure can be 2 or 3, depending upon the RF setting for Oplog.

Measurement	Description	Measurement Unit	Interpretation
Compression enabled	Indicates whether/not compression is enabled for this container.	Number	The Nutanix Capacity Optimization Engine (COE) is responsible for performing data transformations to increase data efficiency on disk. Currently compression is one of the key features of the COE to perform data optimization. DSF provides both inline and offline flavors of compression to best suit the customer's needs and type of data. Inline compression will compress sequential streams of data or large I/O sizes (>64K) in memory before it is written to the Extent Store (SSD + HDD). This includes data draining from OpLog as well as sequential data skipping it. Offline compression will initially write the data as normal (in an un-compressed state) and then leverage the Curator framework to compress the data cluster wide. When inline compression is enabled but the I/Os are random in nature, the data will be written un-compressed in the OpLog, coalesced, and then compressed in memory before being written to the Extent Store. The Google Snappy compression library is leveraged which provides good compression ratios with minimal computational overhead and extremely fast compression / decompression rates. This measure reports the value *On* if compression is enabled for this container, and the value *Off* if compression is disabled. The numeric values that correspond to these

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>measure values are listed below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>On</td><td>1</td></tr><tr><td>Off</td><td>0</td></tr></table>
Note:
By default, this measure reports the values listed in the Measure Value column to indicate whether/not compression is enabled for a container. In the graph of the measure however, the same will be represented using the numeric equivalents only.</td></tr>
<tr><td>On disk deduplication</td><td>Indicates whether/not on-disk deduplication is enabled for this container.</td><td></td><td>Deduplication is the process of eliminating duplicate data, so as to increase the effective capacity in the disk tier, and the system’s RAM and flash tiers.

Nutanix delivers two types of data deduplication to accelerate application performance and to optimize storage capacity. These are namely, Performance tier deduplication and Capacity tier / MapReduce / On-disk deduplication. On-disk deduplication, if enabled for a container, reduces repetitive data in the capacity tier to increase the effective storage capacity of a cluster. This type of deduplication is global and distributed across all nodes in the cluster, minimizing any performance overhead. MapReduce deduplication is particularly useful for virtual desktops with full clones.Performance-trier deduplication can be used without Capacity- tier deduplication but not the other way around.</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>This measure reports the value On if deduplication is enabled for this container, and the value Off if deduplication is disabled. The numeric values that correspond to these measure values are listed below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>On</td><td>1</td></tr><tr><td>Off</td><td>0</td></tr></table>
Note:
By default, this measure reports the values listed in the Measure Value column to indicate whether/not on-disk deduplication is enabled for a container. In the graph of the measure however, the same will be represented using the numeric equivalents only.</td></tr>
<tr><td>Erasure coding</td><td>Indicates whether/not erasure coding is enabled for this container.</td><td></td><td>The Nutanix platform leverages a replication factor (RF) for data protection and availability. This method provides the highest degree of availability because it does not require reading from more than one storage location or data re- computation on failure. However, this does come at the cost of storage resources as full copies are required.
To provide a balance between availability while reducing the amount of storage required, DSF provides the ability to encode data using erasure codes (EC).
Similar to the concept of RAID (levels 4, 5, 6, etc.) where parity is calculated, EC encodes a strip of data blocks on different nodes and calculates parity. In</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>
the event of a host and/or disk failure, the parity can be leveraged to calculate any missing data blocks (decoding). In the case of DSF, the data block is an extent group and each data block must be on a different node and belong to a different vDisk.

The number of data and parity blocks in a strip is configurable based upon the desired failures to tolerate. The configuration is commonly referred to as the number of <data blocks>/<number of parity blocks>.

For example, “RF2 like” availability (e.g., N+1) could consist of 3 or 4 data blocks and 1 parity block in a strip (e.g., 3/1 or 4/1). “RF3 like” availability (e.g. N+2) could consist of 3 or 4 data blocks and 2 parity blocks in a strip (e.g. 3/2 or 4/2).

