Hands-On Enterprise Automation on Linux

Efficiently perform large-scale Linux infrastructure automation with Ansible

Packt
www.packt.com

James Freeman

Hands-On Enterprise Automation on Linux

Efficiently perform large-scale Linux infrastructure automation with Ansible

James Freeman

BIRMINGHAM - MUMBAI

Hands-On Enterprise Automation on Linux

Commissioning Editor: Vijin Boricha
Acquisition Editor: Rohit Rajkumar
Content Development Editor: Alokita Amanna
Senior Editor: Rahul Dsouza
Technical Editor: Prachi Sawant
Copy Editor: Safis Editing
Project Coordinator: Vaidehi Sawant
Proofreader: Safis Editing
Indexer: Pratik Shirodkar
Production Designer: Nilesh Mohite

First published: January 2020

Production reference: 1240120

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham
B3 2PB, UK.

ISBN 978-1-78913-161-1

www.packt.com

This book is dedicated to everyone who has inspired me to follow my dreams, my passions, and live my truth, especially Lyndon Rees, Eleonora Guantini, Elane Slade, and the late Sirdar Khan.

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 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 for more details.

At 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.

Foreword

Few would disagree when I say that the world of technology has grown ever more complex over the last couple of decades since the internet came to prominence. More and more products have arrived, promising us solutions to tame the growing complexity. Along with the promises come a raft of experts, there to help us through what is actually yet more complexity.

2012 saw the first release of Ansible. By 2013, it was gaining significant traction since its promise of power through simplicity was not an empty one. Here was a technology rooted in a simple truth—solving problems with technology really means solving problems for people. Therefore, people matter. A tool that is easy to pick up and learn? What an amazing thought! Early adopters were those who saw through the functionality list to realize that here was a people-pleasing game changer.

I first met James at one of his technical Ansible talks a few years ago. It was still relatively early days for Ansible, although we'd just been acquired by Red Hat. At that first meeting, I realized that here was a fellow who understood the link between people and Ansible's powerful simplicity. I've been lucky enough to see James speak on a number of occasions since, with two standout talks coming to mind.

At AnsibleFest 2018 in Austin, Texas, James gave a great talk about a client engagement where he presided over a business-critical database upgrade—on a Friday afternoon. What's the golden rule we all tout in tech? Don't make business-critical changes on a Friday! Yet James's charismatic storytelling had the audience enthralled. The second occasion was more recent, at an Ansible London meetup. Taking a very different approach to the usual tech-heavy talks, James presented the audience with a tale of positive psychology, a story that had Ansible as the underlying tool supporting people. It turned out to be a great success, sparking a lively interaction across the audience during the Q&A session that followed.

Scalability isn't just about a technology; it is about people. If you want a technology to scale, it must be easy for people to adopt, to master, and to share. James is a model of scalability himself, as he so readily shares his knowledge. He also shows in this book that Ansible is an orchestrator, a conductor of the symphony if you like, with the ability to span an enterprise. I'm sure you'll enjoy reading it as much as I've enjoyed every interaction I've had with James.

Mark Phillips
Product Marketing Manager, Red Hat Ansible

I've worked alongside James for several years and consider him to be one of the foremost Ansible experts in the world. I've been witness to his help in the digital modernization efforts of large and small organizations with the help of automation and DevOps practices.

In *Hands-On Enterprise Automation on Linux*, James generously shares his experience with a practical, no-nonsense approach to managing heterogeneous Linux environments. If you learn best through a hands-on approach, then this is the book for you. James provides plenty of in-depth examples in each chapter so that you can cement your understanding and feel prepared to take Ansible into a live environment.

Ready to become an automation rockstar and revolutionize your IT ops team? Then read on!

Ben Strauss
Security Automation Manager, MindPoint Group

Contributors

About the author

James Freeman is an accomplished IT consultant and architect with over 20 years' experience in the technology industry. He has more than 7 years of first-hand experience of solving real-world enterprise problems in production environments using Ansible, frequently introducing Ansible as a new technology to businesses and CTOs for the first time. He has a passion for positive psychology and its application in the world of technology and, in addition, has authored and facilitated bespoke Ansible workshops and training sessions, and has presented at both international conferences and meetups on Ansible.

