



Red Hat Insights 2021

Comparing System Configurations and Baselines in Red Hat Insights Inventory

Using drift service comparisons and baselines for system analysis

Red Hat Insights 2021 Comparing System Configurations and Baselines in Red Hat Insights Inventory

Using drift service comparisons and baselines for system analysis

Legal Notice

Copyright © 2021 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

Compare system configurations of a system over time, or to other systems and baselines to identify discrepancies in your environment, and perform drift analysis.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	3
PROVIDING FEEDBACK ON RED HAT HYBRID CLOUD CONSOLE DOCUMENTATION	4
CHAPTER 1. USING COMPARISONS IN YOUR INSIGHTS FOR RHEL INVENTORY TO PERFORM DRIFT ANALYSIS	5
1.1. COMPARING SYSTEM CONFIGURATIONS IN YOUR INSIGHTS FOR RHEL INVENTORY	5
CHAPTER 2. INSTALLING AND ACCESSING THE DRIFT SERVICE	6
2.1. INSTALLING THE DRIFT SERVICE	6
2.2. ACCESSING THE DRIFT SERVICE	6
CHAPTER 3. CREATING AND WORKING WITH COMPARISONS	8
3.1. USING COMPARISONS OF SYSTEMS IN YOUR INSIGHTS FOR RHEL INVENTORY	8
3.1.1. Creating and adding to comparisons	8
3.1.2. Adding systems manually to your comparisons	9
CHAPTER 4. REVIEWING AND MANAGING BASELINES	11
4.1. THE ROLE OF FACTS IN BASELINES	11
4.2. CREATING BASELINES	11
4.2.1. Creating baselines from scratch	11
4.2.2. Copying an existing baseline to create a new baseline	12
4.2.3. Copying an existing system to create a new baseline	13
4.2.4. Copying an historical system profile to create a new baseline	13
4.3. EDITING OR DELETING BASELINES	14
4.3.1. Editing a baseline	14
4.3.2. Deleting a baseline	15
4.4. COMPARING SYSTEM CONFIGURATIONS TO BASELINES	15
4.5. EXPORTING BASELINES	16
4.5.1. Exporting individual baselines, facts, and categories	16
4.5.2. Exporting all system baselines	16
CHAPTER 5. USING SYSTEM FACTS IN COMPARISONS	17
5.1. HOW FACTS ARE USED IN COMPARISONS	17
5.2. THE ROLE OF SYSTEM FACTS IN COMPARISONS	17
5.3. AVAILABLE FACTS AND THEIR FUNCTIONS	17
5.3.1. Filtering system facts	20
5.3.2. Sorting system facts	21
CHAPTER 6. COMPARING SYSTEMS AND SYSTEM PROFILES	23
6.1. COMPARING SYSTEM PROFILES	23
6.2. COMPARING SYSTEMS TO A REFERENCE POINT	23
6.2.1. Comparing systems to a single point of reference	23
6.3. USING THE HISTORY OF SYSTEM PROFILES	23
6.4. EXPORTING SYSTEM COMPARISON OUTPUT	24
6.4.1. Exporting system comparison output to a CSV file	24

MAKING OPEN SOURCE MORE INCLUSIVE

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](#).

PROVIDING FEEDBACK ON RED HAT HYBRID CLOUD CONSOLE DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so, create a Bugzilla ticket:

1. Go to the [Bugzilla](#) website.
2. As the Component, use **Documentation**.
3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
4. Click **Submit Bug**.

CHAPTER 1. USING COMPARISONS IN YOUR INSIGHTS FOR RHEL INVENTORY TO PERFORM DRIFT ANALYSIS

You can analyze and troubleshoot components of your Insights for RHEL inventory and perform drift analysis by examining and comparing your system configurations.

1.1. COMPARING SYSTEM CONFIGURATIONS IN YOUR INSIGHTS FOR RHEL INVENTORY

The drift service enables you to compare the system configuration of one system to other systems in your Insights for RHEL inventory. You can also compare the system configurations over time to track and analyze changes. In addition, you can set a reference point and compare all systems against that reference. If needed, you can use any of these comparisons to troubleshoot systems at various points in time.

When comparing across different systems, you can filter profile facts to highlight the ones that match, and those that are different, and also see where information is missing or considered problematic.

You can also define and manage baselines. You can use defined baselines in your Insights for RHEL inventory to compare baselines and system configurations. In addition, you can generate CSV output of the systems and baselines you are comparing.

