# Linux Is Now IPv6 Ready

Hideaki YOSHIFUJI

USAGI/WIDE Project / Keio University

Copyright (C)2005 USAGI/WIDE Project

# Topics

- Introduction
- USAGI Improvements
- IPv6 Ready Logo
- Recent Activities
- Future Directions
- Conclusion

# Little History of IPv6

- End of '80s
  - Issues
    - address space
    - routing scalability
    - ...
- 1992
  - first full-scale discussion at IETF
    - IPng
- 1995
  - RFC1883
    - Proposed Standard
- 1998
  - RFC2460
    - Draft Standard

# IPv6 Features

- Huge Address Space
  - 128bits (> 32bits)
- Simpler architecture for high speed network
  - addressing architecture
  - aggregatable address
  - extension header
  - no fragmentations
    - except for end nodes
    - Built-in autoconfiguration
    - Stateless Address Autoconfiguration
- Security
- Mobility

# Linux IPv6 Development

- Since 2.1.x
  - by Pedro Roque (1996)
  - It works but...
- Stale
  - specification conformity
    - NDP, Autoconf, ...
    - API
  - missing features
    - IPsec, Mobile IPv6, ...

# USAGI Project

- Universal Playground for Ipv6
- http://www.linux-ipv6.org
- Goals
  - production quality IPv6 stack
  - whatever makes us happier with IPv6

cf) KAME <http://www.kame.net> for BSD variants

# USAGI Project (cont.)

- Established
  - 2000
- Project Leader
  - Jun Murai, Keio University
- Members
  - 10(+x) from 2 universities and 8 companies
- Office
  - Keio University, Kawasaki, Japan
- Collaboration
  - WIDE
    - KAME, TAHI, IPsec
  - Go/Core (Helsinki University of Technology)

# USAGI Improvements

- IPsec
  - FreeS/WAN based
  - XFRM based
- API
  - sin6_scope_id
    - extra member in sockaddr_in6{}
    - backward compatibility
  - IPV6_V6ONLY
    - split port space if requested

# USAGI Improvements (cont.)

- NDP, Addrconf
  - incorrect state transition
  - timer
  - Most bugs were spot by TAHI Conformance Test Suite
    - http://www.tahi.org
- Routing
  - default routes

# IPv6 Ready Logo Phase-1

- International certification program for IPv6 products
  <http://www.ipv6ready.org>
  - Phase-1
    - since 2003/09
    - Self Test + Interop
    - Basic functions
      - Spec, ND, Addrconf, ICMP
      - (most of) "MUST"
    - ~180 products from various countries

# IPv6 Ready Logo Phase-1 in Linux

- USAGI products
  - 2.4, 2.6
- Knoppix IPv6 Edition
- EFI Linux
- Linux 2.6.11-rc2
  - Host
  - Router
    - w/ patched rtadvd
      - backported to rtadvd distribution
  - later versions are expected to pass
    - retest required
    - unexpected regression is possible

# IPv6 Ready Phase-2

- Phase-2
  - since 2005/02
  - Self Test + Interop
  - More “advanced” tests
    - Core Protocols
      - Spec, ND, Addrconf, PMTU, ICMP
      - “MUST” +” SHOULD”
    - IPsec
      - Basic, Advanced
    - Mobile IPv6
      - MN, HA, CN
    - And more...
      - MLDv2, Transition, IKE, ...

# IPv6 Ready Logo Phase-2 in Linux

- Issues
  - PMTU
    - always add fragment header after receiving TOOBIG w/ tiny mtu value
  - Metric Inheritance
    - hop limit etc.
  - Persistent on-link (prefix route)
    - manual deletion did not delete prefix route
  - Several redirect issues
  - RA configuration and interface configuration interaction
    - sysctl
- Fixes included in 2.6.12 and later
  - Interop not yet

# Recent Activities

- TAHI Automatic Running System
- Mobile IPv6
- Connection Tracking
- IPsec
- Statistics
- Advanced API
- IPv6-Fix

# TAHI Automatic Running System

- Development continues, but maintaining quality is important!
  - watching quality by less human resource
    - we're not robots...
  - consistent result
  - spotting "bad" change

