

SAP Monitoring

SUSE® Linux Enterprise Server for SAP Applications · SUSE Linux Enterprise High Availability

This article shows monitoring solutions for SAP administrators to efficiently monitor their SAP systems. The solutions that is described here works for SUSE® Linux Enterprise Server 12 SP3 to 15 SP2.

Publication Date: September 28, 2024

Contents

- 1 Introduction 2
- 2 Terminology 2
- 3 Installing exporters 2
- 4 SAP HANA database exporter 3
- 5 High Availability cluster exporter 10
- 6 SAP Host exporter 20
- 7 Additional information 23
- 8 Legal notice 23
- A GNU licenses 23

1 Introduction

The exporters described here make it possible to export metrics that can be combined and integrated with Prometheus and Grafana to produce advanced dashboards.

SUSE supports Prometheus and Grafana through SUSE Manager 4.0. Some Grafana dashboards for SAP HANA, SAP S/4HANA, SAP NetWeaver, and the cluster monitoring are provided by SUSE via Grafana community dashboards.

2 Terminology

Grafana

An interactive visualization and analytics Web application. It provides methods to visualize, explore, and query your metrics, and trigger alerts.

Prometheus

A toolkit that collects and evaluates metrics, displays the result, and triggers possible alerts when an observed condition occurs. Metrics can be collected from different targets at specified intervals.

3 Installing exporters

Installation of an exporter follows the procedure below.

PROCEDURE 1: EXPORTER INSTALLATION

1. Install the package. All package are available in SUSE Linux Enterprise Server for SAP Applications.
2. *(Optional)* Copy the configuration file to `/etc/EXPORTER_DIR`. The exact folder name is different for each exporter. This step depends on the exporter. If you skip this step, the default configuration is used.
3. Start the daemon:

```
systemctl start NAME_OF_DAEMON
```

The above procedure is automatically performed by each of the Salt formulas described in *Article "SAP Automation"*.

4 SAP HANA database exporter

SAP HANA database exporter makes it possible to export SAP HANA database metrics. The tool can export metrics from more than one database and tenant if the `multi_tenant` option is enabled in the configuration file (enabled by default).

The labels `sid` (system identifier), `insnr` (instance number), `database_name` (database name) and `host` (machine hostname) are exported for all the metrics.

4.1 Prerequisites

- A running and reachable SAP HANA database (single or multi-container). It is recommended to run the exporter on the same machine with the SAP HANA database. Ideally, each database should be monitored by one exporter.
- One of the following SAP HANA connectors:
 - `dbapi` (official SAP connector) (<https://help.sap.com/viewer/1e-fad1691c1f496b8b580064a6536c2d/Cloud/en-US/39eca89d94ca464ca52385ad50fc7dea.html>)
 - `pyhdb` (unofficial/open source) (<https://github.com/SAP/PyHDB>)
- Certain metrics are collected in the SAP HANA monitoring views by the SAP Host agent (https://help.sap.com/saphelp_nwpi711/helpdata/en/21/98c443122744e-fae67c0352033691d/frameset.htm). To have access to all the monitoring metrics, make sure that the SAP Host agent is installed and running.

4.2 Metrics file

The exporter relies on a metrics file to determine what metrics to export. When the metrics file uses the JSON format, you can use the options listed below.

- enabled (boolean, optional). Determines whether the query is executed or not. If set to false, the metrics for this query are not executed.
- hana_version_range (list, optional). The SAP HANA database versions range where the query is available ([1.0.0] by default). If the current database version is not within the specified range, the query is not executed. If the list has only one element, all versions beyond the specified value (including the defined one) are queried.
- metrics (list) A list of metrics for the query.
- name (string) A name for the exported metrics.
- description (string) A description of the metrics.
- labels (list) A list of labels used to split the records.
- value (string) A name of the column for the exported value (must match with one of the columns of the query).
- unit (string) Used unit for the exported value (for example, `mb`).
- type (enum{gauge}) Defines the type of the exported metric (gauge is the only available option).

