From systemd to journalctl

initializing and inspecting the OS

Canton Linux Enthusiasts Meetup
Thu 2020-08-27 1900 EDT

Sponsor

This presentation was sponsored by

Skedzy provides and develops tools that give clients the power to manage and administer open source, Linux based infrastructures. As a result, clients enjoy the benefits of secure, reliable, high performance, cost-effective architectures matched to their current needs as well as future peak demands.

info@skedzy.com https://skedzy.com 330.754.2004

Two daemons + two commands

daemon	systemd	journald
command	systemctl	journalctl

What is systemd?

systemd is a system and service manager for Linux operating systems. When run as first process on boot (as PID 1), it acts as init system that brings up and maintains userspace services. Separate instances are started for logged-in users to start their services.

– man systemd(1) Fedora 32 systemd 245

set point for this presentation

This presentation is largely based upon inspection and study of

- systemd 245
- running on Fedora 32

systemd timeline

1970s origin of Unix 1983 origin of GNU 1991 origin of Linux

2009-11 first commit to systemd repo by Lennart Poettering

2011-05 Fedora 15 adopted systemd

2012-09 openSUSE 12.2 adopted systemd

2014-04 centOS 7.14.04 adopted systemd

2014-06 Red Hat Enterprise Linux 7.0 adopted systemd

2014-10 SUSE Linux Enterprise Server 12 adopted systemd

2015-04 Debian 8, Ubuntu 15.05 adopted systemd

– Benno Rice, *The tragedy of systemd,* 2019

systemd invocation

- system manager (init, PID=1)
- user manager (one per user, after log in)
- initial RAM disk (detects /etc/initrd-release)

systemd invocation (2)

- see man bootup(7)
- overview of bootup order
- parallel start up paths

systemd concepts

- a dependency system between various entities
- the entities are called “units”
- there are 11 basic unit types (man systemd.unit)
- each unit has a systemd man page

  automount *device* *mount* *path* *scope* *service*

  slice *socket* *swap* *target* *timer*

unit states

- *activating*
- activated
- *deactivating*
- deactivated
- failed

  there may be sub-states, though all map to the above five states

systemd architecture

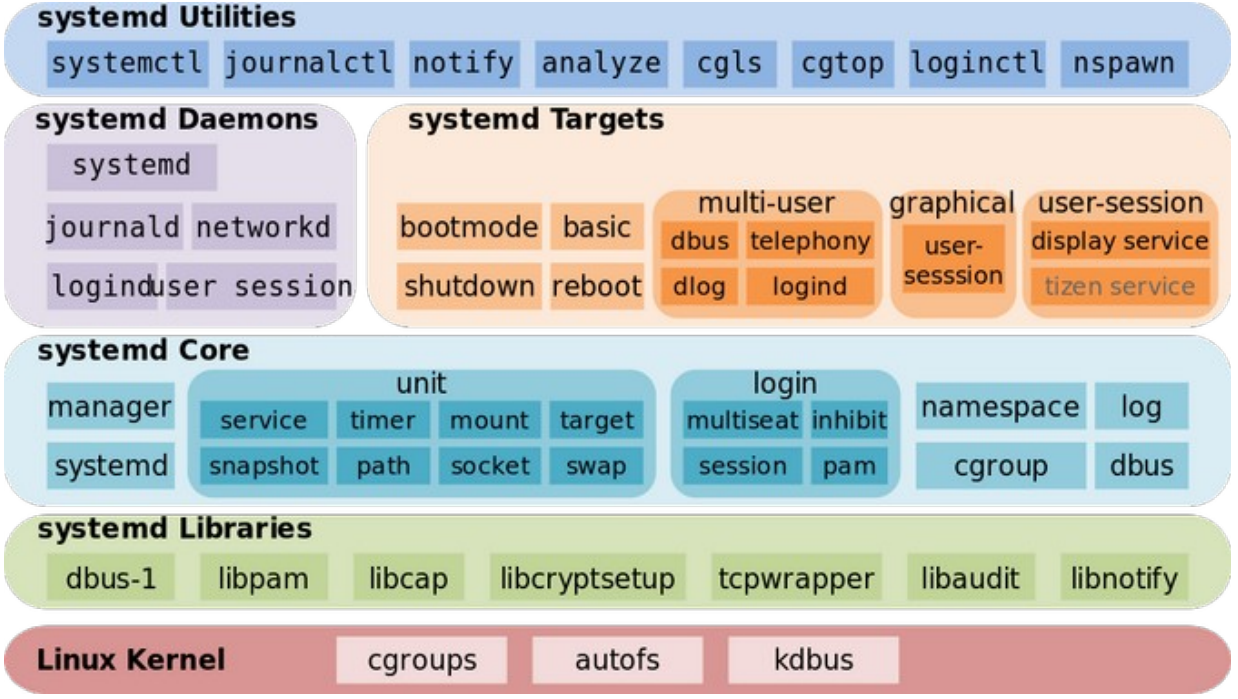

Shmuel Csaba Otto Traian (CC BY-SA 3.0)

unit configuration files

- /usr/lib/systemd/system (default, system)
- /usr/lib/systemd/user (default, user)
- man systemd.unit(5) – unit configuration