CHAPTER 2. INSTALLING AND ACCESSING THE DRIFT SERVICE

2.1. INSTALLING THE DRIFT SERVICE

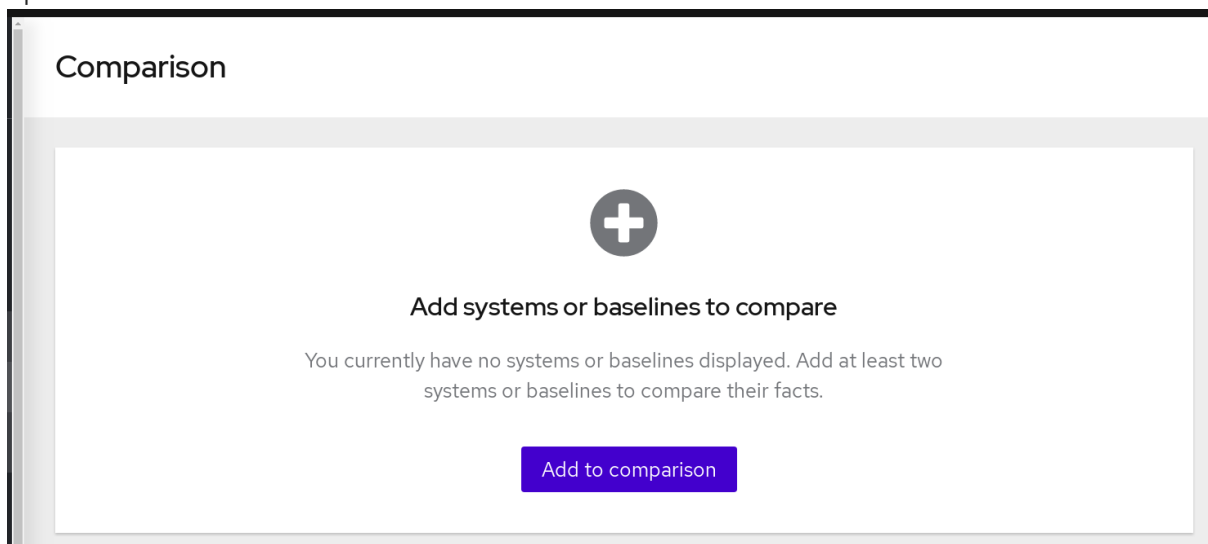
The drift service uses the Insights client and does not require any additional tool to install the service or command to run it.

2.2. ACCESSING THE DRIFT SERVICE

The drift service is part of Insights for Red Hat Enterprise Linux. Access this service via [Red Hat Hybrid Cloud Console](#).

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Comparison](#) page. The Comparison screen opens.



2. Click **Add to comparison**. The Add to comparison screen opens, where you can add systems or baselines to compare.

Add to comparison ✕

Systems Baselines

Name 1-1 of 1 < >

Status Fresh ✕ Stale ✕ Source Insights ✕ Clear filters

Name	Tags	Last seen	Historical profiles
<input type="checkbox"/> rhel8.localdomain	0	⚠ 3 days ago	

1-1 of 1 << < 1 of 1 > >>

[Cancel](#)

- On the **Systems** tab, you can add any systems you want to compare. This screen also lists any baselines that already exist in your Insights for RHEL inventory.
- On the **Baselines** tab, you can create baselines. This screen also lists any baselines that already exist in your Insights for RHEL inventory.

Baselines

Filter by name

Name	Last updated
<input type="checkbox"/> demo2-baseline	5 months ago

3. Click **Submit** to add your systems or baselines for comparison.

Once you add them, you can start comparing the facts of systems and baselines in your Insights for RHEL inventory.

**NOTE**

For steps on how to install the Insights client and register your systems to Insights for RHEL, see [Insights for Red Hat Enterprise Linux, Get Started instructions](#).

CHAPTER 3. CREATING AND WORKING WITH COMPARISONS

With Insights for RHEL, you can analyze and troubleshoot characteristics and performance of components in your Insights for RHEL inventory by viewing and making comparisons between them.

3.1. USING COMPARISONS OF SYSTEMS IN YOUR INSIGHTS FOR RHEL INVENTORY

You use system facts to create the comparisons. Analyzing the facts about components in your Insights for RHEL inventory gives you a view of your system, or several of your systems, at a particular time. You can also create comparisons of your systems over time to learn about and analyze changes in your systems.