This measure reports the value On if erasure coding is enabled for this container, and the value Off if it is disabled. The numeric values that correspond to these measure values are listed below:
<table>
<tr><th>Measure Value</th><th>Numeric Value</th></tr>
<tr><td>On</td><td>1</td></tr>
<tr><td>Off</td><td>0</td></tr>
</table>
Note:

By default, this measure reports the values listed in the Measure Value column to indicate whether/not erasure coding is enabled for a container. In the graph of the measure however, the same will be represented using the numeric equivalents only.
</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Is marked for removal?</td><td>Indicates whether/not the container is marked for removal.</td><td></td><td>The values that this measure can report and their corresponding numeric values are listed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>On</td><td>1</td></tr><tr><td>Off</td><td>0</td></tr></table>
Note:
By default, this measure reports the values listed in the Measure Value column to indicate whether/not the container is marked for removal. In the graph of the measure however, the same will be represented using the numeric equivalents only.</td></tr>
<tr><td>Attached VMs</td><td>Indicates the number of VMs that are attached to this container.</td><td>Number</td><td>Use the detailed diagnosis of this measure to know which VMs are attached to this container.</td></tr>
<tr><td>Total I/O latency</td><td>Indicates the average time taken by this container to process I/O requests.</td><td>Secs</td><td>Ideally, the value of this measure should be very low. A high value or a steady increase in this value could indicate an I/O processing bottleneck on the container. In such a case, compare the value of the Read IO latency and Write IO latency measures to figure out where the slowness is worst - when processing read requests? or write requests?</td></tr>
<tr><td>Read IO latency</td><td>Indicates the average time taken this container to process read I/O requests.</td><td>Secs</td><td rowspan="2">If the Total I/O latency measure reports an abnormally high value, then compare the value of these measures to figure out where the slowness is maximum - when processing read requests? or write requests?</td></tr>
<tr><td>Write IO latency</td><td>Indicates the average time taken by this container to process write I/O requests.</td><td>Secs</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
Total IO bandwidth	Indicates the bandwidth per second used by this container when processing I/O requests.	KB/Sec	A high value for this measure denotes that the container is processing bandwidth- intensive I/O. In such situations, you may want to compare the value of the *Read IO bandwidth* and *Write IO bandwidth m*easures to know what type of I/O requests are truly contributing to the excessive bandwidth consumptions - read requests? or write requests?
Read IO bandwidth	Indicates the bandwidth per second used by this container when processing read I/O requests.	KB/Sec	If the value of the *Total IO bandwidth* measure is high, then you m*ay want to compare the value of the* Read IO bandwidth and *Write IO bandwidth* measures to know what type of I/O requests are truly contributing to the excessive bandwidth consumption - read requests? or write requests?
Write IO bandwidth	Indicates the bandwidth per second used by this container when processing write I/O requests.	GB	
Total IOPS	Indicates the number of I/O operations performed currently on this container.	Number	This measure is a good indicator of the level of I/O activity on the container. A steady and significant increase in the value of this measure could indicate a potential I/O overload. In such situations, you may want to compare the value of the *Read IOPS* and *Write IOPS* measures of the container to know what type of IO operations are contributing to the overload.
Read IOPS	Indicates the number of read I/O operations performed currently on this container.	Number	If the value of the *Total IOPS* measure is unusually high, then compare the value of these measures for that container to know what is contributing to the unusual I/O activity levels - read requests? or write requests?
Write IOPS	Indicates the number of write I/O operations	Number	

Measurement	Description	Measurement Unit	Interpretation
	performed currently on this container.		
Max capacity	Indicates the maximum capacity configured for this container .	GB	The maximum capacity value reflects total available storage regardless of how many containers are defined. Therefore, when you have two containers, it can appear you have twice as much capacity because the field values for both containers show the full amount. This will normally match the storage pool size.
Used space	Indicates the amount of storage space in this container that is being used currently.	GB	Ideally, the value of this measure should be low.
Free space	Indicates the amount of storage space in this container that is still unused.	GB	Ideally, the value of this measure should be high.
Space usage	Indicates the percentage of storage space in this container that is in use.	Percent	A value close to 100% is a cause for concern, as it indicates that the container is rapidly running out of free space. You may want to consider allocating more space to the container to avoid loss of data.
Free space	Indicates the percentage of storage space in this container that is available for use.	Percent	A value less than 50% could be a cause for concern, as it indicates that the storage space in the container is being over-utilized. You may want to consider allocating more space to the container to avoid loss of data.
Reserved usage	Indicates the amount of space reserved for the use of this container.	GB	If storage space is reserved for a container, then that container is

Measurement	Description	Measurement Unit	Interpretation
			guaranteed the availability of the reserved capacity. For instance, if the maximum usable capacity of a container is 40TB, but 10TB is its reserved capacity, then this 10TB of space will be available for the use of the container at any given point in time. Since containers are thin-provisioned and space is consumed on a first come first serve basis, there is no guarantee that storage will be available to this container once the reserved capacity of 10TB is consumed.
Logical usage	Indicates the amount of logical space used by this container.	GB	
Disk physical usage	Indicates the total amount of physical storage space used in the container.	GB	
Unreserved own usage	Indicates the amount of unreserved storage that can be used by the container	GB	The value of this measure is the difference between the value of the *Disk physical usage* and *Reserved usage* measures.

