

A

- Abelson, Harold 141
- abstraction layer 111
- Accenture 88
- access control 76–80
 - billing validation 76
 - identity verification 77
 - key 78
 - key pair 79–80
 - sign-in credentials 77
 - X.509 certificates 78–79
- ACID 32
- ACL 83
- Adobe
 - Flex, characteristics 206
 - use of Hadoop 146
- agile development, and Ruby on Rails 43
- agility, benefits of cloud computing 7
- AICPA 76
 - SAS 70 171
- Ajax
 - and load testing 164
 - history 191
- Akamai 39, 129, 175
- AlertSite 185
- Alexa 129
- Amazon
 - and cloudbursting 102
 - and Open Cloud Manifesto 203
 - APIs 28, 202–203
 - RESTful API 202
 - SOAP API 202
 - as content provider 129
 - auto scaling 28
 - AWS cloud, and Eli Lilly 70
 - AWS Start-Up Challenge 66
 - cloud services 38–39
 - cloud storage 125
 - CloudFront 39, 70
 - cost to cloud customers 96
 - data center
 - build-out 12
 - Dalles, OR 22
 - mega 20
 - security 75
 - EBS 28, 58, 207
 - e-commerce platform 2
 - Elastic Compute Cloud. *See* EC2
 - elastic load balancing 28
 - Large Instance option 57
 - limit on 500 XL instances 86
 - pricing 128
 - public cloud offering 2
 - S3 38, 80
 - shared-nothing partitioning 105
 - SimpleDB 39, 202
 - SQS 39, 202
 - startups as customers 195
 - use of Hadoop 146
 - virtual VPN 58
- Amazon VPC 89, 92
 - security solution 94
- AMD
 - and Open Cloud Manifesto 203
 - and parallelization 157
- AMI 27
 - and cloudbursting 119
 - connecting to 28
 - paid 28

- private 28
- public 28
- shared 28
- standard flow 28
- AOL 10
 - scaling issues 36
- Apache Struts, characteristics 206
- Apex 44
- API endpoint 82–83
- App Engine 42–43
 - and cloudbursting 121
 - and Force.com 44
 - automatic elasticity 42
 - datastore 43
 - development environment 42–43
 - evaluating 47
 - Java runtime environment 43
 - lock-in 47
 - MegaStore 42
 - pricing 43
 - programming languages 42
 - Python runtime environment 43
 - sandbox 42
- AppEngine 16
 - evolution into FaaS 210
 - operational parameters 182
- AppExchange 44
- Appirio 88
- Appistry 88–89
 - and Sprint 95
- Apple
 - and Browser Wars 165
 - data center, mega 20, 200
 - eMate 192
 - iPad 192
 - Newton 192
- application
 - ad hoc 136
 - corporate, expanding into the cloud 94
 - customizations 103
 - dependencies 105
 - deployment models 4
 - development
 - cost no longer a barrier 212
 - predictions about evolution of 205–212
 - adoption and growth of mashups 208–210
 - evolution of development tools to build mashups 210–211
 - higher-level services with unique APIs 208
 - PaaS and FaaS as predominant tools 210
 - rapid evolution for different storage mechanisms 207
 - role of application frameworks 205–206
 - second and third tiers running in the cloud 206–207
 - stronger options to protect sensitive data 207–208
 - success of non-Western developers 212
 - distributed 136
 - failure point, determining 156
 - frameworks
 - and growth of cloud computing 205–206
 - popular 206
 - high-scale 101
 - legacy, and the cloud 87
 - lifetime, and cloud deployment 60
 - mission-critical 63
 - nonstrategic 62
 - overprovisioned 117
 - owner, roles and responsibilities 83
 - patterns, best for the cloud 101–103
 - performance, testing 156
 - real-time 63
 - robustness 174
 - scale 175–176
 - scaling 24, 31
 - standardization on browser platform 190–191
 - stingy 174
 - application logic tier, running in the cloud 206
 - application server 133
 - and N-tier architecture 133
 - application service provider, hype 2
 - application stack, and functional tests 159

- Apps.gov 97
 - Aptana, and Ruby on Rails 44
 - architect 150
 - architecture
 - and automated deployment 154
 - cloudbursting 119
 - clustered 133
 - definition 103
 - deployment 138
 - distributed 32
 - multitier 133
 - Army, use of cloud 204
 - ARP cache poisoning 82
 - array, associative 32
 - artificial intelligence, hype 2
 - ASP 13
 - Atlassian Bamboo 159
 - AT&T Business Solutions, cloudbursting
 - scenarios 122
 - attribute-value 34
 - auditability 74
 - AUT 156
 - authentication 219
 - billing validation 76
 - multifactor 77, 208
 - single-factor 77
 - two-factor 75
 - authorization 219
 - autocomplete 165
 - automation 3, 5, 84, 154
 - auto-scaling 36
 - availability
 - 9s 172
 - and network connectivity 174
 - definition 171
 - in a private cloud 86
 - measuring 172
 - from multiple global locations 185
 - planning for failure 173
 - unplanned 173
 - AWS 34
 - access control 76–80
 - certificate, generating 79
 - identity verification 77
 - key 78
 - key pair 79
 - sign-in credentials 77
 - accessing bastion host 81
 - evaluating 46
 - host-based firewalls 82
 - Management Console 79, 184
 - Secret Access Key 83
 - AWS Start-Up Challenge 66
 - Azure 39–42
 - API 40–42, 121
 - AppFabric Access Control SLA 180
 - architecture 39
 - Compute SLA 179
 - Connectivity Downtime 179
 - Monthly Connectivity Uptime Percentage 179
 - evaluating 47
 - lock-in 47
 - operational dashboard 183
 - PaaS 40
 - pricing 40
 - SLAs 179–180
 - SQL Server SLA 180
 - scheduled maintenance 180
 - Storage SLA 180
 - error rate 180
- B**
-
- backup
 - and cloud computing 68
 - data storage service 62
 - remote 68
 - Baidu, use of Hadoop 146
 - bandwidth
 - elastic 153
 - utilization 55
 - BatchPutAttributes 34
 - Bechtel
 - cost and efficiency, vs. competitors 96
 - private cloud 96
 - best-effort basis 63
 - BI 67
 - as SaaS 67