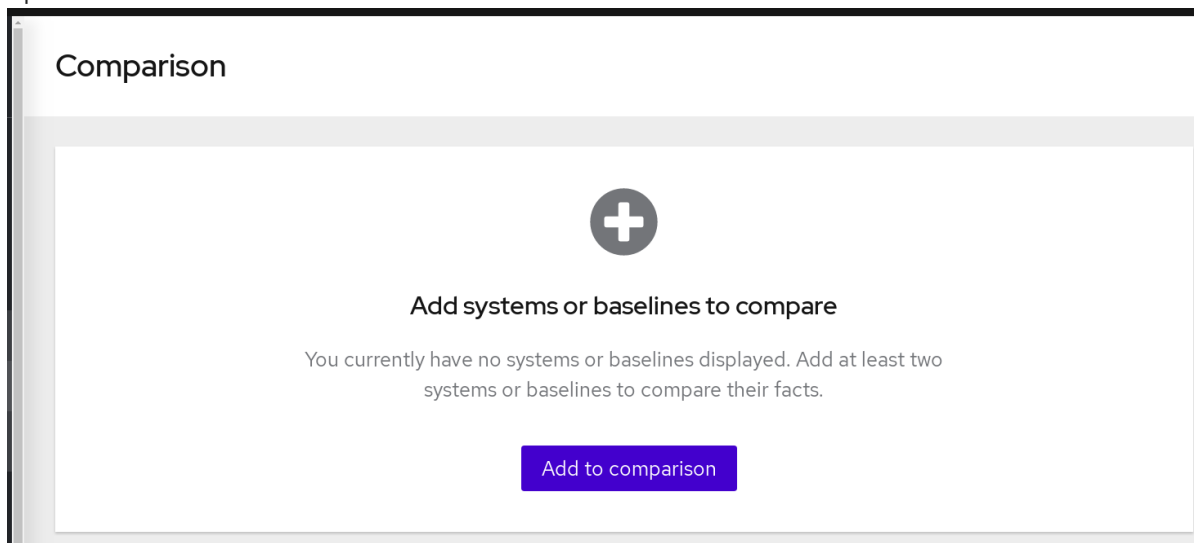
Once you create or add to comparisons, you can start comparing the facts of systems and baselines in your Insights for RHEL inventory.

3.1.1. Creating and adding to comparisons

You can create or add to comparisons you use when comparing systems in your Insights for RHEL inventory.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Comparison](#) page. The Comparison screen opens.



2. Click **Add to comparison**. The Add to comparison screen opens, where you can add systems or baselines to compare.

Add to comparison ✕

Systems **Baselines**

Name Filter by name

1-1 of 1 < >

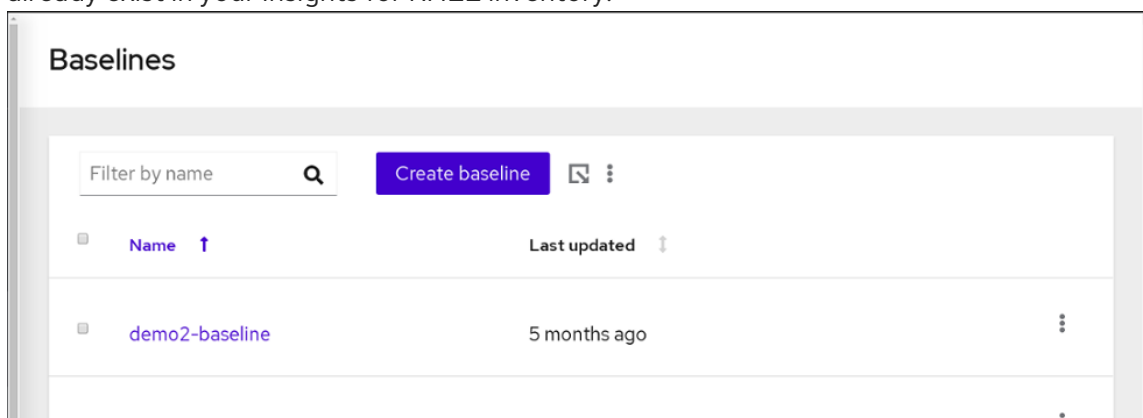
Status
Fresh ✕
Stale ✕
Source
Insights ✕
Clear filters

	Name ↑	Tags	Last seen ↓	Historical profiles
<input type="checkbox"/>	rhel8.localdomain	0	⚠ 3 days ago	🔄

1-1 of 1
<<
<
1
>
>>

Submit
Cancel

- On the **Systems** tab, you can add any systems you want to compare. This screen also lists any baselines that already exist in your Insights for RHEL inventory.
- On the **Baselines** tab, you can create baselines. This screen also lists any baselines that already exist in your Insights for RHEL inventory.



3. Click **Submit** to add your systems or baselines for comparison.

3.1.2. Adding systems manually to your comparisons


Within the drift service, you can add systems registered in your Insights for RHEL inventory.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Comparison](#) page. The Comparison screen opens.
2. Click **Add to comparison**.
3. Select the systems to compare from the list. Alternatively, enter the system name in the search box to find by name, and then select the system.
4. Click **Submit**.



