

**Mastering**

# GitLab 12

Implement DevOps culture and repository management solutions



Joost Evertse

**Packt** >

[www.packt.com](http://www.packt.com)

# Mastering GitLab 12

Implement DevOps culture and repository management solutions

**Joost Evertse**

**Packt**

BIRMINGHAM - MUMBAI

# Mastering GitLab 12

Copyright © 2019 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing or its dealers and distributors, will be held liable for any damages caused or alleged to have been caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

**Commissioning Editor:** Pavan Ramchandani

**Acquisition Editor:** Rohit Rajkumar

**Content Development Editor:** Drashti Panchal

**Senior Editor:** Rahul Dsouza

**Technical Editor:** Komal Karne

**Copy Editor:** Safis Editing

**Project Coordinator:** Jagdish Prabhu

**Proofreader:** Safis Editing

**Indexer:** Pratik Shirodkar

**Production Designer:** Aparna Bhagat

First published: August 2019

Production reference: 2201219

Published by Packt Publishing Ltd.

Livery Place

35 Livery Street

Birmingham

B3 2PB, UK.

ISBN 978-1-78953-128-2

[www.packt.com](http://www.packt.com)

*For my family, who supported me throughout the entire effort of writing this book.*



Packt.com

Subscribe to our online digital library for full access to over 7,000 books and videos, as well as industry leading tools to help you plan your personal development and advance your career. For more information, please visit our website.

## Why subscribe?

- Spend less time learning and more time coding with practical eBooks and Videos from over 4,000 industry professionals
- Improve your learning with Skill Plans built especially for you
- Get a free eBook or video every month
- Fully searchable for easy access to vital information
- Copy and paste, print, and bookmark content

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at [www.packt.com](http://www.packt.com) and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at [customercare@packtpub.com](mailto:customercare@packtpub.com) for more details.

At [www.packt.com](http://www.packt.com), you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on Packt books and eBooks.

# Contributors

## About the author

**Joost Evertse** is an all-round professional with over 20 years of experience in IT in the financial and telecom sectors. He has worked for big and small organizations and has lived in different worlds, including Unix, Oracle, Java, and Windows. Creating order from chaos has been a big focus during his system-engineering years. After 10 years of system administration, he moved into software development and started using CI/CD tools, including GitLab.

At the end of 2016, he started at a significant financial company in the GitLab team, shifting his focus more toward the entire CI/CD pipeline, with the mission of making the CI/CD platform more stable and highly available. His team eventually migrated GitLab to a private cloud and improved release cycles.

## About the reviewer

**Orlando Monreal** is a software engineer with over 12 years of experience, currently working at HCL Technologies Mexico, as part of the Source Code Management team in his project account. He has worked with GitLab applications as an administrator and contact for application support queries, handling upgrade processes and troubleshooting performance and configuration-related issues with the application.

## Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit [authors.packtpub.com](https://authors.packtpub.com) and apply today. We have worked with thousands of developers and tech professionals, just like you, to help them share their insight with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

# Table of Contents

<b>Preface</b>	1
<hr/>	
<b>Section 1: Install and Set Up GitLab On-Premises or in the Cloud</b>	
<hr/>	
<b>Chapter 1: Introducing the GitLab Architecture</b>	8
<b>Technical requirements</b>	9
<b>The origins of GitLab</b>	9
<b>Exploring GitLab editions – CE and EE</b>	10
<b>The core system components of GitLab</b>	11
NGINX	12
Debugging NGINX	12
Unicorn	13
Debugging Unicorn	13
Timeouts in Unicorn logs	13
Unicorn processes disappear	14
Other kinds of errors or 100% CPU load	14
Sidekiq	15
Debugging Sidekiq	16
Sidekiq processes disappear	16
A Sidekiq process is seemingly doing nothing	16
Other kind of errors or 100% CPU load	17
GitLab Shell	18
Debugging GitLab Shell	18
Redis	20
Basic data operations in Redis	20
Gitaly	22
Debugging Gitaly	24
GitLab Workhorse	25
Debugging GitLab Workhorse	25
Database	26
Debugging PostgreSQL	29
<b>GitLab CI</b>	29
Pipelines and jobs	31
<b>GitLab Runners</b>	32
Issues with the old runner	33
Switching to Go	33
<b>Cloud native</b>	36
<b>Summary</b>	37
<b>Questions</b>	37
<b>Further reading</b>	38



<b>Chapter 2: Installing GitLab</b>	39
<b>Technical requirements</b>	39
<b>Installing GitLab using omnibus packages</b>	40
Omnibus structure	40
Project definition	41
Individual software definitions	41
A GitLab configuration template	41
Chef components	41
Runit recipe	42
Tests	42
gitlab-ctl commands	42
Running the installer	43
Browsing to the external URL and login	44
Upgrade using the omnibus-gitlab package	44
<b>Running from source</b>	44
Operating system – Debian 10	45
Required basic software packages	45
Required programming languages	47
Ruby	47
Go	48
Node.js	48
System users	49
SQL database	50
Redis memory database	52
GitLab	53
Installing GitLab Shell	56
Installing GitLab-Workhorse	56
Installing Gitaly	56
Initializing the database and activating advanced features	57
Final steps for preparing the system	57
Preparing to serve	58
Compiling GetText PO files	58
Compiling assets	59
Starting your GitLab instance	59
NGINX	60
<b>Using it from Docker</b>	61
Running the image directly	62
Configuring GitLab after startup	63
Starting the container with configuration settings as input	64
Upgrading GitLab	65
Run GitLab CE on a different IP address	66
Debugging the container	66
Install GitLab using Docker Compose	66
Updating GitLab using Docker Compose	67
<b>Deploying GitLab using Kubernetes</b>	68
GitLab Runner Helm chart	68
Deploying of a GitLab Runner to Kubernetes	69
GitLab Helm chart	69
Deploying GitLab to Kubernetes	70

Monitoring the deployment	71
Initial login	71
Outgoing email	72
Updating GitLab using the Helm chart	72
Uninstalling GitLab using the Helm chart	73
Creating droplets on DigitalOcean	73
<b>Summary</b>	76
<b>Questions</b>	76
<b>Further reading</b>	77
<b>Chapter 3: Configuring GitLab Using the Web UI</b>	78
<b>Technical requirements</b>	78
<b>Configuring GitLab settings at the instance level</b>	79
Menu options	80
Monitoring	80
Messages	81
System hooks	82
Plugins	83
Applications	83
Abuse reports	84
License	85
Kubernetes	86
Push rules	86
Geo	86
Deploy Keys	87
Service templates	87
Appearance	88
Settings	89
General	89
Visibility and access controls	89
Account and limit	90
Diff limits	90
Sign-up restrictions	90
Sign-in restrictions	91
Terms of service and privacy policy	92
External authentication	92
Web Terminal	92
Web IDE	92
Integrations	93
Elasticsearch	93
PlantUML	94
Third-party offers	94
Snowplow	95
Repository	95
Repository mirror	95
Repository storage	95
Repository maintenance	96
Templates	97
CI/CD	97

Auto DevOps settings	97
Shared runner settings	98
Container registry	99
Reporting	99
Spam and anti-bot protection	99
Abuse reports	99
Error reporting and logging	100
Metrics and profiling	100
Metrics – InfluxDB	100
Metrics – Prometheus	100
Profiling – Performance Bar	101
Usage statistics	101
Pseudonymizer data collection	101
Network	102
Performance optimization	102
User and IP rate limits	102
Outbound requests	102
Geo	102
Preferences	102
Email	102
Help page	103
Pages	104
Real-time features	105
Gitaly	105
Localization	105
<b>Configuring GitLab settings at the group level</b>	105
<b>Configuring GitLab settings at the project level</b>	107
General	107
Naming, topics, avatar	108
Visibility, project features, permissions	109
Merge requests	109
<b>Summary</b>	111
<b>Questions</b>	111
<b>Further reading</b>	111
<b>Chapter 4: Configuring GitLab from the Terminal</b>	112
<b>Technical requirements</b>	112
<b>Configuring omnibus and GitLab installations from the terminal</b>	113
<b>Configuring source installations</b>	114
GitLab app settings	114
Storing big files	117
Using object storage	118
GitLab pages	120
Mattermost	121
Gravatar	121
Sidekiq	122
GitLab Registry	122
GitLab CI settings	123

