

# **Software-Defined Data Infrastructure Essentials**

**Cloud, Converged, and Virtual  
Fundamental Server Storage  
I/O Tradecraft**

---

# Preface

---

This book follows from my previous books, *Resilient Storage Networks: Designing Flexible Scalable Data Infrastructures* (aka “The Red Book”), *The Green and Virtual Data Center: Enabling Sustainable Economic, Optimized, Efficient and Productive IT Environments* (aka “The Green Book”), and *Cloud and Virtual Data Storage Networking: Your Journey to Efficient and Effective Information Services* (aka “The Yellow, or Gold, Book”). *Software-Defined Data Infrastructures Essentials* is more than a follow-up to these previous works; it looks in various directions—up, down, left, right, current and emerging—and extending into various adjacent data infrastructure topic areas.

*Software-Defined Data Infrastructures Essentials* provides fundamental coverage of physical, cloud, converged, and virtual server storage I/O networking technologies, trends, tools, techniques, and tradecraft skills. Software-defined data centers (SDDC), software data infrastructures (SDI), software-defined data infrastructures (SDDI), and traditional data infrastructures support business applications including components such as a server, storage, I/O networking, hardware, software, services, and best practices, among other management tools. Spanning cloud, virtual, container, converged (and hyper-converged) as well as legacy and hybrid systems, data infrastructures exist to protect, preserve, and serve data and information.

With a title containing terms such as tradecraft, essentials, fundamentals, and advanced emerging topics, some will assume that the material in this book is for the beginner or newbie, which it is. However, being focused on fundamentals and essentials of data infrastructure topics, where there is constant change (some evolutionary, some revolutionary), there are also plenty of “new” essential fundamentals to expand or refresh your tradecraft (i.e., “game”). By “game” I refer to your skills, experience, and abilities to play in and on the data infrastructure game field, which also means being a student of the IT data infrastructure game.

Regardless of whether you are a new student of IT or a specific focus area in a college, university, or trade school, or are going through a career transition, or simply moving from one area of IT to another, this book is for you. Why? Because it converges various data infrastructure topics, themes, trends, techniques, and technologies that can be used in various ways.

If, on the other hand, you are a seasoned pro, veteran industry expert, guru, or “unicorn” with several years or decades of experience, you will find this book useful as well. From web-scale, software-defined, containers, database, key-value store, cloud, and enterprise to small or

## xvi Software-Defined Data Infrastructure Essentials

medium-size business, there are plenty of themes, technologies, techniques, and tips to help develop or refine your server storage I/O tradecraft (game). Thus, there is plenty of new material or new ways to use new and old things even if you are a new or old expert ;).

One of the main and recurring themes in this book is the importance of understanding and recognizing context about server and storage I/O fundamentals. This means gaining as well as expanding (and sharing) your experience with technologies, tools, techniques, and trends in what is also known as your tradecraft.

To some people, the phrase “server storage I/O network tradecraft and fundamentals” will mean how I/O is done, from applications to servers to storage devices. Other people may assume a hardware focus or a software focus. Still others may see a focus on physical machines (PM) or bare metal, virtual, container, cloud, or converged solutions. Context is a recurring theme throughout this book; after all, sometimes words have two or more meanings.

“Tradecraft” refers to the skills and experiences needed and obtained from practicing or being a part of a particular trade. For example, if you were (or are) a successful intelligence spy,

### Key Themes and What You Will Learn in This Book Include:

Theme and Topic	What You Will Learn
Data infrastructure supporting information infrastructure	Server, storage, I/O networking hardware, software, services, resources defined by policies and best practices
SDDC, SDI, SDDI Data infrastructure essentials Converged CI/HCI	Server, storage, I/O networking hardware, software, services, processes, procedures, and best practice fundamentals along with advanced and emerging topics for today and tomorrow
Context matters	Clarify different terms that have various meanings
Decision making, strategy, planning, and management	How you can create a strategy, plan, or manage and decide what you do not know. Gain insight into applications and technology to avoid “flying blind.”
Everything is not the same	Different environments have various attributes and needs
Fast applications need fast data resources	Fast applications need robust underlying data infrastructures: fast servers, storage, I/O resources, and policies
Hybrid home run	Hybrid tools, technology, techniques, and services are the IT home run as they adapt to environments different needs
Knowing your toolbox	Know the tools and technologies in your toolbox as well as techniques on when and how to use them in new ways
Protect, preserve, and serve IT information and assets	Data infrastructure exists in physical, virtual, cloud, and containers to support business and information applications
Server storage I/O tradecraft	Enhancing, expanding, or refreshing your skills and “game”
Software defined	Hardware, software, and services all get defined by software algorithms and data structures (programs) that require some hardware existing somewhere in the technology stack
Various usage examples	While everything is not the same, there are some common usage and deployment scenarios to use as examples
What's in your toolbox	Understanding various hardware, software, and services tools, technologies, and techniques and where to use them

your tradecraft would be the skills, techniques, and experiences of using different tools and techniques to accomplish your job.

