

Brought to you by:



# Modern Infrastructure

for  
**dummies**<sup>®</sup>  
A Wiley Brand

Recognize modern  
storage challenges



Understand storage  
infrastructure options



Get your #GetModern  
assessment



Lawrence C. Miller

Dell EMC #GetModern  
Special Edition



# Modern Infrastructure

Dell EMC #GetModern Special Edition

**by Lawrence C. Miller**

for  
**dummies**<sup>®</sup>  
A Wiley Brand

# Modern Infrastructure For Dummies®, Dell EMC #GetModern Special Edition

Published by: **John Wiley & Sons Singapore Pte Ltd.**, 1 Fusionopolis Walk, #07-01 Solaris South Tower, Singapore 138628, [www.wiley.com](http://www.wiley.com)

© 2017 by John Wiley & Sons Singapore Pte Ltd.

Registered Office

John Wiley & Sons Singapore Pte Ltd., 1 Fusionopolis Walk, #07-01 Solaris South Tower, Singapore 138628

All rights reserved No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted by the Copyright Act of Singapore 1987, without the prior written permission of the Publisher. For information about how to apply for permission to reuse the copyright material in this book, please see our website <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Dell EMC and the Dell EMC logo are registered trademarks of Dell EMC. All other trademarks are the property of their respective owners. John Wiley & Sons Singapore Pte Ltd., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: WHILE THE PUBLISHER AND AUTHOR HAVE USED THEIR BEST EFFORTS IN PREPARING THIS BOOK, THEY MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS BOOK AND SPECIFICALLY DISCLAIM ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. IT IS SOLD ON THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING PROFESSIONAL SERVICES AND NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. IF PROFESSIONAL ADVICE OR OTHER EXPERT ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL SHOULD BE SOUGHT.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact [info@dummies.biz](mailto:info@dummies.biz), or visit [www.wiley.com/go/custompub](http://www.wiley.com/go/custompub). For information about licensing the *For Dummies* brand for products or services, contact [BrandedRights&Licenses@Wiley.com](mailto:BrandedRights&Licenses@Wiley.com).

ISBN 978-0-730-34866-5 (pbk); ISBN 978-0-730-34868-9 (ebk)

Printed in Singapore by C.O.S. Printers Pte Ltd.

10 9 8 7 6 5 4 3 2 1

## Publisher's Acknowledgments

Some of the people who helped bring this book to market include the following:

**Project Editor:** Carrie A. Burchfield

**Acquisitions Editor:** Katie Mohr

**Editorial Manager:** Rev Mengle

**Business Development**

**Representative:** CJ Hwu

**Production Editor:** Siddique Shaik

# Table of Contents

INTRODUCTION .....	1
About This Book .....	2
Foolish Assumptions .....	2
Icons Used in This Book .....	3
Beyond the Book .....	3
Where to Go from Here .....	3
CHAPTER 1: <b>Recognizing Modern Storage Challenges</b> .....	5
Digital Transformation .....	5
Explosive Data Growth .....	6
Constrained Infrastructure Budgets .....	7
Security and Compliance Mandates .....	8
CHAPTER 2: <b>Understanding Storage Infrastructure Options</b> .....	11
All-Flash and Hybrid Storage Arrays .....	11
Converged and Hyper-Converged Systems .....	12
Software-Defined Storage .....	13
Cloud-based Storage .....	14
CHAPTER 3: <b>Addressing Data Protection Requirements</b> .....	17
Snapshots and Replication .....	18
Data Encryption .....	19
Backup and Recovery .....	20
Long-Term Retention .....	21
Global Oversight for Modern Data Management .....	21
Business Continuity and Disaster Recovery .....	22
CHAPTER 4: <b>Exploring Dell EMC All-Flash Storage Solutions</b> .....	25
Dell EMC VMAX All Flash .....	25
Dell EMC XtremIO .....	27
Dell EMC Unity .....	28
Dell EMC SC Series .....	29
Dell EMC Isilon .....	30

<b>CHAPTER 5:</b>	<b>Learning About Dell EMC Converged and Hyper-Converged Solutions.....</b>	<b>33</b>
	Dell EMC VxBlock.....	34
	Dell EMC VxRail.....	35
	Dell EMC VxRack.....	36
<b>CHAPTER 6:</b>	<b>Introducing Dell EMC Software-Defined Storage Solutions.....</b>	<b>37</b>
	Dell EMC ScaleIO.....	38
	Massive scalability.....	38
	Extreme performance.....	38
	Unparalleled flexibility.....	39
	Supreme elasticity.....	39
	Essential features.....	40
	Dell EMC Elastic Cloud Storage.....	41
	Global content repository.....	42
	Modern archive.....	42
	Storage for next-generation applications.....	44
<b>CHAPTER 7:</b>	<b>Getting to Know Dell EMC Data Protection Solutions.....</b>	<b>45</b>
	Dell EMC Data Domain.....	45
	Dell EMC Data Protection Suite.....	46
	Dell EMC Integrated Data Protection Appliance.....	48
<b>CHAPTER 8:</b>	<b>Ten Questions to Assess Your IT Transformation Maturity.....</b>	<b>49</b>
	IT Transformation Assessment.....	49
	Dell EMC #GetModern Assessment.....	51

# Introduction

**M**odern businesses are increasingly embracing digital transformation — leveraging new technologies to reinvent core processes, business models, product offerings, and the customer experience. As businesses embark on the digital transformation journey, they need to modernize their underlying technology infrastructure to enable a more agile, customer-focused, flexible, and innovative digital workplace to drive greater human and operational efficiencies.

Digital transformation initiatives require a modern IT infrastructure to support modern business requirements. Legacy siloed infrastructure prevents organizations from using their data holistically; performance bottlenecks negatively impact user experience, and rigid and/or proprietary architectures limit options when business requirements evolve. The unfortunate result is that IT is increasingly seen as a barrier to the business, rather than an enabler of new business opportunities — and is thus becoming largely irrelevant.

Many modern organizations have fully embraced cloud strategies — with many going so far as to implement a “cloud first” or “cloud only” policy. But the reality is that most organizations will require a hybrid approach for the foreseeable future, consisting of on-premises and public, private, and hybrid cloud environments. To effectively integrate these diverse operating environments, legacy on-premises infrastructure (and architecture) must be modern and agile to support the business with the same speed and agility as in the public cloud.

Modern infrastructure is comprised of compute, storage, and network components. Most organizations maintain regular refresh cycles for their compute and network infrastructure. While some proprietary technology and vendor lock-in issues exist in compute and network infrastructure, these refreshes are relatively seamless. Storage infrastructure refreshes, on the other hand, are far more complex. Storage architectures tend to grow over time as more capacity and performance is needed. As a result, storage infrastructure has become the most expensive capital IT investment in the data center. Vendor lock-in due to the use of proprietary technologies is yet another challenge. Finally,

the sheer volume of data that must be migrated when storage infrastructure is upgraded is daunting and often means significant risk and downtime for the organization.

## About This Book

In this book, you discover the innovative storage technologies that are available to help you modernize your infrastructure and address these challenges to enable your organization's digital transformation journey.

