



Red Hat Insights 1-latest

Assessing and Monitoring RHEL Resource Optimization with Insights for Red Hat Enterprise Linux

Understanding RHEL resource-usage statistics

Red Hat Insights 1-latest Assessing and Monitoring RHEL Resource Optimization with Insights for Red Hat Enterprise Linux

Understanding RHEL resource-usage statistics

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

Install and begin using the Insights for RHEL resource optimization service. This new service helps manage your public cloud systems. Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see our CTO Chris Wright's message.

Table of Contents

CHAPTER 1. THE RESOURCE-OPTIMIZATION SERVICE FOR PUBLIC-CLOUD SYSTEMS	3
1.1. RESOURCE OPTIMIZATION SERVICE CORE CONCEPTS	3
1.1.1. The resource optimization service performance rules	3
1.1.2. Data security guarantee for the resource optimization service	4
1.1.3. Performance metrics for resource optimization	4
1.1.4. Access usage metrics for the resource optimization service	4
CHAPTER 2. INSTALLING AND CONFIGURING THE RESOURCE-OPTIMIZATION COMPONENTS	5
2.1. INSTALLING RESOURCE-OPTIMIZATION COMPONENTS	5
2.1.1. Installing Ansible and running the resource-optimization installation playbook	5
2.1.2. Installing resource optimization when Ansible is already installed	6
2.1.3. Installing resource optimization without installing or using Ansible	7
2.2. ENABLING KERNEL PRESSURE STALL INFORMATION (PSI)	8
2.3. ENABLING NOTIFICATIONS AND INTEGRATIONS IN THE RESOURCE OPTIMIZATION SERVICE	9
CHAPTER 3. VIEWING RESOURCE OPTIMIZATION REPORTS	10
3.1. VIEWING HISTORICAL UTILIZATION DATA	10
3.2. DOWNLOADING RESOURCE OPTIMIZATION SERVICE REPORTS	10
CHAPTER 4. DISABLING THE RESOURCE OPTIMIZATION SERVICE	12
4.1. REMOVING RESOURCE OPTIMIZATION FILES AND DATA	12
4.2. DISABLING KERNEL PRESSURE STALL INFORMATION (PSI)	13
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	14

CHAPTER 1. THE RESOURCE-OPTIMIZATION SERVICE FOR PUBLIC-CLOUD SYSTEMS

The Red Hat Insights for Red Hat Enterprise Linux resource optimization service enables RHEL customers to assess and monitor their public RHEL cloud usage and optimization. The service shows metrics for the following:

- CPU
- Memory
- Disk-usage

It analyzes those metrics and compares them to resource limits recommended by your public cloud provider. Leveraging data from the past day, the resource optimization service considers each resource parameter in several distinct ways and returns actionable data. This data enables better resource allocation and helps you to save money on your public cloud investment.

Features

The service reveals the following information:

- Utilization and optimization data for existing systems in the Insights for Red Hat Enterprise Linux inventory.
- Range of systems running in the public cloud.
- Overview of system characteristics.
- Highlights potential issues.
- Formulates suggestions for issue resolution.

1.1. RESOURCE OPTIMIZATION SERVICE CORE CONCEPTS

1.1.1. The resource optimization service performance rules

Use the resource optimization service to view performance metrics from your managed hosts that run in the supported public cloud, Amazon Web Services (AWS). The service uses a framework called the Performance Co-Pilot (PCP) toolkit to record performance metrics. These metrics empower you to make better business decisions.

Insights performance rules

The performance rules are sets of conditions that are applied to the data collected by PCP. They identify the following system states:

- **Undersized.** The undersized state is determined by examining CPU, RAM and disk input/output (I/O) usage, and combining that with CPU idle time, over a period of 24 hours. If that results in a high score, the resource optimization service labels the system as too small for its workload. A system will be reported as undersized whenever any of the dimensions are undersized.
- **Oversized.** The oversized state is determined by examining CPU, RAM and disk I/O usage, and combining that with CPU idle time, over a period of 24 hours. If that results in a low score, the resource optimization service labels the system as too big for its workload. A system will be reported as oversized only if all of the dimensions are oversized.

