



Subscription Central 1-latest

Getting Started with the Subscriptions Service

Subscription Central 1-latest Getting Started with the Subscriptions Service

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

XFS[®] is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack[®] Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This guide is for users who want to understand how the subscriptions service reports usage data for their Red Hat subscriptions at the Red Hat account level. Procurement, operational, and technical teams can use the subscriptions service to help them understand where Red Hat technology is being used, how much of it is being used, and whether they can use more or need to purchase more.

Table of Contents

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	5
PART I. ABOUT THE SUBSCRIPTIONS SERVICE	6
CHAPTER 1. WHAT IS THE SUBSCRIPTIONS SERVICE?	7
CHAPTER 2. WHAT ARE THE BENEFITS OF THE SUBSCRIPTIONS SERVICE?	8
CHAPTER 3. WHAT DOES THE SUBSCRIPTIONS SERVICE TRACK?	9
3.1. RED HAT ENTERPRISE LINUX	9
3.1.1. RHEL with a traditional Annual subscription	9
3.1.2. RHEL with a pay-as-you-go On-Demand subscription	9
3.2. RED HAT OPENSIFT	10
3.2.1. Impact of various factors on Red Hat OpenShift tracking	10
3.2.2. Core-based usage counting workflow for self-managed Red Hat OpenShift products	11
3.2.3. Understanding the subscribed cluster size compared to the total cluster size	11
3.2.4. Determining the subscribed cluster size	12
3.2.5. Red Hat OpenShift Container Platform with a traditional Annual subscription	13
3.2.6. Red Hat OpenShift Container Platform or Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription	13
3.2.7. Red Hat OpenShift Service on AWS Hosted Control Planes with a pre-paid plus On-Demand subscription	14
3.3. RED HAT CLOUD SERVICES	14
3.3.1. Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription	14
3.3.2. Red Hat Advanced Cluster Security for Kubernetes with a pay-as-you-go On-Demand subscription	14
PART II. REQUIREMENTS AND YOUR RESPONSIBILITIES	16
CHAPTER 4. REQUIREMENTS	17
4.1. RED HAT ENTERPRISE LINUX	17
4.2. RED HAT OPENSIFT	17
4.3. RED HAT CLOUD SERVICES	17
CHAPTER 5. HOW TO SELECT THE RIGHT DATA COLLECTION TOOL	19
5.1. RED HAT INSIGHTS	19
5.2. RED HAT SUBSCRIPTION MANAGEMENT	19
5.3. RED HAT SATELLITE	19
5.4. RED HAT OPENSIFT MONITORING STACK AND OTHER TOOLS FOR RED HAT OPENSIFT DATA COLLECTION	20
5.5. RED HAT OPENSIFT MONITORING STACK AND OTHER TOOLS FOR RED HAT CLOUD SERVICES DATA COLLECTION	20
5.6. CLOUD INTEGRATIONS FOR RED HAT SERVICES	20
CHAPTER 6. HOW TO SET SUBSCRIPTION ATTRIBUTES	22
6.1. SETTING SUBSCRIPTION ATTRIBUTES FOR RHEL	22
6.2. SETTING SUBSCRIPTION ATTRIBUTES FOR RED HAT OPENSIFT	23
6.3. SETTING SUBSCRIPTION ATTRIBUTES FOR RED HAT CLOUD SERVICES	24
CHAPTER 7. YOUR RESPONSIBILITIES	25
PART III. SETTING UP THE SUBSCRIPTIONS SERVICE FOR DATA COLLECTION	26
CHAPTER 8. DEPLOYING RED HAT INSIGHTS	27
8.1. INSTALLING RED HAT INSIGHTS	27
8.2. WHAT DATA DOES RED HAT INSIGHTS COLLECT?	27

CHAPTER 9. INSTALLING THE SATELLITE INVENTORY UPLOAD PLUGIN	29
CHAPTER 10. REGISTERING SYSTEMS TO RED HAT SUBSCRIPTION MANAGEMENT	31
CHAPTER 11. CONNECTING RED HAT OPENSIFT TO THE SUBSCRIPTIONS SERVICE	32
CHAPTER 12. CONNECTING CLOUD INTEGRATIONS TO THE SUBSCRIPTIONS SERVICE	34
PART IV. ACTIVATING AND OPENING THE SUBSCRIPTIONS SERVICE	35
CHAPTER 13. DETERMINING WHETHER MANUAL ACTIVATION OF THE SUBSCRIPTIONS SERVICE IS NECESSARY	36
CHAPTER 14. ACTIVATING THE SUBSCRIPTIONS SERVICE	37
CHAPTER 15. LOGGING IN TO THE SUBSCRIPTIONS SERVICE	39
CHAPTER 16. VERIFYING ACCESS TO THE SUBSCRIPTIONS SERVICE	40
PART V. VIEWING AND UNDERSTANDING THE SUBSCRIPTIONS SERVICE DATA	41
CHAPTER 17. HOW DOES THE SUBSCRIPTIONS SERVICE SHOW MY SUBSCRIPTION DATA?	42
17.1. HOW TO USE THE SUBSCRIPTION DATA IN THE VIEWS	42
17.2. MEASUREMENT OF USAGE AND CAPACITY	48
17.2.1. Measurement of usage and capacity for Red Hat Enterprise Linux	48
17.2.2. Measurement of usage and capacity for Red Hat OpenShift	51
17.2.3. Measurement of usage and capacity for Red Hat Cloud Services	54
17.3. UNITS OF MEASUREMENT	56
17.3.1. Units of measurement for Red Hat Enterprise Linux	56
17.3.2. Units of measurement for Red Hat OpenShift	57
17.3.3. Units of measurement for Red Hat Cloud Services	59
17.4. FILTERING	59
CHAPTER 18. WHAT DATA DOES THE SUBSCRIPTIONS SERVICE STORE?	62
CHAPTER 19. HOW THE SUBSCRIPTIONS SERVICE GETS AND REFRESHES DATA	63
PART VI. TROUBLESHOOTING AND COMMON QUESTIONS	67
CHAPTER 20. TROUBLESHOOTING: CORRECTING OVER-REPORTING OF VIRTUALIZED RHEL	68
CHAPTER 21. TROUBLESHOOTING: CORRECTING PROBLEMS WITH FILTERING	70
CHAPTER 22. HOW IS THE SUBSCRIPTION THRESHOLD CALCULATED?	71
CHAPTER 23. HOW IS CORE HOUR USAGE DATA CALCULATED?	72
23.1. AN EXAMPLE FOR RED HAT OPENSIFT ON-DEMAND SUBSCRIPTIONS	72
23.1.1. Units of measurement for Red Hat OpenShift On-Demand subscriptions	72
23.1.2. Example core hour usage calculation	73
23.1.3. Resolving questions about core hour usage	76
CHAPTER 24. HOW DO VCPUS, HYPER-THREADING, AND SUBSCRIPTION STRUCTURE AFFECT THE SUBSCRIPTIONS SERVICE USAGE DATA?	78
24.1. IMPROVED CALCULATIONS FOR X86-64 ARCHITECTURES WITH SIMULTANEOUS MULTITHREADING	78
24.2. IMPROVED ANALYSIS OF SUBSCRIPTION CAPACITY FOR CERTAIN SUBSCRIPTIONS	79
24.3. ISOLATING RED HAT OPENSIFT CLUSTER MANAGER OPERATIONAL METRICS FROM METRICS FOR THE SUBSCRIPTIONS SERVICE	79
CHAPTER 25. WHY IS A RHEL SERVER APPEARING IN AN UNEXPECTED PLACE IN THE SUBSCRIPTIONS	

SERVICE USAGE? 80

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your feedback on our documentation. To provide feedback, open a Jira issue that describes your concerns. Provide as much detail as possible so that your request can be addressed quickly.

Prerequisites

- You have a Red Hat Customer Portal account. This account enables you to log in to the Red Hat Jira Software instance. If you do not have an account, you will be prompted to create one.

Procedure

To provide your feedback, use the following steps:

1. Click the following link: [Create Issue](#)
2. In the **Summary** text box, enter a brief description of the issue.
3. In the **Description** text box, provide more details about the issue. Include the URL where you found the issue.
4. Provide information for any other required fields. Allow fields that contain default information to remain at the defaults.
5. Click **Create** to create the Jira issue for the documentation team.

A documentation issue will be created and routed to the appropriate documentation team. Thank you for taking the time to provide feedback.

PART I. ABOUT THE SUBSCRIPTIONS SERVICE

The subscriptions service in the Hybrid Cloud Console provides a visual representation of the subscription experience across your hybrid infrastructure in a dashboard-based application. The subscriptions service is intended to simplify how you interact with your subscriptions, providing both a historical look-back at your subscription usage and an ability to make informed, forward-facing decisions based on that usage and your remaining subscription capacity.



NOTE

The April 2021 release of the subscriptions service includes the following changes for how you access the subscriptions service:

- The subscription watch tool has a new name, and is now known as the subscriptions service.
- The primary navigation for the Hybrid Cloud Console at cloud.redhat.com has been redesigned. The subscriptions service has been relocated within the navigation tree for the individual product portfolios that it works with, Red Hat Enterprise Linux, Red Hat OpenShift, and Red Hat Cloud Services. Product page views generated by the subscriptions service are located within the **Subscriptions** submenu. This **Subscriptions** submenu might also include other subscription-related pages that are not directly related to the subscriptions service.

Learn more

- To learn more about the subscriptions service, see the following information:
 - [What is the subscriptions service?](#)
- To learn more about the benefits that the subscriptions service offers, see the following information:
 - [What are the benefits of the subscriptions service?](#)
- To learn more about the current capabilities of the subscriptions service, see the following information:
 - [What does the subscriptions service track?](#)

CHAPTER 1. WHAT IS THE SUBSCRIPTIONS SERVICE?

The subscriptions service provides reporting of subscription usage information across the constituent parts of your hybrid infrastructure, including physical and virtual technology deployments; on-premise and cloud environments; and cluster, instance, and workload use cases for select Red Hat product portfolios.

Currently the subscriptions service supports the following product portfolios:

- Unified reporting of Red Hat Enterprise Linux subscription usage information for physical, virtual, hypervisor, and public cloud based usage. This unified reporting model enhances your ability to consume, track, report, and reconcile your RHEL subscriptions with your purchasing agreements and deployment types.
- Reporting of Red Hat OpenShift Container Platform subscription usage information. The subscriptions service uses data available from Red Hat internal subscription services, in addition to data from Red Hat OpenShift reporting tools, to show aggregated cluster usage data in the context of different Red Hat OpenShift subscription types.
- Reporting of Red Hat Cloud Services subscription usage information. The subscriptions service also uses data available from some of the Red Hat OpenShift reporting tools to show usage of these services. These services consume resources differently, so the tracking of usage can vary per product. In general, usage can be represented as a combination of one or more metrics such as data transfer and data storage for workload activities, and instance availability for the consumption of control plane resources. Usage can also be represented as cluster usage data for the virtual cores.

The simplified, consistent subscription reporting experience shows your account-wide Red Hat subscriptions compared to your total inventory across all deployments and programs. It is an at-a-glance impression of both your account's remaining subscription capacity measured against a subscription threshold and the historical record of your software usage.

The subscriptions service provides increased and ongoing visibility of your subscription usage. By implementing it, you might be eligible to shift away from the challenges of the current content enforcement model for subscriptions. This older model can be error-prone and inconvenient for your operational workload requirements, while the newer model of content access and consumption results in fewer barriers to content deployment. The simple content access tool enables this shift to the newer model.

You can choose to use neither, either, or both of these services. However, the subscriptions service and simple content access are designed as complementary services and function best when they are used in tandem. Simple content access simplifies the subscription experience by allowing more flexible ways of consuming content. The subscriptions service provides account-wide visibility of usage across your subscription profile, adding governance capabilities to this flexible content consumption.

To learn more about the simple content access tool and how you can use it with the subscriptions service, see the [Getting Started with Simple Content Access](#) guide.



NOTE

As of April 2021, simple content access is now available to customers who manage subscriptions through Red Hat Satellite or Red Hat Subscription Management. Previously, simple content access was available only to Satellite customers. In addition, the previous restrictions that limited the use of simple content access to certain geographical regions during the early development of simple content access have now been lifted. Customers in all geographical regions can now use simple content access.

CHAPTER 2. WHAT ARE THE BENEFITS OF THE SUBSCRIPTIONS SERVICE?

The subscriptions service provides these benefits:

- Tracks selected Red Hat product usage and capacity at the fleet or account level in a unified inventory and provides a daily snapshot of that data in a digestible, filterable dashboard at cloud.redhat.com.
- Tracks data over time for self-governance and analytics that can inform purchasing and renewal decisions, ongoing capacity planning, and mitigation for high-risk scenarios.
- Helps procurement officers make data-driven choices with portfolio-centered reporting dashboards that show both inventory-occupying subscriptions and current subscription limits across the entire organization.
- With its robust reporting capabilities, enables the transition to simple content access tooling that features broader, organizational-level subscription enforcement instead of system-level quantity enforcement.

CHAPTER 3. WHAT DOES THE SUBSCRIPTIONS SERVICE TRACK?

The subscriptions service currently tracks and reports usage information for Red Hat Enterprise Linux, some Red Hat OpenShift products, and some Red Hat Cloud Services services.

The subscriptions service identifies subscriptions through their stock-keeping units, or SKUs. Only a subset of Red Hat SKUs are tracked by the subscriptions service. In the usage reporting for a product, the tracked SKUs in your account contribute to the maximum capacity information, also known as the subscription threshold, for that product.

For the SKUs that are not tracked, the subscriptions service maintains an explicit deny list within the source code. To learn more about the SKUs that are not tracked, you can view this deny list in the code repository.

3.1. RED HAT ENTERPRISE LINUX

The subscriptions service tracks RHEL Annual subscription usage on physical systems, virtual systems, hypervisors, and public cloud. For a limited subset of subscriptions, currently Red Hat Enterprise Linux Extended Life Cycle Support Add-On on Amazon Web Services (AWS), it tracks RHEL pay-as-you-go On-Demand subscription usage for instances running in public cloud providers.

If your RHEL installations predate certificate-based subscription management, the subscriptions service will not track that inventory.

3.1.1. RHEL with a traditional Annual subscription

The subscriptions service tracks RHEL usage in sockets, as follows:

- Tracks physical RHEL usage in CPU sockets, where usage is counted by socket pairs.
- Tracks virtualized RHEL by the installed socket count for standard guest subscriptions with no detectable hypervisor management, where one virtual machine equals one socket.
- Tracks hypervisor RHEL usage in CPU sockets, with the socket-pair method, for virtual data center (VDC) subscriptions and similar virtualized environments. RHEL based hypervisors are counted both for the copy of RHEL that is used to run the hypervisor and the copy of RHEL for the virtual guests. Hypervisors that are not RHEL based are counted for the copy of RHEL for the virtual guests.
- Tracks public cloud RHEL instance usage in sockets, where one instance equals one socket.
- Additionally, tracks Red Hat Satellite to enable the visibility of RHEL that is bundled with Satellite.

3.1.2. RHEL with a pay-as-you-go On-Demand subscription

The subscriptions service tracks metered RHEL in vCPU hours, as follows:

- Tracks pay-as-you-go On-Demand instance usage in virtual CPU hours (vCPU hours), a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHEL pay-as-you-go On-Demand subscription usage, availability for computational activity is the availability of the RHEL instance over time.

**NOTE**

Currently, Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on is the only RHEL pay-as-you-go On-Demand subscription offering that is tracked by the subscriptions service.

The subscriptions service ultimately aggregates all instance vCPU hour data in the account into a monthly total, the unit of time that is used by the billing service for the cloud provider marketplace.

3.2. RED HAT OPENSIFT

Generally, the subscriptions service tracks Red Hat OpenShift usage as cluster size on physical and virtual systems. The *cluster size* is the sum of all subscribed nodes. A *subscribed node* is a compute or worker node that runs workloads, as opposed to a control plane or infrastructure node that manages the cluster.

However, beyond this general rule, tracking is dependent on several factors:

- The Red Hat OpenShift product
- The type of subscription that was purchased for that product
- The version of that product
- The unit of measurement for the product, as defined by the subscription terms, that determines how cluster size and overall usage is calculated
- The structure of nodes, including any labels used to assign node roles and the configuration of scheduling to control pod placement on nodes

3.2.1. Impact of various factors on Red Hat OpenShift tracking

The subscriptions service tracks and reports usage for fully managed and self-managed Red Hat OpenShift products in both physical and virtualized environments. Due to changes in the reporting models between Red Hat OpenShift major versions 3 and 4, usage data for version 3 is reported at the node level, while usage data for version 4 is reported and aggregated at the cluster level. The following information is more applicable to the version 4 reporting model, with data aggregated at the cluster level.

Much of the work for the counting of Red Hat OpenShift usage takes place in the monitoring stack tools and OpenShift Cluster Manager. These tools then send core count or vCPU count data, as applicable, to the subscriptions service for usage reporting. The core and vCPU data is based on the subscribed cluster size, which is derived from the cluster nodes that are processing workloads.

For fully managed Red Hat OpenShift products, such as Red Hat OpenShift Dedicated or Red Hat OpenShift AI, usage counting is generally time-based, measured in units such as core hours or vCPU hours. The infrastructure of the Red Hat managed environment is more consistently available to Red Hat, including the monitoring stack tools and OpenShift Cluster Manager. Data for subscribed nodes, the nodes that can accept workloads, is readily discoverable, as is data about cores, vCPUs and other facts that contribute to usage data for the subscriptions service.

For self-managed Red Hat OpenShift products, such as Red Hat OpenShift Container Platform Annual and Red Hat OpenShift Container Platform On-Demand, usage counting is generally based on cores. The infrastructure of a customer-designed environment is less predictable, and in some cases facts pertinent to usage calculations might be less accessible, especially in a virtualized, x86-based environment.

Because some of these facts might be less accessible, the usage counting process contains assumptions about simultaneous multithreading, also known as hyperthreading, that are applied when analyzing and reporting usage data for virtualized Red Hat OpenShift clusters for x86 architectures. These assumptions are necessary because some vendors offer hypervisors that do not expose data about simultaneous multithreading to the guests.

