

vRealize Operations Tenant App 2.6

for Metering and Chargeback

Table of contents

| | |
|---|----|
| vRealize Operations Tenant App Overview | 3 |
| Licensing | 4 |
| Architecture and Dataflow | 5 |
| Pricing Policies | 5 |
| Base Settings | 5 |
| CPU Rate | 6 |
| Memory Rate | 8 |
| Price based on Sizing Policy..... | 8 |
| Storage Rate | 8 |
| Network Rate..... | 9 |
| Advanced Network Rate | 11 |
| Guest OS Rate | 11 |
| Cloud Director Availability | 11 |
| VCenter Tag Rate | 11 |
| A flag to enable/disable the publishing particular tag specific prices as metrics on VMs..... | 12 |
| VCD Metadata Rate | 12 |
| One Time Fixed Cost | 13 |
| Rate Factors..... | 13 |
| Additional Fixed Cost..... | 14 |
| Reports and Bills | 14 |
| Alerts | 15 |
| Differences between Tenant App and Chargeback Manager / Usage Meter | 15 |
| Appendix | 15 |
| vROps metrics for Metering the Infrastructure..... | 15 |
| Advanced Network Metrics..... | 25 |
| Tenant App - API usage | 27 |
| Get an Authorization Token | 27 |
| Resources query API | 28 |
| Stats Query API | 32 |

vRealize Operations Tenant App Overview

VMware supports thousands of partners to host and sell clouds built on VMware technology using vCenter and Cloud Director (VCD). While vCenter provides the core of virtualization, Cloud Director provides needed constructs for segmenting the virtual infrastructure appropriately and offering it as a service to tenants of these partners.

As can be imagined, there are several variants of infrastructure that can be sold by these partners such as “Pay as you go”, “Raw capacity” aka “Allocation based”, “Raw capacity with minimum guarantee” aka “Reservation based”, “Flex” i.e. “a combination of others”. Potentially combination of these could be offered to same tenants. It becomes challenging to track usage over a period and charge appropriately due to this complexity. In addition, tenants demand for a transparency in this billing, and it is imperative that service providers offer it.

vRealize Operations Tenant App (TA) sets out to solve these problems by metering the infrastructure. Further, it provides options to configure different models for pricing this metered infrastructure. Finally, it closes the loop by providing tenant specific views that help tenants validate their charges by looking at usage.

Architecture of TA is based on

- Data collector that collects data from vCenter, VCD and NSX,
- Pricing engine, that applies the charging policies defined by provider on the collected data

The collector and the pricing engine reside inside vRealize Operations Manager (vROps). TA appliance acts as a simple interface for creating pricing policies and storing generated “Bill”. TA appliance also provides a plugin to expose Tenant views inside VCD.

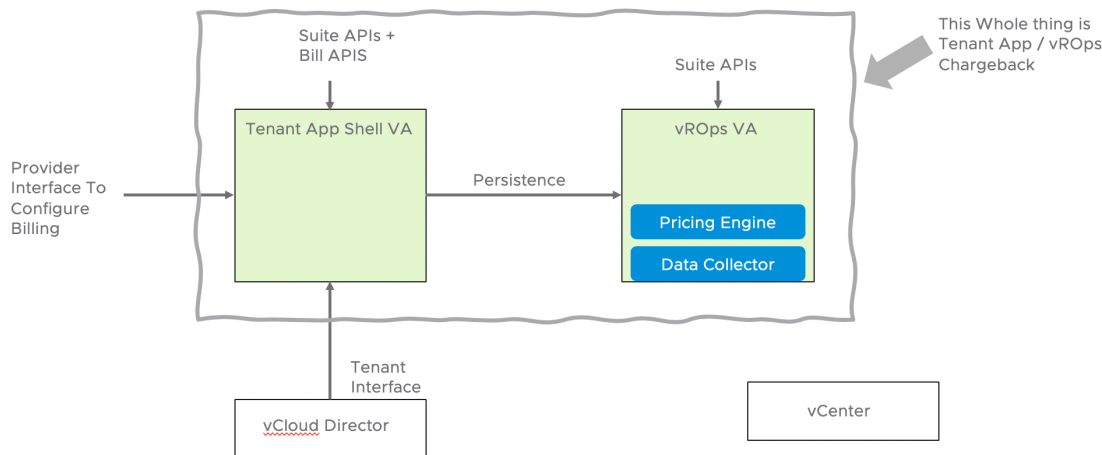
Because of these criterion, deployment of TA needs two Virtual Appliances namely the TA Virtual Appliance and the vROps Virtual Appliance. Partners who are already running vROps will be able to just connect Tenant App and begin, and those who do not can start with deploying the Chargeback Licensed edition of vROps.

Configuration of TA post deployment involves configuring vROps, Cloud Director, vCenter, and NSX. vROps connection is used for metering of collected data and bill generation. Cloud director connection is used for user management, providing Tenant view, and for collection of data from VCD using vROps VA. vCenter connection is used for collecting inventory changes and metering data about virtual machines using vROps. NSX connection is used for collecting usage information on network elements such as edges, load balancers, firewalls, etc. using vROps data collectors.

A functional flow of using TA involves

- Defining pricing policies
- Associating these pricing policies to Org- VDCs coming from VCD
- Generating bills and reports (showback) using this combination
- Exposing these bills and other usage data to Tenants
- Providing an isolated view of the infrastructure and services for each tenant

Functional Flow Diagram



Pricing policies have several variants such as their infrastructure scope of application (e.g.: Allocation pool, PAYG, Reservation pool), charge rates for different resources (e.g.: 2\$/VCPU/Day, 1\$/GB RAM configuration/Day) and conditions for applying the rates (E.g.: Only when powered on, slabbed rate or overage). Associating pricing policies to org VDCs sets up a combination that can be used henceforth for all metering and billing. Bills are created on-demand based on provider’s request and contain a detailed result of applying a pricing policy on and Org-VDC. Reports are customizable entities containing historical data about pricing and usage. They can be scheduled and exported to formats such as CSV and PDF. A plugin to VCD allows these bills and usage data to be exposed to tenants. In addition, Tenant App provides a set of dashboards that can be exposed to tenants.

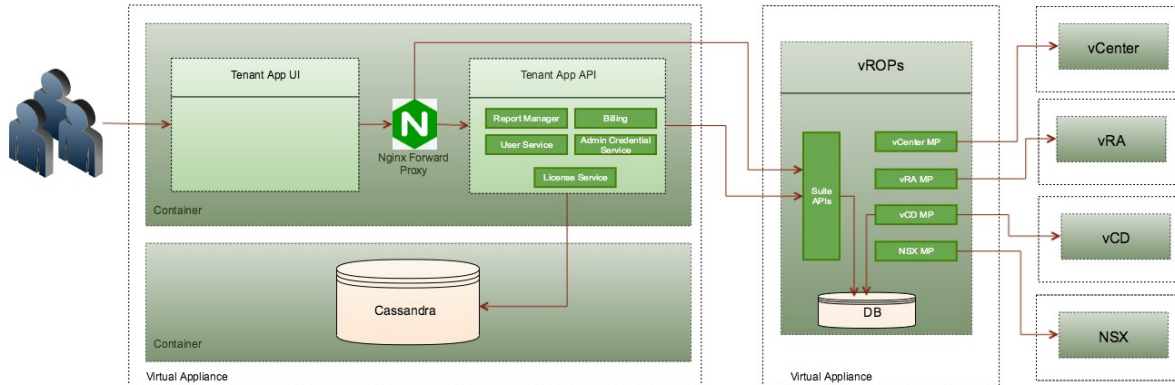
Rest of this paper expands on the above notions and details out each of them.

Licensing

Tenant App works in one of the two license modes.

- Chargeback License: This license mode enables basic functionalities of Tenant App including metering for Cloud Director, vCenter and NSX based infrastructure. The license key is applied from vROps UI, post which the UI of vROps is turned off and can be used only as a headless data collection and pricing engine.
- vROps entitled mode: If the partner is entitled to any variant of vROps already, he can continue to use the same vROps as a backend for Tenant App. In such cases, the partner gets additional flexibility of using vROps UI and perform customizations such as custom reports and dashboards based on his entitlement.

