IPv6 Development Status 2005

YOSHIFUJI Hideaki
<yoshfuji@linux-ipv6.org>
Keio University / USAGI/WIDE Project

Agenda

- Recent Activities
  - IPv6 Ready Logo Phase-1 / Phase-2
  - TAHI Automatic Test Running System
  - Statistics
  - Advanced API (RFC3542)
  - IPsec
  - Mobile IPv6
  - Connection Tracking
- Misc. TODOs / Future Plans

IPv6 Ready Logo

- We are finally "IPv6 Ready!"
  - 2.6.11-rc2 is IPv6 Ready Logo Phase-1 certified
    - with special version of radvd (Router Advertisement Daemon)
  - 2.6.12 is available with improvements to pass the Self Test for IPv6 Ready Logo Phase-2, Core Protocols.
    - with special version of radvd
    - interop is not tested yet...
- http://www.ipv6ready.org

TAHI Automatic Test Running System (aka USAGI Testlab)

Chinen Mitsuru @IBM Japan, Ltd.
Hideaki Yoshifuji @Keio University

USAGI Testlab

- Test every new snapshots automatically
- Target
  - IPv6 Ready Logo Phase-1
    - Host
    - Router
  - IPv6 Ready Logo Phase-2, Core Protocols
    - Host
    - Router
- <http://testlab.linux-ipv6.org> (IPv6 Only)

USAGI Testlab TODOs

- Increase test targets
  - IPv6 Ready Logo Phase-2
    - IPsec
    - Mobile IPv6
  - The Original TAHI Cts?
  - Even now, it takes so looooooooooooong time...
- Description / Documentation
  - kernel? userspace?
- Requests?

Statistics

Koichi Kunitake @Anchor Technology, Inc.
Hideaki Yoshifuji @Keio University

Statistics

- Try not to get refcnt everytime by using rt6i_idev in rt6_info{} (rt6).
- Issue
  - dst (rt6) is unavailable in ipv6_rcv().
  - Is it safe to use __in6_dev_get() in ipv6_rcv()?
    - Yes: okay, do it.
    - No: Hmm...

      1)in6_dev_get()?

      2)put critical stats in net_device()?!
      - IPSTATS_MIB_INRECEIVES

Statistics TODO / Future Plans

- clean-ups and check races
  - will be submitted after this
- The “HC" Counters
  - plan to implement in not-too-generic manner

Advanced API

Hideaki YOSHIFUJI @Keio University

Advanced API (RFC3542) (1)

- Basic
  - Rename access to extension headers
    - IPV6_RECVHOPOPTS, IPV6_RECVDSTOPTS, ...
    - OLD: IPV6_HOPOPTS, IPV6_DSTOPTS, ...
  - Split the "sticky" option
    - IPV6_HOPOPTS, IPV6_DSTOPTS, ...
    - OLD: IPV6_PKTOPTIONS

Advanced API (2)

- Additional
  - PMTU discovery
    - IPV6_USE_MIN_MTU (like IPV6_MTU)
      - disable for multicast (default) / always / disable
  - Fragmentation
    - IPV6_DONTFRAG
  - PMTU notification / query
    - IPV6_RECVPMTU, IPV6_PMTU
    - sockets w/ IPV6_RECVPMTU receives PMTU information for all destinations
    - check current PMTU by IPV6_PMTU (connected socket)

Advanced API

- Collaborating with David L. Stevens @IBM
- Implemented (most of) basic part
  - Preserving old API by renaming old ones.
    - IPV6_2292xxx
  - Reallocates ipv6_txoptions{}
  - Not tested yet; will do, of course.

Advanced API Issues (1)

- We assume the order of extension headers
  - we record the offset to each extension header in skb->cb in extension header handlers, assuming the standard order.
  - we need to preserve the order.
- Solution?
  - remember last offset to the pointer and parse the packet again

Advanced API Issues (2)

- Is it Okay to access skb->nh.raw + offset in recvmsg()? No!
  - need to implement put_cmsg_skb() to cope with cmsg, userspace pointers, and non-linear skb
  - will do.