Auth settings	124
Advanced settings	126
Rack Attack	128
<b>Reconfiguring GitLab Docker containers</b>	129
<b>Changing GitLab in a Kubernetes environment</b>	130
Basic configuration	130
Configuring TLS	130
Configuring outgoing emails	131
Other settings	131
<b>Summary</b>	132
<b>Questions</b>	132
<b>Further reading</b>	133
<hr/> <b>Section 2: Migrating Data from Different Locations</b> <hr/>	
<b>Chapter 5: Importing Your Project from GitHub to GitLab</b>	135
<b>Technical requirements</b>	136
<b>Using the GitHub integration feature</b>	136
Preparing GitHub for export	137
Preparing GitLab for import	140
Running the import	141
<b>Using a GitHub token</b>	144
Preparing GitHub for export	144
Running the import	146
<b>Using a GitLab rake task</b>	148
Preparing GitLab for import	148
Running the import	148
<b>Summary</b>	151
<b>Questions</b>	151
<b>Further reading</b>	151
<b>Chapter 6: Migrating from CVS</b>	152
<b>Technical requirements</b>	153
<b>CVS versus Git</b>	153
Filesets versus changesets	154
Git branching	155
Creating repositories	156
Atomic operations	157
Object naming or referencing versions	157
Keyword substitution	157
Binary blobs	158
Amending commits	159
A central repository	160
Accompanying toolset	160
Detecting file renames	161
Commit before merge	161

<b>Preparing to migrate from CVS to Git</b>	161
Preparing for a conversion using cvs-fast-export	162
Preparing for a conversion using cvs2git	164
<b>Running the conversion</b>	165
Converting data using cvs-fast-export	165
Converting data using cvs2git	167
<b>Cleaning up after migration</b>	168
<b>Summary</b>	169
<b>Questions</b>	169
<b>Further reading</b>	170
<b>Chapter 7: Switching from SVN</b>	171
<b>Technical requirements</b>	171
<b>The difference between SVN and Git</b>	172
Security and access control	173
Space requirements and references	174
Branching	175
Handling binaries with SVN and Git	176
<b>Mirroring SVN and GIT</b>	178
No sync, just convert	182
<b>Using svn2git to migrate in one cut</b>	183
<b>Summary</b>	185
<b>Questions</b>	186
<b>Further reading</b>	186
<b>Chapter 8: Moving Repositories from TFS</b>	187
<b>Technical requirements</b>	187
<b>TFS versus Git</b>	188
<b>Centralized or decentralized</b>	189
Handling changes	189
Branching and merging capacity	190
History	191
Traceability	192
File handling	192
<b>The git-tfs tool</b>	193
Preparing to migrate	194
<b>Summary</b>	200
<b>Questions</b>	201
<b>Further reading</b>	201
<b>Section 3: Implement the GitLab DevOps Workflow</b>	
<b>Chapter 9: GitLab Vision - the Whole Toolchain in One Application</b>	203
<b>Technical requirements</b>	203
<b>The Agile Manifesto</b>	204

The initial model – waterfall	204
Royce's model	206
The sashimi model	207
The V-model	208
DSDM	209
Timeboxing	210
MoSCoW	211
Scrum	212
Crystal methods	213
<b>XP</b>	214
Fine-scale feedback	214
Planning game	215
Release planning	215
Iteration planning	218
Pair programming	219
Test Driven Development	220
Continuous processes	222
Continuous integration	222
Refactoring	223
Short iterations	224
Everybody owns the code	224
Shared understanding	224
Coding standards	225
Simple design	225
System metaphor	225
<b>The DevOps movement</b>	226
History of the movement	227
Four Quadrant Model	228
Four levels of competence	230
<b>The toolchain</b>	231
<b>Summary</b>	235
<b>Questions</b>	235
<b>Further reading</b>	236
<b>Chapter 10: Create Your Product, Verify, and Package it</b>	237
<b>Technical requirements</b>	237
<b>The GitLab workflow</b>	238
<b>DevOps phase – manage your ideas</b>	239
Cycle analytics	239
<b>DevOps phase – plan your feature</b>	241
Issues	241
Content	241
Status of the issue	241
Meta information	242
Discussions	243
Milestones	246
Epics	248

Time tracking	249
Quick actions	251
The Project Issue board	252
Todos	255
<b>DevOps phase – create it</b>	256
Projects and groups	256
Snippets	260
Web IDE	261
Wiki	261
Protected branches	263
Merge requests	265
<b>DevOps phase – verify your product</b>	270
Code Quality reports	271
Review apps	273
<b>DevOps phase – package it for use</b>	278
GitLab container registry	278
<b>Summary</b>	281
<b>Questions</b>	281
<b>Further reading</b>	281
<b>Chapter 11: The Release and Configure Phase</b>	282
<b>Technical requirements</b>	282
<b>Continuous Deployment</b>	283
<b>Auto DevOps</b>	289
Configuring Auto DevOps	289
Build step	291
Code quality scan	292
Container scanning	292
Dependency scanning	294
License management	294
Static application security testing (sast)	295
The final test step	295
Production	296
Performance	299
<b>Summary</b>	300
<b>Questions</b>	301
<b>Further reading</b>	301
<b>Chapter 12: Monitoring with Prometheus</b>	302
<b>Technical requirements</b>	302
<b>Setting up Prometheus</b>	303
Using an external Prometheus host	309
Enabling the external dashboard link	311
<b>Customizing monitoring</b>	311
<b>The static analysis of security vulnerabilities</b>	313
<b>Dynamic Application Security Testing</b>	317

<b>Dependency checking</b>	320
<b>Summary</b>	324
<b>Questions</b>	325
<b>Further reading</b>	325
<b>Chapter 13: Integrating GitLab with CI/CD Tools</b>	326
<b>Technical requirements</b>	326
<b>Using Jira with GitLab</b>	327
<b>Connecting Jenkins to GitLab</b>	334
<b>Integrating with Mattermost</b>	344
<b>Using webhooks for events</b>	354
<b>Summary</b>	359
<b>Questions</b>	360
<b>Further reading</b>	360
<b>Section 4: Utilize GitLab CI and CI Runners</b>	
<b>Chapter 14: Setting Up Your Project for GitLab Continuous Integration</b>	362
<b>Technical requirements</b>	363
<b>Pipelines</b>	363
<b>Jobs</b>	368
<b>Creating .gitlab-ci.yml</b>	369
<b>Configuring a runner</b>	370
GitLab Runner features	371
<b>Summary</b>	374
<b>Questions</b>	374
<b>Further reading</b>	374
<b>Chapter 15: Installing and Configuring GitLab Runners</b>	375
<b>Technical requirements</b>	375
<b>The Runner client architecture</b>	376
Basic architecture	376
<b>Creating a basic Runner with the shell executor</b>	379
Installing Runner on Linux	380
Using a package manager	380
Using a manual installation	382
Updating a manually installed runner binary	383
Installing on Mac	383
The manual way of installing a runner	384
Installing and using the Homebrew package manager	384
Updating a manually installed runner binary	384
Installing on Windows	385
<b>Registering a runner</b>	386
The interactive way of registering a runner	386
The non-interactive way of registering a runner	387
<b>Running the nightly version</b>	389



