

Resource Virtualization

Reading List

Note: We will not read all of these papers in class. The syllabus is the final word on the specific papers that we shall read in class.

Most of these papers are available from the web (use <http://scholar.google.com> and <http://citeseer.ist.psu.edu> to find them. I will make photocopies of the older, non-web papers available as needed.)

Book

J. Smith, R. Nair, *Virtual Machines: Versatile Platforms For Systems And Processes*, Morgan Kaufmann, 2005.

Collection

- R. Figueiredo, P. Dinda, J. Fortes, Editors, *Resource Virtualization Renaissance*, IEEE Computer, May, 2005. This issue includes these papers
- J. Smith, R. Nair, *The Architecture of Virtual Machines*
 - M. Rosenblum, T. Garfinkel, *Virtual Machine Monitors: Current Technology and Future Trends*
 - R. Uhlig, et al, *Intel Virtualization Technology*
 - A. Whitaker, et al, *Rethinking the Design of Virtual Machine Monitors*,
 - P. Ruth, et al, *Virtual Distributed Environments in a Shared Infrastructure*
 - P. Magnusson, *The Virtual Test Lab*

Virtualization-based Computing Environments

1. F. Corbato, V. Vyssostsky, *Introduction and Overview of the Multics System*, Fall Joint Computer Conference, 1965. Note: Multics was **not** a VM-based computing environment. Nonetheless, it gives you an idea of the “vision” of many of these environments.
2. C. Waldsburger, T. Hogg, *Spawn: A Distributed Computational Economy*, IEEE Transactions on Software Engineering, 18(2), 1992. Note: not VM-based, but an excellent paper on computation economies.
3. R. Figueiredo, P. Dinda, J. Fortes, *A Case for Grid Computing on Virtual Machines*, ICDCS 2003.
4. S. Hand, et al, *Controlling the Xenoserver Open Platform*, OpenARCH 2003.
5. K. Fraser, et al, *The Xenoserver Computing Infrastructure*, Technical Report UCAM-CL-TR-552, University of Cambridge, 2003.
6. C. Sapuntzakis, et al, *Virtual Appliances for Deploying and Maintaining Software*, LISA 2003.

7. T. Garfinkel, et al, *Terra: A Virtual Machine-based Platform for Trusted Computing*, SOSP 2003.
8. X. Jiang, D. Xu, *SODA: A Service-on-demand Architecture for Application Service Hosting Utility Platforms*, HPDC 2003.
9. A. Shoykhet, et al, *Virtuoso: A System for Virtual Machine Marketplaces*, Technical Report NWU-CS-04-39, Northwestern University, 2004.
10. R. Potter, *One-click Distribution of Preconfigured Linux Runtime State*, USENIX VM 2004 (WIP).
11. A. Bavier, et al, *Operating System Support for Planetary-scale Network Services*, NSDI 2004.
12. S. Adablala, et al, *From Virtualized Resources to Virtual Computing Grids: The In-VIGO System*, Future Generation Computer Systems 21:6, April, 2005
13. R. Chandra, et al, *The Collective: A Cache-based System Management Architecture*, NSDI 2005.
14. P. Ruth, et al, *VioCluster: Virtualization for Dynamic Computational Domains*, Cluster 2005.
15. M. Satyanarayanan, et al, *Towards Seamless Mobility on Pervasive Hardware*, Pervasive and Mobile Computing 1:2, June, 2005. The home page of the ISR deployment at CMU is at <http://isr.cmu.edu/>
16. K. Keahey, et al, *Virtual Workspaces in the Grid*, Euro-Par 2005.
17. A. Maccabe, et al, *Panel on Whole System Virtualization In High End Computing Systems*, HPDC 2005. (slides available from <http://www.caip.rutgers.edu/hpdc2005/program.html>)

Traditional OS-level Virtual Machines

18. R. Goldberg, *Architecture of Virtual Machines*, Proceedings of the Workshop on Virtual Computer Systems, Cambridge, Mass, 1973, pp 74-11
19. R. Goldberg, *Survey of Virtual Machine Research*, IEEE Computer, June 1974, pp 34-45.
20. L. Seawright, R. MacKinnon, *VM/370: A Study of Multiplicity and Usefulness*, IBM Systems Journal 18(1), 1979.
21. R. Creasy, *The Origin of the VM/370 Time-Sharing System*, IBM Systems Journal 25(5), 1981.
22. VM/370: The subject of IBM Systems Journal 18(1), 1979 (9 articles)
23. VM/386 (<http://www.igcinc.com/vm386.htm>)
24. Virtual 8086 mode: <http://www.openbg.net/sto/os/xml/v86.html>
25. P. Karger, et al, *A Retrospective on the VAX VMM Security Kernel*, IEEE Transactions on Software Engineering, November, 1991.
26. K. Lawton, *Running Multiple Operating Systems Concurrently on an IA32 PC Using Virtualization Techniques*.
27. Plex/86: <http://plex86.sourceforge.net/>
28. J. Nieh, O. Leonard. *Examining VMWare*, Dr. Dobb's Journal, August 2000.
29. C. Waldsburger, *Memory Resource Management in VMWare ESX Server*, OSDI 2002.
30. VMWare, Inc., *Introducing the VMWare Virtual Platform*

