

# XORP and other open multicast router implementations

Stig Venaas  
UNINETT  
venaas@uninett.no

# Introduction

---

- Heavily based on info from Pavlin Radoslavov
- Only looking at UNIX based PIM implementations
- Will look at history followed by current status
- Info about implementations at <http://netweb.usc.edu/pim/>

# Implementations(1)

---

- History begins sometime pre 97
  
- People at USC try to implement PIM-SMv1
  
- Ahmed Helmy at USC implements PIM-SMv2 (1996 I think)
  - may have used code from previous v1 attempts
  - described in draft-helmy-pim-sm-implement-00.txt from January 1997
  - <http://netweb.usc.edu/pim/pimd/docs/draft-helmy-pim-sm-implement-00.txt>
  
- He is also one of the authors of RFC 2117 (PIM-SMv2) from June 1997

# Implementations(2)

---

- Pavlin Radoslavov in summer 1997 as intern at SUN reimplemented and updated large parts
  - Based on what became RFC 2362 in 1998
  - High level design was kept from old version
  - Worked on FreeBSD and Solaris
- This is the daemon normally used on Linux (and \*BSD I think)
  - Also supported on other platforms with kernel support
  - More about that later
- Several problems with design and implementation and some serious bugs (according to Pavlin)
- At about the same time Rusty Eddy at USC implemented PIM-SMv2 for GateD 5.0
  - Unclear how free the implementation is
  - PIM code itself is available at <http://www.isi.edu/~eddy/pim/pim.html>
  - Seems Linux supports it

# Implementations(3)

---

- The previous implementations are for IPv4
  
- In 1998 (I think) Mickaël Hoerdt at University Louis Pasteur ported Pavlin's daemon to IPv6
  - Maintained by KAME since then
  - KAME has added features, fixed bugs etc
  - The new PIM draft draft-ietf-pim-sm-v2-new-09.txt fixes some problems in RFC 2362, not sure to what extent the changes have been incorporated into the pim6sd daemon
  - Also several patches around for BSR filtering, embedded-RP etc.

# Implementations(4)

---

- Over time Pavlin found it very hard to maintain and extend the daemon
- When joining XORP in Summer 2001, he decided to implement PIM-SM from scratch
- From beginning based on the new PIM draft
  - draft-ietf-pim-sm-v2-new-00.txt written in 2000
  - Draft has been through several revisions. RFC RSN
- He believes the XORP implementation is more correct
  - Realised several problems with old implementation while working on this
- Only open-source dual IPv4/IPv6 implementation
  - Most of the code is common for IPv4 and IPv6
  - Makes it easier to maintain
  - Available at <http://www.xorp.org/>

# Implementations(5)

---

- PIM-DM implementations also exist
- GateD has PIM-DM
- Implementation by Kurt Windisch
  - then at University of Oregon
- MRTD at <http://www.mrtd.net/>
- Not sure about GateD, the others are not maintained

# Multicast kernel support (1)

---

- Kernel forwards multicast packets according to MFC
  - Multicast Forwarding Cache
- If no matching entry, message passed to daemon
- PIM/IGMP parsed by kernel and message passed to daemon
  
- Ahmed did original implementation for IRIX and SunOS
- Ported to FreeBSD by Pavlin
- NetBSD port by Hitoshi Asaeda (KAME, now INRIA)
- KAME ported to IPv6
  - IPv6 code appeared in \*BSD long before IPv4
- Linux IPv4 support loosely based on Ahmed's work
- In Fall 2003 IPv4 code added to FreeBSD 4.9/5.1

# Multicast kernel support (2)

---

- FreeBSD IPv4 API now more advanced than others
- Support for
  - in-kernel PIM Register encapsulation
  - kernel bandwidth monitoring of multicast dataflows and upcalls
  - etc
- More info in multicast(4) on FreeBSD
- Hitoshi Asaeda has expressed interest to again port from FreeBSD to NetBSD
- If interest, Pavlin might at some point port to OpenBSD
- Pavlin thinks XORP might be the only to use the advanced API