Below is an example of a metrics file:

```
{
  "SELECT TOP 10 host, LPAD(port, 5) port, SUBSTRING(REPLACE_REGEXPR('\n' IN
statement_string WITH ' ' OCCURRENCE ALL), 1,30) sql_string, statement_hash sql_hash,
execution_count, total_execution_time + total_preparation_time total_elapsed_time FROM
sys.m_sql_plan_cache ORDER BY total_elapsed_time, execution_count DESC;":
  {
    "enabled": true,
    "hana_version_range": ["1.0"]
    "metrics": [
      {
        "name": "hanadb_sql_top_time_consumers",
        "description": "Top statements time consumers. Sum of the time consumed in all
executions in Microseconds",
        "labels": ["HOST", "PORT", "SQL_STRING", "SQL_HASH"],
        "value": "TOTAL_ELAPSED_TIME",
        "unit": "mu",
        "type": "gauge"
      },
      {
        "name": "hanadb_sql_top_time_consumers",
```

```

        "description": "Top statements time consumers. Number of total executions of the
SQL Statement",
        "labels": ["HOST", "PORT", "SQL_STRING", "SQL_HASH"],
        "value": "EXECUTION_COUNT",
        "unit": "count",
        "type": "gauge"
    }
]
}
}

```

4.3 Installing the SAP HANA database exporter

Use the `zypper install prometheus-hanadb_exporter` command to install the exporter.

You can find the latest development repositories at [SUSE's Open Build Service \(https://build.opensuse.org/package/show/network:ha-clustering:sap-deployments:development/prometheus-hanadb_exporter\)](https://build.opensuse.org/package/show/network:ha-clustering:sap-deployments:development/prometheus-hanadb_exporter).

To install the exporter from the source code, make sure you have Git and Python 3 installed on your system. Run the following commands to install the exporter with the PyHDB SAP HANA connector:

```

git clone https://github.com/SUSE/hanadb_exporter
cd hanadb_exporter # project root folder
virtualenv virt
source virt/bin/activate
pip install pyhdb
pip install .

```

4.4 Configuring the exporter

Use the following example of the `config.json` configuration file as a starting point.

```

{
  "listen_address": "0.0.0.0",
  "exposition_port": 9668,
  "multi_tenant": true,
  "timeout": 30,
  "hana": {
    "host": "HOSTNAME",
    "port": 30013,
    "user": "SYSTEM",

```

```

    "password": "PASSWORD",
    "ssl": false,
    "ssl_validate_cert": false
  },
  "logging": {
    "config_file": "PATH/logging_config.ini",
    "log_file": "PATH/hanadb_exporter.log"
  }
}

```

Below is a list of key configuration options.

- `listen_address` IP address of the Prometheus exporter (0.0.0.0 by default).
- `exposition_port` Port through which the Prometheus exporter is accessible (9968 by default).
- `multi_tenant` Export the metrics from other tenants. This requires a connection to the system database (port 30013).
- `timeout` Timeout to connect to the database. The app fails if connection is not established within the specified time (even in daemon mode).
- `hana.host` Hostname or IP address of the SAP HANA database host.
- `hana.port` Port through which the SAP HANA database is accessible.
- `hana.userkey` Stored user key (see [Section 4.5, "Using the stored user key"](#)). Use this option if you do not want to store the password in the configuration file. The `userkey` and `user/password` are mutually exclusive. If both are set, `hana.userkey` takes priority.
- `hana.user` Existing user with access right to the SAP HANA database.
- `hana.password` Password of an existing user.
- `hana.ssl` Enable SSL connection (`false` by default). Only available for the `dbapi` connector.
- `hana.ssl_validate_cert` Enable SSL certification validation. This option is required by SAP HANA cloud. Only available for the `dbapi` connector.
- `hana.aws_secret_name` Secret name containing the username and password (see [Section 4.6, "Using AWS Secrets Manager"](#)). Use this option when SAP HANA database is stored on AWS. `aws_secret_name` and `user/password` are mutually exclusive. If both are set, `aws_secret_name` takes priority.

- `logging.config_file` Python logging system configuration file (by default, WARN and ERROR level messages are sent to the syslog).
- `logging.log_file` Logging file (`/var/log/hanadb_exporter.log` by default)

The logging configuration file follows the [Python standard logging system style \(https://docs.python.org/3/library/logging.config.html\)](https://docs.python.org/3/library/logging.config.html).