*Modern Infrastructure For Dummies*, Dell EMC #GetModern Special Edition, consists of eight short chapters that explore

- » Storage challenges that today's businesses must address (Chapter 1)
- » Modern storage infrastructure options and technologies (Chapter 2)
- » Important data protection capabilities to look for in a modern storage solution (Chapter 3)
- » All-flash, converged, hyperconverged, and software-defined storage solutions available from Dell EMC (Chapters 4 through 6)
- » Data protection solutions available from Dell EMC (Chapter 7)
- » Ten important criteria for evaluating modern infrastructure solutions (Chapter 8)

## Foolish Assumptions

It's been said that most assumptions have outlived their usefulness, but I assume a few things nonetheless!

Mainly, I assume that you're a technology professional (such as a data center manager or storage architect) or decision maker (such as a CIO or IT director) with some knowledge of data storage challenges and technologies. As such, this book is written primarily for technical readers, but just in case you're not too technical (or

perhaps a bit rusty), I explain any technical terms and concepts that come up in this book.

If any of these assumptions describe you, then this book is for you! If none of these assumptions describe you, keep reading anyway. It's a great book and when you finish reading it, you'll have enough stored knowledge about modern storage infrastructure to be dangerous!

## Icons Used in This Book

Throughout this book, I occasionally use special icons to call attention to important information. Here's what to expect:



REMEMBER

This icon points out information you should commit to your non-volatile memory, your gray matter, or your noggin' — along with anniversaries and birthdays!



TECHNICAL  
STUFF

You won't find a map of the human genome here, but if you seek to attain the seventh level of NERD-vana, perk up! This icon explains the jargon beneath the jargon!



TIP

Tips are appreciated, never expected — and I sure hope you'll appreciate these tips! This icon points out useful nuggets of information.

## Beyond the Book

There's only so much we can cover in 56 short pages, so if you find yourself at the end of this book, thinking "Gosh, this is a great book; where can I learn more?" just go to [www.emc.com](http://www.emc.com).

## Where to Go from Here

With our apologies to Lewis Carroll, Alice, and the Cheshire cat:

"Would you tell me, please, which way I ought to go from here?"



“That depends a good deal on where you want to get to,” said the Cat — err, the Dummies Man.

“I don’t much care where . . . ,” said Alice.

“Then it doesn’t matter which way you go!”

That’s certainly true of *Modern Infrastructure For Dummies*, Dell EMC #GetModern Special Edition, which, like *Alice in Wonderland*, is also destined to become a timeless classic!

If you don’t know where you’re going, any chapter will get you there — but Chapter 1 might be a good place to start! However, if you see a particular topic that piques your interest, feel free to jump ahead to that chapter. Each chapter is written to stand on its own, so you can read this book in any order that suits you (though we don’t recommend upside down or backwards).

I promise you won’t get lost falling down the rabbit hole!

## IN THIS CHAPTER

- » Embracing digital transformation
- » Looking at the exponential growth of data
- » Dealing with storage budget limitations
- » Meeting regulatory compliance requirements

# Chapter 1

# Recognizing Modern Storage Challenges

In this chapter, you learn how modern business trends are creating data storage challenges for legacy storage infrastructure.

## Digital Transformation

As today's businesses undertake digital transformation initiatives to increase business agility, productivity, and competitiveness, IT organizations are struggling to keep pace with legacy infrastructure. Figure 1-1 illustrates how IT transformation maturity directly impacts business outcomes.



TIP

Use the free maturity self-assessment tool at [www.dellemc.com/en-us/storage/it-transformation-assessment/index.htm](http://www.dellemc.com/en-us/storage/it-transformation-assessment/index.htm) to see where your organization stands in the IT transformation spectrum and what you can do to accelerate your progress.

Legacy storage infrastructures built on traditional storage-area networks (SAN) and network-attached storage (NAS) are often proprietary and complex, costly to upgrade, maintain, or replace,

and risky and time-consuming to migrate. These systems create “islands of storage” that effectively silo information that must then be duplicated across multiple systems in order to be consumed by the users and applications that need access to the data.

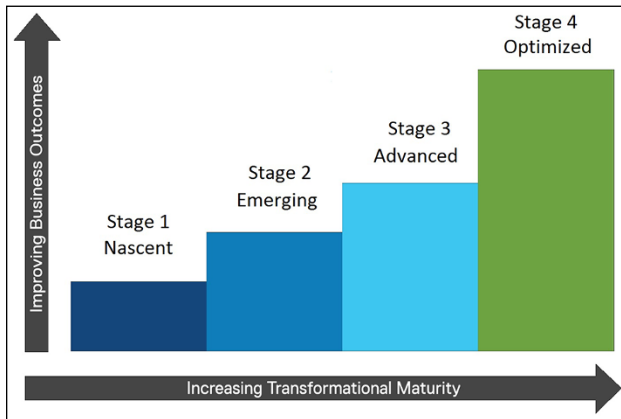


FIGURE 1-1: IT transformation maturity.

The resulting inefficiency is costly, not only in terms of poor utilization of available capacity but also

- » Lost productivity and increased end-user frustration due to slow application performance and storage provisioning processes
- » Missed business opportunities due to slow decision making with incomplete or old information
- » Increased risk of data loss and security breaches due to inadequate protection of redundant data copies, self-service “shadow” IT services, and ineffective data protection tools

## Explosive Data Growth

Data — both structured and unstructured — is everywhere and continues to grow at a stunning pace. Every day, approximately 25 quintillion bytes — 2,500 petabytes (PB) — of new data is generated worldwide, and the total amount of digital data doubles approximately every two years. New applications that leverage big



data, mobile, social, cloud and Internet of Things (IoT) trends are driving this massive growth of new (primary) data.

*Structured data* refers to data that's organized, for example, in a database. *Unstructured data* refers to data that doesn't have a defined model or framework — for example, video files.

Another driver of this exponential data growth is copy data. Copy data is created whenever primary data is copied for secondary purposes, including data protection, analytics, operations, and testing and development. Copy data includes anything from a backup or DR copy to a snapshot for data protection. It also includes analytics copies as well as copies for IT operations like patch testing and sand boxing and copies for test/dev.

This “copy sprawl” consumes storage capacity and increases costs — between 45 and 60 percent of total storage capacity is consumed by copy data at a cost of more than \$50 billion by 2018, according to IDC. And, 82 percent of businesses have at least ten copies of any one production database instance.

This explosion of copy data is being driven by “shadow IT” and the shift to user self-service, which has led to siloes of copy creation across the enterprise without global oversight. Solving this copy data problem requires modern data management, which focuses on optimizing infrastructure efficiency, streamlining operations, and consistently meeting service levels across the enterprise.

## Constrained Infrastructure Budgets

While data storage needs are growing exponentially, storage budgets are increasing only one to five percent annually, creating an ever-widening gap between data growth and storage spending (see Figure 1-2).

Ironically, the cost per gigabyte of storage — both hard disk drives (HDDs) and flash solid-state drives (SSDs) — continues to fall, but data storage requirements and increasingly complex storage architectures are rising more quickly.

