



Community Experience Distilled

Nginx HTTP Server

Adopt Nginx for your web applications to make the most of your infrastructure and serve pages faster than ever

Clément Nedelcu

[PACKT] open source*
PUBLISHING community experience distilled

Nginx HTTP Server

Adopt Nginx for your web applications to make the most of your infrastructure and serve pages faster than ever

Clément Nedelcu

[PACKT] open source 
PUBLISHING community experience distilled

BIRMINGHAM - MUMBAI

Nginx HTTP Server

Copyright © 2010 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, and its dealers and distributors will be held liable for any damages caused or alleged to be 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.

First published: July 2010

Production Reference: 1140710

Published by Packt Publishing Ltd.
32 Lincoln Road
Olton
Birmingham, B27 6PA, UK.

ISBN 978-1-849510-86-8

www.packtpub.com

Cover Image by Vinayak Chittar (vinayak.chittar@gmail.com)

Credits

Author

Clément Nedelcu

Reviewers

Pascal Charest

Manlio Perillo

Acquisition Editor

Usha Iyer

Development Editor

Wilson D'souza

Technical Editor

Kartikey Pandey

Copy Editor

Leonard D'Silva

Indexers

Hemangini Bari

Tejal Daruwale

Editorial Team Leader

Aanchal Kumar

Project Team Leader

Lata Basantani

Project Coordinator

Jovita Pinto

Proofreader

Lynda Sliwoski

Graphics

Geetanjali Sawant

Production Coordinator

Aparna Bhagat

Cover Work

Aparna Bhagat

About the Author

Clément Nedelcu was born and raised in France, and studied in U.K., French, and Chinese universities. He is now a computer science teacher at Jiangsu University of Science and Technology in Zhenjiang, a southwestern city of China. He also works as technology consultant in France, specialized in web and Microsoft .NET development as well as Linux server administration. Since 2005, he has been administering a major network of websites in his spare time. This eventually led him to discover Nginx: it made such a difference that he started his own blog about it. One thing leading to another...

The author's blog can be visited at <http://cnelcu.net> and contains articles about Nginx and other web development topics.

I would like to express my gratitude to my girlfriend, my family and my friends who have been very supportive all along the writing stage. This book is dedicated to Martin Fjordvald for originally directing me to Nginx when my servers were about to kick the bucket. Special thanks to Maxim Dounin, Jérémie Bertrand, Shaun James, Zhang Yichun, Brendan, and all the folks on the #Nginx IRC channel on Freenode.

About the Reviewers

Pascal Charest works as senior principal consultant for Les Laboratoires Phoenix – an information system performance consulting firm based in Canada. Working with leading-edge algorithms and free software, he is called as subject matter expert to manage infrastructure projects, lead operations, and execute process validation.

Over the last year, sample mandates includes redesigning storage system (glusterfs) for a large North American investment group and managing the carrier-grade, international network of a prominent member of the telecommunication industry. He is also leading operations for quite a few local startups and answers their scalability needs through custom cloud computing solution / network infrastructure.

He is also a free software/society advocate and often speaks in conference about scalability issues in information systems.

He can be reached at pascal.charest@labsphoenix.com.

Thanks to Catherine, my love, for everything you've done so I did not have to do it.

Manlio Perillo lives in Italy, in the Irpinia region, near Naples.

He currently works as a freelance programmer, mainly developing web applications using Python and Nginx.

In 2008, he began working on a WSGI (Python Web Server Gateway Interface) implementation for Nginx. It is available on <http://bitbucket.org/mperillo/>, along with some other open source projects.