<b>Summary</b>	390
<b>Questions</b>	390
<b>Further reading</b>	390
<b>Chapter 16: Using GitLab Runners with Docker or Kubernetes</b>	391
<b>Technical requirements</b>	392
<b>Runner client architecture</b>	393
<b>Creating your own Dockerized GitLab Runner</b>	395
<b>Using a prebuilt Docker container to deploy GitLab Runners</b>	401
<b>Using a Kubernetes cluster to spawn GitLab Runners</b>	404
<b>Summary</b>	408
<b>Questions</b>	408
<b>Further reading</b>	408
<b>Chapter 17: Autoscaling GitLab CI Runners</b>	409
<b>Technical requirements</b>	410
<b>Runner client architecture</b>	410
<b>Setting up the environment</b>	412
Preparing a bastion host	412
Deploying the GitLab Runner software	412
Installing Docker Machine	413
<b>Configuring the Runner</b>	414
Off-peak time mode configuration	414
Distributed runners caching	415
Setting the cache globally	415
Setting the cache at the project level	415
Distributed container registry mirroring	416
Installing and running a proxy container registry and a caching server	417
Proxy container registry	417
Caching server	418
Creating an S3 bucket in Amazon Web Services	418
Creating your own MinIO server	419
Scaling your runners	420
Using Docker Machines with a local VirtualBox instance	420
Using docker machines that have been created on Amazon Web Services (EC2)	421
<b>Summary</b>	423
<b>Questions</b>	424
<b>Further reading</b>	424
<b>Chapter 18: Monitoring CI Metrics</b>	425
<b>Technical requirements</b>	426
<b>Enabling monitoring for Runners</b>	426
<b>Editing the GitLab Runner configuration file</b>	428
<b>Runner business logic metrics</b>	433
Key metrics to watch	433

<b>General process metrics</b>	438
Key metrics to watch	439
Alert management	443
<b>Summary</b>	448
<b>Questions</b>	448
<b>Further reading</b>	449
<b>Section 5: Scale the Server Infrastructure (High Availability Setup)</b>	
<hr/>	
<b>Chapter 19: Creating a Basic HA Architecture Using Horizontal Scaling</b>	451
<b>Technical requirements</b>	452
<b>The underlying architecture of this solution</b>	453
Amazon services	455
Elastic compute cloud (EC2)	455
Classic load balancer	456
Virtual private cloud and subnets	456
SGs	456
Terraform	457
Installing Terraform	457
Ansible	458
Installing the Ansible Terraform provider	458
Starting with the code	459
vpc.tf	460
subnet.tf	461
instance.tf	462
ansible_host.tf	463
route_table.tf	465
security_group.tf	465
variable.tf	466
keypair.tf	467
lb.tf	468
providers.tf	470
Preparing to run Terraform to deploy the virtual hardware	470
Running the deployment	471
<b>Setting up the bastion hosts</b>	473
<b>Configuring the database nodes</b>	474
Contents of the gitlab.rb.postgres.j2 template	474
Configuring the consul nodes	477
Contents of gitlab.rb.consul.j2	477
Configuring the PgBouncer node	478
Contents of gitlab.rb.pgbouncer.j2	478
<b>The Redis configuration</b>	480
Contents of gitlab.rb.redis.j2	480
<b>Connecting the shared filesystem</b>	481
Contents of nfs_exports.j2	482

<b>Setting up the application servers</b>	482
Contents of gitlab.rb.j2	482
Running all the Ansible playbooks	484
<b>Summary</b>	488
<b>Questions</b>	488
<b>Further reading</b>	489
<b>Chapter 20: Managing a Hybrid HA Environment</b>	490
<b>Technical requirements</b>	491
<b>The basic architecture of this solution</b>	491
<b>A renewed Terraform configuration</b>	493
instance.tf	493
ansible_host.tf	495
<b>Splitting application components into frontend and middleware tiers</b>	498
Splitting Sidekiq from the frontend	498
Contents of gitlab.rb.middleware.j2	498
Creating a monitoring instance	500
Contents of gitlab.rb.prometheus.j2	500
Creating a monitoring dashboard with Grafana	501
Contents of gitlab.rb.grafana.j2	501
<b>Connecting the shared filesystem</b>	502
Contents of the gitlab.rb.gitally.j2 file	502
<b>Changes in Ansible files</b>	503
<b>Script enhancements</b>	504
<b>Summary</b>	505
<b>Questions</b>	505
<b>Further reading</b>	506
<b>Chapter 21: Making Your Environment Fully Distributed</b>	507
<b>Technical requirements</b>	507
<b>The basic architecture of this solution</b>	508
<b>Performing changes to the Terraform configuration</b>	510
instance.tf	510
ansible_host.tf	512
<b>Splitting more application components</b>	514
The third application server for Git SSH	515
Contents of gitlab.rb.frontend_ssh.j2	515
The middleware layer – Sidekiq	516
The ASAP Sidekiq instance	516
Contents of gitlab.rb.sidekiq_asap.j2	517
The real-time Sidekiq instance	517
Contents of gitlab.sidekiq.realtime.j2	518
The pipeline Sidekiq instance	518
Contents of gitlab.sidekiq.pipeline.j2	519
The normal Sidekiq instance	519

<b>The clustered Redis/Sentinel configuration</b>	519
Contents of gitlab.rb.redis-cluster.j2	520
<b>Changes in Ansible files</b>	521
<b>Script enhancements</b>	522
<b>Summary</b>	523
<b>Questions</b>	523
<b>Further reading</b>	523
<b>Chapter 22: Using Geo to Create Distributed Read-Only Copies of GitLab</b>	524
<b>Technical requirements</b>	524
<b>The basic architecture of this solution</b>	525
<b>Preparing the infrastructure</b>	526
The root module explained – main.tf	528
Utilizing the keypair module – modules/services/keypair/main.tf	529
Variables for the keypair module – modules/services/keypair/vars.tf	530
Outputs from the keypair module – modules/services/keypair/output.tf	530
Explaining the GitLab module	530
The main module file – modules/services/gitlab/main.tf	530
The variable file – modules/services/gitlab/vars.tf	532
The outputs for the module – modules/services/gitlab/outputs.tf	533
<b>Setting up Geo</b>	534
Installing the GitLab software and license	534
Contents of gitlab.rb.j2	535
Preparing the database replicas	535
Contents of gitlab.rb.primary.j2	536
Contents of gitlab.rb.primary.j2	537
Changing the SSH key lookup method to the database	538
Adding the secondary node via the web UI	538
Activating hashed storage	539
Checking the status of the secondary node	541
<b>Summary</b>	542
<b>Questions</b>	542
<b>Further reading</b>	542
<b>Assessments</b>	543
<b>Other Books You May Enjoy</b>	552
<b>Index</b>	555

---

# Preface

GitLab is a tool to enhance the workflow of teams and enable parts of the DevOps life cycle. It started out as a tool only for source code management, but today, GitLab can offer help ranging from managing an initial idea to building and testing source code, all the way from development to production.

You'll learn ways to use all of the features available in GitLab to enhance your business via the integration of all phases of the development process. You'll benefit from lower friction by creating one platform on-premises or in the cloud, increase collaboration, and drive competitive advantage with more efficient operations.

## Who this book is for

This book is for developers and DevOps professionals who want to master the software development workflow in GitLab and boost their productivity by putting their teams to work on GitLab via an on-premise installation or cloud-based infrastructure.

## What this book covers

Chapter 1, *Introducing the GitLab Architecture*, provides a short introduction to the company and the people that created the product, along with a high-level overview of GitLab and its components.

Chapter 2, *Installing GitLab*, shows you how to install and configure GitLab via several different methods. This can be done from scratch, or via the Omnibus installer. Special attention is given to Docker and Kubernetes when outlining containerized solutions. Finally, a cloud installation via the DigitalOcean infrastructure is taken as an example.

Chapter 3, *Configuring GitLab Using the UI*, explains the options in the GitLab web UI that can be configured after installation. This chapter also covers the administration pages where these instance-level options are situated.

Chapter 4, *Configuring GitLab from the Terminal*, looks at the different ways of configuring GitLab. The first approach is by using the Omnibus package installer provided by GitLab, which automates most of the installation. The chapter continues with configuring a source installation. Configuring Docker containers and managing a Kubernetes installation are also covered.