NOTE

- At any time, you can add more systems by clicking the **Add to comparison** button near the top of the Comparison screen.
- Similarly, you can remove a particular system under comparison by clicking the **X** symbol on the upper-right corner of the individual system name, or you can remove all systems under comparison by clicking the options menu () located at the top, then clicking **Clear all comparisons** to start again.

CHAPTER 4. REVIEWING AND MANAGING BASELINES

There are many things you can do when reviewing and managing baselines, such as creating, copying, editing, deleting, exporting, or comparing them.

4.1. THE ROLE OF FACTS IN BASELINES

Baselines are configurations (sets of name / value facts) that can be created from scratch, as a copy of an existing baseline, a copy of an existing system configuration, or a copy of a historical system profile. You use defined baselines in the drift service to compare baselines and system configurations.

You can manage system profile baseline definitions for your organization. That is, you can edit baselines by changing values on facts, or by deleting them. You can use the drift user interface to create, copy, edit, delete, and export baselines.



NOTE

You can also use the baselines API to create and manage baselines. Refer to the Insights for Red Hat Enterprise Linux API documentation for information about using the REST API to query and edit baselines.

4.2. CREATING BASELINES

You can use a variety of system facts and choose the method you want to use to create a baseline. You can:

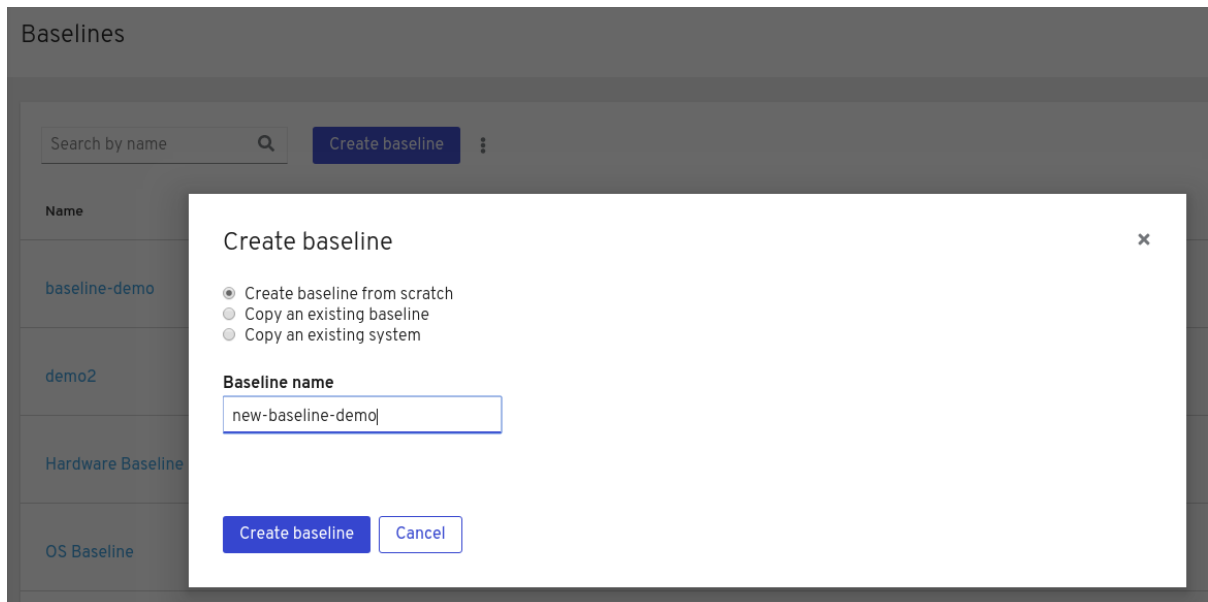
- Create a baseline from scratch
- Copy an existing baseline
- Copy an existing system
- Copy an historical system profile


4.2.1. Creating baselines from scratch

You can create baselines from scratch.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Click **Create baseline**. The Create baseline screen opens.



3. Select **Create baseline from scratch**.
4. In the **Baseline name** field, enter a name for the baseline you want to create.
5. Click **Create baseline**. The Edit screen for the new baseline opens.
6. Click **Add fact or category**. The **Add fact** screen for the new baseline opens.
7. If this is a category (parent fact) under which you will add sub fact(s), select **This is a category** and only enter the **Category name**. Then click **Save**.
8. Otherwise, enter the **Fact name** and **Value**, then click **Save**.
9. To add sub facts under a category, click the more options menu  next to a category and select **Add sub fact**. You can also edit a category or delete a category, using the more options menu.



