

Limits setzen/überschreiben mit systemd

Wie kann ich Limits für Services setzen, die via systemd gestartet werden? Meine Einstellungen in `/etc/security/limits.conf` oder `/etc/security/limits.d/*.conf` werden ignoriert, da diese nur von `pam_limits.so` verwendet werden, was `systemd` nicht nutzt.

Um die Limit anzupassen muss das systemd unit angepasst werden, z.B. für MySQL:

```
$ systemctl edit mysql.service
# einfügen und speichern:
[Service]
LimitNOFILE=500000

# Service neu starten
$ systemctl restart mysql.service
```

Im `systemctl status` sieht man jetzt den Override

```
$ systemctl status mysql.service
● mysql.service - Percona Server
   Loaded: loaded (/lib/systemd/system/mysql.service; enabled; vendor preset: enabled)
   Drop-In: /etc/systemd/system/mysql.service.d
            └─override.conf
```

Limits für alle Prozesse überschreiben

```
mkdir -p /etc/systemd/system.conf.d/
cat >/etc/systemd/system.conf.d/10-filelimit.conf <<EOF
[Manager]

DefaultLimitNOFILE=500000
EOF
systemctl daemon-reload
## ggf. Reboot!
```

Folgende Limits können überschrieben werden:

Directive	ulimit equivalent	Unit
ulimit -u	ulimit -u	processes
ulimit -m	ulimit -m	memory
ulimit -s	ulimit -s	stack
ulimit -v	ulimit -v	virtual memory
ulimit -f	ulimit -f	file descriptors
ulimit -n	ulimit -n	open files
ulimit -l	ulimit -l	locked memory
ulimit -p	ulimit -p	pipe size
ulimit -r	ulimit -r	real-time priority
ulimit -t	ulimit -t	time limit
ulimit -c	ulimit -c	core file size
ulimit -x	ulimit -x	shared library size
ulimit -y	ulimit -y	shared object size
ulimit -z	ulimit -z	shared object size

Notes

The diagram illustrates the relationship between different time limits on a timeline. A horizontal line represents time, with vertical tick marks indicating specific points. Below the line, labels are placed at these ticks: 'LimitCPU=' at the first tick, 'ulimit -t' at the second tick, 'Seconds' at the third tick, and '-' at the fourth tick. A horizontal line segment is drawn below the main timeline, starting from the origin and ending at the 'ulimit -t' tick, indicating that the ulimit time limit is less than or equal to the Seconds limit.

The diagram consists of a horizontal line with four vertical tick marks. Below the line, the text "LimitFSIZE=" is aligned under the first tick mark, "ulimit -f" under the second, "Bytes" under the third, and "-" under the fourth. A horizontal line segment is drawn below the main line, starting from the left edge and ending at the second tick mark.

LimitDATA=	ulimit -d	Bytes	Don't use.
This limits the allowed address range, not memory			
			use! Defaults
to unlimited and should not be lowered. To limit			
			memory use,
see MemoryMax= in systemd.resource-control(5).			

The diagram consists of two horizontal lines. The top line has four vertical tick marks. The bottom line has one vertical tick mark. Below the top line, the text 'LimitCORE=' is positioned under the first tick mark, 'ulimit -c' is under the second, 'Bytes' is under the third, and '-' is under the fourth. Below the bottom line, a single vertical tick mark is positioned under the 'ulimit -c' label.

Diagram illustrating memory limits:

- LimitRSS=
- ulimit -m
- Bytes
- Don't use. No

effect on Linux.

LimitNOFILE=	ulimit -n	Number of File Descriptors	Don't use. Be
careful when raising the soft limit above 1024,			

			since
select(2) cannot function with file descriptors above			1023 on
Linux. Nowadays, the hard limit defaults to 524288, a			very high
value compared to historical defaults. Typically			applications
should increase their soft limit to the hard			limit on
their own, if they are OK with working with file			descriptors
above 1023, i.e. do not use select(2). Note that			file
descriptors are nowadays accounted like any other form of			memory, thus
there should not be any need to lower the hard			limit. Use
MemoryMax= to control overall service memory use,			including
file descriptor memory.			

	LimitAS=	ulimit -v	Bytes	Don't use.
This limits the allowed address range, not memory				use! Defaults
to unlimited and should not be lowered. To limit				memory use,
see MemoryMax= in systemd.resource-control(5).				

	LimitNPROC=	ulimit -u	Number of Processes	This limit is
enforced based on the number of processes				belonging to
the user. Typically it's better to track				processes per
service, i.e. use TasksMax=, see				
systemd.resource-control(5).				

LimitMEMLOCK=	ulimit -l	Bytes	-
LimitLOCKS=	ulimit -x	Number of Locks	-
LimitSIGPENDING=	ulimit -i	Number of Queued Signals	-
LimitMSGQUEUE=	ulimit -q	Bytes	-
LimitNICE=	ulimit -e	Nice Level	-
LimitRTPRIO=	ulimit -r	Realtime Priority	-
LimitRTTIME=	ulimit -R	Microseconds	-

Weitere Infos in den manpages

```
man 5 systemd.exec
```

```
man 5 systemd.resource-control
```

Revision #2

Created 17 October 2023 07:36:46 by magenbrot

Updated 17 October 2023 07:53:53 by magenbrot