Ongoing analysis and customer feedback have resulted in incremental improvements, both to the subscriptions service and to the associated data pipeline, that have enhanced the accuracy of usage counting for hyperthreading use cases. The foundational assumption currently used in the subscriptions service reporting is that simultaneous multithreading occurs at a factor of 2 threads per core. Internal research has shown that this factor is the most common configuration, applicable to a significant majority of customers. Therefore, assuming 2 threads per core follows common multithreading best practices and errs in favor of the small percentage of customers, approximately 10%, who are not using multithreading. This decision is the most equitable for all customers when deriving the number of cores from the observed number of threads.



NOTE

A limited amount of self-managed Red Hat OpenShift offerings are available as socket-based subscriptions. For those socket-based subscriptions, the hypervisor reports the number of sockets to the operating system, usually Red Hat Enterprise Linux CoreOS, and that socket count is sent to the subscriptions service for usage tracking. The subscriptions service tracks and reports socket-based subscriptions with the socket-pair method that is used for RHEL.

3.2.2. Core-based usage counting workflow for self-managed Red Hat OpenShift products

For self-managed Red Hat OpenShift products such as Red Hat OpenShift Container Platform Annual and Red Hat OpenShift Container Platform On-Demand, the counting process initiated by the monitoring stack tools and OpenShift Cluster Manager works as follows:

1. For a cluster, node types and node labels are examined to determine which nodes are subscribed nodes. A *subscribed node* is a node that can accept workloads. Only subscribed nodes contribute to usage counting for the subscriptions service.
2. The chip architecture for the node is examined to determine if the architecture is x86-based. If the architecture is x86-based, then simultaneous multithreading, also known as hyperthreading, must be considered during the usage counting.
 - a. If the chip architecture is not x86-based, the monitoring stack counts usage according to the cores associated with the subscribed nodes and sends that core count to the subscriptions service.
 - b. If the chip architecture is x86-based, the monitoring stack counts usage according to the number of threads on the subscribed nodes. Threads equate to vCPUs, according to the Red Hat definition of vCPUs. This counting method applies whether the multithreading data can accurately be detected, the multithreading data is ambiguous or missing, or multithreading data is specifically set to a value of false on the node. Based on a global assumption of multithreading at a factor of 2, the number of threads is divided by 2 to determine the number of cores. The core count is then sent to the subscriptions service.

3.2.3. Understanding the subscribed cluster size compared to the total cluster size

For Red Hat OpenShift, the subscriptions service does not focus merely on the total size of the cluster

and the nodes within it. The subscriptions service focuses on the subscribed portion of clusters, that is, the cluster nodes that are processing workloads. Therefore, the subscriptions service reporting is for the *subscribed cluster size*, not the entire size of the cluster.

3.2.4. Determining the subscribed cluster size

To determine the subscribed cluster size, the data collection tools and the subscriptions service examine both the node type and the presence of node labels. The subscriptions service uses this data to determine which nodes can accept workloads. The sum of all noninfrastructure nodes plus master nodes that are schedulable is considered available for workload use. The nodes that are available for workload use are counted as subscribed nodes, contribute to the subscribed cluster size, and appear in the usage reporting for the subscriptions service.

The following information provides additional details about how node labels affect the countability of those nodes and in turn affect subscribed cluster size. Analysis of both internal and customer environments shows that these labels and label combinations represent the majority of customer configurations.

Table 3.1. How nodes contribute to the subscribed cluster size

Node label	Usage counted	Exceptions
worker	yes	Unless there is a combination of the worker label with an infra label
worker + infra	no	See Note
custom label	yes	Unless there is a combination of the custom label with the master, infra, or control plane label
custom label + master, infra, control plane (any combination)	no	
master + infra + control plane (any combination)	no	Unless there is a master label present <i>and</i> the node is marked as schedulable
schedulable master + infra, control plane (any combination)	yes	



NOTE

A known issue with the Red Hat OpenShift monitoring stack tools can result in unexpected core counts for Red Hat OpenShift Container Platform versions earlier than 4.12. For those versions, the number of worker nodes can be artificially elevated.

For OpenShift Container Platform versions earlier than 4.12, the Machine Config Operator does not support a dual assignment of infra and worker roles on a node. The counting of worker nodes is correct in OpenShift Container Platform according to the principles of counting subscribed nodes, and this count will display correctly in the OpenShift Container Platform web console.

However, when the monitoring stack tools analyze this data and send it to the subscriptions service and other services in the Hybrid Cloud Console, the Machine Config Operator ignores the dual roles and sets the role on the node to worker. Therefore, worker node counts will be elevated in the subscriptions service and in OpenShift Cluster Manager.

3.2.5. Red Hat OpenShift Container Platform with a traditional Annual subscription

The subscriptions service tracks Red Hat OpenShift Container Platform usage in CPU cores or sockets for clusters and aggregates this data into an account view, as refined by the following version support:

- RHOCP 4.1 and later with Red Hat Enterprise Linux CoreOS based nodes or a mixed environment of Red Hat Enterprise Linux CoreOS and RHEL based nodes
- RHOCP 3.11

For RHOCP subscription usage, there was a change in reporting models between the major 3 and 4 versions. Version 3 usage is considered at the node level and version 4 usage is considered at the cluster level.

The difference in reporting models for the RHOCP major versions also results in some differences in how the subscriptions service and the associated services in the Cloud Services platform calculate usage. For RHOCP version 4, the subscriptions service follows the rules for examining node types and node labels to calculate the subscribed cluster size as described in [Determining the subscribed cluster size](#). The subscriptions service recognizes and ignores the parts of the cluster that perform overhead tasks and do not accept workloads. The subscriptions service recognizes and tracks only the parts of the cluster that do accept workloads.

However, for RHOCP version 3.11, the version 3 era reporting model cannot distinguish the parts of the cluster that perform overhead tasks and do not accept workloads, so the reporting model cannot find the subscribed and nonsubscribed nodes. Therefore, for RHOCP version 3.11, you can assume that approximately 15% of the subscription data reported by the subscriptions service is overhead for the nonsubscribed nodes that perform infrastructure-related tasks. This percentage is based on analysis of cluster overhead in RHOCP version 3 installations. In this particular case, usage results that show up to 15% over capacity are likely to still be in compliance.

3.2.6. Red Hat OpenShift Container Platform or Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription

- RHOCP or OpenShift Dedicated 4.7 and later

The subscriptions service tracks RHOCP or OpenShift Dedicated 4.7 and later usage from a pay-as-you-go On-Demand subscription in core hours, a measurement of cluster size in CPU cores over a range of time. For an OpenShift Dedicated On-Demand subscription, consumption of control plane resources

by the availability of the service instance is tracked in instance hours. The subscriptions service ultimately aggregates all cluster core hour and instance hour data in the account into a monthly total, the unit of time that is used by the billing service for Red Hat Marketplace.

As described in the information about RHOCP 4.1 and later, The subscriptions service recognizes and tracks only the parts of the cluster that contain compute nodes, also commonly called worker nodes.

3.2.7. Red Hat OpenShift Service on AWS Hosted Control Planes with a pre-paid plus On-Demand subscription

The subscriptions service tracks Red Hat OpenShift Service on AWS Hosted Control Planes (ROSA Hosted Control Planes) usage from a pre-paid plus On-Demand subscription in vCPU hours and in control plane hours.

- A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms) for a total of one hour, measured to the granularity of the meter that is used. For ROSA Hosted Control Planes, availability for computational activity is the availability of the vCPUs for the ROSA Hosted Control Planes subscribed clusters over time. A subscribed cluster is comprised of subscribed nodes, which are the noninfrastructure nodes plus schedulable master nodes that are available for workload use, if applicable. Note that for ROSA Hosted Control Planes, schedulable master nodes are not applicable, unlike other products that also use this measurement. The vCPUs that are available to run the workloads for a subscribed cluster contribute to the vCPU hour count.
- A control plane hour is a measurement of the availability of the control plane. With ROSA Hosted Control Planes, each cluster has a dedicated control plane that is isolated in a ROSA Hosted Control Planes service account that is owned by Red Hat.

3.3. RED HAT CLOUD SERVICES

Because the services in the Red Hat Cloud Services portfolio consume different types of resources while handling different types of workloads, the subscriptions service tracks usage of these services in different ways.

3.3.1. Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription

The subscriptions service tracks Red Hat OpenShift AI (RHOAI) in vCPU hours, a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHOAI pay-as-you-go On-Demand subscription usage, the availability for computational activity is the availability of the cluster over time.

The subscriptions service ultimately aggregates all cluster vCPU hour data in the account into a monthly total, the unit of time that is used by the billing service for the cloud provider marketplace.

3.3.2. Red Hat Advanced Cluster Security for Kubernetes with a pay-as-you-go On-Demand subscription

The subscriptions service tracks Red Hat Advanced Cluster Security for Kubernetes (RHACS) in vCPU hours, a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHACS pay-as-you-go On-Demand subscription usage, the availability for computational activity is the availability of the cluster over time.

The subscriptions service aggregates all cluster vCPU hour data and then sums the data for each cluster where RHACS is running into a monthly total, the unit of time that is used by the billing service for the cloud provider marketplace.

Additional resources

- For more information about the purpose of the subscriptions service deny list, see the [What subscriptions \(SKUs\) are included in Subscription Usage?](#) article.
- For more information about the contents of the subscriptions service deny list, including the specific SKUs in that list, see the [deny list source code](#) in GitHub.

PART II. REQUIREMENTS AND YOUR RESPONSIBILITIES

Before you start using the subscriptions service, review the hardware and software requirements and your responsibilities when you use the service.

Learn more

- Review the general requirements for using the subscriptions service:
 - [Requirements](#)
- Review information about the tools that you must use to supply the subscriptions service with data about your subscription usage:
 - [How to select the right data collection tool](#)
- Review information about improving the subscriptions service results by setting the right subscription attributes:
 - [How to set subscription attributes](#)
- Review information about your responsibilities when you use the subscriptions service:
 - [Your responsibilities](#)

CHAPTER 4. REQUIREMENTS

To begin using the subscriptions service, you must meet the following software requirements. For more complete information about these requirements, contact your Red Hat account team.

4.1. RED HAT ENTERPRISE LINUX

You must meet at least one of the following requirements for Red Hat Enterprise Linux management:

- RHEL managed by Satellite.
 - The minimum Satellite version is 6.9 or later (versions that are under full support).
- RHEL managed by Red Hat Insights.
- RHEL managed by Red Hat Subscription Management.
- Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on with a pay-as-you-go On-Demand subscription requires configuration of a cloud integration between a cloud provider and the cost management service in the Hybrid Cloud Console.

4.2. RED HAT OPENSIFT

You must meet the following requirements for Red Hat OpenShift management, based on your product version and subscription type:

- Red Hat OpenShift Container Platform with an Annual subscription
 - RHOCP version 4.1 or later managed with the monitoring stack tools and OpenShift Cluster Manager.
 - RHOCP version 3.11 with RHEL nodes managed by Insights, Satellite, or Red Hat Subscription Management.
- RHOCP with a pay-as-you-go On-Demand subscription
 - RHOCP version 4.7 or later managed with the monitoring stack tools and OpenShift Cluster Manager.
- Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription
 - OpenShift Dedicated version 4.7 or later. The monitoring stack tools and OpenShift Cluster Manager are always in use for OpenShift Dedicated.

4.3. RED HAT CLOUD SERVICES

The Red Hat Cloud Services portfolio includes managed services that rely on Red Hat infrastructure. Part of that infrastructure is the Red Hat OpenShift monitoring stack tools that, among other jobs, supply data about subscription usage to the subscriptions service.

**NOTE**

Some of the services in the Red Hat Cloud Services portfolio might also gather and display their own usage data that is independent of the data that is gathered by the Red Hat OpenShift monitoring stack tools and displayed in the subscriptions service. The data displayed in these service-level dashboards is designed more for the needs of the owners of individual clusters, instances, and so on. However, the Red Hat OpenShift platform core capabilities provided by the monitoring stack tools typically gather and process the data that is used in the subscriptions service.

For the following services, no user setup of the monitoring stack tools is necessary:

- Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription
- Red Hat Advanced Cluster Security for Kubernetes with a pay-as-you-go On-Demand subscription

CHAPTER 5. HOW TO SELECT THE RIGHT DATA COLLECTION TOOL

To display data about your subscription usage, the subscriptions service requires a data collection tool to obtain that data. The various data collection tools each have distinguishing characteristics that determine their effectiveness in a particular type of environment.

It is possible that the demands of your environment require more than one of the data collection tools to be running. When more than one data collection tool is supplying data to the services in the Cloud Services platform, the tools that process this data are able to analyze and deduplicate the information from the various data collection tools into standardized facts, or *canonical facts*.

The following information can help you determine the best data collection tool or tools for your environment.

5.1. RED HAT INSIGHTS

Insights as a data collection tool is ideal for the always-connected customer. If you fit this profile, you are interested in using Insights not only as a data collection tool, but also as a solution that provides analytic, threat identification, remediation, and reporting capabilities.

With the inclusion of Insights with every Red Hat Enterprise Linux subscription beginning with version 8, and with the availability of Red Hat Insights for Red Hat OpenShift in April 2021, the use of Insights as your data collection tool becomes even more convenient.

However, using Insights as the data collection tool is not ideal if the Insights agent cannot connect directly to the `cloud.redhat.com` website or if Red Hat Satellite cannot be used as a proxy for that connection. In addition, it cannot be used as the sole solution if hypervisor host-guest mapping is required for virtual data centers (VDCs) or similar virtualized environments. In that case, Insights must be used in conjunction with Satellite.

5.2. RED HAT SUBSCRIPTION MANAGEMENT

Red Hat Subscription Management is an ideal data collection tool for the connected customer who uses the Subscription Manager agent to send data to Red Hat Subscription Management on the Red Hat Customer Portal.

For customers that are using the subscriptions service, Red Hat Subscription Management automatically synchronizes its data with the Cloud Services platform tools. Therefore, in situations where Red Hat Subscription Management is in use, or required, such as with RHEL 7 or later, it is being used as a data collection tool.

5.3. RED HAT SATELLITE

The use of Satellite as the data collection tool is useful for customers who have specific needs in their environment that either inhibit or prohibit the use of the Insights agent or the Subscription Manager agent for data collection.

For example, you might be able to connect to the Cloud Services platform directly, but you might find the connection and maintenance of a per-organization Satellite installation is more convenient than the per-system installation of Insights. The use of Satellite also enables you to inspect the information that is being sent to the Cloud Services platform on an organization-wide basis instead of a system-only basis.

As another example, your Satellite installation might not be able to connect directly to the Cloud Services platform because you are running Satellite from a disconnected network. In that case, you must export the Satellite reports to a connected system and then upload that data to the Cloud Services platform. To do this, you must use a minimum of Satellite 6.9 or later (versions that are under full support). You must also install the Satellite inventory upload plugin on your Satellite server.

Finally, you might have a need to view the subscriptions service results for RHEL usage from a virtual data center (VDC) subscription or similar virtualized environments. To do so, you must obtain accurate hypervisor host-guest mapping information as part of the data that is collected for analysis. This type of data collection requires the use of Satellite in combination with the Satellite inventory upload plugin and the virt-who tool.

5.4. RED HAT OPENSIFT MONITORING STACK AND OTHER TOOLS FOR RED HAT OPENSIFT DATA COLLECTION

The data collection for Red Hat OpenShift usage is dependent on several tools, including tools developed by the Red Hat OpenShift development team. One tool is Red Hat OpenShift Cluster Manager. Another set of tools is known as the monitoring stack. This set of tools is based on the open source Prometheus project and its ecosystem, and includes Prometheus, Telemetry, Thanos, Observatorium, and others.

The subscriptions service is designed to work with customers who use Red Hat OpenShift 4.1 and later products in connected environments. For the Red Hat OpenShift version 4.1 and later products that the subscriptions service can track, Red Hat OpenShift Cluster Manager and the monitoring stack tools are used to gather and process cluster data before sending it to Red Hat Subscription Management. Red Hat Subscription Management provides the relevant usage data to the Cloud Services platform tools such as inventory and the subscriptions service.

Customers with disconnected environments can use the Red Hat OpenShift data collection tools by manually creating each cluster in Red Hat OpenShift Cluster Manager. This workaround enables customers with disconnected environments to simulate an account-level view of their Red Hat OpenShift usage. For example, an organization with disconnected clusters distributed across several departments might find this workaround useful.

For Red Hat OpenShift Container Platform version 3.11, data collection is dependent on an older, RHEL based reporting model. Therefore, data collection is dependent upon the connection of the RHEL nodes to one of the RHEL data collection tools, such as Insights, Red Hat Subscription Management, or Satellite.

5.5. RED HAT OPENSIFT MONITORING STACK AND OTHER TOOLS FOR RED HAT CLOUD SERVICES DATA COLLECTION

The Red Hat Cloud Services portfolio includes managed services that rely on Red Hat infrastructure. Part of that infrastructure is the monitoring stack tools that, among other jobs, supply data about subscription usage to the subscriptions service. No additional user action is necessary to set up these data collection tools for the following managed services:

- Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription
- Red Hat Advanced Cluster Security for Kubernetes with a pay-as-you-go On-Demand subscription

5.6. CLOUD INTEGRATIONS FOR RED HAT SERVICES

The data collection for some pay-as-you-go On-Demand subscriptions requires a connection known as a cloud integration, configured with the integrations service of the Hybrid Cloud Console.

A *cloud integration* on the Red Hat Hybrid Cloud Console is a connection to a service, application, or provider that supplies data to another Hybrid Cloud Console service. Through a cloud integration, the connected service can connect with and use data from public cloud providers and other services or tools to collect data for that service.

The following products require the configuration of a cloud integration to enable data collection for the subscriptions service:

- Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on

Additional resources

- For additional information about registering version 4.1 disconnected clusters in Red Hat OpenShift Cluster Manager, see the chapter about cluster subscriptions and registration in the [Managing Clusters](#) guide.

CHAPTER 6. HOW TO SET SUBSCRIPTION ATTRIBUTES

Red Hat subscriptions combine technology with use cases to help procurement and technical teams make the best purchasing and deployment decisions for their business needs. When the same product is offered in two different subscriptions, these use cases differentiate between the options. They inform the decision-making process at the time of purchase and remain associated with the subscription throughout its life cycle to help determine how the subscription is used.