# Autorun System Components

- The Scripts
- InitScript
  - select various settings by boot parameter
    - interface
    - address
    - forwarding
    - radvd (Router Advertise)
- Test Summarizer
  - multiple results in one table
  - coloring

Requires:

- TAHI Conformance Test Suite
- GRUB

# TAHI Autorun System Diagram

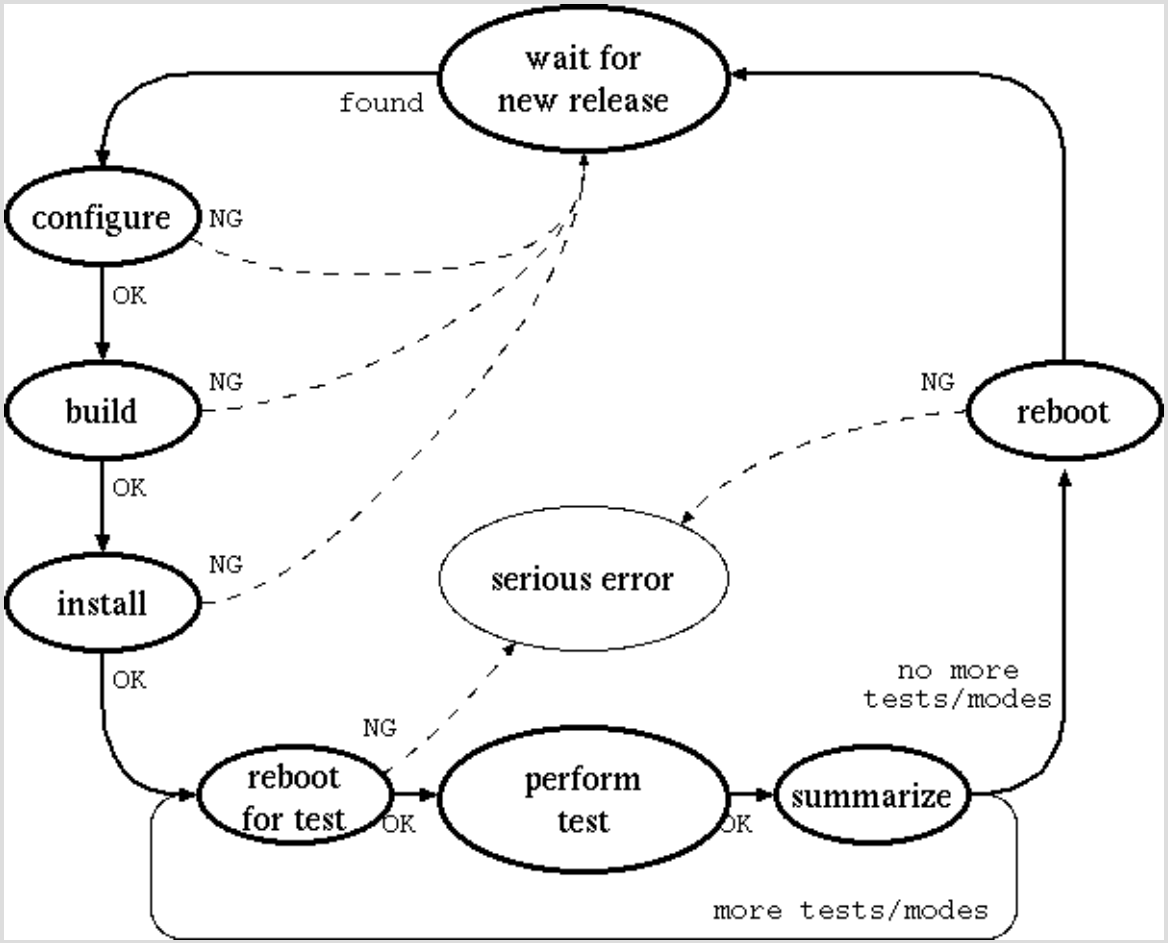

# USAGI Testlab

- http://testlab.linux-ipv6.org
- FreeBSD 4
  - 1000BASE-T (em)
  - Serial and USB Serial (cuaa, ucom)
  - 140GB disk for results
  - Testsets
    - IPv6 Ready Logo Phase-1
    - IPv6 Ready Logo Phase-2 Core Protocols
    - 400MB
- Test every snapshots