Chapter 5, *Importing Your Project from GitHub to GitLab*, outlines the process of migration from GitHub via a hands-on lab. It starts by exploring settings that should be altered in your GitHub project. After this, the settings necessary in GitLab to prepare an import are shown, and finally, the procedure for running the import is addressed.

Chapter 6, *Migrating from CVS*, begins with a comparison of the fundamentally different systems of CVS and Git. It then provides directions on preparing for migration. Actual conversion is addressed, as is the cleaning up of artifacts not needed anymore.

Chapter 7, *Switching from SVN*, begins by explaining the subtle and not-so-subtle differences between SVN and Git. The reader is shown how to migrate using two different methods: mirroring with SubGit and using the `svn2git` tool.

Chapter 8, *Moving Repositories from TFS*, first deals with the differences between TFS and Git. Subsequently, the act of migrating information from a TFS project to Git is shown via the use of the `git-tfs` tool.

Chapter 9, *GitLab Vision - the Whole Toolchain in One Application*, explains GitLab's vision of providing the whole DevOps toolchain to the developer, looking at the origins of XP and the Agile manifesto. The emergence of the DevOps paradigm is also explored, and the toolchain that GitLab provides is summarized.

Chapter 10, *Create Your Product, Verify It, and Package It*, shows how the product vision for GitLab and its workflow is centered around the idea of providing a complete toolchain to create a product. This chapter focuses on the different phases and explains the relevant concepts with examples.

Chapter 11, *The Release and Configure Phase*, discusses one of the big features of GitLab: the ability to offer the complete journey to production with different, easy-to-design stages. This way, you can create different environments and, ultimately, automate the whole pipeline for a product.

Chapter 12, *Monitoring with Prometheus*, handles ways of monitoring your GitLab environment by using the built-in Prometheus feature and default scripting languages. The second part of this chapter explains the different security tests that are available.

Chapter 13, *Integrating GitLab with CI/CD Tools*, explains how, although GitLab aims to provide a complete toolchain in the real world, there will always be a need for integration. This chapter explains some of the bigger possible integrations that are configurable out of the box. It closes with a section on how webhooks provide a general way to consume information from GitLab.

Chapter 14, *Setting Up Your Project for GitLab Continuous Integration*, describes GitLab CI concepts that are present on the application server and can be fine-tuned and customized per project. The second part of the chapter mainly focuses on how to get your project ready to use these CI concepts and set up a runner for it to use.

Chapter 15, *Installing and Configuring GitLab Runners*, explains the way GitLab runners work, by installing them. The next step is creating an example project and building it with a shell executor.

Chapter 16, *Using GitLab Runners with Docker or Kubernetes*, examines the architecture of Docker-based runners and runners using the Kubernetes API, using the same examples as in earlier chapters.

Chapter 17, *Autoscaling GitLab CI Runners*, demonstrates the architecture of runners using autoscaling. The number of runners required will decrease and increase based on demand. The example shown uses VirtualBox and **Amazon Web Services (AWS)** to deploy instances.

Chapter 18, *Monitoring CI Metrics*, deals with monitoring specific GitLab runners. Using a lab, we demonstrate how to enable monitoring inside the runner. After this introduction, the specific functional and system metrics are explained.

Chapter 19, *Creating a Basic HA Architecture by Using Horizontal Scaling*, visualizes the way in which different components interact. Secondly, the preparation of databases is shown, as well as several all-in-one application servers. Finally, the shared filesystem for repositories and Redis caching in this **high availability (HA)** setup is explained. We will use Terraform and Ansible to create the demonstration environment.

Chapter 20, *Managing a Hybrid HA Environment*, builds on the earlier architecture of horizontal HA, but continues to grow in complexity. The main difference is that the application servers combined several components that are now split into new tiers.

Chapter 21, *Making Your Environment Fully Distributed*, builds on earlier chapters. A fully distributed architecture aims to create more fault tolerance by again splitting components into new tiers. There is now an SSH node and several sidekiq tiers.

Chapter 22, *Using Geo to Create Distributed Read-Only Copies of GitLab*, starts with an explanation of the GEO product, which is part of the Enterprise Edition license. Using the same tools as in earlier chapters from Section 5 of this book 'Scale the Server Infrastructure (High Availability Setup)', we will explain how to set up GEO to create replication between two different geographical locations.

## To get the most out of this book

To get the most out of this book, you should have access to a Linux or macOS machine, have an internet connection, and have Amazon AWS, Google, and Microsoft Azure accounts. These are all necessary to run the examples.

Some basic IT knowledge is necessary to read this book. The subjects you need experience in are as follows:

- Linux
- Shell scripting
- Basic programming skills in Ruby and JavaScript
- A basic understanding of Docker containers
- A basic understanding of using Terraform to create infrastructure as code
- A basic understanding of Ansible

## Download the example code files

You can download the example code files for this book from your account at [www.packt.com](http://www.packt.com). If you purchased this book elsewhere, you can visit [www.packtpub.com/support](http://www.packtpub.com/support) and register to have the files emailed directly to you.

You can download the code files by following these steps:

1. Log in or register at [www.packt.com](http://www.packt.com).
2. Select the **Support** tab.
3. Click on **Code Downloads**.
4. Enter the name of the book in the **Search** box and follow the onscreen instructions.

Once the file is downloaded, please make sure that you unzip or extract the folder using the latest version of:

- WinRAR/7-Zip for Windows
- Zipeg/iZip/UnRarX for Mac
- 7-Zip/PeaZip for Linux

The code bundle for the book is also hosted on GitHub at <https://github.com/PacktPublishing/Mastering-GitLab-12>. In case there's an update to the code, it will be updated on the existing GitHub repository.



We also have other code bundles from our rich catalog of books and videos available at <https://github.com/PacktPublishing/>. Check them out!

## Code in Action

Visit the following link to see the code being executed:

<http://bit.ly/2KirIoO>

## Download the color images

We also provide a PDF file that has color images of the screenshots/diagrams used in this book. You can download it here: [https://static.packt-cdn.com/downloads/9781789531282\\_ColorImages.pdf](https://static.packt-cdn.com/downloads/9781789531282_ColorImages.pdf).

## Conventions used

There are a number of text conventions used throughout this book.

**CodeInText**: Indicates code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. Here is an example: "Let's continue with installing web documents in `/usr/local/www`."

A block of code is set as follows:

```
server {  
  listen 8080;  
  server_name localhost;
```

When we wish to draw your attention to a particular part of a code block, the relevant lines or items are set in bold:

```
server {  
  listen 8080;  
  server_name localhost;
```

Any command-line input or output is written as follows:

```
$mkdir /usr/local/www  
$chmod 755 /usr/local/www  
$cd /usr/local/www
```

**Bold:** Indicates a new term, an important word, or words that you see onscreen. For example, words in menus or dialog boxes appear in the text like this. Here is an example: "You can do this by clicking the **Choose File** button near the **Logo** section."



Warnings or important notes appear like this.



Tips and tricks appear like this.

## Get in touch

Feedback from our readers is always welcome.

**General feedback:** If you have questions about any aspect of this book, mention the book title in the subject of your message and email us at [customercare@packtpub.com](mailto:customercare@packtpub.com).

**Errata:** Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you have found a mistake in this book, we would be grateful if you would report this to us. Please visit [www.packtpub.com/support/errata](http://www.packtpub.com/support/errata), selecting your book, clicking on the Errata Submission Form link, and entering the details.

**Piracy:** If you come across any illegal copies of our works in any form on the Internet, we would be grateful if you would provide us with the location address or website name. Please contact us at [copyright@packt.com](mailto:copyright@packt.com) with a link to the material.

**If you are interested in becoming an author:** If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, please visit [authors.packtpub.com](http://authors.packtpub.com).

## Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions, we at Packt can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about Packt, please visit [packt.com](http://packt.com).

# 1

## Section 1: Install and Set Up GitLab On-Premises or in the Cloud

This section will give you a solid understanding of GitLab deployment options and GitLab component architecture, leaving you able to install and configure GitLab on-premises and in the cloud.