Red Hat provides a method for you to associate use case information with products through the application of subscription attributes. These subscription attributes can be supplied at product installation time or as an update to the product.

The subscriptions service helps you to align your software deployments with the use cases that support them and compare actual consumption to the capacity provided by the subscription profile of your account. Proper, automated maintenance of the subscription attributes for your inventory is important to the accuracy of the subscriptions service reporting.

Subscription attributes can generally be organized into the following use cases:

technical use case

Attributes that describe how the product will be used upon deployment. Examples include role information for RHEL used as a server or alternatively used as a workstation.

business use case

Attributes that describe how the product will be used in relation to your business environment and workflows. Examples include usage as part of a production environment or alternatively as part of a disaster recovery environment.

operational use case

Attributes that describe various operational characteristics such as how the product will be supported. Examples include a service level agreement (SLA) of premium, or a service type of L1-L3.

The subscription attributes might be configured from the operating system or its management tools, or they might be configured from settings within the product itself. Collectively, these subscription attributes might be known as system purpose, subscription settings, or similar names across all of these tools.

Subscription attributes are used by the Cloud Services platform tools such as the inventory tool to build the most accurate usage profile for products in your inventory. The subscriptions tool uses the subscription attributes found and reported by these other tools to filter data about your subscriptions, enabling you to view this data with more granularity. For example, filtering your RHEL subscriptions to show only those with an SLA of premium could help you determine the current usage of those premium subscriptions compared to your overall capacity for premium subscriptions.

The quality of subscription attribute data can greatly affect the accuracy and usefulness of the subscriptions service data. Therefore, a best practice is to ensure that these attributes are properly set, both for current use and any possible future expansion of subscription attribute use within the subscriptions service.

6.1. SETTING SUBSCRIPTION ATTRIBUTES FOR RHEL

You can set subscription attributes for the RHEL product from RHEL, Red Hat Subscription Management, or Satellite.

You should set the subscription attributes from only one tool. If you use multiple tools, there is a possibility for mismatched settings. Because these tools report data to the Cloud Services platform

tools at different intervals, or *heartbeats*, and because the subscriptions service shows its results as a once-per-day snapshot based on last-reported data, adding subscription attributes to more than one tool could potentially affect the quality of the subscriptions service data.

Setting the subscription attributes from RHEL

For RHEL 8 and later, you can use a few different methods to set subscription attributes. These methods, which include using the **syspurpose** command line tool, are described in a few different contexts in the RHEL 8 documentation. For more information, see the following links:

- See the section about completing post-installation tasks in the [Performing a Standard RHEL 8 Installation](#) guide.
- See the section about configuring system purpose in the [System Design Guide](#).



NOTE

The **syspurpose** command line tool has also been added to RHEL 7.7 and later.

Setting the subscription attributes from Red Hat Subscription Management

For Red Hat Subscription Management, the methods to set subscription attributes are contained in the section for registering a system and the descriptions of register commands, but are more fully described in the section related to using system purpose. For more information, see the following link:

- See the section about managing subscription usage in the [Using Red Hat Subscription Management](#) guide.

Setting the subscription attributes from Satellite

For Satellite, the methods to set subscription attributes are described in instructions for creating a host and editing the system purpose of a host. For more information, see the following link:

- See the section about administering hosts in the [Managing Hosts](#) guide.

6.2. SETTING SUBSCRIPTION ATTRIBUTES FOR RED HAT OPENSIFT

You can set subscription attributes from Red Hat OpenShift Cluster Manager for version 4. For version 3, you use the same reporting tools as those defined for RHEL.

Setting the subscription attributes for Red Hat OpenShift 4

You can set subscription attributes at the cluster level from Red Hat OpenShift Cluster Manager, where the attributes are described as subscription settings.

1. From the Clusters view, select a cluster to display the cluster details.
2. Click **Edit Subscription Settings** on the cluster details page or from the **Actions** menu.
3. Make any needed changes to the values for the subscription attributes and then save those changes.

Setting the subscription attributes for Red Hat OpenShift 3

You can set subscription attributes at the node level by using the same methods that you use for RHEL, setting these values from RHEL itself, Red Hat Subscription Management, or Satellite. As described in

that section, set subscription attributes by using only one method so that the settings are not duplicated.

If your subscription contains a mix of socket-based and core-based nodes, you can also set subscription attributes that identify this fact for each node. As you view your Red Hat OpenShift usage, you can use a filter to switch between cores and sockets as the unit of measurement.

To set this subscription attribute data, run the applicable command for each node:

- For core-based nodes:

```
# echo '{"ocm.units":"Cores/vCPU"}' | sudo tee /etc/rhsm/facts/openshift-units.facts
```

- For socket-based nodes:

```
# echo '{"ocm.units":"Sockets"}' | sudo tee /etc/rhsm/facts/openshift-units.facts
```

6.3. SETTING SUBSCRIPTION ATTRIBUTES FOR RED HAT CLOUD SERVICES

The current offerings for Red Hat Cloud Services services, for example, Red Hat OpenShift AI or Red Hat Advanced Cluster Security for Kubernetes, are of one subscription type only. Therefore, setting subscription attributes for these services is not required.

CHAPTER 7. YOUR RESPONSIBILITIES

The subscriptions service and the features that make up this service are new and are rapidly evolving. During this rapid development phase, you have the ability to view, and more importantly contribute to, the newest capabilities early in the process. Your feedback is valued and welcome. Work with your Red Hat account team, for example, your technical account manager (TAM) or customer success manager (CSM), to provide this feedback. You might also be asked to provide feedback or request features from within the subscriptions service itself.

As you use the subscriptions service, note the following agreements and contractual responsibilities that remain in effect:

- Customers are responsible for monitoring subscription utilization and complying with applicable subscription terms. The subscriptions service is a customer benefit to manage and view subscription utilization. Red Hat does not intend to create new billing events based on the subscriptions service tooling, rather the tooling will help the customer gain visibility into utilization so it can keep track of its environment.

PART III. SETTING UP THE SUBSCRIPTIONS SERVICE FOR DATA COLLECTION

To set up the environment for the subscriptions service data collection, connect your Red Hat Enterprise Linux and Red Hat OpenShift systems to the Cloud Services platform services through one or more data collection tools.

After you complete the steps to set up this environment, you can continue with the steps to activate and open the subscriptions service.

Do these steps

1. To gather Red Hat Enterprise Linux usage data, complete at least one of the following steps to connect your Red Hat Enterprise Linux systems to the Cloud Services platform by enabling a data collection tool. This connection enables subscription usage data to show in the subscriptions service.
 - a. Deploy Insights on every RHEL system that is managed by Red Hat Satellite:
 - [Deploying Red Hat Insights](#)
 - b. Ensure that Satellite is configured to manage your RHEL systems and install the Satellite inventory upload plugin:
 - [Installing the Satellite inventory upload plugin](#)
 - c. Ensure that Red Hat Subscription Management is configured to manage your RHEL systems:
 - [Registering systems to Red Hat Subscription Management](#)
 - d. For pay-as-you-go On-Demand subscriptions for metered RHEL, ensure that a cloud integration is configured in the Hybrid Cloud Console for collection of the metering data.
 - [Connecting cloud integrations to the subscriptions service](#)
2. To gather Red Hat OpenShift usage data, complete the following step for Red Hat OpenShift data collection on the Cloud Services platform.
 - a. Set up the connection between Red Hat OpenShift and the subscriptions service based upon the operating system that is used for clusters:
 - [Connecting Red Hat OpenShift to the subscriptions service](#)

CHAPTER 8. DEPLOYING RED HAT INSIGHTS

If you are using Red Hat Insights as the data collection tool, deploy Red Hat Insights on every RHEL system that is managed by Red Hat Satellite.

Do these steps

1. To install Red Hat Insights, see the following information:
 - [Installing Red Hat Insights](#)

Learn more


- To learn more about what data Red Hat Insights collects and your options for controlling that data, see the following information:
 - [What data does Red Hat Insights collect?](#)

8.1. INSTALLING RED HAT INSIGHTS

Install Red Hat Insights to collect information about your inventory.

Procedure

1. Install the Insights client on every RHEL system that is managed by Red Hat Satellite by using the following instructions:
 - [Red Hat Insights on Satellite Red Hat Enterprise Linux \(RHEL\)](#)



NOTE

The Insights client is installed by default on RHEL 8 and later systems unless the minimal installation option was used to install RHEL. However, the client must still be registered, as documented in the client installation instructions.

8.2. WHAT DATA DOES RED HAT INSIGHTS COLLECT?

When the Red Hat Insights client is installed on a system, it collects data about that system on a daily basis and sends it to the Red Hat Insights cloud application. The data might also be shared with other applications on the Cloud Services platform, such as inventory or subscription watch. Insights provides configuration and command options, including options for data obfuscation and data redaction, to manage that data.

For more information, see the [Client Configuration Guide for Red Hat Insights](#), available with the Red Hat Insights product documentation.

You might also want to examine the types of data that Insights collects and sends to Red Hat or add controls to the data that is sent. For additional information that supplements the information available in the product documentation, see the following articles:

- For more information about the use of the **insights-client --offline** command to generate an offline dump of the data before you register a system to Insights, see [How can I see what data is collected by Red Hat Insights?](#)

- For more information about the use of the **insights-client --no-upload** command to run a test data collection process, see [System Information Collected by Red Hat Insights](#) .
- For more information about the use of the **remove.conf** file and its options to exclude specific data from collection based on file, command, pattern, and keyword settings, see [Opting Out of Sending Metadata from Red Hat Insights Client](#).

CHAPTER 9. INSTALLING THE SATELLITE INVENTORY UPLOAD PLUGIN

When you are using Red Hat Satellite as a data collection tool for the subscriptions service, you must use the Satellite inventory upload plugin in combination with the `virt-who` tool for accurate reporting of hypervisor host-guest mapping information for Red Hat Enterprise Linux virtual data center (VDC) subscriptions and similar virtualized environments. The use of this plugin enables the uploading of host-based data from Satellite into the inventory service and enables the association of hosts with guests and with the Satellite instance that manages them. If the plugin is not enabled, the subscriptions service cannot accurately report usage for RHEL virtualized subscriptions.



NOTE

The plugin is also required by other Hybrid Cloud Console applications. In addition to the requirement for the inventory service of Red Hat Insights to enable the uploading of host inventory information, the plugin is also required for the remediations service of Red Hat Insights to enable remediation actions from both Satellite and Red Hat Insights.

After the Satellite inventory upload plugin is enabled, it is active for every Satellite organization, including existing and newly created organizations.

Prerequisites

Red Hat Satellite 6.9 or later (versions that are under full support)

Procedure

For a new installation of Satellite

For a new installation of Satellite, the Satellite inventory upload plugin is installed and enabled by default. No action is required to enable it.

For upgraded Satellite

For Satellite that is upgraded to a currently supported version, the Satellite inventory upload plugin is installed. However, you might have to enable the plugin.

- If the plugin was previously enabled for Satellite before the upgrade, it remains enabled. No action is required to enable it.
- If the plugin was not enabled for Satellite before the upgrade, then you must enable it.

To enable the Satellite inventory upload plugin, use these steps:

1. In the Satellite web interface, expand the **Configure** options and select **Red Hat Inventory**.
2. Follow the instructions on the Red Hat Inventory page to enable the **Automatic inventory upload** option for Satellite.

Usage tips

When the **Automatic inventory upload** option is enabled, the Satellite inventory upload plugin automatically reports once per day by default. You can also manually send data and view the status of the extract and upload actions for individual Satellite organizations.

The Satellite inventory upload plugin includes reporting settings that you can use to address data privacy concerns. Use the options on the Red Hat Inventory page to configure the plugin to exclude certain packages, obfuscate host names, and obfuscate host addresses.

CHAPTER 10. REGISTERING SYSTEMS TO RED HAT SUBSCRIPTION MANAGEMENT

If you are using Red Hat Subscription Management as the data collection tool, register your RHEL systems to Red Hat Subscription Management. Systems that are registered to Red Hat Subscription Management can be found and tracked by the subscriptions service.

Some RHEL images can use the autoregistration feature of the RHEL management bundle and do not have to be manually registered to Red Hat Subscription Management. However, the following specific requirements must be met:

- The image must be based on RHEL 8.4 and later or 8.3.1 and later.
- The image must be an Amazon Web Services (AWS) or Microsoft Azure cloud services image.
- The image can be a Cloud Access Gold Images image or a custom image, such as an image built with Image Builder. If it is a custom image, the **subscription-manager** tool in the image must be configured to use autoregistration.
- The image must be associated with an AWS or Azure integration, as configured from the **Settings > Integrations** option in the Hybrid Cloud Console, with the RHEL management bundle selected for activation.



NOTE

The integrations service was formerly known as the sources service in the Hybrid Cloud Console.

- The image must be provisioned after this integration is created.

RHEL systems that do not meet these requirements must be registered manually to be tracked by the subscriptions service.

Procedure

1. Register your RHEL systems to Red Hat Subscription Management, if not already registered. For more information about this process, see the following information:
 - Information in the Red Hat Subscription Management product documentation, including information about registering and unregistering systems in the [Quick Registration for RHEL](#) guide.
 - Supplemental information about registering systems in the [How do I register a system to Customer Portal Subscription Management?](#) article.

CHAPTER 11. CONNECTING RED HAT OPENSIFT TO THE SUBSCRIPTIONS SERVICE

If you use Red Hat OpenShift products, the steps you must do to connect the correct data collection tools to the subscriptions service depend on multiple factors. These factors include the installed version of Red Hat OpenShift Container Platform and Red Hat OpenShift Dedicated, whether you are working in a connected or disconnected environment, and whether you are using Red Hat Enterprise Linux, Red Hat Enterprise Linux CoreOS, or both as the operating system for clusters.

The subscriptions service is designed to work with customers who use Red Hat OpenShift in connected environments. One example of this customer profile is using RHOCP 4.1 and later with an Annual subscription with connected clusters. For this customer profile, Red Hat OpenShift has a robust set of tools that can perform the data collection. The connected clusters report data to Red Hat through Red Hat OpenShift Cluster Manager, Telemetry, and the other monitoring stack tools to supply information to the data pipeline for the subscriptions service.

Customers with disconnected RHOCP 4.1 and later environments can use Red Hat OpenShift as a data collection tool by manually creating each cluster in Red Hat OpenShift Cluster Manager.

Customers who use Red Hat OpenShift 3.11 can also use the subscriptions service. However, for Red Hat OpenShift version 3.11, the communication with the subscriptions service is enabled through other tools that supply the data pipeline, such as Insights, Satellite, or Red Hat Subscription Management.



NOTE

For customers who use Red Hat OpenShift Container Platform or Red Hat OpenShift Dedicated 4.7 and later with a pay-as-you-go On-Demand subscription (available for connected clusters only), data collection is done through the same tools as those used by Red Hat OpenShift Container Platform 4.1 and later with an Annual subscription.

Procedure

Complete the following steps, based on your version of Red Hat OpenShift Container Platform and the cluster operating system for worker nodes.

For Red Hat OpenShift Container Platform 4.1 or later with Red Hat Enterprise Linux CoreOS

For this profile, cluster architecture is optimized to report data to Red Hat OpenShift Cluster Manager through the Telemetry tool in the monitoring stack. Therefore, setup of the subscriptions service reporting is essentially confirming that this monitoring tool is active.

1. Make sure that all clusters are connected to Red Hat OpenShift Cluster Manager through the Telemetry monitoring component. If so, no additional configuration is needed. The subscriptions service is ready to track Red Hat OpenShift Container Platform usage and capacity.

For Red Hat OpenShift Container Platform 4.1 or later with a mixed environment with Red Hat Enterprise Linux CoreOS and Red Hat Enterprise Linux

For this profile, data gathering is affected by the change in the Red Hat OpenShift Container Platform reporting models between Red Hat OpenShift major versions 3 and 4. Version 3 relies upon RHEL to report RHEL cluster usage at the node level. This is still the reporting model used for version 4 RHEL nodes. However, the version 4 era reporting model reports Red Hat Enterprise Linux CoreOS usage at the cluster level through Red Hat OpenShift tools.

The tools that are used to gather this data are different. Therefore, the setup of the subscriptions service reporting is to confirm that both tool sets are configured correctly.

1. Make sure that all clusters are connected to Red Hat OpenShift Cluster Manager through the Red Hat OpenShift Container Platform Telemetry monitoring component.
2. Make sure that Red Hat Enterprise Linux nodes in all clusters are connected to at least one of the Red Hat Enterprise Linux data collection tools, Insights, Satellite, or Red Hat Subscription Management. For more information, see the instructions about connecting to each of these data collection tools in this guide.

For Red Hat OpenShift Container Platform version 3.11

Red Hat OpenShift Container Platform version 3.11 reports cluster usage based on the Red Hat Enterprise Linux nodes in the cluster. Therefore, for this profile, the subscriptions service reporting uses the standard Red Hat Enterprise Linux data collection tools.

1. Make sure that all Red Hat Enterprise Linux nodes in all clusters are connected to at least one of the Red Hat Enterprise Linux data collection tools, Insights, Satellite, or Red Hat Subscription Management. For more information, see the instructions about connecting to each of these data collection tools in this guide.

CHAPTER 12. CONNECTING CLOUD INTEGRATIONS TO THE SUBSCRIPTIONS SERVICE

Data collection for certain pay-as-you-go On-Demand subscriptions requires a connection known as a cloud integration, configured with the integrations service of the Hybrid Cloud Console.

A *cloud integration* on the Red Hat Hybrid Cloud Console is a connection to a service, application, or provider that supplies data to another Hybrid Cloud Console service. Through a cloud integration, the connected service can connect with and use data from public cloud providers and other services or tools to collect data for that service.

The following products require the configuration of a cloud integration to enable data collection for the subscriptions service:

- Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on

The configuration of a cloud integration for RHEL for Third Party Linux Migration with ELS includes a creating a connection between a cloud provider and the cost management service in the Hybrid Cloud Console. This cloud integration ensures that usage data from the cloud provider and from the cost management service is used to calculate metered usage in the subscriptions service and that usage data is returned to the cloud provider for billing purposes.