Table of Contents

Preface	1
Chapter 1: Preparing your Work Environment	7
Setting up a terminal emulator	7
Finding and downloading PuTTY	8
Creating a session	8
Working with PuTTY and the shell	10
Basic shell commands	11
File and directory management	11
User and group management	15
Superuser account	15
User accounts	15
Group management	17
Programs and processes	18
Starting an application	18
System services	19
Process management	20
Discovering the Linux filesystem	22
Directory structure	22
Special files and devices	25
Device types	25
Pseudo devices	26
Mounting a storage device	27
Files and inodes	28
EXT3 filesystem specifications	29
Filenames	29
Inodes	29
Atime, ctime, and mtime	30
Symbolic and hard links	31
File manipulation	32
Reading a file	33

Editing a file	34
Compression and archiving	35
System administration tools	37
Running a command as Superuser	37
Su command	37
Sudo command	38
System verification and maintenance	39
Disk Free	39
Disk Usage	39
Free memory	40
Software packages	40
Package managers	40
Downloading and installing packages manually	41
Building from source	42
Files and permissions	43
Understanding file permissions	43
Directory permissions	43
Octal representation	44
Changing permissions	44
Changing ownership and group	45
Summary	46
Chapter 2: Downloading and Installing Nginx	47
Setting up the prerequisites	47
GCC — GNU Compiler Collection	48
PCRE library	49
zlib library	50
OpenSSL	50
Downloading Nginx	51
Websites and resources	51
Version branches	52
Features	53
Downloading and extracting	54
Configure options	55
The easy way	55
Path options	56
Prerequisites options	58
Module options	59
Modules enabled by default	59
Modules disabled by default	60
Miscellaneous options	61
Configuration examples	62
About the prefix switch	63
Regular HTTP and HTTPS servers	63
All modules enabled	64

Mail server proxy	64
Build configuration issues	65
Make sure you installed the prerequisites	65
Directories exist and are writable	65
Compiling and installing	66
Controlling the Nginx service	67
Daemons and services	67
User and group	68
Nginx command-line switches	68
Starting and stopping the daemon	69
Testing the configuration	69
Other switches	70
Adding Nginx as a system service	71
System V scripts	71
What is an init script?	73
Creating an init script for Nginx	73
Installing the script	75
Debian-based distributions	76
Red Hat-based distributions	76
Summary	77
Chapter 3: Basic Nginx Configuration	79
<hr/>	
Configuration file syntax	79
Configuration Directives	80
Organization and inclusions	81
Directive blocks	83
Advanced language rules	84
Directives accept specific syntaxes	84
Diminutives in directive values	85
Variables	86
String values	86
Base module directives	86
What are base modules?	87
Nginx process architecture	87
Core module directives	88
Events module	93
Configuration module	95
A configuration for your profile	95
Understanding the default configuration	95
Necessary adjustments	96
Adapting to your hardware	97
Testing your server	99
Creating a test server	99

Performance tests	100
Httpperf	101
Autobench	102
OpenWebLoad	103
Upgrading Nginx gracefully	105
Summary	106
Chapter 4: HTTP Configuration	107
HTTP Core module	107
Structure blocks	108
Module directives	109
Socket and host configuration	110
Paths and documents	114
Client requests	117
MIME Types	121
Limits and restrictions	123
File processing and caching	125
Other directives	127
Module variables	130
Request headers	130
Response headers	131
Nginx generated	132
The Location block	133
Location modifier	133
Search order and priority	136
Case 1:	137
Case 2:	138
Case 3:	138
Summary	139
Chapter 5: Module Configuration	141
Rewrite module	141
Reminder on regular expressions	142
Purpose	142
PCRE syntax	142
Quantifiers	144
Captures	145
Internal requests	146
error_page	147
Rewrite	148
Infinite loops	149
Server Side Includes (SSI)	150
Conditional structure	151
Directives	153
Common rewrite rules	156

Performing a search	156
User profile page	156
Multiple parameters	156
Wikipedia-like	157
News website article	157
Discussion board	157
SSI module	157
Module directives and variables	158
SSI Commands	160
File includes	160
Working with variables	162
Conditional structure	163
Configuration	163
Additional modules	164
Website access and logging	164
Index	164
Autoindex	165
Random index	166
Log	166
Limits and restrictions	168
Auth_basic module	168
Access	168
Limit zone	169
Limit request	169
Content and encoding	170
Empty GIF	170
FLV	171
HTTP headers	171
Addition	172
Substitution	172
Gzip filter	173
Gzip static	175
Charset filter	175
Memcached	176
Image filter	178
XSLT	179
About your visitors	179
Browser	179
Map	180
Geo	180
GeoIP	181
UserID filter	181
Referer	182
Real IP	183
SSL and security	183
SSL	183
Setting up an SSL certificate	185
Secure link	186