IPsec

IPsec Team
Mitsuru Kanda @Toshiba Corp. (emertus)
Kazunori Miyazawa @Yokogawa Electric Corp.

IPsec Status (1)

- Score of IPv6 Ready Logo Phase-2 IPsec Self Test (host , aka endnode) is not so bad.
  - AES-128-XCBC-96 is missing
- router, aka sgw, will be tested

IPsec Status (2): RFC3566

- RFC3566 AES-128-XCBC-96
  - Keyed-MAC algorithm
  - Required for IPv6 Ready Logo Phase-2 IPsec
  - We implemented it like hmac extension against md5 and/or sha1.
  - It works; tested with test vectors in the RFC by using tcrypt.
  - Will be submitted shortly

IPsec Status (3): racoon2

- WIDE/IPsec working group released racoon2 IKEv2 and KINK.
  - racoon2 runs on both Linux and NetBSD.
  - design and implementation of racoon2 is different from ones of racoon.
  - They will implement IKEv1 on racoon2 architecture.
- Maintenance of racoon1 has been tossed to sourceforge.
  - KAME is trying to pass over things to community.

IPsec Question (1): Asymmetry

XFRM Policy Asymmetry

- FAQ
- We think the interface is confusing.
- We would like to know the reason of asymmetry.
  - Why is there fwd or isn't fwd-out for outbound?
- We / people prefer a symmetric interface.

IPsec Question (2): Interprotocol tunnel

- We are interested in IPv6 over IPv4 IPsec tunnel.
- Status?

Mobile IPv6

MIP6 Team “U-MIP”

Noriaki Takamiya @NTT Software Corporation
Masafumi Aramoto @Sharp Corporation
Masahide Nakamura @Hitachi Communication Technologies, Ltd.
Shinta Sugimoto @Nippon Ericsson K.K.

MIPL2: Mobile IPv6 for Linux

- What is MIPL2?
- Interaction between MIPv6 and IPsec/IKE
- Kernel design
- Development status
- Future plan

What is MIPL2?

- A Mobile IPv6(MIPv6) stack targets 2.6 kernel
  - support MIPv6 basic specification RFC377{5,6}
  - support IPsec/IKE interaction
- USAGI/WIDE Project is making joint effort with Helsinki University of Technology(HUT) for MIPL2
  - UMIP Team
    - Noriaki Takamiya @NTT Software Corporation
    - Masafumi Aramoto @Sharp Corporation
    - Masahide Nakamura @Hitachi Communication Technologies, Ltd.
    - Shinta Sugimoto @Nippon Ericsson K.K.

Mobile IPv6 Kernel Design

- New daemon is defined and it controls kernel status
- The daemon handles
  - signaling
  - master information (e.g. binding cache)
  - movement detection
  - IPsec interaction
- not so large lines required for kernel

Mobile IPv6 & IPsec/IKE Interaction

- IPsec tunnel established between the MN and HA needs to be updated whenever the MN changes its CoA
- IPsec/IKE should also be aware of ‘movement’ of MN because:
  - IPsec Security Policy Database (SPD) needs to be updated in accordance with CoA change.
  - IPsec Security Association (SADB) needs to be updated in accordance with CoA change.
  - IKE needs to update the IKE connection (K-bit).
- A solution:
  - Make an interface between Mobile IPv6 and IPsec/IKE by extending PF_KEY framework (PF_KEY MIGRATE message)
  - “PF_KEY Extension as an Interface between Mobile IPv6 and IPsec/IKE”, draft-sugimoto-mip6-pfkey-migrate-00
  - Implemented in MIPL2.0 RC2

IPv6 Kernel Functional Block(linux-2.6 mainline)