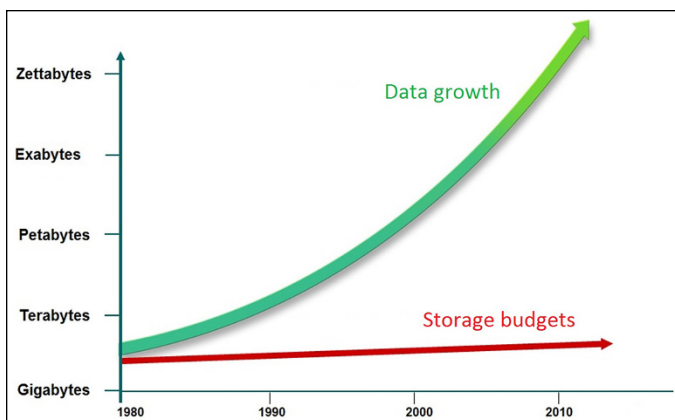


FIGURE 1-2: Storage requirements are outpacing storage investments.

## Security and Compliance Mandates

An ever-growing and dizzying number of security and compliance mandates contribute to the data growth problem, and the urgent need for modern infrastructure to address these challenges and requirements. Different regulatory requirements mandate various — and sometimes conflicting — requirements for data protection, privacy, and retention. Some examples include

- » **Australia:** Telecommunications (Interception and Access) Amendment (Data Retention) Bill 2015
- » **Canada:** Personal Information Protection and Electronic Documents Act (PIPEDA) and Privacy Act
- » **European Union:** Data Retention Directive (Directive 2006/24/EC), General Data Protection Regulation (GDPR)
- » **Japan:** Act on the Protection of Personal Information (APPI)
- » **Singapore:** Personal Data Protection Act (PDPA)
- » **United Kingdom:** Data Retention and Investigatory Powers (DRIP) Act
- » **United States:** Federal Information Security Management Act (FISMA), Gramm-Leach-Bliley Act (GLBA), Health Insurance Portability and Accountability Act (HIPAA), and Sarbanes-Oxley (SOX)

The inevitable result is that organizations everywhere — and in every industry — must protect and retain digital data for ever longer periods — often for as long as 30, 50, or even 100 years. To put that into perspective, most organizations today didn't even produce digital data 30 years ago (if the organization even existed 30 years ago). How confident are you that you can restore data from a backup tape created 6 months ago, let alone 10, 20, or 30 years ago? Do you still have a compatible (and working) tape drive and the right backup/restore software version to use for the restore?



## IN THIS CHAPTER

- » Choosing flash-y or hybrid (storage, not cars)
- » Getting hyper about converged storage
- » Defining software-defined storage
- » Taking your storage to the cloud

# Chapter 2

# Understanding Storage Infrastructure Options

In this chapter, you learn about the storage technology options that comprise a modern infrastructure.

## All-Flash and Hybrid Storage Arrays

Over the past decade, the cost of flash storage technology has rapidly decreased, making both all-flash and hybrid flash storage arrays viable options for most businesses today. Flash storage technology is significantly faster than traditional hard disk drive (HDD) technology. For example, read and write operations in flash are measured in microseconds (compared to milliseconds in HDDs) and input/output operations per second (IOPS) are measured in tens of thousands to millions (compared to hundreds in HDDs).



TECHNICAL  
STUFF

A millisecond (ms) is equal to one-thousandth of a second. A microsecond ( $\mu$ s) is equal to one-millionth of a second.



Changing market dynamics are driving many IT organizations to adopt an “all-flash for primary storage” strategy. According to IDC, all-flash storage can significantly lower total cost of ownership (TCO) in the following ways:

- » **Reduced hardware requirements:** Solid-state drives (SSDs) support workloads with 60 to 90 percent fewer drives than HDDs. For example, a 200,000 IOPS system requires 10 SSDs versus 1,000 HDDs enabling enterprises to achieve desired application performance levels with 5 to 30 percent fewer servers.
- » **Lower software licensing costs:** Fewer servers means fewer expensive software licenses (including operating systems, databases, and applications) to purchase and maintain.
- » **Decreased energy and floor space requirements:** For HDD-based systems requiring at least 80 terabytes (TB) of provisioned storage capacity, all-flash power savings over a four- to five-year life cycle can be \$50,000 to \$100,000 or more. All-flash configurations also typically require less than half a rack versus several full racks for an HDD-based system.
- » **Lower administration costs:** All-flash automates common administrative tasks for ease-of-use that can significantly reduce the time your storage administrators spend addressing performance issues and tuning systems.

Despite the decreasing cost of flash technology, most organizations today deploy a mix of all-flash and hybrid flash storage arrays. The key to hybrid flash storage arrays is intelligent storage management software that can dynamically and predictively move active (“hot”) data to faster flash storage and less active (“warm” or “cool”) data to slower HDDs in a tiered storage architecture.

## Converged and Hyper-Converged Systems

Converged and hyper-converged systems integrate disparate infrastructure elements including servers, data storage devices, networking functions, virtualization, and management software,

in a storage pool or cluster. Intelligent storage management software ensures that data is automatically distributed (balanced) across new nodes that are added to the storage architecture for scale-out capacity and performance.



TIP

Customers using converged and hyper-converged systems can realize many *business* benefits:


- » Improved organizational agility
- » Faster application development
- » Increased innovation
- » Greater employee productivity

*Technical* benefits include the following:

- » Lower costs
- » Greater levels of utilization
- » Reduced downtime

## Software-Defined Storage

Software-defined storage (SDS) uses standard off-the-shelf compute and storage hardware, and performs storage and management functions in intelligent software. SDS essentially separates storage management software from the underlying hardware, in much the same way that virtualization technology abstracts application workloads from the underlying server hardware, pools the storage resources together, and automatically allocates and balances resources based on each application's need (see Figure 2-1).



**FIGURE 2-1:** SDS abstracts, pools, and automates storage resources.

In addition to off-the-shelf hardware, SDS leverages

- » Storage virtualization
- » Massive scale-out architecture
- » Automated policy-driven administration for storage management functions



TIP

The benefits of modern infrastructure leveraging SDS include the following:

- » **Agility:** Legacy storage systems based on proprietary technologies come with a high total cost of ownership (TCO). SDS provides flexible deployment options, often leveraging existing compute and storage components, with massive scale-out capacity.
- » **Control:** Automated, policy-driven storage management seamlessly moves data to the appropriate storage tier (flash or HDD) and manages common storage functions such as disk caching, snapshots, and replication.
- » **Efficiency:** SDS leverages standard off-the-shelf server, storage, and networking hardware and scales out easily and efficiently when additional capacity or performance is needed.

## Cloud-based Storage

Cloud-based storage typically leverages hyper-scale technology, which decouples the compute and storage infrastructure elements so that they can be scaled independently to achieve greater performance or capacity, as needed. Like hyper-converged systems and SDS, cloud-based hyper-scale storage uses standard off-the-shelf servers and storage elements.

At a more granular level, object storage in the cloud is a different model for data storage that can address the challenge of massive growth in data volume, while enabling enterprises to exploit a range of data-driven, value-generating opportunities — such as cloud computing, Big Data analytics, and the Internet of Things (IoT).