Other miscellaneous modules	187
Stub status	187
Google-perftools	187
WebDAV	188
Third-party modules	189
Summary	190
Chapter 6: PHP and Python with Nginx	191
Introduction to FastCGI	192
Understanding the mechanism	192
Common Gateway Interface (CGI)	193
Fast Common Gateway Interface (FastCGI)	194
Main directives	195
FastCGI caching	201
Upstream blocks	204
Module syntax	205
Server directive	206
PHP with Nginx	207
Architecture	207
PHP-FPM	208
Setting up PHP and PHP-FPM	208
Downloading and extracting	208
Patching	209
Requirements	209
Building PHP	209
Post-install configuration	210
Running and controlling	210
Nginx configuration	211
Python and Nginx	212
Django	212
Setting up Python and Django	213
Python	213
Django	213
Starting the FastCGI process manager	214
Nginx configuration	215
Summary	215
Chapter 7: Apache and Nginx Together	217
Nginx as reverse proxy	217
Understanding the issue	218
The reverse proxy mechanism	219
Advantages and disadvantages	220
Nginx Proxy module	221
Main directives	222

Caching, buffering, and temporary files	225
Limits, timeouts, and errors	228
Other directives	229
Variables	230
Configuring Apache and Nginx	230
Reconfiguring Apache	231
Configuration overview	231
Resetting the port number	231
Accepting local requests only	232
Configuring Nginx	233
Enabling proxy options	233
Separating content	235
Advanced configuration	237
Additional steps	238
Forwarding the correct IP address	238
SSL issues and solutions	239
Server control panel issues	239
Summary	240
Chapter 8: From Apache to Nginx	241
Nginx versus Apache	241
Features	242
Core and functioning	242
General functionality	243
Flexibility and community	244
Performance	244
Usage	245
Conclusion	246
Porting your Apache configuration	246
Directives	246
Modules	249
Virtual hosts and configuration sections	250
Configuration sections	250
Creating a virtual host	251
htaccess files	254
Reminder on Apache .htaccess files	254
Nginx equivalence	255
Rewrite rules	257
General remarks	257
On the location	257
On the syntax	258
RewriteRule	259

WordPress	259
MediaWiki	261
vBulletin	262
Summary	263
Appendix A: Directive Index	265
<hr/>	
Appendix B: Module Reference	287
Access	287
Addition*	287
Auth_basic module	288
Autoindex	288
Browser	288
Charset	288
Core	289
DAV*	289
Empty GIF	289
Events	289
FastCGI	290
FLV*	290
Geo	290
Geo IP*	290
Google-perftools*	291
Gzip	291
Gzip Static*	291
Headers	291
HTTP Core	292
Image Filter*	292
Index	292
Limit Requests	292
Limit Zone	293
Log	293
Map	293
Memcached	293
Proxy	294
Random index*	294
Real IP*	294
Referer	294
Rewrite	295
Secure Link*	295
SSI	295
SSL*	295

Stub status*	296
Substitution*	296
Upstream	296
User ID	296
XSLT*	297
Appendix C: Troubleshooting	299
<hr/>	
General tips on troubleshooting	299
Checking access permissions	299
Testing your configuration	300
Have you reloaded the service?	300
Checking logs	300
Install issues	301
403 Forbidden custom error page	301
Location block priorities	302
If block issues	303
Inefficient statements	303
Unexpected behavior	304
Index	305

Preface

It is a well-known fact that the market of web servers has a long-established leader: Apache. According to recent surveys, as of October 2009 over 45 percent of the World Wide Web is served by this fifteen years old open source application. However, for the past few months the same reports reveal the rise of a new competitor: Nginx, a lightweight HTTP server originating from Russia – pronounced "engine X". There have been many interrogations surrounding the pronounced newborn. Why has the blogosphere become so effervescent about it? What is the reason causing so many server administrators to switch to Nginx since the beginning of year 2009? Is this apparently tiny piece of software mature enough to run my high-traffic website?

