

[-n SKIP_PLUGINS]
[--nodes NODES]
[--no-pkg-check]
[--no-local]
[--primary PRIMARY]
[--image IMAGE]
[--force-pull-image TOGGLE, --pull TOGGLE]
[--registry-user USER]
[--registry-password PASSWORD]
[--registry-authfile FILE]
[-o ONLY_PLUGINS]
[-p SSH_PORT]
[--password]
[--password-per-node]
[--preset PRESET]
[--skip-commands COMMANDS]
[--skip-files FILES]
[-s|--sysroot SYSROOT]
[--ssh-user SSH_USER]
[-t|--threads THREADS]
[--timeout TIMEOUT]
[--transport TRANSPORT]
[--tmp-dir TMP_DIR]
[-v|--verbose]
[--verify]
[-z|--compression-type COMPRESSION_TYPE]

DESCRIPTION

`collect` is an `sos` subcommand to collect `sosreports` from multiple nodes and package them in a single useful tar archive.

`sos collect` can be run either on a workstation that has SSH key authentication setup for the nodes in a given cluster, or from a "primary" node in a cluster that has SSH keys configured for the other nodes.

Some `sosreport` options are supported by `sos-collect` and are passed `di?`

rectly to the sosreport command run on each node.

OPTIONS

-a, --alloptions

Enables all sosreport options.

This does NOT enable all sos collect options.

-b, --become

Become the root user on the remote node when connecting as a non-root user.

--batch

Run in non-interactive mode. This will skip prompts for user input, with the exception of a prompt for the SSH password.

--all-logs

Sosreport option. Collects all logs regardless of size.

Default: no

-c CLUSTER_OPTIONS

Specify options used by cluster profiles. The format is 'profile.file.option_name=value'.

For example, for the ovirt plugin if you wanted to restrict node enumeration to a specific cluster you would use '-c ovirt.cluster=example_cluster'.

Available cluster options can be listed by running 'sos collect -l'.

--chroot CHROOT

Sosreport option. Set the chroot mode. When --sysroot is used commands default to executing with SYSROOT as the root directory. This can be overridden by setting --chroot to "always" (always chroot) or "never" (always run in the host namespace).

--case-id CASE_ID

Sosreport option. Specifies a case number identifier.

--cluster-type CLUSTER_TYPE

When run by itself, sos collect will attempt to identify the type of cluster at play. This is done by checking package or configuration information against the localhost, or the primary

node if "--primary" is supplied.

Setting `--cluster-type` skips this step and forcibly sets a particular profile.

Using a value of `none` or `jbon` (just a bunch of nodes) will effectively disable all cluster-specific checks, and cause `sos collect` to only use the nodes specified by the `--nodes` option. Note that in this scenario, regex string(s) for node names will be ignored.

Example: `sos collect --cluster-type=kubernetes` will force the `kubernetes` profile to be run, and thus set `sosreport` options and attempt to determine a list of nodes using that profile.

`--container-runtime` RUNTIME

`sos report` option. Using this with `collect` will pass this option thru to nodes with `sos` version 4.3 or later. This option controls the default container runtime plugins will use for collections. See `man sos-report`.

`-e` ENABLE_PLUGINS, `--enable-plugins` ENABLE_PLUGINS

`Sosreport` option. Use this to enable a plugin that would otherwise not be run.

This option supports providing a comma-delimited list of plugins.

`--encrypt-key` KEY

Encrypts the resulting archive that `sos collect` produces using GPG. KEY must be an existing key in the user's keyring as GPG does not allow for keyfiles. KEY can be any value accepted by `gpg`'s 'recipient' option.

Note that the user running `sos collect` must match the user owning the keyring from which keys will be obtained. In particular this means that if `sudo` is used to run `sos collect`, the keyring must also be set up using `sudo` (or direct shell access to the account).

Users should be aware that encrypting the final archive will result in `sos` using double the amount of temporary disk space -

the encrypted archive must be written as a separate, rather than replacement, file within the temp directory that sos writes the archive to. However, since the encrypted archive will be the same size as the original archive, there is no additional space consumption once the temporary directory is removed at the end of execution.

This means that only the encrypted archive is present on disk after sos finishes running.

If encryption fails for any reason, the original unencrypted archive is preserved instead.