NOTE


Any changes performed are saved automatically. Therefore, once you have completed adding all facts, sub facts and their values, you can view and manage the baseline you have created by navigating to [Red Hat Enterprise Linux > Drift > Baselines](#).

4.2.2. Copying an existing baseline to create a new baseline

You can copy an existing baseline to create a new baseline.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Click **Create baseline**. The Create baseline screen opens.
3. Select **Copy an existing baseline**.
4. In the **Baseline name** field, enter a name for the new baseline.

5. From the list of existing baselines, choose the baseline you want to copy from and click **Create baseline**. The screen that opens has the newly created baseline with all facts populated.
6. Edit or delete existing facts using the more options menu  next to a fact.
7. You can also add a new fact or category by clicking **Add fact or category**. Similarly, you can add a value to a category.

Once you have completed adding or editing facts, and their values, you can view and manage the baseline you have created by navigating to the [Red Hat Enterprise Linux > Drift > Baselines](#) page.




NOTE


Based on typical expectations about fact behavior, drift alerts users to facts that may require attention. It does not flag facts that differ from each other if the differences are expected or are of no importance. It flags only important differences or similarities that may be problematic. You can then address these exceptions to the typical fact behavior.

4.2.3. Copying an existing system to create a new baseline

You can copy an existing system to create a new baseline.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Click **Create baseline**. The Create baseline screen opens.
3. Select **Copy an existing system**.
4. In the Baseline name field, enter a name for the new baseline.
5. From the existing list of systems, select the system you want to copy from and click **Create baseline**. The screen that opens has the newly created baseline with all facts populated.
6. Edit or delete existing facts using the more options menu  next to a fact.
7. You can also add a new fact or category by clicking **Add fact or category**. Similarly, you can add a value to a category.



NOTE



Based on typical expectations about fact behavior, drift alerts users to facts that may require attention. It does not flag facts that differ from each other if the differences are expected or are of no importance. It flags only important differences or similarities that may be problematic. You can then address these exceptions to the typical fact behavior in a baseline.

Once you have completed adding or editing facts, sub facts and their values, you can view and manage the baseline you have created by navigating to **drift → Baselines**.

4.2.4. Copying an historical system profile to create a new baseline

You can copy an historical system profile to create a new baseline.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Click **Create baseline**. The Create baseline screen opens.
3. Select **Copy an existing system**
4. In the Baseline name field, enter a name for the new baseline.
5. From the existing list of systems, select the system you want to copy from. In the Historical profile column, click the historical profile icon ().
6. From the drop-down menu, select the profile you want to copy, based on the date and time stamp. After selection, drift displays the timestamp of the profile you selected under the name of the system in the **Name** column. If you change your mind and want to select a different historical profile, click the **X** next to the profile timestamp, and select a different historical profile.
7. When you are satisfied that you have the particular profile you want to copy, click **Create baseline**. The screen that opens has the newly created baseline with all facts populated.
8. Edit or delete existing facts using the more options menu  next to a fact.
9. You can also add a new fact or category by clicking **Add fact or category**. Similarly, you can add a value to a fact.

Once you have completed adding or editing facts, and their values, you can view and manage the baseline you have created by navigating to [Red Hat Enterprise Linux > Drift > Baselines](#) .



NOTE

Based on typical expectations about fact behavior, drift alerts users to facts that may require attention. It does not flag facts that differ from each other if the differences are expected or are of no importance. It flags only important differences or similarities that may be problematic. You can then address these exceptions to the typical fact behavior in a baseline.

4.3. EDITING OR DELETING BASELINES

You can edit baselines by renaming the baseline, and changing, adding, or deleting facts, categories, and sub facts.

4.3.1. Editing a baseline

You can edit an existing baseline.

Procedure

1. To edit a baseline, click the baseline name. The edit screen for the chosen baseline opens.

2. Click **Add fact or category** to add facts or categories to the baseline. The Add fact screen opens.
3. If this is a category (parent fact) under which you will add sub fact(s), select **This is a category** and only enter the **Category name**. Then click **Save**.
4. To edit facts, select the item you want to edit and click **Edit fact** on the more options on the right side of the line containing the fact. The Edit fact screen opens. Change the name and/or value and click **Save**.
5. To edit a baseline by deleting facts or categories, select the items you want to delete and click **Delete fact** on the more options on the right side of the line containing the fact.
6. For fact categories, click the more options menu on the right side of the line containing the fact category. Click to **Add sub fact**, **Edit category**, or **Delete category**, and the appropriate screen opens. Provide or edit the information requested and click **Save**.

4.3.2. Deleting a baseline


You can delete a baseline.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Use the drop-down check box to select any or all baselines. When you make a selection, the



field indicates the number of baselines you have selected.