Procedure

For Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on

The post-purchase enablement steps for RHEL for Third Party Linux Migration with ELS include information about setting up the required cloud integration, in addition to other setup information required for the subscription. To ensure that your cloud integration is configured correctly for use by the subscriptions service, review the following information and confirm that the cloud integration configuration steps were completed:

- For more information about the RHEL for Third Party Linux Migration with ELS post-purchase enablement steps, including the step to set up a cloud integration, see the [Getting Started with Red Hat Enterprise Linux for Third Party Linux Migration](#) Customer Portal support article.
- For more information about cloud integrations, see [Configuring cloud integrations for Red Hat services](#)
- For more information about setting up a cloud integration for the cost management service and a specific cloud platform, see [Adding integrations to cost management](#) in the cost management documentation.

PART IV. ACTIVATING AND OPENING THE SUBSCRIPTIONS SERVICE

After you complete the steps to set up the environment for the subscriptions service, you can go to cloud.redhat.com to request the subscriptions service activation. After activation and the initial data collection cycle, you can open the subscriptions service and begin viewing usage data.

Do these steps

1. To find out if the subscriptions service activation is needed, see the following information:
 - [Determining whether manual activation of the subscriptions service is necessary](#)
2. To log in to cloud.redhat.com and activate the subscriptions service, see the following information:
 - [Activating the subscriptions service](#)
3. To log in to cloud.redhat.com and open the subscriptions service after activation, see the following information:
 - [Logging in to the subscriptions service](#)
4. If you cannot activate or log in to the subscriptions service, see the following information:
 - [Verifying access to the subscriptions service](#)

CHAPTER 13. DETERMINING WHETHER MANUAL ACTIVATION OF THE SUBSCRIPTIONS SERVICE IS NECESSARY

The subscriptions service must be activated to begin tracking usage for the Red Hat account for your organization. The activation process can be automatic or manual.

Procedure

Review the following tasks that activate the subscriptions service automatically. If someone in your organization has completed one or more of these tasks, manual activation of the subscriptions service is not needed.

- Purchasing a pay-as-you-go On-Demand subscription for Red Hat OpenShift Container Platform or Red Hat OpenShift Dedicated through Red Hat Marketplace. As the pay-as-you-go clusters begin reporting usage through OpenShift Cluster Manager and the monitoring stack, the subscriptions service activates automatically for the organization.
- Purchasing a Red Hat Cloud Services pay-as-you-go On-Demand subscription through a cloud provider marketplace, such as Red Hat Marketplace or Amazon Web Services (AWS). Examples of these types of products include Red Hat OpenShift AI or Red Hat Advanced Cluster Security for Kubernetes. As these products begin reporting usage through the monitoring stack, the subscriptions service activates automatically for the organization.
- Creating an Amazon Web Services integration through the integrations service in the Hybrid Cloud Console with the RHEL management bundle selected. The process of creating the integration also activates the subscriptions service.
- Creating a Microsoft Azure integration through the integrations service in the Hybrid Cloud Console with the RHEL management bundle selected. The process of creating the integration also activates the subscriptions service.



NOTE

The integrations service was formerly known as the sources service in the Hybrid Cloud Console.

These tasks, especially purchasing tasks, are frequently performed by a user that has the Organization Administrator (org admin) role in the Red Hat organization. The integration creation tasks must be performed by a user with the Sources administrator role in the role-based access control (RBAC) system for the Hybrid Cloud Console.

CHAPTER 14. ACTIVATING THE SUBSCRIPTIONS SERVICE

If the subscriptions service is not activated by one of the tasks that include automatic activation, then the subscriptions service must be manually activated. Tasks that include automatic activation are purchasing an On-Demand subscription through Red Hat Marketplace or creating an Amazon Web Services or Microsoft Azure integration that includes the RHEL management bundle through the integrations service in the Hybrid Cloud Console.

If manual activation is needed, the subscriptions service must be activated by a user with access to the Red Hat account and organization through a Red Hat Customer Portal login. This login is not required to be a Red Hat Customer Portal organization administrator (org admin). In addition, that user must also have the Subscriptions administrator role or the Subscriptions user role in the user access role-based access control (RBAC) system for cloud.redhat.com.



NOTE

If a Red Hat Customer Portal login is associated with an organization that does not have an account relationship with Red Hat, then the subscriptions service cannot be activated.

When the subscriptions service is activated, the Cloud Services platform tools begin analyzing and processing data from the data collection tools for display in the subscriptions service.



NOTE

The following procedure guides you through the steps to activate the subscriptions service from cloud.redhat.com. If the subscriptions service is not already activated, you can also access the activation page at the conclusion of the subscriptions service tour or from an option on the [Subscription Central](#) page.

Procedure

1. In a browser window, go to cloud.redhat.com.
2. If prompted, enter your Red Hat Customer Portal login credentials.
3. In the Hybrid Cloud Console navigation menu, click either **Red Hat Enterprise Linux** or **OpenShift**.
4. Expand **Subscriptions**. Then click one of the following options, depending on the product name that you clicked in the previous step.
 - For **Red Hat Enterprise Linux**, click **All RHEL**.
 - For **OpenShift**, click **Container Platform**
5. Complete one of the following steps, depending on the status of the subscriptions service activation:
 - If the subscriptions service is not yet active for the account, the activation page displays. Click **Activate Subscriptions**.
 - If the subscriptions service is activated but not yet ready to display data, the subscriptions service application opens, but it displays an empty graph. Try accessing the subscriptions service later, typically the next day.

- If the subscriptions service is activated and the initial data processing is complete, the subscriptions service application opens and displays data on the graph. You can begin using the subscriptions service to view data about subscription usage and capacity for the account.

Verification steps

Data processing for the initial display of the subscriptions service can take up to 24 hours. Until data for the account is ready, only an empty graph will display.

CHAPTER 15. LOGGING IN TO THE SUBSCRIPTIONS SERVICE

You access the subscriptions service from the Hybrid Cloud Console after logging in to your Red Hat Customer Portal login.

Procedure

1. In a browser window, go to cloud.redhat.com.
2. If prompted, enter your Red Hat Customer Portal login credentials.
3. In the Hybrid Cloud Console navigation menu, click either **Red Hat Enterprise Linux** or **OpenShift**.
4. Expand **Subscriptions**. Then click one of the following options, depending on the product name that you clicked in the previous step.
 - For **Red Hat Enterprise Linux**, click **All RHEL** or click one of the specific architectures to view more detailed information.
 - For **OpenShift**, click **Container Platform** or **Dedicated (On-Demand)**.
5. If the subscriptions service is activated and the initial data processing is complete, the subscriptions service opens and displays data on the graph. You can begin using the subscriptions service to view data about subscription usage and capacity for the account.



NOTE

If the subscriptions service opens but displays an empty graph, then the subscriptions service is activated but the initial data processing is not complete. Try accessing the subscriptions service later, typically the next day.

CHAPTER 16. VERIFYING ACCESS TO THE SUBSCRIPTIONS SERVICE

User access to cloud.redhat.com services, including the subscriptions service, is controlled through a role-based access control (RBAC) system. User management capabilities for this RBAC system are granted to the organization administrators (org admins) for an organization, as configured through access.redhat.com. Org admins then manage the cloud.redhat.com RBAC groups, roles, and permissions for the other members in the organization. This management can include the assignment of the User Access administrator role to additional members in the organization. The org admins and user access administrators can manage user access by using the **Settings > User access** option at cloud.redhat.com.

The predefined role **Subscriptions user** controls the ability to activate and access the subscriptions service. By default, every user in the organization has this role. However, if your org admin has made changes to user access roles and groups, you might not be able to access the subscriptions service.



NOTE

Beginning in September 2021, the RBAC roles for the subscriptions service are changed. The former Subscription Watch administrator role is renamed to the Subscriptions administrator role. This role contains every available permission for the subscriptions service. The Subscriptions user role, a new role with a subset of the permissions in the Subscriptions administrator role, now exists for users in the organization who do not require all the permissions for the subscriptions service. An example of this type of user is one who only needs to view report data.

After this change to the subscriptions service user access roles, then by default all users for organizations that activate the service and new users for organizations that currently use the service will be assigned the Subscriptions user role. However, the default behavior for role assignments is affected by how an organization is using RBAC groups to manage user access. If custom groups are in use instead of the Default access group, the org admin or another user with the User Access administrator RBAC role must manually update these groups to contain the new roles and manage any default assignment of them to the users in the organization.

Procedure

1. If you cannot activate or access the subscriptions service, contact your organization administrator. Your org admin can provide information about the status of the subscriptions service for your organization.

Additional resources

- For more information about cloud.redhat.com user access, see the [User Access Configuration Guide for Role-Based Access Control](#).

PART V. VIEWING AND UNDERSTANDING THE SUBSCRIPTIONS SERVICE DATA

After you set up the environment for the subscriptions service, such as setting up data collection tools or other data sources, completing any additional required subscriptions service activation steps, and waiting for the initial data ingestion, analysis, and processing to be complete (usually no longer than 24 hours), you can begin viewing subscription usage and capacity data in the subscriptions service.

Learn more

- To learn more about how the subscriptions service displays information about your subscription usage and capacity, see the following information:
 - [How does the subscriptions service show my subscription data?](#)
- To learn more about what data the subscriptions service stores, see the following information:
 - [What data does the subscriptions service store?](#)
- To learn more about how your data gets to the subscriptions service and how frequently this data refreshes, see the following information:
 - [How the subscriptions service gets and refreshes data](#)

CHAPTER 17. HOW DOES THE SUBSCRIPTIONS SERVICE SHOW MY SUBSCRIPTION DATA?

The subscriptions service shows subscription data for Red Hat offerings such as software products or product sets, organized by the Red Hat software portfolio options in the Hybrid Cloud Console. Currently, the subscriptions service shows data for the Red Hat Enterprise Linux, Red Hat OpenShift, and Red Hat Cloud Services (also referred to as Application and Data Services) software portfolios.



NOTE

Only a subset of Red Hat products in each portfolio, identified by their stock-keeping units (SKUs), are tracked by the subscriptions service. The subscriptions service maintains an explicit deny list within the source code for untracked products.

- For more information about the SKUs that are not tracked, see the [deny list source code](#) in GitHub.
- For more information about the purpose of the subscriptions service deny list, see the [What subscriptions \(SKUs\) are included in Subscription Usage?](#) article.

For each software portfolio, the **Subscriptions** menu shows options for navigating to the subscriptions service product pages for the available products within the selected portfolio. The **Subscriptions** menu might also contain options for viewing other subscription-related data or for functions that are not part of the subscriptions service.

Each product page for the subscriptions service offers multiple views. These views enable you to explore different aspects about your subscriptions for that product. When combined, the data from these views can help you recognize and mitigate problems or trends with excess subscription usage, organize subscription allocation across all of your resources, and improve decision-making for future purchasing and renewals.

For all of these activities, and for other questions about your subscription usage, the members of your Red Hat account team can provide expertise, guidance, and additional resources. Their assistance can add context to the account data that is reported in the subscriptions service and can help you understand and comply with your responsibilities as a customer. For more information, see [Your responsibilities](#).

17.1. HOW TO USE THE SUBSCRIPTION DATA IN THE VIEWS

The subscriptions service views can be grouped generally into the graph view and the table view.

The graph view is a visual representation of the subscription usage and capacity for your organization, where your organization is also a Red Hat account. This view helps you track usage trends and determine utilization, which is the percentage of deployed software when measured against your total subscriptions.

The table view can contain one or more tables that provide more details about the general data in the graph view. The current instances table, also referred to as the current systems table, provides details about subscription usage on individual components of your environment, for example, systems in your inventory or clusters in your cloud infrastructure or restricted network. The current subscriptions table provides details about individual subscriptions in your account. The table view helps you to find where Red Hat software is deployed in your environment, to understand how individual subscriptions contribute to your overall capacity for usage of similar types of subscriptions, to resolve questions you might have about subscription usage, and to refine plans for future deployments.



NOTE

For some product pages, the table view data is derived from data in the Cloud Services platform inventory service. User access to subscriptions, inventory, and other services is controlled independently by a role-based access control (RBAC) system for the Cloud Services platform tools, where individual users belong to groups and groups are associated with roles. More specifically, user access to the inventory service is controlled through the Inventory administrator role.

When the Inventory administrator RBAC role is enabled for the group or groups for your organization, information in the current instances table for the subscriptions service can display as links, where you can open a more detailed record in the inventory application for the listed systems or instances, as applicable. Otherwise, current instances table information displays as nonlinked information. For more information about RBAC usage in your organization, contact the organization administrator for your account.

The usage and utilization graph view

The graph view shows you your total subscription usage and capacity over time in a graph form. It provides perspective on your account's subscription threshold, current subscription utilization, and remaining subscription capacity, along with the historical trend of your software usage. The graph view might contain a single graph or multiple graphs, depending upon how subscription usage for a product is measured.

The usage and capacity calculations that appear in the graph are based on data snapshots that are provided periodically as the Hybrid Cloud Console processing tools analyze information from the various data collection tools and data sources. The data snapshots for Annual subscriptions generally update once every 24 hours. The data snapshots for On-Demand subscriptions can be more frequent, updating multiple times per day.

- *Usage* is the measurement of the consumption of Red Hat products installed on physical hardware or its equivalent. Usage is measured with a unit of measurement that is defined within the terms of a subscription.

Units of measurement differ according to the type of product and the type of subscription. The terms of Annual subscriptions determine usage as the physical hardware that is consumed, such as sockets or cores, or equivalent physical hardware that is consumed, such as a cloud platform instance that is equal to a socket. The terms of On-Demand subscriptions, such as pay-as-you-go subscriptions, can determine usage by a combination of metrics that measure consumed resources. One type of these metrics might be a compound unit, or derived unit. Examples of derived units can be a certain amount of physical hardware that is consumed during a specific period of time, such as core hours, or the availability of a Red Hat service instance, such as instance hours.

Usage is represented by a line or area graph, with different types of usage, for example, Red Hat Enterprise Linux physical, virtual, public cloud, and hypervisor usage, represented by different colors.

For Annual subscriptions, usage fluctuates over time as you install and uninstall the software contained in your subscriptions. For On-Demand subscriptions, usage fluctuates as you consume more or less of the resources that are measured by the terms of that subscription.

- *Capacity* is the upper limit of usage for a subscription, expressed in the unit of measurement and then summed for similar subscriptions across all of the contracts in your account. Similar subscriptions can be all products in a certain product portfolio, such as all RHEL subscriptions. The sum of capacity for all of your active subscriptions, the maximum capacity, is also known as the *subscription threshold*. This value is represented by a dashed line in the usage and utilization

graph for a product. Two primary reasons could prevent a subscription threshold from appearing in the graph. If a product page includes a subscription that is sold with unlimited capacity as part of its sales terms, the subscription threshold is not shown. Also, for On-Demand subscriptions or similar subscriptions that are billed for monthly usage, no capacity is set, so a subscription threshold is not shown. If filter selections remove unlimited subscriptions from a view, then the subscription threshold would appear for those filtered results.

The capacity of an individual subscription does not change over time. The subscription threshold fluctuates over time as new subscriptions are activated and old subscriptions expire, affecting the maximum capacity.

- *Utilization* is the percentage of the maximum capacity, as indicated by the subscription threshold, that is exhausted through the deployment and usage of Red Hat software in your account. In simple terms, utilization is the usage divided by the maximum capacity. If capacity is not applicable to a certain type of subscription present in the account, such as an unlimited subscription, utilization as a percentage of the maximum capacity also does not apply. Subscription utilization fluctuates over time due to the interaction of the changes to the usage and the subscription threshold.

Although the graph shows trends over a selected time interval, you can also view more specific information for the graph. For example, if the selected time interval is **Weekly**, you can hover over the graph near a date to see more specific data for a particular week.


You can also use the available filters, which can vary by product, to change the usage data that displays in the graph. For example, you can filter by the time interval, the unit of measurement, or by the subscription attribute filters such as service level agreement (SLA), as applicable. Additional filter options might include a filter by the type of usage measurement, such as physical or virtual, or a filter by variants for that product, such as a supporting architecture of x86.

The graph view: example graph

The following image shows an example RHEL usage and utilization graph in the subscriptions service. For other product pages, the graph view will contain differences in design, depending on how those products are sold and measured.

For the graph, the time filter is set to a daily view, and the graph displays a month of RHEL usage.

Figure 17.1. Usage and utilization graph example



1. A tooltip displays when you hover over a point in the graph. In this example, the tooltip displays

more information about the subscription usage and the subscription threshold for a specific day, January 6. For this day, physical RHEL is consuming 30 sockets, virtual RHEL is consuming 30 sockets, public cloud RHEL is consuming 30 sockets, and hypervisor RHEL is consuming 30 sockets, with a total of 120 sockets for all usage types. This usage total is less than the subscription threshold of 150 sockets.

2. The maximum capacity of RHEL usage, based on a unit of measurement of sockets, displays as the dashed subscription threshold line. This example shows an increase in the subscription threshold sometime between January 11 and January 16. The increase in the available capacity in this Red Hat account is due to the activation of additional RHEL subscriptions in the account.
3. The RHEL subscription usage, based on a unit of measurement of sockets, displays as four different colors for RHEL installed in physical, virtual, public cloud, and hypervisor environments. The example shows how all of these types of usage fluctuate over time. Usage fluctuates according to subscription activity, such as installation and uninstallation on physical systems or launch and termination of instances in the public cloud.

The table view: current instances table

The current instances table shows you details about usage on individual components in your environment, taken from the most recent daily snapshot of the usage data. This table provides information that can help you correlate the aggregated usage totals in the graph with the current software deployments on individual components across your organization. The components and data shown in the table vary by product because of the different ways that usage is tracked for products, by socket count, core count, core hours, and so on. Also, a component that displays as an "instance" or a "system" in the table can be a physical or virtual machine, or it can be another object such as a cluster or instance. Therefore, generic references to this table as either the current instances table or the current systems table are for convenience only.