Other examples of tradecraft are if you are or were a medical doctor, your tradecraft would be that of providing health care, or if you are an airline pilot, your tradecraft would include flying the plane but also navigating, managing systems, and knowing the various technologies, procedures, routes, and tricks to get your job done. If your field is sales, marketing, finance, or engineering (among others), you possess fundamental tradecraft skills, knowledge, experiences, and practices from those disciplines that can be leveraged while learning new techniques, trends, and topics.

Regarding trends, I often hear people tell me that this year (whenever you read this) is the most exciting time ever, with more happening in the server, storage I/O networking, and related technologies. I agree, as my usual response is that every year for the past several decades has been exciting, each with something new. Since becoming involved with servers, storage, I/O hardware, software, and services, from applications to systems to components, I have found that there is always something new. Some things are evolutionary and prompt a sense of *déjà vu*—of having seen or experienced them in the past. Some are revolutionary, new or first-time experiences, while others can be technolutionary (a blend of new, revolutionary along with evolutionary).

While there are plenty of new things, sometimes those new things get used in old ways; and sometimes old things can get used in new ways. As you have probably heard before, the one thing that is constant is change, yet something else that occurs is that as things or technologies change, they get used or remain the same. A not-so-bold prophecy would be to say that next year will see even more new things, not to mention old things being used in new ways.

For example, many technology changes or enhancements have occurred from the time I started writing this book until its completion. There will be more from the time this goes to the publisher for production, then until its release and you read it in print or electronically. That is where my companion website, [www.storageio.com](http://www.storageio.com), along with my blog, [www.storageioblog.com](http://www.storageioblog.com), and Twitter @StorageIO come into play. There you can further expand your tradecraft, seeing what's current, new, and emerging, along with related companion content to this book.

In terms of buzzwords, let's play some buzzword bingo. Here are some (among others) of the trends, technologies, tools, and techniques that are covered in this book: software-defined, containers, object, cloud, and virtual, physical, virtual server infrastructure (VSI), virtual desktop infrastructure (VDI) and work spaces, emerging, legacy server, micro-servers (and services), along with context matters. This includes server, storage I/O networking hardware as well as software, and services, tips, and techniques.

I also discuss converged, CI, HCI, SDDC, VMware, Hyper-V, KVM, Xen converged and hyper-converged, cluster-in-box or cloud-in-box, Azure and Azure Stack, AWS, OpenStack, Ceph, Mesos, Kubernetes, Hadoop, Hortonworks, and hive, as well as “big data” items. Let's not forget little data, big fast data, structured, unstructured, SQL, and NoSQL Cassandra, MongoDB, and other database or key-value stores. Also scale-out object storage, S3, Swift, backup/data protection as well as archiving, NFV and SDN, MySQL, SQL Server, benchmarking, capacity planning, IoT, artificial intelligence (AI), BC/BR/DR, strategy and acquisition decision making.

In terms of context, there is SRM, which can stand for storage or system resource management and monitoring as well as VMware Site Recovery Manager; telemetry, system resource

## xviii Software-Defined Data Infrastructure Essentials

analysis (SRA) and analytics along with CDM; Splunk, Kudu, Windows Nano, Linux LXC, CoreOS, Docker and other containers, S2D, VHDX, VMDK, VVOL, TBW and DWPD, Reed-Solomon (RS), erasure codes, and LRC, along with mirroring and replication. How about storage class memory (SCM), Flash SSD, NBD, NVM, NVMe, PCIe, SAS, SATA, iSCSI, RoCE, InfiniBand, 100 GbE, IP, MPLS, file systems, NFS, HDFS, SDI, SMB/CIFS, Posix, de-dupe, thin provisioning, compression, and much more.

Say “BINGO” for your favorite buzzwords; that is how buzzword bingo is played!

In case you have not connected the dots yet, the cover ties in themes of using new and old things in new ways, existing and emerging technology spanning hardware, software, services, and techniques. Also, for fun, the cover color combined with my other books represents the primary colors and wavelengths of the rainbow (Red, Green, Blue and Yellow) that are also leveraged in modern high-density fiber optic communications, and high-definition video. With that being said, let’s get into *Software-Defined Data Infrastructure Essentials: Cloud Converged Virtual Fundamental for Server Storage I/O Networking*.