About the reviewers

Gareth Coffey is an automation consultant for Cachesure, based in London, developing bespoke solutions to enable companies to migrate services to public and private cloud platforms. Gareth has been working with Unix/Linux-based systems for over 15 years. During that time, he has worked with a multitude of different programming languages, including C, PHP, Node.js, and various automation and orchestration tool sets. As well as consulting, Gareth runs his own start-up – Progressive Ops, developing cloud-based services aimed at helping start-up companies deploy resources across multiple cloud providers, with a focus on security.

Thanks to my wife and daughter for putting up with the late nights and early mornings.

Iain Grant is a senior engineer with over 20 years' experience as an IT professional, in both small and enterprise companies, where he has held a wide variety of positions, including trainer, programmer, firmware engineer, and system administrator. During this time, he has worked on multiple operating systems, ranging from OpenVMS, through Windows, to Linux, where he has also contributed to the Alpha Linux kernel. He currently works in an enterprise environment looking after over 300 Linux servers, with responsibility for their automation and management.

I would recommend this book as standard reading for any professional or senior engineer working with Linux. The areas covered provide you with excellent guidance and examples of a controlled build, as well as a managed and secure environment, resulting in an easier life for anyone looking after small or large Linux estates.

Packt is searching for authors like you

If you're interested in becoming an author for Packt, please visit `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

Preface

Welcome to *Hands-On Enterprise Automation on Linux*, your guide to a collection of the most valuable processes, methodologies, and tools for streamlining and efficiently managing your Linux deployments at enterprise scale. This book will provide you with the knowledge and skills required to standardize your Linux estate and manage it at scale, using open source tools including Ansible, AWX (Ansible Tower), Pulp, Katello, and OpenSCAP. You will learn about the creation of standard operating environments, and how to define, document, manage, and maintain these standards using Ansible. In addition, you will acquire knowledge of security hardening standards, such as the CIS Benchmarks. Throughout the book, practical, hands-on examples will be provided for you to try for yourself, on which you can build your own code, and to demonstrate the principles being covered.

Who this book is for

This book is for anyone who has a Linux environment to design, implement, and care for. It is intended to appeal to a wide range of open source professionals, from infrastructure architects through to system administrators, including professionals up to C level. Proficiency in the implementation and maintenance of Linux servers and familiarity with the concepts involved in building, patching, and maintaining a Linux server infrastructure are assumed. Prior knowledge of Ansible and other automation tools is not essential but may be beneficial.

What this book covers

`Chapter 1`, *Building a Standard Operating Environment on Linux*, provides a detailed introduction to standardized operating environments, a core concept that will be referred to throughout this hands-on book, and which is essential understanding in order for you to embark on this journey.

`Chapter 2`, *Automating Your IT Infrastructure with Ansible*, provides a detailed, hands-on breakdown of an Ansible playbook, including inventories, roles, variables, and best practices for developing and maintaining playbooks; a crash course enabling you to learn just enough Ansible to begin your automation journey.

`Chapter 3`, *Streamlining Infrastructure Management with AWX*, explores, with the help of practical examples, the installation and utilization of AWX (also available as Ansible Tower) so as to build good business processes around your Ansible automation infrastructure.

`Chapter 4`, *Deployment Methodologies*, enables you to understand the various methods available in relation to large-scale deployments in Linux environments, and how to leverage these to the best advantage of the enterprise.

`Chapter 5`, *Using Ansible to Build Virtual Machine Templates for Deployment*, explores the best practices for deploying Linux by building virtual machine templates that will be deployed at scale on a hypervisor in a practical and hands-on manner.

`Chapter 6`, *Custom Builds with PXE Booting*, looks at the process of PXE booting for when the templated approach to server builds may not be possible (for example, where bare-metal servers are still being used), and how to script this to build standard server images over the network.

`Chapter 7`, *Configuration Management with Ansible*, provides practical examples of how to manage your build once it enters service, so as to ensure that consistency remains a byword without limiting innovation.