# Mobile IPv6

- MIPL (Mobile IPv6 for Linux) 2
  - GO/Core and USAGI collaboration
- Targets 2.6.x
- basic specification RFC377{5,6}
- IPsec/IKE interaction

# Mobile IPv6 (2)

- Try to minimize kernel implementation
  - using XFRM infrastructure
- MIP6 daemon
  - signaling
  - master information
  - movement detection
  - IPsec interaction

# MIPL2 Diagram

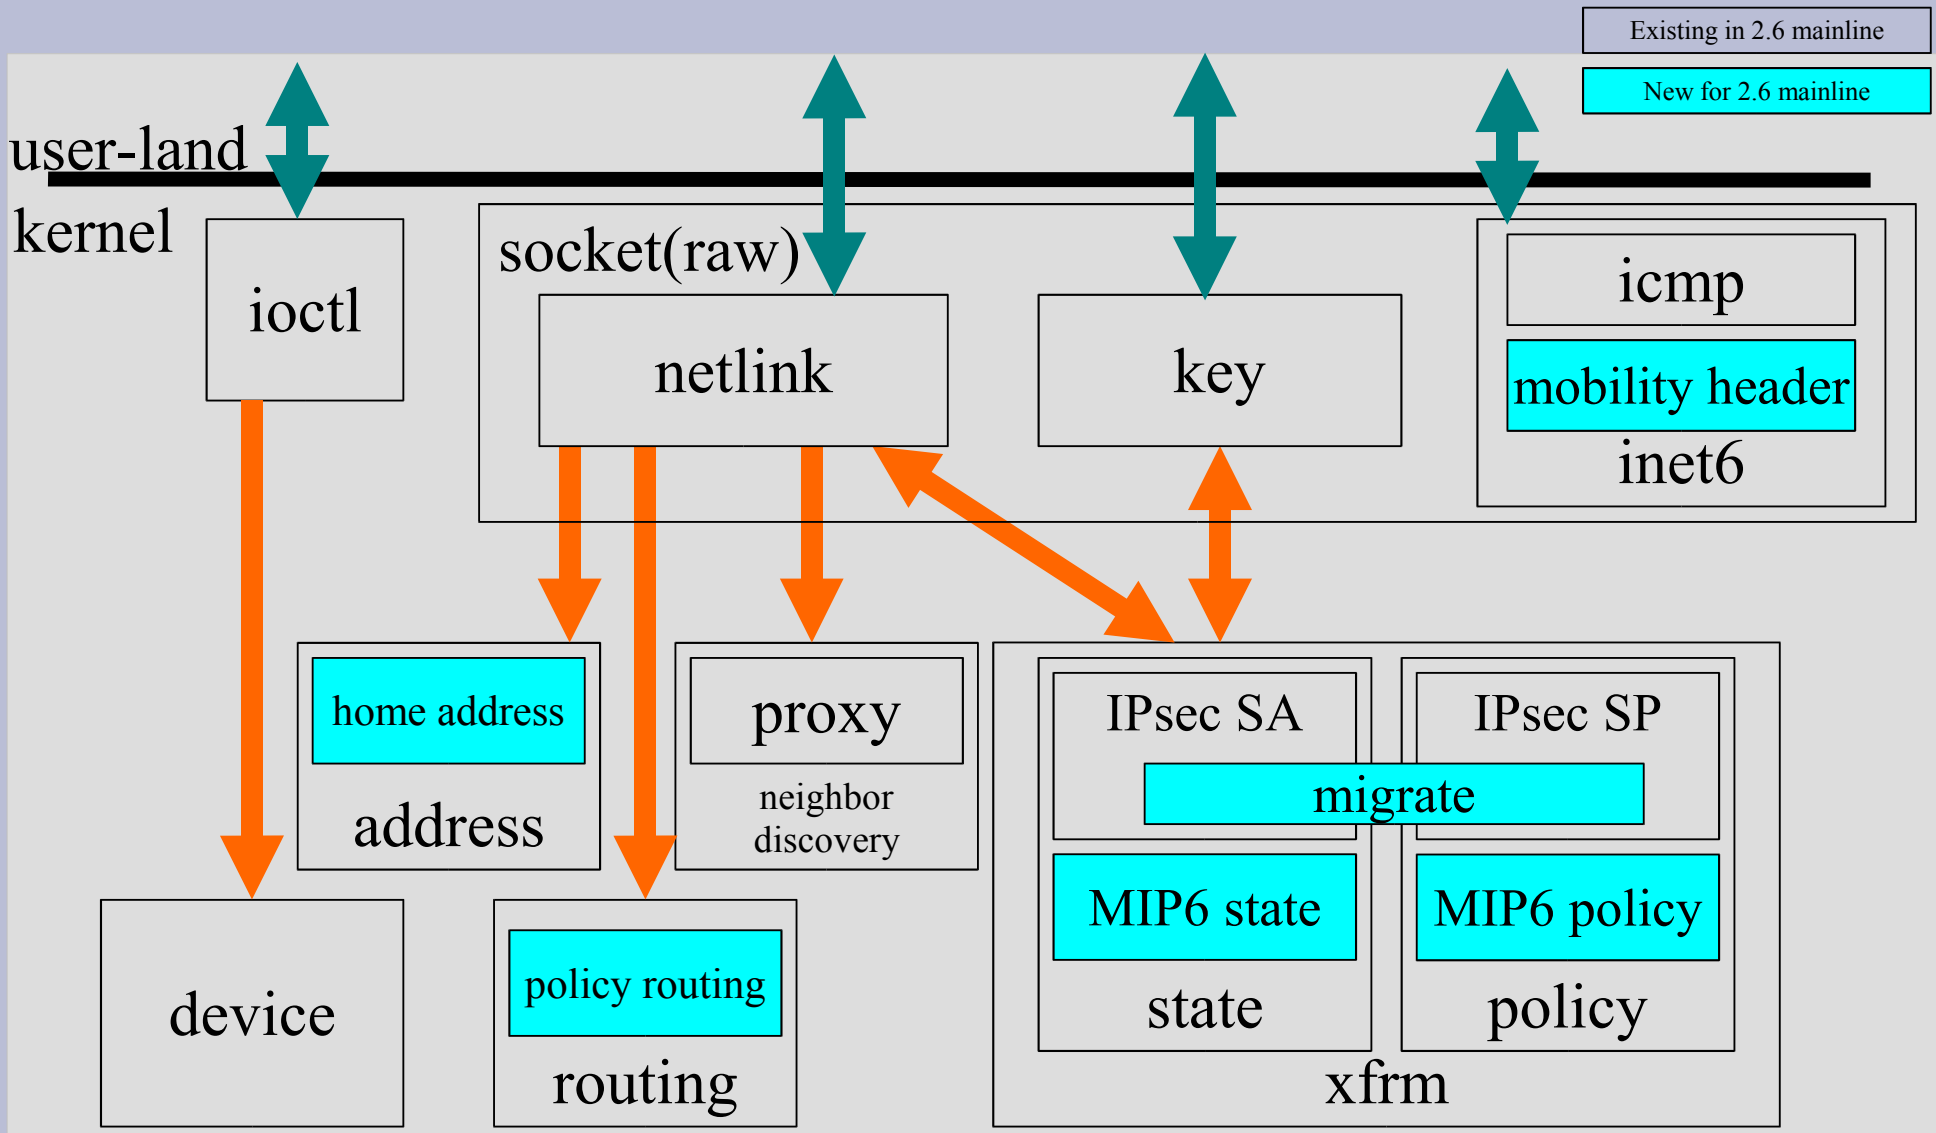

# MIPL2 Status

- "RC2" was released in May 2005
  - based on 2.6.8.1
- Working on 2.6.11
  - Almost stable
- Release
  - Targetting full scores in IPv6 Ready Logo Phase-2 Mobile IPv6 Tests
  - coming very soon

# Connection Tracking: nf_conntrack

- Background
  - Linux did not support stateful connection tracking for IPv6
  - ip6_conntrack was implemented
    - lots of duplicate code
- what it is
  - Version independent connection tracking
- Status
  - Basic part already available at netfilter.org
  - IPv4 NAT...

