Mastering KVM Virtualization

Second Edition

Design expert data center virtualization solutions with the power of Linux KVM

Vedran Dakic

Humble Devassy Chirammal

Prasad Mukhedkar

Anil Vettathu

Packt>

BIRMINGHAM—MUMBAI

Mastering KVM Virtualization
Second Edition

Commissioning Editor: Vijin Boricha
Acquisition Editor: Shrilekha Inani
Senior Editor: Arun Nadar
Content Development Editor: Nihar Kapadia
Technical Editor: Soham Amburle
Copy Editor: Safis Editing
Project Coordinator: Neil D'mello
Proofreader: Safis Editing
Indexer: Priyanka Dhadke
Production Designer: Aparna Bhagat

First published: June 2019
Second edition: October 2020

Production reference: 2250920

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

ISBN 978-1-83882-871-4
www.packt.com

25 years ago, a colleague suggested that I should write what he called "a Linux book". I liked the idea and I promised I would. Years rolled by, and here I am, a quarter of a century later, acting on a promise. As Steve Jobs once said, 'Ideas without action aren't ideas. They're regrets.'

To my family – my mother, father, and brother, for putting up with me over the course of the last 25 years – which led to writing this book. To my TA, Jasmin, for both helping me to improve and offering insights into various topics covered in this book.

To my son, Luka, for showing me how young people can be both talented and focused, especially when faced with problems that require innovative solutions.

To my partner, Sanja, for driving me on in everything that I do.

Here's to not having any regrets.

– Vedran Dakic

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 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.

Contributors

About the authors

Vedran Dakic has a master's in electrical engineering and computing and is an IT trainer, covering system administration, cloud, automatization, and orchestration courses. He is a certified Red Hat, VMware, and Microsoft trainer. He is currently employed as the head of department of operating systems at Algebra University College in Zagreb. As part of his job, he lectures in relation to 3- and 5-year study programs in systems engineering, programming, and multimedia tracks. He also does a lot of consulting and systems integration in relation to his clients' projects – something he has been doing for the past 20 years. His approach is simple – bring real-world experience to all of the courses that he is involved with as this will provide added value for his students and customers.

Humble Devassy Chirammal is a senior software engineer in the Storage Engineering team at Red Hat. He has more than 15 years of IT experience, and his area of expertise is in understanding the full stack in an ecosystem, with emphasis on architecting solutions based on demand. These days he primarily concentrates on Ceph and GlusterFS and its integration to container orchestrator systems like Kubernetes. He has hands-on experience of emerging technologies, such as IaaS and PaaS solutions in Cloud and Containers. In the past, he has worked on intrusion detection systems, Clustering solutions, and Virtualization. As an open source advocate, he is a core contributor to many open source projects like Kubernetes. He actively organizes meetups on Openshift/Kubernetes, Virtualization, GlusterFS, CentOS. His twitter handle is @hchiramm and his website is `https://www.humblec.com`.

This book is dedicated to the loving memory of my parents, C.O. Devassy and Elsy Devassy, whose steady, balanced, and loving guidance has given me the strength and determination to be the person I am today. I would like to thank my wife, Anitha, for standing beside me throughout my career, and for the effort she put into taking care of our son, Heaven, and our daughters, Hail Mariya and Hanna Mariya, while I was writing this book. I would like to thank my brothers, Sible and Fr. Able Chirammal, as well, without whose constant support this book would not have been possible.

Finally, a special thanks to Ulrich Obergfell for being an inspiration that helped me enrich my knowledge in Virtualization.

Prasad Mukhedkar is a specialist cloud solution architect at Red Hat India with over 10 years of experience in helping customers in their journey to Virtualization and Cloud adoption. He is a Red Hat Certified Architect and has extensive experience in designing and implementing high performing cloud infrastructure. His areas of expertise are Red Hat Enterprise Linux 7/8 performance tuning, KVM virtualization, Ansible Automation, and Red Hat OpenStack. He is a huge fan of the Linux "GNU screen" utility.

Anil Vettathu began his association with Linux while in college and began his career as a Linux System Administrator soon after. He is a generalist, with an interest in open source technologies. He has hands-on experience in designing and implementing large scale virtualization environments using open source technologies and has extensive knowledge in libvirt and KVM. These days he primarily works on Red Hat Enterprise Virtualization, containers, and real time performance tuning. Currently, he is working as a Technical Account Manager for Red Hat. His website is `http://anilv.in`.

I'd like to thank my wife, Chandni, for her unconditional support. She took on the pain of looking after our two naughtiest kids, while I enjoyed writing this book. I'd like like to thank my parents, Dr. Annieamma and Dr. George Vettathu, for their guidance and for pushing me hard to study something new. Finally, I would like to thank my sister, Dr. Wilma, for her guidance, and my brother, Vimal.

About the reviewer

Ranjith Rajaram is employed as a senior principle technical support engineer at a leading open source Enterprise Linux company. He began his career by providing support to web hosting companies and managing servers remotely. Ranjith has also provided technical support to end customers. Early in his career, he worked on Linux, Unix, and FreeBSD platforms.

For the past 15 years, he has been continuously learning something new. This is what he likes and admires about technical support. As a mark of respect to all his fellow technical support engineers, he has included "developing software is humane, but supporting it is divine" in his email signature.

At his current organization, he is involved in implementing, installing, and troubleshooting Linux environment networks. Aside from this, he is also an active contributor to the Linux container space (Docker, Podman), Kubernetes, and OpenShift.

Apart from this book, he has reviewed the first editions of *Mastering KVM Virtualization* and *Learning RHEL Networking*, both available from Packt.

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

Section 2: libvirt and ovirt for Virtual Machine Management

3

Installing KVM Hypervisor, libvirt, and oVirt

4

Libvirt Networking