

Index

- 3D support, 257
- 8086, 27
- ACPI, *see* Advanced Configuration and Power Interface
- Adding new devices, 187
- Address Space ID, 238
- Administration tools, 200
- Advanced Configuration and Power Interface, 259
- Advertising devices, 187
- AMD-V, 13
- ASID, *see* Address Space ID
- Asynchronous notification, 119
- Atropos scheduler, 219
- Behavior sensitive instructions, 4
- Binary rewriting, 10
- BIOS, 47, 244
- blkif, *see* Virtual block device
- blkif_front_ring_t, 163
- blkif_request_segment, 165
- blkif_request_t, 165
- blkif_sring_t, 163
- Block cache, 166
- Block device, 36
 - connecting, 163
 - initializing, 162
 - loading data, 167
 - storing data, 165
 - supporting CDs, 177
 - XenStore nodes, 162
- Boot firmware, 47
- Boot trampoline, 40
- Booting, 27
- Borrowed virtual time scheduler, 219
- Breakpoints, 10
- C bindings, 200
- Cache flushing, 92
- Calling convention, 30
- Calling convention, hypercall, 12
- CD drives, 177
- CIM, *see* Common Information Model
- CIM-XML, 210
- Common Information Model, 209
- Compartmentalization, 7
- Console, 49
- Console device driver, 112
- Console driver, 133
- Console interrupt, 123
- Context switch, 4
- Control sensitive instructions, 4
- Core devices, 161
- CPU architectures, 256
- CPU Virtualization, 4
- CPUID in HVM mode, 236
- Credit scheduler, 219, 222
- Cryptographic coprocessors, 267
- DEC Alpha, 4
- Desktop Xen, 257
- DEV, *see* Device Exclusion Vector
- Device drivers, 99
 - block device, 161

- console, 112
- framebuffer, 178
- network interface, 169
- PCI, 184
- TPM, 183
- USB, 186
- XenStore, 150
- Device Exclusion Vector, 14, 239
- Device multiplexing, 100
- Device scheduling, 224
- Direct Memory Access, 6
- Distributed Management Task Force, 209
- DMA, 6
- DMTF, *see* Distributed Management Task Force
- dom0, *see* Domain 0
- Domain, 19
- Domain 0, 19
- Domain 0 devolution, 262
- Domain builder, 273
- Domain U, 19
- Domain virtual time, 53
- domctl hypercall, 228
- domU, *see* Domain U, *see* unprivileged domain
- DOS, 9
- Driver domains, 102
- E2V, *see* Emulated to virtual
- Earliest deadline first, 221
- EDF, *see* Earliest deadline first
- EFI, *see* Extended Firmware Interface
- Emulated devices, 245
- Emulated to virtual, 255
- Emulation, 3
- Event channel
 - bitfields, 51
- Event channels
 - assigning to a VCPU, 127
 - binding, 125
 - closing, 129
 - masking, 130
 - polling, 133
 - querying status, 129
 - signalling, 128
- Event ports, 33
- Event trampoline, 134
- Event types, 123
- Event upcall, 137
- event_channel_op hypercall, 124
- Events, 111, 119
- Extended Firmware Interface, 18
- Firmware, boot, 47
- Flush TLB, 92
- Flushing cache, 92
- GART, *see* Graphics Address Remapping Table
- GDT, *see* Global descriptor table
- gettimeofday(), 54
- Global descriptor table, 76
- GMFN, *see* Guest machine frame number
- gnttab_copy_t, 64
- gnttab_map_grant_ref_t, 62
- gnttab_transfer_t, 64
- GPFN, *see* Guest page frame number
- Grant reference, 35, 61
- Grant table operations, 66
- Grant tables, 34
- grant_entry_t, 68
- grant_table_op hypercall, 61
- Graphics Address Remapping Table, 7
- Guest loader, 39
- Guest machine frame number, 79
- Guest page frame number, 79
- Handling events, 134
- Hardware assisted
 - device pass-through, 239
 - DMA safety, 239
 - page tables, 238

