

Building the Network of the Future

Getting Smarter, Faster,
and More Flexible with
a Software Centric
Approach

John Donovan and Krish Prabhu

 **CRC Press**
Taylor & Francis Group
A CHAPMAN & HALL BOOK

Building the Network of the Future

**Getting Smarter, Faster, and More Flexible
with a Software Centric Approach**



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Building the Network of the Future

Getting Smarter, Faster, and More Flexible
with a Software Centric Approach

by

John Donovan and Krish Prabhu



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

A CHAPMAN & HALL BOOK

Trademarks or registered trademarks of others referenced in this book are used for informational purposes only.

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2017 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

Printed on acid-free paper

International Standard Book Number-13: 978-1-1386-3152-6 (Hardback)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged, please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Library of Congress Cataloging-in-Publication Data

Names: Donovan, John, 1960- author. | Prabhu, Krish, 1954- author.
Title: Building the network of the future : getting smarter, faster, and more flexible with a software centric approach / by John Donovan, Krish Prabhu.
Description: Boca Raton : Taylor & Francis, CRC Press, 2017. | Includes bibliographical references and index.
Identifiers: LCCN 2017009175 | ISBN 9781138631526 (hardback : alk. paper)
Subjects: LCSH: Computer networks--Technological innovations. | Cloud computing. | Virtual computer systems.
Classification: LCC TK5105.5 .B8454 2017 | DDC 004.6--dc23
LC record available at <https://lcn.loc.gov/2017009175>

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

Foreword	vii
Acknowledgments.....	ix
Authors.....	xi
Chapter 1 The Need for Change	1
<i>John Donovan and Krish Prabhu</i>	
Chapter 2 Transforming a Modern Telecom Network—From All-IP to Network Cloud	9
<i>Rich Bennett and Steven Nurenberg</i>	
Chapter 3 Network Functions Virtualization.....	25
<i>John Medamana and Tom Siracusa</i>	
Chapter 4 Network Functions Virtualization Infrastructure	49
<i>Greg Stiegler and John DeCastra</i>	
Chapter 5 Architecting the Network Cloud for High Availability.....	67
<i>Kathleen Meier-Hellstern, Kenichi Futamura, Carolyn Johnson, and Paul Reeser</i>	
Chapter 6 Software-Defined Networking	87
<i>Brian Freeman and Han Nguyen</i>	
Chapter 7 The Network Operating System: VNF Automation Platform.....	103
<i>Chris Rice and Andre Fuetsch</i>	
Chapter 8 Network Data and Optimization	137
<i>Mazin Gilbert and Mark Austin</i>	
Chapter 9 Network Security.....	171
<i>Rita Marty and Brian Rexroad</i>	
Chapter 10 Enterprise Networks.....	201
<i>Michael Satterlee and John Gibbons</i>	
Chapter 11 Network Access.....	223
<i>Hank Kafka</i>	

Chapter 12	Network Edge.....	249
	<i>Ken Duell and Chris Chase</i>	
Chapter 13	Network Core	265
	<i>John Paggi</i>	
Chapter 14	Service Platforms	293
	<i>Paul Greendyk, Anisa Parikh, and Satyendra Tripathi</i>	
Chapter 15	Network Operations	331
	<i>Irene Shannon and Jennifer Yates</i>	
Chapter 16	Network Measurements.....	353
	<i>Raj Savoor and Kathleen Meier-Hellstern</i>	
Chapter 17	The Shift to Software	389
	<i>Toby Ford</i>	
Chapter 18	What's Next?	401
	<i>Jenifer Robertson and Chris Parsons</i>	
Index		407

Foreword

It's easy to forget just how thoroughly mobile connectivity has changed our relationship with technology. Any new piece of hardware is judged not just for the processor, memory, or camera it contains but for how well it connects. Wi-Fi, fiber, 4G—and soon 5G—have become the new benchmarks by which we measure technological progress.

In many ways, this era of high-speed connectivity is a lot like the dawn of the railroads or the superhighways in the nineteenth and twentieth centuries, respectively. Like those massive infrastructure projects, connectivity is redefining society. And it is bringing capabilities and knowledge to all facets of society, enabling a new era of collaboration, innovation, and entertainment.

Yet most people take networks for granted because so much of what happens on them is hidden from view. When we use our phones or tablets or stream video to our connected cars, we expect the experience to be simple and seamless. The last thing we should care about is how they are connecting or onto what networks. And all the while we continue to place ever-heavier loads on these networks, which now include 4K video and virtual reality.

Keeping up with these demand curves is a monumental task. In fact, knowing how to manage, secure, transport, and analyze these oceans of data might be the single greatest logistical challenge of the twenty-first century.

Needless to say, this is a relatively new phenomenon. It used to be that the biggest test for any telephone network was the ability to handle the explosion of calls on Mother's Day. Now, any day can be the busiest for a network—when an entire season of a popular show is released or when the latest smartphone OS upgrade goes live. A network today must always be prepared.

This paradigm shift requires a new kind of network, one that can respond and adapt in near-real time to constant and unpredictable changes in demand while also being able to detect and deflect all manner of cyber threats.

The old network model completely breaks down in the face of this. That's why software-defined networking (SDN) and network function virtualization (NFV) are moving from concept to implementation so rapidly. This is the network of the future: open-sourced, future-proofed, highly secure, and flexible enough to scale up to meet any demand.

This book lays out much of what we have learned at AT&T about SDN and NFV. Some of the smartest network experts in the industry have drawn a map to help you navigate this journey. Their goal is not to predict the future but to help you design and build a network that will be ready for whatever that future holds. Because if there's one thing the last decade has taught us, it's that network demand will always exceed expectations. This book will help you get ready.

Randall Stephenson

Chairman and Chief Executive Officer, AT&T



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Acknowledgments

The authors would like to acknowledge the support from and the following contributions by Mark Austin, Rich Bennett, Chris Chase, John DeCastra, Ken Duell, Toby Ford, Brian Freeman, Andre Fuetsch, Kenichi Futamura, John Gibbons, Mazin Gilbert, Paul Greendyk, Carolyn Johnson, Hank Kafka, Rita Marty, John Medamana, Kathleen Meier-Hellstern, Han Nguyen, Steven Nurenberg, John Paggi, Anisa Parikh, Chris Parsons, Paul Reeser, Brian Rexroad, Chris Rice, Jenifer Robertson, Michael Satterlee, Raj Savor, Irene Shannon, Tom Siracusa, Greg Stiegler, Satyendra Tripathi, and Jennifer Yates.

The authors would also like to acknowledge the contributions of Tom Restaino and Gretchen Venditto in the editing of the book.

All royalties will be donated to the AT&T Foundation and will be used to support STEM education.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Authors

John Donovan is chief strategy officer and group president AT&T Technology and Operations (ATO). Donovan was executive vice president of product, sales, marketing, and operations at VeriSign Inc., a technology company that provides Internet infrastructure services. Prior to that, he was chairman and CEO of inCode Telecom Group Inc. Mr. Donovan was also a partner with Deloitte Consulting, where he was the Americas Industry Practice director for telecom.

Krish Prabhu retired as president of AT&T Labs and chief technology officer for AT&T, where he was responsible for the company's global technology direction, network architecture and evolution, and service and product design. Krish has an extensive background in technology innovation from his time at AT&T Bell Laboratories, Alcatel, Morgenthaler Ventures, and Tellabs. He also served as COO of Alcatel, and CEO of Tellabs.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

1 The Need for Change

John Donovan and Krish Prabhu

CONTENTS

References..... 8

Although the early history of electrical telecommunication systems dates back to the 1830s, the dawn of modern telecommunications is generally viewed as Alexander Graham Bell’s invention of the telephone in 1876. For the first 100 years, the telecom network was primarily “fixed”—the users were restricted to being stationary (other than some limited mobility offered by cordless phones) in locations that were connected to the telecom infrastructure through fixed lines. The advent of mobile telephony in the 1980s, made it possible to move around during a call, eventually ushering in the concept of “personal communication services,” allowing each user to conduct his or her own call on his or her own personal handset, when and where he or she wants to. To facilitate this, the fixed network was augmented with special mobility equipment (towers with radio antennas, databases that tracked the movement of users, etc.) that interoperated with the existing infrastructure.

The invention of the browser launched the Internet in the mid-1990s. Like mobile networks, the Internet also builds on top of the fixed telecom network infrastructure, by augmenting it with special equipment (modems, routers, servers, etc.). However, the development of the Internet has not only made it possible to bring communication services to billions of people, but has also dramatically transformed the business and the social world. Today, the widespread use of the Internet by people all over the world and from all walks of life has made it possible for Internet Protocol (IP) to be not only the data network protocol of choice, but also the technology over which almost all forms of future networks are built, including the traditional public telephone network.