Using the default configuration file, redirects the logs to the file specified in the JSON configuration file and to the syslog (only logging level up to WARNING).

4.5 Using the stored user key

Use this option to keep the database secure (you can use `user/password` with the `SYSTEM` user for development, as it is faster to set up). To use the `userkey` option, the `dbapi` must be installed (normally stored in `/hana/shared/SID/hdbclient/hdbcli-N.N.N.tar.gz` and installable with `pip3`). The key is stored in the client itself. To use a different client, you must create a new stored user key for the user running Python. To do that, use the following command (note that the `hdbclient` is the same as the `dbapi` Python package):

```
/hana/shared/PRD/hdbclient/hdbuserstore set USER_KEY host:30013@SYSTEMDB hanadb_exporter
pass
```

4.6 Using AWS Secrets Manager

Use the AWS Secrets Manager to store the login credentials outside the configuration file when the SAP HANA database is stored on AWS EC2 instance.

- Create a JSON `secret` (https://docs.aws.amazon.com/secretsmanager/latest/userguide/manage_create-basic-secret.html) file that contains two key-value pairs. The first pair contains the `username` key and the actual database user as the value. The second pair has the `password` key and the actual password as the value. For example:

```
{
  "username": "DATABASE_USER",
  "password": "DATABASE_PASSWORD"
}
```

Use the actual secret as the secret name, and pass it in the configuration file as a value for the `aws_secret_name` entry.

- Configure read-only access from EC2 IAM role to the secret by attaching a [resource-based policy](https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and-access_resource-based-policies.html) (https://docs.aws.amazon.com/secretsmanager/latest/userguide/auth-and-access_resource-based-policies.html) to the secret. For example:

```
{
  "Version" : "2012-10-17",
  "Statement" : [
    {
      "Effect": "Allow",
      "Principal": {"AWS": "arn:aws:iam::123456789012:role/EC2RoleToAccessSecrets"},
      "Action": "secretsmanager:GetSecretValue",
      "Resource": "*",
    }
  ]
}
```

Tips and recommendations:

- Set `SYSTEMDB` as the default database for the exporter to get the tenants data.
- Do not use the stored user key created for the backup, because the key is created using the `sidadm` user.
- Instead of the `SYSTEM` user, use an account limited to accessing the monitoring tables only.
- In case you use a user account with the monitoring role, this user must exist in all the databases (`SYSTEMDB` and tenants).

4.7 Create a new user with the monitoring role

Run the following commands to create a user with the monitoring roles (the commands must be executed in all the databases):

```
su - prdadm
hdbsql -u SYSTEM -p pass -d SYSTEMDB #(PRD for the tenant in this example)
CREATE USER HANADB_EXPORTER_USER PASSWORD MyExporterPassword NO
FORCE_FIRST_PASSWORD_CHANGE;
CREATE ROLE HANADB_EXPORTER_ROLE;
GRANT MONITORING TO HANADB_EXPORTER_ROLE;
GRANT HANADB_EXPORTER_ROLE TO HANADB_EXPORTER_USER;
```


4.8 Running the exporter

Start the exporter with the `hanadb_exporter -c config.json -m metrics.json` command. If the `config.json` configuration file is stored in the `/etc/hanadb_exporter` directory, the exporter can be started with the following command (note that the identifier matches with the `config.json` file without extension):

```
hanadb_exporter --identifier config
```

4.9 Running as a service

To run the `hanadb_exporter` as `systemd` service, install the exporter using the RPM package as described in [Section 4.3, "Installing the SAP HANA database exporter"](#).

Next, create the configuration file as `/etc/hanadb_exporter/my-exporter.json`. You can use the example file above as a starting point (the example file is also available in the `/usr/etc/hanadb_exporter` directory).

You can use the example `/usr/etc/hanadb_exporter/metrics.json` metrics file.

Adjust the default logging configuration file `/usr/etc/hanadb_exporter/logging_config.ini`.