- real mode, 239
- virtual CPUs, 238
- virtual interrupt routing, 240
- Hardware page tables, 75
- Hardware virtual machine, 22, 23, 29, 235
- Heterogeneous multicore, 266
- Hibernate, 260
- HVM, *see* Hardware virtual machine, *see* Hardware virtual machine
- HVM hypercalls, 236
- `hvm_function_table`, 248
- `hvm_op` hypercall, 246
- Hybrid virtualization, 14, 240
- Hypcall calling convention, 12
- Hypercall, 11, 30
- Hypercall API, 197
- Hypercall macro, 30
- Hypercall page, 12, 30
- Hypercall page setup, 248
- Hypercalls
 - `domctl`, 228
 - `event_channel_op`, 124
 - `grant_table_op`, 61
 - `hvm_op`, 246
 - `memory_op`, 84, 93
 - `mmu_update`, 89
 - `mmuext_op`, 91
 - `multicall`, 277
 - `sched_op`, 132
 - `set_gdt`, 93
 - `update_va_mapping`, 90
 - `update_va_mapping_otherdomain`, 90
 - `vm_assist`, 83
- Hypercalls, HVM, 236, 243
- Hypervisor-based copy, 64
- I/O rings, 65, 103, 164, 188
- I/O virtualization, 36
- IDT, *see* Interrupt descriptor table
- Infiniband, 175
- Input/Output Memory Management Unit, 6, 102, 239
- Interdomain communication, 59
- Interdomain events, 123
- Interprocess communication, 34, 59
- Interprocessor interrupts, 124
- Interrupt 80h, 11, 120
- Interrupt 82h, 30
- Interrupt descriptor table, 94, 120
- Interrupt handlers, 274
- Interrupt vector, 28
- Interrupts, 111, 119
- Intradomain events, 124
- Invalidate TLB entry, 92
- IOMMU, *see* Input/Output Memory Management Unit
- IPC, *see* interprocess communication
- IPIs, *see* Interprocessor interrupts
- IVT, *see* Virtualization Technology for x86
- Jumbo frames, 173
- Kernel header, 39
- KQEMU, 236
- Laptops, 259
- LDT, *see* Local descriptor table
- `libcurl`, 205
- `libvirt`, 200
- `libxen`, 201
- `libxml`, 206
- Local descriptor table, 76
- Lockless ring buffer, 103
- Mach ports, 33
- Machine frame number, 79
- Mapping memory, 61
- Mapping the XenStore, 150
- Masking events, 130
- memory, 59
- Memory assists, 82
- Memory barriers, 107

- Memory Management Unit, 5, 75, 238
- Memory model, 75
- Memory model, Xen, 78
- Memory pages
 - copying, 34, 64, 174
 - exchanging, 86
 - granting access, 66
 - mapping, 61
 - protecting, 77
 - sharing, 34
 - transferring, 34, 63, 174
- Memory protection, 77
- memory_op hypercall, 84, 93
- Message passing, 34
- MFN, *see* Machine frame number
- Migration, 94
- Minix, 256
- mmap(), 60
- MMU, *see* Memory Management Unit
- mmu_update hypercall, 89
- mmuext_op, 91
- mmuext_op hypercall, 91
- Mobile systems, 259
- Model-specific register, 243
- Mouse tracking, 181
- MSR, *see* model-specific register
- multicall hypercall, 277
- multicall_entry, 277
- MULTICS, 28
- Multitasking, 7, 217

- Native device drivers, 184
- Nemesis Exokernel, 264
- Nested Page Tables, 14, 238
- NetBSD, 69
- NetChannel2, 174
- netif, *see* Virtual network interface
- netif_extra_info, 172
- netif_rx_response_t, 173
- netif_tx_request_t, 171
- Network device
 - initializing, 169
- Network interface
 - receiving, 173
 - transmitting, 171
- New devices, adding, 187
- Non-Uniform Memory Architecture, 265
- Nonlocal I/O, 266
- NPT, *see* Nested Page Tables
- NUMA, *see* Non-Uniform Memory Architecture