At a high level, a telecom network can be thought of as comprising two main components: a core network with geographically distributed shared facilities interconnected through information transport systems, and an access network comprising dedicated links (copper, fiber, or radio) connecting individual users (or their premises) to the core network. The utility provided by the infrastructure is largely dependent on availability of transport capacity. The Communications Act of 1934 called for “rapid, efficient, Nation-wide, and world-wide wire and radio communication service with adequate facilities at reasonable charges” to “all the people of the United States.” This principle of universal service helped make telephone service ubiquitous. Large capital investments were made to deploy transport facilities to connect customers who needed telephone service, across wide geographic areas. The deployed transport capacity was much above what was needed for a basic voice service, and as new technology solutions developed (digital subscriber, optical fiber, etc.), it became clear that the network was capable of doing much more than providing voice services. The infrastructure that supported ubiquitous voice service, evolved to support data services and provide broadband access to the Internet. Today, according to the Federal Communications Commission (FCC), high-speed (broadband) Internet is an essential communication technology, and ought to be as ubiquitous as voice. The quality of political, economic, and social life is so dependent on the telecom network that it is rightly deemed as a nation’s critical infrastructure.

The history of mobile telephony and Internet has been one of continuous evolution and ever increasing adoption. The pace of innovation has been exponential—matching the explosion of service offerings with those of new applications. Starting with the 1G mobility system that were first

introduced in the 1970s and reached a peak subscriber base of 20 million voice customers, today's 4G systems serve a global base of nearly 2 billion smartphone customers.¹

“The transformative power of smartphones comes from their size and connectivity.”² Size makes it easy to carry them around; connectivity means smartphones can not only connect people to one another, but also can deliver the full power of online capabilities and experiences, spawning new business models as with Uber and Airbnb. In addition, the fact that there are 2 billion smartphones spread out all over the world, and each one of them can “serve as a digital census-taker,” makes it possible to get real-time views of what people like and dislike, at a very granular level. But the power of this framework is limited by the ability to interconnect the devices and web sites—at the necessary interconnection speeds and manage the latency requirements for near-instantaneous services. This is made possible by a transport layer, with adequate bandwidth and throughput between any two end points.

Transport technology has continually evolved from the early local access and long-distance transport networks. Fundamental advances in the copper, radio, and fiber technologies have provided cheaper, more durable, and simpler to manage transport. The traffic being transported has evolved from analog to digital (time division multiplexed) to IP. Layered protocols have enabled many different services to share a common transport infrastructure. The hub and spoke, fixed bandwidth, fixed circuits, and relatively long duration switched sessions (an architecture that grew with the growth of voice services) are no longer the dominant patterns but one of many different traffic patterns driven by new applications such as video distribution, Internet access, and machine-to-machine communication, with mobility being an essential part of the offering. This evolution has now made it possible for the global telecom network to be the foundation for not just voice, data, and video communication among people, but also for a rich variety of social and economic activity transacted by billions of people all over the world. The next thing on the horizon is the Internet of Things (IoT)—driven by the ability to connect and control tens of billions of remote devices, facilitating the next round of productivity gains as operations across multiple industries get streamlined.

Traditionally, the users of carrier networks were fairly static and predictable, as was the traffic associated with them. As a result, the underlying architecture that carriers relied on at the time was, more often than not, based on fixed, special purpose, networking hardware. Since the operating environment within a Telecom operator followed an organizational approach that separated the IT and Network groups, the growth in traffic and the variety of services that needed to be supported, created an environment where the ability to respond to business needs was increasingly restricted. The networks of today support billions of application-heavy, photo-sharing, video-streaming, mobile devices, and emerging IoT applications, all of which require on-demand scaling, high resiliency, and the capability to do on-the-fly service modifications.

Network connectivity is at the core of every innovation these days, from cars to phones to video and more. Connectivity and network usage are exploding—data traffic on AT&T's mobile network alone grew by nearly 150,000% between 2007 and 2014. This trend will continue, with some estimates showing wireless data traffic growing by a factor of 10 by 2020. Video comprises about 60% of AT&T's total network traffic—roughly 114 petabytes per day as of this writing. To put that into perspective, 114 petabytes equals about 130 million hours of HD video. As it stands right now, sending a minute of video takes about 4 megabytes of data. In comparison, sending a minute of virtual reality (VR) video takes hundreds of megabytes. With VR becoming increasingly popular, having an easy and convenient way to access this data-hungry content is essential. At the same time, IoT is also generating mountains of data that have to be stored and analyzed. By 2020, there will be an estimated 20–50 billion connected devices.

Public networks and especially those that cover large areas have always been expensive to deploy. To enable the growth and the use of networks (the value of a network increases as more users get on it), operators have relied on technology advances to continue to improve performance and drive costs down. Compared to a public network in the twentieth century, today's all-IP network with advanced services radically changes many of the requirements. The number of different services or applications available at any one time has grown by several orders of magnitude and services are

constantly evolving, driven by mobile device capabilities, social media, and value created by available content and applications.

To keep pace with these new applications, and to be able to provide reliable networking services in the face of growing traffic and diminishing revenue, we need a new approach to networking. Traditional telecom network design and implementation has followed a tried and tested approach for several decades. When new network requirements are identified, a request for proposals (RFPs) is issued by the carrier. The vendors respond with proprietary solutions that meet interoperability standards, and each network operator then picks the solution they like. The networks are built by interconnecting physical implementation of functions. Today’s telecom networks have over 250 distinct network functions deployed—switches, routers, access nodes, multiplexors, gateways, servers, etc. Most of these network functions are implemented as stand-alone appliances—a physical “box” with unique hardware and software that implements the function and conforms to standard interfaces to facilitate interoperability with other boxes. For operational ease, network operators prefer to use one or two vendors, typically, for a given class of appliances (routers, for example). The appliance vendor often uses custom hardware to optimize the cost/performance; the software is mated with hardware, and the product can be thought of as “closed.” This creates vendor lock in, and since most deployed appliances are seldom replaced, the result is a platform lock in, with limited options for upgrading with technology advances. [Figure 1.1](#) qualitatively depicts the evolution of cost over a 10-year period. For purposes of illustration, the unit cost signifies the cost for performing one unit of the network function, (e.g., packet processing cost for 1 GB of IP traffic). Due to platform lock in, any cost improvement comes from either hardware redesigns at the plug-in level with cheaper components, or price concessions made by the vendor to maintain market position. On the other hand, technology advances characterized by advances in component technology (Moore’s Law) and competitive market dynamics, facilitate a much faster decline in unit cost, more in-line with what is needed to keep up with the




FIGURE 1.1 Cost evolution.

growth in traffic. A second problem with the current way of network design and implementation is that in the face of rampant traffic growth, network planners deploy more capacity than what is needed (since it could take several months to enhance deployed capacity, beyond simple plug add-ons), thus creating a low utilization rate.

We need to transition from a hardware-centric network design methodology to one that is software-centric. Wherever possible, the hardware deployed ought to be one that is standardized and commoditized (cloud hardware, for example) and can be independently upgraded so as to benefit from technology advances as and when they occur. The network function capability ought to be largely implemented in software running on the commodity hardware. The hardware is shared among several network functions, so that we get maximum utilization. Network capacity upgrades occur in a fluid manner, with continuous deployment of new hardware resources and network function software as and when needed. In addition, the implementation of software-defined networking (SDN) would provide three key benefits: (1) separation of the control plane from the data plane allowing greater operational flexibility; (2) software control of the physical layer enabling real-time capacity provisioning; and (3) global centralized SDN control with advanced and efficient algorithms coupled with real-time network data allowing much better multilayer network resource optimization on routing, traffic engineering, service provisioning, failure restoration, etc. This new network design methodology (depicted in Figure 1.2) has a much lower cost structure in both capital outlays and ongoing operating expenses, greater flexibility to react to traffic surges and failures, and better abilities to create new services.

In order to become a software-centric network, and tap into the full power of SDN, several steps need to be taken—rethinking the IT/Network separation, disaggregation of hardware and software, implementing network functions predominantly in software and capable of executing on a commodity cloud hardware platform, and a high degree of operational automation. Add to this the requirement that the implementation be done on open source platforms with open Application Programming Interfaces (APIs).