Start the exporter as a daemon. Because there are multiple `hanadb_exporter` instances running on one machine, you need to specify the name of the created configuration file, for example:

```
# systemctl start prometheus-hanadb_exporter@my-exporter
# systemctl status prometheus-hanadb_exporter@my-exporter
# systemctl enable prometheus-hanadb_exporter@my-exporter
```

Important: Configure the Prometheus server

The exporter only exposes a port, without pushing the data to the Prometheus server. This means that the Prometheus server must be configured to periodically pull the data from the exporter. This is done by either adding the `hanadb_exporter` job to the Prometheus server configuration, or by adding `hanadb_exporter` to an existing job. For example:

```
- job_name: hana_db
  static_configs:
    - targets:
      - "HOSTNAME:PORT"
```



Important: Configure firewall

Use the following command to open the port for `hanadb_exporter`.

```
# firewall-cmd --zone=ZONE --add-port=PORT/tcp --permanent
# firewall-cmd --reload
# firewall-cmd --list-all --zone=ZONE
```

Replace `ZONE` with the actual interface used for the exporter, and `PORT` with the actual port number of `hanadb_exporter` (default is 9968).

5 High Availability cluster exporter

Enables monitoring of Pacemaker, Corosync, SBD, DRBD and other components of High Availability clusters. Collects metrics to easily monitor cluster status and health.

Link: https://github.com/ClusterLabs/ha_cluster_exporter ↗.

EXPORT METRICS IN THE PROMETHEUS FORMAT

- Pacemaker cluster summary, nodes and resources stats
- Corosync ring errors and quorum votes
- Health status of SBD devices.
- DRBD resources and connections status.

5.1 Installation

To install the High Availability cluster exporter on SUSE Linux Enterprise, run the `zypper install prometheus-ha_cluster_exporter` command.

5.1.1 Enabling systemd service

The High Availability cluster exporter RPM packages comes with the `ha_cluster_exporter.service` `systemd` service. To enable and start it, use the following command:

```
systemctl --now enable prometheus-ha_cluster_exporter
```

5.2 Using High Availability cluster exporter

You can run the exporter on any of the cluster nodes. Although it is not strictly required, it is advisable to run the exporter on all nodes.

The generated metrics are stored in the `/metrics` path. By default, the metrics can be accessed through the web interface on port 9664.

Although the exporter can run outside an High Availability cluster node, it cannot export any metric it is not able to collect. In this case, the exporter displays a warning message.

5.3 Configuring High Availability cluster exporter

Before you proceed, make sure that the Prometheus server and the firewall are configured as described in *Important: Configure the Prometheus server* and *Important: Configure firewall*

The provided default configuration is designed specifically for the latest version of SUSE Linux Enterprise. If necessary, any of the supported parameters can be modified either via command-line flags or via a configuration file. Use the `ha_cluster_exporter --help` command for more details on configuring parameters from the command line. Refer to the `ha_cluster_exporter.yaml` file for an example configuration.

It is also possible to specify CLI flags via the `/etc/sysconfig/prometheus-ha_cluster_exporter` file.

GENERAL FLAGS

`web.listen-address`

Address to listen on for web interface and telemetry (default 9664).

`web.telemetry-path`

Directory for storing metrics data (default `/metrics`).

`web.config.file`

Path to a the web configuration file (default `/etc/ha_cluster_exporter.web.yaml`).

`log.level`

Logging verbosity (default `info`).

`version`

Print version information.

COLLECTOR FLAGS

crm-mon-path

Path to the `crm_mon` executable (default `/usr/sbin/crm_mon`).

cibadmin-path

Path to the `cibadmin` executable (default `/usr/sbin/cibadmin`).

corosync-cfgtool-path

Path to the `corosync-cfgtool` executable (default `/usr/sbin/corosync-cfgtool`).

corosync-quorumtool-path

Path to the `corosync-quorumtool` executable (default `/usr/sbin/corosync-quorum-tool`).

sbd-path

Path to the `sbd` executable (default `/usr/sbin/sbd`).

sbd-config-path

Path to the `sbd` configuration (default `/etc/sysconfig/sbd/`).

drbdsetup-path

Path to the `drbdsetup` executable (default `/sbin/drbdsetup`).

drbdsplitbrain-path

Path to the `drbd splitbrain` hooks temporary files (default `/var/run/drbd/split-brain`).