- Operating system support, 255

- P2V, *see* Physical to virtual
- Pacifica, 13
- Page directory, 78
- Page directory base register, 78
- Page directory entry, 78
- Page fault handling, 94
- Page faults, 241
- Page frame number, 79
- Page table entry, 78
- Page table management, 275
- Page table updates, 89
- Page tables, 78
- Page tables, nested, 14
- Page tables, shadow, 14, 82
- Page tables, writable, 82
- Paged real mode, 239
- PALCode, 4
- Paravirtual I/O on HVM, 246
- Paravirtualization, 10
- Partial relocation, 267
- PCI devices, 184
- PDA's, 259
- PDE, *see* Page directory entry
- Physical IRQs, 123
- Physical to virtual, 253
- Pin page table entry, 92
- Plan 9, 256
- Platform PCI device, 23
- Popek and Goldberg, 4

- Porting to Xen, 273
- Porting Xen, 255
- Power management, 259
- PowerTOP, 261
- Privilege rings, 28
- Privileged instructions, 4, 28
- Privileged operation, 30
- Protected memory, 78
- Protected mode, 17
- PTE, *see* Page table entry
- Python bindings, 200
- Python tools, 207
- pyxen, 201

- QEMU, 236
- QEMU Accelerator, 236

- RDP, 178, *see* Remote Display Protocol
- Real mode, 17, 27
- Realtime scheduling, 220
- Remote attestation, 183
- Remote Display Protocol, 182
- Requesting events, 124
- Resume, 94
- Ring buffers, 36, 65, 103, 164, 188
- ring.h, 65
- Round robin scheduler, 219

- S/360, 8
- sched_op hypercall, 132
- Scheduler
 - adding, 229
 - adding domains, 225
 - adding virtual CPUs, 225
 - API, 218
 - Atropos, 220
 - borrowed virtual time, 219
 - configuring, 228
 - credit scheduler, 222
 - defining, 224
 - hypercall interaction, 228
 - initialising CPUs, 225
 - interface, 218
 - realtime, 220
 - running, 225
 - Simple EDF, 221
 - SMP support, 225
 - stub domains, 263
 - work conserving, 221
- scheduler, 218
- Scheduler domains, 264
- Scheduler operations, 132
- Secure Virtual Machine, 237
- SEDF, *see* Simple EDF scheduler
- Segment registers, 76
- Segmentation offload, 171
- Segmented memory, 80
- Sending events, 128
- Sensitive instructions, 4
- set_gdt hypercall, 93
- Shadow page tables, 14, 82, 238
- Shared info page, 95
- Shared memory, 34, 59
- Shared memory buffers, 103
- shared_info_t, 51
- shmget(), 60
- Signals, 33
- Simple EDF scheduler, 219
- SimpleKernel, 157
- Simplest kernel, 38
- Single user virtualization, 258
- Sleep, 260
- Solarflare, 175
- SPARC, 6
- Split driver model, 35
- Split drivers, 100
- SPUs, *see* Synergistic Processing Units
- Start info page, 47
- start_info_t, 48
- Strongly ordered CPU, 107
- Stub domain scheduling, 263
- Stub domains, 245, 263
- Suspend, 94, 260
- SVM, *see* Secure Virtual Machine