3.3 The Nutanix Clusters Layer

The test mapped to this layer discovers and monitors each cluster managed by the Prism, if the Prism Central is being monitored. In case of a Prism Element, only a single cluster will be monitored by the test.

The resources available to the cluster and how well the cluster uses the resources are tracked and reported, so that administrators can decide whether the cluster's resource capacity is to be expanded.

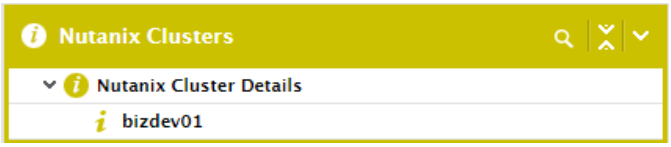

Figure 3.4: The test mapped to the Nutanix Clusters layer

3.3.1 Nutanix Cluster Details Test

Nutanix software runs as a virtual machine (VM) on industry-standard hypervisors, including VMware vSphere and Microsoft Hyper-V, as well as Nutanix's built-in AHV virtualization technology (based on KVM). Servers running Nutanix software are clustered together in one or more clusters. The Nutanix software aggregates the local data storage resources on each node and forms a distributed storage fabric that is presented back to the hypervisor as NFS, SMB3 or iSCSI depending on the hypervisor deployed. A global data system architecture integrates each new node into the cluster, allowing you to scale the solution to meet the needs of your infrastructure.

To ascertain the needs of the infrastructure and to determine whether/not additional nodes need to be added to the cluster, Nutanix administrators should:

- Monitor each cluster and determine its composition and configuration;
- Study how each cluster is utilizing the pooled compute and storage resources;

This visibility is provided by the Nutanix Cluster Details test. This test auto-discovers the clusters managed by the monitored Nutanix Prism, and reports the following for each cluster:

- The number and names of VMs and hypervisors managed by each cluster;
- The total compute and storage capacity of every cluster:
- How every cluster is using the available capacity;
- The I/O load on each cluster;
- How each cluster is handling the I/O load;

In the process, the test pinpoints those clusters that many not have the storage and processing power to handle their current and anticipated load, and are hence candidates for expansion.

Target of the test : A Nutanix Acropolis Prism

Agent deploying the test : A remote agent

Outputs of the test : One set of results for every cluster managed by the Nutanix Acropolis Prism. If the Prism Element is being monitored, then this test will report metrics only for a single cluster. If the Prism Central is being monitored, then this test will report metrics for multiple clusters.

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** - The port at which the specified **HOST** listens. By default, this is NULL.
4. **NUTANIX PRISM USER** and **NUTANIX PRISM PASSWORD** - To connect to the Nutanix Prism and collect metrics from it, the eG agent should be configured with the credentials of a Prism user with the *Viewer* role. The steps for creating such a user are detailed in the Pre-requisites for Monitoring Nutanix Prism topic.

5. **CONFIRM PASSWORD** - Confirm the **NUTANIX PRISM PASSWORD** by retyping it here.
6. **SSL** - By default, the Nutanix Prism server is SSL-enabled. Accordingly, the SSL flag is set to **Yes** by default. This indicates that the eG agent will communicate with the Prism server via HTTPS by default.
7. **WEBPORT** - By default, the Nutanix Prism server listens on port 9440. This implies that to monitor a Nutanix Prism server, the eG agent connects to the server via port 9440.
8. **DD FREQUENCY** - Refers to the frequency with which detailed diagnosis measures are to be generated for this test. For a Nutanix Acropolis Prism server, this is set to *1:1* by default. This indicates that, by default, detailed measures will be generated every time this test runs, and also every time the test detects a problem. It is recommended that you do not change the default setting of this parameter.
9. **DETAILED DIAGNOSIS** - To make diagnosis more efficient and accurate, the eG suite embeds an optional detailed diagnostic capability. With this capability, the eG agents can be configured to run detailed, more elaborate tests as and when specific problems are detected. To enable the detailed diagnosis capability of this test for a particular server, choose the **On** option. To disable the capability, click on the **Off** option.