31. Devine, S., et al, *Virtualization System Including A Virtual Machine Monitor For a Computer With A Segmented Architecture*, U.S. Patent #6397242. (VMWare Patent)
32. J. Honneycutt, *Microsoft Virtual PC 2004 Technical Overview*.

Paravirtualized Virtual Machines

33. D. Engler, et al, *Exokernel: An Operating System Architecture for Application-Level Resource Management*, SOSP 1995.
34. E. Bugnion, et al, *Disco: Running Commodity Operating Systems on Scalable Multiprocessors*, ACM Transactions on Computer Systems, 18(3), August 2000.
35. J. Dike, *A User-mode Port of the Linux Kernel*, Linux Showcase and Conference 2000. (see also <http://user-mode-linux.sourceforge.net/>)
36. H. Hoxer, et al, *Implementing a User-mode Linux with Minimal Changes from the Original Kernel*, 2002 Linux System Technology Conference.
37. Whitaker, et al, *Scale and Performance in the Denali Isolation Kernel*, OSDI 2002.
38. P. Barham, et al, *Xen and the Art of Virtualization*, SOSP 2003.
39. S. Hand, et al, *Are Virtual Machine Monitors Microkernels Done Right?*, HOTOS X (2005).

Architectural Issues

40. E. Mallach, *On the Relationship Between Virtual Machines and Emulators*, Proceedings of the Workshop on Virtual Computer Systems, Cambridge, Mass, 1973, pp 117-126.
41. G. Popek, R. Goldberg, *Formal Requirements for Virtualizable Third Generation Architectures*, Communications of the ACM, 17(7), July, 1974, pp 413-421.
42. J. Hall, P. Robinson, *Virtualizing the VAX Architecture*, Proceedings of ISCA 1991, pp 380-389
43. J. Robin, C. Irvine, *Analysis of the Intel Pentium's Ability to Support a Secure Virtual Machine Monitor*, Proceedings of of the 9th USENIX Security Symposium, Denver, CO, August 2000.
44. Intel Corporation, *Intel Vanderpool Technology for IA-32 Processors (VT)*, <http://www.intel.com/technology/computing/vptech/>
45. AMD Corporation, *Pacifica Virtualization Extensions*, <http://enterprise.amd.com/Enterprise/serverVirtualization.aspx>

Emulation and Binary Translation

46. E. Mallach, *On the Relationship Between Virtual Machines and Emulators*, Proceedings of the Workshop on Virtual Computer Systems, Cambridge, Mass, 1973, pp 117-126.
47. R. Sites, et al, *Binary Translation*, Communications of the ACM, 36(2), February, 1993.
48. P. Magnusson, et al, *Simics: A Full System Simulation Platform*, IEEE Computer 35(2), 2002.
49. Microsoft Virtual PC for Mac
50. Bochs IA-32 Emulator: <http://bochs.sourceforge.net/>
51. F. Bellard, *QEMU: A Fast and Portable Dynamic Translator*, USENIX 2005.

Language-level Virtual Machines

52. W. Dobrusky, T. Steel, *UNCOL: Universal Computer-Oriented Language*, Communications of the ACM 4(3), 1961. (abstract only)
53. P. Bagley, *Principles and Problems of a Universal Computer-Oriented Language*, The Computer Journal 4(4), 1962.
54. UCSD p-system: <http://www.threedee.com/jcm/psystem/>,
<http://www.ics.uci.edu/~archive/documentation/p-system/p-system.html>
55. T. Lindholm, F. Yellin, *The Java Virtual Machine Specification, 2nd Edition*. Available from <http://java.sun.com/docs/books/vmspec/2nd-edition/html/VMSpecTOC.doc.html>
56. A. Adl-Tabatabai et al, *Fast, Effective Code Generation in a Just-in-Time Java Compiler*, PLDI 1998.
57. E. Meijer, J. Gough, *Technical Overview of the Common Language Runtime*.
58. Microsoft .NET/CLR documentation:
<http://msdn.microsoft.com/netframework/>
59. W. Vogels, *HPC.NET: Are CLI-based Virtual Machines Suitable for High Performance Computing?*, SC 2003.