5.4 TLS and basic authentication

The High Availability cluster exporter supports TLS and basic authentication. To use TLS or basic authentication, specify a configuration file using the `--web.config.file` parameter. The format of the file is described in <https://github.com/prometheus/exporter-toolkit/blob/master/docs/web-configuration.md>.

5.5 Metrics specification

The following provides an overview of metrics generated by the High Availability cluster exporter.

Pacemaker. The Pacemaker subsystem collects an atomic snapshot of the High Availability cluster directly from the XML CIB of Pacemaker using `crm_mon`.

PACEMAKER

`ha_cluster_pacemaker_config_last_change`

A Unix timestamp in seconds converted to a floating number, corresponding to the last time Pacemaker configuration changed.

`ha_cluster_pacemaker_fail_count`

The fail count per node and resource ID.

`ha_cluster_pacemaker_location_constraints`

Resource location constraints.

LABELS

- `constraint` A unique string identifier of the constraint
- `node` The node the constraint applies to
- `resource` The resource the constraint applies to
- `role` The resource role the constraint applies to (if any)

`ha_cluster_pacemaker_migration_threshold`

The number of migration threshold for each node and resource ID set by a Pacemaker cluster.

`ha_cluster_pacemaker_nodes`

The status of each node in the cluster (one line for the status of every node). 1 indicates the node is in the status specified by the status label, 0 means it is not.

LABELS

- `node` The name of the node (normally the hostname)
- `status` Possible values: `standby`, `standby_onfail`, `maintenance`, `pending`, `unclean`, `shutdown`, `expected_up`, `dc`
- `type` Possible values: `member`, `ping`, `remote`

`ha_cluster_pacemaker_node_attributes`

This metric exposes in its labels raw, opaque, cluster metadata, called node attributes that often leveraged by Resource Agents. The value of each line is always 1.

LABELS

- node The name of the node (normally the hostname)
- name The name of the attribute
- value The value of the attribute

ha_cluster_pacemaker_resources

The status of each resource in the cluster (one line for the status of each resource). 1 means the resource is in the status specified by the status label, 0 means that it is not.

LABELS

- agent The name of the resource agent for the resource
- clone The name of the clone this resource belongs to (if any)
- group The name of the group this resource belongs to, (if any)
- managed Can be either true or false
- node The name of the node hosting the resource
- resource The unique resource name
- role Possible values: started, stopped, master, slave or one of starting, stopping, migrating, promoting, demoting

ha_cluster_pacemaker_stonith_enabled

Whether or not stonith is enabled in the cluster. The value is either 1 or 0.

Corosync. The Corosync subsystem collects cluster quorum votes and ring status by parsing the output of corosync-quorumtool and corosync-cfgtool.

COROSYNC

ha_cluster_corosync_member_votes

The number of votes each member node has contributed to the current quorum.

LABELS

- node_id The internal corosync identifier associated with the node
- node The name of the node (normally the hostname)
- local Indicates whether the node is local

ha_cluster_corosync_quorate

Indicates whether the cluster is quorate. The value is either 1 or 0

ha_cluster_corosync_quorum_votes

Cluster quorum votes (one line per type).

LABELS

- type Possible values: expected_votes, highest_expected, total_votes, quorum.

ha_cluster_corosync_ring_errors

The total number of faulty Corosync rings.

ha_cluster_corosync_rings

The status of each Corosync ring. 1 is healthy, 0 is faulty.

LABELS

- ring_id The internal Corosync ring identifier (normally corresponds to the first member node to join)
- node_id The internal Corosync identifier of the local node
- number The ring number
- address the IP address locally linked to this ring

SBD. The SBD subsystems collect statistics of each device by parsing its configuration and the output of **sbd --dump**.

SBD

ha_cluster_sbd_devices

The SBD devices in the cluster (one line per device). The line is either absent or has the value of 1.

LABELS

- device The path of the SBD device
- status Possible values: healthy, unhealthy

ha_cluster_sbd_timeouts

The SBD timeouts for each SBD device.

LABELS

- device The path of the SBD device
- type Possible values: watchdog, msgwait

DRBD. The DRDB subsystem runs a special **drbdsetup** command to get the current status of a DRDB cluster in the JSON format.

