Impressions from the 2006 USENIX Annual Technical Conference

Muli Ben-Yehuda

`muli@il.ibm.com`

IBM Haifa Research Lab

Summary

- USENIX ’06: USENIX Annual Technical Conference
  - May 30-June 3, 2006, Boston, MA, USA.
  - Held in conjunction with FAST-OS PI Meeting and Workshop
  - `http://www.usenix.org/events/usenix06/tech/`
- About 700 attendees
- 194 paper submissions, 36 acceptances
- 3 tracks: Systems Practice and Experience Refereed Papers, Invited Talks, Guru Is In Sessions
- Atul Adya, Microsoft and Erich Nahum, IBM T.J. Watson Research Center, Program Co-Chairs
- USENIX ’07 will be June 17-22, Santa Clara, CA.

Optimizing Network Virtualization in Xen

- *Aravind Menon, EPFL; Alan L. Cox, Rice University; Willy Zwaenepoel, EPFL*
- Retains the usual dom0/domU split (no direct access)
- Redefine the virtual network interfaces to incorporate high-level network offload features (even if HW does not support) — checksum offload, S/G IO, TCP segmentation offload
- Optimize data transfer path - no page flipping on tx, replace flipping by copying on rx
- Utilize superpages and and global page mappings.
- Overall improvement in transmit performance of guest domains by a factor of 4.4!

High Performance VM-Bypass IO in VMs

- *Jiuxing Liu, IBM Watson; Wei Huang, Ohio State; Bulent Abali, IBM Watson, Dhabaleswar K. Panda, Ohio State*
- Direct hardware access for faster and more efficient IO from virtual machines, utilizing IB hardware, including direct userspace access.
- Nearly the same raw performance as InfiniBand running in a non-virtualized environment.
- Frontend/Backend drivers for Xen/dom0 control over privileged operations, data path operations carried out directly against the adapter.
- Source code available:

  `http://xenbits.xensource.com/ext/xen-smartio.hg`

Panel: Universities and Industry

- "Is University Systems Teaching and Research Relevant to Industry?"
- *Moderator: Gernot Heiser — Andy Tanenbaum, Orran Krieger, Margo Seltzer, Sun, Intel Research, HP Labs, others?*
- Answer is not obvious at all. . .
- Universities should be teaching students to think, not to work in industry – which they don't
- Are Open Source and Linux alternative education sources?
- Universities in funding crunch, running after the money – too much influence by industry on curriculum (industry agreed)

Invited Talk: Gold and Fool's Gold

- Successes, Failures, and Futures in Systems Research
- *Butler Lampson, Microsoft Research*
- Failure: we didn't invent the web! old, wasteful, flaky and doesn't scale. . .
- Catastrophe mode - write software that expects to fail and handles it gracefully
- Conclusions for Engineers:
  - Understand Moore's law
  - Aim for mass markets — Computers are everywhere
  - Learn how to deal with uncertainty
  - Learn how to avoid catastrophe

Invited Talk: Pixar Technology

- Why Mr. Incredible and Buzz Lightyear Need Better Tools: Pixar and Software Development
- *Greg Brandeau, VP of Technology, Pixar Animation Studios*
- “Cars” rendered on a cluster of x86-64 machines running Linux (Fedora Core)
- Moved to x86-64 to break the 4GB wall
- Using Linux because it’s the best they have, but it’s by no means good enough
- Development tools are awful — gdb was someone’s MSc thesis 20 years ago — can we have dtrace for Linux please?
- Perfectly happy to pay for improvements, but can’t find anyone to do it

Replay Debugging for Distributed Apps

- *Dennis Geels, Gautam Altekar, Scott Shenker, and Ion Stoica, University of California, Berkeley*
- `liblog` is new lightweight replay and debugging tool for distributed C/C++ applications
- `liblog` logs and replays deterministically the execution of distributed applications
- Consistent replay of arbitrary subsets of nodes
- Mixed environment of logging and non-logging processes
- No hardware of kernel patches needed — LD_PRELOAD based

Runtime Hardware Reconfiguration

- System- and Application-level Support for Runtime Hardware Reconfiguration on SoC Platforms
- *Dimitris Syrivelis and Spyros Lalis, University of Thessaly, Hellas*
- Enable programs running on a reconfigurable SoC to modify the underlying FPGA at runtime
- Applications may request he addition and/or removal of softcore devices at any point in time
- System reconfiguration via a fast suspend-resume mechanism — entire FPGA is reconfigured, including softcore CPU!
- Reconfiguration is transparent to the application; device drivers deal with the disappearing and reappearing devices

IP Only Server

- *Muli Ben-Yehuda, Oleg Goldshmidt, Elliot K. Kolodner, Zorik Machulsky, Vadim Makhervaks, Julian Satran, Marc Segal, Leah Shalev, and Ilan Shimony, IBM Haifa Research Laboratory*
- Poster and short paper presented by Oleg to a very warm reception
- “If this was a product, I’d buy it”
- “Google has several departments who would be very interested in this — and by the way, we’re hiring”
- Great feedback, USENIX is a good venue to present at
- Let’s do it again next year...