Architecture and Dataflow



As can be seen in the diagram, combination of Tenant App appliance and vROps appliance together help service providers meter their infrastructure. vROps acts as the data collection engine talking to endpoints such as vCenter and Cloud Director using its management packs. Installation and configuration of these management packs are covered in Administration Guide of Tenant App. Tenant App adds a pricing layer on top of these collected metrics by allowing configuration of the pricing engine. The pricing engine itself runs inside the vROps appliance and outputs the pricing metrics back into vROps. So in essence for a provider, Tenant App VA acts as the configuration interface for pricing policies, whereas vROps acts as the single source of metrics for all collected and processed information. Full list of collected and processed metrics can be found in the addendum at the end of this white paper.

Tenant App VA also acts as the User interface that can be configured and plugged into Cloud Director to be accessed by Tenant Administrators. It provides summary information of usage and metered charges to Tenants of a provider.

Virtual Machines are priced depending on how long they live in the environment. Tenant App and vROps can track VMs to a minimum granularity of 5 minutes and charge them i.e. even if the VM is alive only for 5 minutes, pricing engine will be able to associate a charge with this VM and roll it up.

Pricing Policies

This section provides details on various elements of pricing policy and how each of the components work in different scenarios, computation algorithm is described in detail.

Base Settings

Policy Name: Service provider admin will provide a relevant name to policy; one policy can be eventually assigned to multiple organization VDC. Each Organization in Cloud Director corresponds to one tenant, so pricing policy provides flexibility at much granular level (Org-VDC) so that differential charging can be done for different tenants (organizations) as well as multiple Org-VDCs within same tenant (organization).

Pricing Policy Type: This corresponds to types of allocation models (Allocation pool, Reservation pool and Pay-as-you-go) in VCD. This field is just used as an identifier only, so that when you create a policy you can assign a label to it so as to later associate it with appropriate organization VDC of same allocation model selected.

The choice here decides how you charge for CPU and memory. For details, refer to Table 1.1

Currency: This is the currency in which pricing is setup and bills are generated, this currency is same as set in vROps. In vROps navigate to Administration > Management > Global Settings > Currency to set this value. Please note that this can be set only once per installation and cannot be modified once set.

Policy Description: Any additional details to be mentioned about pricing policy

CPU Rate

Charge CPU based on: This is applicable when admin selects policy type as PAYG, where pricing happens at virtual machine level, and admin can choose to charge CPU based on either gHZ or based on vCPU count.

Charge Period: A specific price can be applied with different periodicity such as Hourly, Daily and Monthly. The base rate and fixed cost mentioned will get applied based on this frequency.

Charge Based on: For each allocation model in VCD admin can define specific rules based on which chargeback can happen, it can be based on resource allocation, usage or reservation (guaranteed resource) or max of two entities (usage vs allocation, usage vs reservation). The details of which value will be considered in computation based on this set rule is described in the below table.

Table 1.1

| Pricing Policy Type / Allocation Model | Charge Based on | Power State | CPU | Memory |
|--|--------------------------|----------------------------|--|--|
| Allocation Pool | Charged on | | Org-VDC | Org-VDC |
| | Allocation | Not Applicable / No Impact | CPU allocation (GHZ) | Memory allocation (GB) |
| | Reservation | | CPU resources guaranteed (GHZ) | Memory resources guaranteed (GB) |
| | Usage | | CPU Allocation Used (GHZ) | Memory Allocation Used (GHZ) |
| | Max (Allocation, Usage) | | Usage can go beyond what is allocated, so take whichever is higher | Usage can go beyond what is allocated, so take whichever is higher |
| | Max (Reservation, Usage) | | Usage can go beyond guaranteed, so take whichever is higher | Usage can go beyond guaranteed, so take whichever is higher |
| Reservation Pool | Charged on | | Org-VDC | Org-VDC |
| | Reservation | | Same as allocation, as 100% is guaranteed | Same as allocation, as 100% is guaranteed |
| | Usage | | CPU Reservation used (GHZ) | Memory Reservation used (GB) |
| | Max (Allocation, Usage) | | Usage can go beyond what is allocated, so take whichever is higher | Usage can go beyond what is allocated, so take whichever is higher |

| | | | | |
|---------------|--------------------------------|--|---|---|
| | Max (Reservation, Usage) | | Usage can go beyond guaranteed, so take whichever is higher | Usage can go beyond guaranteed, so take whichever is higher |
| Pay As You Go | Charged on | | VM | VM |
| | Allocation | | Amount of configured CPU ghz at a VM level | Amount of configured Memory GB at a VM level |
| | Usage | | CPU demand at the VM level | Memory demand at the VM level |

Default Base rate: Specific the rate or price at which you would like to charge the resource under consideration, this value will be multiplied accordingly with resource metered value to arrive at final charge to tenant in the bill generated.

Charge Based on Power State: This rule is applicable for PAYG type of policies, where charging is done at virtual machine level and chargeback can be done based on three available options: Always, Only when powered-on, Powered-on at least once.

- **Always:** do not consider VM uptime, charge irrespective of that
- **Only when powered-on:** Charge will be prorated based on uptime

E.g. if daily charge is 10\$ per gHZ for CPU and specific VM was only 20 min a day, then charge will be prorated based on the fraction (20/1440), where 1440 is total minutes in day.

- **Powered-on at least once:** In this case full charge is applied even if VM is switched on at least once for minimum duration of minute.

Charge Overage: This rule is applicable only to allocation pool model, where guaranteed resources are some % of total available resources. In this case till guaranteed %, normal base rate will be applied. If usage goes beyond guaranteed %, the rate mentioned in overage will be applied (for the delta usage which is higher than guaranteed).

E.g. in allocation pool model for specific Org-VDC if allocated CPU is 10 gHZ and guaranteed is 50%, which is 5 gHZ, the normal base rate is 3\$ per gHZ and overage base rate is 4\$, then if usage is found to be 6.5 gHZ, then for 1.5 gHZ the overage rate of 4\$ is applicable and below which normal base rate of 3\$ will be applicable.

Base Rate Slab: The service provider may want to charge differently based on the volume of resources consumed.

E.g. If usage is under 2 vCPUs, then default Base Rate \$4 per vCPU will be applied for whole usage. If usage is greater than 2 vCPU, then Base Rate \$6 per vCPU will be applied for whole usage.

Fixed Cost: This price is applied to resource under consideration irrespective of resource metering. The difference between base rate and fixed cost is that base rate gets multiplied with resource metered quantity, if there is no resource consumption/allocation base rate will not be applied whereas fixed cost gets applied irrespective of metered resource. Fixed cost follows the charge period.

E.g. charge period = If the per vCPU rate is \$2 and a fixed cost of \$10 is specified for a VM that has 4 vCPUs, charges will be computed as $4\text{vCPUs} * \$2 + \$10 = \$18$

Memory Rate

Please refer to the “CPU rate” section, as all the fields and rules are similar in nature except that resource under consideration will be memory metered in MBs.

Price based on Sizing Policy

There may be Service Providers that do not want to charge per vCPU or per GB RAM, instead they want standard pricing templates for charging their customers (similar to the concept in AWS where they have Small, Medium, Large, T2 Small, R2 Large etc.)

Sizing policies are predefined templates that are present in Cloud Director wherein you create a combo of vCPU and memory sizes and enter a rate against it.

In this, the price is based on VM Size instead of per CPU or Memory GB Charge individually.

Storage Rate

Charge Storage based on: Admin has two options here, Storage policies and Default rate.. When ‘storage policies’ option is selected, admin can do a ‘differential’ pricing based on storage tiers as set in VCD, for specific Org-VDC. When ‘Default rate’ is selected, for all the datastore same base rate mentioned will be applicable irrespective of storage tiers (Default Rate is a legacy option and will be deprecated soon).

E.g. if the storage policies created in Cloud Director are Gold, Silver and Bronze tiers, a differential price of 4\$, 3\$ and 2\$ per GB per month can be applied based on storage tiers.

Storage charges are not only applicable at the Org-VDC and VM level. But, they are also applicable for independent disks, media files, vApp templates.

Charge Based on: VMs can be charged based on configured storage by selecting ‘Limit’ as an option or based on the actual storage consumed by them based on ‘Usage’ as the option in pricing policy.

