

Evolution of Linux networking infrastructure

Pavel Šimerda
pavlix@pavlix.net

LinuxDays 2012, Prague

<http://data.pavlix.net/linuxdays/2012/>

- Kernel configures the loopback and link-local addresses
- The `ip` command can tweak network interface configuration
- Suitable for testing
- Usable for very simple and very static environments

Script-based static networking

- A tiny layer over iproute2 or ifconfig
- Various configuration file formats
- Possibility to activate/deactivate individual connections
- Suitable only for manual (aka static) addressing
- Very limited support for dynamic configuration

Script-based static networking – examples

- ifcfg variants (Fedora, openSUSE, Mandriva, ...)
- ifupdown (Debian, Ubuntu, ...)
- ifnet (Gentoo, ...)
- uci network (OpenWRT)
- ...

Consequences of misusing DHCP clients as network daemons

- You cannot properly configure two interfaces using DHCP
- Interoperability with other network tools is broken
- There's no central point to make policy decisions
- The same applies to all IPv6 dynamic configuration methods
- VPN software packages are also affected

Network configuration daemon

- Provides a central coordination point
- Can aggregate much of the configuration
- Communicates with other tools over IPC

Everyone wants one

- NetworkManager (used by most distributions)
- connman (came from Intel's Meego project)
- Wicd (a network daemon written in Python)
- netcfg (Archlinux)
- netifd (new OpenWRT development)
- wicked (is not Wicd)
- ...?

NetworkManager, the Fedora's network daemon of choice

- Long history of desktop and especially laptop usage
- Works well with IPv6 (use 0.9.6 or later)
- Is under active development
- Influences kernel development as well as wpa_supplicant
- May influence IETF standards
- Users who avoid NM still benefit from our work

- Was primarily designed for desktop usecases
- Lacks some important server features
- Has a history of bad regressions

Platform interaction redesign

- The whole platform interaction code is being rewritten
- Hacks for kernel bugs will be in one place
- Compatibility hacks for libnl versions likewise
- The platform interaction code will be testable
- Fake platform code will be used to test NM behavior
- Most regressions will be fixed before a release

Connection configuration redesign

- NetworkManager will support runtime configuration
- APIs will be created to save/restore configurations
- On-disk configuration will be also manipulated through API
- IP configuration on unknown devices will be supported
- Modifications from other tools will be acted upon
- Take-over of connections will be extended and fixed

The big picture of Linux networking

- DNSSEC support is still being evaluated
- Kernel ignores RA router preference
- Kernel breaks TCP on preferred address expiry
- RA DNS information is not cached by kernel
- GLIBC currently breaks both IPv4 and IPv6 in specific cases
- Avahi still doesn't support IPv6 link-local addresses
- IPv6 dynamic reconfiguration is not fully supported in NM
- Classic ping command still can't handle IPv6

NetworkManager target use cases

- Desktop and laptop (we're not abandoning you!)
- Simple server (especially command-line interface)
- Virtualization host (the most complicated)
- Network-enabled initramfs (e.g. for NFS or iSCSI boot)
- Simple routing use cases (e.g. connection sharing)
- Controlled environment (NM used as an API)

We ask the community to...

- Follow development news on NetworkManager
- Use NetworkManager for the new use cases
- Report bugs to the project bugzilla
- Submit patches and help with development
- Avoid writing new and new network daemons
- We can work together, can't we?

<http://data.pavlix.net/linuxalt2011/>

psimerda@redhat.com
pavlix@pavlix.net