- **Idling.** The idling state is determined by examining CPU, RAM and disk I/O usage, and combining that with CPU idle time, over a period of 24 hours. If that results in very low utilization, the resource optimization service labels the system as appropriate for its workload but underused. The idling condition can be viewed as a needs improvement scenario.
- **Optimized.** The optimized state is determined by examining CPU, RAM and disk I/O usage, and combining that with CPU idle time, over a period of 24 hours. If that results in a middle point, the resource optimization service labels the system as optimized.
- **Under pressure.** This state is only active when Kernel Pressure Stall Information (PSI) has been enabled. Systems are labeled as under pressure when they are optimized utilization-wise, but some pressure condition persists.

The resource optimization service measures the system's state and the desired performance criteria that you have set, in order to assign a score to the system.

Additional resources

For more information about the PCP toolkit and registering PAYG, visit the following links:

- [PCP toolkit website](#)
- [Monitoring performance with PCP toolkit](#)
- [PCP and sysstat comparison](#)
- [How to register a Red Hat Enterprise Linux system running on AWS to Insights](#)

1.1.2. Data security guarantee for the resource optimization service

The resource optimization service adheres to the data and application security practices for Red Hat Insights for Red Hat Enterprise Linux services. For more details see [Security](#).

1.1.3. Performance metrics for resource optimization

The resource optimization service installs the **pcp** package on your system and runs two services, **pmcd** and **pmlogger**. Both are part of the Performance Co-Pilot (PCP) toolkit, which monitor and process specific metrics on your system. Metrics are stored in an archive, which the Insights client uploads to Red Hat Insights for Red Hat Enterprise Linux.

1.1.4. Access usage metrics for the resource optimization service

The resource optimization service captures data from the previous day and provides system utilization metrics after 24 hours. By default, the archive is uploaded to Insights for Red Hat Enterprise Linux at 12:00am +/- 1 hour, local system time. However, the time when this data is uploaded can be configured in the Performance Co-Pilot (PCP) toolkit configuration.

CHAPTER 2. INSTALLING AND CONFIGURING THE RESOURCE-OPTIMIZATION COMPONENTS

Installing resource optimization involves installing packages, configuring settings and enabling local services. This can be done manually, or with an Ansible playbook provided by Red Hat.



NOTE

Pay as you go (PAYG) customers need to register the Insights client with subscription-manager (RHSM). There are two ways to register with subscription-manager:

- Using activation keys (recommended)
- Using your user name and password

For more information about how to register the Insights client, refer to [Client Configuration Guide for Red Hat Insights](#).

Table 2.1. Compatibility information

RHEL Versions	Cloud Provider	Resource Optimization Compatibility
8.x-9.x	AWS	Yes (x86_64 and ARM 64-bit)
7.7-7.9	AWS	Yes (x86_64 and ARM 64-bit)
7.0-7.6	AWS	No
6.x	AWS	No

Prerequisites

The following applications and configurations need to be installed or confirmed before the resource optimization service can be used:

- Cloud marketplace RHEL instance is configured.
- The Insights client is installed on the system and is operational.
- If you want to use Ansible to install or uninstall the resource optimization service:
 - The Ansible repository is enabled and the Ansible client is installed on each system.
 - The system administrator can run Ansible Playbooks.


2.1. INSTALLING RESOURCE-OPTIMIZATION COMPONENTS

There are a few options for installing resource-optimization components. Choose whichever works with your Ansible workflow.

2.1.1. Installing Ansible and running the resource-optimization installation playbook

The use of Ansible is recommended to expedite the installation process. This procedure installs the Ansible client and runs the Ansible Playbook on your system.

Cloud marketplace images on Amazon Web Services (AWS) are configured to use repositories hosted by the cloud provider. Currently, these repositories do not contain the Ansible client, so you must perform the following steps to enable the Ansible repository on your cloud marketplace - managed RHEL system.



NOTE

On RHEL 8.6 and later, and RHEL 9.0, Red Hat recommends using Ansible Core. For more information, see [Updates to using Ansible in RHEL 8.6 and 9.0](#) .

Prerequisites

- On RHEL 8, the Ansible repository is enabled.