This section comprises the following chapters:

- Chapter 1, *Introducing the GitLab Architecture*
- Chapter 2, *Installing GitLab*
- Chapter 3, *Configuring GitLab Using the Web UI*
- Chapter 4, *Configuring GitLab from the Terminal*

# 1

## Introducing the GitLab Architecture

Understanding the context of the GitLab project will help us to appreciate the choices that were made with regard to the design of the GitLab workflow. The GitLab project started out as a small, open source project, and has grown to be an organization of 400 people and thousands of volunteers. It is currently available in two versions, a free **Community Edition (CE)** and an **Enterprise Edition (EE)** with a proprietary license. There are several tiers of support for the enterprise version. Although it is proprietary licensed, the source code for that version is publicly available from GitLab.

To master GitLab, it is necessary to have a solid understanding of its individual components. In this chapter, we will look at the basic components of a GitLab installation, paying special attention to GitLab **Continuous Integration (CI)** and the accompanying runners. As the different components can be distributed across servers or even cloud providers, we will also provide an overview of those providers and how GitLab views them.

In this chapter, we will be covering the following topics:

- The origins of GitLab
- GitLab CE or EE
- The core components of GitLab
- GitLab CI
- GitLab Runners
- Cloud native

## Technical requirements

To follow along with the instructions in this chapter, please download the Git repository with examples, commands and instructions, available at GitHub: <https://github.com/PacktPublishing/Mastering-GitLab-12/tree/master/Chapter01>. Look in the `Readme.md` file for a general explanation of the content of the directory.

To run or install software used in this chapter you need one of the following platforms:

- Debian 10 Linux codename 'Buster'
- CentOS 7.x or RHEL (Red Hat Enterprise Linux) 7.x
- macOS Sierra or later

## The origins of GitLab

The story began in 2011, when Dimitri Zaporozhets, a web programmer from Ukraine, was faced with a common problem. He wanted to switch to Git for version management and GitHub to collaborate, but that was not allowed in his company. He needed a tool that did not hinder him in developing code and was easy to use. Like many developers, he had issues with the collaboration tool that he was obliged to use. To get around those issues, he created his side project in Ruby on Rails: GitLab. Together with his colleague, Valery Sizov, he developed this project alongside his regular work.

After this initiative, the project grew enormously:

Date	Fact
2011	Sytze Sybrandij, the future CEO of GitLab, is impressed by the GitLab project and code, and offers Zaporozhets the opportunity to try to commercialize it via <a href="https://about.gitlab.com/">https://about.gitlab.com/</a> .
2012	GitLab was announced to a broader audience via Hacker News ( <a href="https://news.ycombinator.com/item?id=4428278">https://news.ycombinator.com/item?id=4428278</a> ).
2013	Dimitri Zaporozhets decides to work full-time on GitLab and joins the company.
2015	GitLab becomes part of the Y Combinator class and received VC funding that year.
2018	GitLab receives another \$100 million of VC funding and is valued at \$1 billion.
2019	The GitLab company employs over 600 employees.

The initial idea of GitLab was to earn money from open source technology by offering support services. However, what happened was that companies started to bring in consultants only to upgrade GitLab, and then they would stop the service contract. It became clear that going for a 100% open source was not going to be competitive. Instead of this, therefore, they chose **open core**. Under open core, a company releases a core software system under an open source license. A different version of the software is sold under a commercial license and contains more features.

So, GitLab was split up into two editions: an open source version, and an enterprise version.

## Exploring GitLab editions – CE and EE

The core of the GitLab software is called the CE. It is distributed under the MIT license, which is a permissive free software license created at the Massachusetts Institute of Technology. You are allowed to modify the software and use it in your creations.

No feature that ever made it to CE will ever be removed, or moved to a closed source version. When GitLab EE was created in 2013, it was, at its core, GitLab CE, but it had additional enterprise features, such as **Lightweight Directory Access Protocol (LDAP)** groups. Those features are not open source, per se, but can be added to the core version if they are perceived by the company as a core feature. The idea was that companies should also contribute as much as possible to solving problems and creating new features.

In 2016, the GitLab EE product was divided into three tiers: Starter, Premium, and Ultimate. Each tier is about five times more expensive than the previous one and contains more features and support options, as mentioned in the following table:

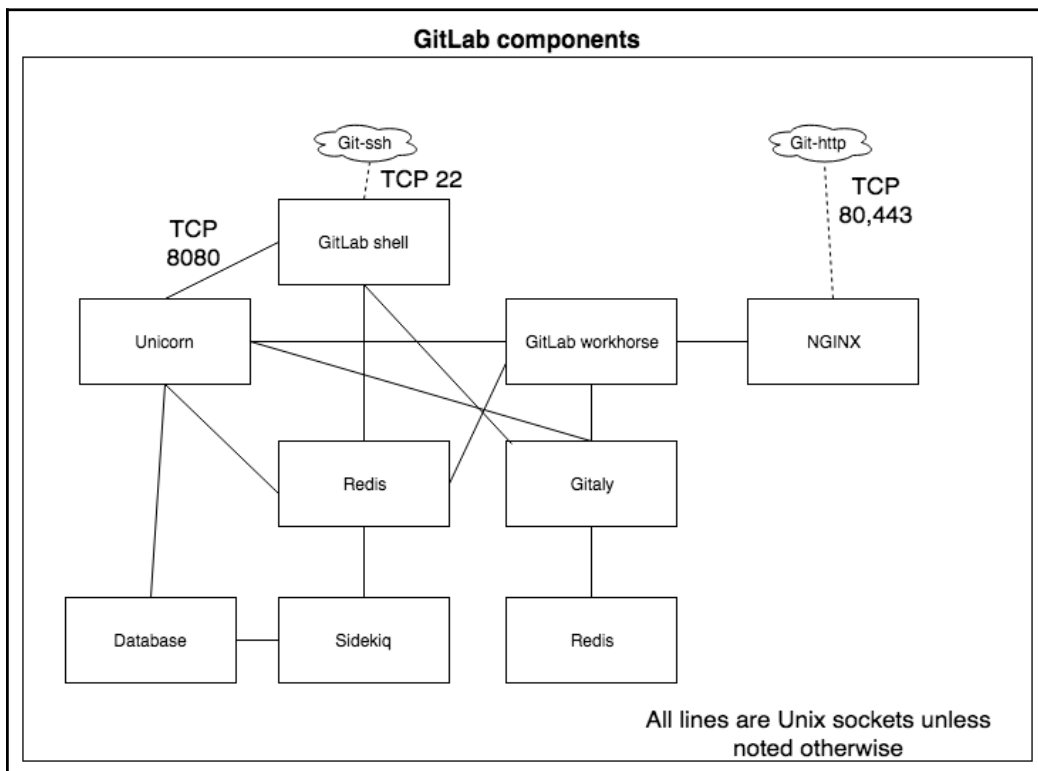
Version	Features (short list)
Starter	Everything on core GitLab CE: <ul style="list-style-type: none"> <li>• CI/CD</li> <li>• Project Issue Board</li> <li>• Mattermost integrations</li> <li>• Time tracking</li> <li>• GitLab pages</li> </ul>
Premium	More enterprise features such as the following: <ul style="list-style-type: none"> <li>• Maven and NPM repository functionality</li> <li>• Protected environments</li> <li>• Burndown charts</li> <li>• Multiple LDAP servers and Active Directory support</li> </ul>

Ultimate	All options, including the following: <ul style="list-style-type: none"> <li>• All security scanning tools</li> <li>• Epics</li> <li>• Free guest users</li> <li>• Web terminal for the web IDE</li> </ul>
----------	--

GitLab has a lot of features, but let's concentrate first on the basic building blocks.