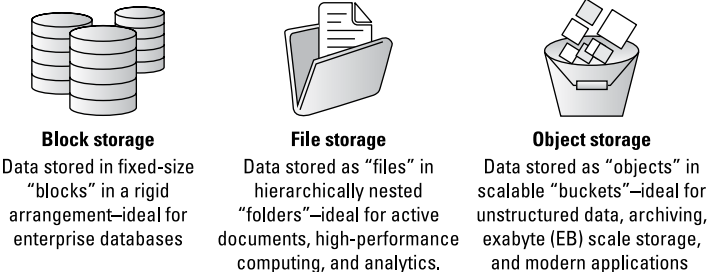
Traditional data storage models include *block* storage and *file* storage. Block storage became increasingly popular as servers in the data center moved from each using its own directly-attached storage device, to sharing a single large storage device or array — a storage area network (SAN).

File storage, utilizes a metaphor from the traditional paper-based office — representing data as files (such as documents, presentations, and spreadsheets), arranged in a series of nested folders.

Although block and file storage have many advantages, the changing nature and explosive growth of data also create new challenges for these storage types. More and more of the data that's created today is unstructured — individual items of information, such as documents, images, and videos. Because this type of data isn't organized in a structured format, such as a database, block storage isn't well suited for it.

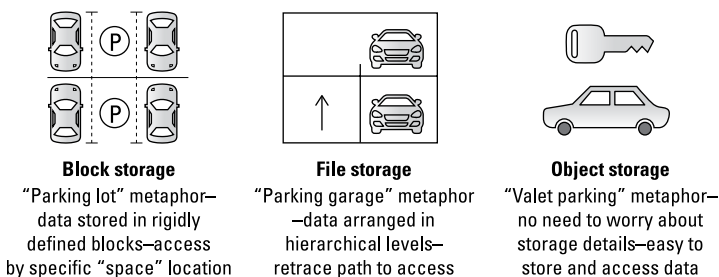
Some of the unstructured data growth is in the form of files and documents — so file storage can be used. However, the sheer volume of data being created and stored is beyond the volume that most file systems were designed to handle. The hierarchical nature of file storage becomes ever more complex to sustain and manage as the amount of data grows.

In the late 1990s, a new kind of storage was developed to address these emerging challenges — object storage. In this storage architecture, data items aren't stored in blocks, files, or folders. Instead, they're stored in flexibly sized containers known as objects. Placing data items in object storage is quick and easy — like putting items in a bucket. This analogy has led to “bucket” becoming the generic term for a volume of object storage (see Figure 2-2).



**FIGURE 2-2:** The three pillars of modern enterprise data storage.

A common analogy used to explain how object storage works is parking a car. Imagine your car is an item of unstructured data and you want to store or “park” it. Block storage is like a parking lot — you can park it in a rigidly defined space, but the available space in the lot is limited. File storage is like a parking garage, arranged in hierarchical levels, but again physically limited. Object storage is like valet parking. You simply drop off your keys, and the valet takes care of parking your car — you don’t need to worry about whether it’s ultimately parked in an adjacent parking lot or garage, or a couple of blocks away, as long as you know it’s safe and secure, and available when you need it (see Figure 2-3).



**FIGURE 2-3:** “Parking your data” — storage models compared.

Large enterprises have been early adopters of object storage architectures and have used it to gain a competitive advantage in their industries. With the latest generation of object storage solutions, companies of all sizes can now realize these advantages. Many organizations initially deploy object storage because of the huge cost savings realized by tiering warm and cold file data to object storage for long-term retention. At the same time, object storage supports new workloads like modern application development, such as cloud-native applications. Ultimately organizations are leveraging object storage to analyze their data and gain new insights about their customers.



**REMEMBER**

Object storage enables cloud technologies that businesses can leverage to gain real competitive advantage.

## IN THIS CHAPTER

- » Creating redundant copies of data with snapshots and replication
- » Ensuring the confidentiality and integrity of data
- » Backing up and restoring data
- » Archiving data and meeting compliance requirements
- » Eliminating copy data sprawl with intelligent copy oversight and analytics
- » Keeping the lights on after an outage or disaster

# Chapter 3

# Addressing Data Protection Requirements

Companies are rethinking their data protection strategies due to exponential data growth, regulatory compliance, strict service-level agreements (SLAs), and shrinking backup windows. Additionally, IT teams face challenges associated with virtualization, cloud and mobile computing, agile development and rapid deployment (DevOps) initiatives, and the need to better protect data throughout the enterprise.

In this chapter, you find out about next-generation data protection technologies to help you address rapidly evolving data protection requirements in the enterprise.

# Snapshots and Replication

Having a redundant copy of data is crucial to protect against outages which can interrupt data availability. Snapshots and replication help to ensure a redundant copy of data is available in the event of an outage or disaster.

A *snapshot* is a point-in-time replica of a storage resource that is usually stored within the storage array or in another array within the data center. Snapshots in modern storage infrastructure can be scheduled to occur at regular intervals. They can require relatively little disk space to store. After the initial snapshot, subsequent snapshots only write the changed blocks of data to storage. However, if the original snapshot is lost, all subsequent snapshots are lost. Also, snapshot versions become difficult to manage and aren't recommended as a sole data protection strategy.

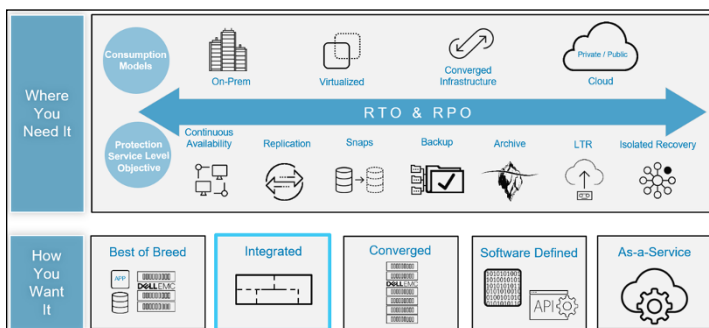
*Replication* protects against a storage system outage by creating a copy of data on a remote system. Replication is a software feature in storage systems that synchronizes data to a remote system within the same site or a different location. Data replication provides redundancy and safeguards against storage system failures. Having a remote disaster recovery (DR) site protects against system and site-wide outages, and provides a remote location to resume production and minimize downtime due to a disaster.

*Synchronous replication* ensures that each block of data written to a storage resource is first saved locally and to a remote image before the write is acknowledged to the host. This ensures that in the event of a disaster, zero data is lost. However, each data write needs to be saved locally and remotely, which adds response time to each transaction. This response time increases as distance increases between remote images. Synchronous replication has a distance limitation based on latency between systems. For synchronous replication, latency between the local and remote system needs to be less than 10 milliseconds.

*Asynchronous replication* is primarily used to replicate data over long distances, but can also be used to replicate resources between storage pools within the same system. Asynchronous replication doesn't impact host input/output I/O latency — host writes are acknowledged once they're saved to the local storage resource. Because write operations aren't immediately replicated to a

destination resource, all writes are tracked on the source. This data will be replicated during the next synchronization.

Recovery time objectives (RTO) and recovery point objectives (RPO) are functions of two dimensions: the *time* it takes to recover data, which is RTO, and the *amount of data loss* that may occur if you need to recover from a protection copy, which is RPO. IT consumption models range from on-premises to cloud, and protection service level objectives (SLOs) can be achieved for all models through a combination of data protection solutions, such as continuous availability, replication, backup, archive, long-term retention, and isolated recovery (see Figure 3-1).