Procedure on RHEL 8

1. Install Ansible:

```
# yum install ansible-core -y
```

Procedure on RHEL 7

1. Enable the Subscription-Manager repository and register the system

```
# subscription-manager config --rhsm.manage_repos=1  
# subscription-manager register
```

2. Optionally, attach your system to a subscription pool

```
# subscription-manager attach --pool xxxxxxxx
```

3. Enable the required Ansible repository.

```
# subscription-manager repos --enable=rhel-7-server-ansible-2.9-rpms
```

4. Install Ansible:

```
# yum install ansible -y
```

5. If you are using RHEL PAYG and want to use RHUI update servers only, disable the Subscription-Manager repository:

```
# subscription-manager config --rhsm.manage_repos=0
```

2.1.2. Installing resource optimization when Ansible is already installed

Once Ansible is installed, proceed to complete the installation of the resource optimization service.

Procedure

1. Download the Ansible Playbook with the following command:

```
$ curl -O https://raw.githubusercontent.com/RedHatInsights/ros-backend/v2.0/ansible-playbooks/ros_install_and_set_up.yml
```

2. Set localhost in Ansible inventory by appending the line **localhost** to **/etc/ansible/hosts**.
3. Run the Ansible Playbook:

```
# ansible-playbook -c local ros_install_and_set_up.yml
```

The system will show in Insights immediately in a "Waiting for data" state, and data and suggestions will be available the day after registering.

Verification step

Data files with a timestamp will appear under **/var/log/pcp/pmlogger/ros** and after a few minutes, you can verify metrics are being collected:

```
$ ls -l /var/log/pcp/pmlogger/ros
$ pmllogsummary /var/log/pcp/pmlogger/ros/
```

2.1.3. Installing resource optimization without installing or using Ansible

Procedure

If you choose not to use Ansible for installation, use the following manual installation procedure: . Ensure the latest version of insights-client is installed.

```
$ yum update insights-client
```

1. Set **core_collect=True** in **/etc/insights-client/insights-client.conf**
2. Install the Performance Co-Pilot (PCP) toolkit.

```
$ sudo yum install pcp
```


3. Create the PCP configuration file **/var/lib/pcp/config/pmlogger/config.ros** with this content:

```
log mandatory on default {
  hinv.ncpu
  mem.physmem
  mem.util.available
  disk.dev.total
  kernel.all.cpu.idle
  kernel.all.pressure.cpu.some.avg
  kernel.all.pressure.io.full.avg
  kernel.all.pressure.io.some.avg
  kernel.all.pressure.memory.full.avg
  kernel.all.pressure.memory.some.avg
}
[access]
```

```
disallow .* : all;
disallow .* : all;
allow local:* : enquire;
```

- To configure pmlogger to gather the metrics required by resource optimization, add this line to **/etc/pcp/pmlogger/control.d/local**:

```
LOCALHOSTNAME n n PCP_LOG_DIR/pmlogger/ros -r -T24h10m -c config.ros -v 100Mb
```



NOTE

In previous versions of this procedure, this line began with **LOCALHOSTNAME n y**. The procedure now advises that you use **LOCALHOSTNAME n n**, which disables the usage of **pmsocks**. For more information about **pmsocks**, refer to the **man** page for **pmsocks**.

- Start and enable the required PCP services.

```
$ sudo systemctl enable pmcd pmlogger
$ sudo systemctl start pmcd pmlogger
```

- Re-register insights-client and upload the archive. The system will show in Insights immediately in a "Waiting for data" state, and data and suggestions will be available the day after registering.

```
$ sudo insights-client --register
```

Verification step

Data files with a timestamp will appear under **/var/log/pcp/pmlogger/ros** and after a few minutes, you can verify metrics are being collected:

```
$ ls -l /var/log/pcp/pmlogger/ros
$ pmlogsummary /var/log/pcp/pmlogger/ros/
```

2.2. ENABLING KERNEL PRESSURE STALL INFORMATION (PSI)

PSI provides a canonical way to see resource pressure increases as they develop. There are pressure metrics for three major resources: memory, CPU, and input/output (I/O). PSI is available on RHEL 8 and newer versions, and is disabled by default.

When PSI is enabled, the resource optimization service can augment its findings and provide more details and better suggestions. Enabling PSI is strongly recommended to identify peaks.