Today, SDN is commonly perceived to be a new networking paradigm that arose from the virtualization of data-center networks, which in turn was driven by the virtualization of compute and storage resources. However, SDN techniques have been used at AT&T (and other carriers) for more

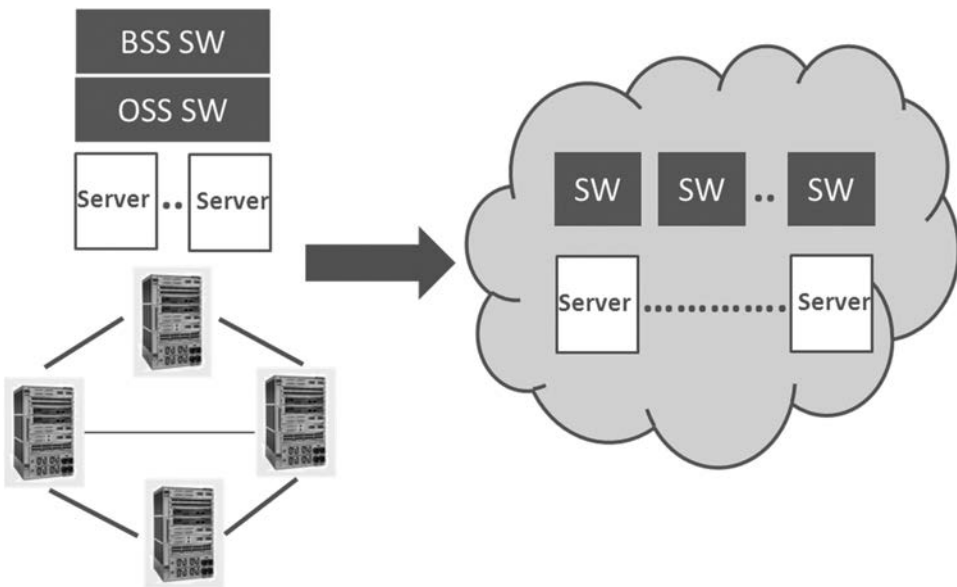


FIGURE 1.2 Transforming the network.

than a decade. These techniques have been embedded in the home grown networking architecture used to create an overlay of network-level intelligence provided by a comprehensive and global view of not only the network traffic, resources, and policies, but also the customer application's resources and policies. The motivation was to enable creation of customer application-aware, value-added service features that are not inherent in the vendor's switching equipment. This SDN-like approach has been repeatedly applied successfully to multiple generations of carrier networking services, from traditional circuit-switching services through the current IP/Multi-Protocol Label Switching (MPLS) services.

As computer technology evolved, the telecommunication industry became a large user of computers to automate tasks such as setting up connections. This also launched the first application of software in telecom—software that allowed the control of the telecommunications specific hardware in the switches. As the number and complexity of the routing decisions grew, there came a need to increase the flexibility of the network—to make global decisions about routing a call that could not be done in a distributed fashion. The first “software-defined network controller” was invented, namely the Network Control Point (NCP). Through a signaling “API,” the controller could make decisions in real time about whether a request for a call should be accepted, which end point to route the call to, and who should be billed for the call. These functions could not be implemented at the individual switch level but needed to be done from a centralized location since the subscriber, network, and service data were needed to decide how to process the call.

Initially, the NCP was a relatively static service-specific program. It had an application to handle collect calls, an application for toll-free 800 services, an application for blocking calls from delinquent accounts, etc. It soon became clear that many of the service application shared a common set of functions. The “service logic” was slightly different but there were a set of primitive functions that the network supported, and a small set of logical steps that tied those functions together. The data varied and the specific parameters used for a service might vary but there was a reusable set. The programmable network and, in particular, the programmable network controller was born.³

The Direct Services Dialing Capability (DSDC) network was a new way to do networking. Switches, known as action points, had an API that could be called over a remote transport network (CCIS6/SS7) and controllers, known as control points, could invoke the programs in a service execution environment. The service execution environment was a logical view of the network. Switches queried controllers, controllers executed the program and called APIs into the switches to bill and route calls. The controller had a function to set billing parameters and to route calls for basic functions. Advanced functions used APIs to play announcements and collect digits, to look up digits for validation and conversion, and to temporarily route to a destination and then reroute the call. These functions allowed different services to be created.

The Toll Free Service (i.e., 800-service) that translated one dialed number into another destination number and set billing parameters to charge the called party for the call was another example of using SDN-like techniques in a pre-SDN world. The decisions on which destination number was controlled by the customer for features such as area code routing, time of day routing, etc. were implemented using software that controlled the switches. An entire multibillion dollar industry grew out of the simple concept of reversing the charges and the power of letting the customer decide which data center to route the call to, for reducing their agent costs or providing better service to their end customers.

For enterprise customers, AT&T provided a trademarked service called AT&T SDN ONENET, which provided private dialing plans, authentication codes and departmental charging capability as a virtual Time-Division Multiplexing (TDM) network for enterprise customers. With or without a Private Branch Exchange (PBX), the service exposed a programmable virtual network to enterprises. The DSDC network also provided the ability to add or remove resources to the call for advanced functions and mechanisms for dealing with federated controllers. Service Assist was a mechanism to service-chain an announcement and digit collection node into a call; this grew into Interactive Voice Response (IVR) technology with speech recognition but it was an early form of service chaining for virtual networks. When more than one controller was needed to handle the load

(or to perform the entire set of desired features) NCP transfer was used, which today we would see as controller-to-controller federation. Many variations of these types of service were created, across the globe, as both the ability for software control and the availability of computing to do more processing, allowed the service providers to let their customers manage their business.

These SDN-like networks in the circuit-switched world were followed by a similar capability of providing network-level intelligent control in the packet-routed world, where the routers replace the switches. The development and implementation of SDN for AT&T's IP/MPLS network—generally known as the Intelligent Routing Service Control Platform (IRSCP)—was launched in 2004.⁴ It was driven with the goal of providing AT&T with “competitive service differentiation” in a fast-maturing and commoditized general IP networking services world. The IRSCP design uses a software-based network controller to control a set of specially enhanced multiprotocol Border Gateway Protocol (BGP) (MP-BGP) route reflectors and dynamically distribute selective, fine-grained routing and forwarding controls to appropriate subsets of IP/MPLS routers. The essential architectural construct of IRSCP enables on-demand execution of customer-application specific routing and forwarding treatment of designated traffic flows, with potential control triggers based on any combination of the network's and the customer-application's policies, traffic and resource status, etc. This yields a wide range of value-added features for many key networking applications such as Content Distribution Network (CDN), networked-based security (Distributed Denial of Service [DDoS] mitigation), Virtual Private Cloud, etc.

Apart from SDN, the other key aspect of the transformation is network function virtualization with transitioning to an architecture that is cloud-centric. AT&T started its own cloud journey in late 2010 with multiple efforts, each focused on different objectives—a VMware-based cloud for traditional IT business applications, a Compute-as-a-Service (CaaS) offering for external customers, and an OpenStack-based effort for internal and external developers. These disparate clouds were subsequently merged into a common cloud platform, which evolved into the AT&T Integrated Cloud (AIC). Today, this is the cloud environment in which AT&T's business applications run. It is also the platform for the cloud-centric network function virtualization effort. By 2020, it is anticipated that 75% of the network would be virtualized and running in the cloud.

The methods and mechanisms for deploying applications on a dedicated server infrastructure are well known, but a virtualized infrastructure has properties, such as scalability and active reassignment of idle capacity, which are not well understood. If applications are not structured to make use of these capabilities, they will be more costly and less efficient than the same application running on a dedicated infrastructure. Building services that are designed around a dedicated infrastructure concept and deploying them in a virtualized infrastructure fails to exploit the capabilities of the virtualized network. Furthermore, building a virtualized service that makes no use of SDN to provide flexible routing of messages between service components adds significantly to the complexity of the solution when compared to a dedicated application. Virtualization and SDN are well defined, but the key is to enable using virtualization and SDN together to simplify the design of virtualized services. Integration of virtualization and SDN technologies boils down to understanding how decomposition, orchestration, virtualization, and SDN can be used together to create, manage, and provide a finished service to a user. It enables services to be created from modular components described in recipes where automated creation, scaling, and management are provided by Open Network Automation Platform (ONAP), an open source software platform that powers AT&T's SDN. ONAP enables

- Independent management of applications, networking, and physical infrastructure
- A service creation environment that is not limited by a fixed underlying network or compute infrastructure
- The automatic instantiation and scaling of components based on real-time usage
- The efficient reuse of modular application logic
- Automatic configuration of network connectivity via SDN
- User definable services

ONAP will benefit service providers by driving down operations costs. It will also give providers and businesses greater control of their network services becoming much more “on demand.” At the end of the day, customers are the ones who will benefit the most. The idea of creating a truly personalized secure set of services on demand will change the way we enable consumer applications and business services.

Today’s telecom networks are rich treasure troves of data—especially pertaining to the world of mobility. On an average day, AT&T measures 1.9 billion network-quality check points from where our wireless customers are actually using their service. This data go into network analytics to better understand the customers’ true network experience. The program is internally referred to as Service Quality Management; sophisticated analytics is used to make sense of this massive volume of network data and discern what customers are experiencing. These technologies have transformed how we manage our network—driving smarter decisions and resolving issues quicker than ever before.⁵

For instance, if two areas have disrupted cell towers—restoring service to all customers quickly is a key operational objective. Based on the real-time data and historical measurements that are available, AT&T’s home developed algorithm called TONA—the Tower Outage and Network Analyzer—factors in current users, typical usage for various times of the day, population and positioning of nearby towers, to assess which towers can respond by off-loading the traffic from the disrupted towers. TONA has created a 59% improvement in identifying customer impact, and it shortens the duration of network events for the greatest number of customers. This is but one example of the power of using real-time data to improve network performance.