system mode unit search paths

$SYSTEMD_UNIT_PATH
/etc/systemd/system.control/*
/run/systemd/system.control/*
/run/systemd/transient/*
/run/systemd/generator.early/*
/etc/systemd/system/*
/etc/systemd/systemd.attached/*
/run/systemd/system/*
/run/systemd/systemd.attached/*
/run/systemd/generator/*
/usr/local/lib/systemd/system/*
/usr/lib/systemd/system/*
/run/systemd/generator.late/*

user mode unit search path

~/.config/systemd/user.control/*
$XDG_RUNTIME_DIR/systemd/user.control/*
$XDG_RUNTIME_DIR/systemd/transient/*
$XDG_RUNTIME_DIR/systemd/generator.early/*
/run/systemd/transient/*
/run/systemd/generator.early/*
~/.config/systemd/user/*
/etc/systemd/user/*
$XDG_RUNTIME_DIR/systemd/user/*
/run/systemd/user/*
$XDG_RUNTIME_DIR/systemd/generator/*
~/.local/share/systemd/user/*
/usr/local/lib/systemd/user/*
/usr/lib/systemd/user/*
$XDG_RUNTIME_DIR/systemd/generator.late/*

$XDG_CONFIG_HOME used if set
otherwise ~/.config is used

$XDG_DATA_HOME
$XDG_DATA_DIRS

actual search paths

- systemd-analyze --system unit-paths
- systemd-analyze --user unit-paths
- systemctl link … creates symlinks from additional directories

unit file format

- text files similar in format to Windows init files
- sections: [SECTION_LABEL]
- properties (also called directives): PROP=value
- comments: #
- man systemd.directives lists 362 directives

unit file sections

- man pages for different unit types articulate the sections for that type, as well as the directives for that section
- man systemd.service: [Service]
- man systemd.unit: [Unit] [Install]
- man systemd.network: [Match] [Link] ...

unit file section [Unit]

Properties (also called directives)

- Description=
- Documentation=
- Before= After=
- Requires= RequiredBy=
- Wants= WantedBy=

unit file section [Service]

- Type=
- BusName= (a D-Bus name)
- ExecStart=
- ExecStartPre= ExecStartPost=
- ExecCondition=
- ExecReload=

unit file section [Timer]

- AccuracySec=
- OnActiveSec=
- OnBootSec=
- OnStartupSec=
- OnUnitActiveSec=
- OnUnitInactiveSec=

in depth: logrotate service, timer

- systemctl list-timers
- systemctl show logrotate.service
- systemctl show logrotate.timer
- systemctl show logrotate.target

special unit files

- man systemd.special
- special units are handled internally by systemd, hence their unit files have no Exec properties

  basic.target bluetooth.target blockdev@.target boot-complete.target getty.target graphical.target halt.target hibernate.target machines.target network.target swap.target system.slice time-set.target time-sync.target dbus.service dbus.socket

cgroups – Linux control groups

Cgroups are a Linux kernel feature which allow processes to be organized into hierarchical groups whose usage of various types of resources can then be limited and monitored.

– man cgroups(7)

cgroups version 2

- a unified hierarchy against which all controllers are mounted
- “internal” processes are not permitted
- active groups specified via files: cgroup.controllers and cgroup.subtree_controller

show cgroup subsystems

- cat /proc/cgroups
- cat /proc/PID_NUM/cgroup

systemd vulnerabilities

- CVE-2018-16864
- CVE-2018-16865
- CVE-2018-16866

  could allow unprivileged local attackers or malicious programs to gain root access

CVE-2018-16864,5,6

- memory corruption issues (64,65)
- out-of-bounds read systemd-journald (66)

CVE-2018-16864,5,6

Linux distros not affected by the flaws because their userspace code is compiled with GCC’s fstack clash protection

- SUSE Linux Enterprise 15
- openSUSE Leap 15.0
- Fedora 28 and 29

CVE-2018-16864

- existed in systemd’s codebase since April 2013 (systemd 203) and became exploitable in Feb 2016 (systemd 230)

– Qualys

CVE-2018-16865

- was introduced in Dec 2011 (systemd 38) and became exploitable in Apr 2013 (systemd 201)

  – Qualys

CVE-2018-16866

- was introduced in Jun 2015 (systemd 221) but was inadvertently fixed in Aug 2018

https://cve.mitre.org

https://thehackernews.com/2019/10/linux-systemd-exploit.html

Wang Wei, 2019-01-10

systemctl

- command to query and manager systemd
- systemctl list-units (only active units)
- systemctl show UNIT_NAME
- systemctl list-timers
- systemctl status UNIT_NAME
- systemctl start UNIT_NAME
- systemctl enable UNIT_NAME

systemctl (user)

- systemctl –user list-units

journald

- the systemd-journald.service maintains a journal of systemd activity
- output goes to binary files
- tamper resistant
- automatic file switching
- automatic purge of old journals

journald log files

- one chain of log files for system side
- another chain of log files for each user

journald log file retention

- /etc/systemd/journald.conf
- SystemMaxUse=
- back up inactive files from /var/log/journal/{128_BIT_ID}
- active files

  system.journal user-1000.journal

journalctl – query the systemd logs

- journactl -b -o short-iso-precise

  show all logs from the current boot session with microsecond precision

- journactl -b -1 -o short-iso-precise

  show all logs from previous boot session

Supplementary

Resources for this presentation

plus presentation slides

see

https://skedzy.com > blog > systemd resources

Questions
Comments
Discussion