Procedure

- Edit the **/etc/default/grub** file and append **psi=1** at the end of the GRUB_CMDLINE_LINUX line (mind the quotes).
- Regenerate the grub configuration file.

```
$ sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

3. Reboot the system.

**NOTE**

Enabling PSI incurs in a slight (<1%) performance hit.

Verification step

When PSI is enabled, files for CPU, memory and IO appear under `/proc/pressure`.

2.3. ENABLING NOTIFICATIONS AND INTEGRATIONS IN THE RESOURCE OPTIMIZATION SERVICE

You can enable the notifications service on Red Hat Hybrid Cloud Console to send notifications whenever the resource optimization service detects an issue and generates a suggestion. Using the notifications service frees you from having to continually check the Red Hat Insights for Red Hat Enterprise Linux dashboard for recommendations.

For example, you can configure the notifications service to automatically send an email message whenever the resource optimization service generates a suggestion.

Enabling the notifications service requires three main steps:

- First, an Organization Administrator creates a User access group with the Notifications administrator role, and then adds account members to the group.
- Next, a Notifications administrator sets up behavior groups for events in the notifications service. Behavior groups specify the delivery method for each notification. For example, a behavior group can specify whether email notifications are sent to all users, or just to Organization administrators.
- Finally, users who receive email notifications from events must set their user preferences so that they receive individual emails for each event.

In addition to sending email messages, you can configure the notifications service to pull event data in other ways:

- Using an authenticated client to query Red Hat Insights APIs for event data.
- Using webhooks to send events to third-party applications that accept inbound requests.
- Integrating notifications with applications such as Splunk to route resource optimization recommendations to the application dashboard.

Additional resources

- For more information about how to set up notifications for resource optimization recommendations, see [Configuring notifications on the Red Hat Hybrid Cloud Console](#) and [Integrating the Red Hat Hybrid Cloud Console with third-party applications](#) .

CHAPTER 3. VIEWING RESOURCE OPTIMIZATION REPORTS

Historical data reports are available to help you assess your level of optimization over time, in order to make informed decisions about your future public cloud investment.

3.1. VIEWING HISTORICAL UTILIZATION DATA

The resource optimization service enables you to see how your system utilization scores have been trending over the last 7-45 days. The service displays a bar chart that indicates **CPU Utilization** and **Memory Utilization** percentages on a daily basis.

Complete the following steps to view, filter, and sort system historical utilization data:

Procedure

1. Navigate to the [Business > Resource Optimization](#) page. The system states screen opens.
2. Click on the **Name** header on the left side of the page to filter by **Name, State or Operating system**. Use the sort arrow to the right of each column name to sort by **OS, CPU, Memory Utilization, I/O Output, Suggestions, State, and Last Reported**. Clicking once sorts the column so that optimized systems are displayed first. Clicking a second time sorts the column so that systems categorized as **Waiting for data** are displayed first.
3. Systems that have been analyzed render in blue. Click on the blue system name for a more detailed view.
4. Click on the **Actions** dropdown to see the system's **properties in Inventory**, such as operating system, infrastructure, configuration, BIOS and other data.
5. By default, the resource optimization system displays 7 days of utilization results. Click on the dropdown labeled **Last 7 Days** to view 45 days of utilization data. To view specific days and the utilization scores for those days, use the mouse wheel and buttons to pan and zoom across the bar chart.
6. Scroll down to see specific suggestions for that system.

3.2. DOWNLOADING RESOURCE OPTIMIZATION SERVICE REPORTS

You can download the resource optimization reports for all registered systems. The report identifies the following data gathered over the last 7- 45 days:

- **Registered systems.** This section details the number of systems that are optimal, non-optimal, and stale. The optimized state is determined by examining CPU, RAM, and disk I/O usage, combined with CPU idle time, over a period of 24 hours. If the calculation, based on the examination of the three factors, results in a middle point, the resource optimization service labels the system as optimized. A stale system is defined as one that has not submitted data to the resource optimization service in 7 days.
- **Kernel pressure stall information (PSI)** This is an analysis of the number of systems that have PSI enabled and the number of systems that have NOT enabled PSI. PSI allows you to receive better system recommendations since it can identify resource pressure increases as they develop.
- **System performance issues.** Specific performance issues such as RAM or CPU related peaks are identified along with the number of occurrences.

- **Most used current instance types** The service will evaluate and display your top 5 most frequently used instance types across all registered systems.
- **Suggested instance types.** The service identifies the top 5 frequently suggested instance types based on the most recent utilization metrics. This may indicate that a change is necessary for better resource allocation.
- **Suggested instance types in 45 days** This metric displays the top 5 frequently suggested instance types based on 45 days of historical data. You can also view the effectiveness of changes you have made in the recent past.

Prerequisites

The following prerequisites and conditions must be met to create a PDF of the executive report:

- The Insights client is installed on the system and is operational.
- Performance Co-Pilot is installed and correctly configured.
- At least one system is registered and sending data to the resource optimization.



NOTE

The longer your systems have been sending information to the resource optimization, the more accurate and valuable the recommendations will be.

Procedure

1. Navigate to [Business > Resource Optimization](#).
2. In the top right corner, click on **Download executive report**
3. You will see a dialog box with the message, **Export successful** and notice the PDF file in your taskbar.

Additional Resources

- See section 3.5 Enabling Kernel Pressure Stall Information (PSI)
- PCP toolkit website: [PCP website](#)

CHAPTER 4. DISABLING THE RESOURCE OPTIMIZATION SERVICE

4.1. REMOVING RESOURCE OPTIMIZATION FILES AND DATA

Using Ansible to disable the resource optimization service

Perform the following steps on each system to disable and uninstall the resource optimization service.

Procedure

1. Download the Ansible Playbook with the following command:

```
$ curl -O https://raw.githubusercontent.com/RedHatInsights/ros-backend/v1.0/ansible-playbooks/ros_disable.yml
```

2. Run the Ansible Playbook using command:

```
# ansible-playbook -c local ros_disable_and_clean_up.yml
```

Uninstalling the playbook does not stop or remove the Performance Co-Pilot (PCP) toolkit. Note that PCP may support multiple applications. If you are using PCP exclusively for the resource optimization service, and desire to remove PCP as well, there are a couple options. You can stop and disable the **pmlogger** and **pmcd** services, or remove PCP completely by uninstalling the **pcp** package from the system.

Manually disabling the resource optimization service without the use of Ansible

The use of Ansible is recommended to expedite the uninstallation process. If you choose to not use Ansible, use the manual procedure that follows:

Procedure

1. Disable resource optimization service metrics collection by removing this line from **/etc/pcp/pmlogger/control.d/local**

```
LOCALHOSTNAME n y PCP_LOG_DIR/pmlogger/ros -r -T24h10m -c config.ros -v 100Mb
```

2. Restart PCP so that resource optimization service metrics collection is effectively stopped:

```
$ sudo systemctl pmcd pmlogger
```

3. Remove the resource optimization service configuration file

```
$ sudo rm /var/lib/pcp/config/pmlogger/config.ros
```

4. Remove the resource optimization data from the system

```
$ sudo rm -rf /var/log/pcp/pmlogger/ros
```

5. If you are not using PCP for anything else, you can remove it from your system


```
┆ $ sudo yum remove pcp
```

4.2. DISABLING KERNEL PRESSURE STALL INFORMATION (PSI)

Procedure

1. Edit the **/etc/default/grub** file and remove **psi=1** from the GRUB_CMDLINE_LINUX line.
2. Regenerate the grub configuration file.

```
┆ [user]$ sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

3. Reboot the system.

Verification step

When PSI is disabled, **/proc/pressure** does not exist.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate and prioritize your feedback regarding our documentation. Provide as much detail as possible, so that your request can be quickly addressed.

Prerequisites

- You are logged in to the Red Hat Customer Portal.

Procedure

To provide feedback, perform the following steps:

1. Click the following link: [Create Issue](#)
2. Describe the issue or enhancement in the **Summary** text box.
3. Provide details about the issue or requested enhancement in the **Description** text box.
4. Type your name in the **Reporter** text box.
5. Click the **Create** button.

This action creates a documentation ticket and routes it to the appropriate documentation team. Thank you for taking the time to provide feedback.