Base Rate Slab: The service provider may want to charge differently based on the storage consumed by the VM. So, in Storage also, we can charge based on the slab of consumption. Many Service providers provide a discounted rate if the consumption by VMs is higher

E.g. If you want to charge a VM consuming 100GB of storage at a lower rate (\$1 per month) compared to a VM consuming 50GB of storage that is charged at \$1.5 per GB per month. In such a case, base rates will be configured as “greater than or equal to:50GB”, “base rate:1”, “default base rate:1.5”. In such a case, a VM having 150GB of storage will be charged \$150 a month and a VM having 30GB of storage will be charged \$45 a month.

Create Pricing Policy

- 1 Base Settings
- 2 CPU Rate
- 3 Memory Rate
- 4 Storage Rate
- 5 Network Rate
- 6 Advanced Network Rate
- 7 Guest OS Rate
- 8 vCenter Tag Rate
- 9 VCD Metadata Rate
- 10 One Time Fixed Cost
- 11 Rate Factors
- 12 Review

Storage Rate

Charge storage based at Charge storage based at Storage Policy Default Rate

Base Rate

Storage Policy Name: Storage1

Any values entered below will be made part of pricing policy only after clicking the "Add" Button

Charge Period: Monthly

Charge Based on: Usage

Charge Based on Power State: Always

Default Base Rate: 1.5 per gb

Base Rate Slab: 0 Greater than or equal 0 Base Rate ADD SLAB

| Greater than or equal | Base Rate (₹) | Explanation | |
|-----------------------|---------------|---|----|
| 50 | 1 | If usage is 50 gb and Above, then Base Rate 1 will be applied for whole usage | 🗑️ |

ⓘ Any range of usage not covered in the above set of slabs will be applied with default base rate of 1.5

ADD

CANCEL
BACK
NEXT

Pricing Settings: There are two types of Storage Pricing Policy Settings

1. Aggregate storage charges from Storage profiles for PAYG Org-VDCs - this is the default mode for Storage Policy settings, and it includes overhead storage such as swap disks and log disks
2. Aggregate storage charges from VMs, media, vApp templates and independent disks for PAYG Org-VDCs – When the provider doesn't want to charge the overheads to the users

Please refer to the "CPU rate" section for other configuration parameters and rules, as all the fields and rules are similar in nature except that resource under consideration will be storage metered in GBs.

Network Rate

Tenant app charges network only at edge gateway level and not at the individual VM level.

Tenant App can charge for Edge gateways that are backed either by NSXV or NSXT. The same pricing rate card fields are used for charging either of these networks. The details of which metrics will be used in each case are given in the metrics table towards the end of this document.

External Network Transmit: Tenant app captures the egress data from edge gateway which in turn is associated with specific Org-VDC. Admin can apply a specific base rate 'per mb' of data transfer. The data is measured per edge gateway. All the traffic going through the network is charged here regardless of whether it's going through the internet or not

External Network Receive: Tenant app captures the ingress data from edge gateway which in turn is associated with specific Org-VDC. Admin can apply a specific base rate 'per mb' of data transfer. The data is measured per edge gateway. All the traffic going through the network is charged here regardless of whether it's going through the internet or not

Note: Charging VMs for the network consumption is not available today. But, it's a part of the roadmap

Network Transmit Rate (Bandwidth):

There are 3 ways of charging bandwidth here

1. **Conservative** – In this this case the Service Providers charge for the average bandwidth consumed by the customer. This is more in favour of the customer and less in that of the Service Provider

E.g. - You have a 10 GBPS line and the rate is \$10 per GBPS. But your average bandwidth usage is 5GBPS. Here you'll be charged \$50

2. **Aggressive** – Here, the Service Provider charges the whole usage at the peak bandwidth usage. This is more in favour of the Service Provider and less in that of the Customer

E.g. - The consumption rate is \$10 per GBPS. If the average bandwidth usage is 5GBPS. But, the peak bandwidth used is 8GBPS. In this case, the tenant be charged \$80 and not \$50

3. **Moderate** – In this case, the Service Provider would charge at the 95th Percentile value of the bandwidth.

E.g. – The user has had 100 usage samples of Bandwidths 1 to 100 GBPS each respectively. In this case, the peak bandwidth is 100 GBPS and average bandwidth is 50 GBPS. But, in a moderate pricing mechanism, the charge would be calculated on the basis of 95th percentile i.e. with respect to 95 GBPS.

Network Receive Rate (Bandwidth): Similar to Network Transmit rate, 95th percentile of ingress data is considered to charge for network receive bandwidth. These stats are retrieved via NSX manager APIs

There are 3 ways of charging bandwidth here

1. **Conservative** – In this this case the Service Providers charge for the average bandwidth consumed by the customer. This is more in favour of the customer and less in that of the Service Provider

E.g. - You have a 10 GBPS line and the rate is \$10 per GBPS. But your average bandwidth usage is 5GBPS. Here you'll be charged \$50

2. **Aggressive** – Here, the Service Provider charges the whole usage at the peak bandwidth usage. This is more in favour of the Service Provider and less in that of the Customer

E.g. - The consumption rate is \$10 per GBPS. If the average bandwidth usage is 5GBPS. But, the peak bandwidth used is 8GBPS. In this case, the tenant be charged \$80 and not \$50

3. **Moderate** – In this case, the Service Provider would charge at the 95th Percentile value of the bandwidth.

E.g. – The user has had 100 usage samples of bandwidths 1 to 100 GBPS each respectively. In this case, the peak bandwidth is 100 GBPS and average bandwidth is 50 GBPS. But, in a moderate pricing mechanism, the charge would be calculated on the basis of 95th percentile i.e. with respect to 95 GBPS.

Advanced Network Rate

Edge Services: Apart from the basic data transfer, there are additional value added services offered in Cloud Director in combination with NSX.

All the network services associated with specific edge such as HA, DHCP, IPV6, IP Sec, Load Balancer, NAT, SSL VPN, L2 VPN, Firewall, Static Routing, BGP Routing, OSPF Routing are considered for charging based on these services are 'Enabled' or not. If services are enabled for specific day and base rate is applied for that service, then that particular service gets charged for that specific day. If the service is disabled on any day then base rate will not get applied.

IP Count: These are the unique IP counts available on the external network of the Org-VDC. Pricing can be performed based on the count of these IPs.

Edge Gateway Sizes: Tenant app understands the size of the edge gateway (Compact, Large, Extra Large and Quad Large) and differential price can be assigned based on the edge size.

Guest OS Rate

Guest OS for a specific VM can be charged to apply based on licensing costs. The Guest OS name should be mentioned as displayed in the VCD for specific VM.

In future, Guest OS name will be a dropdown in tenant app policy creation for the admin to select from the available list.

Cloud Director Availability

Cloud Director availability services can be priced on the basis of replication objects that are created in a Cloud Director availability solution.

Replication objects occupy certain amount of storage and have special qualities dependent on SLA profile that they belong to.

E.g. If the storage policies created in Cloud Director are Gold, Silver and Bronze tiers, a differential price of 4\$, 3\$ and 2\$ per GB per month can be applied based on storage tiers. These prices for the replication objects will be the same depending on the storage profile they belong to.

Replication objects are also charged on the basis of the SLA profiles that they belong to. SLA profile defines attributes of these replication objects such as how frequently is the backup taking place, how many back-ups are preserved etc. Based on these SLA profiles, you can add additional costs on top of it.

E.g. If it is a 'High frequency' SLA profile, then you may want to charge an additional \$100 per month for the given replication object.

These replication objects will now be visible in the bill with their corresponding charges.

VCenter Tag Rate

If a VM is tagged with specific key (tag category) and value (tag value) in vCenter then this key-value can be specified to charge a VM or bunch of VMs differentially. Any VM having the specific tag will get this charge applied based on charge period.

vCenter tag based charges that are added can be individually saved or obtained based on a setting 'meteringTagDetailedMetricsEnabled'. When this setting is enabled, you can see each of the charges added here in this vCenter tag based section as an individual metric in vROps and hence in Tenant app as well.

How to use the config file

In EACH NODE within the vROps cluster should be performed the following steps:

1. Connect the vROps node via SSH.
2. Open and add/edit appropriate configuration in "\$ALIVE_BASE/user/conf/analytics/advanced.properties" file.
3. Restart the analytics service via "*service vmware-vcops restart analytics*" command.