`Chapter 8`, *Enterprise Repository Management with Pulp*, looks at how to perform patching in a controlled manner to prevent inconsistencies re-entering even the most carefully standardized environment through the use of the Pulp tool.

`Chapter 9`, *Patching with Katello*, builds on our work involving the Pulp tool by introducing you to Katello, providing even more control over your repositories whilst providing a user-friendly graphical user interface.

`Chapter 10`, *Managing Users on Linux*, provides a detailed look at user account management using Ansible as the orchestration tool, along with the use of centralized authentication systems such as LDAP directories.

`Chapter 11`, *Database Management*, looks at how Ansible can be used both to automate deployments of databases, and to execute routine database management tasks, on Linux servers.

`Chapter 12`, *Performing Routine Maintenance with Ansible*, explores some of the more advanced on-going maintenance that Ansible can perform on a Linux server estate.

`Chapter 13`, *Using CIS Benchmarks*, provides an in-depth examination of the CIS server hardening benchmarks and how to apply them on Linux servers.

`Chapter 14`, *CIS Hardening with Ansible*, looks at how a security hardening policy can be rolled out across an entire estate of Linux servers in an efficient, reproducible manner with Ansible.

`Chapter 15`, *Auditing Security Policy with OpenSCAP*, provides a hands-on look at the installation and use of OpenSCAP to audit Linux servers for policy violations on an on-going basis, since security standards can be reversed by either malicious or otherwise well-meaning end users.

`Chapter 16`, *Tips and Tricks*, explores a number of tips and tricks to keep your Linux automation processes running smoothly in the face of the ever-changing demands of the enterprise.

To get the most out of this book

To follow the examples in this book, it is recommended that you have access to at least two Linux machines for testing on, though more may be preferable to develop the examples more fully. These can be either physical or virtual machines—all examples were developed on a set of Linux virtual machines, but should work just as well on physical ones. In `Chapter 5`, *Using Ansible to Build Virtual Machine Templates for Deployment*, we make use of nested virtualization on a KVM virtual machine to build a Linux image. The exact hardware requirements for this are listed at the beginning of this chapter. This will require either access to a physical machine with the appropriate CPU to run the examples on, or a hypervisor that supports nested virtualization (for example, VMware or Linux KVM).

Please be aware that some examples in this book could be disruptive to other services on your network; where there is such a risk, this is highlighted at the beginning of each chapter. I recommend you try out the examples in an isolated test network unless/until you are confident that they will not have any impact on your operations.

Although other Linux distributions are mentioned in the book, we focus on two key Linux distributions—CentOS 7.6 (though if you have access to it, you are welcome to use Red Hat Enterprise Linux 7.6, which should work just as well in most examples), and Ubuntu Server 18.04. All test machines were built from the official ISO images, using the minimal installation profile.

As such, where additional software is required, we take you through the steps needed to install it so that you can complete the examples. If you choose to complete all the examples, you will install software such as AWX, Pulp, Katello, and OpenSCAP. The only exception to this is FreeIPA, which is mentioned in `Chapter 10`, *Managing Users on Linux*. Installing a directory server for your enterprise is a huge topic that sadly requires more space than we have in this book—hence, you may wish to explore this topic independently.

The text assumes that you will run Ansible from one of your Linux test machines, but Ansible can actually be run on any machine with Python 2.7 or Python 3 (versions 3.5 and higher) installed (Windows is supported for the control machine, but only through a Linux distribution running in the **Windows Subsystem for Linux** (**WSL**) layer available on newer versions of Windows. Supported operating systems for Ansible include (but are not limited to) Red Hat, Debian, Ubuntu, CentOS, macOS, and FreeBSD.

This book uses the Ansible 2.8.x.x series release, although a few examples are specific to Ansible 2.9.x.x, which was released during the course of writing. Ansible installation instructions can be found at `https://docs.ansible.com/ansible/intro_installation.html`.

Download the example code files

You can download the example code files for this book from your account at `www.packt.com`. If you purchased this book elsewhere, you can visit `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`.
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/Hands-On-Enterprise-Automation-on-Linux`. 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!