# IPsec: IPv6 Ready

- Filling gaps between Linux and IPv6 Ready Logo
  - end-node is not bad; just a few failures
    - investigating
- Missing feature
  - AES-128-XCBC-96
    - keyed-MAC
    - already implemented
    - will be merged soon

# IPsec: Racoon2

- by WIDE IPsec WG
- multi protocol support
  - (IKEv1)
  - IKEv2
  - KINK
- Linux, NetBSD
- redesigned

Note: racoon1 has beed tossed to sourceforge

- KAME is trying to pass over things to community

# Statistics

- Per-interface statistics
- basic infrastructure available
  - /proc/net/dev_snmp6/<ifname>
  - refcnt reduced than before
- most part already implemented
  - merged soon
- patch for net-snmp also available

# Advanced API

- structures / constants / enums
- access to extension headers
- other advanced controls
  - fragmentation
  - ...

# Advanced API (cont.)

- RFC2292->RFC3542
  - access to extension headers (on receipt)
    - (OLD) IPV6_HOPOPTS, ...
    - (NEW) IPV6_RECVHOPOPTS, ...
  - sticky option
    - (OLD) IPV6_PKTOPTINOS
    - (NEW) IPV6_HOPOPTS, IPV6_DSTOPTS, ...
- both conflict
  - old options will be renamed
    - IPV6_2292xxx
- Mostly implemented
  - will be merged after testing

# IPv6-Fix Project

- http://v6fix.net
- People sometimes recognize some regressions w/ IPv6-enabled Oss.
  - One day a guy tried to connect the Internet via the RJ-45 in his room. But his OS could not show any web pages. The hotel staff advised him to type “ipv6 uninstall.”
  - Mozilla on Linux sometimes take time to to show web page, which is not available via IPv6. Some people advice that to disable Ipv6.

# IPv6-Fix Project (cont.)

- Most are not of IPv6 itself in fact, but they may prevent people from using IPv6.
- IPv6-Fix Project launched to "fix" such issues.
  - spec
  - implementation (incl. work arounds)
  - network operation
  - ...
- Please send any symptoms / information at <contact at v6fix.net>.
  - product information highly appreciated.

# Future Directions

- Mobility
  - HMIP
    - reduction of signaling
    - location privacy
  - MIPv6 in different address space
    - IPv4 traversal
    - Ipv4 home-address
  - NEMO (Network Mobility)
  - FMIP
    - Fast handover

# Future Directions (2)

- Statistics
  - HC (64bit) counters
- Advanced API
  - fragmentation, PMTU reporting
- Policy routing / source address selection
  - based on implementation by MIPL2
- Multicast Forwarding
  - based on implementation by Michael Hoerdt

# Future Directions (3)

- Tunnel
  - ISATAP
    - based on obsolete implementation for 2.4 by Fred Templin
  - IPv4-in-IPv6 Tunnel
  - tunl/sit integration
- Performance
  - RCU (Read-copy-update)
  - Offloading
- More Tests
  - IPsec, Mobile IPv6, MLDv2, ...
- Whatever that makes us happier!
  - multihoming
  - prefix deprecation

# Thing in the Wild

- XCAST (Explicit Multi-unicast)
  - http://www.xcast.jp
  - multicast-like communication for small group
    - video conference etc.

# Conclusion

- Linux Is Now IPv6 Ready!
  - This is the beginning.
- Future Directions
  - Mobile IPv6
  - Connection Tracking
  - Routing
  - Statistics
  - Advanced API
  - IPsec
  - Performance
  - More tests

Slides will be available at
http://www.linux-ipv6.org/materials/200507-OLS/

# Thank You