

NetworkManager on servers

Pavel Šimerda
pavlix@pavlix.net

42. konference EurOpen.CZ, Třešť

<http://data.pavlix.net/euopen/42/>

Server use case

Alternatives

Features

...

Questions

Would you use NetworkManager on servers?

Server use case

Alternatives

Features

...

Questions

Why would you want to do that?

Server use case

Alternatives

Features

...

Questions

Unified network configuration API

- Desktops and laptops
- Servers and virtualization
- Embedded and initramfs
- Multipurpose systems
- Event-based API
- Runtime and permanent configuration (devel)

Imagine a development laptop with VPN and virtualized servers, connected to various networks with autoconfiguration.

Independent network interface configuration

- Each interface configuration is handled separately
- Configuration files and runtime information
- Data from configuration protocols – DHCP and RA
- All information sources taken into account
- LAN connectivity and split DNS
- Default route and default DNS

Multi-interface dynamic configuration needs a coordination point.

Alternatives?

Network configuration scripts

- ifcfg variants – Fedora, openSUSE, Mandriva
- ifupdown – Debian, Ubuntu
- ifnet – Gentoo
- UCI-based network scripts – OpenWRT
- Custom scripts using `iproute` tool

Static-only configuration unless a daemon is involved.

Misused DHCP clients

- Support for multi-interface configuration
- Some configuration daemon features
- Integration suffers
- IPv4-only configuration

Neither fish nor fowl.

Network configuration daemons

- connman – came from Intel's Meego project
- Wicd – a network daemon written in Python
- netcfg – Archlinux
- netifd – OpenWRT
- wicked – is not Wicd

OpenWRT netifd is not ported to other distributions.

Talk is cheap, show me the features!

Supported connection types

- Wired and wireless Ethernet connections (including 802.1x)
- ADSL connections, mobile broadband and bluetooth PAN
- OLPC mesh
- Wimax connections
- Infiniband

- VLAN interfaces
- Bridges and bonds for wired Ethernet
- Team driver integration (wishlist)
- Macvlan, macvtap and more → 0.9.10
- VPN plugin interface + several plugins available

- Generic device support → 0.9.10
- Extensive refactoring, port to nm-platform → 0.9.10

Address and route configuration

- Consolidated IPv4/IPv6 format
- Dynamic IPv4 with DHCP
- Dynamic IPv6 configuration with RD and DHCP
- Removed time-based polling in IPv6 → 0.9.10
- DUID from /etc/machine-id or similar → 0.9.10
- DHCPv6 info-only soft failure → 0.9.10
- Using dhclient \geq 4, stale support for dhcpcd
- Switching from kernel autoconf to libndp (wishlist)

DNS configuration

- resolvconf/netconfig support (to be removed)
- Split DNS using dnsmasq
- Split DNS and DNSSEC using unbound (wishlist)
- Arbitrary scripts for DNS management (wishlist)

Customization (aka server features)

- Configuration autoload optional → 0.9.10
- Autocreated connections optional → 0.9.10
- Default route management optional → 0.9.10
- Carrier detection optional → 0.9.10
- /etc/resolv.conf management optional → 0.9.10

Temporary connections / connection take over

- Separate runtime and persistent configuration (wishlist)
- Save, restore and modification APIs (wishlist)
- Accepting modifications from other tools (wishlist)
- Taking over connections during startup (to be improved)
- Explicit take-over at runtime (wishlist)

Command-line interface and API

- Library and dbus API improvements
- CLI redesign → 0.9.10
- Connection add → 0.9.10
- Guided configuration → 0.9.10
- Proper line between connection and device → 0.9.10
- Extensive documentation (not only for nmcli)

https://bugzilla.gnome.org/show_bug.cgi?id=682056

NetworkManager in distributions

- Integration with network scripts
- Native configuration using keyfiles
- Reads ifcfg, ifupdown and ifnet
- Writes ifcfg and ifupdown
- Special DNS/hostname handling for SUSE
- Service management
- Session management
- Patch-free packaging

Notable dependencies

- Linux \geq 3.1 (3.9 recommended)
- libnl \geq 3.2.8
- glib \geq 2.32
- dbus-glib \geq 0.100 (daemon optional)
- gudev \geq 165
- polkit-gobject \geq 0.97 (optional)
- dhclient \geq 4

Testing strategy

- Code coverate and valgrind integrated to autotools
- New nm-platform layer → 0.9.10
- Kernel and libnl workarounds centralized or removed
- Tests for the nm-platform, kernel and libnl
- Tests for NetworkManager core behavior (wishlist)
- Possible separate library or contribution to libnl
- Possibility of using LNST for network tests

<https://fedoraproject.org/wiki/Features/NetworkManagerTestSuites>

Open issues

- Connection sharing is IPv4 only
- IPv4 and IPv6 autoconf scheduling
- Cleanup level on exit/unmanage/disconnect
- Persistence of networking and radio switches
- /etc/init.d/network style boot process
- Proper autoconnection of VPNs
- Multiple VPN support



Every presentation needs a cat, right?

Questions?

<http://data.pavlix.net/euopen/42/>

pavlix@pavlix.net
psimerda@redhat.com

<http://fedoraproject.org/wiki/User:Pavlix>

<http://fedoraproject.org/wiki/Networking>