

Hypervisor details in SUSE Customer Center

This article describes how to provide customers with details about the hypervisors they use. The collected data is available under the customers' SUSE Customer Center (SCC) account.

Publication Date: April 17, 2025

Contents

- 1 What is **scc-hypervisor-collector**? 2
- 2 Benefits 2
- 3 Requirements 2
- 4 How does **scc-hypervisor-collector** work? 2
- 5 Collected data 3
- 6 Configuration 4
- 7 **scc-hypervisor-collector** systemd services 6
- 8 Related topics 7

1 What is **scc-hypervisor-collector**?

scc-hypervisor-collector is a command-line tool that collects details about customer hypervisor solutions and uploads them to SCC.

2 Benefits

Collecting details about hypervisors and providing our customers with them has the following benefits:

- Customers are proactively informed about their subscription consumption. This allows them to stay compliant and purchase more subscriptions if needed.
- Customers can avoid license auditing which is a time consuming and exhausting experience.

3 Requirements

To collect hypervisor data and upload it to SCC, you need:

- An existing VM Host Server running a supported hypervisor—VMware or libvirt.
- To install the scc-hypervisor-collector package and configure the collector to match your environment.

4 How does **scc-hypervisor-collector** work?

The **scc-hypervisor-collector** command collects details (see [Section 5, “Collected data”](#)) about the hypervisor solutions it has been configured to collect data from.

The collected data, when uploaded to SCC, is correlated with existing SCC registrations and subscription entitlements.

You can either run **scc-hypervisor-collector** manually or enable a systemd timer to run it regularly.

5 Collected data

scc-hypervisor-collector collects the following data:

- The system UUID of a libvirt host—it is the primary piece of information used to correlate hypervisor hosts and virtual machines with existing SCC registrations
- The system UUIDs of the managed virtual machines
- Basic hardware capabilities, such as RAM and CPUs, of the hypervisor host
- Type of the hypervisor—VMware or libvirt
- The state of each discovered virtual machine—whether it is running or stopped/shut off

The following example shows a JSON representation of hypervisor data that is being collected.

```
{
  "virtualization_hosts": [
    {
      "identifier": "SCC_06bc59754...",
      "group_name": "group.example.com",
      "properties": {
        "name": "host.example.com",
        "arch": "x86_64",
        "cores": 64,
        "sockets": 1,
        "threads": 128,
        "ram_mb": 515233,
        "type": "KVM"
      },
      "systems": [
        {
          "uuid": "47fc43cc-c704-4e08-82ee-0e1694XXXXYZ",
          "properties": {
            "vm_state": "running",
            "vm_name": "vm1"
          }
        },
        {
          "uuid": "a914d996-fb27-4d74-8ba7-fde99eXXXXYZ",
          "properties": {
            "vm_state": "running",
            "vm_name": "vm2"
          }
        }
      ]
    }
  ],
}
```

```
    ]
  }
  ]
}
```

6 Configuration

You need to configure **scc-hypervisor-collector** and specify the hypervisors that it needs to query, and the credentials required to upload the collected details to SCC. **scc-hypervisor-collector** checks configuration files with the `.yaml` or `.yml` suffix in the `/var/lib/sc-chvc/.config/scc-hypervisor-collector/` directory.



Tip

To include YAML configuration files from a different directory, specify it with the `--config-dir` option. To include a single YAML configuration file, specify it with the `--config` option.

scc-hypervisor-collector can be run as any user with appropriate configuration settings. If you intend to use the `systemd` integration (see [Section 7, “scc-hypervisor-collectorsystemd services”](#)) to run it automatically on a regular basis, configure it to run under the default restricted `scchvc` user account.

The configuration file must include the following top-level entries:

credentials

Includes a collection of credentials that are used by the tool.



Tip

The provided credentials only need the privileges to retrieve the specified details from the hypervisor solution. It is highly recommended that they do not have capabilities to modify the state of the hypervisor.

backends

Lists hypervisors that are queried to obtain the relevant details.

You can verify if the configuration is correct by using the `--check` option.

The following example configuration file allows querying libvirt and VMware hypervisors, and specifies credentials for uploading the data to SCC:

```
credentials:
  scc:
    username: 'SCC_USERNAME'
    password: 'SCC_PASSWORD'

backends:
  - id: 'vcenter1'
    module: 'VMware'
    hostname: 'dc1-vcenter.example.com'
    port: 443
    username: 'VC1_USERNAME'
    password: 'VC1_PASSWORD'

  - id: 'kvmhost1'
    module: 'Libvirt'
    uri: 'qemu+ssh://someuser@kvmhost1.example.com/system'
```



Tip

For more details about configuration options, refer to the corresponding man page [man 5 scc-hypervisor-collector](#).

6.1 Separate collection and upload phases

You can separate the collection and upload phases by manually running the [scc-hypervisor-collector](#) command:

- The [--output](#) option lets you specify the path to a file to store the results of querying the hypervisors in the YAML format. This option disables uploading results to the SCC.
- The [--input](#) option lets you specify the path to the file that contains the previously collected results of querying the hypervisors in the YAML format. These results are uploaded to SCC instead of querying the current state of hypervisors.

This separation helps to fulfil the following scenarios:

- Customer systems that have the necessary access to collect the data from the configured hypervisor systems may not have access to the Internet to upload the collected data to SCC. Conversely, systems that have Internet access to upload data to SCC may not have sufficient access to collect the data from the customer hypervisor solutions.
- You want to sanitize the collected data before it is uploaded to SCC, for example, change host names or virtual machine names. Note that certain values, such as UUID identifiers, cannot be changed.

7 **scc-hypervisor-collector** systemd services

The `scc-hypervisor-collector.timer` timer triggers the `scc-hypervisor-collector.service` service that runs the **scc-hypervisor-collector** command on a daily basis. These `systemd` units are disabled by default.



Tip: Test **scc-hypervisor-collector** manually

Before enabling `scc-hypervisor-collector.timer`, run **scc-hypervisor-collector** manually via the `scchvc` user account, then verify that it runs correctly when triggered by `scc-hypervisor-collector.service`. When the result is satisfactory, consider enabling the timer.

To enable the collector's service and timer, run the following commands:

```
> sudo systemctl enable scc-hypervisor-collector.service
> sudo systemctl enable scc-hypervisor-collector.timer
```

7.1 Customizing systemd unit files

To modify the default behavior of the collector's `systemd` unit files, create `systemd` unit *drop-in* files to override the default behavior, using the **systemctl edit** mechanism.

To change the user under which the `scc-hypervisor-collector.service` service runs the **scc-hypervisor-collector** command, create a drop-in snippet with the following content:

```
[Service]
```

```
User=EXAMPLE_USER  
Group=EXAMPLE_GROUP
```

To change the timer to trigger the `scc-hypervisor-collector.service` service to a weekly cadence with a randomized skew of up to 6 hours rather than 15 minutes, create a drop-in snippet with the following content:

```
[Timer]  
OnCalendar=weekly  
RandomizedDelaySec=6h
```

8 Related topics

- Detailed configuration of `scc-hypervisor-collector` is described in the man page `man 5 scc-hypervisor-collector`.
- Detailed information about `scc-hypervisor-collector` `systemd` units is described in the man page `man 8 scc-hypervisor-collector`.
- Creating `systemd` drop-in files is described in the man page `man 5 systemd.unit`, search for the `Overriding vendor settings` title.
- Virtualization limits and support information are described in <https://documentation.suse.com/sles/html/SLES-all/cha-virt-support.html>.