This transformation is not limited to the network and its associated technologies. The transformed network also needs a software-centric workforce. Reskilling the employees in a variety of software and data specialties is as important as transforming the architecture of the network. Besides, the traditional RFP approach, used today to procure new products or introduce new vendors, gets replaced by a more iterative and continuous approach for specifying and procuring software. It also calls for active involvement with the open source community.⁶ Open source speeds up innovation, lower costs, and helps to converge quickly on a universal solution, embraced by all. Initially conceived to be something people can modify and share, because the design is publicly accessible, “open source” has come to represent open exchange, collaborative participation, rapid prototyping, transparency, meritocracy, and community-oriented development.⁷ The phenomenal success of the Internet can be attributed to open source technologies (Linux, Apache Web server application, etc.) Consequently, anyone who is using the Internet today benefits from open source software. One of the tenets of the open source community is that you do not just take code. You contribute to it, as well.

AT&T strongly believes that the global telecom industry needs to actively cultivate and nurture open source. To spur this, we have decided to contribute the ONAP platform to open source. Developed by AT&T and its collaboration partners, the decision to release ONAP into open source was driven by the desire to have a global standard on which to deliver the next generation of applications and services. At the time of this writing, there are at least four active open source communities pertaining to the Network Functions Virtualization (NFV)/SDN effort:

- OpenStack
- ON.Lab
- OpenDaylight
- OPNFV

AT&T has been active in all four, and will continue to do so. Many other telecom companies, as well as equipment vendors, have also been equally active.

Any significant effort to re-architect the telecom network to be software-centric cannot ignore the needs of wireless (mobile) systems. The fourth generation wireless system, Long-Term Evolution (LTE), which was launched in late 2009, made significant advances in the air interface, as well as the architecture of the core network handling mobile traffic. The driver behind this was the desire to

build a network that was optimized for smartphones and the multitude of applications that runs on them. As of this writing, fifth generation mobile networks are being architected and designed and their rollout is expected in 2020. Compared to the existing 4G systems, 5G networks offer

- Average data rate to each device in the range of 50 Mb/s, compared to the current 5 Mb/s
- Peak speeds of up to 1 Gb/s simultaneously to many users in the same local area
- Several hundreds of thousands of simultaneous connections for massive wireless sensor network
- Enhanced spectral efficiency for existing bands
- New frequency bands with wider bandwidths
- Improved coverage
- Efficient signaling for low latency applications

New use cases such as the IoT, broadcast-like services, low latency applications, and lifeline communication in times of natural disaster will call for new networking architectures based on network slicing (reference). NFV and SDN will be an important aspect of 5G systems—for both the radio access nodes as well as the network core.

The chapters that follow explore the concepts discussed in this chapter, in much more detail. The transformation that is upon us will achieve hyper scaling, not unlike what the web players have done for their data centers. We think this is one of the biggest shifts in networking since the creation of the Internet.

REFERENCES

1. <https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/>
2. <http://www.economist.com/news/leaders/21645180-smartphone-ubiquitous-addictive-and-transformative-planet-phones>
3. 800 service using SPC network capability—special issue on stored program controlled network, *Bell System Technical Journal*, Volume 61, Number 7, September 1982.
4. J. Van der Merwe et al. Dynamic connectivity management with an intelligent route service control point, *Proceedings of the 2006 SIGCOMM Workshop on Internet Network Management*. Pisa, Italy, pp. 147–161.
5. <http://about.att.com/innovationblog/121014datapoints>
6. <http://about.att.com/innovationblog/061714hittingtheopen>
7. <https://opensource.com/resources/what-open-source>

The Need for Change

<https://www.statista.com/statistics/330695/number-of-smartphone-users-worldwide/>

<http://www.economist.com/news/leaders/21645180-smartphone-ubiquitous-addictive-and-transformative-planet-phones>

800 service using SPC network capability—special issue on stored program controlled network, Bell System Technical Journal, Volume 61, Number 7, September 1982.

J. Van der Merwe Dynamic connectivity management with an intelligent route service control point, Proceedings of the 2006 SIGCOMM Workshop on Internet Network Management. Pisa, Italy, pp. 147–161.

<http://about.att.com/innovationblog/121014datapoints>

<http://about.att.com/innovationblog/061714hittingtheopen>

<https://opensource.com/resources/what-open-source>

Transforming a Modern Telecom Network—From All-IP to Network Cloud

<https://tools.ietf.org/html/rfc6020>

<http://www.etsi.org/technologies-clusters/technologies/nfv>

http://www.etsi.org/deliver/etsi_es/282000_282099/282007/02.00.00_50/es_282007v020000m.pdf

Network Functions Virtualization

Anderson, L. 2006. Framework for Layer 2 Virtual Private Networks (L2VPNs), RFC 4664. Retrieved from IETF: <https://www.ietf.org/rfc/rfc4664.txt>

Apache . n.d. OpenContrail. www.opencontrail.org

Central, S. n.d. Service Chaining. <https://www.sdxcentral.com/sdn/network-virtualization/definitions/what-is-network-service-chaining/>

Cisco . n.d. LISP Overview. http://lisp.cisco.com/lisp_over.html

Comer, D. E. 2014. Computer Networks and Internets. Upper Saddle River, New Jersey: Pearson Education, Inc. ISBN: 978-0133587937.

Foundation, O. N. n.d. <https://www.opennetworking.org/sdn-resources/openflow>

IETF . n.d. NETCONF. <https://tools.ietf.org/html/rfc6241>

IETF . n.d. Segment Routing. <https://tools.ietf.org/html/draft-ietf-spring-segment-routing-04>

IETF . n.d. YANG. <https://tools.ietf.org/html/rfc6020>

Intel . n.d. SR-IOV. <http://www.intel.com/content/www/us/en/pci-express/pci-sig-sr-iov-primer-sr-iov-technology-paper.html>

International, S. R. n.d. SNMP. <http://www.snmp.com/>

Jun Wu, Z. Z. n.d. Cloud Radio Access Network (C-RAN): A Primer.

<http://www.ntu.edu.sg/home/ygwen/Paper/WZHW-WCM-15.pdf>

KVM . n.d. Live Migration. Retrieved from <http://www.linux-kvm.org/page/Migration>

Openstack . n.d. Neutron. <http://docs.openstack.org/developer/neutron/>

Rekhter, K. K. 2007. Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling (RFC 4761). Retrieved from IETF: <https://tools.ietf.org/html/rfc4761>

Rosen, E. 2006. BGP/MPLS IP Virtual Private Networks (VPNs), RFCs 4364 and 2547. Retrieved from <https://tools.ietf.org/html/rfc4364>

Rosenblum, M. 2004. The reincarnation of virtual machines. ACM Queue. July–August. 2(5): 34–40.

Sajassi, A. 2015. BGP MPLS-Based Ethernet VPN. Retrieved from IETF: <https://tools.ietf.org/html/rfc7432>

Silberschatz, A. , Galvin, P. B. , and Gange, G. 2012. Operating System Concepts. Hoboken, New Jersey: John Wiley and Sons. ISBN: 978-0470128725.

Simpson, W. 1995. IP in IP Tunneling, RFC 1853. Retrieved from <https://tools.ietf.org/html/rfc1853>

Smith, P. 2011. Network resilience: A systematic approach. IEEE Communications Magazine. 49(7): 88–97.

Spurgeon, C. E. and Zimmerman, J. 2000. Ethernet: The Definitive Guide: Designing and Managing Local Area Networks, 2nd Edition. Sebastopol, California: O'Reilly Media. ISBN: 978-1449361846.

VMWare . n.d. Introduction to VMware DRS and VMware HA Clusters. Retrieved from VMware vSphere 5 Documentation Center: <https://pubs.vmware.com/vsphere-55/index.jsp>