IMPORTANT: As of this version of sos collect, only the final archive on the local machine running sos collect will be encrypted. The individual sos reports that are collected on the nodes will be collected unencrypted.

`--encrypt-pass PASS`

The same as `--encrypt-key`, but use the provided PASS for symmetric encryption rather than key-pair encryption.

`--group GROUP`

Specify an existing host group definition to use.

Host groups are pre-defined settings for the cluster-type, primary node, and nodes options saved in JSON-formatted files under `/var/lib/sos collect/<GROUP>`.

If `cluster_type` and/or `primary` are set in the group, sos collect behaves as if these values were specified on the command-line.

If `nodes` is defined, sos collect extends the `--nodes` option, if set, with the nodes or regexes listed in the group.

Note that sos collect will only write group definitions to `/var/lib/sos collect/` however the GROUP value may be a filename for any group definitions that exist outside of the default location. If you are manually writing these files, use the value `null` when a python `NoneType` is expected. Caveat: use string `'none'` if setting `cluster_type` to `none`.

`--save-group GROUP`

Save the results of this run of sos collect to a host group def?

initiation.

sos-collector will write a JSON-formatted file with name GROUP

to /var/lib/sos collect/ with the settings for cluster-type,

primary, and the node list as discovered by cluster enumeration.

Note that this means regexes are not directly saved to host

groups, but the results of matching against those regexes are.

-j JOBS --jobs JOBS

Specify the number of concurrent node collections that should be

run.

If the number of nodes enumerated exceeds the number of JOBS,

then sos collect will start collecting from the first X number

of nodes and then continue to iterate through the remaining

nodes as sosreport collection finishes.

Defaults to 4.

--nopasswd-sudo

Use this option when connecting as a non-root user that has

passwordless sudo configured.

If this option is omitted and a bogus sudo password is supplied,

collection of sosreports may exhibit unexpected behavior and/or

fail entirely.

-k PLUGIN_OPTION, --plugin-option PLUGIN_OPTION

Sosreport option. Set a plugin option to a particular value.

This takes the form of plugin_name.option_name=value.

Example: To enable the kubernetes "all" option in sosreport use

-k kubernetes.all=on.

--label LABEL

Specify a label to be added to the archive names. This label

will be applied to both the sos collect archive and the sosre?

port archives.

If a cluster sets a default label, the user-provided label will

be appended to that cluster default.

--log-size SIZE

Places a limit on the size of collected logs and output in MiB.

Note that this causes sos to capture the last X amount of the file or command output collected.

By default, this is set to 25 MiB and applies to all files and command output collected with the exception of journal collections, which are limited to 100 MiB.

Setting this value to 0 removes all size limitations, and any files or commands collected will be collected in their entirety, which may drastically increase the size of the final sos report tarball and the memory usage of sos during collection of commands, such as very large journals that may be several GiB in size.

`-n SKIP_PLUGINS, --skip-plugins SKIP_PLUGINS`

Sosreport option. Disable (skip) a particular plugin that would otherwise run. This is useful if a particular plugin is prone to hanging for one reason or another.

This option supports providing a comma-delimited list of plugins.

`--nodes NODES`

Provide a comma-delimited list of nodes to collect sosreports from, or a regex string to be used to compare discovered node names against. If using a regex, only nodes matching the regex will be used - i.e. it can be used as a whitelist but not a blacklist.

This option can be handed multiple regex strings separated by commas. Additionally, both whole node names/addresses and regex strings may be provided at the same time.

`--no-pkg-check`

Do not perform package checks. Most cluster profiles check against installed packages to determine if the cluster profile should be applied or not.

Use this with `--cluster-type` if there are rpm or apt issues on the primary/local node.

--no-local

Do not collect a sosreport from the local system.

If --primary is not supplied, it is assumed that the host running sosreport is part of the cluster that is to be collected.

Use this option to skip collection of a local sosreport.

This option is NOT needed if --primary is provided.

--primary PRIMARY

Specify a primary node IP address or hostname for the cluster.

If provided, then sos collect will check the primary node, not localhost, for determining the type of cluster in use.

--image IMAGE

Specify an image to use for the temporary container created for collections on containerized host, if you do not want to use the default image specified by the host's policy. Note that this should include the registry.

--force-pull-image TOGGLE, --pull TOGGLE

When collecting an sos report from a containerized host, force the host to always pull the specified image, even if that image already exists on the host. This is useful to ensure that the latest version of that image is always in use. Disabling this option will use whatever version of the image is present on the node, and only attempt a pull if there is no copy of the image present at all.