   The option to selectively enable/disable the detailed diagnosis capability will be available only if the following conditions are fulfilled:

   - The eG manager license should allow the detailed diagnosis capability
   - Both the normal and abnormal frequencies configured for the detailed diagnosis measures should not be 0.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Is it a cloud cluster?	Indicates whether/not this cluster is deployed on the cloud.		The values that this measure can report and their corresponding numeric values are detailed in the table below: **Measure Value** / **Numeric Value** Yes – 1 No – 0 **Note:** By default, this measure reports the **Measure Value**s listed in the table above to indicate whether/not the cluster is on the cloud. In the graph of this measure however, the same is represented using the numeric equivalents only.

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td>Are shadow clones enabled?</td><td>Indicates whether/not shadow clones are enabled for this cluster.</td><td></td><td>Shadow Clones significantly improve performance by caching virtual machine data across a Nutanix cluster. Unique to Nutanix, Shadow Clones benefit scenarios where there are multiple VMs reading a single source of data, such as deployment servers and repositories. VDI deployments, where many linked clones forward read requests to a central master (e.g., Citrix MCS Master VM or VMware View replica disks), are an ideal example.

With Shadow Clones, Nutanix actively monitors vDisk access trends. If there are requests originating from more than two remote Controller VMs (CVMs), as well as the local CVM, and all of the requests are read I/O, the vDisk will be marked as immutable. Once the disk has been marked immutable, the vDisk is then cached locally by each CVM so read operations are now satisfied locally by direct-attached storage resources.

Shadow clones are enabled by default, and can be disabled if required.

If shadow clones are enabledfor this cluster, then this measure will report the value Yes. If they are disabled, then this measure will report the value No.

The numeric values that correspond to these measure values are as follows:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above</td></tr>
</table>

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>to indicate whether/not shadow clones are enabled for this cluster. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td>Is lockdown mode enabled?</td><td>Indicates whether/not the lockdown mode has been enabled for this cluster.</td><td></td><td>Cluster lockdown is the ability to disable password based CVM access and/or only allow key based access.

The values that this measure can report and their corresponding numeric values are listed in the table below:
<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Yes</td><td>1</td></tr><tr><td>No</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed in the table above to indicate whether/not lockdown mode is enabled. In the graph of this measure however, the same is represented using the numeric equivalents only.</td></tr>
<tr><td>Host hypervisors in the cluster</td><td>Indicates the number of hypervisors in this cluster.</td><td>Number</td><td>Use the detailed diagnosis of this measure to know which hypervisors have been added to which cluster. The name, IP address, and type of hypervisor is provided as part of the detailed metrics.</td></tr>
<tr><td>Total CPU capacity</td><td>Indicates the total CPU capacity of this cluster.</td><td>GHz</td><td></td></tr>
<tr><td>CPU used</td><td>Indicates the amount of CPU used by this cluster currently.</td><td>GHz</td><td>Ideally, the value of this measure should be low.</td></tr>
<tr><td>Available CPU</td><td>Indicates the amount of</td><td>GHz</td><td>Ideally, the value of this measure should</td></tr>
</table>

Measurement	Description	Measurement Unit	Interpretation
	CPU that is currently unused by this cluster.		be high.
CPU utilization	Indicates the percentage of CPU utilized by this cluster.	Percent	A value close to 100% indicates excessive utilization of CPU by the cluster, and points to a potential CPU contention on the cluster. If the consumption increases consistently, you may want to add more nodes to the cluster to ensure that more CPU resources are available to it. Compare the value of this measure across clusters to know which cluster is consuming the maximum CPU.
Available CPU	Indicates the percentage of CPU unused by this cluster.	Percent	If the value of this measure slides steadily, it is a cause for concern, as it indicates that CPU resources are being over-consumed by the cluster. If the problem persists, you may want to add more nodes to the cluster to ensure that more CPU resources are available to it.
Total IOPS	Indicates the total number of I/O operations on this cluster.	Number	This is a good indicator of the level of I/O activity on a cluster. A consistent increase in the value of this measure for a cluster could indicate a probable I/O overload on that cluster. In such situations, you may want to compare the value of the *Read IOPS* and *Write IOPS* measures of the cluster to know what type of IO operations are contributing to the overload.
Read IOPS	Indicates the number of read I/O operations performed currently on this cluster.	Number	If the value of the *Total IOPS* measure is unusually high, then compare the value of these measures for that cluster to know what is contributing to the unusual I/O activity levels - read requests? or write requests?
Write IOPS	Indicates the number of	Number	