NOTE

For some products such as RHEL, the data in the current instances table view contains some data that is available from the Hybrid Cloud Console inventory application, with the following differences:

- The inventory application shows significantly more system data. The current instances table view is a small subset of this data.
- Data in the inventory application can be more current because of the methods that are used to update the data. The current instances table view in subscriptions is based on a daily snapshot, so that data could be up to 24 hours old.
- Consumption of sockets or cores in the inventory application is represented as actual consumption. Usage in subscriptions is represented as normalized consumption, bound by the terms of subscription. For example, usage of a physical RHEL subscription is measured by socket pair, so a socket count for that type of system is always rounded to the next higher even number.

The information in the current instances table generally shows the name of the instance or system, the type of the system, the usage total for that system according to the unit of measurement, and the date that the system was last seen. However, the available columns in the table might differ according to the types of data that are relevant for that product. Columns in the table are sortable.

For the **Name** column that contains the name of the system, the *system* is the machine, either physical or virtualized, on which the product or product set is deployed. A system can also be a different

component, such as a Red Hat OpenShift cluster or an instance of a Red Hat Cloud Services service. The system is usually represented by either its display name or its universally unique ID (UUID). For multi-guest systems such as hypervisors, you can expand the system to see more information about individual guests. For some objects in the **Name** column, you can also click the system name to open the full system record in a different resource, for example, in the Hybrid Cloud Console inventory application.



NOTE

Currently for the display of Red Hat OpenShift Container Platform and Red Hat OpenShift Dedicated pay-as-you-go On-Demand subscription data, the **Name** column uses the inventory UUID. This ID is not the same as the cluster ID that is used for the cluster in Red Hat OpenShift Cluster Manager. In addition, the inventory UUID in the **Name** column does not provide a link to the cluster record in Red Hat OpenShift Cluster Manager. However, in both the subscriptions service and Red Hat OpenShift Cluster Manager you can use the available search filters to cross-reference these IDs.

For the **Guests** column that displays in the table when hypervisor usage can be tracked, the guest count is for the number of guests that are under the management of that hypervisor system. For other types of usage, the double dashed line represents a null value for that system.

For the **Type** column that contains the type of the system, the type is the *infrastructure type* on which the product or product set is deployed. A system can be a physical host, hypervisor, individual virtual machine, or other form of virtual deployment such as a public cloud instance. The information in this column might not be applicable to all products, so for some products the **Type** column might not appear.

For the column where the *usage total* for that system is displayed, the column label will vary according to how product usage is measured. For subscriptions where usage is measured with multiple metrics, multiple columns will display. The usage is the actual or equivalent amount of physical hardware that the product or product set is consuming on that system. Usage is counted according to the applicable *unit of measurement*, which in turn is determined by the terms of the subscription. For example, for a subscription that is sold by sockets, the usage total is the number of sockets, also known as subscribed sockets, that are consumed by a system. Other subscriptions such as On-Demand subscriptions are sold with different terms, such as by core hours, or might include multiple metrics in the terms, such as data transfer, data storage, and instance hours.



NOTE

The data for the usage total is based on the update, or heartbeat, cycles for the subscriptions service. For Annual subscriptions, the value that displays for the usage total is based on the 24-hour snapshot of usage for the most recently tallied day. For On-Demand subscriptions, the value is the most recently tallied data that is available to the subscriptions service, data that could be from the current day.

For the **Last seen** column that contains a date, that *last seen date* is the date that the system was last found by the Cloud Services platform tools, such as the inventory service or Red Hat OpenShift Cluster Manager and other tools in the monitoring stack. As part of the underlying tasks that subscriptions and other tools perform to calculate usage, the inventory service and the monitoring stack help to identify and deduplicate system data that is gathered by the various data collection tools.

As with the usage and utilization graph, you can use the filters to change the data that displays in the current instances table. However, a change to the time interval, such as changing from days to weeks, has no effect on the current instances table. The data displayed is from the most recent snapshot, so it is usually no more than 24 hours old.

You can also search the current instances table for a specific system name or a group of similarly named systems by using the search field. Exact and partial strings are accepted, but common wildcard characters are treated as literal characters, not special character wildcards.

The table view: current subscriptions table

The current subscriptions table shows you details about your currently active subscriptions, taken from the most recent daily snapshot of this data. This table contains information that can help you understand the maximum capacity for your usage of that product within your account. The maximum capacity is displayed as the subscription threshold in the usage and utilization graph view.

The table shows the capacity for each subscription in the unit of measurement by which that subscription is sold, for example, sockets or cores. The sum of the capacity for all rows equals the subscription threshold.

By using the data in the current subscriptions table, you can more fully understand how individual subscriptions are contributing to the subscription threshold. This information can help you plan for any future purchasing decisions, such as adjusting the amount of existing subscriptions or purchasing different subscriptions that are more suited to your usage profile. You can also use the information in the table to anticipate upcoming events that could affect your business activities in relation to purchasing and renewals, such as contract expiration.



NOTE

Currently, some On-Demand subscriptions such as Red Hat OpenShift Dedicated On-Demand are restricted to one subscription per account. Therefore, the current subscriptions table does not display for these types of products.

The information in the current subscriptions table generally shows the name of the product subscription, the service level agreement (SLA) for the subscription, the quantity of the subscription, the capacity of that subscription according to the unit of measurement, and the next renewal event for the subscription. All columns in the table are sortable.

The **Product** column lists unique product subscriptions that are currently active in your account. Future-dated subscriptions that are not yet active do not appear in the table. Expired subscriptions that are not renewed are removed from the table.

In general, subscriptions that share the same stock-keeping unit (SKU) appear on a single row. Subscriptions that can be grouped on the same row include these characteristics:

- Subscriptions with the same SKU, whether purchased in the same or different contracts or purchased at the same or different times.
- Subscriptions with the same SKU but with other minor differences to attributes, such as differences in quantity, that do not result in the creation of a new SKU.

However, some subscriptions might display multiple times in the Product column. These subscriptions include these characteristics:

- Subscriptions that have different SKUs but identical description text: The text that displays for a subscription is derived from the SKU description text. In some cases, this text might be identical for different SKUs. For example, two subscriptions could differ in one major attribute such as the SLA, resulting in a different SKU for the changed SLA.
- Subscriptions that have the same SKU but that are purchased through different marketplaces: Some Red Hat subscriptions are made available through multiple cloud provider marketplaces,

such as Red Hat Marketplace and AWS Marketplace. A subscription of this type would have a single SKU, even though it is available in multiple locations. In the current subscriptions table, these subscriptions would appear on different rows to help clarify which marketplace was used for the purchase.

The **Service level** column contains the service level agreement (SLA) for a subscription, as defined within the terms of the subscription. Examples include Premium, Standard, or Self-Support. This information can sometimes help you distinguish between two subscriptions in the Product column that have identical descriptions.

The **Quantity** column contains the number of active subscriptions for a SKU. For example, a single table row might contain multiples of the same SKU purchased in the same transaction. It might also contain multiples of the same SKU purchased in different transactions.

For the column where the *capacity* for a subscription is displayed, the column label will vary according to how product usage is measured. For example, RHEL is sold in socket pairs, so the capacity column for RHEL has the label **Sockets**. This capacity column measures the maximum amount of available usage for the subscriptions in each table row. Usage is counted according to the applicable *unit of measurement*, which in turn is determined by the terms of the subscription. When summed, the total for all rows in the table represents the maximum possible capacity of usage for all subscriptions of that product. This value is also the subscription threshold in the graph view.



NOTE

When a row includes a subscription that is sold with unlimited capacity, the capacity value for that row will show the infinity symbol to represent the unlimited capacity.

For some product pages, the capacity column might be replaced with a different column if capacity is not applicable. For example, for an On-Demand subscription, a **Subscription type** column might display the type of subscription, such as Annual or On-Demand.

The **Next renewal** column lists the next pending renewal event for any subscription that is included in that row.

17.2. MEASUREMENT OF USAGE AND CAPACITY

Currently, the subscriptions service tracks certain types of Red Hat Enterprise Linux, Red Hat OpenShift, and Red Hat Cloud Services products. The data that is displayed for usage and capacity varies by product.

Overall usage and capacity trends display on the usage and utilization graph. The information in the current instances table provides additional detail about the most recent day of data from the graph. Where applicable, the information in the current subscriptions table provides additional detail about your currently active subscriptions.

17.2.1. Measurement of usage and capacity for Red Hat Enterprise Linux

The measurement of RHEL usage is based on the type of subscription.

Red Hat Enterprise Linux with a traditional Annual subscription

For Red Hat Enterprise Linux with a traditional Annual subscription, measurement of usage is based on the consumption of sockets, according to the terms of your subscription.



NOTE

As of 6 September 2023, the subscriptions service has changed the number of reported system types in the usage data for the RHEL for x86 variant. The following system types are no longer reported in the graph and table data: RHEL Desktop, RHEL Workstation, and RHEL Compute Node. Instead, only the RHEL Server system type is reported in the graph and table for the RHEL for x86 variant. You can still use the API to access data for RHEL Workstation and RHEL Compute Node, but the data for RHEL Desktop has been removed. The change is applicable for future reporting, but not for historical data.

This reporting change is intended to streamline current and future workflows in the subscriptions service user interface by emphasizing the technical variants of the RHEL platform. Beginning with release 8, subscriptions for limited subsets of the subscriptions service content began using the same overall source repositories, eliminating technical differences between them. Because RHEL Server subscribers significantly outnumber RHEL Workstation and RHEL Compute Node subscribers, removing identical-source RHEL offerings that are not technical variants reduces the burden of viewing and understanding RHEL Server usage in the user interface.

With this change, you might notice that your overall usage numbers of RHEL for x86 are reduced in the user interface graph and table data, by an amount that previously represented the RHEL Desktop, RHEL Workstation, and RHEL Compute Node systems. You can use the **rhsm-subscriptions-api** API if you want to continue to monitor RHEL usage for the RHEL Workstation and RHEL Compute Node systems.

Use the following endpoints to view this system usage data. For additional information about these endpoints, see the associated API documentation:

- Use the `/instances/product/{product_id}` endpoint with the correct product ID for the `{product_id}` variable to see the RHEL Workstation or RHEL Compute Node instances. For more information, see the [API documentation](#).
- Use the `/tally/products/{product_id}/{metric_id}` endpoint with the correct product ID for the `{productID}` variable and a value of `sockets` for the `{metric_id}` variable to see the RHEL Workstation or RHEL Compute Node usage. For more information, see the [API documentation](#).

Usage: RHEL with a traditional Annual subscription

Usage is measured in CPU sockets. Data is aggregated for all supported architectures and can be filtered by those architectures, including the supported IBM and x86 architectures. You can view usage data that is specific to each architecture by selecting from the options in the **Variant** filter. You can view aggregated data for all architectures by clearing the filter options.

The usage data in the graph is divided into four sections, based on RHEL on physical systems, virtual systems, public cloud systems, or hypervisors.

Capacity: RHEL with a traditional Annual subscription

To measure capacity, the socket contribution of each RHEL subscription is totaled.

For some Red Hat products, RHEL is included with and is installed to support that product. This inclusion is sometimes also referred to as being "bundled with" another product. For example, RHEL is included with Red Hat Satellite and is tracked separately in a Satellite page view. Bundled RHEL is not tracked or counted against the total RHEL usage or capacity that is intended for production workloads or similar purposes.

For Red Hat Satellite, similar to RHEL, measurement of usage is based on the consumption of sockets, according to the terms of your subscription.

Usage: Satellite

Usage is measured in CPU sockets. Data is aggregated for all Satellite and can be filtered by those products, Satellite Server and Capsule Server. You can view usage data that is specific to each product by selecting from the options in the **Variant** filter. You can view aggregated data for all products by clearing the filter options.

Satellite usage is shown to help you identify the presence of RHEL that is installed in support of Satellite. An example of this type of subscription is a RHEL subscription with the Smart Management bundle. If you have known RHEL systems that are running this type of bundled RHEL, you can use the Satellite view to find those systems.

Capacity: Satellite

Capacity is not an applicable metric for Satellite. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

RHEL that is installed in support of Satellite is not counted with your overall RHEL capacity. Instead, that RHEL subscription is specifically intended to run Satellite and is not intended to be used for other workload or development purposes.

Red Hat Enterprise Linux with a pay-as-you-go On-Demand subscription

For Red Hat Enterprise Linux with a pay-as-you-go On-Demand subscription, measurement of usage is based on vCPU hours, based on the terms of your subscription.



NOTE


Currently, Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on is the only RHEL pay-as-you-go On-Demand subscription offering that is tracked by the subscriptions service.

Usage: RHEL with a pay-as-you-go On-Demand subscription

Usage of a pay-as-you-go On-Demand subscription of RHEL is measured with a single metric, in virtual CPU hours (vCPU hours). A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHEL pay-as-you-go On-Demand subscription usage, availability for computational activity is the availability of the RHEL instance over time. To obtain usage in vCPU hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation.

The vCPU hour based usage data for the RHEL instances is summed and then displays as daily usage in the usage and utilization graph. Because of the monthly billing cycle for an On-Demand subscription, the default time interval for the graph is one month, the current calendar month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand usage because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from.

Adjacent to the graph, daily and monthly totals display for the most current snapshot of usage.



NOTE

The vCPU hour usage data for the account and for individual instances that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. The usage values that are displayed in different locations in the interface might show slight discrepancies due to this rounding. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places, and is not from the displayed values.

Capacity: RHEL with a pay-as-you-go On-Demand subscription

Capacity is not an applicable metric for a pay-as-you-go On-Demand subscription. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

17.2.2. Measurement of usage and capacity for Red Hat OpenShift

For Red Hat OpenShift, measurement of usage is based upon the size of clusters. More specifically, the measurement is based on the subscribed cluster size. The unit of measurement that is used to measure the subscribed cluster size depends upon the subscription terms and type of subscription for the product.

The *subscribed cluster size* is the sum of the size of all the subscribed nodes, which are the nodes that process workloads. In the versions of Red Hat OpenShift where the facts about node types and node labels can be obtained, all noninfrastructure nodes plus master nodes that are schedulable are considered available for workload use. For each of the subscribed nodes, the kernel is queried for the number of sockets, the number of cores on each socket, and the number of threads supported by each core. Then the total number of threads is divided by the threads per core to determine the number of cores on the node (physical or virtual machine).



NOTE

For Red Hat OpenShift version 4.1 and later (including the 4.7 versions of Red Hat OpenShift Container Platform and OpenShift Dedicated for On-Demand subscriptions), the subscriptions service is able to use node type and node label data to find the subscribed nodes. In the aggregation of usage data based on cluster size for these versions of Red Hat OpenShift, the nonsubscribed node usage, such as control plane node usage, is ignored. However, for OpenShift Dedicated On-Demand, the control plane usage is tracked as instance hours, based upon the availability of clusters.

The subscriptions service is not able to make this same distinction for earlier versions of Red Hat OpenShift Container Platform, so data for subscribed and nonsubscribed nodes is displayed and counted. Analysis of cluster data indicates that approximately 15% of data displayed for earlier versions of Red Hat OpenShift Container Platform is nonsubscribed node overhead. Therefore, if your subscription profile includes Red Hat OpenShift Container Platform version 3, it is possible that you can exceed your Red Hat OpenShift subscription threshold by up to 15% but still be in compliance with your subscriptions.

For additional details about how the subscriptions service uses subscribed node and subscribed cluster size data, see the following information: [What does the subscriptions service track?](#)

For additional details about improvements to Red Hat OpenShift usage tracking in the subscriptions service, see the following information: [How do vCPUs, hyper-threading, and subscription structure affect the subscriptions service usage data?](#)

After the cluster size information is obtained, usage and capacity information is calculated according to the product and type of subscription. For more information, see the following descriptions of each product and subscription type.

Red Hat OpenShift Container Platform

Usage: Red Hat OpenShift Container Platform with an Annual subscription

Usage of an Annual subscription of Red Hat OpenShift Container Platform is measured in CPU cores or sockets. Data displays as an account-level view that is a sum of usage across active clusters.

Capacity: Red Hat OpenShift Container Platform with an Annual subscription

To measure capacity, the core or socket contribution (as applicable) of each subscription is added to a total for Annual subscriptions.

Usage: Red Hat OpenShift Container Platform with a pay-as-you-go On-Demand subscription

Usage of a pay-as-you-go On-Demand subscription of Red Hat OpenShift Container Platform is measured in core hours. A *core hour* is a unit of measurement for computational activity on one core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. To obtain usage in core hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation.

The core hour based usage data for all clusters is summed and then displays as daily usage in the usage and utilization graph. Because of the monthly billing cycle for an On-Demand subscription, the default time interval for the graph is one month, the current calendar month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand usage because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from.

A cumulative **core hours used** value also displays for the most recent snapshot of the usage for that month if there is accumulated usage to display.



NOTE

The core hour usage data for the account and for individual clusters that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. The usage values that are displayed in different locations in the interface might show slight discrepancies due to this rounding. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places, and is not from the displayed values.

Capacity: Red Hat OpenShift Container Platform with a pay-as-you-go On-Demand subscription

Capacity is not an applicable metric for a pay-as-you-go On-Demand subscription. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

Red Hat OpenShift Dedicated

Usage: Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription

Usage of a pay-as-you-go On-Demand subscription of Red Hat OpenShift Dedicated is measured with two units of measurement, core hours and instance hours. Therefore, the usage and utilization graph includes a dual y-axis, also known as a primary y-axis and secondary y-axis.

- A *core hour* is a unit of measurement for computational activity on one core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For Red Hat OpenShift Dedicated On-Demand, core hours measure the workload

usage on the compute machines.

- An *instance hour* is a unit of measurement for the availability of a Red Hat service instance, during which it can accept and execute customer workloads. For Red Hat OpenShift Dedicated On-Demand, instance hours use your cluster availability data to measure the control plane usage on the control plane machines (in older versions of Red Hat OpenShift, the master machines). This data is used to calculate the control plane cost, also known as the cluster fee, that is included in your Red Hat Marketplace invoice.

To obtain usage in core hours and instance hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation. This process samples usage multiple times per hour, normalizes the samples for a specific time interval, aggregates the normalized samples into a daily total, and then sums each day into a total that is determined by the billing terms of the subscription.