Virtual Servers

60. Ensim, Inc, *Virtual Private Servers*
61. Free VSD: <http://www.freevsd.org>
62. Linux V-server: <http://www.linux-vserver.org/>
63. Berkeley (BSD) Jails
64. Sun Solaris Zones

Virtual Devices

65. J. Sugarman, G. Venkitachalam, B-H Lim, *Virtualizing I/O Devices on VMWare Workstation's Virtual Machine Monitor*, USENIX 2001.

66. S. King, et al, *Operating System Support for Virtual Machines*, USENIX 2003.
67. K. Fraser, et al, *Reconstructing I/O*, Technical Report UCAM-CL-TR-596, University of Cambridge.
68. J. LeVasseur, et al, *Unmodified Device Driver Reuse and Improved System Dependability Via Virtual Machines*, OSDI 2004.
69. A. Whitaker, et al, *Constructing Services With Interposable Virtual Hardware*, NSDI 2004.
70. A. Menon, et al, *Diagnosing Performance Overheads in the Xen Virtual Machine Environment*, VEE 2005.
71. L. Cherkasova, R. Gardner, *Measuring CPU Overhead for I/O Processing in the Xen Virtual Machine Monitor*, USENIX 2005.

Virtual Storage / SANs

72. G. Alvarez, et al, *MINERVA: An Automated Resource Provisioning Tool for Large-Scale Storage Systems*, ACM Transactions on Computer Systems 19(4), 2001.
73. M. Beck, et al, *An End-to-end Approach to Globally Scalable Network Storage*, SIGCOMM 2002. (see also <http://loci.cs.utk.edu/>)
74. C. Lumb, et al, *Façade: Virtual Storage Devices with Performance Guarantees*, USENIX FAST 2003.
75. L. Huang, et al, *Multidimensional Storage Virtualization*, SIGMETRICS 2004.
76. Wikipedia, *Storage Area Networks*, http://en.wikipedia.org/wiki/Storage_area_network. Brief Intro. Includes pointer to the excellent and highly detailed IBM Redbook on SANs.

Virtual Services

77. J. Weissman, B. Lee, *The Virtual Service Grid: an architecture for delivering high-end network services*, Concurrency: Practice and Experience, 14(4), 2002.
78. M. Zhao, et al, *Distributed File System Support For Virtual Machines in Grid Computing*, HPDC 2004
79. I. Foster, et al, *The Open Grid Services Architecture, Version 1.0*, Global Grid Forum Informational Draft, 2005.

Virtual Networking and Overlays

80. P. Ferguson, G. Huston, *What is a VPN?*, Technical Report, Cisco Systems, March 1998.
81. B. Gleeson, et al, *A Framework for IP-based Virtual Private Networks*, IETF RFC 2764, February 2000.
82. G. Italiano, et al, *Restoration Algorithms for Virtual Private Networks in the Hose Model*, INFOCOM 2002.

83. IEEE 802.1Q Working Group, *802.1q: Virtual LANs*, IEEE, 2001.
84. J. Jannotti, et al, *Overcast: Reliable Multicasting With An Overlay Network*, OSDI 2000.
85. D. Andersen, et al, *Resilient Overlay Networks*, SOSP 2001.
86. Y-H Chu, et al, *A Case For End-System Multicast*, IEEE JSAC 20:8, 2002.
87. S. Shi, J. Turner, *Routing In Overlay Multicast Networks*, INFOCOM 2002.
88. J. Rosenberg, et al, *STUN: Simple Traversal Of User Datagram Protocol (UDP) Through Network Address Translators (NATs)*, IETF RFC 3489, 2003.
89. S. Banerjee, et al, *Resilient Multicast Using Overlays*, SIGMETRICS 2003.
90. A. Sundararaj, P. Dinda, *Towards Virtual Networks for Virtual Machine Grid Computing*, USENIX VM 2004.
91. X. Jiang, D. Xu, *VIOLIN: Virtual Internetworking on Overlay Infrastructure*, Technical Report CSD TR 03-027, Department of Computer Science, Purdue University.
92. M. Walfish, et al, *Middleboxes No Longer Considered Harmful*, OSDI 2004.
93. L. Subramanian, I. Stoica, *OverQoS: An Overlay-based Architecture for Enhancing Internet QoS*, NSDI 2004.
94. S. Birrer, et al, *Building a Resilient Multisource Multicast Fat Tree*, WCCD 2004.

