

# Architectural Changes in NetworkManager

Pavel Šimerda  
pavlix@pavlix.net

LinuxAlt 2013, Brno

<http://data.pavlix.net/linuxalt/2013/>

# Me and NetworkManager

- Regular developer from May 2012 to July 2013
- Participated on a major redesign
- Still an active contributor and bug triager

Disclaimer:

The information and opinions presented here are subject to my personal **point of view**. I'm just **one of upstream developers**.

## Focusing on

- Linux kernel interoperability
- Address and routing configuration
- Keyfile configuration (upstream default)
- Testing, coverage and valgrind

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

# About NetworkManager

# Purpose of network configuration service

- Read configuration from files
- API for getting/setting the configuration
- Change notification API
- Configure networking infrastructure
- Manage auxiliary processes

# NetworkManager branches and releases

- 0.9.8 (the stable branch)
  - Important fixes
  - Bluez 5 support (requested by Gnome)
  - Occasional releases
- 0.9.10 (the development branch)
  - New development (git master)
  - Major changes
  - Needs a lot of testing
- Most of the talk will be about 0.9.10 development!

## Redesign and inner workings

# Why redesign?

- Server and virtualization use cases
- Desktop behavior needs to be optional
- Breaking other tools configuration
- Support for initramfs and security updates
- Staying compatible is not an easy task, though



From the beginnings. . .

# Initialization

- Read on-disk configuration (via plugins)
- Start tracking network interfaces
- Initialize known network interfaces

# Device initialization and hotplugging

- Check existing kernel configuration
- Generate a runtime connection profile
- Search for a matching connection profile
- **Decide about device management**
- Perform device-specific initialization actions
- Generate a default profile (optional)
- Search for autoactivatable profile
- **Launch profile activation**

# Supported devices

- Wired and wireless Ethernet (including 802.1x)
- Modems via ModemManager
- Wimax
- Infiniband
  
- VLAN interfaces
- Bridges, bonds and **teams** for wired Ethernet
- tun/tap, GRE, macvlan/macvtap
  
- **Generic device support**
- Extensive refactoring, port to nm-platform

## Runtime connection profiles

[https://fedoraproject.org/wiki/  
Networking/Ideas/NetworkManagerRuntimeConfiguration](https://fedoraproject.org/wiki/Networking/Ideas/NetworkManagerRuntimeConfiguration)

# Runtime state information

- Connection profiles
- Active connections
  - Connected devices
  - VPNs
- Disconnected devices
- Unmanaged devices

# Configuration files

- On-disk counterparts to runtime profiles
- Different due to runtime modifications
- Automatic reload now optional
  - Surprisingly a problem for libvirt!
- Explicit reload (also individually?)
- Saving the runtime state to disk

# External changes

- Added/removed interfaces
- Link state changes
- Low-level device configuration
- Addressing and routing
  
- Logging
- Updating the runtime profile



# Configuration take over

- Between initramfs and real system
- When restarting NetworkManager
- After another tool?
  
- Static configurations
- DHCPv4, RA+DHCPv6
- Software devices (bridges, bonds, etc)
- Unmanaged devices and external configuration

## Address and route configuration

# Address configuration information

- Interface (by ifindex)
- Family
- Address data
- Prefix length
- **Gateway (to be removed)**
- **Lifetime (also determines dynamic addresses)**
- Preferred lifetime
- **Label/alias information (mostly for ifcfg-rh)**

# Route configuration information

- Interface (by ifindex)
- Family
- Network (address and prefix length)
- Optional gateway
- Lifetime (in planning)
- Metric, MSS
  
- Only single-hop routes
- **Separate default gateway information**

# Policy routing

- Answer connections via alternative default routes
- Per-interface routing tables
- Per-address routing rules
- Backup access to servers
- Prerequisite for multipath TCP

## Dynamic configuration

- Configuration method redesign
- DHCP lease time propagation
- IPv6 router discovery in userspace
- Link-local IPv6 left for the kernel
- IPv4/IPv6 method consolidation
- DHCPv6 DUID support (contribution by cyphermox)
  
- libndp in git submodule
- DHCP library would be nice
- dhclient  $\geq 4$ , dhcpd is IPv4 only

[http://fedoraproject.org/wiki/  
Networking/Ideas/NetworkManagerMethods](http://fedoraproject.org/wiki/Networking/Ideas/NetworkManagerMethods)  
[http://fedoraproject.org/wiki/  
Networking/Ideas/AutomaticConfiguration](http://fedoraproject.org/wiki/Networking/Ideas/AutomaticConfiguration)

# Connection sharing

- IPv4 connection sharing over Ethernet
- IPv4 sharing over Ad-hoc Wifi and Hotspot mode
- Firewalld integration (in planning)
- Considering DHCPv6 PD-based sharing
- IPv6 masquerade seems to be too fragile
- Bridging as an alternative
- Sharing based on NDP proxy is also a possibility

## DNS configuration support



# DNS backends

- Option to turn off resolver configuration
- Direct resolv.conf writer
- unbound support (in planning)

# DNSSEC support

- Locally connected zones (in planning)
  - Option to turn DNSSEC on
  - Configure a certificate
- Global DNSSEC using dnssec-trigger
  - Fallback methods
  - Hotspot detection
- dispatcher.d versus plugin system

## Offtopic: Network name resolution library

- Avoids issues and limitations of glibc's resolver
- Asynchronous public and backend APIs
- Optional callback API
- Configurable backends
- Compatibility with nsswitch backends
- Validation information (DNSSEC)
- Fake `getaddrinfo()` for LD\_PRELOAD

## API and tools

# Configuration, API and tools

- Private socket support, D-Bus daemon made optional
- Connection normalization
- ActiveConnection API (between device and connection)
- Thorough CLI redesign
- CLI for connection management
- Guided configuration CLI/TUI
- Library/API redesign in planning (danw, #680675)
- Documentation

# Community

# Community

- Four regular developers (at Red Hat)
- Distribution maintainers
  - Debian
  - SUSE
  - Gentoo
- Individual contributors
- Bug triagers
- Users, testers and bug reporters
- Documentation and websites

`git://anongit.freedesktop.org/  
NetworkManager/NetworkManager.git`

`http://bugzilla.gnome.org/`

# Questions?

<http://data.pavlix.net/linuxalt/2013/>

[pavlix@pavlix.net](mailto:pavlix@pavlix.net)  
[psimerda@redhat.com](mailto:psimerda@redhat.com)

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

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