Big Switch, The (Nicholas Carr) 98
 BigTable 33, 42, 104
 in Google App Engine 33
 scalability 33
 structure 33
 usage examples 33
 billing
 automated 26
 metered 3, 6, 84
 validation 76
 biometric 75
 Black Friday 86
 Blogger, I/O as bottleneck 106
 bootstrap 64
 bottleneck
 in traditional deployment 150
 master as 109
 broker, and SOA 137
 browser
 evolution into operating system 191
 platform, application standardization on
 190–191
 browser farm 160
 Browser Wars 165
 BrowserCam 166
 BrowserLab 166
 BrowserMob 148, 164
 B-tree 33
 bucket
 adding objects to 126
 deleting 126
 location constraint 126
 modifying 83
 new, creating 125
 request payment configuration 126
 Requester Pays 125
 S3 30, 125
 buffer-overflow attacks 208
 build agent 159
 Bureau of Alcohol, Tobacco, Firearms, and
 Explosives, use of cloud 204
 burst compute 102–103
 bursting 116
 business continuity, plus cloudbursting 123

business criticality, and choosing a public
 cloud provider 170
 Business Insights, predictions about
 mashups 210
 business intelligence 35
 business logic 132
 business service 172
 businesses, small and medium, using the
 cloud for 67–69

C

candidate key 32
 C#, and Selenium 160
 capacity
 on demand 115–124
 variability 60
 CAPEX. *See* capital expenses
 capital expenses
 definition 51
 in various application deployment
 models 4
 shift to operational expenses 4
 economic benefits of 6
 Carr, Nicholas 116
 The Big Switch 98
 CC 90
 CDMI 29
 CDN 70, 175
 Certificate Authority 79
 Chinese wall 208
 Chrome 191
 and Selenium 160
 CI 158–159
 adding build agents 159
 CIDR block 81
 ciphertext 220
 Citrix, and open source cloud 202
 CLC 90
 client
 requests 27
 stateless 132
 client-server 132
 client-server era 10
 Clinton, Bill 70

- Clinton, Hillary 70
- cloud
 - adoption
 - government leadership 204
 - slowed by security concerns 73–75
 - and confidential data 63
 - and consumer internet 189–194
 - and entrepreneurs 64
 - and legacy systems 63
 - and nonstrategic applications 62
 - and real-time/mission-critical applications 63
 - and scale volatility 60–62
 - and short-term applications 60
 - API 19, 27–28
 - as engine of growth for early adopters 199
 - bandwidth, cost 57
 - best-effort basis 63
 - billing model 53
 - classification 18, 37–46
 - compliance 176
 - database 19, 30–35
 - drawbacks 35
 - dedicated services 90
 - deployment 152–156
 - economics 24
 - economies of scale 85, 87
 - elasticity 19, 36–37, 176
 - file systems 128
 - geographic dependence 174
 - hybrid. *See* hybrid cloud
 - improving production operations with 152–155
 - in the enterprise 194–197
 - infrastructure
 - expertise required to run 87
 - scale 5
 - lack of up-front costs 64
 - load balancing, cost 58
 - managing predictable traffic peaks 61
 - metaphor, origin 8
 - model
 - cost structures 4
 - for IT 53
 - lack of startup fees 4
 - operations
 - measuring 181–186
 - visibility 181–185
 - through third-party providers 185–186
 - private 16, 20, 44–46, 72, 195–196
 - and control over security and compliance 195
 - and legacy apps 195
 - automation 84
 - availability 86
 - concerns 45
 - contrarian view 45
 - definition 84
 - deployment
 - concerns about 87–88
 - hosted 88–89
 - open source 88–89
 - options 88–92
 - proprietary 88–89
 - system integrator 88
 - elasticity 84
 - in practice 95–98
 - long-term viability 98
 - open source 90–92
 - pay-for-what-you-use 45
 - primary considerations 85
 - rate of innovation behind public clouds 88
 - rationale 84–92
 - security 85–86
 - small scale 87
 - user community 85, 87
 - virtualization 84
- provider
 - accountability 181
 - and SLAs 178–181
 - built-in verification 80
 - competitive factors 178
 - connections to internet backbone 174
 - cost sensitivity 201
 - costs, vs. cost of data centers 199
 - choosing 46–48
 - economies of scale 198
 - ease of switching between 177–178

- generating certificate 79
- incident response 80
- notification page 181
- operational dashboard 182–184
- performance, tracking 181–186
- public, economic advantages 87
- security 177
- staffing 201
- standardized platforms 201
- stored data verification 82
- transparency 181–185
- variable pricing 86
- provisioning, elastic, vs. data-center
 - provisioning 117
- public 16
 - access control 76–80
 - and legacy apps 195
 - and SAS 70 76
 - API endpoint 82
 - business attributes 84
 - co-mingling security 82
 - data center security 75–76
 - data storage security 83
 - information security 73–83
 - network and data security 80–83
 - network security 81–82
 - operating system security 81
 - principles 84
 - privatizing 89
 - provider, choosing 170
 - spending, in perspective 85
 - system control security 82
- ready for? 154
- reliability 139–146
- scale
 - definition 100
- scaling 19, 36–37
- security 198
 - vs. traditional security 73
- smallest possible system 90
- SOA as precursor 132–139
- storage 19, 124–129
 - API 125–128
 - attributes 125
 - cost 29, 128
 - definition 124
 - geographic control 129
 - latency 129
 - on demand 29
 - S3 125
 - saving persistent data 29–30
 - standard 29
 - unstructured 30
 - uses 29
 - virtualized 29
- support for innovation 199
- technological underpinnings 19–37
- usability, future improvement 199
- vendor taxonomy 13
- virtual private 72, 92–95
 - API 93–94
 - and disaster recovery 95
 - elastically scaling websites into 94
 - expanding corporate apps into 94
 - how it works 92–93
 - usage scenarios 94–95
- virtual VPN 58
- virtualized 19
 - when not to use 63–64
 - when to use 59–63
- cloudbursting 92, 102, 115–124
 - appliance 121
 - architecture 119–120
 - business case 117–118
 - data-access problem 122–124
 - definition 116
 - implementing 120–121
 - independent clusters 122
 - internal data center plus cloud 116–117
 - introduction 116
 - need for standards 121
 - on-demand data placement 123
 - plus business continuity/disaster
 - recovery 123
 - pre-positioned data 123
 - remote access to consolidated data 123
- cloud computing
 - agility benefits 7
 - and deadline-driven, large compute
 - problems 70