**FIGURE 3-1:** Modern data protection — a comparison of RTO and RPO among protection strategies.



REMEMBER

An organization’s RPO affects the amount and frequency of data that needs to be replicated during each asynchronous replication. RPO is the acceptable amount of data, measured in units of time, which may be lost due to a failure or disaster. RTO is the acceptable amount of time it takes to recover data during a failure or disaster. For example, an RPO of four hours means the most recent backup or replica of data can be no older than four hours. An RTO of four hours means that IT must be able to restore a backup copy of lost data in four hours or less.

## Data Encryption

Data encryption ensures the confidentiality and integrity of data in backups and storage (in flight and at rest). Confidentiality is ensured by rendering the data unreadable if it’s improperly



accessed. Integrity is ensured though cryptographic hashes of data sets. If data is improperly modified — even a single byte — the new hash won't match the original hash.

For data encryption to be effective, it must use a strong encryption algorithm, such as the industry standard 128-bit or 256-bit Advanced Encryption Standard (AES), and the cryptographic key must be safeguarded. Data encryption and decryption can be performed in either software or hardware, but hardware encryption is generally preferred for performance reasons.



REMEMBER

Data encryption is important not only in primary storage but also on backup storage. Most backup storage systems provide inline data encryption.

## Backup and Recovery

Most organizations manage their applications and data in multiple places, from traditional on-premises infrastructure and virtualized environments to hybrid and public clouds. This reality calls for a modern data protection strategy that acknowledges the current landscape and protects data across all consumption models, because it's likely that you'll need to protect multiple environments for the foreseeable future. At the same time, protection and availability should be based on the business value of the data and service levels that align to your business objectives.

Traditional solutions are inefficient because they repeatedly backup everything — including duplicate files and sub-file data segments that exist across servers, desktops, laptops, and different offices. When combined with traditional daily incremental and weekly full backups, the amount of duplicate data is staggering and often leads to decreased performance and increased costs. Searching for backups to recover data in large environments becomes a daunting task.

Modern data protection solutions protect data no matter where it is: on client devices, in on-premises data centers, or in the cloud. They address the spectrum of service level requirements and leverage data deduplication capabilities to reduce storage and other environmental and infrastructure costs. They also provide

advanced search capabilities for finding and restoring data across all consumption models.



REMEMBER

Most organizations will use a combination of data protection methods to fully protect their data against outages, data corruption, deletion, and threats such as cyberattacks and ransomware.

## Long-Term Retention

Organizations that lack robust data archiving processes risk reduced user productivity, along with increased storage and backup costs. Organizations also require fast, easy access to archived email, files, and other content in support of high-volume discovery, regulatory compliance, and secure legal holds.

As I discuss in Chapter 1, increasingly stringent regulatory requirements for long-term retention of business data creates technical challenges in siloed legacy storage infrastructure. Archiving (or vaulting) to tape is still quite common for organizations, but cloud-based archiving of aging or inactive data is becoming an increasingly viable option for long-term retention needs.

Modern storage and backup solutions provide cloud-tiering capabilities for archiving cold data to the private or public cloud for long-term retention. While this option leverages cloud economies, note that public cloud solutions may impact service level objectives during data recovery needs.



REMEMBER

Along the RTO/RPO spectrum, long-term retention (LTR) leveraging object storage provides the best storage economics.

## Global Oversight for Modern Data Management

As companies start to move from traditional server-centric backup strategies designed around centralized control, and progress towards providing more modern data protection designed for self-service, the amount of copy data that is created has exploded. Copy data, which includes copies for data protection, backups,

analytics, IT operations (such as patch testing and sandboxing), and testing and development, can be challenging to track and manage in a self-service model without global oversight.



TIP

Recent IDC research reveals that up to 82 percent of businesses have at least ten copies of any one production database instance.

This proliferation of copy data is being driven by the move to the modern data center, which drives copy creation across the enterprise. This evolution of the data center needs intelligent, consolidated oversight of data and service levels across the business to ensure protection service level objectives (SLO) compliance is met. In addition, solving this gap in this self-service model also requires modern data management, which focuses on optimizing infrastructure efficiency and streamlining operations.

A modern data protection infrastructure includes a copy data management platform as an integral component to manage all data copies, allowing managers to gain insight into the health and optimization of protection infrastructure.

## Business Continuity and Disaster Recovery

Business continuity and disaster recovery planning (BCP/DRP) is a critical activity for every organization. Plans that call for RPOs and RTOs of four days or more are no longer realistic. Businesses can't afford to be out of commission for a week or more due to an outage or disaster — the modern competitive environment doesn't permit it. However, organizations that are straddled with legacy storage infrastructure are struggling to simply meet the normal day-to-day needs of the business, such as keeping up with application performance requirements, providing sufficient storage capacity, and completing regular system and data backups in ever shorter backup windows due to 24/7 business operations. An outage or disaster — for many reasons — is unthinkable.

And no matter what your industry and organization, cyberattacks are on the rise. The number of corporate breaches is growing exponentially and is led by hacking and malware. This type of breach can cripple an organization, leading to revenue loss, negative publicity, and lasting customer distrust.

Modern infrastructure addresses these challenges in both primary and secondary storage, in many cases, enabling RPOs and RTOs of a minute or less, and failover for an entire data center in a matter of minutes — without minimal business disruption. And with the right setup, organizations can create an air-gap to prevent all copies of data being sabotaged from cybercriminals, allowing for an isolated recovery.



- » **Maxing out all-flash performance for mission-critical workloads**
- » **Enabling extreme storage performance in the modern data center**
- » **Learning about midrange storage options**
- » **Getting attached to Dell EMC network-attached storage (NAS) solutions**

# Chapter 4

## Exploring Dell EMC All-Flash Storage Solutions

**A**dopting all-flash technology is key to keeping pace with and exceeding customer demands in today's digital world. While flash is the foundation that drives your modern infrastructure, the secret to outperforming your competition lies in innovative software features and flash-designed architectures that come together to deliver higher performance, lower total cost of ownership (TCO), and better business outcomes for you and your company. In this chapter, you learn about Dell EMC's portfolio of all-flash storage solutions.

### Dell EMC VMAX All Flash

All-flash arrays are accelerating the pace of business transformation as IT professionals search for the most relevant technologies to modernize their operation and drive down operational and capital expenditures. As flash prices rapidly decline, capacity points exceed spinning disk, and data reduction techniques advance, more organizations are evaluating, testing, and deploying all-flash solutions to tackle the most demanding mixed workloads that span across the modern data center.

Dell EMC VMAX All Flash arrays are architected to solve the challenge of embracing a modernized flash-centric data center and hybrid cloud while simultaneously simplifying, automating, and consolidating IT operations. VMAX All Flash is engineered for the latest high-density flash technology and to specifically exploit the rich set of data services of VMAX All Flash. These data services address the new requirements of the modern data center while continuing to deliver reliability and “6 nines” mission-critical availability.

VMAX All Flash enables organizations to dynamically grow, easily share, and cost-effectively manage massive amounts of open systems and mainframe storage. VMAX All Flash maintains consistently high-performance levels, while running thousands of mixed workloads concurrently on a single VMAX All Flash array, to deliver predictable and responsive service — even at massive scale.