A flag to enable/disable the publishing particular tag specific prices as metrics on VMs

| | |
|-----------------|---|
| Config name | meteringTagDetailedMetricsEnabled |
| Default value | false |
| Possible values | true/false |
| Example | <i>meteringTagDetailedMetricsEnabled=true</i> |

These can be used when you want to apply additional charge for the applications running on the VM

E.g. – You want to charge additional \$10 if there is an SQL server installed (across all VMs). In this case you can set a tag like “SQL Server = True” for these VMs in vCenter. Then, in Tenant App’s Pricing Policy, add a tag “SQL Server” with the value “True” and a rate of “\$10” per month. This would apply an additional charge of \$10 per month for all the VMs with the tag “SQL Server = True”

Alternate Pricing Policy: Within vCenter Tag rate, there is an ability to configure ‘Alternate Pricing Policy’, that allows the user to configure an additional charge for each vCPU that may have value added services enabled.

E.g. – For all VMs that have SQL server, you want to charge a different rate per vCPU per month. Say, a VM is charged at \$1 per vcpu per month and you want to charge the SQL enabled VMs at \$2 per vcpu per month. Here, you can create an alternate pricing policy like ‘SQL Server pricing policy’ that charges \$2 per vCPU and \$2 per GB RAM. In Tenant App Pricing Policy, you can apply a condition that if ‘SQL Server’ tag = ‘True’ then the Alternate pricing policy should be ‘SQL server pricing policy’

VCD Metadata Rate

Very similar to VM Tags (as defined in the vCenter), admin can make use of VCD specific tags which are known as VCD Metadata. Objects with which specific metadata is associated will get the charge based on pricing attached to it in VCD metadata rate section of pricing policy.

These can be used when you want to apply additional charge for the applications running on the VM

Metadata based charges that are added can be individually saved or obtained based on a setting ‘meteringTagDetailedMetricsEnabled’. When this setting is enabled, you can see each of the charges added here in this metadata based section as an individual metric in vROps and hence in Tenant app as well. <x ref to section above>

E.g. – You want to charge additional \$10 if there is an SQL server installed (across all VMs). In this case you can use a VCD metadata like “SQL Server = True” for these VMs in VCD. Then, in Tenant App pricing policy, add the metadata tag “SQL Server” with the value “True” and a rate of “\$10” per month. This would apply an additional charge of \$10 per month for all the VMs with the tag “SQL Server = True”

Alternate Pricing Policy: Within VCD Metadata, there is an ability to configure ‘Alternate Pricing Policy’, that allows the user to configure an additional charge for each vCPU that may have value added services enabled.

E.g. – For all VMs that have SQL server, you want to charge a different rate per vCPU per month. Say, a VM is charged at \$1 per vcpu per month and you want to charge the SQL enabled VMs at \$2 per vcpu per month. Here, you can create an alternate pricing policy like ‘SQL Server pricing policy’ that charges \$2 per vCPU and \$2 per GB RAM. In Tenant App Pricing Policy, you can apply a condition that if ‘SQL Server’ tag = ‘True’ then the Alternate pricing policy should be ‘SQL server pricing policy’

Note: This mechanism can effectively be used when charging for FLEX Org-VDCs

One Time Fixed Cost

One time fixed cost is a charge levied by the Service Providers for the events (service or resource provisioning) that happen at a point in time (not regularly or periodically).

E.g. The charge of a new VM creation can be added as a 'one time fixed cost'.

Or, the Service Provider may address a support ticket that can be charged additionally, every time there's a ticket that is addressed. This can also be configured with tags i.e. whenever you address a service request, you can add the tag (or VCD metadata)'SR addressed', the value 'true' on the VM. In Tenant App, if you have configured the Pricing Policy to add a 'one time fixed cost' when 'SR Addressed' = 'True' and one time fixed cost = '\$50', then the VM will get an additional 'one time fixed cost' of '\$50' in the bill. This one time fixed cost will be added every time the VM gets a tag 'SR Addressed' = 'True'

Note: In the above example, if the Service Provider intends to add 'SR Addressed' = 'True' cost every time an SR is addressed, then the recommended way is to set the tag 'SR addressed =True', leave it for a period of time (e.g. 2 hours) and then remove the tag.

Rate Factors

Rate factors are used to either bump up or discount the prices either against individual resources consumed by the Virtual Machines, or whole charges against the Virtual Machine.

E.g. 1. Discount overall charge on VM by 50% (Rate Factor 0.5) for all VMs tagged with Promo=True. So, if the cost of VM as per the base rate is \$100 then after applying the Rate Factor for promotional VMs (or Promo = True), it'll be \$50

VCD Metadata

For any resource with Metadata Promo True

Change the price of Total ▾

By applying a factor of 0.5 ⬆️⬇️⬆️

ADD

E.g. 2. Increase Storage rate by 100% (Rate Factor 2) for all VMs backed up by Avamar with tag 'Avamar_Backed_Up = True'. In this case, the storage cost of VM as per the base rate is \$100, then after applying the mentioned tag, it'll become \$200.

vCenter Tag

For any resource with Tag Avamar_Backed_Up True

Change the price of Storage ▾

By applying a factor of 2 ⬆️⬇️⬆️

ADD

Additional Fixed Cost

This section is used to add any other costs that are to be applied at Org-VDC level. This can be used for charges such as overall tax, overall discounts etc. These can be applied to selective Org-VDCs based on Org-VDC metadata

Create Pricing Policy

- 1 Base Settings
- 2 CPU Rate
- 3 Memory Rate
- 4 Storage Rate
- 5 Network Rate
- 6 Advanced Network Rate
- 7 Guest OS Rate
- 8 vCenter Tag Rate
- 9 VCD Metadata Rate
- 10 One Time Fixed Cost
- 11 Rate Factors
- 12 Additional Fixed Cost
- 13 Review

Additional Fixed Cost ✕

This section is used to add any other costs that are to be applied at Org-VDC level. This can be used for charges such as overall tax, overall discounts etc. These can be applied to selective Org-VDCs based on Org-VDC metadata

Fixed Cost

Charge Period Monthly ▾

Fixed cost 50

VCD Metadata

Metadata Snapshots Enabled True

Charge Period Daily ▾

Base Rate 10

ADD

VCD Metadata One Time

Metadata Installation Charge True

One time fixed cost 100

ADD

CANCEL BACK NEXT

E.g.

1. In the above screenshot, any Org-VDC with the above Pricing Policy enabled will be charged with \$50 as an additional fixed cost per month
2. Additionally, if the Org-VDC is tagged with the metadata 'Snapshots Enabled = True' will be charged \$10 per day
3. Also, if the Org-VDC has the 'Installation Charge = True' tag enabled, then an additional charge of \$100 will be applied as a one-time fixed cost.

Reports and Bills

Once these rates are set using the billing policy, they can be applied to organization VDCs in two different ways.

1. You can assign a pricing policy to an organization VDC once, and vROps starts automatically calculating price metrics for the given Org-VDC and its resources using the assigned pricing policy. Subsequently, either customized reports can be created in vROps containing usage and price metrics, or prebuilt reports can be accessed from Tenant App. The advantage of setting

up reports is that they can be scheduled and exported to formats such as PDF and CSV. You can also generate reports for the tenants to view directly at their end and send an email notification when reports become available.

2. You can generate an On-Demand Bill by selecting an organization VDC and a pricing policy. This calculates the prices of Org-VDC and associated resources as per the values of the Pricing Policy at that point in time and presents a Bill inside Tenant App. This bill can be also shared with tenants.
3. You can schedule bill generation, by selecting an Organization-VDC, a pricing policy and enable Schedule Bill setting. Then define the Recurrence (Monthly or Weekly), and detailed days of the billing period.

Note: This bill is typically used by the Service Providers via an API to read the details and feed to their own billing systems

Alerts

You can enable and select the type of alerts to be shared with the customer to view at their end. An email notification is also triggered to the tenants on the alerts enabled for them.

Differences between Tenant App and Chargeback Manager / Usage Meter

vRealize Operations Tenant App today does not present rollups of usage in units of MB.hours, but in terms of \$ cost. This can be worked around by setting a cost of \$1/MB.hr. vRealize Operations has all the individual data points, which can be accessed through the Tenant App APIs. A detailed description of the metrics used for the bill creation and API usage examples can be found in the Appendix.

However, the collection interval, collection methods and rollup methods of vRealize Operations differs from Chargeback Manager / Usage Meter.