VMWare . n.d. VMware NSX. <http://www.vmware.com/products/nsx.html>

Network Functions Virtualization Infrastructure

n.d. Retrieved from <http://searchsdn.techtarget.com/definition/virtual-network-functions>
n.d. Retrieved from <http://searchsdn.techtarget.com/tip/NFV-Infrastructure-What-really-lies-beneath>
n.d. Retrieved from <http://searchsdn.techtarget.com/tip/NFV-Infrastructure-What-really-lies-beneath>
n.d. Retrieved from [https://en.wikipedia.org/wiki/Server_\(computing\)](https://en.wikipedia.org/wiki/Server_(computing))
n.d. Retrieved from https://en.wikipedia.org/wiki/Central_processing_unit
n.d. Retrieved from https://en.wikipedia.org/wiki/Graphics_processing_unit
n.d. Retrieved from https://en.wikipedia.org/wiki/Networking_hardware
n.d. Retrieved from https://en.wikipedia.org/wiki/Network_interface_controller
n.d. Retrieved from <http://searchsdn.techtarget.com/definition/control-plane-CP>
n.d. Retrieved from [https://en.wikipedia.org/wiki/Router_\(computing\)](https://en.wikipedia.org/wiki/Router_(computing))
n.d. Retrieved from <http://searchstorage.techtarget.com/definition/block-storage>
n.d. Retrieved from <http://searchstorage.techtarget.com/definition/object-storage>
n.d. Retrieved from <https://www.sdxcentral.com/nfv/definitions/virtualized-infrastructure-manager-vim-definition/>
n.d. Retrieved from <https://en.wikipedia.org/wiki/OpenStack>
n.d. Retrieved from <https://www.openstack.org/foundation/>
n.d. Retrieved from <https://en.wikipedia.org/wiki/Virtualization>
n.d. Retrieved from https://en.wikipedia.org/wiki/Virtual_machine
n.d. Retrieved from <http://searchservvirtualization.techtarget.com/definition/container-based-virtualization-operating-system-level-virtualization>
n.d. Retrieved from <http://searchservvirtualization.techtarget.com/definition/hypervisor>
n.d. Retrieved from <http://searchstorage.techtarget.com/definition/storage-area-network-SAN>
n.d. Retrieved from https://en.wikipedia.org/wiki/Storage_area_network
n.d. Retrieved from <http://searchstorage.techtarget.com/definition/storage-virtualization>
n.d. Retrieved from https://en.wikipedia.org/wiki/Storage_virtualization
n.d. Retrieved from <https://wiki.openstack.org/wiki/Swift>
n.d. Retrieved from <https://wiki.openstack.org/wiki/Cinder>
n.d. Retrieved from <https://www.sdxcentral.com/nfv/definitions/what-is-nfv-orchestration/>
n.d. Retrieved from <https://www.sdxcentral.com/nfv/definitions/nfv-orchestrator-nfv-definition/>
n.d. Retrieved from https://en.wikipedia.org/wiki/Agile_software_development
n.d. Retrieved from <https://en.wikipedia.org/wiki/DevOps>
n.d. Retrieved from <http://searchsoftwarequality.techtarget.com/definition/continuous-integration>
n.d. Retrieved from <http://searchdatacenter.techtarget.com/photostory/4500255104/Data-center-terminology-that-will-get-you-hired/3/Define-continuous-integration-and-delivery-as-Dev-meets-Ops>
n.d. Retrieved from <http://searchitoperations.techtarget.com/definition/continuous-delivery-CD>
n.d. Retrieved from <https://www.sdxcentral.com/wp-content/uploads/2014/09/opnfv-nfv-open-source-network-functions-virtualization-diagram.jpg>
n.d. Retrieved from <http://stackalytics.com/>

Architecting the Network Cloud for High Availability

3GPP . 2016. 3GPP. Retrieved from 3GPP: www.3gpp.org
ADC . 2006. TIA-942 Data Center Standards Overview—102264AE—Accu-Tech, ADC Telecommunications Whitepaper. www.accu-tech.com/hs-fs/hub/54495/file-15894024-pdf/docs/102264ae.pdf
Amazon . 2016. AWS Global Infrastructure. Retrieved from Amazon: <https://aws.amazon.com/about-aws/global-infrastructure/>
Amazon Web Services . 2013. Amazon EC2 Service Level Agreement. <https://aws.amazon.com/ec2/sla/>
Bennett, C. and Tseitlin, A. 2012. Chaos Monkey Released into the Wild, Netflix Tech Blog, July 30, 2012. <http://techblog.netflix.com/2012/07/chaos-monkey-released-into-wild.html>
Butler, B. 2016. And the cloud provider with the best uptime in 2015 is.... Network World. January 7. <http://www.networkworld.com/article/3020235/cloud-computing/and-the-cloud-provider-with-the-best-uptime-in-2015-is.html>
Ceph . 2016. Ceph. Retrieved from www.ceph.com
Choudhury, G. 2004. Models for IP/MPLS routing performance: Convergence, fast reroute, and QoS impact. In Performance, Quality of Service and Control of Next-Generation Communications Networks II (Proceedings of SPIE). Philadelphia, PA: SPIE—The International Society for Optical Engineering, pp. 1–10.
Giloth, G. F. 1987. The evolution of fault tolerant switching systems in AT&T. In H. K. A. Avizienis (ed.), The Evolution of Fault Tolerant Computing (pp. 37–54). Baden, Austria: Springer-Verlag/Wien.

Healy, J. 2016. Network Outage Reporting System (NORS). <https://www.fcc.gov/network-outage-reporting-system-nors>

Hoeflin, D. A. 2005. An integrated defect tracking model for product deployment in telecom services. In Proceedings of the 10th IEEE Symposium on Computers and Communication (ISCC 2005) (pp. 927–932). IEEE. June 27–30, 2005, Cartagena, Murcia, Spain.

Hoff, T. 2015. The Stunning Scale of AWS and What It Means for the Future of the Cloud. <http://highscalability.com/blog/2015/1/12/the-stunning-scale-of-aws-and-what-it-means-for-the-future-o.html>

Martersteck, K. E. and Spencer, A. E. Jr. 1985. The 5ESS switching system—Introduction. AT&T Technical Journal, 64(6), part 2, 1305–1314.

IEEE Computer Society . 2010. IEEE Std. 1044-2099 (Revisions of IEEE Std. 1044-1993) IEEE Standard Classification for Software Anomalies. New York: IEEE Computer Society.

Johnson, C. R. 2004. VoIP reliability: A service provider's perspective. IEEE Communications, July, pp. 48–54.

Klincewicz, J. G. 2013. Designing an IP link topology for a metro area backbone network. International Journal of Interdisciplinary Telecommunications and Networking, 5(1), 26–42.

Krishnan, K. D. 1994. Unified models of survivability for multi-technology networks. In Proceedings of ITC 14. New York: Elsevier Science B.V, Volume 1A, pp. 655–665.

Lancaster, W. G. 1986. Carrier grade: Five nines, the myth and the reality. Pipeline (www.pipelinepub.com), 3(11), 6.

Lyu, M. R. 1996. Handbook of Software Reliability Engineering. Los Alamitos, CA: IEEE Computer Society Press and McGraw-Hill.

Martersteck, K. E. 1981. The 4ESS system evolution. Bell System Technical Journal, 60(6), part 2, 1041–1228.

Metz, C. 2002. IP anycast: Point-to-(any) point communication. IEEE Internet Computing, 6(2), 94–98.

Miller, R. 2015. Inside Amazon's Cloud Computing Infrastructure. <http://datacenterfrontier.com/inside-amazon-cloud-computing-infrastructure/>

Musa, J. 1999. Software Reliability Engineering. New York: McGraw-Hill.

Rackspace . 2015. Creating Resilient Architectures on Microsoft Azure. Windcrest, TX: Rackspace.

Govindan Ramesh (Google, USC); Minei Ina , Kallahalla Mahesh , Koley Bikash , Vahdat Amin (Google). 2016. Evolve or die: High availability design principles drawn from Google's network infrastructure. Proceedings of the 2016 ACM Conference on Special Interest Group on Data Communication, pp. 58–72. August 22–26, Florianopolis, Brazil.

RCloud . 2016. RCloud. Retrieved from www.rcloud.social/index.html

Taylor, M. 2016. Telco-grade service availability from an it-grade cloud. IEEE CQR, May 9–12 (p. 21).

Stevenson, WA: IEEE.

Tortorella, M. 2005a. Service reliability theory and engineering, I: Foundations. Quality Technology and Quantitative Management, 2(1), 1–16.

Tortorella, M. 2005b. Service reliability theory and engineering, II: Models and examples. Quality Technology and Quantitative Management, 2(1), 17–37.

Tseitlin, C. B. 2012. Chaos Monkey Released into the Wild. <http://techblog.netflix.com/2012/07/chaos-monkey-released-into-wild.html>

Wirth, P. A. 1988. A unifying approach to performance and reliability objectives. ITC 12, June 1998. Torino, Italy: International Advisory Council of the International Teletraffic Congress, (Vol 4, pp. 4.2B2.1–4.2B2.7).

Wu, T.-H. 1992. Fiber Network Survivability. Boston, MA: Artech House, Inc.

Zhang, X. 2009. Software risk management. In Sherif, M. H. (ed.), Handbook of Enterprise Integration. Boca Raton, FL: CRC Press. Chapter 10, pp. 225–270.

Software-Defined Networking

YANG —A Data Modeling Language for the Network Configuration Protocol (NETCONF), IETF RFC6020, <https://tools.ietf.org/html/rfc6020>.

S. Agrawal , V. Narasayya , B. Yang , Integrating vertical and horizontal partitioning into automated physical database design, Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data, June 13–18, 2004, Paris, France.

Node-RED —<http://www.nodered.org/>

PCEP —Path Computation Element (PCE) Communication Protocol (PCEP) RFC5440, <https://tools.ietf.org/html/rfc5440>.