3. To delete a baseline in the list, click **Delete** on the more options menu  at the top of the screen. An alert will open with a reminder that deleting a baseline cannot be undone. If you still wish to proceed, click **Delete baseline**.

4.4. COMPARING SYSTEM CONFIGURATIONS TO BASELINES


You can compare system configurations to one or more baselines to identify discrepancies in your environment and perform drift analysis.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Comparison](#) page. The Comparison screen opens.
2. Click **Add to comparison**.
3. On the **Systems** tab, select one or more systems from the list. Alternatively, you can search for a system by name, and then select the system.
4. On the **Baselines** tab, select one or more baselines from the list. Alternatively, you can search for a baseline by name, and then select the baseline.
5. Click **Submit**.

At any time, you can add more systems and baselines by clicking the **Add System** button on the right side of the systems already added for comparison.

Similarly, to remove a particular system or baseline under comparison, click the **X** sign in the upper-right corner of the individual system or baseline name panel.

To remove all systems and baselines under comparison, click **Clear all comparisons** from the options menu  located at the top, and start again.

You can view the displayed comparison result and filter as necessary by fact name, comparison state, and category. To export the result, along with any current selections such as filters, to a comma-separated values (CSV) file for further analysis, click the **Export as CSV** option in the menu.


4.5. EXPORTING BASELINES


You can export all system baselines or individual baseline facts or categories to a CSV file and analyze them externally. Note that any search filters applied are preserved in the downloaded CSV file.

4.5.1. Exporting individual baselines, facts, and categories

You can export individual baselines, facts, and/or categories.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. From the list of baselines, select a baseline from the list.
3. On the Edit screen for the baseline you selected, click the Export to CSV icon () next to the **Add fact or category** button.
4. Click **Export to CSV**.




NOTE

Use this same procedure to export individual facts and/or categories to CSV.

4.5.2. Exporting all system baselines

You can export all system baselines to a CSV file.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Baselines](#) page. The Baselines screen opens.
2. Click the Export to CSV icon () next to the **Create baseline** button.
3. Click **Export to CSV**.

CHAPTER 5. USING SYSTEM FACTS IN COMPARISONS

System facts are important components that help you understand your system comparisons. Examining them reveals detailed information about the performance and changes in your Insights for RHEL system inventory. System facts also alert you to system components whose state is unknown, as well as identifying parts of your systems that require attention.

5.1. HOW FACTS ARE USED IN COMPARISONS

Comparison states based on observed fact values provide guidance in managing your system. The drift service indicates facts whose behavior differs from expectations, identifies facts whose state is unknown, and alerts users to facts that require attention.

5.2. THE ROLE OF SYSTEM FACTS IN COMPARISONS

Comparison states based on observed fact values provide guidance in managing your system. The application indicates facts whose behavior differs from expectations, identifies facts whose state is unknown, and alerts users to facts that require attention.

The drift service displays the observed fact values in different colors.

- A red icon indicates an issue you should examine.



- A green icon denotes an expected state or value.



- A state shown in black, marked with a question mark icon, indicates the expected state is unknown.



Some facts are system-specific and are considered unique. Their state is marked in red if values are equal for a given comparison. This is the case for **fqdn** and IP addresses. When marked in red, these facts require your attention.

The drift service expects other facts, such as **last_boot_time**, to be different for all compared systems. For such facts, it does not highlight the differences, and it marks the comparison state as unknown (no opinion).

5.3. AVAILABLE FACTS AND THEIR FUNCTIONS

The table below displays the system facts for use in system comparisons.

Table 5.1. System Facts

Fact Name	Description	Example Value
arch	System architecture	x86_64

Fact Name	Description	Example Value
bios_release_date	BIOS release date; typically MM/DD/YYYY	01/01/2011
bios_vendor	BIOS vendor name	LENOVO
bios_version	BIOS version	1.17.0
cloud_provider	Cloud vendor. Values are google, azure, aws, alibaba , or empty	google
cores_per_socket	Number of CPU cores per socket	2
cpu_flags	Category with a list of CPU flags. Each name is the CPU flag (ex: vmx), and the value is always enabled .	vmx , with a value of enabled .
enabled_services	Category with a list of enabled services. Each name in the category is the service name (ex: crond), and the value is always enabled .	crond , with a value of enabled .
fqdn	System Fully Qualified Domain Name	<i>system1.example.com</i>
infrastructure_type	System infrastructure; common values are virtual or physical	virtual
infrastructure_vendor	Infrastructure vendor; common values are kvm, vmware, baremetal , etc.	kvm
installed_packages	List of installed RPM packages. This is a category.	bash , with a value of 4.2.46-33.el7.x86_64 .
installed_services	Category with a list of installed services. Each name in the category is the service name (ex: crond), and the value is always installed .	crond , with a value of installed .
kernel_modules	List of kernel modules. Each name in the category is the kernel module (ex: nfs), and the value is enabled .	nfs , with a value of enabled .
last_boot_time	The boot time in YYYY-MM-DDTHH:MM:SS format. Informational only; we do not compare boot times across systems.	2019-09-18T16:54:56
network_interfaces	List of facts related to network interfaces.	