Appendix

vROps metrics for Metering the Infrastructure

vROps and Tenant App Appliance together enable metering and charging of infrastructure offered by service providers to their consumers. vROps collects a series of metrics from vCenter with the help of vCenter adapter, and from Cloud Director with the help of Cloud Director management pack. These metrics are typically used by service providers that want raw metering metrics to feed into their own billing systems.

Context

vROps and Tenant App Appliance together enable metering and charging of infrastructure offered by service providers to their consumers. vROps collects a series of metrics from vCenter with the help of vCenter adapter, and from Cloud Director with the help of Cloud Director management pack. These metrics are typically used by service providers that want raw metering metrics to feed into their own billing systems.

Tenant App Appliance adds charging on top of these metrics wherein a rate card can be assigned against these collected metrics and Bills can be generated. These are used by service providers who want to perform billing in addition to metering the infrastructure using VMware solutions.

The third category of service providers consume the data from billing metrics but as a summary instead of using raw metrics using what is called \$1 rate mechanism. For example, if a service provider wanted to know the number of vCPU hours used under an Org-VDC within a month, only when VMs were powered on, first category of service providers would depend on the raw data of used vCPU metric and power state metric to create this number. However, using \$1 rate card mechanism, a service provider can configure \$1/hour as the vCPU charge based on power state, and simply use the output dollars to mean the number of vCPU hours.

Below we first describe general mechanisms to query a resource in vROps and obtain any metric on the resources. Next, the table describes typical metrics used by service providers for metering the infrastructure. These serve as definitive sources to the meaning of the metric and mechanism to obtain them for metering purposes.

[Resource Query API - How to Query Objects in vROps \(e.g.: VMs, Org-VDCs, Orgs\)](#)

[Stats Query API - How to Query metrics on the Objects in vROps and Summarize them](#)

| S L N O | METRIC | ENTITY | BASED ON | DESCRIPTION | VROPS STATS KEY | METRIC OR PROPER TY | SOURCE VC OR VCD OR NSX | STATS QUERY API RESPONSE (TIMESTAMPS & DATA) | COMMENTS | SUMMARY PARAMETER |
|------------------|--------|-----------------|-------------|--|-----------------------------|---------------------|-------------------------|--|--|-------------------|
| 1 | CPU | Virtual Machine | Reservation | CPU Effective limit (MHz) | cpuleff_limit | Metric | VC | 5 mins Interval | The response contains 5 mins sample measured by averaging 20 sec samples collected directly from ESXi box. This is the minimum of allocated MHz, Reserved MHz at VM level, Reserved MHz at resource pool of VM. This is the actual CPU limit for the VM. It's minimum of VM configured CPU capacity and CPU limit. | Average (AVG) |
| | | | Allocation | Total Capacity (MHz) | cpulvm_capacity_provisioned | Metric | VC | 5 mins Interval | This represents the allocate no of CPU MHz for the VM measured at 5 mins interval. CPU MHz is derived based on the no of vCPUs allocated to the VM & the CPU speed of ESXi hosting the VM. Configured capacity in MHz based on nominal (static) frequency of the CPU. | |
| | | | Usage | CPU Usage (MHz) | cpulusage_mhz | Metric | VC | 5 mins Interval | The response contains 5 mins sample measured by averaging 20 sec samples collected directly from ESXi box . The amount of CPU resources descendant virtual machines would use if there were no CPU contention or CPU limit. | |
| 2 | vCPU | Virtual Machine | Usage | Configurat ion Hardw are Numb er of CPUs (vCPUs) | confighardware_num_cpu | Metric | VC | 5 mins Interval | This is configured the No of vCPU for the VM collected at 5 mins interval. It counts both the vSocket and vCore. A VM with 2 vSockets x 4 vCores each has 8 vCPU. | Average (AVG) |
| 3 | Memory | Virtual Machine | Reservation | Memory Effective limit (KB) | memleff_limit | Metric | VC | 5 mins Interval | The response contains 5 mins sample measured by averaging 20 sec samples collected directly from ESXi box. This is the minimum of allocated MB, Reserved MB at VM level, Reserved MB at resource pool of VM. This is the actual memory limit for the VM. It's minimum of VM configured memory capacity and memory limit. | Average (AVG) |
| | | | Allocation | Memory Total | memlguest_p | Metric | VC | 5 mins Interval | This represents the allocate amount of memory for the VM (in KB) measured at 5 mins interval. | |

| | | | | | | | | | | |
|---|---------|-----------------|-------------------------------------|---|---|--------|------|------------------|---|---------------|
| | | | | Capacity (KB) | Provisioned | | | | | |
| | | | Usage | Memory Consumed (KB) | memory consumed | Metric | VC | 5 mins Interval | The response contains 5 mins sample measured by averaging 20 sec samples collected directly from ESXi box. Amount of memory utilized by the Virtual Machine. Reflects the guest OS memory required (for certain vSphere and VMTools versions) or Virtual Machine consumption. | |
| 4 | Storage | Virtual Machine | Usage | Disk Space Virtual Machine used(GB) | diskspace used | Metric | VC | 5 mins Interval | Used disk space by the VM as measured by the hypervisor (Not Guest OS). Space used by virtual machine files. | |
| | | | Allocation | Configuration Hardware Disk Space (GB) | config hardware diskSpace | Metric | VC | 5 mins Interval | Configured disk space for the VM. | |
| | | | Allocation Split by Storage Profile | Storage Profiles:<Storage Profile Name> usage | Storage Profiles:<Storage Profile Name> usage | Metric | VC D | 10 mins Interval | Configured disk space per storage profile allocated to the VM. | Average (AVG) |
| 5 | CPU | Org VDC | Usage | CPU Used (MHz) | cpu used | Metric | VC D | 5 mins Interval | This is the sum of no of vCPUs per VM in use with in this OVDC multiplied by 1000 MHz. For example, if an OVDC has 5 VMs of 2 vCPU each and only 3 of them are powered on at the point of collection, value of this metric at the OVDC level will be 3x2x1000 MHz. | Average (AVG) |
| | | | Allocation | CPU Allocation (MHz) | cpu allocation | Metric | VC D | 5 mins Interval | Allocated CPU MHz to this OVDC in Cloud Director. For the PAYG OVDC is always 0. | Average (AVG) |
| | | | Reservation | CPU Reserved (MHz) | cpu reserved | Metric | VC D | 5 mins Interval | Reserved CPU MHz to this OVDC in Cloud Director. For the PAYG OVDC is always 0. | Average (AVG) |
| 6 | vCPU | Org VDC | Usage | vCPU Used | vcpu used | Metric | VC D | 5 mins Interval | This is the sum of no of vCPUs per VM in use with in this OVDC. For example, if an OVDC has 5 VMs of 2 vCPU each and only 3 of them are powered on at the point of collection, value of this metric at the OVDC level will be 3x2. | Average (AVG) |
| | | | Usage | Memory Used (MB) | memory used | Metric | VC D | 5 mins Interval | This is the sum of amount of memory MB configured per VM in use with in this OVDC. For | Average (AVG) |

| | | | | | | | | | | |
|----|--------------------------------|-----------------|-------------|---|--|--------|------|-----------------|---|---------------|
| 7 | Memory | Org VDC | | | | | | | example, if an OVDC has 5 VMs of 2GB Memory each and only 3 of them are powered on at the point of collection, value of this metric at the OVDC level will be 3x2x1024 MB. | |
| | | | Allocation | Memory Allocation (MB) | memory allocation | Metric | VC D | 5 mins Interval | Allocated memory MB to this OVDC in Cloud Director. For the PAYG OVDC is always 0. | Average (AVG) |
| | | | Reservation | Memory Reserved (MB) | memory reserved | Metric | VC D | 5 mins Interval | Reserved memory MB to this OVDC in Cloud Director. For the PAYG OVDC is always 0. | Average (AVG) |
| 8 | Storage | Storage Profile | Usage | Storage Used (MB) | storage used | Metric | VC D | 5 mins Interval | Under every ODVC one Storage profile object is created for every storage profile OVDC has access to. This metric represents storage used by the OVDC on this storage profile. | Average (AVG) |
| | | | Limit | Storage Limit (MB) | storage limit | Metric | VC D | 5 mins Interval | It's just the storage Limit we collect, not at the profile level. | Average (AVG) |
| 15 | Count Of Networks | Org VDC | Usage | General Number of Organization VDC Networks | others Number of Organization VDC Networks | Metric | VC D | 5 mins Interval | Total no of OVDC network configured on this OVDC. | |
| 16 | Media | Media | Storage | General Storage (bytes) | General storageB | Metric | VC D | 5 mins Interval | Total Storage used by media file | Average (AVG) |
| 17 | vApp Template | vAPP Template | Storage | General Storage (Kilo bytes) | General storageKB | Metric | VC D | 5 mins Interval | Total Storage used by vApp Template | Average (AVG) |
| 18 | Enabled IPSec VPN Tunnel Count | Edge Gateway | Usage | Services:ipseclEnabled | service:ipseclenabled | Metric | VC D | 5 mins Interval | Count of VPN Tunnels on edge device | |
| 19 | NAT Service | Edge Gateway | Usage | Services:nat Enabled | service:nat enabled | Metric | VC D | 5 mins Interval | Boolean stating NAT enabled on edge device | |

