

Akamai OS War Stories

Bruce Maggs

Carnegie Mellon



15-410 Gratuitous Quote of the Day

Well you're not hardcore
(No you're not hardcore)
Unless you live hardcore
(Unless you live hardcore)
But the legend of the rent was way hardcore

- School of Rock


My role at Akamai

- Joined the company before it was a company (still at MIT, Fall 1998)
- First leader of engineering organization (“VP, Research and Development”)
 - Engineering grew from 10 to 140 employees in under a year
- Now “VP, Research”
 - “Data Driven Design of Distributed Systems”

Major Akamai Customers

- Microsoft (Windows Update)
- Apple (iTunes)
- Anti-worm/virus software vendors
- Images for Yahoo!, MSNBC, Amazon ...
- Google (DNS)

Network Deployment



Akamai Operating Systems

- Started with Red Hat Linux, 2.0.34 kernel (October 1998)
- Deployed Windows 2000 Server, early 2000
- Linux performance optimizations
- “SecureOS” derived from Linux (2003)
 - Battle hardened

Why more than one OS?

- Original plan: half Linux, half FreeBSD
- Windows later added
 - Windows Media Server runs on no other platform
 - Major effort underway to port more services to Windows
- Use of socket-based interactions

Optimizations and Security

- Manage disk and disk cache directly
(Optimize dedicated server for its application)
- Optimize network kernel for short transactions
- Run services in user mode!
- Only one access method: ssh

Reliance on GNU software

- GNU GPL (GNU Public License) excerpts:
Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. **The act of running the Program is not restricted**, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.
- 3. **You may copy and distribute the Program** (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

Themes

- It's our fault if the client's machine doesn't work!
- It's not easy to convince a vendor that their OS is broken.
- Be prepared to fix it yourself.

Steve can't see the new Powerbook

- Steve's assistant Eddie explains the problem
- I spend all night poring through the logs
- Eddie sneaks into Steve's office
- Mystery solved

David is a Night Owl

- Your servers aren't responding!
- Why don't you support half-closed connections?
- Why don't you support "transactional" TCP?
- (Why would transactional TCP be bad for us?)

The “Magg Syndrome”

- We “hijack” a customer’s site?
- I become the most hated person on the Internet
- We isolate the problem (nine months of work)
- Nobody cares?

Don't do this at home

- Irrate end user threatens to go to police
- Akamai is attacking my home system!
- It's in the logs.
- It all began in a Yahoo! chat room
- Have your lawyers call our lawyers

Packet of Death

- Akamai servers take care of each other
- A router in Malaysia is taking down our whole system!
- The mysterious MTU
- The “final” Linux kernel isn’t so final
 - 2.0.36 (Nov. 1998) \Rightarrow 2.0.37 (June 1999)

BIND Miseries

- Open-source DNS server code
- Messy, buggy implementations
- Our customers still run old versions!
- BIND 4.8 TTL issue
 - Refresh attempt when 15 minutes left
 - Success if new list of IP's overlaps with old list of IP's
 - Otherwise, refuse to resolve for next 15 minutes!

We're under attack!

- Someone has cracked our authentication scheme!
- But they haven't got the format of control messages quite right.
- Wait a minute. That's one of ours!
- Where DO these servers disappear to?

What's coming?

- Customers are running IBM Websphere applications on our servers!
- Physical security is now more of an issue.
- Isolation between customer applications?
- Database caching?