

Specification of the storage-related journal entries

This is a specification of the *journal* entries reporting events and actions related to the file systems and storage. For some background and rationale please read the two blog posts published in Spring 2017: [Reporting And Monitoring Storage Events](#) and [Reporting And Monitoring Storage Actions](#).

Dictionary

device: a block device

subsystem: a storage subsystem/technology like LVM, MD RAID, btrfs, SMART,...

Specifications

There are two kinds of entries (each detailed in the respective blog post):

1. device discoveries, failures or recoveries, i.e. *state changes*,
2. *actions* performed on devices.

1. State changes

All the fields described below are required unless explicitly marked as **optional**.

MESSAGE_ID **Must have the value** 3183267b90074a4595e91daef0e01462.

This identifies the entry as a storage state change report.

DEVICE Complete name of the device that changed state.

Usually this is the path of one of the device's symlinks without the `/dev` prefix (e.g. `vgname/lvname` for an LVM LV, or `md/raidname` for an MD RAID device with a name). However, some devices don't have pretty/persistent names in which case the name of the device file in `/dev/` should be used.

See the note about messages from kernel describing an exception to the above.

DEVICE_ID (optional) Unique and persistent identifier of the device. Either UUID, WWN/WWID, SERIAL or some other similar value. **If some kind of such an identifier exists, it must be specified.** It is only optional for devices that don't have any such identifier.

STATE The new state of the device. If possible, the value should be one of the following well-known states: `discovered`, `initialized`, `failing`, `degraded`, `failed`, `missing`. If none of those are applicable, the value should describe the new state of the device, preferably with a single word.

SOURCE The subsystem reporting the state change.

SOURCE_MAN (optional) Manual page providing extra information about the *SOURCE* in the `MAN_NAME (MAN_SECTION)` format, for example `smartd(8)`.

DETAILS Extra information about what happened and/or about the new state of the device. For example if a device is failing this should describe the symptoms like `too many sector reallocations` or similar.

PRIORITY An integer priority value between 0 (emergency) and 7 (debug) formatted as a decimal string. This field is compatible with syslog's priority concept and describes the importance of the state change.

PRIORITY_DESC A one-word description of *PRIORITY*. Possible values are: `emergency, alert, critical, error, warning, notice, info, debug`.

MESSAGE An arbitrary free-form message describing what happened.

Note

If the journal entry comes from the kernel, the `_KERNEL_` prefix is added to some of the above fields, namely: `DEVICE`, `DEVICE_ID`, `SOURCE`, `SOURCE_MAN`, `DETAILS`, `PRIORITY_DESC` due to how journald handles structured logging from kernel.

Moreover, the `_KERNEL_DEVICE` field is very different, because kernel uses a special form of the device identifier of the form `+SUBSYSTEM:SYSNAME`, for example `+scsi:0:0:0:0` for the first SATA/SCSI disk in the system which is usually known as `sda`.

2. Actions

TBD