## The core system components of GitLab

GitLab is not a monolithic application. It tries to follow the Unix philosophy, which means that a software module should do only one particular thing, and do it well. The components that GitLab is made of are not as small and elegant as Unix's `awk` and `sed`, but each component has a single purpose. You can find a high-level overview of these components in the following diagram:



Gitlab started as a pure Ruby on Rails application, but some components were later redesigned using Go. Ruby on Rails is a development framework built on top of the Ruby programming language. It implements a model-view-controller pattern and offers methods to connect to different databases (for example, ActiveRecord). It values convention over configuration and **don't-repeat-yourself (DRY)** programming. It is very well suited to rapid development, and at the same time, it is highly performant and has many features.

Let's dive a little deeper into those components in order to understand their roles.

## NGINX

The Unicorn web component cannot be used directly as it does not offer all the features for handling clients. The reverse proxy that is bundled by default is NGINX. It is also possible to use Apache as a frontend for GitLab, but it is preferable to use NGINX. There are many web servers available that could be installed in front of Unicorn, but in the end, there are basically two types, which are as follows:

- Process-based (forking or threaded)
- Asynchronous

NGINX and lighttpd are probably the two most-well known asynchronous servers. Apache is without a doubt the de facto standard process-based server. The biggest difference between the two types is how they handle scalability. For a process-based server, any new connections require a thread, while an event-driven, asynchronous server such as NGINX only needs a few threads (or, theoretically, only one). For lighter workloads, this does not matter much, but you will see a big difference when the number of connections grows, especially in terms of RAM. When serving tens of thousands of simultaneous connections, the amount of RAM used by NGINX would still hover around a couple of megabytes. Apache would either use hundreds, or it would not work at all. This is why NGINX is the better choice.

## Debugging NGINX

The first thing you will want to look at are the log files which by default are called `error.log` and `access.log`. In a GitLab environment installed from source these log files will typically reside in `/var/log/nginx/` and in a GitLab omnibus install in `/var/log/gitlab/nginx`.



Following is an example of the error log:

```
2019/09/08 20:45:14 [crit] 2387#2387: *95 connect() to
unix:/var/www/gitlab-app/tmp/sockets/unicorn.sock failed (2: No such file
or directory) while connecting to upstream, client: 127.0.0.1, server:
localhost, request: "GET /-/metrics HTTP/1.1", upstream:
"http://unix:/var/www/gitlab-app/tmp/sockets/unicorn.sock:/-/metrics",
host: "127.0.0.1:8080"
```

## Unicorn

Unicorn is an HTTP server for applications that deal with well-performing clients on connections that show low latency and have enough bandwidth. It takes advantage of features that are present in the core of Linux-like systems. It is called a **Rack HTTP server** because it implements HTTP for Rack applications. Rack, in turn, is actually a Ruby implementation of a minimal interface to deal with web requests, which you can use in your code.



You can find the project at <https://rack.github.io>.

Unicorn runs as a daemon server in Unix and is programmed in Ruby and the C programming language. Using Ruby means that it can also run a Ruby on Rails application such as GitLab.

## Debugging Unicorn

Maybe installing Unicorn produced errors, or you are experiencing bad performance that you suspect is caused by Unicorn not working properly.

There are several ways to find the cause. The log files can point you in the right direction.

### Timeouts in Unicorn logs

The following output is what a Unicorn worker timeout looks like in `unicorn_stderr.log`. This is not necessarily bad; it just means that a new worker is spawned:

```
[2015-06-05T10:58:08.660325 #56227] ERROR -- : worker=10 PID:53009 timeout
(61s > 60s), killing
```

```
[2015-06-05T10:58:08.699360 #56227] ERROR -- : reaped #<Process::Status:
pid 53009 SIGKILL (signal 9)> worker=10
[2015-06-05T10:58:08.708141 #62538] INFO -- : worker=10 spawned pid=62538
[2015-06-05T10:58:08.708824 #62538] INFO -- : worker=10 ready
```

It could be that there are just not enough Unicorn workers available to respond to the requests at hand. NGINX buffers a lot of requests so we must check on the handover socket whether Unicorn can keep up. To do this, a little nifty script is available here: <https://github.com/jahio/unicorn-status>.

It can be called with the following command:

```
$ ruby unicorn_status.rb /var/opt/gitlab/gitlab-rails/sockets/gitlab.socket
10
Running infinite loop. Use CTRL+C to exit.
-----
Active Requests Queued Requests
20 11
```

The first argument here is the `unicorn_status.rb` script, the second is the socket to connect to `./socket`, and the last argument is the poll interval (10).

## Unicorn processes disappear

On Linux, there is a mechanism called **Out-of-Memory (OOM) Killer** that will free up memory if the system is running low on memory, and you don't have any swap memory left. It might kill Unicorn if it is using too much memory.

Use `dmesg | egrep -i 'killed process'` to search for OOM events:

```
[102335.3134488] Killed process 5567 (ruby) total-vm:13423004kB, anon-
rss:554088kB
```

## Other kinds of errors or 100% CPU load

The ultimate way to debug Unicorn processes is to run `strace` on them:

1. Run `sudo gdb -p (PID)` to attach to the Unicorn process.
2. Run `call (void) rb_backtrace()` in the GDB console and find the generated Ruby backtrace in `/var/log/gitlab/unicorn/unicorn_stderr.log`:

```
from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/bundler-1.16.2/lib/bu
ndler/cli/exec.rb:28:in `run'
from
```

```
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/bundler-1.16.2/lib/bu
ndler/cli/exec.rb:74:in `kernel_load'
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/bundler-1.16.2/lib/bu
ndler/cli/exec.rb:74:in `load'
  from /opt/gitlab/embedded/bin/unicorn:23:in `>'<br/> from /opt/gitlab/embedded/bin/unicorn:23:in `load
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/unicorn-5.1.0/bin/uni
corn:126:in `>'
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/unicorn-5.1.0/lib/uni
corn/http_server.rb:132:in `start'
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/unicorn-5.1.0/lib/uni
corn/http_server.rb:508:in `spawn_missing_workers'
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/unicorn-5.1.0/lib/uni
corn/http_server.rb:678:in `worker_loop'
  from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/unicorn-5.1.0/lib/uni
corn/http_server.rb:678:in `select'
```

3. When you are done, leave GDB with `detach` and `q`.

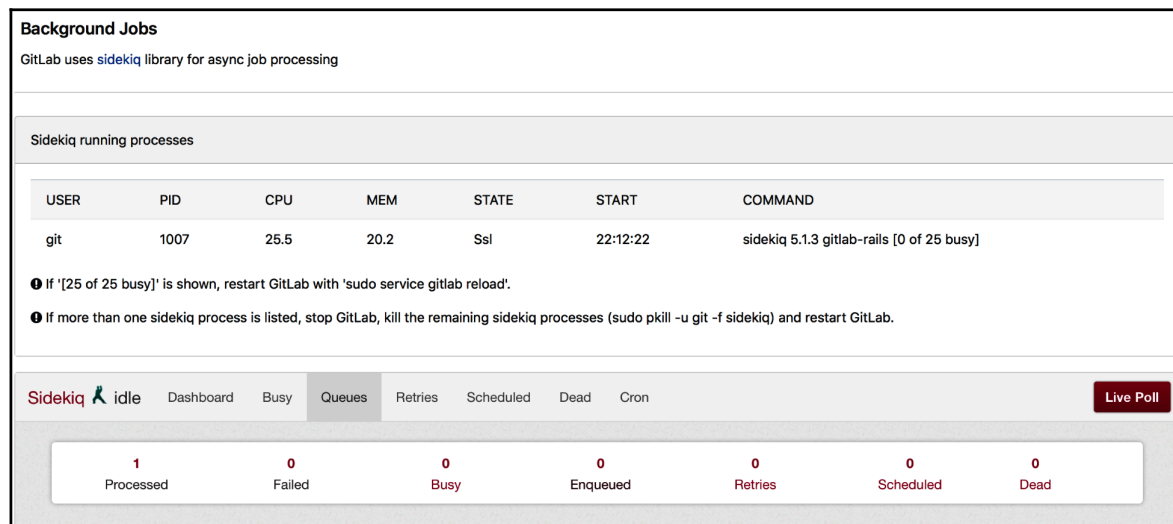
## Sidekiq

**Sidekiq** is a framework for background job processing. It allows you to scale your application by performing work in the background. For more information on Sidekiq, consult the following website: <https://github.com/mperham/sidekiq/wiki>.