Fact Name	Description	Example Value
	There are six facts for each interface: ipv6_addresses , ipv4_addresses , mac_address , mtu , state and type . The two address fields are comma-separated lists of IP addresses. The state field is either UP or DOWN . The type field is the interface type (ex: ether , loopback , bridge , etc.).	
	Each interface (ex: lo , em1 , etc) is prefixed to the fact name. For example, em1's mac address would be the fact named em1.mac_address .	
	Most network interface facts are compared to ensure they are equal across systems. However, ipv4_addresses , ipv6_addresses , and mac_address are checked to ensure they are different across systems. A subexception for lo should always have the same IP and mac address on all systems.	
number_of_cpus	Total number of CPUs	1
number_of_sockets	Total number of sockets	1
os_kernel_version	Kernel version	4.18.0
os_release	Kernel release	8.1
running_processes	List of running processes. The fact name is the name of the process, and the value is the instance count.	crond , with a value of 1 .
sap_instance_number	SAP instance number	42
sap_sids	SAP system ID (SID)	A42
sap_system	Boolean field that indicates if SAP is installed on the system	True
sap_version	SAP version number	2.00.052.00.1599 235305
satellite_managed	Boolean field that indicates is a system is registered to a Satellite server.	FALSE

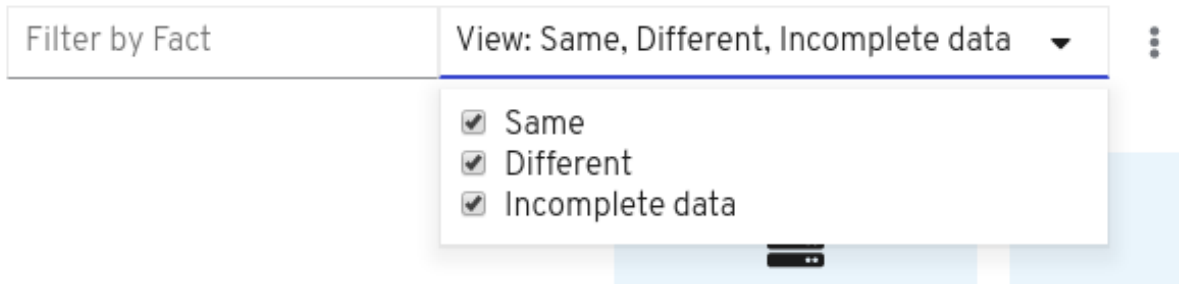
Fact Name	Description	Example Value
selinux_current_mode	Current SELinux mode	enforcing
selinux_config_file	SELinux mode set in the config file	enforcing
system_memory	Total system memory in human-readable form	3.45 GiB
tuned_profile	Current profile resulting from the command tuned-adm active	desktop
yum_repos	List of yum repositories. The repository name is added to the beginning of the fact. Each repository has the associated facts base_url,enabled, and gpgcheck.	Red Hat Enterprise Linux 7 Server (RPMs).base_ur I would have the value https://cdn.redhat.com/content/dist/rhel/server/7/\$releasever/\$basearch/os

5.3.1. Filtering system facts

You can filter system facts in several ways: by the fact comparison state, by fact name, and/or by fact category.

Procedure

1. Navigate to the [Red Hat Enterprise Linux > Drift > Comparison](#) page. The Comparison screen opens.
2. On the Comparison screen, click **Add to comparison**.
3. On the Add to Comparison screen, select the **Systems** tab to see the systems that are checked into your Insights for RHEL inventory.
4. In the Name column, select the checkbox for two or more systems and click **Submit**. The Comparison screen opens, showing the state of facts in the systems.
5. Filtering by comparison state: Click the **View** drop-down list and select **Same** to show only the facts for which values are the same, **Different** to show only the facts that are different, or **Incomplete Data** to show only the facts where the information is incomplete. You can also select a combination of **Same**, **Different**, and **Incomplete data** states, and clear selections as necessary. When you first add systems for comparison, all three options are selected by default.



Filtering by fact name

Enter the fact name in the search box at the top to filter by specific fact name. For example, entering **kernel** as a filter displays all facts containing **kernel** in their name. Or enter **installed_packages** in the search box to view all packages.

