Hype and Virtue

Timothy Roscoe, Kevin Elphinstone and Gernot Heiser

Presented by Neal H. Walfield

July 18, 2007

Question

- *Are hypervisors a disruptive force in OS research?*

Claim

- No.
- *But*
  - Enabler for new ideas
  - Source of a new class of problem

Question

- *Are hypervisors a disruptive force in OS research?*

Claim

- No.
- *But*
  - Enabler for new ideas
  - Source of a new class of problem

Implicit Question

- What are the goals of OS research?

Hypervisor Research

- Novel applications
- Better hypervisors
  - Sharing and Protection
  - Communication
  - Abstraction

Novel Applications

- Replay debugging
- Security mechanisms

- How to solve at OS level?
- Requirements for OS design?

Novel Applications

- Replay debugging
- Security mechanisms

- How to solve at OS level?
- Requirements for OS design?

Sharing and Protection

- Sharing: anything new here?
- Protection: VM essentially an address space

Communication

- Define problem away
  - VMs are largely self-contained
  - No shared servers
  - Ethernet suffices
- Counter-examples
  - Driver servers
  - Parallax (storage server)

Communication

- Define problem away
  - VMs are largely self-contained
  - No shared servers
  - Ethernet suffices
- Counter-examples
  - Driver servers
  - Parallax (storage server)

Abstraction

- VM = explicit kernel representation of an application
- VM = abstraction of an application

- Virtual appliances
  - VM contains and isolates a single application

Abstraction

- VM = explicit kernel representation of an application
- VM = abstraction of an application
- Virtual appliances
  - VM contains and isolates a single application

Virtual Hardware as an API

- A simple interface?
  - Implementation complex
  - Interface complex
  - Semantic bottleneck

Implementation Complexity

- Harder to program $\Longrightarrow$ more code
- Duplicated code
- Harder to keep up to date

Interface Complexity

- Hard to specify
- Hard to reason about

Semantic Bottleneck

- How to efficiently express complex requests?
  - Loss of semantic information

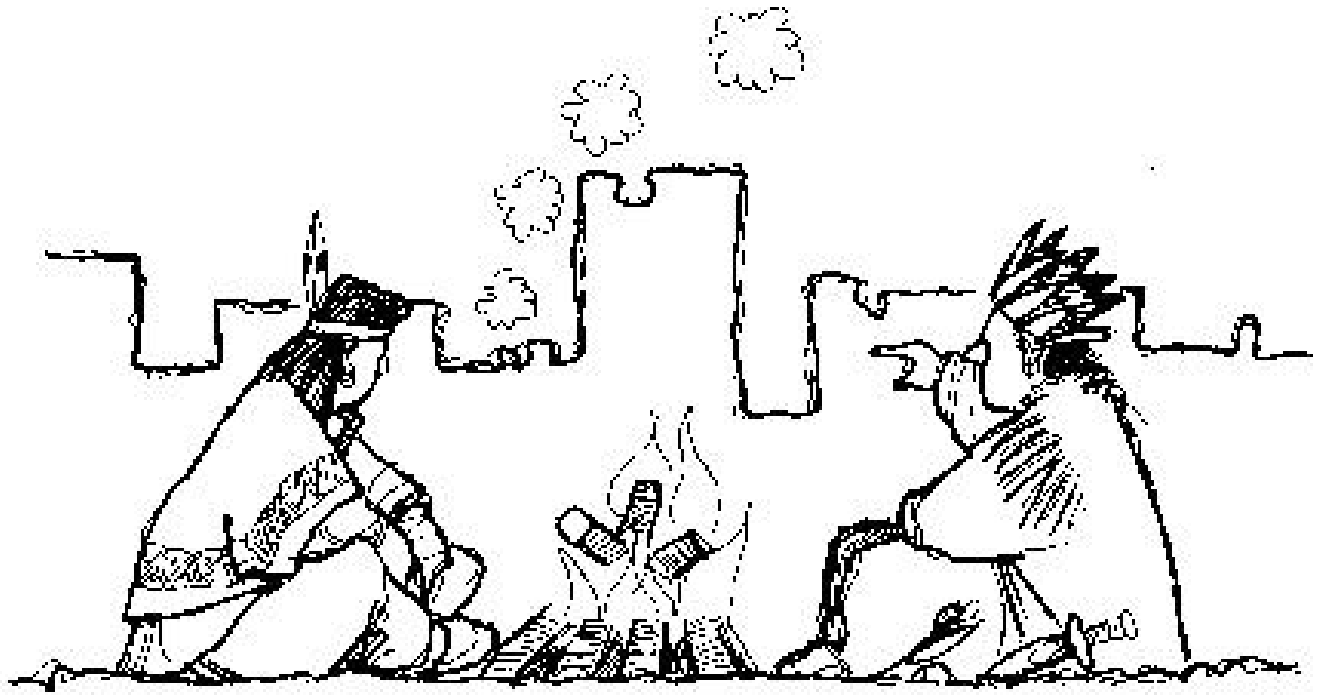

http://tscp.open.ac.uk/t183/images/l2cartoon2.gif

Summary

- Hypervisors expose clumsy, hardware ABI
- Research doesn't try to formulate OS interfaces

What Hypervisors are Good For

- Legacy Support
- Device Drivers

Virtualization: an Enabler

- Experiment with new interfaces
- Gain practical experience
- Return to basic research
  - System organization
  - Effective multiplexing of machine resources

Research Possibilities

- Shared interface to MMU
  - Many parties managing shared physical resources
  - GC
- Formal Verification
  - seL4: Haskell runtime ported to kernel interface
- Killer-applications
- Metrics

Conclusion

- Goal of OS research ought to be
  - Understanding requirements
  - Formulating better OS API
- Research on hypervisors ignores these questions
- Hypervisors are enablers

Discussion Points

- Agree or Disagree?
  - Systems community too focused on Hypervisors?
  - Goal of research is better APIs?
  - What are the long-term research goals?
  - Is system research dead?