ODL —<http://www.opendaylight.org/>

OpenContrail —<http://www.opencontrail.org/>

ONOS —<http://onosproject.org/>

Network Data and Optimization

Choosing a NoSQL Database—Technology Comparison Matrix, October 14, <http://accionlabs.com/2014/10/16/choosing-a-nosql-database-technology-comparison-matrix/>

Android Developer Package Index, Android API's, <https://developer.android.com/reference/packages.html>

The Apache Software Foundation, <https://www.apache.org/>

Nanocubes: Fast Visualization of Large Spatiotemporal Datasets, <http://nanocubes.net/>

Cai, L. , Zhu, Y. 2015. The challenges of data quality and data quality assessment in the big data era. *Data Science Journal*, 14, p. 2. DOI: <http://doi.org/10.5334/dsj-2015-002>

Cask Data Application Platform, <http://cask.co/products/cdap/>

AT&T Deploying SON Technology as Part of Latest Network Upgrades, <http://www.cellular-news.com/story/Operators/53297.php>

Cloudera, <http://www.cloudera.com/>

Understanding Big Data: The Seven V's, May 22, 2014, <http://dataconomy.com/seven-vs-big-data/>

Deng, L. Recent advances in deep learning for speech recognition at Microsoft, ICASSP 2013—2013 IEEE International Conference on Acoustics, Speech and Signal Processing, May 26–31, Vancouver, BC, Canada.

Desai, A. , Nagegowda, K.S. , Ninikrishna, T. 2016. A framework for integrating IoT and SDN using proposed OF-enabled management device, 2016 International Conference on Circuit, Power, and Computing Technologies (ICCPCT), March 18–19, Tamilnadu, India.

Apache Hadoop Turns 10: The Rise and Glory of Hadoop, February 2016, <https://www.dezyre.com/article/apache-hadoop-turns-10-the-rise-and-glory-of-hadoop/211>

Sqoop vs. Flume Battle of the Hadoop ETL Tools, October 15, <https://www.dezyre.com/article/sqoop-vs-flume-battle-of-the-hadoop-etl-tools-176>

ECOMP (Enhanced Control, Orchestration, Management and Policy) Architecture White Paper, 2016, <https://www.google.com/#q=ecomp+white+paper>

“Big Data” is No Longer Enough: It's Now All about “Fast Data”, May 2016, <https://www.entrepreneur.com/article/273561>

5G What is it? <http://www.ericsson.com/res/docs/2014/5g-what-is-it.pdf>

Lewis, J. , Fowler, M. 2014. Microservices, <http://www.martinfowler.com/articles/microservices.html>

Comprehensive Online Tracking is Not Unique to ISPs, May 20, 2016, <https://fpf.org/2016/05/20/14382/>

Gilbert, M. , Jana, R. , Noel, E. , Gopalakrishnan, V. 2016. Control loop automation management platform, 5th IEEE Global Conference on Signal and Information Processing—IEEE GlobalSip 2016, Washington, DC.

DeepMind AlphaGo, 2016, <https://gogameguru.com/tag/deepmind-alphago-lee-sedol>

Ghemawat, S. , Gobiuff, H. , Leung, S. , 2003, The Google File System. <https://static.googleusercontent.com/media/research.google.com/en//archive/gfs-sosp2003.pdf>

Dean, J. , Ghemawat, S. MapReduce: Simplified data processing on large clusters, OSDI'04: Sixth Symposium on Operating System Design and Implementation, December 2004.

SDN, Big Data, and the Self-Optimizing Network, InfoWorld|October 23, 2013, <http://www.infoworld.com/article/2612727/sdn/sdn--big-data--and-the-self-optimizing-network.html>

The 7 V's of Big Data, <https://www.impacradius.com/blog/7-vs-big-data/>

Apache Hadoop, <http://hortonworks.com/>

iOS Developer Library, <https://developer.apple.com/library/index.html>

Apache Kafka: A Distributed Streaming Platform, <https://kafka.apache.org/>

Kotsiantis, S.B. 2007. Supervised machine learning: A review of classification techniques, *Informatica*. 31, 2007, 249–268. <http://www.informatica.si/index.php/informatica/article/viewFile/148/140>

Lambda Architecture, <http://lambda-architecture.net/>

Laney, D. 2001. 3D Data Management: Controlling Data Volume, Velocity and Variety, META Group, Stamford, CT. <http://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>

Lecun, Y. , Bengio, Y. , Haffner, P. 1998. Gradient based learning applied to document recognition, *Proceedings of the IEEE*. 86, 11.

Light Reading's Big Telecom Event, The Role of Big Data and Advanced Analytics in SDN/NFV, https://accedian.com/wp-content/uploads/2015/06/BTE15_The-Role-of-Big-Data-and-Advanced-Analytics-in-SDN-NFV.pdf

The 7 V's of Big Data, Livingstone Advisory, <https://livingstoneadvisory.com/2013/06/big-data-or-black-hole/> <https://www.mapr.com/>

Mitchell, T. 1997. *Machine Learning*, McGraw-Hill Series in Computer Science. McGraw-Hill, Inc. New York, NY.

Minsky, M. , Papert, S. 1969. *Perceptrons: An Introduction to Computational Geometry*, MIT Press, Cambridge, Massachusetts.

Ng, A. 2016. *Machine Learning*, Coursera Course, <https://www.coursera.org/learn/machine-learning>

NGMN Use Cases Related to Self Organising Network, Overall Description, https://www.ngmn.org/uploads/media/NGMN_Use_Cases_related_to_Self_Organising_Network__Overall_Des

cription.pdf

FCC Releases Proposed Rules to Protect Broadband Consumer Privacy, Paragraph 4,
<https://www.fcc.gov/document/fcc-releases-proposed-rules-protect-broadband-consumer-privacy>

Web Filtering, <https://www.opendns.com/enterprise-security/solutions/web-filtering/>

Gopalan, R. 2016. Predictive Machine Learning for Software-Defined Networking Applications, AT&T Internal White Paper.

Rahim, M. 1991. Artificial Neural Network for Speech Analysis/Synthesis, Chapman and Hall, Cambridge.

Rumelhart, D.E. , Hinton, G.E. , Williams, R.J. 1986. Learning Internal Representations by Error Propagation, MIT Press, Cambridge, Massachusetts.

Intel Helps Developers Blend Machine Learning and IoT, <https://www.sdxcentral.com/articles/news/intel-helps-developers-blend-machine-learning-iot/2016/08/>

Shannon, C.E. 1948. A mathematical theory of communication, Bell Systems Technical Journal, 27, 3, 379–423.

Turning, A.M. 1950. Computing machinery and intelligence. Mind, 59, 236, 433–460.

<http://www.jstor.org/stable/2251299>

Apache STORM, <http://storm.apache.org/releases/current/index.html>

Apache SPARK RDD, https://www.tutorialspoint.com/apache_spark/apache_spark_rdd.htm

Wang, J.L. , Li, H. , Wang, Q. 2010. Research on ISO 8000 series standards for data quality, Standard Science, 12, pp. 44–46.

Big Data, https://en.wikipedia.org/wiki/Big_data

Apache Hadoop, https://en.wikipedia.org/wiki/Apache_Hadoop

Privacy Policy, https://en.wikipedia.org/wiki/Privacy_policy

Wyatt, M. 2016. Microservices Principles, AT&T Internal White Paper.

Apache Hadoop YARN, <https://hadoop.apache.org/docs/r2.7.2/hadoop-yarn/hadoop-yarn-site/YARN.html>

Hadoop Zookeeper Tutorial, <https://www.dezyre.com/hadoop-tutorial/zookeeper-tutorial>

Network Security

E. Amoroso . Rings around Things, AT&T Cyber Security Conference.

Seven Pillars of Carrier Grade Security in the AT&T Global IP/MPLS Network, AT&T, 2011.

AT&T Information and Network Security Customer Reference Guide, January, 2010, Version 4.1.

https://www.nist.gov/sites/default/files/documents/itl/AT-T_APPENDIX_A_-_ATTSecurity_Customer_Reference_Guide-1.pdf.

ECOMP (Enhanced Control, Orchestration, Management, and Policy) Architecture White Paper, AT&T Inc., 2016.

AT&T Project Astra, ISE® Northeast Project Award Winner 2015 and ISE® North America Project Award Winner 2015—Commercial Category.

M. Gillespie . 2015. Intel Trusted Execution Technology—A Primer, Intel Software Network:

<http://software.intel.com>.

T. Kim , M. Peinado , G. Mainar-Ruiz . StealthMEM: System-Level Protection against Cache-Based Side Channel Attacks in the Cloud, USENIX Association, August 8, 2012.