Measurement	Description	Measurement Unit	Interpretation
	write I/O operations performed currently on this cluster.		
Total I/O latency	Indicates the average time taken by the physical disks in this cluster to process I/O requests .	Secs	Ideally, the value of this measure should be very low. A high value or a steady increase in this value could indicate an I/O processing bottleneck on this cluster. In such a case, compare the value of the *Read IO latency* and *Write IO latency* measures to figure out when the slowness is worst - when processing read requests? or write requests?
Read IO latency	Indicates the average time taken by the physical disks in this cluster to process read I/O requests.	Secs	If the *Total I/O latency* measure reports an abnormally high value, then compare the value of these measures to figure out where the slowness is maximum - when processing read requests? or write requests?
Write IO latency	Indicates the average time taken by the physical disks in this cluster to process write I/O requests.	Secs	
Total IO bandwidth	Indicates the bandwidth per second used by the physical disks in this cluster when processing I/O requests.	KB/Sec	A high value for this measure denotes that the cluster is processing bandwidth-intensive I/O. In such situations, you may want to compare the value of the *Read IO bandwidth* and *Write IO bandwidth m* measures to know what type of I/O requests are truly contributing to the excessive bandwidth consumptions - read requests? or write requests?
Read IO bandwidth	Indicates the bandwidth per second used by this clusterwhen processing read I/O requests.	KB/Sec	If the value of the *Total IO bandwidth* measure is high, then you may want to compare the value of the *Read IO*

Measurement	Description	Measurement Unit	Interpretation
			bandwidth and *Write IO bandwidth* measures to know what type of I/O requests are truly contributing to the excessive bandwidth consumption - read requests? or write requests?
Write IO bandwidth	Indicates the bandwidth per second used by this cluster when processing write I/O requests.	GB	
Read IO	Indicates the amount of data		
Read IO	Indicates the amount of data read from this cluster in response to read I/O requests.	KB	These measures are good indicators of the workload on the cluster.
Write IO	Indicates the amount of data written to this cluster in response to write I/O requests.	KB	
Total memory capacity	Indicates the total memory capacity of this cluster.	GB	
Memory used	Indicates the amount of memory that is currently in use in this cluster.	GB	Ideally, the value of this measure should be low.
Available memory	Indicates the amount of memory that is unused in this cluster.	GB	Ideally, the value of this measure should be high.
Memory utilization	Indicates the percentage of memory capacity currently being used in this cluster.	Percent	A value close to 100% indicates excessive memory usage by the cluster. If the high memory usage condition persists, you may want to consider adding more nodes to the cluster to increase its memory capacity. If memory capacity is not expanded, then the cluster will soon

Measurement	Description	Measurement Unit	Interpretation
			exhaust its memory resources, putting the health of both the hypervisors and VMs in the cluster at peril.
Available memory	Indicates the percentage of memory capacity that is currently unused by this cluster.	Percent	If the value of this measure slides consistently, it could mean that the cluster's memory resources are being rapidly depleted. If the condition persists, then the cluster will soon exhaust its memory resources, putting the health of both the hypervisors and VMs in the cluster at peril. Under such circumstances, you may want to consider adding more nodes to the cluster to increase its memory capacity.
Total storage capacity	Indicates the total storage capacity of this cluster.	GB	
Storage used	Indicates the amount of storage space that is currently being used by this cluster.	GB	Ideally, the value of this measure should be low.
Available storage	Indicates the amount of storage space that is currently available for the use of this cluster.	GB	Ideally, the value of this measure should be high.
Storage utilization	Indicates the percentage of storage capacity that is currently utilized byunu this cluser.	Percent	A value close to 100% indicates excessive space usage by the cluster. If the high space sage condition persists, you may want to consider adding more nodes to the cluster to increase its memory capacity. If storage capacity is not expanded, then the cluster will soon exhaust its available storage resources, putting the health of both the hypervisors and VMs in the cluster at peril.

Measurement	Description	Measurement Unit	Interpretation
Available storage	Indicates the percentage of storage capacity that unused in the cluster.	Percent	If the value of this measure slides consistently, it could mean that the cluster's storage resources are being rapidly depleted. If the condition persists, then the cluster will soon exhaust its storage resources, putting the health of both the hypervisors and VMs in the cluster at peril. Under such circumstances, you may want to consider adding more nodes to the cluster to increase its storage capacity.
Total VMs	Indicates the total number of VMs in this cluster.	Number	Use the detailed diagnosis of this measure to know the names, IP address, and operating system of the VMs in the cluster.
Powered on VMs	Indicates the number of VMs in this cluster that are currently powered on.	Number	
Powered off VMs	Indicates the number of VMs in this cluster that are currently powered off.	Number	
Controller VMs	Indicates the number of CVMs in this cluster.	Number	The detailed diagnosis of this measure lists the name, IP address, and operating system of the CVMs, and which hypervisor each CVM runs on.

