



Virtual Routing in the Cloud

ciscopress.com

**Arvind Durai
Stephen Lynn
Amit Srivastava**

Virtual Routing in the Cloud

Arvind Durai, CCIE No. 7016

Stephen Lynn, CCIE No. 5507 & CCDE No. 20130056

Amit Srivastava

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

Virtual Routing in the Cloud

Arvind Durai, CCIE No. 7016

Stephen Lynn, CCIE No. 5507 & CCDE No. 20130056

Amit Srivastava

Copyright© 2016 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

1 16

Library of Congress Control Number: 2016934921

ISBN-13: 978-1-58714-494-3

ISBN-10: 1-58714-494-8

Warning and Disclaimer

This book is designed to provide information about CSR 1000V router and adoption of NFV technology in the cloud environment. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Publisher	Paul Boger
Associate Publisher	Dave Dusthimer
Business Operation Manager, Cisco Press	Jan Cornelssen
Executive Editor	Brett Bartow
Managing Editor	Sandra Schroeder
Development Editor	Ellie Bru
Senior Project Editor	Tonya Simpson
Copy Editor	Kitty Wilson
Technical Editor(s)	Matt Bollick, Ray Wong
Editorial Assistant	Vanessa Evans
Cover Designer	Mark Shirar
Composition	Mary Sudul
Indexer	Brad Herriman
Proofreader	The Wordsmithery LLC



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0812R)

About the Authors

Arvind Durai, CCIE No. 7016, is an advanced services principal architect for Cisco Systems. His primary responsibility in the past 17 years has been in supporting major Cisco customers in the enterprise sector, including financial, retail, manufacturing, e-commerce, state government, utility (smart grid networks), and health-care sectors. Some of his focuses have been on security, multicast, network virtualization, and data center, and he has authored several white papers and design guides on various technologies. He has also been involved in data center design for more than 10 years and has designed many enterprise private cloud data center environments.

Arvind maintains two CCIE certifications: Routing and Switching, and Security. He holds a Bachelor of Science degree in Electronics and Communication, a Master's degree in Electrical Engineering (MS), and a Master's degree in Business Administration (MBA). He is a coauthor of two Cisco press books, *Cisco Secure Firewall Services Module* and *TcL Scripting for Cisco IOS*.

He has coauthored IEEE WAN smart grid architecture and has been a panel member for IEEE publications. Arvind also has presented in many industry forums, such as IEEE and Cisco LIVE.

Stephen Lynn, CCIE No. 5507 (Routing & Switching/WAN/Security) and CCDE No. 20130056, is an architect at Cisco Systems in the U.S. federal area. He has been with Cisco for more than 16 years and is a subject matter expert on enterprise network architecture. His focus is on large-scale network designs, including campus, WAN, and data center. As a recognized expert within Cisco and in the industry, Stephen has been working on large-scale, complex wide-area network designs in an enterprise environment. Stephen's focus has been on architectural designs involving 1,000 nodes to more than 10,000 nodes, leveraging technologies such as DMVPN, GET VPN, and FlexVPN to provide transport encryption and network segmentation over IP transport such as MPLS/Ethernet. Other areas of focus include high availability and convergence, QoS, Performance Routing (PFR), and network virtualization.

Stephen is a well-known speaker who has presented at several conferences and seminars worldwide. He holds a Bachelor of Science in Electrical Engineering from the University of Virginia. Stephen is based out of the Cisco office in Washington, DC.

Amit Srivastava is a senior manager with Equinix, Inc. At Equinix his team is responsible for global network and product fulfillment for Equinix's Cloud Exchange platform. Amit formerly worked as a technical leader with Cisco Systems, Inc. He has developed, tested, and enhanced network software for nearly 14 years. Before joining Cisco, he held positions in software application development, management, and testing.

Amit was involved in developing embedded applications for mobile devices in his engagement with Hughes Networks prior to joining Cisco.

