systemd journal or computer readable logs

Maksim 'max_posedon' Melnikau

Linux Mobile hobbyist
World of Tanks developer

June 9, 2012

the question

Who reads logs?

generating logs

- part of context information is lost
- lossy human language
- not useful for early boot and shutdown

reading logs

- no structured format is defined
- parsing and processing is messy
- many key log operations have a complexity of O(n)
- binary data cannot be logged
- security questions

Linux logs

- syslog*
- utmp/wtmp
- kernel logs
- firmware logs
- multiple application-specific log formats
- /dev/null

the log item

- split message for human/for machine
- UUID as message type/id
- "business" data as key-value dict
- attach all system information

example.c

```
#include <stdlib.h>
#include <systemd/sd-journal.h>
int main() {
    int i=0, j=0;
    while(i<10 && j<10) {
        random()&1 ? ++i : ++j;
        sd_journal_send(
            "MESSAGE=(%d,%d)", i, j,
            "MESSAGE_ID=51141ddad48f4924aef970b1eab2af42",
            "I=%d", i,
            "J=%d", j,
            NULL
        );
    }
    return 0;
}
```

systemd/sd-journal.h

```
int sd_journal_print(int piority, const char *format, ...)
    __attribute__ ((format (printf, 2, 3)));

int sd_journal_printv(
    int priority, const char *format, va_list ap);

int sd_journal_send(const char *format, ...)
    __attribute__ ((sentinel));

int sd_journal_sendv(const struct iovec *iov, int n);

int sd_journal_stream_fd(const char *identifier,
    int priority, int level_prefix);
```

the logs

- different formats
- different data
- efficient search

systemd-journalctl -o export -f J=2 -n 1

```
.realtime=1337895699642000
.monotonic=27122021981
.boot_id=b149fa547201419f93dcfaa7a214dc8b
MESSAGE=(5,2)
MESSAGE_ID=51141ddad48f4924aef970b1eab2af42
I=5
J=2
_TRANSPORT=journal
_PID=9334
_UID=1001
_GID=1001
_COMM=example
_EXE=/home/max_posedon/systemd-journald/example
_CMDLINE=./example
_SYSTEMD_CGROUP=/system/kdm@.service/tty7
_SYSTEMD_UNIT=kdm@tty7.service
_SOURCE_REALTIME_TIMESTAMP=1337895699633622
```

systemd-journalctl _SYSTEMD_UNIT=dbus.service

```
_TRANSPORT=stdout
PRIORITY=6
SYSLOG_FACILITY=3
SYSLOG_IDENTIFIER=dbus-daemon
MESSAGE=**** pci IGNORING ADD \
    /sys/devices/pci0000:00/0000:00:1c.0/0000:07:00.0
_PID=1526
_UID=0
_GID=0
_COMM=dbus-daemon
_EXE=/usr/bin/dbus-daemon
_CMDLINE=/usr/bin/dbus-daemon --system --address=systemd: -
_SYSTEMD_CGROUP=/system/dbus.service
_SYSTEMD_UNIT=dbus.service
_BOOT_ID=f9c7b9b79e584d31a5fe238fe4de16a0
```

the info

- Maksim 'max_posedon' Melnikau maxposedon@gmail.com
- http://www.freedesktop.org/wiki/Software/systemd
- https://docs.google.com/document/pub?id=1IC9yOXj7j6cdLLxWEBAGRL6wl97tFxgjLUEHIX3MSTs