## Who Should Read This Book

*Software-Defined Data Infrastructure Essentials: Cloud, Converged, and Virtual Fundamental Server Storage I/O Tradecraft* is for people who are currently involved with or looking to expand their knowledge and tradecraft skills (experience) of data infrastructures. Software-defined data centers (SDDC), software data infrastructures (SDI), software-defined data infrastructure (SDDI) and traditional data infrastructures are made up of software, hardware, services, and best practices and tools spanning servers, I/O networking, and storage from physical to software-defined virtual, container, and clouds. The role of data infrastructures is to enable and support information technology (IT) and organizational information applications.

Everything is not the same in business, organizations, IT, and in particular servers, storage, and I/O. This means that there are different audiences who will benefit from reading this book. Because everything and everybody is not the same when it comes to server and storage I/O along with associated IT environments and applications, different readers may want to focus on various sections or chapters of this book.

If you are looking to expand your knowledge into an adjacent area or to understand what’s “under the hood,” from converged, hyper-converged to traditional data infrastructures topics, this book is for you. For experienced storage, server, and networking professionals, this book connects the dots as well as provides coverage of virtualization, cloud, and other convergence themes and topics.

This book is also for those who are new or need to learn more about data infrastructure, server, storage, I/O networking, hardware, software, and services. Another audience for this book is experienced IT professionals who are now responsible for or working with data infrastructure components, technologies, tools, and techniques.

For vendors, there are plenty of buzzwords, trends, and demand drivers as well as how things work to enable walking the talk as well as talking the talk. There is a mix of Platform 2 (existing, brownfield, Windows, Linux, bare metal, and virtual client-server) and Platform 3 (new, greenfield, cloud, container, DevOp, IoT, and IoD). This also means that there is a Platform 2.5, which is a hybrid, or in between Platforms 2 and 3, that is, existing and new

emerging. For non-vendors, there is information on different options for usage, and the technologies, tools, techniques, and how to use new and old things in new ways to address different needs.

Even if you are going to a converged or hyper-converged cloud environment, the fundamental skills will help you connect the dots with those and other environments. Meanwhile, for those new to IT or data infrastructure-related topics and themes, there is plenty here to develop (or refresh) your skillsets, as well as help you move into adjacent technology areas.

Student of IT	New to IT and related topics, perhaps a student at a university, college, or trade school, or starting a new or different career.
Newbie	Relatively new on the job, perhaps first job in IT or an affiliated area, as well someone who has been on the job for a few years and is looking to expand tradecraft beyond accumulating certificates of achievement and expand knowledge as well as experiences for a current or potential future job.
Industry veteran	Several decades on the job (or different jobs), perhaps soon to retire, or simply looking to expand (or refresh) tradecraft skills in a current focus area or an adjacent one. In addition to learning new tradecraft, continue sharing tradecraft experiences with others.
Student of the game	Anyone who is constantly enhancing game or tradecraft skills in different focus areas as well as new ones while sharing experiences and helping others to learn.

## How This Book Is Organized

There are four parts in addition to the front and back matter (including Appendices A to G and a robust Glossary). The front matter consists of Acknowledgments and About the Author sections; a Preface, including Who Should Read This Book and How This Book Is Organized; and a Table of Contents. The back matter indicates where to learn more along with my companion sites ([www.storageio.com](http://www.storageio.com), [www.storageioblog.com](http://www.storageioblog.com), and @StorageIO). The back matter also includes the Index.