Amit has been involved in the development cycles of new operating systems such as IOS XR and IOS XE and delivering features such as MPLS-based Layer 2 and 3 VPNs and traffic engineering. With IOS XE, Amit has worked with platforms such as ASR 1000 and CSR 1000V right from their inception, delivering enterprise-level features like IPsec, NAT, firewalls, NetFlow, AVC, and QoS. Amit holds a Bachelor of Science degree in Electrical Engineering.

About the Technical Reviewers

Ray Wong is a technical marketing engineer (TME) for Cisco Systems. In his more than eight years with Cisco, he has worked in multiple roles, from system testing, to solution design and validation, to technical marketing. He was a major contributor in the Cisco Virtual Office (CVO) solution. Together with his TME role for Cisco Cloud Services Router (CSR 1000V), he is also a subject matter expert for IOS VPN, including DMVPN, GET VPN, and FlexVPN.

Ray holds a Bachelor of Science degree in Computer Science and Mathematics from the University of Wisconsin–Madison. He is also a frequent speaker at Cisco Live events.

Matt Bollick has worked in technical marketing at Cisco for the past 19 years, running an obstacle course of technologies, including SNA, ATM and Ethernet switching, service provider aggregation, metro Ethernet, network management, and enterprise branch architectures. He has also worked on a variety of products, including the Cisco 7500, 7200, LS1010, 8540, 7300, and Cisco 10K before finally settling down for the past several years as the platform architect for the ISR series of branch routers. In his spare time, Matt is an avid SCUBA diver in North Carolina.

Dedications

From Arvind:

I am thankful to God for everything. I would like to dedicate this book to my wife, Monica, and my son, Akhill, who have been extremely patient and supportive during my long working hours. I am grateful to my parents for their blessings and for providing me with strong values.

I would also like to thank my parents-in-law, brother and family, and brother-in-law and family for all their support and wishes.

From Stephen:

I would like to dedicate this book to my wonderful and beautiful wife, Angela, and to my two incredible children, Christina and Ethan. Without your love, sacrifice, and support, this book would not have been possible. Thanks for putting up with the late nights and weekends I had to spend behind the computer and on conference calls instead of playing games, building Legos, and doing other fun family activities.

From Amit:

I would like to dedicate this book to my wife, Reshma, my daughter, Aarushi, and my parents. Without their love and support, I would never have been able to work on this. I would also like to thank my parents-in-law and my entire extended family. Their love and support have always been unconditional.

Acknowledgments

Arvind Durai:

Thanks to my wife, Monica, for encouraging me to write my third book. She inspired me and helped keep my spirits up all the time and provided her thoughts in multiple sections of this book. Thank you!!!

It was great working with Amit and Stephen. Their excellent technical knowledge and passion for writing made this writing experience a pleasure. I am looking forward to more years of working together as colleagues and friends.

Stephen Lynn:

A debt of gratitude goes to my coauthors, Arvind and Amit. Your knowledge and dedication to this project are appreciated more than you will ever know.

Acknowledgements for this book wouldn't be complete without mentioning my wife, Angela, who has endured and supported me through all my endeavors.

Amit Srivastava:

Special thanks to Arvind and Stephen, from whom I learned a lot while writing this book. I look forward to their continued support.

Our Acknowledgement

Many people within Cisco have provided feedback and suggestions to make this a great book. Thanks to all who have helped in the process, especially Ray Blair and Matt Falkner, for providing insightful input during the proposal process. A special thank you goes to our technical editors, Ray Wong and Matt Bollick, for your technical accuracy and insight into the technologies. Special thanks to Dimitris Vlassopoulos for providing his NSO lab setup and sharing his insights!

A big thank you goes out to the production team for this book, Brett Bartow, Ellie Bru, and Tonya Simpson, who have been incredibly professional and a pleasure to work with, and for making this book possible.

Contents at a Glance

	Introduction	xv
Chapter 1	Introduction to Cloud	1
Chapter 2	Software Evolution of the CSR 1000	37
Chapter 3	Hypervisor Considerations for the CSR	59
Chapter 4	CSR 1000V Software Architecture	95
Chapter 5	CSR 1000V Deployment Scenarios	141
Chapter 6	CSR Cloud Deployment Scenarios	185
Chapter 7	CSR in the SDN Framework	223
Chapter 8	CSR 1000V Automation, Orchestration, and Troubleshooting	247
Appendix A	Sample Answer File for Packstack	293
	Index	319

Reader Services

Register your copy at www.ciscopress.com/title/9781587144943 for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9781587144943 and click Submit. Once the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Introduction

In today's business environment, enterprise customers are under more pressure than ever to innovate and adapt to new challenges and market conditions. Enterprises want to focus their investments on their core business while reducing IT spending.