DRBD

ha_cluster_drbd_connections

The DRBD resource connections (one line per resource and per peer_node_id). The line is either absent or has the value of 1.

LABELS

- resource The resource the connection is for
- peer_node_id The id of the node this connection is for
- peer_role Possible values: primary, secondary, unknown
- volume The volume number
- peer_disk_state Possible values attaching, failed, negotiating, inconsistent, outdated, unknown, consistent, uptodate

The total number of lines for this metric is the cardinality of resource multiplied by the cardinality of peer_node_id.

ha_cluster_drbd_connections_sync

The DRBD disk connections in sync percentage. Values are floating numbers between 0 and 100.00.

LABELS

- resource The resource the connection is for
- peer_node_id The id of the node this connection is for
- volume The volume number

ha_cluster_drbd_connections_received

Volume of net data received from the partner via the network connection in KiB (one line per resource and per peer_node_id). The value is an integer greater than or equal to 0.

LABELS

- resource The resource the connection is for
- peer_node_id The id of the node this connection is for
- volume The volume number

ha_cluster_drbd_connections_pending

Number of requests sent to the partner that have not yet been received (one line per resource and per peer_node_id). The value is an integer greater than or equal to 0.

LABELS

- resource The resource the connection is for
- peer_node_id The id of the node this connection is for
- volume The volume number

ha_cluster_drbd_connections_unacked

Number of requests received by the partner but have not yet been acknowledged (one line per resource and per peer_node_id). The value is an integer greater than or equal to 0.

LABELS

- resource The resource the connection is for
- peer_node_id The id of the node this connection is for
- volume The volume number

ha_cluster_drbd_resources

The DRBD resources (one line per name and per volume). The line is either absent or has the value of 1.

LABELS

- resource The name of the resource
- role Possible values: primary, secondary, unknown
- volume The volume number
- disk_state Possible values: attaching, failed, negotiating, inconsistent, outdated, outdated, unknown, consistent, uptodate

The total number of lines for the metric is the cardinality of name multiplied by the cardinality of volume.

ha_cluster_drbd_written

Amount of data in KiB written to the DRBD resource (one line per resource and per volume)
The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_read

Amount of data in KiB read from the DRBD resource (one line per resource and per volume)
The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_al_writes

Number of updates of the activity log area of the meta data (one line per resource and per volume). The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_bm_writes

Number of updates of the bitmap area of the metadata (one line per resource and per volume). The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_upper_pending

Number of block I/O requests forwarded to DRBD but not yet answered by DRBD (one line per resource and per volume). The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_lower_pending

Number of open requests to the local I/O sub-system issued by DRBD (one line per resource and per volume). The value is an integer greater than or equal to 0.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_quorum

Quorum status of the DRBD resource according to the configured quorum policies (one line per resource and per volume). The value is 1 when quorate, or 0 when inquorate.

LABELS

- resource The name of the resource
- volume The volume number

ha_cluster_drbd_split_brain

Signals when there is a split brain occurring per resource and volume. The line is either absent or has the value of 1. To make this metric work you must setup a DRBD custom split-brain handler.

LABELS

- resource The name of the resource
- volume The volume number

Scrape. The scrape subsystem is a generic namespace dedicated to internal instrumentation of the exporter itself.

SCRAPE

ha_cluster_scrape_duration_seconds

The duration of a collector scrape in seconds.

LABELS

- collector collector names that correspond to the subsystem they collect metrics from

Example:

```
# TYPE ha_cluster_scrape_duration_seconds gauge
ha_cluster_scrape_duration_seconds{collector="pacemaker"} 1.234
```

ha_cluster_scrape_success

Indicates whether a collector succeeded. Collectors can fail gracefully, but that does not prevent them from running. If certain metrics cannot be scraped, the value of this metric is 0. In this case, the exporter logs for more details.

LABELS

- collector collector names that correspond to the subsystem they collect metrics from

Example:

```
# TYPE ha_cluster_scrape_success gauge
ha_cluster_scrape_success{collector="pacemaker"} 1
```

6 SAP Host exporter

SAP Host exporter is a Prometheus exporter that enables monitoring of SAP systems (SAP NetWeaver applications).