- and evolution of IT 8–13
- and large data sets 69
- and online web presence 70
- and product development 68
- and small and medium businesses 67–69
- and SOA 138
- and zero-capital start-ups 64–67
- as disruptive force in mobile world 193
- automation 5
- backup and file-storage systems 68
- benefits 6–8
- conversion of capital expenses to operational expenses 6
- corporate website 67
- definition 1
- demand variability 86
- economics 51–59
- efficiency benefits 7
- elasticity 5
- enterprise 69–71
- era 10
- first reference to 8
- five main principles 3–6
- future phases 194
 - ascendancy of private cloud 196
 - internal cloud migration 195
 - startups doing it hard and manually 195
 - transition to utility computing 196–197
- hype 2
- metered billing 6
- predictions about evolution of 198–205
 - 500,000 servers costing \$1 billion by 2020 200
 - Amazon’s APIs will lead the way 202–203
 - cheaper, more reliable, more secure, easier to use 198–199
 - engine of growth for early adopters 199
 - government leadership in cloud adoption 204
 - much lower costs than corporate data centers 199
 - open source dominance 201
 - ratio of administrators to servers 201
 - SaaS use of basic web standards 205
 - ultimate ISO cloud standard 203–204
 - private 73
 - reliable 132
 - security 8
 - service types 13–16
 - virtualization 4
 - vs. non-cloud-based, future prediction 197
 - X-as-a-Service 13
- Cloud Computing Interoperability Forum (CCIF), as part of Cloud Standards Coordination Working Group 203
- Cloud Standards Coordination Working Group 203
- CloudFront 39, 70, 129
- CloudTest 163
- CloudWatch 184
 - accessing 184
 - metrics 184
- clustering, and redundancy 140
- ColdFusion 207
- colocation 52
 - 95th percentile billing model 57
 - contracts 52
 - facility 52
- command-line logging 83
- commodity hardware 5
- communication
 - asynchronous 139
 - secret 218
- compatibility 177–178
- competition, and the cloud 198
- competitive advantages from efficiency
 - benefits 7
- compliance 176
- component, failure, handling 133
- compute unit 56
- computing
 - device-based 192–193
 - elastic 24
 - high-performance 70
 - paradigm shifts 10
 - pooled resources 3–4
 - virtualized 3

Compuware, projected computing mix 194
 confidentiality 219, 222
 persistent 223
 connection, limiting 82
 consumer internet, and the cloud 189–194
 consumers and cloud apps 194
 container 29
 content delivery, and combating latency 129
 continuous deployment 168
 Continuous Integration in .NET 159
 corporate website, cloud computing and 67
 cost
 allocation, colocation as proxy for 55
 up front 64
 CPC
 cloud visibility dashboard 185
 measuring performance and
 availability 185
 Cray-1 11
 CreateAutoScalingGroup 37
 CreateCustomerGateway 94
 CreateDomain 34
 CreateLaunchConfiguration 37
 CreateLoadBalancer 36
 CreateOrUpdateScalingTrigger 37
 CreateVpc 94
 CreateVpnConnection 94
 CreateVpnGateway 94
 credentials 27, 77
 CRM
 and transference 101
 Salesforce.com 44
 crowd sourcing 158, 167
 crypto-algorithm 219
 cryptography 218–222
 algorithms 219
 asymmetric 221
 authentication 219
 authorization 219
 confidentiality 219
 definition 218
 digital signature 221
 encryption 221
 identify 219

integrity 219
 key 219
 nonrepudiation 219
 private key 221
 public key 78, 221
 shared-key 220
 symmetric 220
 CSA 203
 as part of Cloud Standards Coordination
 Working Group 203
 Customer Gateway 94
 Cyber Monday, and load testing 162

D

data
 anomalies 107
 archiving in the cloud 29
 backing up to cloud storage 29
 backup storage 62
 centralizing 80
 combating latency 129
 confidential 63
 deleting 64
 consolidated, remote access to 123
 denormalized 108
 encrypting before uploading to
 cloud 83
 hoarders 124
 in transit 83
 integrity 113
 joining 113
 keeping secret 73
 leakage 80
 nonrelational 32
 normalized 107
 parallelized 108
 partitioning 31
 persistent, saving in cloud storage 29–30
 placement, on-demand 123
 pre-positioned 123
 rebalancing 112
 redundant 32
 relations 30
 replicating 33

- replication 31, 109
 - and redundancy 140
- segregating 80
- storage, security 83
- stored, verification 82
- streaming 29, 129
- structured, storing in cloud
 - database 30–35
- table 30
- transactional 30
- uniform allocation 110
- data center 19
 - Amazon 22
 - cooling 20
 - cooling towers 22
 - cost 21
 - vs. costs of cloud providers 199
 - Dalles, OR 21
 - economies of scale 12, 19–23
 - efficiency 24
 - evolution 11–12
 - Google 21
 - hardware 21
 - mega 20, 200
 - and virtualization 26
 - costs 200
 - labor 201
 - Microsoft 23
 - modularity 23
 - network connectivity and bandwidth 20
 - platform complexity 201
 - power 20, 21
 - provisioning, vs. elastic cloud
 - provisioning 117
 - PUE 24
 - scaling 21–22
 - security 75–76
 - physical 75–76
 - structure 20
 - Quincy, WA 22
 - roofless 23
 - security 20
 - by obscurity 75
 - servers 21
 - staffing 201
- Data Center-as-a-Service, private clouds 44–46
- database
 - architecture, traditional, vs. sharding 107–109
 - centralized 104
 - consistency
 - eventual 32
 - guarantees 32
 - decomposition 105
 - item-oriented 32
 - join 33
 - key-value 32
 - master-slave 109
 - scaling 108
 - shared-nothing partitioning 104
 - updates, frequent, and scaling 104
- DDoS attack 82, 208
- decision support 35
- decryption 219
- DeleteAttributes 35
- Dell
 - and mega data centers 199
 - and open source cloud 202
- demand, event-driven spike 61
- denormalizing 108
- Department of Homeland Security, as part of FedRAMP 204
- Department of the Treasury, use of cloud 204
- deployment
 - architecture, traditional 149–150
 - automating 154–155
 - cloud 152–156
 - continuous 168
 - cost 152
 - typical 149–152
- deployment model, cost comparison 53
- depreciation 60
- deprovisioning, automated 116
- design, definition 103
- development
 - accelerating 155–156
 - single-threaded 157
- development environment 149