Enable with true/on/yes or disable with false/off/no

Default: true

--registry-user USER

Specify the username to authenticate to the registry with in order to pull the container image

--registry-password PASSWORD

Specify the password to authenticate to the registry with in order to pull the container image. If no password is required, leave this blank.

--registry-authfile FILE

Specify the filename to use for providing authentication credentials to the registry to pull the container image.

Note that this file must exist on the node(s) performing the pull operations, not the node from which sos collect was run.

`-o ONLY_PLUGINS, --only-plugins ONLY_PLUGINS`

Sosreport option. Run ONLY the plugins listed.

Note that a cluster profile will NOT override this option. This may cause the sosreports generated to not contain the relevant output for a given type of cluster.

This option supports providing a comma-delimited list of plugins.

`--password`

Specifying this option will cause sos collect to prompt the user for an SSH password that will be used to connect to all nodes.

If you have differing passwords for the same user across cluster nodes, you should ideally deploy SSH keys, but the `--password-per-node` option is also available.

`--password-per-node`

When using this option, sos collect will prompt the user for the SSH password for each node that will have an sosreport collected from it individually before attempting to connect to the nodes.

`--preset PRESET`

Specify a sos preset to use, note that this requires sos-3.6 or later to be installed on the node. The given preset must also exist on the remote node - local presets are not used.

If `--preset` is specified and a given node either does not have that preset defined, or has a version of sos prior to 3.6, this option is ignored for that node.

`-p SSH_PORT, --ssh-port SSH_PORT`

Specify SSH port for all nodes. Use this if SSH runs on any port other than 22.

`--skip-commands COMMANDS`

A comma delimited list of commands to skip execution of, but

still allowing the rest of the plugin that calls the command to run. This will generally need to be some form of UNIX shell-style wildcard matching. For example, using a value of hostname will skip only that single command, while using hostname* will skip all commands with names that begin with the string "hostname".

--skip-files FILES

A comma delimited list of files or filepath wildcard matches to skip collection of. Values may either be exact filepaths or paths using UNIX shell-style wildcards, for example /etc/sos/*.

--ssh-user SSH_USER

Specify an SSH user for sos collect to connect to nodes with. Default is root.

sos collect will prompt for a sudo password for non-root users.

-s SYSROOT, --sysroot SYSROOT

Sosreport option. Specify an alternate root file system path.

-t THREADS --threads THREADS

Report option. Specify the number of collection threads to run.

The report process on each node will run THREADS number of plugins concurrently during the collection process.

Defaults to 4.

--timeout TIMEOUT

Timeout for sosreport generation on each node, in seconds.

Note that sosreports are collected in parallel, so you can approximate the total runtime of sos collect via timeout*(number of nodes/jobs).

Default is 180 seconds.

--transport TRANSPORT

Specify the type of remote transport to use to manage connections to remote nodes.

sos collect uses locally installed binaries to connect to and interact with remote nodes, instead of directly establishing those connections. By default, OpenSSH's ControlPersist feature

is preferred, however certain cluster types may have preferences of their own for how remote sessions should be established.

The types of transports supported are currently as follows:

`auto` Allow the cluster type to determine the transport used

`control_persist` Use OpenSSH's ControlPersist feature.

This is the default behavior

`oc` Use a locally configured oc binary to deploy collection pods on OCP nodes

`--tmp-dir TMP_DIR`

Specify a temporary directory to save sos archives to. By default one will be created in /tmp and then removed after sos collect has finished running.

This is NOT the same as specifying a temporary directory for sosreport on the remote nodes.

`-v --verbose`

Print debug information to screen.

`--verify`

Sosreport option. Passes the "--verify" option to sosreport on the nodes which causes sosreport to validate plugin-specific data during collection.

Note that this option may considerably extend the time it takes sosreport to run on the nodes. Consider increasing --timeout when using this option.

`-z COMPRESSION, --compression-type COMPRESSION`

Sosreport option. Override the default compression type.

SEE ALSO

`sos(1)` `sos-report(1)` `sos-clean(1)` `sos.conf(5)`

MAINTAINER

Jake Hunsaker <jhunsake@redhat.com>

April 2020

1

SOS(COLLECT)