

Index

A

Accelerated computing instances

- F1, 91

- G3, 91

- P3 and P2, 89–91

Amazon

- AWS (*see* Amazon Web Services (AWS))

- EC2, 17–18

- network security, 18

- RDS, 16–17

Amazon Cloud

- CloudWatch, 335

- master key creation, Servers

 - admin user, choosing, 296

 - alias and description, 296

 - AWS Services Console, 294

 - create Key button, 295

 - encryption keys, 295

 - IAM users and roles, 297

- monitoring RDS/EC2 instance

 - basic monitoring, 337

 - detailed monitoring, 337

 - enable enhanced

 - monitoring, 336

 - enabling alarms/alerts, 337,

 - 339, 341–344

 - search for CloudWatch, 336

Amazon Web Services (AWS)

- availability zones, 63–64

- backing up

 - EC2 instance, 239–244

 - RDS instance, 236–238

- CloudWatch dashboard,

 - 345–346

- connection pooling, 328–329

- create account, 65–67

- defined, 16

- EC2 (*see* Elastic Compute Cloud (EC2), Amazon)

- read replicas, 274–278

- RDS (*see* Relational Database Service (RDS))

- ANSI C, 33

- AUTOVACUUM launcher process, 41

- Autovacuum settings, 61

- Availability zones (AZ), 63–64

- AWS, *see* Amazon Web Services (AWS)

Azure

- connection pooling, 332–333

- Database, 22

- offers, 22

- PostgreSQL (*see* PostgreSQL software)

INDEX

Azure (*cont.*)
 storage, 188
 subscription, 22
 virtual machine (*see* Virtual machines)
Azure Cloud Services, 305–307

B

Background writer settings, 60
Binary installation, 47
Blob storage, 264
 add new storage, 266–268
 add to create storage, 265
Buckets, 254–257
Burstable performance instances, 82

C

CHECKPOINTER process, 40
Cloud
 -based cost model, 15
 computing
 pros and cons, 12–13
 vs. on-premise computing, 12–14
 migrating to, 24–28
Cloud block storage
 Attach Volume option, 129–132
 block storage volume, 125, 128
 volume details and actions, 128
Cloud security
 Compute Engine, 220

 Microsoft Azure Security, 227–228
 SQL Database, 232–234
 SSH Keys, 226–227
 VM instances, 220, 226
 VM machines, 228, 230–232
Cloud SQL encryption
 cloud proxy, 305
 SSL, 304
CloudWatch, 335–337, 345–346
Commit LOG (CLOG)
 buffers, 39
Compute Engine
 adding a member, 221–222
 available roles, 223–224
 creation, 140
 dashboard, 139
 features, 138, 225
 Google data centers, 137
 IAM console, 221
 OS, 143
 PostgreSQL installation, 148–150
 SSH keys, 148
 virtual machine, 141–142, 145–146
 VM security options, 144
Compute Optimized instances, 84–86
Connection pooling
 AWS (*see* Amazon Web Services (AWS))
 Azure, 332–333
 cost and time consuming, 311

- external connection poolers, 312
- forking, 310–311
- Google Cloud, 330, 332
- PgBouncer (*see* PgBouncer)
- Rackspace, 330

D

- Data at rest, 293
- Databases as a service (DBaaS)
 - in business, 2–3
 - deliver production, 2
 - enabling, 1–2
 - features
 - administration, 9
 - high availability, 10
 - monitoring tools, 10
 - provisioning
 - mechanisms, 7–9
 - scalability, 11
 - security, 11
 - limitations, 8
 - organization
 - cloud vendors, 6
 - collaboration with other team, 6
 - durable database, 4
 - fast-performing database, 4
 - high availability features, 6
 - horizontal scaling, 5
 - no single point of failures, 6
 - redundant, 4
 - reliability, 4

- secured database
 - environment, 3–4
 - testing, 2
- Data security model, 22
- Deployment model, 266

E

- Elastic Compute Cloud (EC2),
 - Amazon, 278–281
 - backup instance, 239–244
 - create customized dashboard, 346–348
 - create instance
 - Add Storage option, 97–98
 - add tag, 98
 - choose AMI, 95
 - choose instance type, 96
 - configure instance
 - details, 97
 - configuring security groups, 99
 - create key pair, 100
 - EC2 dashboard, 94
 - .pem file, 101–102
 - create PostgreSQL instance, 102–104
 - dedicated hosts, 81
 - dedicated instance, 81
 - elastic GPUs
 - accelerated computing, 89–91
 - compute optimized, 84–86
 - M3 instances, 83