- device-based computing 192–193
 - Digg
 - scale of 32
 - shared-nothing partitioning 105
 - digital signature 79, 221
 - and XML Signature 222
 - DISA
 - private cloud 97
 - Rapid Access Computing Environment (RACE) 97
 - disaster recovery 20, 68
 - and virtual private clouds 95
 - plus cloudbursting 123
 - disk storage, elastic 153
 - disk, virtual, synching to cloud 29
 - distributed computing, MapReduce 141
 - distributed grep 145
 - Distributed Management Task Force (DMTF), as part of Cloud Standards Coordination Working Group 203
 - distributed processing, and MapReduce 143
 - DoD
 - as part of FedRAMP 204
 - DISA 97
 - domain, in SimpleDB 34
 - DomainMetadata 34
 - downtime per year 173, 181
 - drug development, and cloud computing 69
 - dual tree 115
 - dynamic scaling 5
- E**
-
- eBay
 - and internet scale 103
 - scale of 32
 - shared-nothing partitioning 105
 - EBS 28, 58, 128, 207
 - ebXML 137
 - EC2 15, 37–39
 - and browser farms 160
 - and EBS 128
 - and Eventseer 119
 - and Xen 201
 - API 28, 121
 - documentation 28
 - API calls, using to creating VPC 94
 - automatic scaling 38
 - co-mingling security 82
 - elasticity 176
 - evaluating 46
 - failover 38
 - lock-in 47
 - operational dashboard 183
 - pricing 39
 - Selenium Grid 160
 - SLA, Region Unavailability 178–179
 - use of Xen paravirtualization 38
 - using in Amazon VPC 92
 - Eclipse, and sharding 113
 - economies of scale 19–23, 26
 - in a private cloud 85, 87
 - efficiency benefits 7
 - elastic people 167
 - elastic storage 102–103
 - elasticity 3, 5, 36–37, 84, 176
 - and celebrity deaths 36
 - electricity, generation, comparison to cloud computing 98
 - Electronic Business using eXtensible Markup Language (ebXML) 137
 - Eli Lilly
 - and Amazon AWD cloud 70
 - cloud computing in the enterprise 69–70
 - EMC 89
 - encryption 219
 - EngineYard, and Ruby on Rails 44
 - Enomaly 88
 - enterprise
 - cloud in 194–197
 - using the cloud for 69–71
 - enterprise IT, survey 74
 - Enterprise Mashup Markup Language (EMML) 210
 - entrepreneur 64
 - bootstrapping 64
 - challenges 64
 - need for venture capital 65

- Eucalyptus 46, 88–89
 - clusters 91
 - elasticity 176
 - interfaces 91
 - private cloud implementation 90–92
 - Xen 90
 - Eventseer.net, cloudbursting 102
 - architecture 119–120
 - business case 118
 - Expedia, data center, build-out 12
- F**
-
- FaaS 16
 - as predominant tool for app construction 210
 - Facebook 193–194
 - and internet scale 101, 103
 - I/O as bottleneck 106
 - scale of 32
 - shared-nothing partitioning 105
 - sharding 106
 - storage 124
 - capacity 106
 - use of Hadoop 146
 - user-created apps 211
 - fat client 191
 - Fawcett, Farrah 36
 - Federal Risk and Authorization Management Program (FedRAMP) 204
 - federated 105
 - file storage, and cloud computing 68
 - Firefox 191
 - and Selenium 160
 - firewall 53, 74
 - public cloud 81
 - five 9s 172
 - Flash 207
 - history 190
 - Flickr 194
 - and internet scale 101, 103
 - database partitioning scheme 113–114
 - I/O as bottleneck 106
 - profile 113
 - reliability strategy 114–115
 - shared-nothing partitioning 105
 - sharding 113–115
 - storage 124
 - FlightCaster 66
 - FLOPS 11
 - Force.com 44
 - and App Engine 44
 - AppExchange 44
 - as example of FaaS 16
 - evaluating 48
 - evolution into FaaS 210
 - lock-in 48
 - pricing 44
 - foreign key 32
 - Fort Knox, Salesforce.com data center 75
 - Framework as a Service. *See* FaaS
 - fraud detection, and Sprint 95
 - Freedom of Information Act 70
 - Friendster, shared-nothing partitioning 105
 - functional testing 156, 159–161
 - automating 160–161
 - parallelization 160–161
 - reducing time to complete 160–161
- G**
-
- Gartner Group
 - estimate of SaaS market size 205
 - private clouds 73
 - GetAttributes 35
 - GetDatabaseFor 107, 111
 - GFS. *See* Google File System
 - Gibson, William 189
 - Gmail 194
 - and visual testing 166
 - history 191
 - GoGrid, and Simple Cloud API 202
 - Gomez
 - measuring performance and availability 185
 - study of cloud performance 175
 - GoodData 66
 - Google
 - Analytic, availability issues 174
 - and Browser Wars 165

- and internet scale 103
 - and Open Cloud Manifesto 203
 - Apps, vs. Microsoft Office 39
 - autocomplete 164
 - BigTable 104
 - commodity hardware 145
 - data center
 - build-out 12
 - Dalles, OR 21
 - mega 20, 200
 - security 75
 - Docs, vs. Microsoft Office 39
 - Maps
 - history 191
 - use in mashups 209
 - public cloud investment 84
 - reliability 145
 - search requests, cloud computing in 8
 - server maintenance 96
 - sharding 104
 - shards, rebalancing 112
 - Googleplex 22
 - government leadership in cloud
 - adoption 204
 - Gramm-Leach-Bliley Act (GLBA), and
 - SAS 70 76
 - green field 150
 - green-screen terminal 10
 - growth and the cloud 199
 - growth, exponential 102
 - GSA 122
 - as part of FedRAMP 204
 - prebidding cloud services 204
- H**
-
- Hadoop 142, 146
 - Hadoop in Action* 146
 - hand-geometry scanner 75
 - hardware
 - as mashup of standard cells 210
 - commodity 5
 - deep security 208
 - failure 154
 - inspection of infrastructure 177
 - Hardware as a Service. *See* IaaS
 - Harkins, Peter 70
 - Harvard, and Facebook 102, 106
 - hashing algorithm, bad 111
 - hash table 32
 - Henderson, Cal 113
 - Heroku, and Ruby on Rails 44
 - Hibernate, and sharding 113
 - High Scalability 113
 - highscalability.com 113
 - Hilton, Perez 36
 - HIPAA 64, 74
 - and compliance 177
 - and SAS 70 76
 - and the public cloud 86
 - horizontal scalability 173
 - host, bastion 81
 - hot cloud standby 121
 - HP
 - and visibility 185
 - Mercury 7
 - HP-UX 63
 - HTML 5, and SaaS 205
 - HTTP
 - GET 128
 - PUT 128
 - Hudson 159
 - Hulu, use of Hadoop 146
 - hybrid cloud 16
 - Hyperic, CloudStatus, measuring
 - performance and availability 185
 - hypermedia 27
 - hypervisor 25, 38
 - and promiscuous mode 82
 - deep security 208
 - separation from guest 81
- I**
-
- IaaS 15, 39, 47
 - and cloudbursting 119, 121
 - AWS 46
 - Azure 39–42
 - EC2 37–39
 - in early stage of cloud computing 195