The cloud offers enterprise customers many benefits, such as lower costs and flexibility. The cloud's elastic model enables a company to increase and decrease infrastructure capacity on demand. The usage-based model offered by the cloud helps governments and enterprises reduce costs while increasing business agility by moving applications to the cloud and consuming infrastructure resources from the cloud. This leads to enterprises looking at consuming network and IT services from the cloud rather than investing in in-house operations.

The enabling technology in unlocking the cloud is virtualization. Virtualization abstracts and isolates the computing hardware and underlying infrastructure into a logical resource pool, allowing key capabilities such as resource sharing, virtual machine (VM) isolation, and load balancing. These capabilities provide the fundamental building blocks for an agile and scalable cloud environment with rapid provisioning, workload sharing, and increased availability.

The surge in applications and IT service consumption moving to the cloud highlights the need for evolved technologies and network elements in the cloud that offer security and visibility to help businesses with performance and compliance verification. Evolved networks and network services enable the provider to offer cloud services with security, performance, and availability. The Cisco Cloud Services Router 1000V (CSR 1000V) is a fully virtualized software router that offers a platform for enterprises to extend the data center to the cloud and to enforce their policies in the cloud.

The Cisco CSR 1000V provides a transparent solution for extending IT services into provider-hosted clouds. The solution offers a rich set of features, including VPN, firewall, Network Address Translation (NAT), application visibility, and WAN optimization. These functions allow enterprise and cloud providers to build highly secure, scalable, and extensible cloud networks. In addition, the Cisco CSR 1000V supports a rich set of application programming interfaces (API), providing robust integration into software-defined networking (SDN) for automated provisioning of these networks and network services and allowing simplified management and orchestration, which help in driving down costs further.

Networks inherently carry vast amounts of information, including user locations, device capabilities, topologies, and end-to-end performance characteristics. When exposed appropriately through well-defined APIs, such information can be consumed by cloud applications to fine-tune and customize their efficient delivery. The future holds the promise of increasingly rich application–network interactions.

The primary objective of this book is to simplify design aspects and architectural details in a unified resource, augmenting Cisco's existing collection of installation and configuration guides for various cloud-related products and solutions. This book covers the key

virtualization technologies used in the cloud; it provides a concise, accessible presentation of cloud network services and the different types of operational environments in the cloud. Cloud networking service and delivery concepts are reinforced with illustrative examples; architecture of SDN orchestration and its connection to Cisco CSR 1000V network services are introduced and elaborated upon. In addition, the book reviews the building blocks of the CSR 1000V, covering its architecture and software design.

This book also explains network design and deployment scenarios for the Cisco CSR 1000V, which influence its pivotal role in the cloud environment. Furthermore, the book distills how intelligent networks help providers simplify cloud service management and reduce costs through efficient scaling and optimized capacity utilization. This book provides architectural knowledge that contextualizes the roles and capabilities of these advanced networks and network services, along with discussions of design factors essential for their insertion into cloud services:

- The book introduces the readers to the cloud and provides an overview to different types of cloud operational environments, including a prelude to the evolution of virtual routers.
- Virtualization is introduced as a pivotal technology in cloud adoption.
- The book covers the details of the operating systems and hypervisors on which virtual routers run. It provides details pertaining to the operational aspects of virtual routing.
- The reader is introduced to the architecture and software design of the Cisco CSR 1000V virtual router. The reader is subsequently introduced to a comprehensive set of APIs that can be leveraged by SDN.
- The book focuses on different designs and use cases and configuration examples for routing, secure extension of enterprises to the cloud, and VM mobility. It illustrates how the CSR 1000V addresses the challenges that an architect faces in migrating toward the cloud.
- This book covers the different management techniques available to simplify operational and monitoring aspects of cloud services.