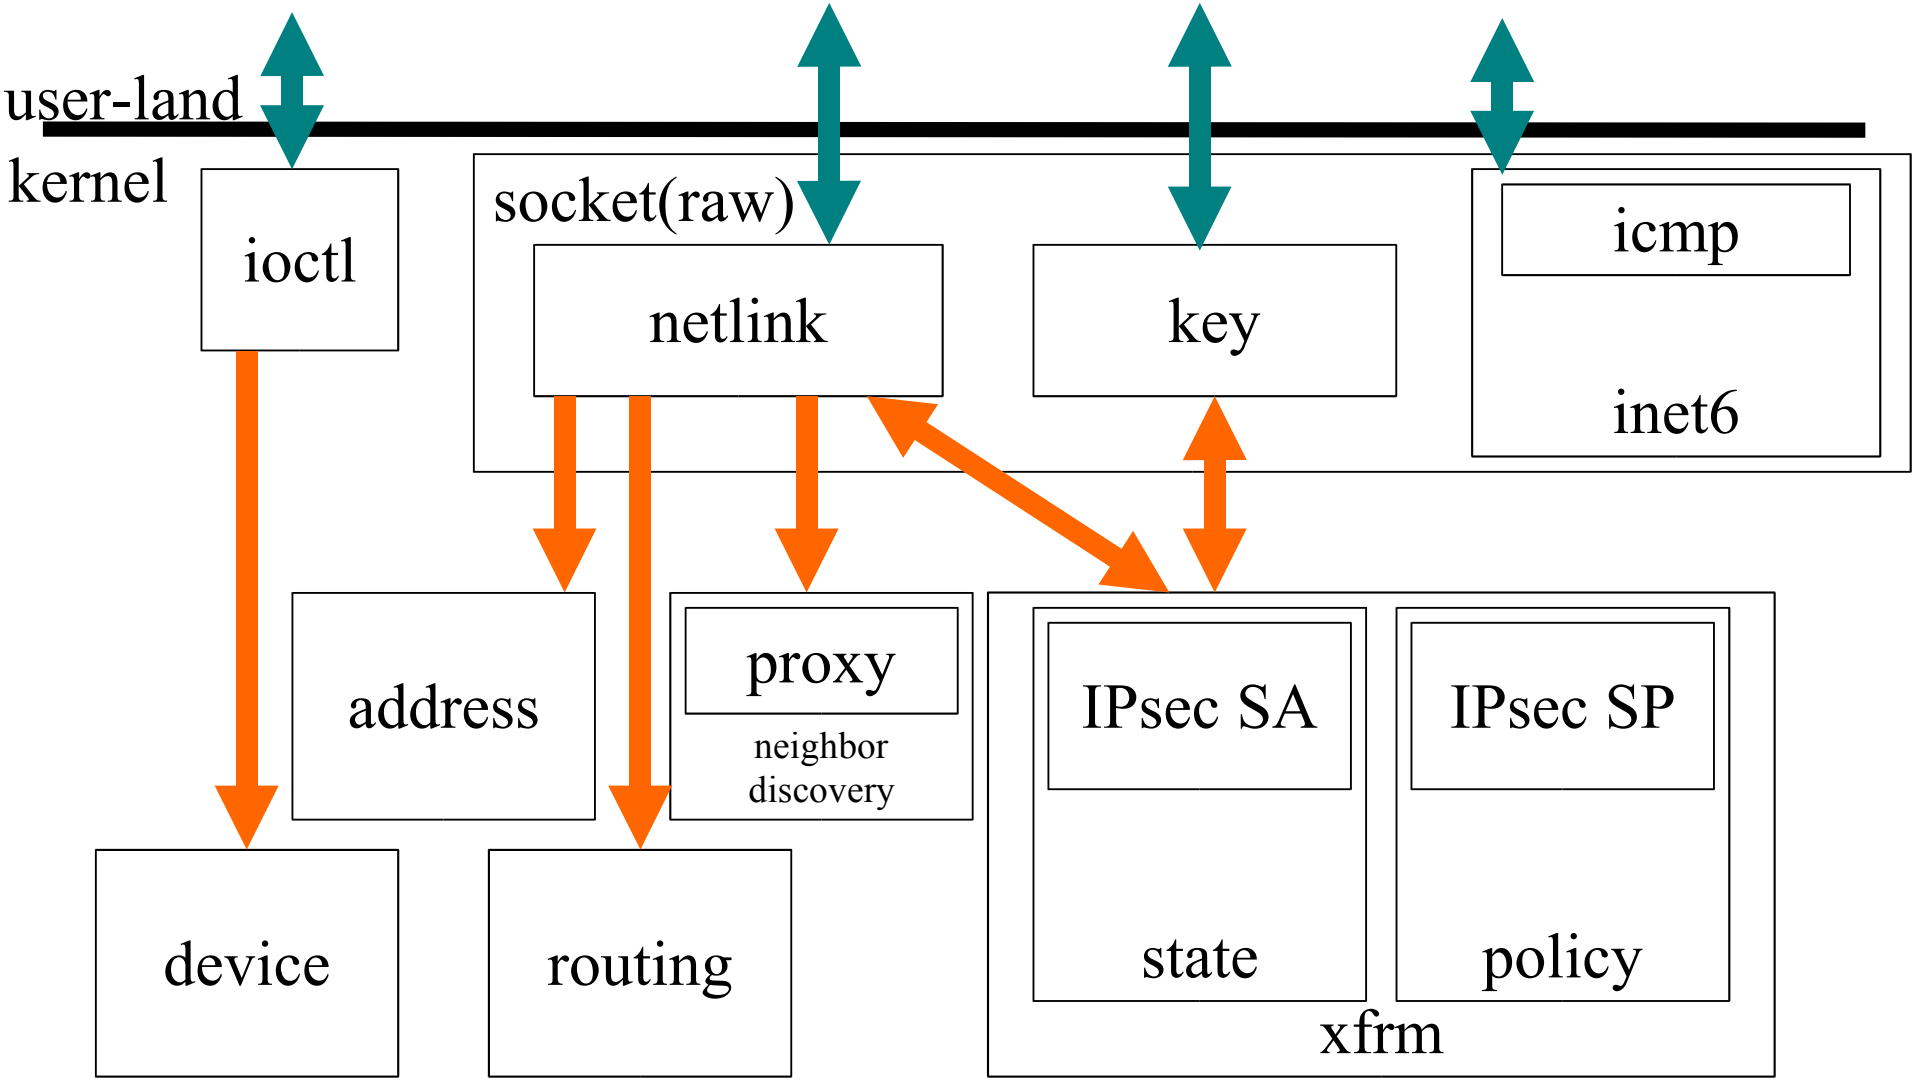

MIPL2 Kernel Functional Block

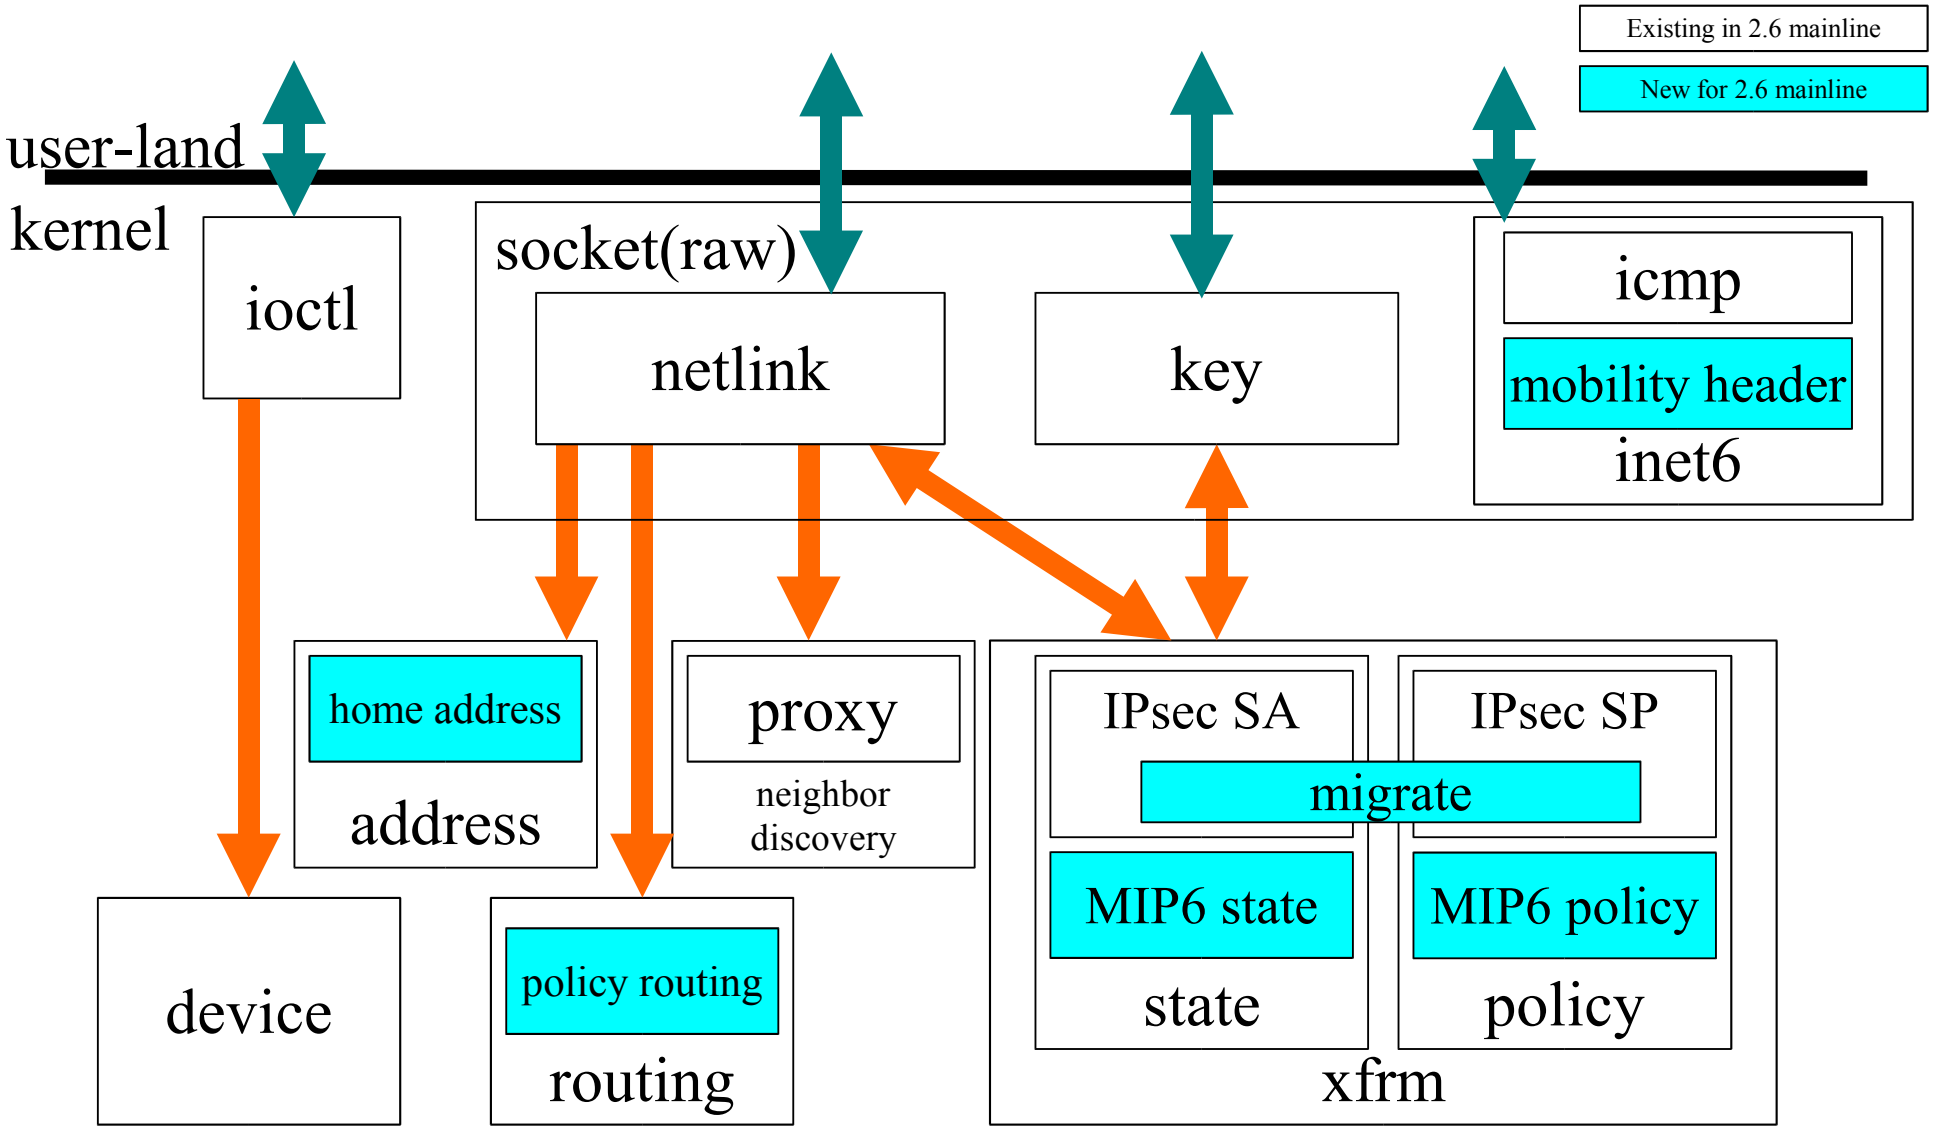

MIPL2: XFRM Modifications (1/3)

- extend “xfrm template” and “xfrm state” to support two more protocols/extension headers (i.e. xfrm is a subsets of binding cache/binding update list in kernel)
  - Destination options header (to carry home address option)
  - Routing header type 2
- Issue
  - It must be prepared one policy which carries both MIPv6 template and IPsec template in advance by user-space when using IPsec and MIPv6 at the same time; MIPL2 daemon must handle IPsec policy to append MIPv6 protocol to it
    - plan: separating policy in kernel and combined them (testing experimental code)

MIPL2: XFRM Modifications (2/3)

- Add xfrm notification (from kernel to the daemon)
  - for Binding Error; It is occurred when packet is dropped by xfrm policy (MIPv6 policy)

MIPL2: XFRM Modifications(3/3)

- Use "Migrate interface" to update endpoint address of IPsec tunnel when MN moves

