



University of
Massachusetts
Amherst

ECE697AA – Lecture 7

Software Router:
Networking in OS Implementations

Tilman Wolf
Department of Electrical and Computer Engineering
09/23/08



Implementation of protocol stack

- Router process packets at Network Layer
 - Operations specified in RFC 1812
- How can we build an IPv4 router?

Network Working Group	F. Baker, Editor
Request for Comments: 1812	Cisco Systems
Obsoletes: 1716, 1809	June 1995
Category: Standards Track	
Requirements for IP Version 4 Routers	
Status of this Memo	
This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.	
PREFACE	
This document is an updated version of RFC 1716, the historical Router Requirements document. That RFC preserved the significant work that went into the working group, but failed to adequately describe current technology for the IESG to consider at a current standard.	
The current editor had been asked to bring the document up to date, so that it is useful as a procurement specification and a guide to implementors. In this, he stands squarely on the shoulders of those who have gone before him, and depends largely on expert contributors for text. Any credit is theirs: the errors are his.	
The content and form of this document are due, in large part, to the working group's chair, and document's original editor and author: Philip Alquist. It is also largely due to the efforts of its previous editor, Frank Mastenholz. Without their efforts, this document would not exist.	
Table of Contents	
1. INTRODUCTION	4
1.1 Reading this Document	4
1.1.1 Organization	4
1.1.2 Requirements	9
1.1.3 Compliance	10
1.2 Relationships to Other Standards	11
1.3 General Considerations	12
1.3.1 Continuing Internet Evolution	12
1.3.2 Robustness Principle	13
1.3.3 Error Logging	14
Baker	Standards Track [Page 1]

Software router

- Hardware
 - PC with multiple NICs
- Software
 - What functions do we need to implement?

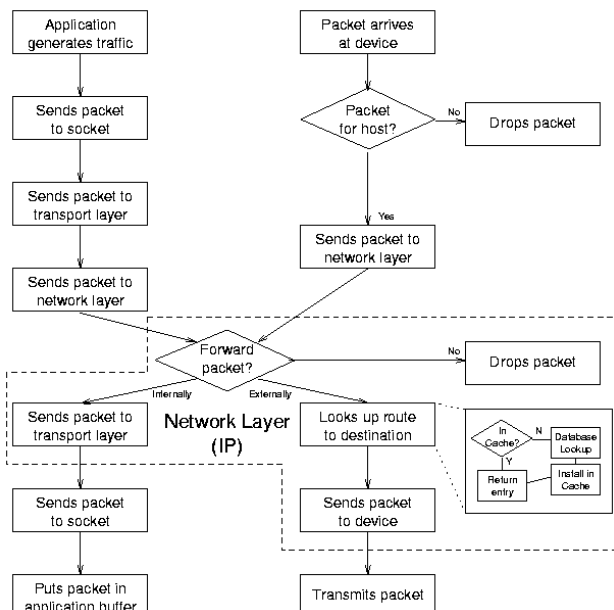


(Figure from Douglas Comer: Network System Design using Network Processors)

Software router

- Can we write a user-space IPv4 application?
 - Possible, but inefficient
- IPv4 as operating system component
 - OS has better access to hardware resources
 - OS can access packet data without copying
- Typical
 - IPv4 processing part of kernel
 - Three steps:
 - » Input processing
 - » Forwarding
 - » Output processing
- What are the technical challenges?

OS implementation

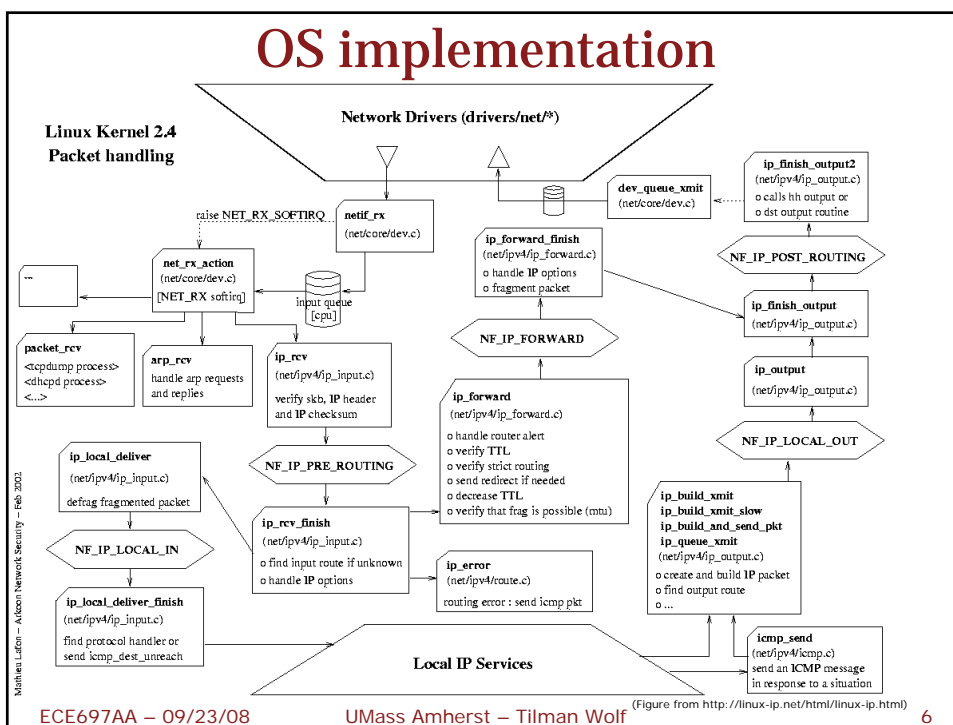


ECE697AA – 09/23/08

UMass Amherst – Tilman Wolf

5

OS implementation



ECE697AA – 09/23/08

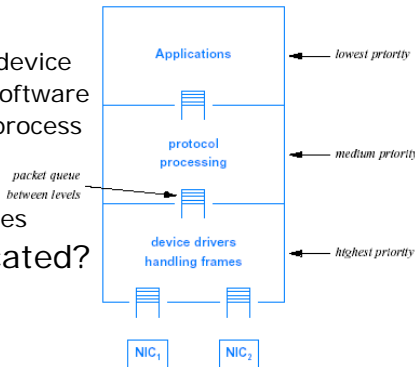
UMass Amherst – Tilman Wolf

(Figure from <http://linux-ip.net/html/linux-ip.html>)

6

Interrupts

- Timing of IPv4 processing is tricky
 - Packet arrivals are asynchronous
- Interrupt triggers processing
 - What is an interrupt?
 - » Event signal
 - » Hardware interrupts: raised by device
 - » Software interrupts: raised by software
 - » "Interrupt handler" is called to process interrupt
 - Interrupt Priorities
 - » Interrupts have different priorities
- How are interrupt levels allocated?
 - Livelock must be avoided



(Figure from Douglas Comer: Network System Design using Network Processors)

Scalability

- What is scalability?
 - A system (design) is scalable if it can easily be extended in "size" and performance
 - » More ports
 - » Faster links
- Scalability important in system design
 - Design process is expensive
 - Ability to easily extend to new requirements is important
 - Performance requirement increase really fast
 - » Moore's Law
 - Systems will eventually be used in a different context
- Is a software router scalable?

Assignments

- Exam I
 - Prepare for next Tuesday
 - Sample problems listed under “assignments” on course web site
- SPARK
 - Assessment quiz