- SVPC WG, *see* System Virtualization, Partitioning, and Clustering Working Group
- Synergistic Processing Units, 267
- System call, 11, 30, 120
- System Virtualization, Partitioning, and Clustering Working Group, 209
- System/360, 8
- Tagged translation lookaside buffer, 238
- task_slice, 224
- Threading models, 217
- Time keeping, 53
- Time-Stamp Counter, 53
- Timer device, 123
- Timer interrupt, 123
- TLB, *see* translation lookaside buffer
- TPM, *see* Trusted Platform Module
- Transferring memory, 63
- Translation lookaside buffer, 75, 90, 238
- Trap table, 94, 120
- trap_infol_t, 121
- Traps, 120
- Trusted Platform Module, 177, 183
- TSC, *see* Time-Stamp Counter
- TTLB, *see* tagged translation lookaside buffer
- Unmodified guests, 235
- Unprivileged domain, 22
- update_va_mapping hypercall, 90
- update_va_mapping_otherdomain hypercall, 90
- Updating the virtual framebuffer, 180
- USB, 186
- USB-over-IP, 187
- Userspace network drivers, 175
- Userspace tools, 200
- V2E, *see* Virtual to emulated
- V2P, *see* Virtual to physical
- VCPU, *see* Virtual CPU
- VGA BIOS, 244
- VIRQ, *see* Virtual IRQs
- Virtual 8086, 9
- Virtual appliances, 8
- Virtual block device, 161
- Virtual CPU, 51, 119, 220
- Virtual device bus, 109
- Virtual devices, 99, 100
- Virtual disk, 161
- Virtual framebuffer
 updating, 180
- Virtual IRQs, 123
- Virtual keyboard, 180
- Virtual Machine Control Block, 250
- Virtual Machine Control Structure, 250
- Virtual machine lifecycle, 37
- Virtual memory, 78
- Virtual network interface, 169
- Virtual pointing device, 180
- Virtual servers, 7
- Virtual sound devices, 264
- Virtual time, 53
- Virtual to emulated, 255
- Virtual to Physical, 253
- Virtual-Processor Identifier, 238
- Virtualization
 CPU, 4, 217, 220
 full, 235
 hardware-assisted, 235
 hybrid, 240
 I/O, 5, 36, 99, 224, 235, 239
 RAM, 4, 78
- Virtualization Technology for devices, 240
- Virtualization Technology for x86, 237
- Virtualization, hybrid, 14
- VM/370, 8
- vm_assist hypercall, 83
- VMCB, *see* Virtual Machine Control Block

- VMCS, *see* Virtual Machine Control Structure
- VMS, 28
- VMWare, 10
- VMX, 13
- VNC, 178
- VPID, 238
- VT, *see* Virtualization Technology for x86
- VT-d, *see* Virtualization Technology for devices
- VT-x, *see* Virtualization Technology for x86

- Wall clock time, 53
- Weakly ordered CPU, 108
- Windows guests, 235
- Work conserving schedulers, 221
- Writable page tables, 82
- WS-Management, 210

- X11, 178
- x86 memory model, 75
- x86 page directory, 78
- x86 privilege rings, 28
- x86 segmentation, 80
- Xen API
 - C bindings, 203
 - console, 203
 - host, 201
 - host CPU, 202
 - metrics, 202
 - physical block device, 202
 - physical network interface, 202
 - session, 201
 - virtual block device, 203
 - virtual machine, 202
 - virtual network interface, 203
 - virtual TPM, 203
- Xen API, 197
- Xen API classes, 201
- Xen daemon, 197, 206
- Xen device drivers, 100
- Xen driver model, 35
- Xen event model, 33
- Xen interface hierarchy, 200
- Xen maintainers, 253
- Xen management API, 197
- Xen master, 208
- Xen master (command line tool), 197
- Xen memory layout, 80
- Xen memory model, 78
- Xen networking, 169
- Xen security, 262
- xen_domctl_scheduler_op, 228
- xen_machphys_mapping_t, 87
- xen_machphys_mfn_list_t, 86
- xen_memory_exchange_t, 86
- xen_memory_reservation_t, 84
- xen_pci_op, 185
- xen_pci_sharedinfo, 184
- xen_translate_gpfn_list_t, 88
- XenBus, 109, 142, 163
- xenbus_device, 109
- xencons_interface, 113
- xenfb, *see* Virtual framebuffer
- xenfb_page, 178
- xenfb_update, 180
- xenkbd_key, 181
- xenkbd_page, 181
- xenkbd_position, 183
- XenSocket, 264
- XenStore, 36, 187
 - device structure, 145
 - mapping, 150
 - message structure, 146
 - message types, 147
 - reading a response, 153
 - userspace tools, 148
 - writing a message, 152
- XenStore device, 145
- XenStore hierarchy, 142
- XenStore interface, 141
- xenstore_domain_interface, 145

xm, *see* Xen master
XML-RPC, 198
XML-RPC Data Types, 198
xsd_sockmsg, 146, 151

z/VM, 8