| | | | | | | | | | | |
|----|----------------------------|--------------|-------|------------------------------------|-----------------------------------|--------|------|-----------------|--|-----------|
| 20 | DHCP Service | Edge Gateway | Usage | Services:dhcp Enabled | services:dhcp enabled | Metric | VC D | 5 mins Interval | Boolean stating DHCP is enabled on edge device | |
| 21 | FIREWALL Service | Edge Gateway | Usage | Services:firewall Enabled | services:firewall enabled | Metric | VC D | 5 mins Interval | Boolean stating Firewall Service is enabled on edge device | |
| 22 | Load Balancer Service | Edge Gateway | Usage | Services:loadbalancer Enabled | services:loadbalancer enabled | Metric | VC D | 5 mins Interval | Boolean stating Load balancer is enabled on edge device | |
| 23 | Static Routing Service | Edge Gateway | Usage | Services:staticRouting Enabled | services:staticRouting enabled | Metric | VC D | 5 mins Interval | Boolean stating Static routing is enabled on edge device | |
| 24 | Gateway HA Enabled | Edge Gateway | Usage | Services:highavailability Enabled | services:highavailability enabled | Metric | VC D | 5 mins Interval | Boolean stating Gateway is enabled on edge device | |
| 25 | Full Gateway Configuration | Edge Gateway | Usage | | | | | | | |
| 26 | External Network Transmit | Edge Gateway | Usage | Statistics Uplink Data OUT (bytes) | statistics uplink out | Metric | VC D | 5 mins Interval | This represents the amount of outward traffic as measured on this edge device at a 5 mins interval | Sum (SUM) |
| 27 | External Network Receive | Edge Gateway | Usage | Statistics Uplink Data IN (bytes) | statistics uplink in | Metric | VC D | 5 mins Interval | This represents the amount of inward traffic as measured on this edge device at a 5 mins interval | Sum (SUM) |

| | | | | | | | | | | |
|----|--------------------------------|-----------------|-------|--|---------------------------------|----------|------|------------------------------------|---|---------------------------|
| 28 | External Network Transmit Rate | Edge Gateway | Usage | Statistics Uplink Bandwidth OUT (MBps) | statistics uplink Bandwidth Out | Metric | VC D | 5 mins Interval | This represents the speed of outward traffic as measured on this edge device at a 5 mins interval | Max/Percentile 95/Average |
| 29 | External Network Receive Rate | Edge Gateway | Usage | Statistics Uplink Bandwidth IN (MBps) | statistics uplink Bandwidth In | Metric | VC D | 5 mins Interval | This represents the speed of inward traffic as measured on this edge device at a 5 mins interval | Max/Percentile 95/Average |
| 30 | Dynamic Routing OSPF | Edge Gateway | Usage | Services:ospfRouting Enabled | services:ospfRouting enabled | Metric | VC D | 5 mins Interval | Boolean stating Dynamic routing OSPF is enabled on edge device | |
| 31 | Dynamic Routing BGP | Edge Gateway | Usage | Services:bgpRouting Enabled | services:bgpRouting enabled | Metric | VC D | 5 mins Interval | Boolean stating Dynamic routing BGP is enabled on edge device | |
| 32 | L2 VPN | Edge Gateway | Usage | Services:l2vpn Enabled | services:l2vpn enabled | Metric | VC D | 5 mins Interval | Boolean L2 VPN is enabled on edge device | |
| 33 | SSL VPN | Edge Gateway | Usage | Services:sslvpn Enabled | services:sslvpn enabled | Metric | VC D | 5 mins Interval | Boolean stating SSL VPN is enabled on edge device | |
| 34 | Distributed Firewall | Edge Gateway | Usage | | | | | | Boolean stating Distributed firewall is enabled on edge device | |
| 35 | Power Status | Virtual Machine | | System Powered ON | sys powered On | Metric | VC | No of records based on Power State | Indicates 0 or 1 depending on power state of the VM | |
| 36 | Creation Date | Virtual | | Configuration Creation Date | config create Date | Property | VC | Creation Date | VM creation date. | |

| | | | | | | | | | | |
|----|------------------|-----------------|-------------|---|---------------------------------|--------|------|------------------|--|--|
| | | Machine | | | | | | | | |
| 37 | Uptime (Running) | Virtual Machine | | Summary Running | summary running | Metric | VC | 30 mins Interval | Indicates if the VM is powered on. The value is 1 if the VM is powered on and 0 otherwise. | |
| 38 | CPU | Organization | Usage | Aggregate Organization VDC CPU Used (MHz) | Aggregate VDC cpu used | Metric | VC D | 5 mins Interval | | |
| | | | Allocation | Aggregate Organization VDC CPU Allocation (MHz) | Aggregate VDC cpu allocation | Metric | VC D | 5 mins Interval | | |
| | | | Reservation | Aggregate Organization VDC CPU Reserved (MHz) | Aggregate VDC cpu reserved | Metric | VC D | 5 mins Interval | | |
| 39 | Memory | Organization | Usage | Aggregate Organization VDC Memory Used (MB) | Aggregate VDC memory used | Metric | VC D | 5 mins Interval | | |
| | | | Allocation | Aggregate Organization VDC Memory Allocation (MB) | Aggregate VDC memory allocation | Metric | VC D | 5 mins Interval | | |
| | | | Reservation | Aggregate Organization VDC Memory Reserved (MB) | Aggregate VDC memory reserved | Metric | VC D | 5 mins Interval | | |
| 40 | Storage | | Usage | Aggregate Organization | Aggregate VDC sto | Metric | VC D | 5 mins Interval | | |

| | | | | | | | | | | |
|---|--------------|--------------|------------|--|----------------------------------|--------|-----|-------------------|---|--|
| | | Organization | | VDC Storage Used (MB) | Usage | | | | | |
| | | | Allocation | Aggregate Organization VDC Storage Allocation (MB) | Aggregate VDC Storage Allocation | Metric | VCD | 5 mins Interval | | |
| | | | Limit | Aggregate Organization VDC Storage Limit (MB) | Aggregate VDC Storage Limit | Metric | VCD | 5 mins Interval | | |
| Metering Metrics (That can be used for Billing or \$1 rate mechanism) | | | | | | | | | | |
| 41 | Organization | Metering | | Summary Metering Total Cost | Summary metering value | Metric | VCD | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval | |
| 42 | Org VDC | Metering | | Summary Metering Additional | Summary metering Additional | Metric | VCD | 24 Hours Interval | Response contains the Additional price of the resource collected at 24 hours interval. Additional cost contains cost of TAGs, Custom properties, Guest OS, Network related items, metadatas, etc | |
| | | | | Summary Metering CPU | Summary metering CPU | | | 24 Hours Interval | Response contains Price of CPU of the resource collected at 24 hours interval. This will be calculated based on the configurations(<i>Charge period, Charge based on and power state</i>) provided in the pricing policy. If vCPU is selected, then base rate will be calculated based on the count of the vCPU or if CPU is selected then base rate will be calculated based on the GHz(used/allocated) for the CPU | |
| | | | | Summary Metering Memory | Summary metering Memory | | | 24 Hours Interval | Response contains the Price of memory of the resource collected at 24 hours interval. This will be calculated based on the configurations(<i>Charge period, Charge based on and power state</i>) provided in the pricing policy | |
| | | | | Summary Metering Partial Price | Summary metering Partial Price | | | 24 Hours Interval | Response shows whether the calculated price is partial for the resource (TRUE or FALSE), collected at 24 hours interval Partial price is always TRUE in scope of VCD | |