Use the detailed diagnosis of the *Host hypervisors in cluster* measure to know which hypervisors have been added to which cluster. The name, IP address, and type of hypervisor is provided as part of the detailed metrics.

Hypervisors details

HOST NAME	HOST IP	HOST TYPE
Sep 23, 2016 13:58:22		
NTNX-13SX35490019-B	10.20.18.3	VMware
NTNX-13SX35490019-A	10.20.18.32	VMware
NTNX-13SX35490019-C.nutanixdc.local	10.20.18.4	VMware
NTNX-13SX35490019-D	10.20.18.5	VMware

Figure 3.5: The detailed diagnosis of the Host hypervisors in cluster measure

Use the detailed diagnosis of the *Total VMs* measure to know the names, IP address, and operating system of the VMs in the cluster.

Details of VMs

VM NAME	VM IP	VM OS	HYPERVISOR NAME
Sep 23, 2016 13:58:22			
sure_ubuntu14_01	-	Ubuntu Linux (64-bit)	NTNX-13SX35490019-D
Management	fe80::b0ec:ffa3:5ec3:88d8/64,169.254.136.216/0	Microsoft Windows Server 2012 (64-bit)	NTNX-13SX35490019-B
vm-490764fd-6f0c-4bd2-a1ed-62d0460ce074	10.20.28.24/24,fe80::250:56ff:fea4:4597/64	Ubuntu Linux (64-bit)	NTNX-13SX35490019-B
Z-ZVM-BizDEv	-	Microsoft Windows Server 2012 (64-bit)	NTNX-13SX35490019-D
vCenterSrv	fe80::e962:3341:d5df:1b1a/64,10.20.28.215/24	Microsoft Windows Server 2012 (64-bit)	NTNX-13SX35490019-B
Pivotal-Jumpbox	fe80::10e:3b54:fbc1:f6b5/64,10.20.18.70/24	Microsoft Windows 10 (64-bit)	NTNX-13SX35490019-C.nutanixdc.local

Page 1 of 10 Displaying 1 - 10 of 92

Figure 3.6: The detailed diagnosis of the Total VMs measure

The detailed diagnosis of the Controller VMs measure lists the name, IP address, and operating system of the CVMs, and which hypervisor each CVM runs on.

Details of Controller VMs

VM NAME	VM IP	VM OS	HYPERVISOR NAME
Sep 23, 2016 13:58:22			
NTNX-13SX35490019-B-CVM	10.20.18.11/24,10.20.18.27/32,10.20.18.29/32,fe80::20c:29ff:fec9:32f/64,192.168.5.2/25,192.168.5.254/32,fe80::20c:29ff:fec9:339/64	CentOS 4/5/6/7 (64-bit)	NTNX-13SX35490019-B
NTNX-13SX35490019-D-CVM	10.20.18.13/24,fe80::20c:29ff:fe9b:7f11/64,192.168.5.2/25,192.168.5.254/32,fe80::20c:29ff:fe9b:7f1b/64	CentOS 4/5/6/7 (64-bit)	NTNX-13SX35490019-D
NTNX-13SX35490019-A-CVM	**10.20.18.10/24,fe80::20c:29ff:fe87:9af2/64,192.168.5.2/25,192.168.5.254/32,fe80::20c:29ff:fe87:9afc/64**	**CentOS 4/5/6/7 (64-bit)**	**NTNX-13SX35490019-A**
NTNX-13SX35490019-C-CVM	10.20.18.12/24,fe80::20c:29ff:fe74:18c9/64,192.168.5.2/25,192.168.5.254/32,fe80::20c:29ff:fe74:18d3/64	CentOS 4/5/6/7 (64-bit)	NTNX-13SX35490019-C.nutanixdc.local

Figure 3.7: The detailed diagnosis of the Controller VMs measure

4

Conclusion

This document has clearly explained how eG Enterprise monitors Nutanix Prism. For more information on eG Enterprise, please visit our web site at www.eginnovations.com or write to us at sales@eginnovations.com.