Filtering by fact category

Enter a fact category in the search box to compare systems by that category. Examples include **installed_packages**, **installed_services**, **kernel_modules**, **network_interfaces**, **yum_repos**, **cpu_flags**, and **enabled_services**.

In the following example screen, you can see the system comparison data filtered by facts that show a difference across systems. Some facts, such as the **fqdn**, are expected to be different for each system, but the installed packages are expected to be the same. Over time, some packages have been upgraded on system 1, but not on system 2 and system 3, as shown by expanding the fact category **installed_packages**.

Comparison

View: Same, Different, Incomplete data ▾
Add to comparison
🔍 ⋮

1 - 24 of 24 ▾ < >

State

Same ✕

Different ✕

Incomplete data ✕

Fact ↑	State ↓	rhel6-chicago.example.com <small>☆ 20 Apr 2020, 13:37 UTC 🔍</small>	rhel6-tokyo.example.com <small>☆ 20 Apr 2020, 13:19 UTC 🔍</small>	soundwave.infra.pnhc <small>☆ 20 Apr 2020, 14:34 UTC 🔍</small>
bios_release_date	!	04/01/2014	01/01/2011	12/12/2018
bios_vendor	!	SeaBIOS	Seabios	Phoenix Technologies LT
bios_version	!	1.11.0-2.el7	0.5.1	6.00

5.3.2. Sorting system facts

You can sort system facts alphabetically. Click the arrow next to **Fact** (Fact ↑) to switch sorting between ascending and descending order. Note that facts are shown in ascending order by default. You can also sort system facts by the comparison **State**. Click the arrow next to **State** (State ↓) to switch to sorting by state.



NOTE

Sorting works in combination with any applied filters. That is, if you have filtered for installed packages or viewing facts by comparison state, the filtered data can be sorted alphabetically or by comparison state.

CHAPTER 6. COMPARING SYSTEMS AND SYSTEM PROFILES

You can compare systems to each other, to profiles, and to a reference point. When you have selected all the systems you want to compare, you can use the Comparison screen to compare the facts in the systems.

6.1. COMPARING SYSTEM PROFILES

After you have added to your inventory all the systems you want to compare, use the Comparison screen to compare the facts in the systems.

6.2. COMPARING SYSTEMS TO A REFERENCE POINT


At times, you may want to compare all systems to a single reference point rather than compare all systems with each other, or compare them as a group. For example, you may need to compare all systems to a baseline, so that all systems are calculated against that baseline. You can also compare a system against time-stamped profiles to understand where and when changes have occurred.

You might also want to invert a comparison. For example, instead of comparing profiles with old time stamps against the latest profile, you may want to compare all profiles against the oldest known working version of a system. Such a comparison enables you to identify changes that deviate from the reference point.

6.2.1. Comparing systems to a single point of reference

You can compare multiple systems or all systems to a single reference point.

Procedure

1. On the Comparison screen, click the  icon on the system header.

The drift service compares each of your selected systems with the fact values in the reference system, displayed in the first position column, with its header highlighted in blue.

The fact values are displayed with every difference highlighted in red.



NOTE


A fact category with multiple values has no highlights until you expand the fact to view details. When you do so, drift shows the specific fact details highlighted in red.

6.3. USING THE HISTORY OF SYSTEM PROFILES

Every time you submit a system for comparison, the submission checks in the profile and marks it with a time stamp. By examining the different profile versions, you can see a view of the system over time.



NOTE

On the Choose systems screen, if a system is marked with a time stamp  icon, you may be able to open that system directly. Otherwise, add the system you want to compare into drift before following step 1.

6.4. EXPORTING SYSTEM COMPARISON OUTPUT

The drift service allows you to export system comparison output, along with any current selections such as filters, to a comma-separated values (CSV) file. You can then open the CSV file with the tool of your choice to compare exported facts and analyze discrepancies in your systems.



NOTE

The exported CSV report preserves all your current selections on the system comparison output, including any filters applied. That is, it follows the WYSIWYG (What You See Is What You Get) paradigm. Therefore, you will need to expand any nested fact categories (**installed_packages**, for example) to be exported in the report.

6.4.1. Exporting system comparison output to a CSV file

You can export your system profile output, along with any filters you used, to a comma-separated values (CSV) file. The filters highlight the differences and incomplete data in your system profiles. You can use this information for researching characteristics about your systems, and troubleshooting any issues you discover.

Procedure

1. On the Comparison output screen for two or more systems, click the **Export to CSV** icon and then select either **Export to CSV**, or **Clear all comparisons**.
2. Open the CSV file with the tool of your choice so that you can compare exported facts and analyze discrepancies in systems.