Web100 Status and Future Path

Matt Mathis
mathis@psc.edu

Updated
1 Aug 2002

Outline

- Part 1
  - Problem Statement
  - Status Overview
  - My goals for this meeting
- Part 2
  - Standards strategy
  - Kernel strategy
  - Vendor strategy
  - Linux Mainline strategy

The Problem - The Wizard Gap

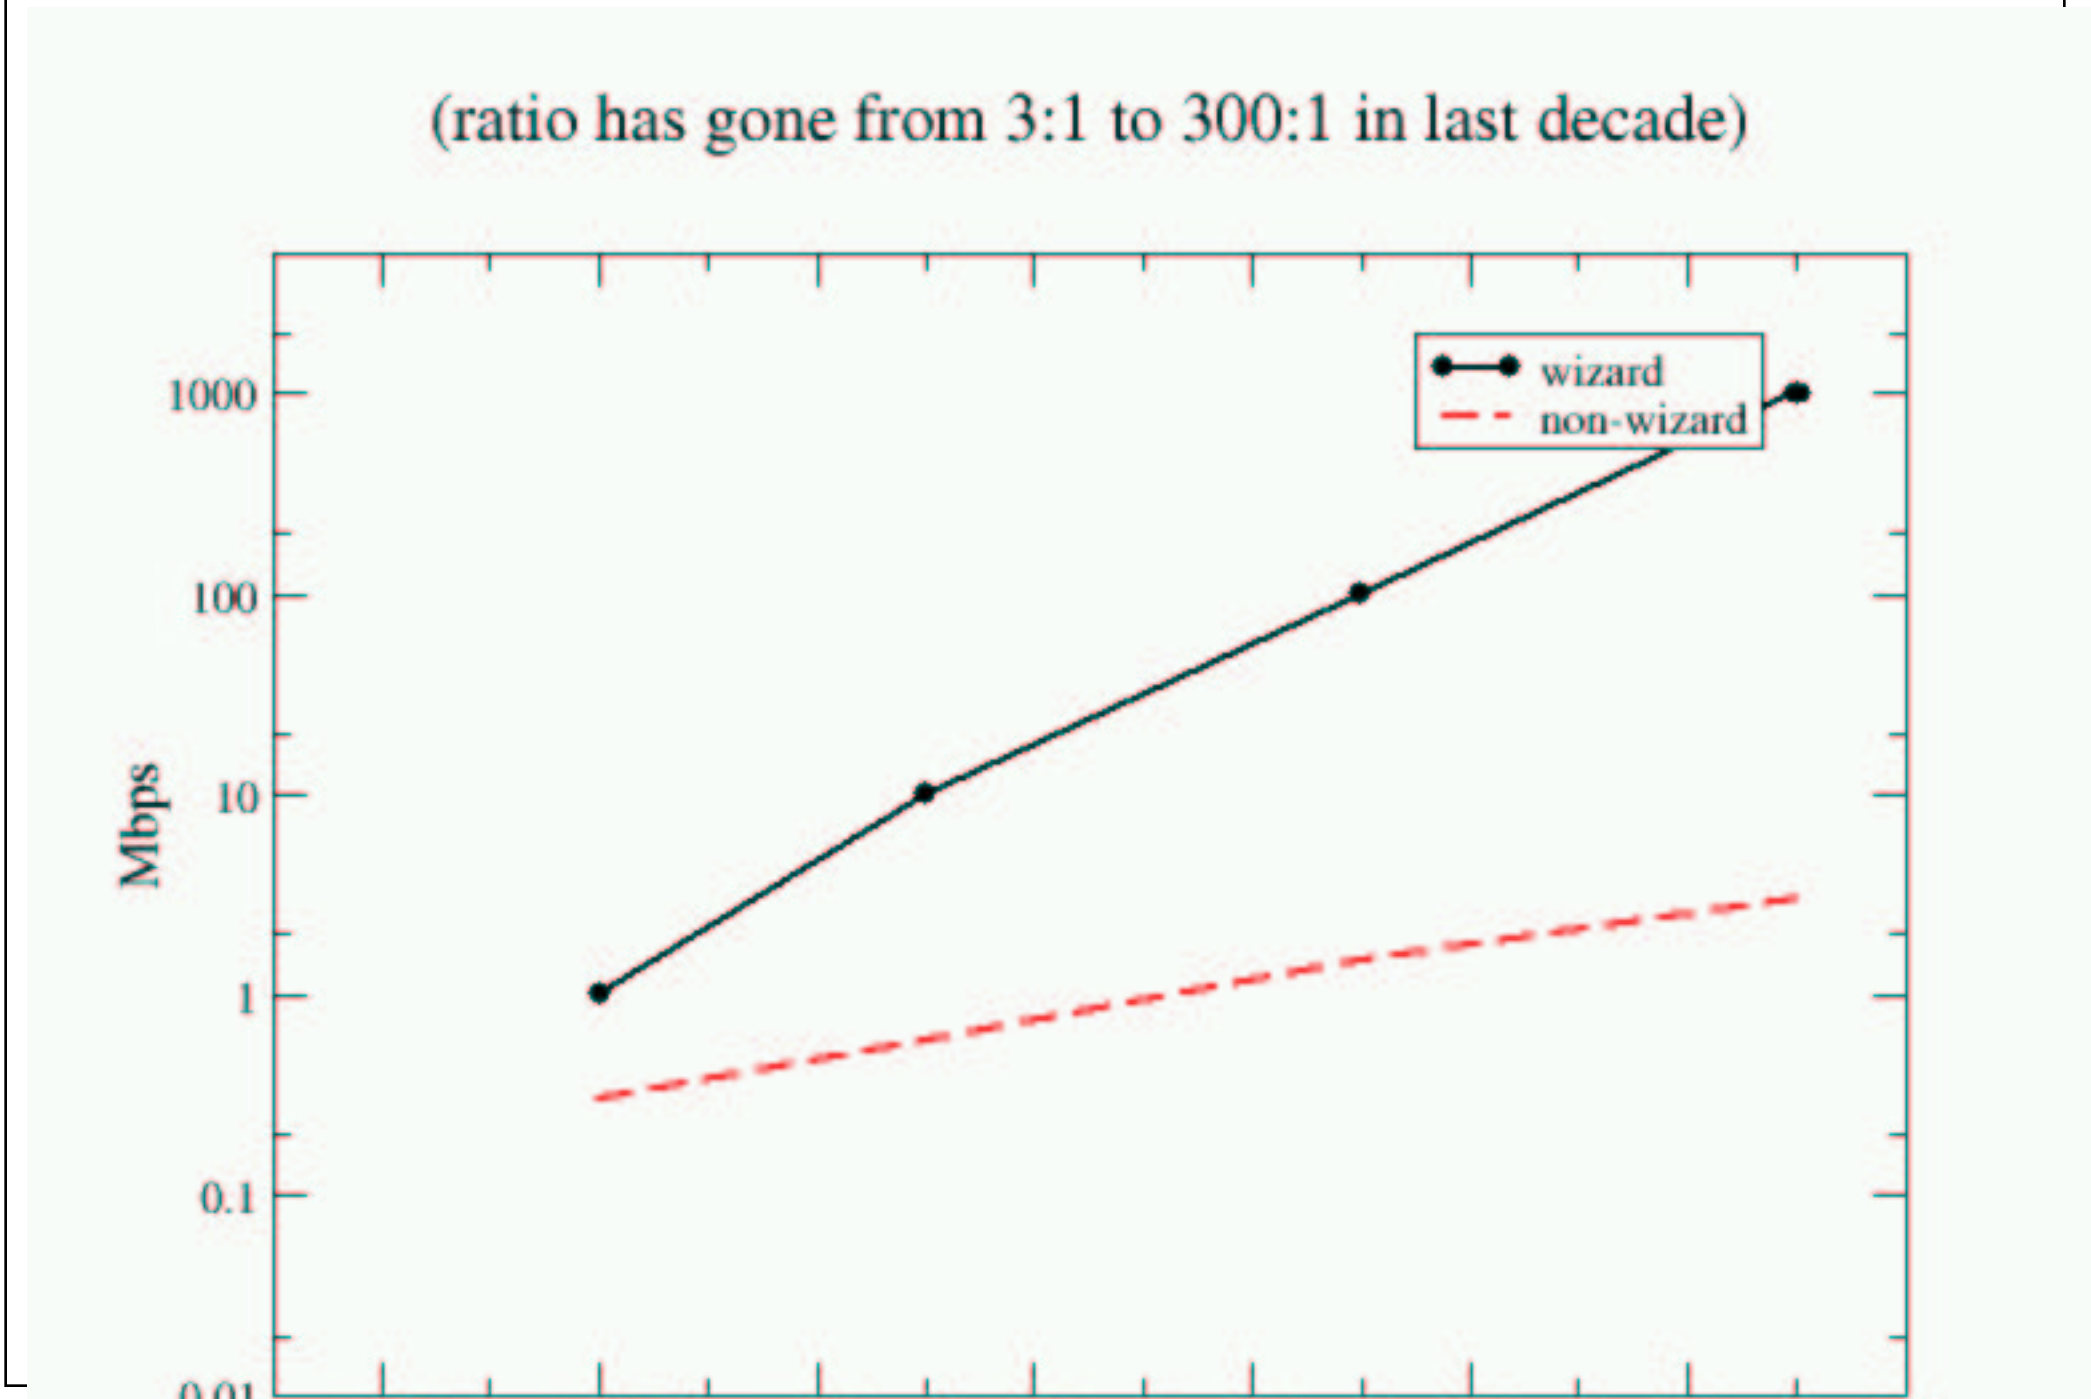

Why? TCP "tuning" requires experts

- TCP/IP hides the net and upper layers for each other
- This is good for the growth of the ’net
  - but it hides all problems
- All bugs have the same symptom: less than expected performance!

TCP tuning is really debugging

- Six classes of bugs limit performance
  - TCP sender or receiver buffer space
  - Packet losses, corruption, congestion, lame HW
  - IP Routing, long round trip times
  - Inefficient applications
  - IP Packet reordering
  - Improper MSS negotiations or MTU discovery

TCP tuning is painful debugging

- Any one problem can mask all other problems confounding all but the best experts
- Akin to finding the weakest link of an invisible chain
- We need better diagnostic visibility

The Web100 Project

- When there is a problem, just ask TCP
  - TCP has an ideal vantage point
  - TCP can identify the bottleneck subsystem
  - TCP already measures the network
  - TCP can measure the application