The usage data for all clusters is summed and displayed as daily usage in the usage and utilization graph. The core hour usage is plotted with the primary y-axis, and the instance hour usage is plotted with the secondary y-axis. Because of the monthly billing cycle for an On-Demand subscription, the default time interval for the graph is one month, the current month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand usage because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from.

A cumulative **core hours used** value also displays for the most recent snapshot of the usage for that month if there is accumulated usage to display.



NOTE

The core hour and instance hour usage data for the account and for individual clusters that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. The usage values that are displayed in different locations in the interface might show slight discrepancies due to this rounding. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places, and is not from the displayed values.

Capacity: Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription

Capacity is not an applicable metric for a pay-as-you-go On-Demand subscription. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

Red Hat OpenShift add-ons

Usage: Red Hat OpenShift Service on AWS Hosted Control Planes with a pre-paid plus On-Demand subscription

Usage of a Red Hat OpenShift Service on AWS Hosted Control Planes (ROSA Hosted Control Planes) with a pre-paid plus On-Demand subscription is measured with two metrics, vCPU hour usage and control plane hour usage.

- A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms) for a total of one hour, measured to the granularity of the meter that is used. For ROSA Hosted Control Planes, availability for computational activity is the availability of the vCPUs for the ROSA Hosted Control Planes subscribed clusters over time. A subscribed cluster is comprised of subscribed nodes, which are the noninfrastructure nodes plus schedulable master nodes that are available for workload use, if

applicable. Note that for ROSA Hosted Control Planes, schedulable master nodes are not applicable, unlike other products that also use this measurement. The vCPUs that are available to run the workloads for a subscribed cluster contribute to the vCPU hour count.

- A *control plane hour* is a measurement of the availability of the control plane that is hosted in the Red Hat account. With ROSA Hosted Control Planes, each cluster has a dedicated control plane that is isolated in a ROSA Hosted Control Planes service account that is owned by Red Hat.

To obtain usage in vCPU hours and control plane hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation. The vCPU hour subscribed cluster usage data is summed and the control plane hour based usage data is summed and then usage data for both displays as daily usage in usage and utilization graphs. Because of the monthly billing cycle for an On-Demand subscriptions, the default time interval for the graph is one month, the current calendar month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand subscriptions usage, but the actual billing cycle from AWS can vary due to the effective date of the contract.

The ROSA Hosted Control Planes usage is measured and compared against the pre-paid portion of your contract with AWS. Usage that does not exceed the pre-paid portion of your contract is recorded, but is not sent to AWS. When usage exceeds the pre-paid portion of your contract, the overage amount is sent to AWS for billing purposes.

The terms of the ROSA Hosted Control Planes contract enable users to increase the amount of pre-paid usage at any time during the contract, including during the middle of a month. If a contract is adjusted mid-month, the subscriptions service can track the contract adjustment.

If you exceed your pre-paid contract usage and begin accumulating pay-as-you-go usage during a calendar month, and you also increase your pre-paid contract usage within that same month, the previously accumulated pay-as-you-go usage does not revert to pre-paid contract usage. However, new pay-as-you-go usage does not accumulate until your usage exceeds the new pre-paid contract usage maximum plus the previously accumulated pay-as-you-go usage. This principle ensures that pay-as-you-go usage is not double-counted.

For example:

- You have a subscription for 100 units of pre-paid contract usage.
- You use 110 units before the calendar month expires. Of those units, 100 units are counted as pre-paid usage and 10 units are counted as pay-as-you-go usage.
- You increase your subscription to 200 units of pre-paid contract usage, and the change takes effect within that same month.
- Additional pay-as-you-go usage is counted only after you exceed 210 units for that month.

Capacity: Red Hat OpenShift Service on AWS Hosted Control Planes with a pre-paid plus On-Demand subscription

Capacity is determined by the number of vCPU hours that are purchased, as shown in the terms of the ROSA Hosted Control Planes contract.

17.2.3. Measurement of usage and capacity for Red Hat Cloud Services

For Red Hat Cloud Services, measurement of usage is based on metrics that generally relate to the consumption of computing resources by the platform that powers the service. These resources might include, but are not limited to, metrics concerning CPU, vCPU, RAM, network traffic, storage volume,

and control plane consumption during the availability of each instance of a service. Because these services perform different jobs and consume different resources, an individual service might be measured by a single metric or a combination of these metrics. In addition, those differences in the services can result in different units of measurement being used for the basic metric types.



NOTE

In the subscriptions service, a Red Hat Cloud Services offering that is purchased multiple times through multiple cloud provider marketplaces will appear as grouped on a single page. Use the filtering function to limit the usage data to an offering that is purchased through a specific cloud provider.

Red Hat OpenShift AI

Usage: Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription

Usage of a pay-as-you-go On-Demand subscription of Red Hat OpenShift AI (RHOAI) is measured with a single metric, in vCPU hours. A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHOAI pay-as-you-go On-Demand subscription usage, availability for computational activity is the availability of the RHOAI cluster over time. To obtain usage in vCPU hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation.

The vCPU hour based usage data for the RHOAI cluster is summed and then displays as daily usage in the usage and utilization graph. Because of the monthly billing cycle for an On-Demand subscription, the default time interval for the graph is one month, the current calendar month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand usage because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from.

Adjacent to the graph, daily and monthly totals display for the most current snapshot of usage.



NOTE

The vCPU hour usage data for the account and for individual clusters that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. The usage values that are displayed in different locations in the interface might show slight discrepancies due to this rounding. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places, and is not from the displayed values.

Capacity: Red Hat OpenShift AI with a pay-as-you-go On-Demand subscription

Capacity is not an applicable metric for a pay-as-you-go On-Demand subscription. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

Red Hat Advanced Cluster Security for Kubernetes

Usage: Red Hat Advanced Cluster Security for Kubernetes (RHACS) with a pay-as-you-go On-Demand subscription

Usage of a pay-as-you-go On-Demand subscription of RHACS is measured with a single metric, in vCPU hours. A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHACS pay-as-you-go On-Demand subscription usage, availability for

computational activity is the availability of the RHACS cluster over time. To obtain usage in vCPU hours, the subscriptions service uses numerical integration, also commonly known as an "area under the curve" calculation.

The vCPU hour based usage data for all clusters managed by RHACS is summed and then displays as daily usage in the usage and utilization graph. Because of the monthly billing cycle for an On-Demand subscription, the default time interval for the graph is one month, the current calendar month. The subscriptions service uses the calendar month interval as a standard for reporting On-Demand usage because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from.

Adjacent to the graph, daily and monthly totals display for the most current snapshot of usage.



NOTE

The vCPU hour usage data for the account and for individual clusters that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. The usage values that are displayed in different locations in the interface might show slight discrepancies due to this rounding. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places, and is not from the displayed values.

Capacity: Red Hat Advanced Cluster Security for Kubernetes (RHACS) with a pay-as-you-go On-Demand subscription

Capacity is not an applicable metric for a pay-as-you-go On-Demand subscription. So capacity is not tracked, nor is a subscription threshold line shown, for this type of subscription.

17.3. UNITS OF MEASUREMENT

The unit of measurement by which product usage is tracked is determined by the terms of the subscription.

17.3.1. Units of measurement for Red Hat Enterprise Linux

Red Hat Enterprise Linux with a traditional Annual subscription

In general, the subscriptions service measures your RHEL usage by sockets for traditional Annual subscriptions. However, because of the inherent differences between physical, virtual, public cloud, and hypervisor offerings and their relation to hardware, the subscriptions service tracking uses different units of measurement, as follows:

Physical usage

The subscriptions service measures your physical RHEL installations by CPU socket pairs. Each system contributes its installed socket count, rounded upwards to the next even number. The value that displays is the total socket count after all of the system-level socket-pair rounding is applied. In the current instances table, on-premise physical hardware displays as physical machines.

Virtual usage

The subscriptions service measures your virtualized RHEL installations by individual sockets, where one socket represents one virtual machine. The systems that are counted with virtual usage are stand-alone systems. In other words, they are virtual machines with no detectable hypervisor management, including no virt-who host-guest mapping data that is sent to the subscriptions service either through the Satellite inventory upload plugin or through Red Hat Subscription Management.

Public cloud usage

The subscriptions service measures public cloud RHEL installations by socket.

The instances launched from public cloud RHEL images are recognized through Desktop Management Interfaces (DMI) fact-value pairs that are present in the image and instance metadata. The values of the DMI facts identify an instance as running in the cloud infrastructure provided by Amazon Web Services (AWS), Microsoft Azure, Google Cloud, and Alibaba Cloud. Other system profile facts are obtained from the cloud provider identity document (for example, the AWS instance identity document) and are combined with the DMI facts during the analysis of the instance. Each running instance is counted as running for the entire day and contributes a single socket to the socket count.

The cloud provider identity document includes system profile facts that help identify where an image was purchased. If an image is identified as being purchased directly from a cloud provider marketplace, the instances for that image do not contribute to the public cloud socket totals, do not count against your subscription threshold, and do not appear in the usage and utilization graph view. However, the systems that contain these types of RHEL instances do appear in the current instances table. In the table, you can identify the systems with cloud provider marketplace instances by a null value in the **Sockets** column, represented by double dashes (--). This null value indicates that these systems do not contribute usage data.

Hypervisor usage

The subscriptions service measures hypervisor RHEL installations by CPU socket pairs. A system that is providing virt-who host-guest mapping data to the subscriptions service either through the Satellite inventory upload plugin or through Red Hat Subscription Management is classified as a hypervisor. Systems that are classified as hypervisors can have multi-guest virtual data center (VDC) subscriptions or have similar virtualized environments. Hypervisor usage is counted in the following ways:

- For a RHEL based hypervisor with RHEL guests, the socket count of the hypervisor is counted twice, with the socket-pair method. One count is for the node's own copy of RHEL that is used as the operating system to run the hypervisor. A second count is for the RHEL that is used by the guest systems.
- For a non RHEL based hypervisor with RHEL guests, the socket count of the hypervisor is counted once, with the socket-pair method. The count is for the RHEL that is used by the guest systems.

Red Hat Enterprise Linux with a pay-as-you-go On-Demand subscription

The subscriptions service measures your RHEL usage in vCPU hours. A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used. For RHEL pay-as-you-go On-Demand subscription usage, availability for computational activity is the availability of the RHEL instance over time.



NOTE

Currently, Red Hat Enterprise Linux for Third Party Linux Migration with Extended Life Cycle Support Add-on is the only RHEL pay-as-you-go On-Demand subscription that is tracked by the subscriptions service.

17.3.2. Units of measurement for Red Hat OpenShift

Red Hat OpenShift Container Platform with an Annual subscription

The subscriptions service measures your Red Hat OpenShift usage in units of CPU cores or CPU sockets. For Red Hat OpenShift 4, the counting is aggregated at the cluster level, and for Red Hat OpenShift 3, the counting is aggregated at the node level. Currently, the subscriptions service cannot display a single, mixed-unit view of Red Hat OpenShift usage in environments that include core-based and socket-based clusters within the same account. You must use filtering to view that data in separate views.

You can use a filter to toggle the usage and capacity data between the two units of measurement. If subscription attributes are set on the cluster (through Red Hat OpenShift Cluster Manager for Red Hat OpenShift 4) or on the node (through the command to set the **ocm.units** value for Red Hat OpenShift 3), then that data can be reported by cores or sockets. If subscription attributes are not set or cannot be set, then the data is included in reports for both core-based and socket-based usage.

Physical usage

The subscriptions service measures your core-based physical Red Hat OpenShift installations by actual core count. Socket-based physical installations are measured by socket pairs, so the count is rounded upwards to the next even number.

In the current instances table, an example of a physical system for Red Hat OpenShift is a Red Hat OpenShift cluster running on bare metal. Another example is a RHEL system reporting as a Red Hat OpenShift 3 cluster node.

Virtual usage

The subscriptions service measures your core-based and socket-based installations by actual core and actual socket count.

In the current instances table, an example of a virtual system for Red Hat OpenShift is a cluster installed in environments such as Red Hat OpenStack Platform, Red Hat Virtualization, VMware vSphere, or on public cloud.

Red Hat OpenShift Container Platform and Red Hat OpenShift Dedicated with a pay-as-you-go On-Demand subscription

The subscriptions service measures your pay-as-you-go On-Demand subscription of Red Hat OpenShift Container Platform or Red Hat OpenShift Dedicated usage in core hours. A *core hour* is a unit of measurement for computational activity on one core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.

Physical usage

The subscriptions service measures your core-based physical Red Hat OpenShift installations by actual core count. Socket-based physical installations are measured by socket pairs, so the count is rounded upwards to the next even number.

Virtual usage

The subscriptions service measures your core-based and vCPU-based virtual installations by actual core count, with vCPUs rationalized to cores using maximum efficiency. Socket-based virtual installations are measured by socket count as reported by your hypervisor. For best reporting, confirm that your hypervisor is reporting accurate socket counts for your virtual machines.

Control plane usage

For Red Hat OpenShift Dedicated On-Demand only, the subscriptions service also measures your cluster availability by instance hour. For Red Hat OpenShift Dedicated On-Demand, this instance hour calculation of control plane usage is based on a cluster hour unit of measurement.

Red Hat OpenShift Service on AWS Hosted Control Planes with a pre-paid plus On-Demand subscription

The subscriptions service measures your pre-paid plus On-Demand subscription of Red Hat OpenShift Service on AWS Hosted Control Planes (ROSA Hosted Control Planes) usage in vCPU hours and in control plane hours.

- A *vCPU hour* is a measurement of availability for computational activity on one virtual core (as defined by the subscription terms) for a total of one hour, measured to the granularity of the meter that is used.
- A control plane hour is a measurement of the availability of the control plane. With ROSA Hosted Control Planes, each cluster has a dedicated control plane that is isolated in a ROSA Hosted Control Planes service account that is owned by Red Hat.

17.3.3. Units of measurement for Red Hat Cloud Services

Red Hat OpenShift AI

The subscriptions service measures your pay-as-you-go On-Demand subscription of Red Hat OpenShift AI (RHOAI) usage in vCPU hours. A *vCPU hour* is a unit of measurement for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.

Red Hat Advanced Cluster Security for Kubernetes

Red Hat Advanced Cluster Security for Kubernetes (RHACS) with a pay-as-you-go On-Demand subscription

The subscriptions service measures your pay-as-you-go On-Demand subscription of RHACS usage in vCPU hours. A *vCPU hour* is a unit of measurement for computational activity on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.

17.4. FILTERING

You can further refine the subscriptions service data by selecting values from the available filters in the interface. When you select a filter option, the graph view (and in some cases, the tables in the table view) generally refreshes to show data that relates to that option. In other words, most of the filters are inclusive, not exclusive, for the selected option.

Filtering by time

For Annual subscriptions, you can filter data by several different time intervals, including daily (the default) and quarterly. For On-Demand subscriptions, you can filter by time interval only for the current month or by any other month in the previous 12 months. Because billing cycles for On-Demand subscriptions might vary across the different cloud provider marketplaces that you might purchase from, the calendar month interval is used in the subscriptions service to report On-Demand usage.

Filtering by time affects only the usage and utilization graph view. The current instances table and current subscriptions table views always show the most recent data that is available, whether from the most recent subscriptions service daily snapshot or from more frequent reporting offered by the monitoring stack tools, and is not affected by the time filter.

**NOTE**

During the rapid development of the subscriptions service, the addition of new features is improving the scope and accuracy of this tool. The subscriptions service does not provide in-application capability to recalculate older usage and capacity data as these new features are being added. Therefore, the selection of a longer time interval could display results that contain inconsistencies.

Filtering by subscription attributes

You can use several filters to filter by subscription attributes, which is data that describes the characteristics and intended usage of subscription. The accuracy of those filters is dependent upon how accurately the subscription attribute data is set.

Subscription attributes might be configured from the operating system or its management tools, or from settings within the product itself. In these various tools, subscription attribute data is also known as system purpose, subscription settings, or similar names. In some cases, subscription attribute values might be derived from the subscription, such as when a subscription is sold either by sockets or cores.

You can use the subscriptions service filters to get a more focused view on usage that meets certain use cases within your subscription profile. For example, filtering your RHEL subscriptions by the **SLA** filter, which is the service level agreement, to show only those with an SLA of **Premium** could help you determine the current usage of premium subscriptions compared to your overall capacity for those premium subscriptions. In turn, this knowledge can inform decisions such as additional deployments, actions to mitigate subscription compliance issues, or future purchasing and renewals.

As another example, selecting a nonspecific value for a filter, such as the **No SLA** or **Unspecified** options, can help show subscriptions that have subscription attribute values that might be missing or that might be less common and not specifically filterable by the subscriptions service. For those subscriptions with missing subscription attributes, adding that data can improve the accuracy and usefulness of the subscriptions service reporting.

The subscriptions service provides the following filters and filter options for RHEL:

- **SLA** (service level agreement): Premium, Standard, Self-Support, No SLA
- **Usage**: Development/Test, Disaster Recovery, Production, Unspecified

The subscriptions service provides the following filter and filter options for Red Hat OpenShift:

- **SLA** (service level agreement): Premium, Standard, Self-Support, No SLA
- **Cores**: Cores (default), Sockets

If an offering is available as one subscription type only, such as with Red Hat OpenShift AI, filtering by subscription attributes is not available.

Filtering by type

You can use the **Type** filter to filter by the type of subscription usage when multiple types of subscription usage are tracked for a product. For example, the subscriptions service tracks RHEL physical, virtual, public cloud, and hypervisor usage to match the different infrastructure types for RHEL product deployment. This filter is different from the Usage filter that corresponds to subscription attribute (system purpose) usage.

**NOTE**

The **Type** filter replaces previously available filtering of the graph display by the values in the legend of the usage and utilization graph.

Filtering by variant

You can use the **Variant** filter to filter on various aspects of a subscription that are appropriate for the product page that is displayed. For example, you can filter on architecture variants for RHEL or product variants for Satellite.

Filtering by name (current instances table)

You can filter the data in the current instances table by the contents of **Name** column, which shows either the display name or universally unique ID (UUID) of each system. To filter by name, use the search field near the **Name** column.

You can search for a specific system name or a group of similarly named systems. Exact and partial strings are accepted, but common wildcard characters are treated as literal characters, not special character wildcards.