It is a stateless exporter that retrieves runtime data from the SAP system via the SAPControl Web interface on each HTTP request. Exported data include start service processes, Enqueue Server stats, and AS Dispatcher work process queue stats.



Important: Note about metrics

The exporter does not export metrics that it cannot collect. However, as it is irrelevant for the exporter which subsystems are present in the monitored target, failing to collect metrics is not considered a critical failure. However, when certain collectors fail to either register or perform collect cycles, a warning is added to the log. The `metrics.md` file that

ships with the RPM package provides information about the exported metrics. The file is also available in the project's [GitHub repository \(https://github.com/SUSE/sap_host_exporter\)](https://github.com/SUSE/sap_host_exporter).

EXPORTED METRICS (FOR SAP S/4HANA, SAP NETWEAVER, OR SAP HANA HOSTS)

- SAP start service process list
- SAP enqueue server metrics
- SAP application server dispatcher metrics
- SAP internal alerts

6.1 Installation

The exporter can be installed using the [RPM](#) from the [Open Build Service \(https://build.opensuse.org/package/show/server:monitoring/prometheus-sap_host_exporter\)](https://build.opensuse.org/package/show/server:monitoring/prometheus-sap_host_exporter). To do this, use the following commands (replace `VERSION` with the exact version of the target system, for example: `SLE_15_SP4`):

```
# zypper addrepo https://download.opensuse.org/repositories/server:/monitoring/VERSION/server:monitoring.repo
# zypper install prometheus-sap_host_exporter
```

The RPM package ships with an accompanying `systemd` unit file that can be enabled using the following command:

```
systemctl --now enable prometheus-sap_host_exporter
```

Alternatively, you can compile and install the exporter from the source code:

```
> git clone https://github.com/SUSE/sap_host_exporter
> cd sap_host_exporter
> make
> sudo make install
```



Note

You can deploy a full SAP NetWeaver cluster via Terraform using [SUSE automated SAP/HA deployments \(https://github.com/SUSE/ha-sap-terraform-deployments\)](https://github.com/SUSE/ha-sap-terraform-deployments). This automatically installs and configures the exporter and the entire Prometheus monitoring stack.

6.2 Using SAP Host exporter

You can run the exporter using the following command:

```
./sap_host_exporter --sap-control-url SAP_HOST:SAP_CONTROL_PORT
```

While it is possible to use the exporter to retrieve data from a remote SAP host, it is recommended to run the exporter locally on the target SAP host, and then connect to the SAPControl Web service via Unix domain sockets:

```
./sap_host_exporter --sap-control-uds /tmp/.sapstream50013
```

By default, the exporter exposes the metrics in the `/metrics` path, on port 9680.

6.3 Configuring SAP Host exporter

You can provide optional runtime parameters using command-line flags or via a configuration file. Run the `sap_host_exporter --help` to view the available flags.

The exporter expects to find the `sap_host_exporter.yaml`, `sap_host_exporter.json`, or `sap_host_exporter.toml` file either in the current working directory, or in one of the following directories: `$HOME/.config`, `/etc`, `/usr/etc`. The first match has precedence, and the command-line flags have precedence over the configuration file.

An example YAML configuration file ships with the RPM package. You can also find the example configuration file in the project's [GitHub repository \(https://github.com/SUSE/sap_host_exporter\)](https://github.com/SUSE/sap_host_exporter).

7 Additional information

- SUSE Manager (<https://documentation.suse.com/suma/4.0/>) ↗
- Grafana Dashboards (<https://grafana.com/grafana/dashboards?orderBy=name&direction=asc>) ↗
- Some `.md` files are included in RPM packages. They contain documentation from upstream sources. This can be helpful in isolated data centers without Internet connection.
- SUSE & Microsoft collaborates to provide SAP monitoring (<https://techcommunity.microsoft.com/t5/running-sap-applications-on-the/suse-amp-microsoft-collaborates-to-provide-sap-monitoring/ba-p/1571926>) ↗

8 Legal notice

Copyright© 2006– 2024 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see <https://www.suse.com/company/legal/> ↗. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

A GNU licenses

This appendix contains the GNU Free Documentation License version 1.2.

GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles. You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <https://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled "GNU
Free Documentation License".
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.