- LAMP stack 38
- IBM 89
 - and Open Cloud Manifesto 203
 - and Simple Cloud API 202
 - and visibility 185
 - cloudbursting appliance 121
 - use of Hadoop 146
- IDC
 - and SaaS 205
 - IT cloud concerns 73
- identity 219
 - verification 77
- IEEE Spectrum 23
- IETF, and XML Signature 222
- image processing, real time 63
- index 107
 - inverted 146
- indirection 111
- information
 - security
 - and social engineering 83
 - in the public cloud 73–83
 - sensitive 64
- Infosys 88
- infrastructure
 - physical vs. virtual 56
 - security, traditional, vs. cloud security 73
- Infrastructure as a Service. *See* IaaS
- initialization fee 55
- innovation, cloud support for 199
- instance 28
 - connecting 28
- integrity 219
- Intel
 - and mega data centers 199
 - and parallelization 157
- interface
 - browser 27
 - constrained 135
 - extensible 135
 - versionable 135
 - well-defined 135
- intermodule dependency 134
- Internet Explorer 191
 - and Selenium 160
- internet scale 101–115, 152
 - application issues that prevent 104
- interoperability, levels of 177–178
- interoperable services 135
- interprocess communication 138–139
- intrusion detection 75
- Intuit
 - data center
 - mega 20
 - security 75
 - load testing 163
- invalidatable 112
- I/O
 - latency 182
 - throughput 182
- I/O system, and scaling 104
- iPhone 193
 - user-created apps 211
- IPtables 82
- ISO cloud standard 203–204
- ISP 10
- ISV, product development 68
- IT
 - deployment
 - cloud 59
 - cost 55–60
 - colocation 55, 59
 - handling peak capacity 61
 - internal IT 59
 - cost 54–55
 - managed service 59
 - cost 55–56
 - evolution of 8–13
 - infrastructure
 - cloud 53
 - colocation 52
 - internal IT 51
 - managed service 52
 - operational costs 54
 - peak utilization 61
 - rental model 55–56
 - shift from self-hosted to outsourcing 4
 - strategy, and public clouds 73

- transformation 189–197
- models
 - comparing 51–53
 - cost comparison 53–59
- IT Roadmap Conference and Expo 96

J

- Jackson, Michael 36, 61
- Java
 - and App Engine 42
 - and Selenium 160
 - JDO/JPA, and App Engine 43
 - runtime, and App Engine 43
 - Servlet standard, and App Engine 43
- JavaScript
 - and autocomplete 164
 - history 190
- JMeter 162
- join operation, avoiding 32
- JSP, and App Engine 43
- JUnit 156
 - and unit testing 157
- just-in-time manufacturing 53
- JVM, and App Engine 43

K

- Kaye, Doug 135
- Kelsey Group 193
- key
 - access 78
 - candidate 32
 - definition 219
 - foreign 32
 - in a sharded database 113
 - keeping secret 219
 - length 220
 - management 208
 - pair 79–80
 - primary 32
 - private 78, 221
 - public 27, 78, 221
 - S3 30
 - secret 78

- SSH 81
 - uniqueness 219
- key-based authentication 83
- key-value pair 32
- Keynote, measuring performance and
 - availability 185
- Kunda, Vivek 96

L

- LA Times*, scaling issues 36
- Lam, Chuck 146
- LAMP stack 38, 54
- latency 129
- legacy application
 - and private clouds 45
 - and the cloud 87
- legacy system 63
- Lightbody, Patrick 148
- Limelight 175
- LinkedIn, I/O as bottleneck 106
- Linux 201
 - and implementing a private cloud 90
 - Eli Lilly use of 70
 - guest OS 83
- Lisp, map and reduce functions 141
- ListDomains 34
- Litmus 166
- LiveJournal, and sharding 113
- load testing 7, 156, 162–165
 - and Pylot 65
 - browser-based, cloud-based 164–165
 - cloud-based, cost 163
 - script, typical 164
 - traditional 162
- load-balancer 53
- load-balancing 103
 - and cloudbursting 117
 - and redundancy 140
- Flickr 115
- LoadRunner 7, 162
- LoadStorm 163
- lock-in, avoiding 46
- loose coupling 133–135
 - and sharding 111

- and SOA 136–137
- critical application characteristics 135
- techniques to create and maintain 135
- Los Angeles, government use of cloud 204

M

- machine, parallel 158–159
- magic bullet 59
- mainframe 10
- man trap 75
- managed service 52
 - contract 52
 - value proposition 56
- management plane 81
- manual testing 156, 167–168
 - crowd sourcing 167
- map function 145
- MapReduce 141–146
 - achieving reliability 141
 - distributed grep 145
 - how it works 143–145
 - inverted index 146
 - map 142
 - master 141
 - parallel programming 141
 - reduce 142–143
 - reverse web-link graph 145
 - roots in Lisp 141
 - term-vector per host 146
 - usage examples 145–146
 - worker 141
- MashQL 210
- mashup 208–211
 - component exchange 210
 - components listed at
 - ProgrammableWeb.com 211
 - development tools 210–212
 - lack of formal standard 210
 - support ecosystem 210
 - tools to build 210
 - use of Google Maps 209
- Mashup API 211
- master-slave 109