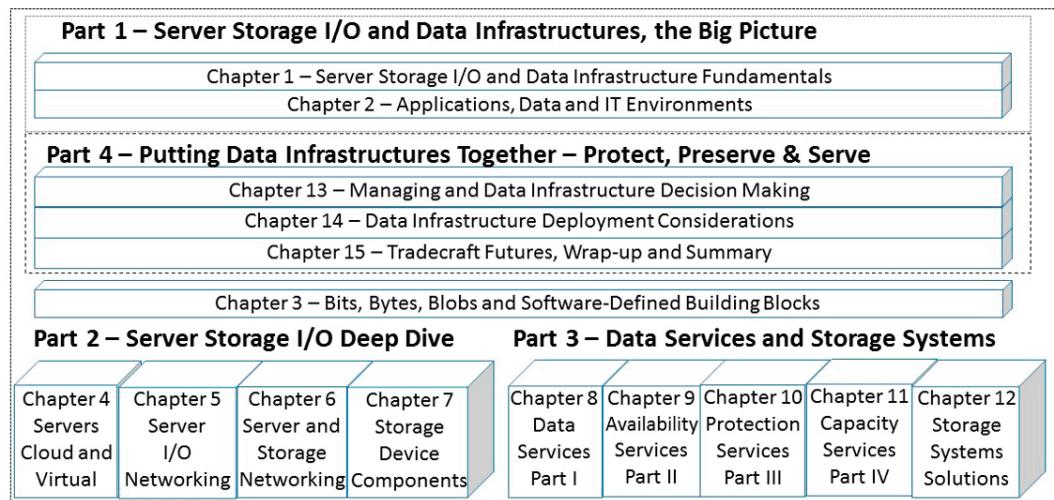
Figure 1 illustrates the organization of this book, which happens to align with typical data infrastructures topics and themes. Using the figure as a guide or map, you can jump around to different sections or chapters as needed based on your preferences.

In between the front and back matter exists the heart of this book: the fundamental items for developing and expanding your server storage I/O tradecraft (experience). There are four parts, starting out with big picture fundaments in Chapter 1 and application, data infrastructures and IT environment items in Chapter 2.

Part Two is a deep dive into server storage I/O, covering from bits and bytes, software, servers, server I/O, and distance networking. Part Three continues with a storage deep dive, including a storage medium and device components, and data services (functionality). Part Four puts data infrastructure together and includes server storage solutions, managing data infrastructures, and deployment considerations, tying together various topics in the book.

## xx Software-Defined Data Infrastructure Essentials

### Contents; Preface: Who Should Read; How Organized; Acknowledgements; About The Author

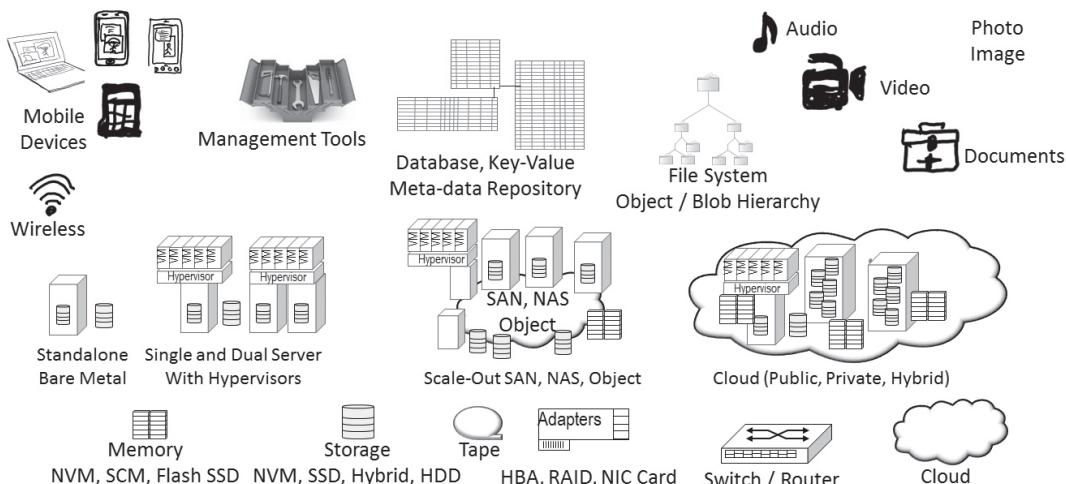


### Appendices A to G; Glossary; Index

**Figure 1** The organization of the book.

In each chapter, you will learn as part of developing and expanding (or refreshing) your data infrastructures tradecraft, hardware, software, services, and technique skills. There are various tables, figures, screenshots, and command examples, along with who's doing what. You will also find tradecraft tips, context matters, and tools for your toolbox, along with common questions as well as learning experiences. Figure 2 shows common icons used in the book.

Feel free to jump around as you need to. While the book is laid out in a sequential hierarchy “stack and layer” fashion; it is also designed for random jumping around. This enables you



**Figure 2** Common icons used in this book.

Preface xxi

to adapt the book's content to your needs and preferences, which may be lots of small, quick reads, or longer, sustained deep reading. Appendix F provides a guide on how to use this book for different audiences who have various focus, interests, and levels of experience.

In case you did not pick up on it, I just described the characteristics of software-defined data infrastructures leveraging server storage I/O technologies for different applications spanning cloud, virtual, container, legacy, and, of course, software—all of which are defined to your needs.