Inference, Measurement, Adaptation, and Reservation

95. W. Tetzlaff, *State Sampling of Interactive VM/370 Users*, IBM Systems Journal 18(1), 1979.
96. A. Gupta, P. Dinda, *Inferring the Topology and Traffic Load of Parallel Programs Running in a Virtual Machine Environment*, JSSPPP 2004.
97. A. Sundararaj, et al, *Increasing Application Performance In Virtual Environments Through Run-time Inference and Adaptation*, HPDC 2005
98. J. Lange, et al, *Automatic Dynamic Run-time Optical Network Reservations*, HPDC 2005.
99. B. Lin, P. Dinda, *VSched: Mixing Batch and Interactive Virtual Machines Using Periodic Real-time Scheduling*, SC 2005.
100. J. Xu, et al, *Towards Autonomic Virtual Applications in the In-VIGO System*, ICAC 2005
101. A. Gupta, et al, *Free Network Measurement for Adaptive Virtualized Distributed Computing*, IPDPS 2006.
102. B. Lin, P. Dinda, *Putting the User in Direct Control of CPU Scheduling*, To Appear.

Security and Reliability

103. S. Madnick, J. Donovan, *Application and Analysis of the Virtual Machine Approach to Information System Security and Isolation*, Proceedings of the Workshop on Virtual Computer Systems, Cambridge, Mass, 1973, pp 210-224.
104. G. Popek, C. Kline, *A Verifiable Protection System*, Proceedings of the International Conference on Reliable Software, LA, CA, 1975, pp 294-304.
105. T. Bressoud and F. Schneider, *Hypervisor-based Fault Tolerance*, ACM Transactions on Computer Systems, 14(1), 1996.
106. P. Chen, B. Noble, *When Virtual is Better Than Real*, HOT-OS 2001.
107. G. Dunlop, et al, *ReVirt: Enabling Intrusion Analysis Through Virtual Machine Logging and Replay*, OSDI 2002.
108. T. Garfinkel, M. Rosenblum, *A Virtual Machine Introspection-based Architecture for Intrusion Detection*, NDSS 2003.
109. P. Dinda, *Addressing the Trust Asymmetry Problem in Grid Computing Using Encrypted Computation*, LCR 2004.
110. S. King, et al, *Debugging Operating Systems With Time-traveling Virtual Machines*, USENIX 2005.
111. T. Hirofuchi, et al, *USB/IP – A Peripheral Bus Extension For Device Sharing Over IP Networks*, USENIX 2005.

Migration of Virtual Machines

112. A. Muthitachoen, et al, *A Low-bandwidth Network File System*, SOSF 2001.
113. M. Kozuch, et al, *Efficient State Transfer for Internet Suspend/Resume*, Technical Report IRP-TR-02-03, Intel Research Pittsburgh, 2002.
114. C. Sapuntzakis, et al, *Optimizing the Migration of Virtual Computers*, OSDI 2002.
115. S. Osman, *The Design and Implementation of Zap: A System for Migrating Computing Environments*, OSDI 2002.
116. T. Boyd, P. Dasgupta, *Process Migration: A Generalized Approach Using a Virtualizing Operating System*, ICDCS 2002.
117. C. Clark, et al, *Live Migration of Virtual Machines*, NSDI 2005.
118. M. Nelson, et al, *Fast Transparent Migration for Virtual Machines*, USENIX 2005.

Remote Display

119. P. Romano, *ITU-T Recommendation T.128 (Application Sharing)*, ITU Technical Report, 1997.
120. T. Richardson, et al, *Virtual Network Computing*, IEEE Internet Computing 2:1, 1998.
121. B. Schmidt, et al, *The Interactive Performance of SLIM: A Stateless Thin-client Architecture*, SOSP 1999.
122. Microsoft, *Remote Desktop Protocol (RDP) Features and Performance*, 2000.
123. A. Lai, J. Nieh, *Limits of Wide-Area Thin-client Computing*, SIGMETRICS 2002.
124. R. Baratto, et al, *THINC: A Virtual Display Architecture for Thin-Client Computing*, SOSP 2005.

General Reference

125. B. Lampson, *Hints for Computer System Design*, ACM SIGOPS 15:5, October, 1983.
126. J. Salzer, et al, *End-to-end Arguments in System Design*, ACM TOCS 2:4, 1984.

Humorous

127. The SCIGen Automatic CS Paper Generator and the story behind it, <http://pdos.csail.mit.edu/scigen/>