Each Sidekiq server process pulls jobs from the queue in Redis and processes them. Like your web processes, Sidekiq boots Rails so that your jobs and workers have the full Rails API available for use, including ActiveRecord. The server will instantiate the worker and call perform with the given arguments. Everything else is up to your code.

## Debugging Sidekiq

As with Unicorn, there are several ways to debug Sidekiq processing. The easiest way is to log in to GitLab as an administrator and view the logs from there, and especially view the queues and jobs on the **Background Jobs** page, as shown in the following screenshot:



**Background Jobs**  
GitLab uses [sidekiq](#) library for async job processing

Sidekiq running processes

USER	PID	CPU	MEM	STATE	START	COMMAND
git	1007	25.5	20.2	Ssl	22:12:22	sidekiq 5.1.3 gitlab-rails [0 of 25 busy]

❗ If '[25 of 25 busy]' is shown, restart GitLab with 'sudo service gitlab reload'.

❗ If more than one sidekiq process is listed, stop GitLab, kill the remaining sidekiq processes (sudo pkill -u git -f sidekiq) and restart GitLab.

Sidekiq idle Dashboard Busy Queues Retries Scheduled Dead Cron Live Poll

1 Processed 0 Failed 0 Busy 0 Enqueued 0 Retries 0 Scheduled 0 Dead

Sometimes, you experience troubles and find situations on your Linux server.

### Sidekiq processes disappear

As mentioned before, in the Unicorn section, the OOM Killer might kill Sidekiq if it is using too much memory.

Use `dmesg | egrep -i 'killed process'` to search for OOM events:

```
[102335.3134488] Killed process 8887 (ruby) total-vm:13523004kB, anon-
rss:5540458kB
```

### A Sidekiq process is seemingly doing nothing

If Sidekiq isn't doing any work and it seems stuck most of the time, this means that the program is waiting for something. A common wait situation is when you are doing remote network calls. If you think this could be the case, you could make Sidekiq processes dump a backtrace to the log by sending it a TTIN signal.

This is what a Sidekiq worker looks like in the log file in

`/var/log/gitlab/sidekiq/current:`

```
{ "severity": "INFO", "time": "2019-06-23T19:00:14.493Z", "class": "RemoteMirrorNotificationWorker", "retry": 3, "queue": "remote_mirror_notification", "jid": "69eb806bfb66b82315bcb249", "created_at": "2019-06-23T19:00:14.461Z", "correlation_id": "toX0HnYW0s9", "enqueued_at": "2019-06-23T19:00:14.461Z", "pid": 471, "message": "RemoteMirrorNotificationWorker JID-69eb806bfb66b82315bcb249: done: 0.03 sec", "job_status": "done", "duration": 0.03, "completed_at": "2019-06-23T19:00:14.493Z" }
```

Since GitLab 12.0, the default output log format for Sidekiq is JSON, this makes it easier to read the log files into a tool like logstash because it is more structured.

## Other kind of errors or 100% CPU load

The ultimate way to debug Sidekiq processes is to make it dump a backtrace via GDB:

1. Run `sudo gdb -p (PID)` to attach to the Sidekiq worker process.
2. Run `call (void) rb_backtrace()` in the GDB console and find the generated Ruby backtrace in `/var/log/gitlab/sidekiq/current:`

```
2018-09-21_19:55:03.48430 from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/redis-3.3.5/lib/redis/connection/ruby.rb:83:in `_read_from_socket'
2018-09-21_19:55:03.48431 from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/redis-3.3.5/lib/redis/connection/ruby.rb:87:in `rescue in _read_from_socket'
2018-09-21_19:55:03.48432 from
/opt/gitlab/embedded/lib/ruby/gems/2.4.0/gems/redis-3.3.5/lib/redis/connection/ruby.rb:87:in `select'
```

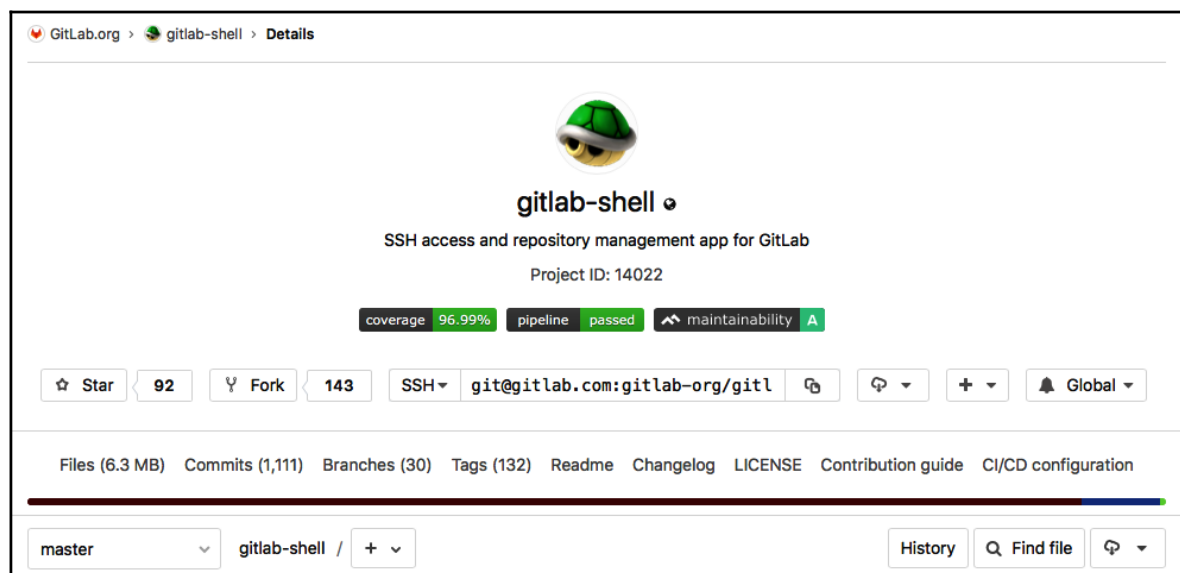
3. It is very hard to read backtraces, but this process was doing network operations while being traced, we can see a (`_read_from _socket`). You can read the source code to check what it is doing (there are line numbers mentioned).
4. When you are done, leave GDB with `detach` and `quit`.

You can also use other tracing tools to examine the behavior of the looping process. On Linux, for instance, `strace -p <pid>` allows you to view the system calls that are being made by the process.

## GitLab Shell

This component is used to provide access to Git repositories through SSH. In fact, for pushes via the `git-http` protocol, it is also called instead of the Rails app. It's essentially a small Ruby wrapper around the Git client. Git, through SSH, uses predefined commands that can be executed on the GitLab server. For authorization, it makes calls to the GitLab API. Before GitLab 5.0, this functionality was delivered by Gitolite and powered by the Perl programming language.

The source code of this project can be found here: <https://gitlab.com/gitlab-org/gitlab-shell>. You can see the following page:



You can install it locally, but it's really only useful when deployed together with other GitLab components. When you have that installed (see [Chapter 2, Installing GitLab](#), for instructions on how), the next section describes a way to debug when you have problems.