CHAPTER 18. WHAT DATA DOES THE SUBSCRIPTIONS SERVICE STORE?

The subscriptions service gathers data to track usage by using several tools, including Red Hat Insights, Red Hat Subscription Management, Red Hat Satellite and the Satellite inventory upload plugin, OpenShift Cluster Manager, the Red Hat OpenShift monitoring stack, and others. The number of tools that are helping to gather data for your account depends on your subscription profile and the products in it, because different tools are used to gather data for different products.

For more information about the data that is gathered and stored by Red Hat Insights, Red Hat Subscription Management, or other products, see that product's documentation.

The subscriptions service uses the data that is gathered in three ways:

- To make sure that inventory is counted only once. Some data is used for deduplication, in both primary and secondary storage.
- To link data submissions to the proper account and to log how and from what resource the data was received. Some quality control data is included.
- To calculate subscription values. Some data indicates the presence of Red Hat software and powers the usage portion of the subscriptions service.

The subscriptions service itself stores only a subset of the data that is collected by Red Hat Insights. The primary data that is stored by the subscriptions service includes information related to installed Red Hat products, system size, and other similar system characteristics.

Additional resources

- For more information about the types of data that the subscriptions service stores and uses, see the following article: [Data Gathered and Used by Subscription Watch](#)
- For more information about Insights data governance and security, see the following article: [Insights 102](#)

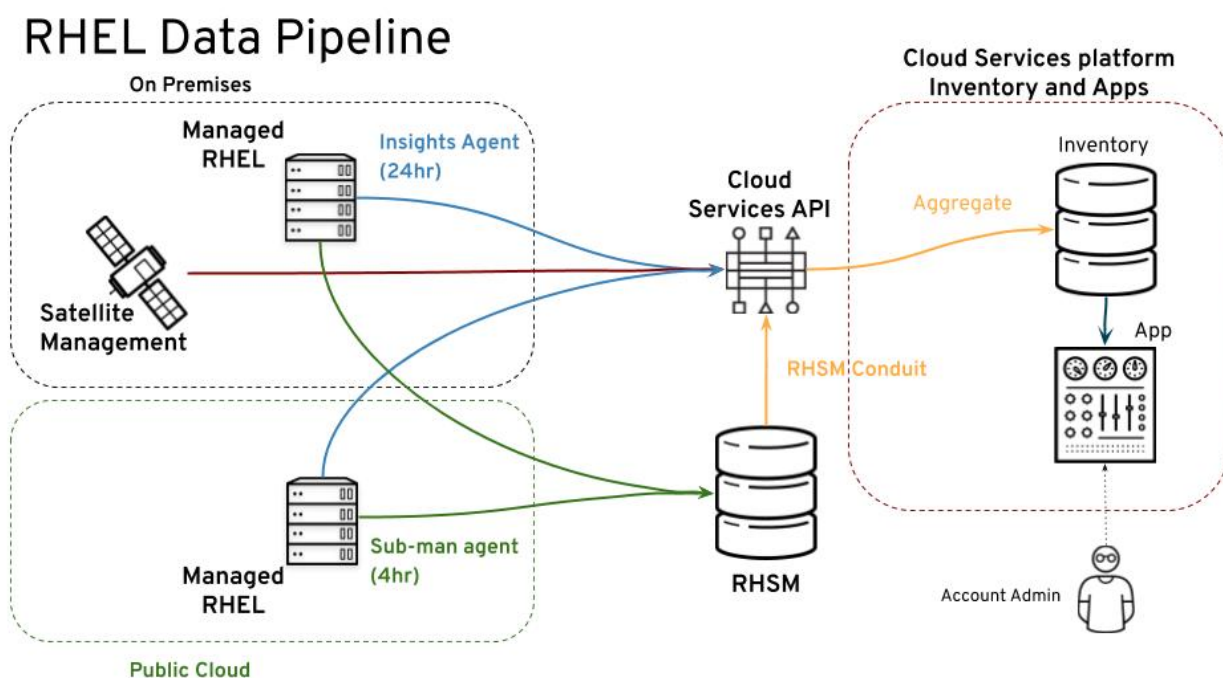
CHAPTER 19. HOW THE SUBSCRIPTIONS SERVICE GETS AND REFRESHES DATA

The data collection tools gather and periodically send data, including data about subscription usage, to the Hybrid Cloud Console tools that analyze and process this data. After the data is processed, the data that is needed for the subscriptions service, including the data related to subscription usage and capacity, is sent to the subscriptions service for display. For Annual subscriptions, this data is sent once per day. For On-Demand subscriptions, this data can be updated more frequently, usually a few times per day. Therefore, the data that displays in the subscriptions service is a tally of the results in the form of a snapshot, either once per day or at a few intervals throughout the day, and is not a real-time, continuous usage monitor.

The Red Hat Enterprise Linux data pipeline

The following image provides additional detail about the data pipeline that moves RHEL data from collection to display in the subscriptions service. The data collection tool, whether you are using Red Hat Insights, Satellite, or Red Hat Subscription Management with the Subscription Manager agent, sends data to the Hybrid Cloud Console processing tools. After data is processed, it is available to Hybrid Cloud Console tools such as the inventory service. The subscriptions service uses a subset of the data that is available to the inventory service to display data about subscription usage and capacity.

Figure 19.1. The RHEL data pipeline for the subscriptions service



The Red Hat OpenShift data pipeline

Red Hat OpenShift can have nodes that are based on Red Hat Enterprise Linux or Red Hat Enterprise Linux CoreOS. Only nodes that are based on RHCOS report data through the tools in the Red Hat OpenShift data pipeline, such as OpenShift Cluster Manager and the monitoring stack. RHEL nodes report through the tools in the RHEL data pipeline, such as Red Hat Insights, Satellite, or Red Hat Subscription Management.

Table 19.1. Node reporting and the data pipelines


Red Hat OpenShift version	Node operating system	Data pipeline used
Version 4	RHCOS	Red Hat OpenShift pipeline Nodes aggregated into cluster reporting Node types and node labels analyzed to determine subscribed nodes
Version 4	RHEL	Red Hat OpenShift pipeline Nodes report individually Node types and node labels analyzed to determine subscribed nodes
Version 3	RHEL	RHEL pipeline Nodes report individually Subscribed nodes cannot be distinguished from nonsubscribed nodes

For Red Hat OpenShift version 4.1 and later data collection, the tools available in the monitoring stack, including Telemetry, Prometheus, Thanos, and others, monitor and periodically sum the CPU activity of all subscribed nodes, while ignoring the activity of nonsubscribed nodes. That data is sent to Red Hat OpenShift Cluster Manager at different intervals for new clusters, resized clusters, and clusters with deleted entities, to maintain currency.

Red Hat OpenShift Cluster Manager then updates the cluster size attribute for existing clusters and creates entries for any new clusters in the Hybrid Cloud Console inventory tool.

Lastly, the subscriptions service analyzes the inventory data and creates account-wide usage information for each Red Hat OpenShift product in the subscription profile. That information is displayed in the the subscriptions service interface, along with capacity data as applicable for the subscription type. For Red Hat OpenShift Container Platform with an Annual subscription, the usage information accounts for both core and socket usage. For Red Hat OpenShift Container Platform or OpenShift Dedicated with an On-Demand subscription, the usage information shows core hour usage.

Figure 19.2. The Red Hat OpenShift data pipeline for the subscriptions service



The Red Hat Cloud Services data pipeline

The managed services in the Red Hat Cloud Services portfolio, such as Red Hat OpenShift AI or Red Hat Advanced Cluster Security for Kubernetes, rely on Red Hat infrastructure. Part of that infrastructure is the monitoring stack tools that, among other jobs, supply data about subscription usage to the subscriptions service. Therefore, the Red Hat Cloud Services services report usage through the tools used in the Red Hat OpenShift data pipeline.

Heartbeats for data collection tools

The frequency at which the data collection tools send data for processing, also known as the *heartbeat*, varies by tool. This variance can affect the freshness of the data that the subscriptions service displays.

The following table shows default heartbeats for the data collection tools. In some cases, these values are configurable within that data collection tool.

Table 19.2. Heartbeats for data collection tools

Tool	Configurable	Heartbeat interval
Insights	No	Daily, once every 24 hours
Red Hat Subscription Management	Yes	Multiple times per day, 4 hour default

Tool	Configurable	Heartbeat interval
Satellite	Yes	<p>Monthly, configurable with the Satellite scheduler function</p> <p>If used, the Satellite inventory upload plugin reports daily, with a manual send option.</p> <p>Additionally, to maintain accurate information about the mapping of virtual guests to hosts, a best practice is to run the virt-who utility daily.</p>
Red Hat OpenShift	No	<p>Several tools are involved in the data pipeline, including tools in the Red Hat OpenShift Container Platform monitoring stack and in the Hybrid Cloud Console, with differing intervals:</p> <p><i>Red Hat OpenShift Container Platform monitoring stack:</i></p> <ul style="list-style-type: none"> New clusters identified every 15 minutes Cluster size updated every 2 hours Cluster cleanup for deleted entities updated every 5 hours <p><i>Red Hat OpenShift Cluster Manager:</i></p> <ul style="list-style-type: none"> New clusters identified to Red Hat Subscription Management every 15 minutes Existing clusters synchronized every 6 hours <p><i>the subscriptions service:</i></p> <ul style="list-style-type: none"> daily, once every 24 hours for Annual subscriptions, multiple times daily for On-Demand subscriptions

PART VI. TROUBLESHOOTING AND COMMON QUESTIONS

When you review the subscriptions service data for your account, you might have additional questions about how those calculations are made or whether the calculations are accurate. Answers for some of the most commonly asked questions might help you understand more about the data that appears in the subscriptions service. Other information can help you troubleshoot some common problems that are experienced by subscriptions service users. In some cases, completing the suggested steps in the troubleshooting information can help you improve the accuracy of the reported data in the subscriptions service.

Troubleshoot

- [Troubleshooting: Correcting over-reporting of virtualized RHEL](#)
- [Troubleshooting: Correcting problems with filtering](#)

Learn more

- [How is the subscription threshold calculated?](#)
- [How is core hour usage data calculated?](#)
- [How do vCPUs and hyperthreading affect the subscriptions service usage data?](#)

CHAPTER 20. TROUBLESHOOTING: CORRECTING OVER-REPORTING OF VIRTUALIZED RHEL

So that the subscriptions service can accurately report Red Hat Enterprise Linux in virtualized environments such as virtual data center (VDC) subscriptions, host-guest mappings must be present in the data that the subscriptions service analyzes. For Red Hat Satellite, the Satellite inventory upload plugin and the virt-who tool gather these mappings for the subscriptions service. For Red Hat Subscription Management, the virt-who tool gathers these mappings. For each of these subscription management options, all necessary tools must be both installed and properly configured to accurately report virtual RHEL usage.


If these tools are not used, virtual usage data cannot be correctly calculated. In that type of scenario, guests are counted, not ignored. Each guest is counted as an individual virtual machine, leading to a rapid escalation of virtualized socket count and unusual deployment over capacity. When multiplied over numerous VDC subscriptions, all running multiple guests, the subscriptions service could easily show RHEL overdeployment that significantly exceeds your subscription threshold.

The following example contains an isolated RHEL usage and utilization graph view from the subscriptions service interface. For the first part of the time period, the graph displays physical, virtual, and public cloud usage, but no hypervisor usage is found. For the second part of the time period, where the Satellite inventory upload plugin and virt-who are correctly configured to find and report host-guest mappings, the graph displays new reporting of hypervisor usage and a substantial drop in the reporting of virtual usage. Throughout the time period displayed, the subscription threshold remains constant. Before correction, total RHEL usage exceeds the subscription threshold. After correction, hypervisor usage can be counted and virtual usage is considerably reduced. The total RHEL usage now falls below the subscription threshold.

Figure 20.1. Corrected RHEL virtual usage and newly displayed RHEL hypervisor usage with virt-who and Satellite data

Red Hat Enterprise Linux

Monitor your Red Hat Enterprise Linux usage by physical, virtual, and public cloud sockets. [Learn more about Subscriptions reporting](#)




Procedure

To correct over-reporting of virtualized RHEL usage in the subscriptions service, make sure that you have completed the following steps:

1. Review your RHEL subscription profile in Red Hat Satellite or Red Hat Subscription Management and determine which subscriptions require virt-who.
 - In the Satellite Web UI, click **Content > Subscriptions**. If needed, use the **Search** field to narrow the list of results. Review the values in the **Requires Virt-Who** column. If any check box is selected, you must configure virt-who.
 - In the Overview page of the Red Hat Subscription Management Customer Portal interface,

click **View All Subscriptions**. If needed, use filtering to narrow the list of results. Select a subscription name to view the details. If **Virt-Who: Required** appears in the **SKU Details**, you must configure virt-who.

2. Confirm that the virt-who tool is deployed on your hypervisors so that host-guest mappings can be communicated. For more information, see the virtualization documentation that is appropriate for your subscription management tool.
 - For Satellite, see the following information as appropriate for your version: [Configuring Virtual Machine Subscriptions in Red Hat Satellite](#)
 - For Red Hat Subscription Management, see the following information: [Configuring Virtual Machine Subscriptions in Red Hat Subscription Management](#)



NOTE

If you are also using simple content access, the workflows for some of the subscription management tools such as virt-who are altered. In addition to reviewing the information in the subscription management documentation, see also the information in the [Simple Content Access](#) article.

3. For Satellite, make sure that the Satellite inventory upload plugin is installed and configured to supply data to the subscriptions service.

CHAPTER 21. TROUBLESHOOTING: CORRECTING PROBLEMS WITH FILTERING

The subscriptions service includes several filters that you can use to sort data by different characteristics. These characteristics include subscription attributes, also known as system purpose or subscription settings, depending on the product. Types of subscription attributes include service level agreement (SLA), usage, and others.

Subscription attributes values must be set on systems to enable filtering by those values on the product-level pages in the the subscriptions service interface. There are different methods to set these values, such as directly in the product or in one of the subscription management tools. Subscription attributes values should be set by only one method to avoid the potential for mismatched values.

In the older entitlement-based subscription model, the system purpose values are used by the subscription management tools such as Red Hat Satellite or Red Hat Subscription Management to help match subscriptions with systems. If a system is correctly matched with a subscription, the system status value (**System Status Details** or **System Purpose Status** in the various tools) shows as **Matched**. However, if you are using simple content access with the subscriptions service, that usage of system purpose is obsolete, because subscriptions are not attached to systems. After you enable simple content access, the system status shows as **Disabled**.



NOTE

The **Disabled** state for the system status means that per-system subscription attachment is not being enforced. It does not mean that system purpose values themselves are unimportant. The subscriptions service filters related to system purpose values will not show reliable data if these values are not set for all systems.

Procedure

If the filters that relate to subscription attributes (system purpose values) are showing unexpected results, you might be able to improve the accuracy of that data by ensuring that the subscription attributes are set correctly:

1. Review system information in your preferred subscription management tool to detect whether there are systems where the subscription attributes are missing.
2. If there are missing values for subscription attributes, set those values. You might be able to use options to set these values in bulk, depending on the type and version of subscription management tool that you are using.

Additional resources

- For more information about how to set system purpose values in bulk in Red Hat Satellite, see the section about editing system purpose for multiple hosts in the [Managing Hosts](#) guide.
- For more information about how to use Ansible and the **subscription-manager** command to set system purpose values in bulk for Red Hat Subscription Management, see the [redhat-subscription](#) module information.

CHAPTER 22. HOW IS THE SUBSCRIPTION THRESHOLD CALCULATED?

In the subscriptions service, the usage and utilization graph for most product pages contains a subscription threshold. This line shows the maximum capacity of similar subscriptions across all of your contracts.



NOTE

Some product pages do not show a subscription threshold on the graph.

- For a product page that includes pay-as-you-go On-Demand subscriptions, that graph does not display a subscription threshold because of the characteristics of that subscription type.
- For an account that includes any subscription with a unit of measurement (UoM) of "Unlimited" as part of the terms, the graph for any product page that includes this subscription does not display a subscription threshold. If filtering is used to exclude this subscription from the views, the graph will display a subscription threshold for the filtered data.

To measure the maximum capacity of an organization's account and plot the subscription threshold line in the graph, the subscriptions service does the following steps:

1. Accesses the Red Hat internal subscription services to gather subscription-related contract data for the account.
2. Analyzes every subscription in the account, including each SKU (stock-keeping unit) that was purchased and the amount of each SKU that was purchased.
3. Determines which products are provided in each SKU that is found.
4. Calculates the maximum amount of technology that is provided by a subscription by multiplying the amount of technology that a SKU allows by the number of that SKU that was purchased in the subscription. The amount of technology that a SKU allows is the unit of measurement for the SKU multiplied by the number of these units (the limit) that the SKU provides.
5. Adds the maximum amount of technology for every subscription to determine the subscription threshold that appears on the graph for every product or product portfolio.
6. Analyzes the available subscription attributes data (also known as system purpose data or subscription settings) to enable filtering of that data with the filters in the subscriptions service.

CHAPTER 23. HOW IS CORE HOUR USAGE DATA CALCULATED?

The introduction of the new pay-as-you-go On-Demand subscription type in 2021 resulted in new types of units of measurement in the subscriptions service, in addition to the units of measurement for sockets or cores. These new units of measurement are compound units that function as derived units, where the unit of measurement is calculated from other base units.

At this time, the newer derived units for the subscriptions service add a base unit of time, so these new units are measuring consumption over a period of time. Time base units can be combined with base units that are appropriate for specific products, resulting in derived units that meter a product according to the types of resources that it consumes.

In addition, for a subset of those time-based units, usage data is derived from frequent, time-based sampling of data instead of direct counting. In part, the sampling method might be used for a particular product or service because of the required unit of measurement and the capabilities of the Red Hat OpenShift monitoring stack tools to gather usage data for that unit of measurement.