  - TCP can adjust itself
- This is the whole point of Web100

Status Overview

- Two years into three years of funding
- Open source in the field
  - 5000 lines of kernel code (Matt)
  - 40000 lines of user code (John E)
- About 130 defined instruments (R Reddy)
- Dissemination
  - Standard TCP-MIB via the IETF (Matt)
  - Vendor adoption (Peter)
  - Web pages and user support (Tanya)
- Evaluator/Collaborator feedback (Jim & 9 teams)

My Goal for this meeting

- Help us maximize our impact beyond the project
  - Set priorities for year 3 and beyond
- Web100 has to grow beyond this team

How?

Part 2

Standards Status Future Path

- The ultimate goal is a widely implemented standard
- I-D introduced prior to each of the last 3 IETFs
  - draft-ietf-tsvwg-tcp-mib-extension-01.txt
- Transport AD's very supportive
  - Scott Bradner and Allison Mankin
- Only 2 comments from other people
  - IPv6 MIB management has noticed
  - Plan to start dialog next month
- Gradual transition to "Kernel follows I-D"

Kernel Development: The big picture

- New autotuning loosely follows Feng's DRS
  - To be presented by John Heffner
- Net100 contribution: the WAD
  - the Work Around Daemon
- Version 2.0 plans
  - Backward compatibility issues
- Kernel mainline strategy