# Multicast kernel support (3)

---

- Linux only has IPv4 multicast forwarding
- Hungarian project is working on IPv6 support
  - See <http://dora.cab.u-szeged.hu/>
- As part of 6NET, liberouter (<http://www.liberouter.org/>) will work on hardware IPv6 multicast support on COMBO6
- This will be useable on at least NetBSD
  - (I think NetBSD was the primary platform)
  - Other platforms are FreeBSD and Linux

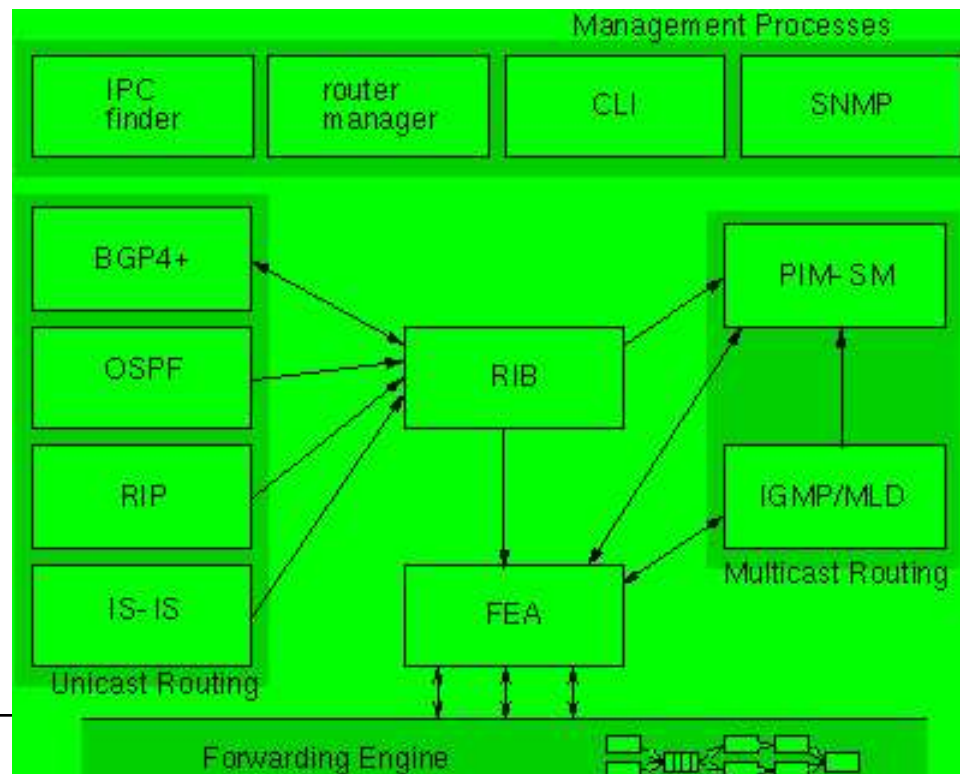
# Multicast kernel support (4)

---

- Today XORP only does IPv4 multicast on Linux, but could use the Hungarian work to support IPv6
  
- The Hungarian work would also make it much easier for liberouter to do IPv6 multicast on Linux
  
- The hardware forwarding offered by COMBO6 should be invisible to the routing daemons, hence it should work with the above PIM implementations
  - (at least that's my understanding)

# XORP

- Modular design
- XRL for communication between modules (RPC)
  - XORP Resource Locator, looks like an URL
- Easy to write custom modules, do simulations etc
- Forwarding engine abstraction
- Centralized management
  - One central config file
  - One CLI
  - One SNMP process



# XORP status

---

- Works on FreeBSD and Linux
- UNIX kernel forwarding, Click later
- PIM-SM and BSR (including scoping) per latest drafts
- Static RP and static MRIB routes
- IGMPv2 and MLDv1
- SSM ok if intermediate router
- IPv4 and IPv6 (not IPv6 multicast on Linux)
- MBGP (including multicast SAFI)
- RIP/RIPng
- OSPF (v4 only I think)
- CLI
- No SNMP in 1.0
- 1.0 to be released by end of month
- <http://www.xorp.org/>