- master-slave replication 151
- McKinsey, private clouds 73
- MD5 hash 80, 82
- mean-time-to-failure 141
- Mechanical Turk 167
 - Turkers 167
- MegaStore 42
- Mercury 162
- message
 - asynchronous 139
 - coarse-grained 135
 - confidentiality 222
 - integrity 222
 - persistent 139
 - persistent confidentiality 223
 - platform-independent 135
 - self-contained 135
 - self-describing 135
 - vendor-independent 135
- message queue 138
- messaging, stateless 135
- metadata, in S3 30
- metered billing 3, 6, 84
- Microsoft
 - and Browser Wars 165
 - and Open Cloud Manifesto 203
 - and Simple Cloud API 202
 - announcement of renting cloud space 9
 - Azure 16
 - data center
 - build-out 12
 - PUE 23
 - Quincy, WA 22
 - security 75
 - mega data centers 200
 - .NET, CLR 40
 - Office, vs. Google Docs 39
 - Tablet PC 192
 - Visual Studio, Hypervisor 39
- middle tier 132
- middleware, Appistry 95
- migrating users to/from shards 111
- mobile revolution 193–194

modulo 110
 Mosaic Mozilla 190
 Mozilla 190
 and Browser Wars 165
 MRI 63
 multicore 157
 multicore server 5
 multifactor authentication 208
 multitenancy 87
 MySpace
 I/O as bottleneck 106
 load testing 163
 MySQL
 and sharding 113
 dual tree 115

N

NASA
 and Rackspace 202
 and redundancy 140
 use of cloud 204
 National Archives 70
 national boundaries, and security 75
 national security, and public cloud
 computing 86
 NC 90
 Nebula Cloud Platform 202
 NetApp 115
 netbook 192
 Netscape Navigator 191
 NetSeer, use of Hadoop 146
 network computer 192
 network security 81–82
 network systems management 185
New York Times, use of Hadoop 146
 Nirvanix, and Simple Cloud API 202
 NIST, definition of cloud computing 116
 nonrepudiation 219, 222
 normalizing 107
 NoSQL 32–35
 distributed architecture 32
 scaling 32
 weak consistency guarantees 32
 NUnit 156

O

OASIS, as part of Cloud Standards
 Coordination Working Group 203
 object, in S3 30, 125
 adding to bucket 126
 copying 126
 deleting 126
 fetching 126
 retrieving information about 126
 Object Management Group (OMG), as part
 of Cloud Standards Coordination
 Working Group 203
 OCC 203
 as part of Cloud Standards Coordination
 Working Group 203
 on-demand, and choosing a public cloud
 provider 170
 on-premises
 and security 45, 82
 transference 101
 Open Cloud Manifesto 203
 Open Grid Forum (OGF), as part of Cloud
 Standards Coordination Working
 Group 203
 Open Mashup Alliance (OMA) 210
 open source
 and implementing a private cloud 90
 client interface compatible with
 Amazon 91
 dominance in future clouds 201
 software
 and App Engine 47
 definition 44
 OpenNebula 88–89
 OpenQA 148
 OpenStack 202
 Opera, and Selenium 160
 Opera Software ASA, browser 191
 operating system, deep security 208
 operational dashboard 182–184
 operational expenses
 definition 51
 in various application deployment
 models 4

- shift from capital expenses 4
 - economic benefits of 6
- OPEX. *See* operational expenses
- OpSource 88
- Oracle 89
- orchestration, and SOA 136
- ORM 35
- OS
 - guest 25, 38, 74, 83
 - Linux 83
 - security 81
 - virtualization 24
- out of band 77

P

- PaaS 16
 - and cloudbursting 121
 - as predominant tool for app
 - construction 210
 - App Engine 42–43, 47
 - Azure 39, 47
 - Force.com 44, 48
 - Ruby on Rails 43, 48
- packet filter 74
- parallel machines 158–159
- parallel processing 133
- parallel programming, MapReduce 141
- parallelization 157–168
- ParaScale 89
- paravirtualization 38, 81
- password 77
- password-based access, disabling 83
- pay-as-you-go 6, 13, 39
- pay-only-for-what-you-use 51
- PCI DSS, and compliance 176
- peak capacity 60–61
- peak load 176
- peak utilization 61
- penetration testing 156, 208
- performance
 - and network connectivity 174
 - definition 171
 - geographic dependence 174
 - measuring 173
 - from multiple global locations 185
 - requirements 69
 - testing 156
- performance engineer 150
- performance testing 156
- Perl
 - and Ruby 44
 - and Selenium 160
- persistent confidentiality 223
- Pew Internet & American Life Project 190
- PHP
 - and Selenium 160
 - and Simple Cloud API 202
 - characteristics 206
- PIN 77
- Platform as a Service. *See* PaaS
- PlayStation, and parallelization 157
- PlentyOfFish 124
- Poneman 8
- pooled resources 3–4, 84
- portal 209
- power outage 20
- predictions
 - about evolution of application
 - development 205–212
 - adoption and growth of mashups
 - 208–210
 - development cost no longer a
 - barrier 212
 - evolution of development tools to build
 - mashups 210–211
 - higher-level services with unique
 - APIs 208
 - PaaS and FaaS as predominant
 - tools 210
 - rapid evolution for different storage
 - mechanisms 207
 - role of application frameworks 205–206
 - second and third tiers running in the
 - cloud 206–207
 - stronger options to protect sensitive
 - data 207–208
 - success of non-Western developers 212
 - about evolution of cloud
 - computing 198–205

500,000 servers costing \$1 billion by 2020 200

Amazon's APIs will lead the way 202–203

cheaper, more reliable, more secure, easier to use 198–199

engine of growth for early adopters 199

government leadership in cloud adoption 204

much lower costs than corporate data centers 199

open source dominance 201

ratio of administrators to servers 201

SaaS use of basic web standards 205

ultimate ISO cloud standard 203–204

computing 189

primary key 32

privacy rules 125

private key 78, 221

privilege escalation 81, 83

processing time 182

product development, and cloud computing 68

product, launching on the cloud 69

production environment 149

typical 150

production operations, improving 152–155

ProgrammableWeb.com, mashup components 211

promiscuous mode 82

prototyping, and internet scale 102

provisioning

- automated 26, 116
- automatic 5
- resources 27
- standards 178

public key 27, 78, 221

public-key cryptography 221

PUE 23

PutAttributes 34

Pylot 7, 65

Python 7

- and App Engine 42–43

- and Selenium 160
- characteristics 206

Q

QA, and cloud computing 68

quality of service (QoS) 6, 29

Quantcast 61–62

query load 106

R

Rackable/SGI 21

Rackspace 149

- and NASA 202
- and Open Cloud Manifesto 203
- and Simple Cloud API 202
- and Xen 201