The WAD (See DoE Net100 project)

- Framework for non-standard features & tuning
  - See Dunigan et al SC2002 paper
- Most require out-of-band parameters or policy
  - No globally suitable default configuration
  - MUST be disabled by default
- Safe controls are/will be in Web100 releases
  - Lmits on cwnd, rwnd, MSS, etc.
  - Disable dupacks, cwnd validation, etc.
  - Hooks for diagnostics
- Unsafe controls are excluded from Web100 release
  - Alter AIMD parameters

Kernel Plans for version 2.0

- IPv6 support (mostly done)
  - Requires a userland "compatibility patch"
- New autotuning (mostly done)
- Bring into alignment with the newest MIB
  - Changes to about 30 instruments
- Move beyond 2.4.16
- Basic CPU profiling (done)
- Document NETLINK event notification (from Net100)
- Various minor fixes and extensions (done)

Vendor strategy

- IBM
- COMPAQ(HP)
- HP
- Sun
- Linux community
- Microsoft

Mainline source strategy

- Linux
  - Recent harsh language is frightening
  - Planned to start dialog last month but slipped
  - We need to do somthing Real Soon Now
- netBSD, freeBSD, *BSD
  - Consider a consortia to re-implement the kernel
  - Solve GPL problems
  - "Second interoperable implementation" for IETF
  - Foster "classic" workstation adoption
  - Put pressure on Linux community