To begin with, Nginx is not as young as one might think. Originally started in 2002, the project was first carried out by a standalone developer, Igor Sysoev, for the needs of an extremely high-traffic Russian website, namely Rambler, which received as of September 2008 over 500 million HTTP requests per day. The application is now used to serve some of the most popular websites on the Web such as WordPress, Hulu, SourceForge, and many more. Nginx has proven to be a very efficient, lightweight yet powerful web server. Along the chapters of this book, you will discover the many features of Nginx and progressively understand why so many administrators have decided to place their trust in this new HTTP server, often at the expense of Apache.

There are many aspects in which Nginx is more efficient than its competitors. First and foremost, speed. Making use of asynchronous sockets, Nginx does not spawn as many times as it receives requests. One process per core suffices to handle thousands of connections, allowing for a much lighter CPU load and memory consumption. Secondly, ease of use – configuration files are much simpler to read and tweak than with other web server solutions such as Apache. A couple of lines are enough to set up a complete virtual host configuration. Last but not least, modularity. Not only is Nginx a completely open source project released under a BSD-like license, but it also comes with a powerful plug-in system – referred to as "modules". A large variety of modules are included with the original distribution archive, and many third-party ones can be downloaded online. All in all, Nginx combines speed, efficiency, and power, providing you the perfect ingredients for a successful web server; it appears to be the best Apache alternative as of today.

Although Nginx is available for Windows since version 0.7.52, it is common knowledge that Linux distributions are preferred for hosting production sites. During the various processes described in this book, we will thus assume that you are hosting your website on a Linux operating system such as Debian, Fedora, CentOS, Mandriva, or other well-known distributions.

What this book covers

Chapter 1, Preparing your Work Environment provides a basic approach of the Linux command-line environment that we will be using throughout this book.

Chapter 2, Downloading and Installing Nginx guides you through the setup process, by downloading and installing Nginx as well as its prerequisites.

Chapter 3, Basic Nginx Configuration helps you discover the fundamentals of Nginx configuration and set up the Core module.

Chapter 4, HTTP Configuration details the HTTP Core module which contains most of the major configuration sections and directives.

Chapter 5, Module Configuration helps you discover the many first-party modules of Nginx among which are the Rewrite and the SSI modules.

Chapter 6, PHP and Python with Nginx explains how to set up PHP and other third-party applications (if you are interested in serving dynamic websites) to work together with Nginx via FastCGI.

Chapter 7, Apache and Nginx Together teaches you to set up Nginx as reverse proxy server working together with Apache.

Chapter 8, From Apache to Nginx provides a detailed guide to switching from Apache to Nginx.

Appendix A, Directive Index lists and describes all configuration directives, sorted alphabetically. Module directives are also described in their respective chapters too.

Appendix B, Module reference lists available modules.

Appendix C, Troubleshooting discusses the most common issues that administrators face when they configure Nginx.

What you need for this book

Nginx is free and open source software running under various operating systems—Linux-based, Mac OS, Windows operating systems, and many more. As such, there is no real requirement in terms of software. Nevertheless in this book and particularly in the first two chapters we will be working in a Linux environment, so running a Linux-based operating system would be a plus. Prerequisites for compiling the application are further detailed in Chapter 2.

Who this book is for

This book is a perfect companion for both Nginx beginners and experienced administrators. For the former, it will take you through the complete process of setting up this lightweight HTTP server on your system and configuring its various modules to get it to do exactly what you need, in a fast and secure way. For the latter, it provides different angles of approach that can help you make the most of your current infrastructure. As the book progresses, it provides a complete reference to all the modules and directives of Nginx. It will explain how to replace your existing server with Nginx or configure Nginx to work as a frontend for your existing server.

Conventions

In this book, you will find a number of styles of text that distinguish between different kinds of information. Here are some examples of these styles, and an explanation of their meaning.

Code words in text are shown as follows: "We can include other contexts through the use of the `include` directive."

A block of code is set as follows:

```
[default]
exten => s,1,Dial(Zap/1|30)
exten => s,2,VoiceMail(u100)
exten => s,102,VoiceMail(b100)
exten => i,1,VoiceMail(s0)
```