VMAX solutions include

- » **Hyper Scale and Consolidation:** Scale to more than 4PB of effective capacity at a fraction of the floor space with high density flash.
- » **Diverse Workload Consolidation:** Support for all primary storage workloads for open systems, mainframe, IBM i, and block and file storage.
- » **Rich Enterprise Data Services:** Enable world-class business continuity and disaster recovery, including copy data management and cloud tiering.
- » **Engineered for Flash:** Maximize your flash investment with greater flash durability and 350 microsecond sustained response times.
- » **Simplified Planning and Deployment:** Grow capacity and performance independently with simple V-Bricks and Flash Capacity Packs.
- » **Deep Ecosystem Integration:** Take advantage of REST APIs, VMware, OpenStack, Oracle, Microsoft, and even integration with leading cloud solutions.

# Dell EMC XtremIO

Dell EMC XtremIO is a purpose-built all-flash array that offers consistently high performance with low latency, storage efficiency with inline, all-the-time data services, rich application integrated copy services, and management simplicity. It provides agility and simplicity for your data center and business through content-aware, in-memory metadata and inline, all-the-time data services, which have made XtremIO the ultimate platform for virtual server and desktop environments and workloads that benefit from efficient copy data management. XtremIO includes many enhancements and new features:

- »» Integrated Copy Data Management (iCDM)
- »» Multi-dimensional scaling
- »» Elegant software-driven performance improvements
- »» Dramatically lower TCO

XtremIO includes inline, all-the-time data services: thin provisioning, deduplication, compression, D@RE, XtremIO Virtual Copies (XVCs), and double solid-state drive SSD failure protection with zero performance impact.

Volumes are always thin-provisioned. XtremIO in-memory deduplication is global—across the entire XtremIO cluster irrespective of the number of X-Bricks in a cluster. This means that XtremIO only writes unique data — data that the entire cluster didn't see in its I/O history — to the SSDs. This inline deduplication not only saves significant capacity but can also improve performance. Inline compression also adds to XtremIO data reduction efficiencies.

XVCs help provision and deploy space-efficient, instant virtual data copies without impacting system performance. Compared to legacy arrays, XtremIO enables average savings of four to twenty times of the amount of physical capacity needed through all these data reduction methods.

XtremIO has developed a flash-based data protection algorithm, known as XtremIO Data Protection (XDP), that offers better than redundant array of independent disks (RAID) 1 performance with



better than RAID 5 capacity savings and protection comparable to RAID 6. It provides dual parity with as little as ten percent capacity overhead.

XtremIO provides consistent performance for all block sizes and all read/write mixes throughout the lifecycle of the array. And iCDM capabilities enable you to create, refresh, and restore thousands of production copies, and run workloads on them, without performance impact or any storage overhead — accelerating business agility.

## Dell EMC Unity

Dell EMC Unity storage solutions are designed for midrange unified storage needs. Dell EMC Unity All-Flash and Hybrid Flash storage platforms optimize solid-state drive (SSD) performance and efficiency, with fully integrated storage area network (SAN) and network-attached storage (NAS) capabilities, and cloud-based storage analytics. Dell EMC Unity solutions include

- » **All-Flash:** Dell EMC Unity All-Flash storage brings high performance and low latency to a broad range of SAN and NAS use cases, including mixed workloads, storage consolidation, server virtualization, and more. Its modern architecture uses high-density SSD technology and inline data reduction. It supports file and block storage, point-in-time snapshots, local and remote data replication, and built-in encryption, and provides deep integration with VMware, Microsoft, and OpenStack ecosystems.
- » **Hybrid Flash:** Dell EMC Unity Hybrid Flash storage systems implement an integrated architecture for block, file, and VMware virtual volumes (VVols), with concurrent support for native NAS, Internet Small Computer Systems Interface (iSCSI), and Fibre Channel protocols. Each system leverages dual storage processors, full 12 gigabit (Gb) serial-attached SCSI (SAS) backend connectivity, and a patented multi-core architected operating environment. Additional storage capacity is added via Disk Array Enclosures (DAEs).
- » **Virtual Storage Appliance (VSA):** Dell EMC Unity VSA allows the advanced unified storage and data management features of the Dell EMC Unity family to be easily deployed

on a VMware ESXi server, enabling you to implement an affordable software-defined solution. This approach provides a low acquisition cost option for hardware consolidation, multi-tenant storage instances, remote/branch office storage, and staging and testing.



TIP

With Dell EMC Unity, you get these benefits:

- » Customer installable and serviceable with CloudIQ
- » Inline compression for balanced cost/performance
- » Unified, transactional, and traditional file workloads with block
- » Converged, software-defined or cloud-enabled DIP upgrades

## Dell EMC SC Series

From Fibre Channel, Internet Small Computer Systems Interface (iSCSI), and Fibre Channel Over Ethernet (FCoE) storage area networks (SANs) to highly-scalable network-attached storage (NAS) solutions, Dell EMC Storage SC Series (formerly Compellent) arrays provide a platform for performance, adaptability and machine-driven efficiency.



TIP

With SC Series, you get these benefits:

- » Best economics in terms of effective \$/GB
- » Self-optimizing architecture, which responds to real-time data use
- » Multi-array federation capability
- » Native auto-failover for business continuity
- » Intelligent deduplication and compression

SC Series software delivers modern features that help you meet aggressive workload demands using the fewest drives necessary. SC Series software includes data progression auto-tiering (which applies diverse media types, including multiple kinds of flash), intelligent block-level compression, deduplication and pervasive “thin” provisioning methods.

With SC Series software, volume movement between arrays is transparent to hosts and optional native auto-failover capability ensures full business continuity during unexpected outages with no additional hardware or software required. Highly modular array components enable easy scalability and hardware reuse in evolving data centers.

The latest SC Series software release, including Storage Center Operating System and Dell Storage Manager (DSM) is a ground-breaking leap forward in the evolution of Dell's SC platform, offering the ideal access point to deploy uniquely adaptable, future-ready storage solutions. DSM offers an intuitive "single-pane-of-glass" interface for one or multiple federated arrays — including unified day-to-day management and cross-platform replication between SC and PS Series (formerly EqualLogic) arrays. With DSM, you can manage your storage at the data center level, ensuring long-term investment protection and efficiency.

## Dell EMC Isilon

Dell EMC Isilon scale-out NAS solutions are powerful yet simple to install, manage, and scale to virtually any size. You can accelerate outcomes with solutions that deliver nine times more IOPS and 18 times more throughput than the closest competitor. With the new ultra-dense design, you can reduce storage footprint by up to 75 percent. Key capabilities include

- » **Performance:** Isilon provides massive room for growth — with up to 68 petabytes (PB) of capacity per cluster. You can scale both capacity and performance in about a minute to meet your specific business needs. You can increase operational flexibility with multiprotocol support.
- » **Efficiency:** Isilon provides you with the management tools to dramatically increase workflow productivity and maximize the value of your enterprise applications and Big Data assets. You can cut costs with over 80 percent utilization and automated tiering.
- » **Protection:** Isilon provides the highest levels of reliability, availability, and serviceability. For fast and efficient data

backup and recovery, you can schedule snapshots as frequently as needed to meet specific recovery-point objectives. For disaster recovery, Isilon provides extremely fast data replication along with push-button failover and fallback simplicity, to further increase the availability of data for mission-critical applications.