Enterprise Networks

<http://www.nongnu.org/quagga>

VMWARE Virtualization: <http://www.vmware.com/solutions/virtualization.html>

Linux Containers: <https://linuxcontainers.org>

Linux Bridge: <https://wiki.linuxfoundation.org/networking/bridge>

Open Virtual Switch: <http://www.openvswitch.org/>

DPDK Poll-mode: http://dpdk.org/doc/guides/prog_guide/poll_mode_drv.html

SR0-IOV: <http://www.intel.com/content/www/us/en/pci-express/pci-sig-sr-iov-primer-sr-iov-technology-paper.html>

Hyperscan: <http://www.intel.com/content/dam/www/public/us/en/documents/solution-briefs/hyperscan-scalability-solution-brief.pdf>

<http://www.ijcaonline.org/research/volume131/number16/sarma-2015-ijca-9076010.pdf>

Network Access

<https://www.fsan.org/task-groups/ngpon/>

http://www.5gamericas.org/files/2214/7257/3276/Final_Mobile_Broadband_Transformation_Rsavy_whitepaper.pdf

http://www.5gamericas.org/files/3214/8833/1313/3GPP_Rel_13_15_Final_to_Upload_2.28.17_AB.pdf

[http://www.semiconductors.org/clientuploads/Research_Technology/ITRS/2015/0_2015%20ITRS%202.0%20Executive%20Report%20\(1\).pdf](http://www.semiconductors.org/clientuploads/Research_Technology/ITRS/2015/0_2015%20ITRS%202.0%20Executive%20Report%20(1).pdf)

<http://www.khronos.org/OpenGL/>

<https://www.broadband-forum.org/>

<http://www.3gpp.org/>

<http://www.cpri.info/>

Network Edge

Bosshart, P. ; P4: Programming protocol-independent packet processors; ACM SIGCOMM Computer Communication Review; Vol. 44, No. 3, 88–95, July 2014.

Clos, C. A study of non-blocking switching networks; Bell System Technical Journal; 32, No. 2, 406–424, March 1953.

Kodeboyina, C. ; An open-source P4 switch with SAI support; 2015 (<http://p4.org/p4/an-open-source-p4-switch-with-sai-support/>).

Kreutz, D. ; Software-defined networking: A comprehensive survey; Proceedings of the IEEE; Vol. 103, No. 1, 14–76, January 2015.

Lemos, R. ; Internet routers hitting 512K limit, some become unreliable; Ars Technica, August 13, 2014 (<https://arstechnica.com/security/2014/08/internet-routers-hitting-512k-limit-some-become-unreliable/>).

Making the Switch ; Disruptive Telecom White Box Collaboration Accelerates and Opens the Platform, Powering Unprecedented Network Performance and Insights; AT&T, April 4, 2017 (http://about.att.com/story/white_box_collaboration.html)

McKeown, N. and Rexford, J. ; Clarifying the differences between P4 and Openflow; 2016 (<http://p4.org/p4/clarifying-the-differences-between-p4-and-openflow/>).

McKeown, N. ; OpenFlow: Enabling innovation in campus networks; ACM SIGCOMM Computer Communication Review; Vol. 38, No. 2, 69–74, April 2008.

Moore, G. ; Gordon Moore: The Man Whose Name Means Progress, The visionary engineer reflects on 50 years of Moore's Law; IEEE Spectrum; March 30, 2015 (<http://spectrum.ieee.org/computing/hardware/gordon-moore-the-man-whose-name-means-progress>).

Open vSwitch ; Linux Foundation Collaborative Project (<http://www.openvswitch.org/>).

Sajassi, A. ; BGP MPLS Based Ethernet VPN; draft-ietf-12vpn-evpn-11; April 2015; and EVPN VPWS Flexible Cross-Connect Service; draft-sajassi-bess-evpn-vpws-fxc-00.txt; January 2016 (ietf.org).

Singh, A. ; Jupiter rising: A decade of Clos topologies and centralized control in Google's datacenter network; Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, pp. 183–197, London, United Kingdom, August 17–21, 2015.

Sivaraman, A. ; Packet transactions: High-level programming for line-rate switches; SIGCOMM 2016 (<http://web.mit.edu/domino/>).

Stallings, W. ; Foundations of Modern Networking: SDN, NFV, QoE, IoT, and Cloud; Pearson Education Inc., Indianapolis, Indiana, 2016.

Subramaniam, K. ; Switch Abstraction Interface (SAI): A Reference Switch Abstraction Interface for OCP; Open Compute Project; 2015 (http://www.opencompute.org/wiki/Networking/SpecsAndDesigns#Switch_Abstraction_Interface).

The Fast Data Project (FD.io) and Vector Packet Processing (VPP); Linux Foundation (<https://fd.io/>).

Network Core

The fiber optics association guide to fiber optics, [Online]. Available:

<http://www.thefoa.org/tech/ref/basic/fiber.html>

ITU-T G.652 Characteristics of a single mode optical fibre and cable, November 2009. [Online]. Available:

<http://www.itu.int/rec/T-REC-G.652-201611-l/en>

ITU-T G.653 Characteristics of a dispersion-shifted single-mode optical fibre and cable, [Online]. Available:

<https://www.itu.int/rec/T-REC-G.653>

ITU-T G.655 Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable, November 2009. [Online]. Available: <https://www.itu.int/rec/T-REC-G.655>

K. Roberts , Electronic dispersion compensation beyond 10 Gb/s, Digest of IEEE/LEOS Summer Topical Meetings, pp. 9–10, 2007. [Online]. Available: <http://ieeexplore.ieee.org/document/4288305/>

M. Birk and K. Tse , Challenges for long haul and ultra-long haul optical networks, in *Optically Amplified WDM Networks*, Academic Press, Chapter 10, p. 277.

Open ROADM home page, March 2016. [Online]. Available: <http://openroadm.org/home.html>

M. Bjorklund , RFC 6020, YANG—A Data Modeling Language for the Network Configuration Protocol (NETCONF), October 2010. [Online]. Available: <https://tools.ietf.org/html/rfc6020>

R. Enns , RFC 4741, NETCONF Configuration Protocol, December 2006. [Online]. Available: <https://tools.ietf.org/html/rfc4741#section-1.1>

Consortium for on Board Optics, [Online]. Available: <http://cobo.azurewebsites.net/>

E. Rosen , A. Viswanathan and R. Callon , RFC 3031, Multiprotocol Label Switching Architecture, January 2001. [Online]. Available: <https://tools.ietf.org/html/rfc3031>

L. Andersson , I. Minei and B. Thomas , RFC 5036, LDP Specifications, October 2007. [Online]. Available: <https://tools.ietf.org/html/rfc5036>

R. Braden , L. Zhang , S. Berson , S. Herzog and S. Jamin , RFC 2205, Resource ReSerVation Protocol, September 1997. [Online]. Available: <https://tools.ietf.org/html/rfc2205>

C. Filstils and K. Michielsen , Introduction to Segment Routing and Its Concepts, [Online]. Available: <http://www.segment-routing.net/tutorials/>

J. Vasseur and J. Le Roux , RFC 5440, Path Computation Element (PCE) Communications Protocol, March 2009. [Online]. Available: <https://tools.ietf.org/html/rfc5440>

H. Gredler , J. Medved , S. Previdi , A. Farrel and S. Ray , RFC 7752, North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP, March 2016. [Online]. Available: <https://tools.ietf.org/html/rfc7752>

Service Platforms

3GPP TS 23.228: IP Multimedia Subsystem (IMS).

3GPP TS 23.402: 3GPP System Architecture Evolution.

IETF RFC 3261: SIP: Session Initiation Protocol.

IETF RFC 6733: Diameter Base Protocol.

3GPP TS 23.060 V13.4.0 (2015-09): General Packet Radio Service (GPRS).

4G Americas' Recommendations 5G Requirements and Solutions, October 2015.

NGMN White Paper, 2-17-2015.

3GPP TS 23.002: Network Architecture.

Network Operations

3GPP TS 32.450 V9.1.0 (2010-06) —Key Performance Indicators (KPI) for Evolved Universal Terrestrial Radio Access Network (E-UTRAN): Definitions (Release 9), ETSI, 2010.

G. Gomez and R. Sanchez , End-to-End Quality of Service over Cellular Networks, John Wiley & Sons, 2005. https://en.wikipedia.org/wiki/Event_correlation.

A. Mahimkar , Z. Ge , J. Wang , J. Yates , Y. Zhang , J. Emmons , B. Huntley and M. Stockert , Rapid detection of maintenance induced changes in service performance, Proceedings of the Seventh Conference on emerging Networking Experiments and Technologies, Article No. 13, Tokyo, Japan, December 6–9, 2011.

A. Mahimkar , H. H. Song , Z. Ge , A. Shaikh , J. Wang , J. Yates , Y. Zhang and J. Emmons , Detecting the performance impact of upgrades in large operational networks, Proceedings of the ACM SIGCOMM 2010 conference, pp. 303–314, New Delhi, India, August 30–September 3, 2010.

A. Mahimkar , Z. Ge , J. Yates , C. Hristov , V. Cordaro , S. Smith , J. Xu and M. Stockert , Robust assessment of changes in cellular networks, Proceedings of the ninth ACM conference on Emerging networking experiments and technologies, pp. 175–186, Santa Barbara, California, December 9–12, 2013.

H. Yan , A. Flavel , Z. Ge , A. Gerber , D. Massey , C. Papadopoulos , H. Shah and J. Yates , Argus: End-to-end service anomaly detection and localization from an ISP's point of view, Proceedings of the 31st annual IEEE International Conference on Computer Communications (INFOCOM), Orlando, Florida, March 25–30, 2012.