INDEX

Elastic Compute Cloud (EC2),
 Amazon (*cont.*)
 M4 instances, 83
 R3 instances, 89
 R4 instances, 88
 storage optimized instances,
 92–94
 T2 instances, 82
 X1e instances, 86
 X1 instances, 87
instances, 79
IOPS, 17
RIs, 80
spot instances, 79

Encryption
 Amazon Cloud Servers
 EC2 instances, 299–300
 enabling, RDS instance,
 297–299
 master key creation, 294
 Azure Cloud Services, 305–307
 cloud SQL, 304–305
 Google cloud storage, 304
 rackspace cloud servers, 301–303

F

Fast-performing database, 4
Forking, 310–311

G

General purpose (SSD) storage
 (gp2), 70
Google Cloud

 connection pooling, 330, 332
 failover, 285
 features, 19
 Google Compute
 Engine, 19–20
 high availability, 283–284
 load balancing, 21
 monitoring instances
 PostgreSQL, 358–361
 project activities, 354–356
 VM, 357
 platform services, 20–21
 read replica, 286–288
 security model, 19
 SQL, 153
 using snapshot option,
 persistent disk, 250–251,
 253–254
 virtual machine, 20

Google Cloud Storage (GCS)
 Cloud SQL, 153
 create buckets and upload files,
 255–257
 features, 152
 PostgreSQL, 153–156
 storage classes, 151–152

Google Compute Engine, 19

Google Cloud Platforms (GCP)
 application servers/web
 servers, 133
 database servers, 134
 Google Cloud Console, 135–136
 project deletion, 136–137
 project quota, 134

H

- High availability cloud, *see*
Replication
- High availability features, 6

I, J

- IaaS, *see* Infrastructure as a
Service (IaaS)
- Identity and access management
(IAM)
 - access-level privileges, 200
 - AWS, 201–203, 205
 - multi-factor
 - authentication, 200
 - policy document, 201
 - types of access, 200
 - user management, 199
- Infrastructure as a
Service (IaaS)
 - cloud vendors, 23
 - in organizations, 23
 - subscribe, 24

K

- Key Management
Services (KMS), 294

L

- Live Migration, Google, 20
- Logical replication, 272

M

- Memory optimized instances
 - R3, 89
 - R4, 88
 - X1, 87
 - X1e, 86–87
 - Microsoft Azure
 - add new storage, 266–268
 - add to create storage, 265
 - blob storage, 264
 - database, 22
 - database for PostgreSQL, 269
 - data security model, 22
 - monitoring instances
 - database for PostgreSQL,
369–371
 - VM, 362–364, 367, 369
 - offers, 22
 - subscription, 22
 - virtual machines, backup on
 - backup policy, 262, 264
 - Backup tab, 260–261
 - options for backing up, 261
 - vaults, 261
 - Microsoft Azure Security, 227–228
 - Multi-regional storage classes, 151
 - Multi-Version Concurrency Control
(MVCC), 35
-
- N**
 - Nearline and Coldline storage
classes, 151

INDEX

O

One-click installers, 49

On-premise computing
cost model, 15
pros and cons, 13–14

P

PgBouncer

configuration, 321–322, 324–326
high availability, 327
installation
file creation, 318–320
PostgreSQL 9.6, 316–318
memory, 314
session pooling, 313
statement pooling, 314
transaction pooling, 313

pgPulse tools, 29

plProfiler Console, 150

Point-in-time recovery (PITR), 154

PostgreSQL service instance

advantages, 153–154
authorizing networks, 165–166
choose, database engine, 160
client-server protocol, 155
Cloud SQL flags, 166
configuration options,
161–162, 164
creation, 157–159
enabling auto backups, 164
features, 154
Google Cloud Console, 157

Google Cloud Dashboard, 157
Google Cloud SQL, 155
Quickstart for Cloud SQL, 159
set maintenance mode, 167
types, 155