- » **Security:** To address compliance and governance requirements, Isilon provides robust security options, including file system auditing and Data at Rest Encryption (DARE) with self-encrypting drives (SEDs). Isilon offers write once, read many (WORM) data protection to prevent accidental or malicious alteration or deletion and help you meet regulatory requirements. You can also leverage role-based access control (RBAC) options and, if needed, create isolated storage pools for specific departments within your organization.

Isilon hardware platforms are built on the Isilon scale-out storage architecture. Platform offerings include flexible product lines that can be combined in a single file system and volume, providing application consolidation tailored for your specific business needs. Isilon All-Flash introduces a revolutionary, highly dense design with four Isilon storage nodes in a compact chassis.

Isilon use case examples include

- » Commercial high-performance computing
- » Electronic design automation
- » Enterprise data lake
- » File archives
- » Financial services
- » Healthcare
- » Large-scale home directories
- » Life sciences
- » Media and entertainment



## IN THIS CHAPTER

- » Discovering Dell EMC VxBlock converged systems
- » Starting small and scaling with Dell EMC VxRail
- » Building the software-defined data center (SDDC) with Dell EMC VxRack

# Chapter 5

# Learning About Dell EMC Converged and Hyper-Converged Solutions

**C**onverged infrastructure brings together disparate infrastructure elements including servers, data storage devices, networking functions, virtualization, management software, orchestration, and applications.

In contrast to traditional converged infrastructure solutions, hyper-converged infrastructure (HCI) is a software-defined architecture with integrated compute, networking, software-defined storage (SDS), and virtualization. It enables compute, storage, and networking functions to be decoupled from the underlying infrastructure and run on a common set of physical resources that are based on industry-standard x86 components. Using hyper-converged infrastructure, customers can start with a small deployment, and then flexibly scale out to support dynamic workloads and evolving business needs.

In this chapter, you learn about Dell EMC's converged and hyper-converged storage solutions.

# Dell EMC VxBlock

Dell EMC VxBlock converged systems integrate enterprise-class technologies — including compute, network, storage, virtualization, and management — into one engineered system. These systems deliver simple IT solutions for a broad range of use cases, significantly reducing costs, time-to-deploy, and on-going management/maintenance-time, thereby enabling your IT staff to spend more time focused on business outcomes and new initiatives, rather than managing the IT infrastructure and keeping the lights on.

VxBlock systems bring together technologies from vendors like Dell EMC, Cisco, VMware, and others. All system elements are pre-integrated, pre-configured, then tested and validated before shipping. Seamless integration allows you to operate and manage your system as a single product, rather than as individual, siloed components. On-going, component-level testing and qualification result in drastically simplified updates and maintenances. The end result is significant time and resource savings throughout the systems life cycle.

VxBlock systems enable you to create a modern data center with all-flash storage, including Dell EMC Unity, Dell EMC XtremIO, and Dell EMC VMAX options (discussed in Chapter 4). All-flash VxBlock systems are ideal for mixed-workload consolidation, applications that require extreme performance and scale, and mission-critical applications.

With multivendor solutions and reference architectures, considerable time and resources are devoted to sourcing, integrating, testing, validating, and ongoing maintenance.

The Dell EMC Release Certification Matrix (RCM) streamlines software release planning and ongoing configuration management across all converged infrastructure components. Dell EMC Quality Assurance tests for interoperability of hardware and software to ensure that your system and data center dramatically reduce downtime due to updates and scheduled maintenance. This unique experience spans across all VxBlock systems.

The intelligence, automation, and visualization in Dell EMC Vision Intelligent Operations facilitates standardized, repeatable IT processes — making it easier to keep your data center/hybrid

cloud environment healthy, stable, optimized, and secure. The software manages compute, network, storage, and virtualization components together as a single system, and multiple systems as a single pool of resources. Functions include health, RCM compliance, and security compliance management.

Dell EMC Vscale Architecture enables scale-up and scale-out data centers by combining a modular grow-as-needed architectural design with the flexibility to add resources incrementally through compute and storage technology extensions. Through the Vscale Fabric, a scalable spine-leaf network fabric, you can connect multiple systems and modular components to create a shared pool of resources.

## Dell EMC VxRail

To support the business as an enabler, your IT organization must break down silos of complexity and streamline operations to keep pace with today's on-demand business culture and users.

Dell EMC VxRail accelerates and simplifies IT through standardization and automation, allowing you to focus on

- » Getting infrastructure operational as quickly as possible while leveraging existing investments and skill sets
- » Enabling (not impeding) business applications
- » Making infrastructure growth transparent to the needs of the application and users

Industry analyst firm, Enterprise Strategy Group, agrees that speed and agility are critical, stating, "A modern data center can't have infrastructure provisioning take weeks or months and expect to stay competitive."

VxRail is a fully integrated, preconfigured, and tested hyper-converged infrastructure (HCI) appliance powered by VMware vSAN. VxRail provides a simple, cost effective hyper-converged solution that solves a wide range of challenges and supports most applications and workloads. It features purpose-built platforms that deliver data services, resiliency, and quality of service (QoS), enabling faster, better, and simpler delivery of virtual desktops, business-critical applications, and remote office infrastructure.



A wide range of platforms and configure-to-order hardware are designed to address any use case. Single node scaling and storage capacity expansion provide a predictable, “pay-as-you-grow” approach for future scale up and out as your business and user requirements evolve. With different optimized models for different use cases and workloads, customers can quickly choose and deploy their hyper-converged infrastructure based on their needs.

## Dell EMC VxRack

Dell EMC VxRack System 1000 consists of hyper-converged rack-scale engineered systems, with integrated networking, to achieve the scalability and management requirements of traditional and cloud native workloads. The VxRack family is designed to enable organizations to quickly deploy Infrastructure as a Service (IaaS) and/or private cloud architectures. The VxRack System tightly integrates the hardware with the software and management layer.

The result is a fully tested, pre-configured, hyper-converged system with automated provisioning, simplified management and robust reporting capabilities at data center and service provider scale. The VxRack System supports the deployment of a variety of application workloads allowing IT to rapidly deliver new services while improving overall agility and efficiency. VxRack System options include

- » **VxRack FLEX:** A rack-scale hyper-converged infrastructure built on Dell PowerEdge servers and Dell EMC ScaleIO software-defined storage (discussed in Chapter 6). Direct-attached storage (DAS) is virtualized into a shared pool of block storage, similar to a storage area network (SAN).
- » **VxRack SDDC:** For the software-defined data center (SDDC), VxRack SDDC is a turnkey hyper-converged solution powered by VMware Cloud Foundation. These systems consist of pre-loaded software and compute, storage, and network components in a hyper-converged stack co-engineered by Dell EMC and VMware.

- » **Scaling out with Dell EMC ScaleIO software**
- » **Enabling cloud-scale storage with Elastic Cloud Storage**

# Chapter **6**

## **Introducing Dell EMC Software-Defined Storage Solutions**

**T**he benefits of server virtualization are well understood in the modern data center. By abstracting, pooling, and automating compute resources, companies have achieved significant savings. In the past decade, web-scale companies like Google and Amazon have demonstrated the ability to operate data centers with ruthless efficiency. Dell EMC ScaleIO Software-Defined Storage applies the principles of abstraction, pooling, and automation to local storage in standard x86 servers, creating a high-performance shared storage service without the need for conventional storage arrays. ScaleIO also enables a highly efficient data center operating model. Combined, these capabilities deliver a radical reduction in storage life cycle costs.