Who Should Read This Book?

This book is targeted for a technical audience responsible for architecture, design, and deployment of data center and enterprise cloud services.

This book also caters to the next generation of cloud network operators to implement enterprise features in the cloud, leveraging the CSR 1000V.

After reading this book, you will have a better understanding of the following:

- Key virtualization concepts and cloud models
- CSR 1000V software architecture and design

- SDN and the CSR 1000V platform and API
- Simplification of data center multitenant design with the CSR 1000V
- Use cases for the CSR 1000V to simplify enterprise routing in the cloud
- Operational visibility, management, and control of an enterprise network in the cloud

How This Book Is Organized

This book is organized into the following chapters.

Chapter 1: Introduction to Cloud

This chapter introduces the concept of cloud computing. It describes the various cloud models available and how virtualization enables the present-day transition to the cloud. Multitenant data center designs are illustrated, and the concept of SDN is introduced here.

Chapter 2: Software Evolution of the CSR 1000

This chapter introduces the software evolution of the Cisco Cloud Services Router (CSR 1000V). It covers the infrastructure requirements and design considerations of a CSR 1000V, and it discusses the features that a CSR 1000V brings to the virtual routing realm.

Chapter 3: Hypervisor Considerations for the CSR

This chapter describes the different hypervisor technologies available on servers to manage the hardware resources for virtual machines. Hypervisor technology selection is an important consideration when deploying the CSR 1000V.

Chapter 4: CSR 1000V Software Architecture

This chapter describes the software design of the CSR 1000V. It details the control-plane and data-plane design of the CSR 1000V. It also describes licensing requirements, software implementation, and packet flow related to the CSR 1000V.

Chapter 5: CSR 1000V Deployment Scenarios

This chapter describes the common deployment scenarios for the CSR 1000V. It depicts these scenarios using configuration examples.

Chapter 6: CSR Cloud Deployment Scenarios

This chapter describes CSR 1000V deployments in the cloud and data center environments.

Chapter 7: CSR in the SDN Framework

This chapter describes SDN components. It also provides an overview of the CSR 1000V in the OpenStack framework. Case studies in this chapter aim to educate the reader on using the APIs for user-defined outcomes.

Chapter 8: CSR 1000V Automation, Orchestration, and Troubleshooting

This chapter provides an overview of CSR 1000V management tools for orchestration, monitoring, and troubleshooting. It also illustrates the operation workflow for deploying a CSR 1000V.

Contents

Introduction xv

Chapter 1 Introduction to Cloud 1

Evolution of the Data Center 1

Data Center Architecture Building Blocks 2

Introduction to Virtualization in the Data Center 4

Evolution of Virtualization 5

Conceptual Architecture of Virtualization 5

Types of Virtualization Technologies 6

Server Virtualization 6

Types of Server Virtualization 8

Storage Virtualization 9

Types of Storage Virtualization 11

Network Virtualization 12

Network Virtualization Evolution 13

Types of Network Virtualization 14

Service Virtualization 15

Introduction to the Multitenant Data Center 16

Introduction to Cloud Services 18

Infrastructure as a Service (IaaS) 18

Platform as a Service (PaaS) 19

Software as a Service (SaaS) 20

Cloud Deployment Models 20

Cloud Design Considerations 21

Domain 1: Infrastructure and Environmental 22

Domain 2: Abstraction and Virtualization 23

Domain 3: Automation and Orchestration 23

Domain 4: Customer Interface 24

Domains 5 and 6: Service Catalog and Financials 24

Domains 7 and 8: Platform and Application 24

Domain 9: Security and Compliance 24

Domain 10: Organization, Governance, and Process 25

Enterprise Connectivity to the Cloud 26

Internet for Transport 26

Direct Connectivity to a Cloud Provider 28

Enterprise Cloud Adoption Challenges 29

Software-Defined Networking	30
Open Networking Foundation	31
OpenDaylight Project	32
Network Function Virtualization	33
OpenStack	34
Summary	35