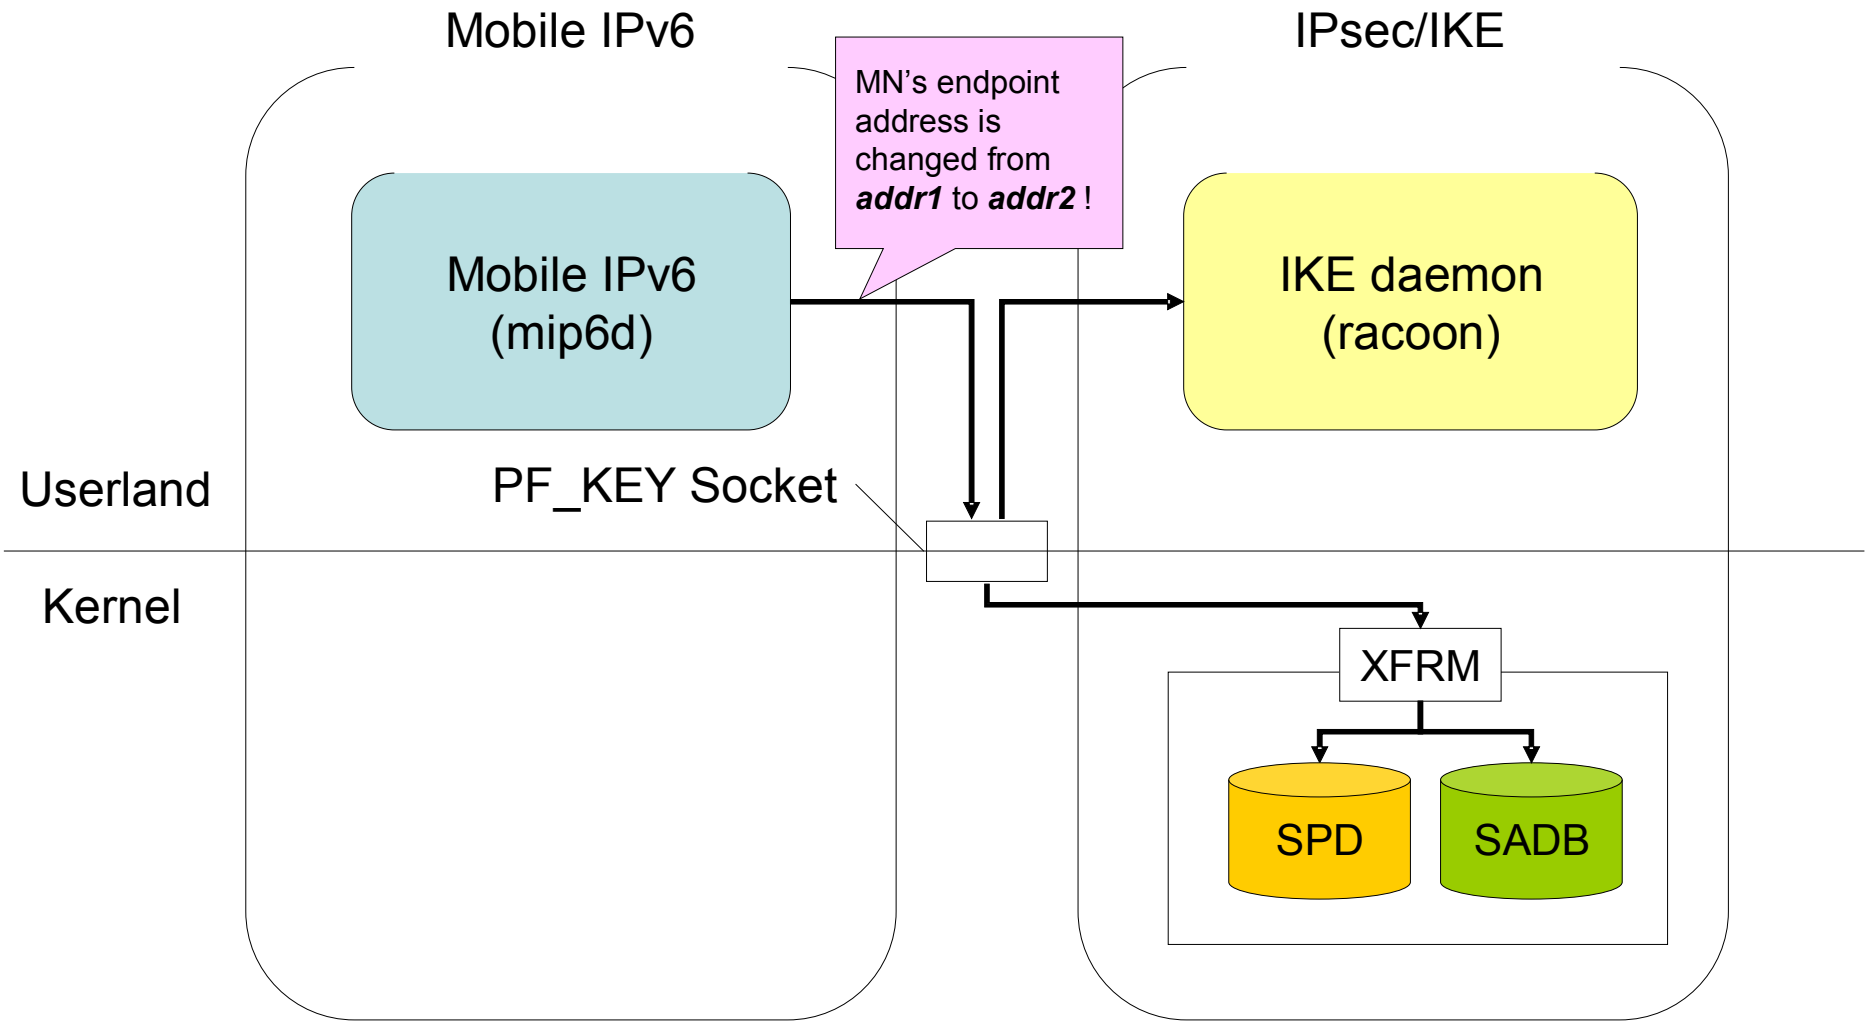

MIPL2: Address and Routing

- Identifying the Home Address
  - add IFA_F_HOMEADDRESS to ifa_flags
- Policy Routing based on source address
  - each source address has each FIB entry based on IPv4 multiple table
  - policy per Home Address for route/interface selection can be used for multiple Home Address support

MIPL2 Status

- Released “RC2” in May 2005
  - based on 2.6.8.1
- Now working on 2.6.11
- Solid and stable performance proved by the successful results of TAHI conformance tests for MN/HA/CN
  - targeting getting full scores
- Release 2.0
  - End of July

MIP6: Future Plans (1)

- Extensions
  - HAHA
    - inter HA protocol for improving reliability
  - Multiple CoA
    - optimized route path selection in multihomed environment
  - Policy Handoff
    - modularized movement detection

MIP6: Future Plans (2)

- HMIP
  - reduction of signaling overhead
  - location privacy
- MIPv6 in different address space
  - IPv4 traversal
  - IPv4 homeaddress support
- NEMO (Nautilus6 Project)