| | | | | | | | | | | |
|--------|------|----------|--|---------------------------------|---------------------------------|--------|-----|-------------------|---|--|
| | | | | Summary Metering storage | summary metering storage | | | 24 Hours Interval | Response contains the Price of storage of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy | |
| | | | | Summary Metering Total Cost | summary metering value | | | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval | |
| 4 3 | vApp | Metering | | Summary Metering additional | summary metering additional | Metric | VCD | 24 Hours Interval | Response contains the Additional price of the resource collected at 24 hours interval. Additional cost contains cost of TAGs, Custom properties, Guest OS, Network related items, metadatas, etc This cost/price will propagate to OVDC cost | |
| | | | | Summary Metering CPU | summary metering CPU | | | 24 Hours Interval | Response contains Price of CPU of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy. If vCPU is selected, then base rate will be calculated based on the count of the vCPU or if CPU is selected then base rate will be calculated based on the GHz(used/allocated) for the CPU This cost/price will propagate to OVDC cost | |
| | | | | Summary Metering memory | summary metering memory | | | 24 Hours Interval | Response contains the Price of memory of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy This cost/price will propagate to OVDC cost | |
| | | | | Summary Metering partialPrice | summary metering partialPrice | | | 24 Hours Interval | Response shows whether the calculated price is partial for the resource (TRUE or FALSE), collected at 24 hours interval Partial price is always TRUE in scope of VCD | |
| | | | | Summary Metering storage | summary metering storage | | | 24 Hours Interval | Response contains the Price of storage of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy This cost/price will propagate to OVDC cost only if "PROPAGATE_STORAGE_PROFILE_PRICE" is set to FALSE in vOps or Tenant App | |

| | | | | | | | | | | |
|----|----|----------|--|----------------------------------|----------------------------------|--------|-----|-------------------|--|--|
| | | | | Summary Metering Total Cost | summary metering value | | | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval This cost/price will propagate to OVDC cost | |
| 44 | VM | Metering | | Summary Metering Additional | summary metering additional | Metric | VCD | 24 Hours Interval | Response contains the Additional price of the resource collected at 24 hours interval. Additional cost contains cost of TAGs, Custom properties, Guest OS, Network related items, metadatas, etc This cost/price will propagate to VAPP cost | |
| | | | | Summary Metering CPU | summary metering CPU | | | 24 Hours Interval | Response contains Price of CPU of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy. If vCPU is selected, then base rate will be calculated based on the count of the vCPU or if CPU is selected then base rate will be calculated based on the GHz (used/allocated) for the CPU This cost/price will propagate to VAPP cost | |
| | | | | Summary Metering Memory | summary metering memory | | | 24 Hours Interval | Response contains the Price of memory of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy This cost/price will propagate to VAPP cost | |
| | | | | Summary Metering Partial Price | summary metering partial Price | | | 24 Hours Interval | Response shows whether the calculated price is partial for the resource (TRUE or FALSE), collected at 24 hours interval Partial price is always TRUE in scope of VCD | |
| | | | | Summary Metering Storage | summary metering storage | | | 24 Hours Interval | Response contains the Price of storage of the resource collected at 24 hours interval. This will be calculated based on the configurations (<i>Charge period, Charge based on and power state</i>) provided in the pricing policy This cost/price will propagate to VAPP cost only if "PROPAGATE_STORAGE_PROFILE_PRICE" is set to FALSE in vOps or Tenant App | |
| | | | | Summary Metering Total Cost | summary metering value | | | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval This cost/price will propagate to VAPP cost | |

| | | | | | | | | | | |
|----|--|-----------------|--|-------------------------------|-------------------------------|--------|-----|-------------------|--|--|
| | | | | StorageProfile instance usage | storageProfile instance usage | Metric | VCD | 24 Hours Interval | configured storage of the VM hard disk under selected storage profile | |
| 45 | Edge Gateway | Metering | | Summary Metering value | summary metering value | Metric | VCD | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval This cost is derived from additional cost - summary metering additional This cost/price will propagate to OVDC cost | |
| 46 | Storage Policy | Metering | | Summary Metering value | summary metering value | Metric | VCD | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval This cost is derived from storage cost - summary metering storage This cost/price will propagate to OVDC cost only if "PROPAGATE_STORAGE_PROFILE_PRICE" is set to TRUE in vROps or Tenant App | |
| 47 | vApp Template | Metering | | Summary Metering value | summary metering value | Metric | VCD | 24 Hours Interval | Response contains the Total price of the resource (Sum of all price components) collected at 24 hours interval This cost is derived from storage cost - summary metering storage This cost/price will propagate to VAPP cost | |
| 48 | Cloud Director availability Actual Usage | Virtual Machine | | StorageProfile actual usage | StorageProfile actual usage | metric | VCD | 5 minutes | The sum of real hard disk usage of a VM per storage profile | |
| 49 | Cloud Director availability Limit (configured) | Virtual Machine | | StorageProfile Configured | storageProfile usage | metric | VCD | 5 minutes | The sum of configured hard disk storage of a VM per storage profile | |

Advanced Network Metrics

Advanced Network Metrics for NSXV

| Edge Service | NSXV Object | Metric | VCD Management Pack Metric |
|----------------|--|-----------------------|--------------------------------|
| HA | Edge Service Gateway | service status | Services:HA Enabled |
| DHCP | DHCPEdgeService | service status | Services:DHCP Enabled |
| IPV6 | A property under NSX Manager > Network | | Services:ipv6 lis_enabled |
| IP Sec | IPSecEdgeService | service status | Services:Ipsec Enabled |
| LB | LoadBalancerEdgeService | service status | Services:Load Balancer Enabled |
| NAT | NATEdgeService | service status | Services:NAT Enabled |
| SSL VPN | SSLVPNEdgeService | service status | Services:SSLVPN Enabled |
| L2 VPN | L2VPNEdgeService | service status | Services:L2VPN Enabled |
| Firewall | FirewallEdgeService | status service_status | Services:firewall Enabled |
| Static Routing | RoutingEdgeService | service status | Services:Routing Enabled |
| BGP Routing | RoutingEdgeService | bgplenabled | Services:BGPI Enabled |
| OSPF Routing | LogicalRouter,EdgeServicesGateway | ospflenabled | Services:OSPF Enabled |

Note: The metrics are available via the NSX-V Management Pack as well as collected by VCD Management Pack. The pricing engine makes use of the metrics collected by VCD Management Pack

Advanced Network Metrics for NSXT

| Edge Service | NSXT Object | Metric |
|--------------|---------------------------|--|
| Data IN | Edge Gateway | Statistics Uplink Data IN |
| Data IN | Edge Gateway | Statistics Uplink Data OUT |
| HA | Logical Router | Services Enabled HA Status Per Transport Node <Transport-NodeID> HA Status |
| DHCP | Logical Router | Currently, metric isn't collected by MP |
| IPV6 | Logical Router | Currently, metric isn't collected by MP |
| IP Sec | Logical Router | Services Enabled IPSEC VPN Enabled |
| LB | Logical Router | Services Enabled Load Balancer Enabled |
| NAT | Router Service(NAT Rules) | Summary NAT Rule count |
| SSL VPN | Not Supported in NSXT | |
| L2 VPN | Logical Router | Services Enabled L2VPN Enabled |
| Firewall | Logical Router | Services Enabled Firewall Enabled |

| | | |
|----------------|-------------------------------|--------------------|
| Static Routing | Router Service (Static Route) | Enabled by default |
| BGP Routing | Router Service(BGP Service) | Summary Status |
| OSPF Routing | Not Supported in NSXT | |

Note: The above metrics are collected by NSX-T Management Pack. Pricing is not supported on these metrics as of Tenant App 2.5

Advanced Network Metrics from Cloud Director

| Edge Service | VCD Object | Metric |
|-------------------|---------------------|-----------------|
| IP Count | vCloud Edge Gateway | Summary IPCOUNT |
| Edge Gateway Size | vCloud Edge Gateway | Summary Size |

Configurations

Charge Period - Charge period can be Hourly/Daily/Monthly. Specified base rate will be applied based on the selected option

Charge Based on - Charge based on can be Usage/Reservation/Max from Usage and Reservation. Specified base rate will be applied based on the selected option

Charge Based on Power State - Power state can be Always/Only when powered on/Powered on at least once. Specified base rate will be applied based on the selected option

Tenant App - API usage

The Tenant App provides a REST API to access pricing policies, bills and other services. The API documentation can be found on <https://{{tenant-app}}/tenant-app-api/swagger-ui.html>.