When the subscriptions service tracks subscription usage with time-based metrics that also use sampling, the metrics used and the units of measurement applied to those metrics are based upon the terms for the subscriptions for these products. The following list shows examples of time-based metrics that also use sampling to gather usage data:

- Red Hat OpenShift Container Platform On-Demand usage is measured with a single derived unit of measurement of core hours. A *core hour* is a unit of measurement for computational activity on one core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.
- Red Hat OpenShift Dedicated On-Demand is measured with two units of measurement, both derived units of measurement. It is measured in core hours to track the workload usage on the compute machines, and in instance hours to track instance availability as the control plane usage on the control plane machines (formerly the master machines in older versions of Red Hat OpenShift). An *instance hour* is the availability of a Red Hat service instance, during which it can accept and execute customer workloads. For Red Hat OpenShift Dedicated On-Demand, instance hours are measured by summing the availability of all active clusters, in hours.
- Red Hat OpenShift AI (RHOAI) On-Demand usage and Red Hat Advanced Cluster Security for Kubernetes (RHACS) On-Demand usage are measured with a single derived unit of measurement of vCPU hours. A *vCPU hour* is a unit of measurement for cluster size on one virtual core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.

23.1. AN EXAMPLE FOR RED HAT OPENSIFT ON-DEMAND SUBSCRIPTIONS

The following information for Red Hat OpenShift On-Demand subscriptions includes an explanation of the applicable units of measurement, a detailed scenario that shows the steps that the subscriptions service and the other Hybrid Cloud Console and monitoring stack tools use to calculate core hour usage, and additional information that can help you understand how core hour usage is reported in the subscriptions service. You can use this information to help you understand the basic principles of how the subscriptions service calculates usage for the time-based units of measurement that also use sampling.

23.1.1. Units of measurement for Red Hat OpenShift On-Demand subscriptions

The following table provides additional details about the derived units of measurement that are used for the Red Hat OpenShift On-Demand products. These details include the name and definition of the unit of measurement along with examples of usage that would equal one of that unit of measurement. In addition, a sample Prometheus query language (PromQL) query is provided for each unit. This example query is not the complete set of processes by which the subscriptions service calculates usage, but it is a query that you can run locally in a cluster to help you understand some of those processes.

Table 23.1. Units of measurement for Red Hat OpenShift Container Platform On-Demand and Red Hat OpenShift Dedicated On-Demand

Unit of measurement	Definition	Examples
core hour	Computational activity on one core (as defined by the subscription terms), for a total of one hour, measured to the granularity of the meter that is used.	For Red Hat OpenShift Container Platform On-Demand and Red Hat OpenShift Dedicated On-Demand workload usage: <ul style="list-style-type: none"> • A single core running for 1 hour. • Many cores running in short time intervals to equal 1 hour.
	<p>Core hour base PromQL query that you can run locally on your cluster:</p> <pre>sum_over_time((max by (_id) (cluster:usage:workload:capacity_physical_cpu_cores:min:5m))[1h:1s])</pre>	
instance hours, in cluster hours	The availability of a Red Hat service instance, during which it can accept and execute customer workloads.	In a cluster hour context, for Red Hat OpenShift Dedicated On-Demand control plane usage: <ul style="list-style-type: none"> • A single cluster that spawns pods and runs applications for 1 hour. • Two clusters that spawn pods and run applications for 30 minutes.
	<p>Instance hour base PromQL query that you can run locally on your cluster:</p> <pre>group(cluster:usage:workload:capacity_physical_cpu_cores:max:5m[1h:5m]) by (_id)</pre>	

23.1.2. Example core hour usage calculation

The following example describes the process for calculating core hour usage for a Red Hat OpenShift On-Demand subscription. You can use this example to help you understand other derived units of measurement where time is one of the base units of the usage calculation and sampling is used as part of the measurement. For example, the vCPU hour calculation for Red Hat OpenShift AI On-Demand is done in the same way, except that the measurement is for virtual cores.

To obtain usage in core hours, the subscriptions service uses numerical integration. Numerical integration is also commonly known as an "area under the curve" calculation, where the area of a complex shape is calculated by using the area of a series of rectangles.

The tools in the Red Hat OpenShift monitoring stack contain the Prometheus query language (PromQL) function **sum_over_time**, a function that aggregates data for a time interval. This function is the foundation of the core hours calculation in the subscriptions service.

```
sum_over_time((max by (_id) (cluster:usage:workload:capacity_physical_cpu_cores:min:5m))[1h:1s])
```



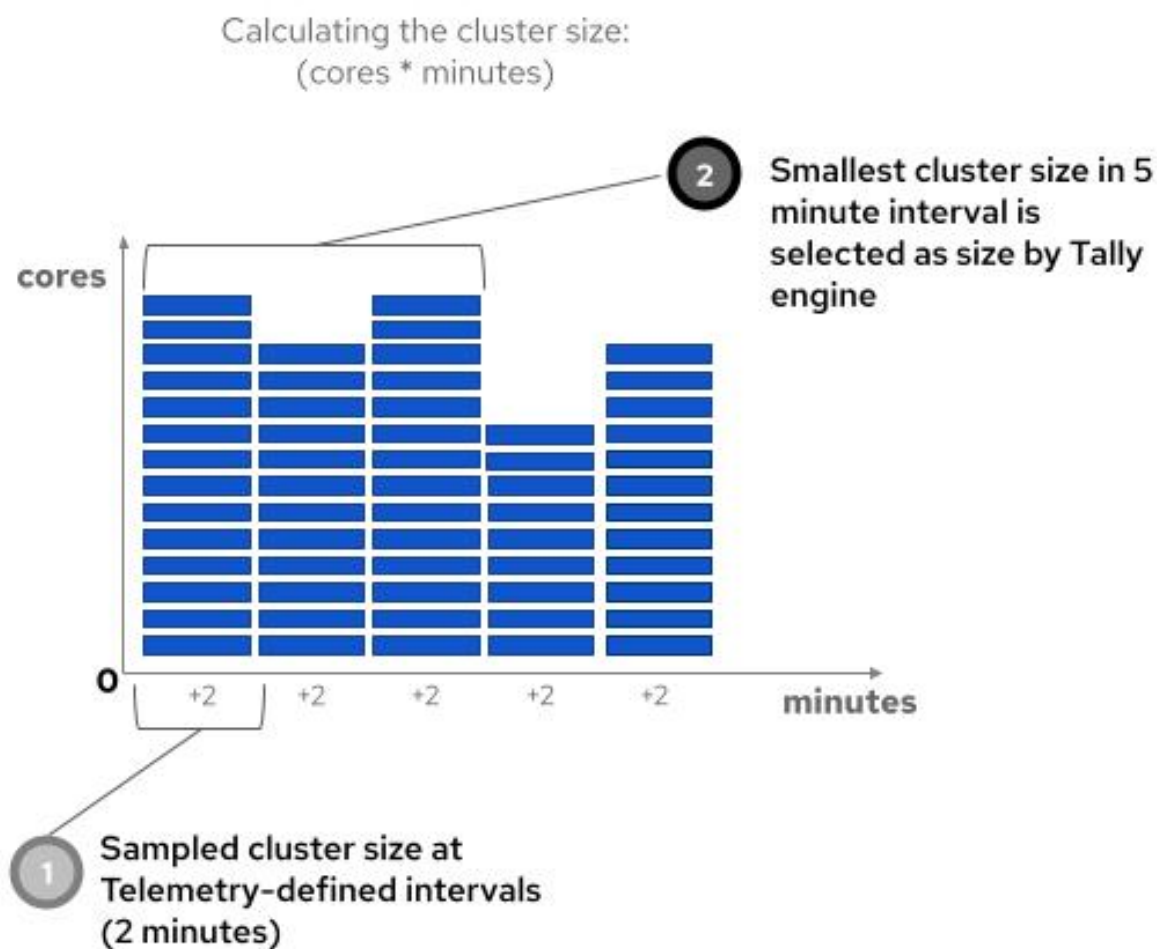
NOTE

You can run this PromQL query locally in a cluster to show results that include the cluster size and a snapshot of usage.

Every 2 minutes, a cluster reports its size in cores to the monitoring stack tools, including Telemetry. One of the Hybrid Cloud Console tools, the Tally engine, reviews this information every hour in 5 minute intervals. Because the cluster reports to the monitoring stack tools every 2 minutes, each 5 minute interval might contain up to three values for cluster size. The Tally engine selects the smallest cluster size value to represent the full 5 minute interval.

The following example shows how a sample cluster size is collected every 2 minutes and how the smallest size is selected for the 5 minute interval.


Figure 23.1. Calculating the cluster size



Then, for each cluster, the Tally engine uses the selected value and creates a box of usage for each 5 minute interval. The area of the 5-minute box is 300 seconds times the height in cores. For every 5 minute box, this core seconds value is stored and eventually used to calculate the daily, account-wide aggregation of core hour usage.

The following example shows a graphical representation of how an area under the curve is calculated, with cluster size and time used to create usage boxes, and the area of each box used as building blocks to create daily core hour usage totals.

Figure 23.2. Calculating the core hours



Every day, each 5 minute usage value is added to create the total usage of a cluster on that day. Then the totals for each cluster are combined to create daily usage information for all clusters in the account. In addition, the core seconds are converted to core hours.

During the regular 24-hour update of the subscriptions service with the previous day's data, the core hour usage information for pay-as-you-go subscriptions is updated. In the subscriptions service, the daily core hour usage for the account is plotted on the usage and utilization graph, and additional **core hours used** information shows the accumulated total for the account. The current instances table also lists each cluster in the account and shows the cumulative number of core hours used in that cluster.



NOTE

The core hour usage data for the account and for individual clusters that is shown in the subscriptions service interface is rounded to two decimal places for display purposes. However, the data that is used for the subscriptions service calculations and that is provided to the Red Hat Marketplace billing service is at the millicore level, rounded to 6 decimal places.


Every month, the monthly core hour usage total for your account is supplied to Red Hat Marketplace for invoice preparation and billing. For subscription types that are offered with a four-to-one relationship of core hour to vCPU hour, the core hour total from the subscriptions service is divided by 4 for the Red Hat Marketplace billing activities. For subscription types that are offered with a one-to-one relationship of core hour to vCPU hour, no conversion in the total is made.

After the monthly total is sent to Red Hat Marketplace and the new month begins, the usage values for the subscriptions service display reset to 0 for the new current month. You can use filtering to view usage data for previous months for the span of one year.

23.1.3. Resolving questions about core hour usage

If you have questions about core hour usage, first use the following steps as a diagnostic tool:

1. In the subscriptions service, review the cumulative total for the month for each cluster in the current instances table. Look for any cluster that shows unusual usage, based on your understanding of how that cluster is configured and deployed.

**NOTE**

The current instances table displays a snapshot of the most recent monthly cumulative total for each cluster. Currently this information updates a few times per day. This value resets to 0 at the beginning of each month.

2. Then review the daily core hour totals and trends in the usage and utilization graph. Look for any day that shows unusual usage. It is likely that unusual usage on a cluster that you found in the previous step corresponds to this day.

From these initial troubleshooting steps, you might be able to find the cluster owner and discuss whether the unusual usage is due to an extremely high workload, problems with cluster configuration, or other issues.

If you continue to have questions after using these steps, you can contact your Red Hat account team to help you understand your core hour usage. For questions about billing, use the support instructions for Red Hat Marketplace.

CHAPTER 24. HOW DO VCPUS, HYPER-THREADING, AND SUBSCRIPTION STRUCTURE AFFECT THE SUBSCRIPTIONS SERVICE USAGE DATA?

The Red Hat OpenShift portfolio contains offerings that track usage with a unit of measurement of cores, but this measurement is obfuscated by virtualization and multithreading technologies. The behavior of these technologies led to the development of the term *vCPUs* to help describe the virtual consumption of physical CPUs, but this term can vary in its meaning. In addition, the structure of Red Hat OpenShift offerings can be complex, making usage data in the subscriptions service difficult to understand.

Red Hat has responded to various customer concerns about Red Hat OpenShift usage data through a series of improvements, both to the subscriptions service itself and to the underlying technologies and methodologies that inform Red Hat OpenShift usage tracking.

24.1. IMPROVED CALCULATIONS FOR X86-64 ARCHITECTURES WITH SIMULTANEOUS MULTITHREADING

October 2021: This change assumes that simultaneous multithreading on x86-64 architectures is enabled, resulting in more accurate usage data within the subscriptions service.

Across different technology vendors, the term *vCPU* can have different definitions. If you work with a number of different vendors, the definition that you use might not match the definition that is used by Red Hat. As a result, you might not be familiar with how Red Hat and the subscriptions service measures usage when *vCPUs* and simultaneous multithreading (also referred to as hyper-threading) are in use within your environment.

Some vendors offer hypervisors that do not expose to guests whether the CPUs of the guests use simultaneous multithreading. For example, recent versions of the VMware hypervisor do not show the simultaneous multithreading status to the kernel of the VM, and always report threads per core as 1. The effect of this counting method is that customers can interpret the subscriptions service reporting of Red Hat OpenShift usage data related to *vCPUs* to be artificially doubled.

To address customer concerns about *vCPU* counting, Red Hat has adjusted its assumptions related to simultaneous multithreading. Red Hat now assumes simultaneous multithreading of 2 threads per core for x86 architectures. For many hypervisors, that assumption results in an accurate counting of *vCPUs* per core, and customers who use those hypervisors will see no change in their Red Hat OpenShift usage data in the subscriptions service.

However, other customers who use hypervisors that do not expose simultaneous multithreading status to the kernel will see an abrupt change in subscriptions service data in October 2021. Those customers will see their related Red Hat OpenShift usage data in the subscriptions service reduced by 50% on the date that this change in counting is implemented. Past data will not be affected.

Customers who encounter this situation will not be penalized. Red Hat requires that the customer purchase enough subscriptions to cover the usage as counted in the subscriptions service only.

In the past, the discrepancies in the definitions for *vCPUs* have resulted in known problems with the interpretation of usage and capacity data for some subscriptions service users. This change in the assumptions for simultaneous multithreading is intended to improve the accuracy of *vCPU* usage data across a wider spectrum of customers, regardless of the hypervisor technology that is deployed.

If you have questions or concerns related to the usage and capacity data that is displayed in the subscriptions service, work with your Red Hat account team to help you understand your data and

account status. For additional information about the resolution of this problem, you can also log in to your Red Hat account to view the following issue: [Bugzilla issue 1934915](#).

24.2. IMPROVED ANALYSIS OF SUBSCRIPTION CAPACITY FOR CERTAIN SUBSCRIPTIONS

January 2022: These changes improved capacity analysis for subscriptions that include extra entitlements or infrastructure subscriptions. These improvements resulted in a more accurate calculation of usage and capacity data for those subscriptions and a more accurate calculation of the subscription threshold within the subscriptions service for the Red Hat OpenShift portion of your Red Hat account.

- Improved accuracy for subscriptions with numerous entitlements: Certain Red Hat OpenShift subscriptions that included a large capacity of cores also included extra entitlements. These entitlements helped to streamline installation by using tools that rely on attached entitlement workflows. However, these extra entitlements were calculated as extra capacity by the subscriptions service, resulting in confusion about how much Red Hat OpenShift could legally be deployed by customers. As of January 2022, counting methods have been revised to remove the extra entitlements from the capacity calculations.
- Infrastructure subscriptions excluded from capacity calculations: For certain purchases of Red Hat OpenShift subscriptions, a particular type of Red Hat OpenShift infrastructure subscription would be added to that purchase automatically. This type of subscription is used to provide infrastructure support for large deployments. Both version 4.1 and later and version 3.11 subscriptions were affected. Normally for Red Hat OpenShift version 4.1 and later, the subscriptions service does not count infrastructure nodes when calculating your Red Hat OpenShift capacity. However, for accounts that received this infrastructure subscription, the improper calculations were occurring at the subscription level, and that data was passed to the subscriptions service. Red Hat OpenShift capacity numbers were artificially inflated, resulting in an incorrect subscription threshold in the subscriptions service. As of January 2022, an added infrastructure subscription is not considered when calculating your Red Hat OpenShift capacity.

24.3. ISOLATING RED HAT OPENSIFT CLUSTER MANAGER OPERATIONAL METRICS FROM METRICS FOR THE SUBSCRIPTIONS SERVICE

December 2023 and March 2024: These changes involved analyzing the metrics used by Red Hat OpenShift Cluster Manager for internal operational purposes and determining whether these metrics were still the optimal metrics to use for subscription usage tracking purposes in the subscriptions service. The use of these metrics for two different purposes was determined to be ineffective, and increased the chances for inaccurate subscription usage reporting. Therefore, the subscriptions service changed to a service-level metric that is designed exclusively to track the subscription obligation in cores. The overall assumption for simultaneous multithreading in virtualized, x86 environments at a factor of 2 threads per core remains in effect. For additional information about these changes, see the details in the [OCP Cluster Size Corrections](#) Customer Portal article.

CHAPTER 25. WHY IS A RHEL SERVER APPEARING IN AN UNEXPECTED PLACE IN THE SUBSCRIPTIONS SERVICE USAGE?

For some Red Hat products, Red Hat Enterprise Linux is included with and is installed to support that product. This inclusion is sometimes also referred to as being "bundled with" another product. For example, RHEL is included with Red Hat Satellite and is tracked separately in a Satellite page view on the RHEL page in the subscriptions service, but does not appear with the main RHEL usage. In addition, RHEL is included with all Red Hat OpenShift subscriptions, but that RHEL does not appear in the usage results on the RHEL page. Bundled RHEL is not tracked or counted against the total RHEL usage or capacity that is intended for production workloads or similar purposes.

In some cases, these bundled RHEL packages might be inadvertently installed in isolation, without the full set of packages included in the bundled subscription. However, these bundled RHEL packages are still associated with the complete set of certificates contained in the bundled subscription. Because the subscriptions service is aware of the associated certificates and uses them to help identify the subscription, RHEL packages that are installed in isolation from their complete bundle could report usage in unexpected places and result in inaccurate usage, capacity, and subscription threshold reporting.

For example, installing a RHEL package from a bundled Red Hat OpenShift Container Platform subscription without the rest of the packages from that bundled subscription could result in a RHEL server appearing as a cluster in the usage results for the Red Hat OpenShift page in the subscriptions service, even if there are no Red Hat OpenShift Container Platform nodes or clusters on that server.

For best results, Red Hat subscriptions and their associated packages should be installed for the purposes for which they were intended. Installing isolated packages from bundled subscriptions is not a best practice, and will lead to poor quality usage results in the subscriptions service.