## Debugging GitLab Shell

In an omnibus installation, the log file for GitLab Shell can be found in the following location:

```
/var/log/gitlab/gitlab-shell/gitlab-shell.log
```

Alternatively, it may be found in the following location, for installations from source:

```
/home/git/gitlab-shell/gitlab-shell.log
```

What you will generally find are log lines that concern the basic operations of GitLab Shell:

- Git commands (such as `git push` and `git pull`).
- Authorization calls to the GitLab Rails API to check whether you are allowed to connect
- Execution of pre-receive hooks
- Actions requested
- Post-receive actions
- Any custom post-receive actions

Here, we have listed some lines from the log file:

```
bash-4.1$ tail gitlab-shell.log
time="2018-09-26T08:59:53+02:00" level=info msg="executing git command"
command="gitaly-upload-pack unix:/var/opt/gitlab/gitaly/gitaly.socket
{"repository":{"storage_name":"default","relative_path":"xxx/xxx.git"},
"git_object_directory":"","git_alternate_object_directories":[],
"gl_repository":"xxx"},"gl_repository":"project-xx",
"gl_id":"key-xx","gl_username":"xxxxxx"}" pid=18855 user="user
with key key-xx"

time="2018-09-26T08:59:53+02:00" level=info msg="finished HTTP request"
duration=0.228132057
method=POST pid=18890 url="http://127.0.0.1:8080/api/v4/internal/allowed"

time="2018-09-26T08:59:54+02:00" level=info msg="finished HTTP request"
duration=0.030036933 method=POST pid=18890
url="http://127.0.0.1:8080/api/v4/internal/pre_receive"

time="2018-09-26T08:59:54+02:00" level=info msg="finished HTTP request"
duration=0.094035804 method=POST pid=18979
url="http://127.0.0.1:8080/api/v4/internal/post_receive"
```

One way to find errors is to look for certain patterns, such as `failed`, as follows. This particular error points to a 500 error from Unicorn while checking whether a user has the right authorization to make a call to the GitLab API.

This error should show up in the Unicorn logs (`production.log`) if you search for an HTTP 500 error:

```
bash-4.1$ grep -i failed gitlab-shell.log
time="2018-09-26T08:05:52+02:00" level=error msg="API call failed"
body="{\"message\": \"500 Internal Server Error\"}" code=500 method=POST
pid=1587 url="http://127.0.0.1:8080/api/v4/internal/allowed"
time="2018-09-26T08:45:13+02:00" level=error msg="API call failed"
body="{\"message\": \"500 Internal Server Error\"}" code=500 method=POST
pid=24813 url="http://127.0.0.1:8080/api/v4/internal/allowed"
```

## Redis

Redis is a caching tool and HTTP session store that allows you to save cached data and session information from your website to an external location. This means that your website doesn't have to calculate everything every time; instead, it can retrieve the data from the cache and load the website much faster. The user sessions are in memory even if the application goes down. Redis is a fast caching tool because it uses memory first. It has several useful advantages:

- Everything is stored in one place, so you only have to flush one cache.
- It is faster than Memcache. This is noticeable when using the websites of large shops.
- Sessions are stored in memory and not in the database.
- The backend becomes faster.

Redis is not merely a cache, but is also a data structure store. It is basically a database and should be viewed conceptually as such. With regard to its operation and how it handles data, it has more in common with a NoSQL database.

## Basic data operations in Redis

We can discover some of the basics of Redis by playing with the data structures. You can install Redis using instructions found at <https://github.com/PacktPublishing/Mastering-GitLab-12/tree/master/Chapter01/InstallingRedis.md>.

Start the `redis-cli` command-line utility, and it will connect to the local Redis server:

```
$redis-cli
127.0.0.1:6379>
```



It is not fair to view Redis as a simple hash database with key values. But still, the five data structures that are provided do actually consist of a key and a value. Let's sum up the five data structures:

- **String:** You can use the `set` command to write a value to Redis. In the case of a simple string, you can simply save the value in the datastore shown as follows. After setting the string value, you can retrieve the value again by issuing the `get` command:

```
$ redis-cli
127.0.0.1:6379> set mykind "Human"
OK
127.0.0.1:6379> get mykind
"Human"
127.0.0.1:6379>
```

- **Hash:** In the same way as the string, you can set an arbitrary number of values to a key. Generally speaking, Redis treats values as a byte array and doesn't care what they are. This makes Redis very handy for representing objects. Again, with the `get` command, you can retrieve the values. GitLab uses this type to store web session information from users:

```
$ redis-cli
127.0.0.1:6379> set programs:tron '{"name": "tron", "kind":
"program"}'
OK
127.0.0.1:6379> get programs:tron
"{\"name\": \"tron\", \"kind\": \"program\"}"
```

- **List:** The list type in Redis is implemented as a linked list. You can add items to the list quite quickly with `rpush` (right push, to the tail of the list) or `lpush` (left push, to the head of the list). On the other hand, accessing an item by index is not that fast because it has to search the linked list. Still, for a queue mechanism, this is a good solution.

```
$ redis-cli
127.0.0.1:6379> rpush specieslist human computer cyborg
(integer) 3
127.0.0.1:6379> rpop specieslist
"cyborg"
127.0.0.1:6379> rpop specieslist
"computer"
127.0.0.1:6379> rpop specieslist
"human"
127.0.0.1:6379> rpop specieslist
(nil)
```

- **Sets:** Another datatype is the set. You add members with the `sadd` command. Don't forget that these sets are unordered, so if you ask for the members with `smembers`, the order will mostly be different to how you entered it:

```
$ redis-cli
127.0.0.1:6379> sadd speciesset human computer cyborg
(integer) 3
127.0.0.1:6379> smembers speciesset
1) "computer"
2) "human"
3) "cyborg"
```

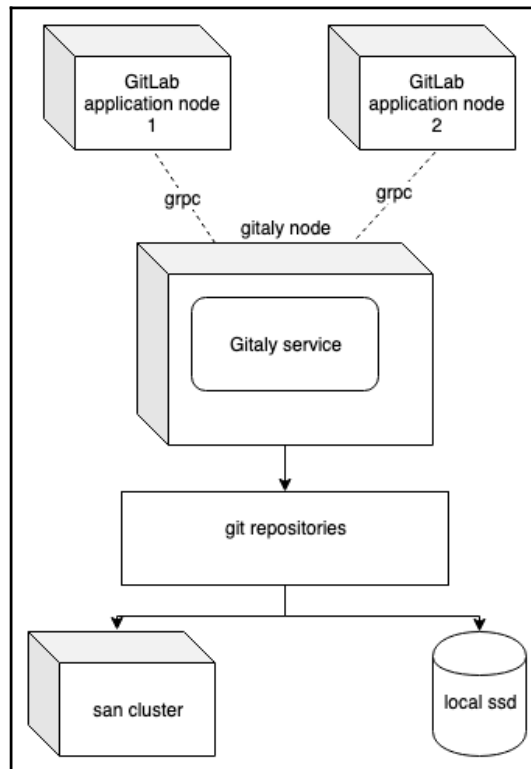
- **Sorted sets:** Fortunately, there is an ordered set as well. It is almost the same, but one difference is that you add a score to the entry, and that will automatically score the sort order, as you can see from the following:

```
127.0.0.1:6379> zadd speciessortedset 1 human
(integer) 1
127.0.0.1:6379> zadd speciessortedset 2 computer
(integer) 1
127.0.0.1:6379> zadd speciessortedset 3 cyborg
(integer) 1
127.0.0.1:6379> zrange speciessortedset 0 -1
1) "human"
2) "computer"
3) "cyborg"
```

## Gitaly

In the first versions of GitLab, all Git operations relied on using a local disk or network share. Gitaly is a project that tries to eliminate reliance on the **Network File System (NFS)**. Instead of calls to a filesystem service, Gitaly provides GitLab with a system based on **Remote Procedure Calls (RPCs)** to access Git repositories. It is written in Go and uses **gRPC Remote Procedure Call (gRPC)**, a cross-platform RPC framework from Google. It has been steadily developing since the beginning of 2017, and since GitLab 11.4, it can replace the need for a shared NFS filesystem.

You can find an overview of Gitaly and its place in the GitLab architecture in the following screenshot:



On a small installation, it runs in the same servers as all other components. In big clustered environments, you can set up dedicated Gitaly servers, which can be used by Gitaly clients such as the following:

- Unicorn
- Sidekiq
- `gitlab-workhorse`
- `gitlab-shell`
- Elasticsearch indexer
- Gitaly as a client

The source code of this project can be found here: <https://gitlab.com/gitlab-org/gitaly>.