It's also possible to access vOps API endpoints through the <https://{{tenant-app}}/suite-api> URL.

These are the steps to get metrics for individual objects:

- Get an Authorization Token
- Get the identifier of the object using the Resources query API
- Get the metrics using the Stats query API

Get an Authorization Token

POST Request: <https://{{tenant-app}}/suite-api/api/auth/token/acquire>

tenant-app}}/ suite-api/api/auth/token/acquire

Request Body:

```
{
  "username" : "{{user}}",
  "authSource" : "local",
  "password" : "{{pass}}"
}
```

Request Response:

```
{
  "token": "5d318488-fed1-4581-914d-ae51a110fe22::fdf29e35-91d6-467b-a976-000e2452f34c",
  "validity": 1588317978927,
  "expiresAt": "Friday, May 1, 2020 7:26:18 AM UTC",
  "roles": []
}
```

The token value can then be used in the following calls:

Resources query API

GET Request: <https://{{tenant-app}}/suite-api/api/resources>

possible request parameters:

- **adapterKind** - vCloud, VMWARE
- **resourceKind** - none, ORG, ORG_VDC, VAPP, VirtualMachine, PRO_VDC
- **name** - "Name of the Entity" (like match)
- **adapterInstanceld** - "ID of Adapter Instances"
- **parentId** - "ID of direct parent"

Examples: To get VMs with name containing (vAPP_91_VM1)

https://{{tenant-app}}/suite-api/api/resources?adapterKind=VMWARE&resourceKind=VirtualMachine&parentId=1b71001a-a643-4a35-bb85-ade906891f72&name=vAPP_91_VM1

RequestHeader:

Authorization: vRealizeOpsToken 5d318488-fed1-4581-914d-ae51a110fe22::fdf29e35-91d6-467b-a976-000e2452f34c

Accept: application/json

Sample Response: (find the highlighted text for Resource ID)

| GET RESOURCE RESPONSE |
|---|
| <pre>{ "pageInfo": { "totalCount": 1, "page": 0, "pageSize": 1000 }, "links": [{ "href": "/suite-api/api/resources?adapterKind=VMWARE&resourceKind=VirtualMachine&parentId=1b71001a-a643-4a35-bb85-ade906891f72&name=vAPP_91_VM1&page=0&pageSize=1000", "rel": "SELF", "name": "current" }, { "href": "/suite-api/api/resources?adapterKind=VMWARE&resourceKind=VirtualMachine&parentId=1b71001a-a643-4a35-bb85-ade906891f72&name=vAPP_91_VM1&page=0&pageSize=1000", "rel": "RELATED", "name": "first" }] }</pre> |

```

    },
    {
      "href": "/suite-api/api/resources?adapterKind=VMWARE&resourceKind=VirtualMachine&parentId=1b71001a-a643-4a35-
bb85-ade906891f72&name=vAPP_91_VM1&page=0&pageSize=1000",
      "rel": "RELATED",
      "name": "last"
    }
  ],
  "resourceList": [
    {
      "creationTime": 1575881661389,
      "resourceKey": {
        "name": "vAPP_91_VM1",
        "adapterKindKey": "VMWARE",
        "resourceKindKey": "VirtualMachine",
        "resourceIdentifiers": [
          {
            "identifierType": {
              "name": "VMEntityInstanceUUID",
              "dataType": "STRING",
              "isPartOfUniqueness": false
            },
            "value": ""
          },
          {
            "identifierType": {
              "name": "VMEntityName",
              "dataType": "STRING",
              "isPartOfUniqueness": false
            },
            "value": "vAPP_91_VM1"
          },
          {
            "identifierType": {
              "name": "VMEntityObjectID",
              "dataType": "STRING",
              "isPartOfUniqueness": true
            },
            "value": "vm-72"
          },
          {
            "identifierType": {
              "name": "VMEntityVCID",
              "dataType": "STRING",
              "isPartOfUniqueness": true
            },

```

```

    "value": "b92a41a0-ef91-4dd9-8420-ff96fa52db1d"
  },
  {
    "identifierType": {
      "name": "VMServiceMonitoringEnabled",
      "dataType": "STRING",
      "isPartOfUniqueness": false
    },
    "value": ""
  }
]
},
"resourceStatusStates": [
  {
    "adapterInstanceId": "b285d5d8-2747-44c6-957b-40be595a6e9f",
    "resourceStatus": "DATA_RECEIVING",
    "resourceState": "STARTED",
    "statusMessage": ""
  },
  {
    "adapterInstanceId": "152b29ff-9e10-4447-b06c-781f47527f1e",
    "resourceStatus": "DATA_RECEIVING",
    "resourceState": "STARTED",
    "statusMessage": ""
  }
],
"resourceHealth": "GREEN",
"resourceHealthValue": 100.0,
"dtEnabled": true,
"badges": [
  {
    "type": "TIME_REMAINING",
    "color": "GREEN",
    "score": 366.0
  },
  {
    "type": "CAPACITY_REMAINING",
    "color": "GREEN",
    "score": 100.0
  },
  {
    "type": "EFFICIENCY",
    "color": "GREEN",
    "score": 100.0
  },
  {

```

```

    "type": "RISK",
    "color": "GREEN",
    "score": 0.0
  },
  {
    "type": "HEALTH",
    "color": "GREEN",
    "score": 100.0
  },
  {
    "type": "WORKLOAD",
    "color": "GREEN",
    "score": 0.0
  },
  {
    "type": "COMPLIANCE",
    "color": "GREY",
    "score": -1.0
  }
],
"relatedResources": [],
"links": [
  {
    "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3",
    "rel": "SELF",
    "name": "linkToSelf"
  },
  {
    "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3/relationships",
    "rel": "RELATED",
    "name": "relationsOfResource"
  },
  {
    "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3/properties",
    "rel": "RELATED",
    "name": "propertiesOfResource"
  },
  {
    "href": "/suite-api/api/alerts?resourceId=3b4c649a-5bed-42eb-8da3-931d5c9971f3",
    "rel": "RELATED",
    "name": "alertsOfResource"
  },
  {
    "href": "/suite-api/api/symptoms?resourceId=3b4c649a-5bed-42eb-8da3-931d5c9971f3",
    "rel": "RELATED",
    "name": "symptomsOfResource"
  }
]

```

```

    },
    {
      "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3/statkeys",
      "rel": "RELATED",
      "name": "statKeysOfResource"
    },
    {
      "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3/stats/latest",
      "rel": "RELATED",
      "name": "latestStatsOfResource"
    },
    {
      "href": "/suite-api/api/resources/3b4c649a-5bed-42eb-8da3-931d5c9971f3/properties",
      "rel": "RELATED",
      "name": "latestPropertiesOfResource"
    },
    {
      "href": "/suite-api/api/credentials/",
      "rel": "RELATED",
      "name": "credentialsOfResource"
    }
  ],
  "identifier": "3b4c649a-5bed-42eb-8da3-931d5c9971f3"
}
]
}

```

You can use this resource id (value of "**identifier**" from the above response) to get the metrics.

Stats Query API

Sample request to get CPU metrics

POST Request: <https://tenant-app/suite-api/api/resources/stats/query>

Payload

```

{
  "resourceId": ["3b4c649a-5bed-42eb-8da3-931d5c9971f3"],
  "statKey": ["cpuldemandmhz"],
  "begin": 1578227562000,
  "end": 1578313962000,
  "metrics": true
}

```


To get the rolled-up metric values "rollUpType" parameter has to be added in the payload to get rolled up metric:

Payload with rollUpType

```
{
  "resourceId": ["3b4c649a-5bed-42eb-8da3-931d5c9971f3"],
  "statKey":["cpuldemandmhz"],
  "begin":1578227562000,
  "end":1578313962000,
  "metrics":true,
  "rollUpType": "SUM" //possible values - SUM,AVG, MIN, MAX, NONE, LATEST, COUNT
}
```

Suite-api Response

```
{
  "values": [
    {
      "resourceId": "3b4c649a-5bed-42eb-8da3-931d5c9971f3",
      "stat-list": {
        "stat": [
          {
            "timestamps": [
              1578227790954,
            ]
          }
        ]
      }
    }
  ]
}
```