Cloud SLA 180–181

Cloud status page 182

resource constraints 86

Sites SLA 180

- use of Hadoop 146

RAID, and redundancy 140

Ramleth, Geir 96

Rapid Access Computing Environment (RACE) 97

R&D, and cloud computing 69

RDBMS 30

- definition 30
- in the cloud 30
- scaling 31

Reagan, Ronald 186

real-time application 63

RedHat 201

reduce function 145

redundancy 54, 109, 140

reengineering 103

referential integrity 31, 113

regulatory requirements, and public cloud computing 86

relational model 30

reliability 133, 139–146

remote-hands capability 52

Remy, Dave, *Securing Web Services with WS-Security* 73

- replication
 - dual master 109
 - lag 114
 - resources
 - and REST 27
 - pooled 84, 153
 - provisioning 27
 - REST 27
 - and mashups 210
 - RESTful
 - cloud storage 124
 - definition 27
 - reverse web-link graph 145
 - risk mitigation 102
 - robustness 174
 - root access 83
 - Rosenberg, Jothy, *Securing Web Services with WS-Security* 73
 - RSA, SecurID 77
 - Ruby, and Selenium 160
 - Ruby on Rails 43
 - characteristics 206
 - evaluating 48
 - lock-in 48
- S**
-
- S3 28, 38
 - and EBS 128
 - and Eventseer 119
 - and storage 153
 - API 29
 - bucket 30, 125
 - cloud storage 125
 - example API 125–128
 - key 30
 - object 30, 125
 - metadata 30
 - pricing 39
 - SLA 178–179
 - Error Rate 179
 - InternalServerError 179
 - Monthly Uptime Percentage 179
 - ServiceUnavailable 179
 - usage 30
 - SaaS 16, 27, 63
 - and HTML 5 205
 - and security 75
 - as requirement for cloud computing 12
 - Azure 39
 - BI as 67
 - evolution 13
 - GoodData as example of 66
 - in early stage of cloud computing 195
 - product development 68
 - Salesforce.com 44
 - use of web standards to grow and stay
 - current 205
 - Safari 191
 - and Selenium 160
 - Salesforce.com 44, 68
 - as example of SaaS 16
 - security 75
 - software 96
 - SAN 153
 - SAS 70 76, 171
 - and the public cloud 86
 - SAS 70 Type II 177
 - Savvis 88–89
 - SC38 203
 - scalability 36
 - high, in nonrelational database 33
 - scale 175–176
 - elastically adjusting 3
 - variability 60
 - scaling
 - application 24
 - horizontally 32, 106, 108
 - importance of 36
 - problems 102
 - vertically 106
 - Schneier, Bruce 219
 - Secret Access Key 83
 - secret communication 218
 - Secure Data Connector 92
 - SecurID 77
 - Securing Web Services with WS-Security* (Jothy Rosenberg and Dave Remy) 73
 - security 8, 177

- by obscurity 75
- co-mingling 82
- concerns, slowing cloud adoption 73–75
- data 80–83
- data storage 83
- in a private cloud 85–86
- logical 20
- network 80–83
- operating system 81
- perimeter 75
- physical 20, 75–76
- requirements 69
- secret communication 218
- standards 177
- stronger options 207–208
- system control 82
- testing 156
- through obscurity 20
- XML Encryption 222
- XML Signature 222
- Selenium 148, 160
 - Grid, and browser farms 160
- self-hosted model 4
 - elasticity 5
- server
 - capacity, measuring in compute units 56
 - commodity 53
 - multicore 5
 - provisioning, and virtualization 26
 - responses 27
 - sizes 20
 - utilization 25
 - world consumption and emissions 22
- service
 - consumer 137
 - developing 137
 - interactions 136
 - interoperable 135
 - provider 137
- service-level agreement. *See* SLA
- Service-Oriented Architecture. *See* SOA
- Servlet Specification, web application 191
- shard
 - dual master 109
 - master-slave 109
- sharding 103–115, 175, 207
 - advantages 105–106
 - and Facebook 106
 - and Flickr 113–115
 - challenges and problems 112–113
 - changes to application 107
 - common partitioning schemes 109–111
 - data
 - denormalized 107–108
 - highly available 109
 - not replicated 109
 - parallelized 108
 - small 108
 - definition 105
 - directory-based 111–113
 - faster queries 105
 - GetDatabaseFor 107
 - hash-based 110
 - high availability 105
 - introduction 104–106
 - joining data 113
 - key-based 110
 - lack of support for 113
 - range-based 110
 - rebalancing data 112
 - referential integrity 113
 - scheme 104
 - simplest model of 105
 - strategy 106
 - unbalanced, avoiding 111
 - vertical 109
 - vs. traditional database architecture 107–109
 - write bandwidth 106
- shared secret 77
- shared-key cryptography 220
- shared-nothing 104–105
- Simple Cloud API 202
- SimpleDB 33, 39
 - API 34
 - domain 34
- single point of failure 132
 - in a shard 109

