# Redhat journald

## Introduction

Mit dem Systemd wurde auch der Journald eingeführt. Er ist das zentrale Logging-Werkzeug. Die Loggdaten des Kernels und vieler Services gehen zum Journald. So wie der Journald auf RedHat standardmässig konfiguriert ist, werden die Logdaten nur im Memory gehalten. Nach einem Reboot des Systems sind die alten Loggdaten weg.

> Die Logdaten unter “**/var/log**” werden vom “**rsyslogd**” geschrieben. Er holt sich die Informationen vom Journald und speichert sie in persistente Dateien.

## Logging rate limits

Auf Systemen mit grossem Logdatenaufkommen kann es vorkommen, dass in der Logdatei plötzlich Daten fehlen. Das ist darauf zurückzuführen, dass der Journald als auch der Rsyslogd je ein “logging rate limit” gesetzt haben. Bei Mailsystemen ist dieser Effekt aber unerwünscht, da man die Logdaten eines jeden Mails haben will.

**Auf https://www.rootusers.com/how-to-change-log-rate-limiting-in-linux/ wird das Problem und die Lösung folgendermassen beschrieben:**

### How To Change Log Rate Limiting In Linux

By default in Linux there are a few different mechanisms in place that may rate limit logging. These are primarily the systemd journal and rsyslog rate limits that are in place by default. Here we cover modifying or removing rate limiting for logging.

---

### Why Rate Limiting?

Rate limitations on logging are in place to prevent logging from using excessive levels of system resources. To log an event, it needs to be written to disk which uses system resources. If there are too many of these events coming in that need to be recorded to disk they can overwhelm a system and cause more important services to respond slowly or fail. For this reason it is generally not recommended to completely disable rate limiting, but to tweak it as required. At the same time we do not want to drop important messages that may be required to generate a critical alert, so a balance needs to be found.

---

## Systemd Journal Rate Limiting

How do we know if the journal limits are actually causing us to drop log messages? Generally you will see similar messages in the log files as below.

```
Jan  9 09:18:07 server1 journal: Suppressed 7124 messages from
/system.slice/named.service
```

In this particular case we have a DNS server running Bind which is logging all DNS queries. 7124 messages were suppressed and dropped (not logged) because they were coming in too fast in this example. By default systemd allows 1,000 messages within a 30 second period.

The limits are controlled in the `/etc/systemd/journald.conf` file.

```
RateLimitInterval=30s
RateLimitBurst=1000
```

If more messages than the amount specified in RateLimitBurst are received within the time defined by RateLimitInterval, all further messages within the interval are dropped until the interval is over. You can modify these values as you see fit, you can completely disable systemd journal logging rate limiting by setting both to 0. If you make any changes to /etc/systemd/journald.conf you will need to restart the systemd-journald service to apply the changes.

```
# systemctl restart systemd-journald
```

---

## Rsyslog Rate Limiting

The systemd journal limit is hit before any default rsyslog limits as its default limits are smaller. By default rsyslog will accept 20,000 messages within a 10 minute period. Therefore if you increase the rate limiting of the systemd journal logging as shown above you may then start to receive similar messages in your syslog logs as shown below.

```
...
Jan  9 22:42:35 server1 rsyslogd-2177: imjournal: begin to drop messages due
to rate-limiting
Jan  9 22:51:26 server1 rsyslogd-2177: imjournal: 143847 messages lost due
to rate-limiting
...
```

The first message states that messages will be dropped as the limit has been reached, and once the interval is over (after 10 minutes by default) the amount of messages that were lost due to rate limiting will then be logged. The limits are controlled in the `/etc/rsyslog.conf` file.

```
$ModLoad imjournal
$imjournalRatelimitInterval 600
$imjournalRatelimitBurst 20000
```

For further information see the imjournal rsyslog documentation. Again you can modify these values as you like, and they can be completely disabled by setting both to 0. If you make any changes to the /etc/rsyslog.conf file you will need to restart the rsyslog service to apply the changes.

```
# systemctl restart rsyslog
```

---

## Summary

As shown we can check our log files to find out if logs are being dropped due to either systemd journal or syslog rate limits. The systemd journal default rate limit is much lower than the syslog default rate limit so it will be triggered first. Once you increase the rate limiting on the systemd journal logging you may then start to experience additional rate limiting by syslog, which can then also be increased if required.

Last update: **2019/03/07 13:16**