Chapter 2 Software Evolution of the CSR 1000 37

IOS Software Architecture	37
IOS XE Architecture	39
The IOS XE Kernel	40
The IOS Daemon	40
The Forwarding Manager	41
The Interface Manager	41
The Platform Manager	41
Cisco ASR 1000 System Architecture Overview	41
Route Processor	42
Embedded Service Processor	42
SPA Interface Processor	43
Cloud Service Router 1000V Overview	44
Deployment Requirements	45
Elastic Performance and Scaling	47
Rapid Deployment and Routing Flexibility in the Cloud	49
CSR 1000V Deployment Examples	50
Secure Cloud VPN Gateway	50
Network Extension from Premises to Cloud	51
Segmentation Within a Cloud	52
CSR 1000V Key Features	52
Summary	57

Chapter 3 Hypervisor Considerations for the CSR 59

Understanding Operating Systems	59
Operating System Design	60
<i>Physical Resource Management</i>	60
<i>Software Access to Physical Resources</i>	62
Kernels	63
<i>Microkernels</i>	63

<i>Hybrid Kernels</i>	64
<i>The Cisco IOS Kernel</i>	64
The Boot Process	66
Linux Memory Management	69
<i>Linux Swap Space and Memory Overcommit</i>	69
<i>Linux Caching</i>	71
Understanding Hypervisors	71
How Does a Hypervisor Compare to an Operating System?	72
Type 1 Hypervisor Design	74
<i>Monolithic Architecture</i>	74
<i>Microkernel Architecture</i>	74
<i>Core Partitioning</i>	75
ESXi Hypervisor	75
Architectural Components of ESXi	75
The VMkernel	75
<i>Components of the VMkernel</i>	76
Processes Running on the VMkernel	77
Device Drivers	78
File Systems	79
Management	80
KVM	82
Architectural Components of KVM/QEMU	84
Guest Emulator (QEMU)	85
Management Daemon (Libvirt)	88
<i>User Tools (virsh, virt-manager)</i>	89
Hyper-V	91
Xen	92
Summary	94
Chapter 4	CSR 1000V Software Architecture
	95
System Design	95
Virtualizing the ASR 1001 into the CSR 1000V	98
CSR 1000V Initialization Process	99
CSR 1000V Data Plane Architecture	100
CSR 1000V Software Crypto Engine	103
Life of a Packet on a CSR 1000V: The Data Plane	103
Netmap I/O	104

	Packet Flow	106
	<i>Device Initialization Flow</i>	106
	<i>TX Flow</i>	107
	<i>RX Flow</i>	108
	<i>Unicast Traffic Packet Flow</i>	109
	Installing the CSR 1000V on a VMware Hypervisor	110
	Bringing Up the VM with the CSR 1000V on ESXi	110
	Installing the CSR 1000V on a KVM Hypervisor	126
	Bring Up the CSR 1000V as a Guest	126
	Performance Tuning of the CSR 1000V	137
	Summary	139
Chapter 5	CSR 1000V Deployment Scenarios	141
	VPN Services	141
	Layer 2 VPNs	141
	Layer 3 VPNs	142
	<i>Site-to-Site VPNs</i>	143
	<i>Remote Access VPNs</i>	147
	Use Cases for the CSR 1000V as a VPN Service Gateway	148
	Enterprise Data Center Network Extension	148
	<i>The CSR 1000V as a VPN Gateway</i>	148
	CSR for Secure Inter-Cloud Connectivity	152
	Remote VPN Access into the Cloud	153
	BGP Route Reflector Use Case for the CSR	155
	The CSR 1000V in a Hierarchical Route Reflector Use Case	157
	Planning for Future Branch Design with the CSR 1000V	162
	Evolution of Branch Virtualization	164
	LISP and CSR	168
	LISP Terminology	169
	The LISP Data Plane	171
	The LISP Control Plane	171
	Typical LISP Use Cases	175
	<i>IP Mobility</i>	175
	<i>IPv6 Migration</i>	175
	<i>Network-to-Network Connectivity</i>	175
	<i>Network-to-Network Interconnection Topology and Configuration</i>	176
	Summary	183