- single-threaded 157
- Skype, shared-nothing partitioning 105
- SLA 26, 173, 178–181
 - violations 178
- Smalltalk, and Ruby 44
- smartphone 193–194
 - adoption 193, 194
 - vs. PC 193
- SOA 135–138
 - and cloud computing 138
 - and loose coupling 136–137
 - and web services 137–138
 - as cloud precursor 132–139
 - as requirement for cloud computing 12
 - hype 2
 - orchestration 136
 - service consumer 137
 - service provider 137
 - services 136
- SOAP, and encryption 222
- SOASTA 163
- social engineering 83
- software
 - licensing, and virtualization 26
 - pricing, and virtualization 26
- SOX 74
 - and SAS 70 76
- space shuttle, and component
 - redundancy 140
- spoofed traffic 82
- Sprint, private cloud 95–96
- SQL 30
 - movement away from 32
- SQS 39, 138–139
 - and cloudbursting 121
 - and Eventseer 120
- squid cache 115
- SSDS 34
- SSL
 - and asymmetric encryption 222
 - and shared-key encryption 221
- staging, cost 155
- staging environment 149
 - defining 150–151
- standards
 - future development 202–203
 - ISO cloud standard 203–204
- startup
 - 2000 vs. 2010 64–65
 - lower barriers to entry 65
- stateless client 132
- stickiness 136
- storage
 - cloud 124–129
 - encrypted 74
 - expanding 124
 - exponentially expanding 124
 - management 102
 - mechanisms, rapid evolution of 207
 - on demand 29
 - structured 32
 - virtualized 29
- Storage Networking Industry Association (SNIA) 29
 - as part of Cloud Standards Coordination Working Group 203
- storage tier, running in the cloud 206
- string, human-readable 135
- strong coupling 134
- Structure and Interpretation of Computer Programs 141
- structured query language. *See* SQL
- sudo 83
- Sun
 - and Open Cloud Manifesto 203
 - network computer 192
 - Project Kenai Cloud 27
- SunGard 88, 89
- Super Bowl 36
 - and load testing 162
- Sussman, Gerald Jay 141
- Sussman, Julie 141
- sweat equity 65
- SYN cookies 82
- synthetic transaction 185
- system
 - distributed 132–133
 - legacy 63
 - loosely coupled 132

T

tablet 192

Target.com, daily U.S. visitors 61

taxonomy, cloud vendors 13

TechCrunch 36

term vector 146

terminal, green-screen 10

Terremark, and USA.gov 122

test automation 156

test lab, virtual, in the cloud 68

test transaction 185

testing

- accelerating 155–156
- acceptance 156
- cost 155
- cloud-based
 - acceleration 156
 - cost savings 155–156
- single-threaded 157
- vulnerability testing 208

testing environment 149

- cloning 156
- defining 150–151

thin client 191

three 9s 172

three-tier 53

throttle 103

tight coupling 133–134

- critical application characteristics 135

time shared 10

time-sharing 24

TM Forum, as part of Cloud Standards
Coordination Working Group 203

TMZ.com 62

- scaling issues 36
- traffic spike 61

traffic

- event-based, and cloudbursting 117
- seasonal, and cloudbursting 117
- spoofed 82

transaction

- and redundancy 140
- synthetic 185
- test 185

transference 101, 103

transit

- inbound 57
- outbound 57

triple modular redundancy with
spares 140

TurboTax, load testing 163

Twitter

- and internet scale 102
- scaling issues 36

two-man-rule 82

U

Ubuntu, and implementing a private
cloud 90

UDDI 137

UEC 90

ultraportable 192

Unisys 89

United States

- Apps.gov 97
- Army, use of cloud 204
- Bureau of Alcohol, Tobacco, Firearms,
and Explosives, use of cloud 204
- CIO, proponent of cloud 204
- Congress, approval of electronic
signatures 222
- Department of Homeland Security, as
part of FedRAMP 204
- Department of the Treasury, use of
cloud 204
- FedRAMP 204
- government leadership in cloud
adoption 204
- private clouds 96–98

unit testing 156–159

- parallelization 158

UNIX, private cloud implementation 88

untethered 193

Uptime Institute 24

usability testing 156

- crowd sourcing 168

usage, event-driven spike 61

USA.gov 122

USA PATRIOT Act 75
 user community, in a private cloud 85, 87
 user interface, verifying 156
 UserTesting 168
 uTest 168
 utility computing 196–197
 definition 188

V

vendor
 availability 171–175
 choosing 170–178
 business considerations 170–171
 questions to ask 170
 technical operational considerations 171–178
 contractual viability 171
 financial viability 170
 operational viability 170
 performance 173–175
 taxonomy 14
 venture capital, necessity of 65–66
 venture capitalist 64
 Victoria's Secret 36
 video surveillance 75
 Virgin Atlantic, cloud computing in the
 enterprise 70
 virtual machine 25
 architecture 25
 automatic creation/deletion 3, 5
 image 24
 instance 24
 virtual machine monitor 24
 virtual memory 24
 Virtual Private Cloud 207
 virtual VPN, and IPsec tunnels 58
 VirtualBox 90
 virtualization 4, 24–26, 84
 and deployment 152
 as requirement for cloud computing 12
 cloud 24, 26
 corporate 26
 disk 82
 effect on corporate data centers 26

 in the enterprise 85
 layer 25
 platform 24
 server 24–25
 time-sharing 24
 virtualized computing 3, 10
 virtualized disk 82
 visibility 181–185
 through third-party providers 185–186
 visual testing 156, 165–166
 VLAN 52, 74
 VMM 25, 38, 81
 VMS 63
 VMware 12, 88, 90, 178
 and Open Cloud Manifesto 203
 virtual machine architecture 25
 voltage/frequency scaling 21
 VPC 28, 46
 VPN
 and virtual private clouds 93
 and Virtual Private Network 46
 Vtravelled.com 70
 vulnerability testing 208

W

W3C
 and XML Encryption 223
 and XML Signature 222
Washington Post, cloud computing in the
 enterprise 70
 Watson, T. J. 189
 Web 2.0, and mashups 210
 web application 191
 architecture 132
 deployment 149
 second and third tiers running in the
 cloud 206
 web service
 and SOA 137–138
 XML Encryption 222
 XML Signature 222
 website, scaling elastically in the cloud 94
 Weinman, Joe 122
 Wikipedia, shared-nothing partitioning 105

Windows Server 2008 39
working set, too large, and scaling 104
writing, as bottleneck 106
WSDL, and loose coupling 135

X

X.509 27
 and Azure 40
X.509 certificate 78–79
 and firewall 82
 and system control security 82
XaaS 13, 15
X-as-a-Service. *See* XaaS
XBox, and parallelization 157
Xen 38, 90, 201
XL instance 86
XML, and SOA 136
XML Encryption 222
XML Signature 222
XML-DSig 222
XMLHttpRequest, history 191
XQuery, and NoSQL 32
xUnit, and unit testing 157

Y

Yahoo!
 and internet scale 103
 data center
 build-out 12
 mega 20
 Quincy, WA 22
 security 75
 use of Hadoop 146
Yahoo! Mail, and visual testing 166
YouTube 194
 and internet scale 101
 cost of bandwidth 96
 shared-nothing partitioning 105
 storage 124

Z

Zend, and Simple Cloud API 202
zero-capital startup, using the cloud for
 64–67
zero-day attack 208
Zillow.com, and transference 101