

Meeting of the Network Stack Cabal Intensive Network Stack Review

Robert N. M. Watson
Ed Maste
George Neville-Neil

16 May 2007

FreeBSD Developer Summit
BSDCan 2007



UNIVERSITY OF
CAMBRIDGE

Introduction

- Intensive network stack reading and review
 - Many complex parts, many complex interactions
- Network stack developers
 - Many specialists, few generalists
- Goals
 - Improve general education about stack
 - Motivate general architectural discussions
- Format based on McKusick code-reading class

Architectural Themes

- Locking strategies and granularity
- Opportunities for increased concurrency
- Cache friendliness
- Scheduling interactions
- Work/dispatch model
- Per-packet overhead
- Copy reduction opportunities
- Context switch reduction
- Device driver API
- Hardware integration (TSO, LRO, TOE, ...)

The Immodest Proposal

Start	Min	Topic	Start	Min	Topic
09:10	5	Introduction	13:40	35	UNIX domain sock
09:15	30	mbufs + allocator	14:15	60	Routing, ARP, ND6
09:45	15	netisr	15:15	15	Break
10:00	20	ifnet	15:30	20	IP input/output
10:20	30	Wireless	15:50	15	IP6 input/output
10:50	10	Break	16:05	30	IPSEC
11:00	30	BPF	16:35	15	Break
11:30	15	pfil/firewalls	16:50	30	Netgraph
11:45	15	pf	17:20	20	accept filters
12:00	30	if_cxgb	17:40	70	TCP
12:30	30	Lunch	18:50	10	Wrap-up
13:00	40	Socket + sockbuf	19:00	...	Dinner

The Plan (cont)

- For each topic area:
 - Introduce key concepts and data structures
 - A guided code-reading of key paths
- Remember to take notes
 - Any bugs or problem areas
 - Architectural concerns
- Will try to keep up pace...
 - ... but will see how day evolves