PostgreSQL software, 4, 288–289

Amazon (*see* Amazon)

architecture, 37

ARCHIVER process, 41

AUTOVACUUM launcher
process, 41

autovacuum settings, 61

availability, 35

Azure database, 189–190,
369–371

background writer settings, 60

back up Cloud SQL, 258–259

CHECKPOINTER process, 40

cluster, 46

configuration file, 53

connection settings, 54

creation, 190–192, 194–196

environment variables, 49

extensions, 30

Google Cloud (*see* Google
Cloud)

installation

binary, 47

one-click installers, 49

RPM, 48

source, 44–46

limitations, 36

log, 57–59

LOGGING COLLECTOR, 41

- memory settings, 55
 - Microsoft Azure (*see* Microsoft Azure)
 - migrating to cloud, 24–28
 - monitoring instances on Google Cloud, 358–361
 - MVCC, 35
 - parameters, 4
 - PGDATA, 42–44
 - portability, 33
 - query planner settings, 56
 - Rackspace (*see* Rackspace)
 - reliability, 34
 - scalability, 34
 - secure, 34–35
 - security and authentication settings, 55
 - service startup on Linux, 50–51, 53
 - shared memory
 - CLOG buffers, 39
 - lock space, 39
 - shared_buffers, 39
 - temp buffers, 39
 - WAL buffers, 39
 - SLRU, 40
 - STATS COLLECTOR process, 41
 - supports, 35
 - tools, 29–30
 - vacuum cost settings, 60
 - WAL, 36
 - WAL WRITER process, 41
 - write ahead log settings, 56–57
 - WRITER process, 40
 - Provisioned IOPS (SSD) Storage, 70
- Q**
- Query planner settings, 56
- R**
- Rackspace
 - backups on cloud
 - check backup details, 249
 - check box, 245
 - choose folder, 249
 - Configure Backup page, 247–248
 - create/restore/delete, 245
 - manage backups, 245
 - schedule, 246–247
 - cloud block storage (*see* Cloud block storage)
 - cloud servers, 301–303, 349, 351–354
 - connection pooling, 330
 - Managed Cloud, 107
 - Managed Hosting
 - dedicated server, 111, 113–117
 - portal dashboard, 108
 - PostgreSQL, 121–124
 - profile settings, 108–109
 - virtual server, 117–120
 - VMWare, 110

INDEX

Rackspace (*cont.*)

- offers, 18
- replication and high availability, 281
- supports, 18

Rackspace cloud security

- Rackspace account
 - barcode, 216
 - identity provider, 218
 - mobile app, 214
 - My Profile & Settings, 213
 - security settings, 213
 - user management, 217
 - user permissions, 217

VM, 218-219

Regional storage classes, 151

Relational Database Service (RDS)

- backing up instance, 236-238
- create PostgreSQL instance
 - advanced settings
 - page, 75, 77
 - backups, 77-79
 - choose database name and port, 77
 - correct range for entry, 74
 - launch DB instance, 71
 - maintenance, 78-79
 - monitoring, 77, 79
 - select engine, 72
 - select use case, 72-73
 - settings, 75
- enables, 68
- gp2, 70
- offers, 68

Provisioned IOPS (SSD)

Storage, 70

subnet groups, 207-211

VPC, 206

Replication

AWS (*see* Amazon Web Services (AWS))

Azure services (*see* Azure)

Google Cloud (*see* Google Cloud)

logical, 272

PostgreSQL, 272

Rackspace cloud, 281

streaming, 272

Reserved Instances (RIs), 80

RPM installation, 48

S

sar tools, 29

Session pooling, 313

Single sign-on (SSO), 200

SLRU system, 40

Snapshot option, 250-251, 253-254

Statement

pooling, 314

STATS COLLECTOR

process, 41

Storage optimized instances

dense-storage (D2), 94

high I/O (I3), 92-93

Streaming replication, 272

Subnet, 206

T, U

Transaction pooling, 313

V

Vacuum cost settings, 60

Virtual machines, 289–290

 account creation, 171

 backup on

 backup policy, 262, 264

 Backup tab, 260–261

 options for backing

 up, 261

 vaults, 261

 creation

 configuration settings,
 175–176

 data and metrics, 183–184

 deployment process,
 182–183

 extensions section, 181–182

 features, 179–180

 locations, 177

 operating systems, 174

 size, 178

 Linux, 184–185

 PostgreSQL, 185–187

 Start Free, 171

Virtual Private Cloud (VPC), 206

W

WAL WRITER process, 41

Write Ahead Logging (WAL), 36

Write ahead log settings, 56–57

WRITER process, 40

X, Y, Z

X1e instances, 86