Z. Ge , M. Kosseifi , M. Osinski , H. Yan and J. Yates , Method and apparatus for quantifying the customer impact of cell tower outages, U.S. patent 9426665.

J. Ramiro and K. Hamied (Editor), *Self-Organizing Networks (SON): Self-Planning, Self-Optimization and Self-Healing for GSM, UMTS and LTE*, Wiley, 2011.

X. Xu , I. Broustis , Z. Ge , R. Govindan , A. Mahimkar , N.K. Shankaranarayanan and J. Wang , Magus: Minimizing cell service disruption during planned upgrades, *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*. Article No. 21, Heidelberg, Germany, December 01–04, 2015.

M. Loukides , *What is DevOps?* O'Reilly, 2012.

G. Kim , K. Behr , G. Spafford , *The Phoenix Project: A Novel about IT, DevOps, and Helping Your Business Win*, IT Revolution Process, 2014.

https://en.wikipedia.org/wiki/Agile_software_development.

C. Larman , *Agile and Iterative Development: A Manager's Guide*, Addison-Wesley, 2004.

A. Chiu , V. Gopalakrishnan , B. Han , M. Kablan , O. Spatscheck , C. Wang , and Y. Xu , EdgePlex: Decomposing the provider edge for flexibility and reliability, *Proceedings of the 1st ACM SIGCOMM Symposium on Software Defined Networking Research*. Article No. 15, Santa Clara, California, June 17–18, 2015.

Network Measurements

Birk, M. , G. Choudhury , B. Cortez , A. Goddard , N. Padi , A. Raghuram , K. Tse , S. Tse , A. Wallace , and K. Xi . Evolving to an SDN-enabled ISP backbone: Key technologies and applications. *IEEE Communications Magazine*, 54(10):129–135, 2016.

Jain, S. , A. Kumar , S. Mandal , J. Ong , L. Poutievski , A. Singh , S. Venkata B4: Experience with a globally-deployed software defined WAN. *ACM SIGCOMM Computer Communication Review*, 43(4):3–14, 2013, ACM. <https://www.opendaylight.org>, accessed June 14, 2016.

Berde, P. , M. Gerola , J. Hart , Y. Higuchi , M. Kobayashi , T. Koide , B. Lantz ONOS: Towards an open, distributed SDN OS. In *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, pp. 1–6, Chicago, Illinois, August 22, 2014, ACM. <https://www.opennetworking.org/sdn-resources/openflow>, accessed on June 14, 2016.

Choudhury, G. , B. Cortez , A. Goddard , N. Padi , A. Raghuram , S. Tse , A. Wallace , and K. Xi . Centralized optimization of traffic engineering tunnels in a large ISP backbone using an SDN controller. *INFORMS Optimization Society Conference*, March 17–19, 2016, Princeton, NJ.

Zhang, Y. , M. Roughan , C. Lund , and D. Donoho . An information theoretic approach to traffic matrix estimation. In *SIGCOMM '03, Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pp. 301–312, Karlsruhe, Germany, August 25–29, 2003.

Feldmann, A. , A. Greenberg , C. Lund , N. Reingold , J. Rexford , and F. True . Deriving traffic demands for operational IP networks: Methodology and experience. *IEEE/ACM Transactions on Networking*, 9(3):265–280, June 2001.

YANG—A Data Modeling Language : <https://tools.ietf.org/html/rfc6020>

Network Configuration Protocol (NETCONF) : <https://tools.ietf.org/html/rfc6241>

Apache AVRO : <http://avro.apache.org/>

Google Remote Procedure Call (gRPC) : <http://www.grpc.io/docs/>

TWAMP Function Configuration : <https://datatracker.ietf.org/doc/draft-cmzrjp-ippm-twamp-yang/>

IEEE Connectivity Fault Monitoring (802.1ag) Standards Specification.

RESTCONF Protocol : <https://tools.ietf.org/html/draft-ietf-netconf-restconf-15>

Framework for Large-Scale Measurement of Broadband Performance (LMAP) : <https://datatracker.ietf.org/doc/rfc7594/>

LMAP Measurement Agent (MA) Collection function : <https://datatracker.ietf.org/doc/draft-ietf-lmap-yang/>

Lorido-Botrán, T. , J. Miguel-Alonso , and J. A. Lozano . A review of auto-scaling techniques for elastic applications in cloud environments. *Journal of Grid Computing*, 559–592, 2014.

T. Lorido-Botrán , J. Miguel-Alonso , and J. A. Lozano . *Auto-Scaling Techniques for Elastic Applications in Cloud Environments*, University of Basque Country, Technical Report EHU-KAT-IK-09-12, 2012.

ETSI, *Network Functions Virtualization (NFV) ; Pre-Deployment Testing; Report on Validation of NFV Environments and Services*, ETSI GS NFV-TST 001 V1.1.1, available at: https://portal.etsi.org/webapp/workProgram/Report_WorkItem.asp?wki_id=46009

Mao, M. , J. Li , and M. Humphrey . Cloud auto-scaling with deadline and budget constraints. In *Proceedings of the 11th IEEE/ACM International Conference on Grid Computing*, pp. 41–48, Brussels, Belgium, October 2010.

AT&T Vendor VNF Guidelines, 2017, <https://wiki.onap.org/pages/viewpage.action?pageId=1015852>

ETSI, *Network Functions Virtualization (NFV) ; Virtual Network Functions Architecture*, ETSI GS NFV-SWA 001 V1.1.1, available at: http://www.etsi.org/deliver/etsi_gs/NFV-SWA/001_099/001/01.01.01_60/gs_NFV-SWA001v010101p.pdf

Anand, A. , J. Lakshmi , and S. Nandy . Virtual machine placement optimization supporting performance SLAs. 2013 IEEE 5th International Conference on Cloud Computing Technology and Science (CloudCom), vol. 1. IEEE, 2013, pp. 298–305. Bristol, United Kingdom, December 2–5, 2013.

The Shift to Software

Andreessen, M. 2011. Why software is eating the world. The Wall Street Journal. August 20, 2011.

Bias, R. 2012. Clouds are complex, but simplicity scales; a winning strategy for cloud builders. Cloudscaling Blog. February 29, 2012.

Salus, P. H. 1994. A Quarter Century of UNIX. Addison-Wesley, Reading, MA.

Voegels , 2007. Dynamo: Amazon's highly available key-value store. 21st ACM Symposium on Operating Systems Principles, October 14–17, 2007, Stevenson, WA.

McIlroy, M. D. , E. N. Pinson , and B. A. Tague . 1978. UNIX time-sharing system: Foreword. Bell System Technical Journal 57 (6): 1899–1904.

Brooks, F. and B. Frederick . 1995. The mythical man-month: After 20 years. IEEE Software 12 (5): 57–60.

Raymond, E. S. 2001. The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary. O'Reilly Media, Inc. ISBN 1-56592-724-9.

Ciao , <https://github.com/01org/ciao>

Choose a License Website , <https://choosealicense.com>

FD.IO , <https://fd.io>

Burgess, M. 2000. Principles of Network and System Administration. Wiley, Hoboken, NJ.

Cheswick, B. 1991. An evening with Berferd: In which a cracker is lured, endured, and studied, USENIX Summer Conference Proceedings, Vol. 1. USENIX. San Francisco, California: The Association. p. 163.

What's Next?

<http://adage.com/article/digitalnext/iot-redefine-data-sharing-privacy/300084/>

Why a Connected Data Strategy is Critical to the Future of Your Data , Hortonworks White Paper. <https://hortonworks.com/info/connected-data-platforms/>.

http://www.nytimes.com/2016/04/07/us/politics/homeland-security-dept-struggles-to-hire-staff-to-combat-cyberattacks.html?_r=0

<http://fortune.com/2016/06/15/data-breach-cost-study-ibm/>

<http://www.foxbusiness.com/features/2016/04/27/cyber-attacks-on-small-businesses-on-rise.html>

Meeker Report.

<http://thomsonreuters.com/en/press-releases/2015/05/tr-analysis-shows-concerning-trend-for-global-innovation.html>

Zimmermann H.-D. and Pucihar A. Open Innovation, Open Data, and New Business Models. (September 1, 2015). Available at SSRN: <https://ssrn.com/abstract=2660692>

Hortonworks PR release for new consortium on genomics with Arizona State University, Baylor and Mayo Clinic. <https://hortonworks.com/press-releases/hortonworks-initiates-precision-medicine-consortium-explore-next-generation-genomics-open-source-platform/>

ibid .

<http://www.economist.com/blogs/economist-explains/2016/07/economist-explains-11>

<http://www.forbes.com/sites/louiscolombus/2016/06/04/machine-learning-is-redefining-the-enterprise-in-2016/#33142a475fc0>

AT&T SHAPE Conference Speech.

http://www.slate.com/articles/technology/future_tense/2016/06/microsoft_ceo_satya_nadella_humans_and_ai_can_work_together_to_solve_society.html

ibid . Hans-Dieter Zimmermann and Andreja Pucihar

ibid . Hans-Dieter Zimmermann and Andreja Pucihar