Practical Networking

Jorge Crichigno
Elie Kfoury
Elias Bou-Harb
Nasir Ghani

# High-Speed Networks

## A Tutorial

Springer

# Practical Networking

**Series Editor**

Zhi-Li Zhang
Sci & Eng, 4-192 EE/CS Bldg
University of Minnesota, Dept of Comp
Minneapolis, MN, USA

Explosive growth in cloud and mobile computing coupled with new advances in systems and networking technologies as well as machine learning and artificial intelligence (AI) have revolutionized how networks and distributed systems are designed, developed, operated and managed. This is epitomized by data center networking where it has spurred a wholesale rethinking and re-designs from network architectures, to physical interconnects to routing, flow management and network application support. New networking paradigms and technologies such as software-defined networking, network function virtualization, smart NICs and software/hardware co-designs have emerged for better designing, operating, managing and evolving networks, and also enabled new visions such as "self-driving networks" and AIOps. The Practical Networking Series is centered on emerging topics in new networking paradigms, architectural designs, algorithms and mechanisms for primarily wired networks (from data center networks, enterprise networks to ISP networks), but also touches on "packet core networks" for emerging 5G and beyond cellular and wireless networks. Books in this series address these topics from both theoretical (e.g., new theoretical foundations, algorithms and performance analysis) and practical (e.g., new network mechanisms, protocols, APIs and standards, software frameworks) perspectives. Relatively short books on a timely and focused topic, research monographs, and textbooks are of interest. The Editor is seeking well written works by well-established researchers and practitioners in the networking field around the world, particularly Asia and North America.
Prospective Authors or Editors:
If you have an idea for a book, we would welcome the opportunity to review your proposal. Should you wish to discuss any potential project further or receive specific information regarding our book proposal requirements, please contact Zhi-Li Zhang or Susan Evans:

Zhi-Li Zhang
Department of Computer Science Department
University of Minnesota
4-192 Keller Hall, 200 Union Street SE
Minneapolis, MN 55455-0159
zhzhang@cs.umn.edu

Susan Evans
Senior Editor
Springer Nature
233 Springe Street
New York, NY 10013 USA
susan.evans@springernature.com

More information about this series at https://link.springer.com/bookseries/16325

Jorge Crichigno • Elie Kfoury • Elias Bou-Harb
Nasir Ghani

# High-Speed Networks

## A Tutorial

Springer

Jorge Crichigno
College of Engineering and Computing
University of South Carolina
Columbia, SC, USA

Elie Kfoury
College of Engineering and Computing
University of South Carolina
Columbia, SC, USA

Elias Bou-Harb
Cyber Center for Security and Analytics
University of Texas at San Antonio
San Antonio, TX, USA

Nasir Ghani
Electrical Engineering
University of South Florida
Tampa, FL, USA

ISSN 2662-1703 ISSN 2662-1711 (electronic)
Practical Networking
ISBN 978-3-030-88840-4 ISBN 978-3-030-88841-1 (eBook)
https://doi.org/10.1007/978-3-030-88841-1



This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

# Preface

Welcome to *High-Speed Networks*: A Tutorial. This book is the result of a journey of the authors who have been designing and deploying high-speed networks for several years, in particular Science Demilitarized Zones (Science DMZs). The Science DMZ is a high-speed network designed to facilitate the transfer of big science data.

As the popularity of high-speed networks and Science DMZs surged, the need for professionals with the skills to operate such infrastructures has increased. However, practitioners have been mostly trained to operate general-purpose networks, which have different characteristics from those of Science DMZs and other high-performance networks. At the time when the authors started designing and operating Science DMZs, the available material was limited to workshops organized by ESnet, the Scientific Networking Division at Lawrence Berkeley National Laboratory in the United States (U.S.).

This book tries to address the above gap. It provides practical knowledge and skills on Science DMZs and high-speed networks in general, which are reinforced with virtual laboratory experiments.

## Audience

This book is for industry professionals and for students in computer science, information technology, and similar programs, who are interested in learning fundamental concepts related to high-speed networks and corresponding implementations. The book assumes minimal familiarity with networking, typically covered in an introductory networking course. It is appropriate for an upper-level undergraduate course, for a first-year graduate course, and for self-pace learning by industry professionals.

## What is Unique About This Book?

The book delves into protocols and devices at different layers, from the physical infrastructure to application-layer tools and security appliances, that must be carefully considered for the optimal operation of Science DMZs and high-speed networks. In contrast to traditional books, the book is accompanied by hands-on virtual laboratory experiments that are conducted on a virtual platform.

## The Virtual Platform and Virtual Laboratory Experiments

The virtual platform enables learners to immediately deploy virtual networks composed of an equipment pod (routers, switches, servers, firewalls, etc.) needed for mastering a topic. Experiments help learners to reinforce concepts and to learn how to optimally configure and manage network devices, based on real measurements and observations. Access to the platform is available for a fee and includes all material required to conduct the experiments. The URL of the virtual platform is: http://highspeednetworks.net/

## Organization

The book follows a bottom-up approach. Chapter "Introduction to High-Speed Networks and Science DMZ" presents the motivation for Science DMZs and high-speed networks. Chapter "Network Cyberinfrastructure Aspects for Big Data Transfers" describes limitations of general-purpose networks when transferring large data sets across a Wide Area Network (WAN), and explores the cyberinfrastructure required to support such transfers. It also discusses different options a network may have to connect to other networks. Chapter "Data-Link and Network Layer Considerations for Large Data Transfers" describes attributes related to routers and switches, which have large impact on performance, including router's buffer size, maximum transmission unit (MTU), and others. Chapter "Impact of TCP on High-Speed Networks and Advances in Congestion Control Algorithms" discusses key features at the transport layer, such as TCP congestion control, pacing, and parallel connections. Chapter "Application and Security Aspects for Large Flows" presents application-layer tools used to support large data transfers. Chapter "Security Aspects" describes security challenges arising in Science DMZs and high-speed networks, and presents best practices. Chapter "Challenges and Open Research Issues" discusses challenges and open research issues.

## Relevance of Networking Tools

The book provides a set of virtual laboratory experiments at the end of most chapters. All equipment pods are implemented with appliances running real protocol stacks. Examples include iPerf3, the Network Emulator (NETEM), traffic control (tc), and Zeek intrusion detection system. Recognizing the impact of Mininet on networking, the authors decided to use this network emulator to create topologies for the laboratory experiments. All tools are based on open-source software, which reflects industry trends. Specifically, over the years, the authors observed that open-source software has been increasingly used to design, build, test, and control networks. For example, since the original publication of the paper describing Mininet in 2010, Mininet has gained wide adoption in the industry and academia. Mininet's paper has received the ACM Test of Time Award.

Columbia, SC, USA — Jorge Crichigno
Columbia, SC, USA — Elie Kfoury
San Antonio, TX, USA — Elias Bou-Harb
Tampa, FL, USA — Nasir Ghani

# Acknowledgement

The authors would like to express their gratitude to the U.S. National Science Foundation (NSF), Office of Advanced Cyberinfrastructure (OAC). This work would not be possible without NSF support. Part of the material was developed under the award numbers 1829698 and 1925484. The authors are also thankful to the Network Development Group (NDG) team who worked with Dr. Crichigno to deploy the virtual laboratory experiments. NDG's president, Richard Weeks, has constantly provided invaluable suggestions.

The first two authors would like to specially thank the members of the Cyberinfrastructure Laboratory at the University of South Carolina (USC) who helped create and test laboratory experiments, and the Department of Integrated Information Technology at USC for the conducive environment for writing this book.

# Contents