systemd-oomd.service

- Description
- System requirements and configuration
- Usage Recommendations
- Options
- Signals
- See Also

systemd

systemd-oomd.service

8

systemd-oomd.service

systemd-oomd

A userspace out-of-memory (OOM) killer

`systemd-oomd.service`

`/usr/lib/systemd/systemd-oomd`

Description

`systemd-oomd` is a system service that uses cgroups-v2 and pressure stall information (PSI) to monitor and take corrective action before an OOM occurs in the kernel space.

You can enable monitoring and actions on units by setting `ManagedOOMSwap=` and `ManagedOOMMemoryPressure=` in the unit configuration, see `systemd.resource-control(5)`. `systemd-oomd` retrieves information about such units from `systemd(1)` when it starts and watches for subsequent changes.

Cgroups of units with `ManagedOOMSwap=` or `ManagedOOMMemoryPressure=` set to `kill` will be monitored. `systemd-oomd` periodically polls PSI statistics for the system and those cgroups to decide when to take action. If the configured limits are exceeded, `systemd-oomd` will select a cgroup to terminate, and send `SIGKILL` to all processes in it. Note that only descendant cgroups are eligible candidates for killing; the unit with its property set to `kill` is not a candidate (unless one of its ancestors set their property to `kill`). Also only leaf cgroups and cgroups with `memory.oom.group` set to `1` are eligible candidates; see `OOMPolicy=` in `systemd.service(5)`.

`oomctl(1)` can be used to list monitored cgroups and pressure information.

See `oomd.conf(5)` for more information about the configuration of this service.

See `org.freedesktop.oom1(5)` and `org.freedesktop.LogControl1(5)` for a description of the D-Bus API.

System requirements and configuration

The system must be running systemd with a full unified cgroup hierarchy for the expected cgroups-v2 features. Furthermore, memory accounting must be turned on for all units monitored by `systemd-oomd`. The easiest way to turn on memory accounting is by ensuring the value for `DefaultMemoryAccounting=` is set to `true` in `systemd-system.conf(5)`.

The kernel must be compiled with PSI support. This is available in Linux 4.20 and above.

It is highly recommended for the system to have swap enabled for `systemd-oomd` to function optimally. With swap enabled, the system spends enough time swapping pages to let `systemd-oomd` react. Without swap, the system enters a livelocked state much more quickly and may prevent `systemd-oomd` from responding in a reasonable amount of time. See "In defence of swap: common misconceptions" for more details on swap. Any swap-based actions on systems without swap will be ignored. While `systemd-oomd` can perform pressure-based actions on such a system, the pressure increases will be more abrupt and may require more tuning to get the desired thresholds and behavior.

Be aware that if you intend to enable monitoring and actions on `user.slice`, `user-$UID.slice`, or their ancestor cgroups, it is highly recommended that your programs be managed by the systemd user manager to prevent running too many processes under the same session scope (and thus avoid a situation where memory intensive tasks trigger `systemd-oomd` to kill everything under the cgroup). If you're using a desktop environment like GNOME or KDE, it already spawns many session components with the systemd user manager.

Usage Recommendations

`ManagedOOMSwap=` works with the system-wide swap values, so setting it on the root slice `-.slice`, and allowing all descendant cgroups to be eligible candidates may make the most sense.

`ManagedOOMMemoryPressure=` tends to work better on the cgroups below the root slice. For units which tend to have processes that are less latency sensitive (e.g. `system.slice`), a higher limit like the default of 60% may be acceptable, as those processes can usually ride out slowdowns caused by lack of memory without serious consequences. However, something like `user@$UID.service` may prefer a much lower value like 40%.

Options

`-n`; `--dry-run`
Do a dry run of `systemd-oomd`: when a kill is triggered, print it to the log instead of killing the cgroup. The short option `-n` was added in version 262.

Signals

`SIGHUP`
Upon reception of the `SIGHUP` process signal `systemd-oomd` will reload its configuration file.

See Also

`systemd(1)`, `systemd-system.conf(5)`, `systemd.resource-control(5)`, `oomd.conf(5)`, `oomctl(1)`, `org.freedesktop.oom1(5)`

□□□ #2
□ □□□ □□□ 2026-10-31 15:00:56 UTC
□ □□□ □□□ 2026-10-31 15:00:57 UTC