- FMIP (RFC4068; Nautilus6 Project)
  - Fast handover

Connection Tracking

Yasuyuki Kozakai @Toshiba Corp.

Nf_conntrack

- IP-independent connection tracking
  - supersedes ip_conntrack (and ip6_conntrack)
- Collaborating Netfilter Project
  - already available there
- Status?
  - git trees?

Misc. TODO / Plans (1)

- Policy Routing
  - Based on work by MIPL2
- ISATAP
  - Based on patch from Fred Templin (author of the draft)
  - Need to refresh
- IPv4-over-IPv6 Tunnel

Misc. TODOs / Plans (2)

- Multicast Forwarding
  - Yuji Sekiya @Univ. Tokyo and myself
  - based on patch from Mickael Hoerdt @Universite' Louis Pasteur, Strasbourg
- Whatever things make us happy

Misc.

- We will get 10Gb connectivity in our office soon!
- Distributed computing
- XCAST (Explicit Multi-unicast
  - <http://www.xcast.jp>
- Netnice
  - <http://www.netnice.org>

Conclusion

- Linux is IPv6 Ready!
  - It is the beginning.
- Things we're working on
  - Statistics
  - Advanced API (RFC3542)
  - IPsec
  - Mobile IPv6
  - Connection Tracking

Thank you