In this chapter, you discover Dell EMC's software-defined storage (SDS) solutions for on-premises and cloud environments.

# Dell EMC ScaleIO

Dell EMC ScaleIO is software that creates a server-based storage area network (SAN) from local application server storage to deliver flexible and scalable performance and capacity on demand. It converges storage and compute resources of commodity hardware into a single-layer architecture — aggregating capacity and performance, simplifying management, and scaling to thousands of nodes.

As an alternative to a traditional SAN infrastructure, ScaleIO combines hard disk drives (HDDs), solid-state drives (SSDs), and Peripheral Component Interconnect Express (PCIe) flash cards, to create a virtual pool of block storage with varying performance tiers. In addition, it provides enterprise-grade protection, multi-tenant capabilities, and add-on enterprise features such as quality of service (QoS), thin provisioning, and snapshots (for more info, see Chapter 3). ScaleIO is hardware-agnostic, supports physical and/or virtual application servers, and delivers significant total cost of ownership (TCO) savings versus a traditional SAN.

## Massive scalability

ScaleIO is designed to massively scale from three to thousands of nodes. Unlike most traditional storage systems, as the number of storage devices grows, so do throughput and input/output operations per second (IOPS). The scalability of performance is linear with regard to the growth of the deployment. Whenever the need arises, additional storage and compute resources (that is, additional servers and/or drives) can be added modularly so that resources can grow individually or together to maintain balance. Therefore, storage growth is always automatically aligned with application needs.

## Extreme performance


Every server in the ScaleIO cluster is used in the processing of I/O operations, making all I/O and throughput accessible to any application within the cluster. Such massive I/O parallelism eliminates bottlenecks. In addition, throughput and IOPS scale in direct proportion to the number of servers and local storage devices added to the system, improving cost and performance rates with growth. Performance optimization is automatic; whenever rebuilds and

rebalances are needed, they occur in the background with minimal or no impact to applications and users. The ScaleIO system autonomously manages performance hot spots and data layout.

## Unparalleled flexibility

ScaleIO provides two flexible deployment options (see Figure 6-1):

- » **Storage-only:** The application and storage (“two layers”) are installed on separate servers in the ScaleIO cluster. This provides efficient parallelism and no single points of failure.
- » **Hyper-converged:** The application and storage are installed on the same servers in the ScaleIO cluster. This creates a single-layer architecture and provides the smallest footprint and lowest cost profile.



**FIGURE 6-1:** Traditional SAN architecture and ScaleIO storage-only and hyper-converged architectures.

ScaleIO provides unmatched choice for these deployment options. The software is infrastructure agnostic, making it a true software-defined storage product. It can be used with mixed server brands, operating systems (physical and virtual), and storage media types (HDDs, SSDs, and PCIe flash cards). In addition, OpenStack commodity hardware can be used as storage and compute nodes.

## Supreme elasticity

With ScaleIO, storage and compute resources can be increased or decreased whenever the need arises. The system automatically rebalances data on the fly with no downtime. Additions and removals can be done in small or large increments. No capacity planning or complex reconfiguration due to interoperability constraints is required, which reduces complexity and cost.

Also, the ScaleIO system reconfigures itself as the underlying resources change; data is rearranged and spread evenly on the servers to optimize performance and enhance resiliency. All of this happens automatically, without operator intervention, eliminating the need for costly and disruptive data migrations.

## Essential features

ScaleIO offers a set of features that enable complete control over performance, capacity, and data location. For both private cloud data centers and service providers, these features enhance system control and manageability — ensuring that QoS performance requirements are met. With ScaleIO, you can limit the amount of performance — IOPS or bandwidth — that selected tenants can consume. The limiter allows for resource distribution to be imposed and regulated, preventing application “hogging” scenarios. CloudLink encryption software can be used to provide added security for sensitive data. ScaleIO offers instantaneous, writeable snapshots for additional data protection.

For improved read performance, dynamic random access memory (DRAM) caching enables you to improve read access by using ScaleIO Data Server (SDS) RAM. Fault sets — a group of SDS nodes that are likely to go down together — can be defined to ensure data mirroring occurs outside the group, improving business continuity. You can also create volumes with thin provisioning, providing on-demand storage as well as faster setup and startup times.

ScaleIO also provides multi-tenant capabilities via protection domains and storage pools. *Protection domains* allow you to isolate specific servers and data sets. This can be done at the granularity of a single customer so that each customer can be under a different SLA. *Storage pools* can be used for further data segregation, tiering, and performance management. For example, data that is accessed very frequently can be stored in a flash-only storage pool for the lowest latency, while less frequently accessed data can be stored in a low-cost, high-capacity pool of spinning disks.



TIP

Whether you prefer a “build it yourself” or “turnkey” solution for your software-defined storage, ScaleIO provides multiple consumption methods that will help you reach your goals quickly and efficiently. You can consume ScaleIO in three different models:

- » **ScaleIO software:** Delivers the maximum flexibility and installs directly on your x86 standard server hardware and infrastructure.
- » **ScaleIO ReadyNode:** Bundles Dell PowerEdge servers and ScaleIO software, enabling you to quickly deploy a fully architected software-defined, scale out, server SAN (see Chapter 5). **Note:** ReadyNodes are only available with Dell PowerEdge servers.
- » **VxRack Flex:** A fully integrated, turnkey, hyper-converged solution which leverages ScaleIO, enabling you to hit the ground running (I discuss this in Chapter 5).

## Dell EMC Elastic Cloud Storage

The unabated growth in unstructured content is driving the need for a simpler storage architecture that can efficiently manage billions and trillions of files and accelerate the development of cloud, mobile, Internet of Things (IoT), and Big Data applications while reducing both storage overhead and cost. In order to fulfill these requirements, IT organizations and service providers have begun to evaluate and utilize low cost, commodity and open-source infrastructures. These commodity components and standard open technologies lower storage costs, but the individual components provide lower performance and reliability — and also require operational expertise to ensure serviceability of the infrastructure.

Dell EMC Elastic Cloud Storage (ECS) provides a complete software-defined object storage platform designed for today's cloud-scale storage requirements. ECS provides the simplicity and low cost benefits without the risk, compliance, and data sovereignty issues of the public cloud. ECS benefits include

- » **Cloud-scale economics:** 60 percent TCO savings versus public cloud services
- » **Simplicity and scale:** Single global namespace, unlimited apps, users, and files
- » **Universal accessibility:** Support for object, file, and Hadoop Distributed File System (HDFS) all on a single platform

- » **Faster app development:** Application programming interface (API) accessible storage and strong consistency accelerates cloud apps and analytics
- » **Turnkey cloud:** Multi-tenancy, self-service access, and metering capabilities

## Global content repository

Unstructured data growth, both in structured and unstructured content are driving the need to store large files (such as images and videos) in high-cost, siloed storage systems. Traditional infrastructure often isolates data, making it difficult and costly to exchange and manage content. This type of environment doesn't scale efficiently and economically, forcing customers to seek out a solution that delivers the benefits of both public and private clouds. ECS enables any organization to consolidate multiple storage systems and content archives into a single, globally accessible and efficient content repository that can host countless applications. Key capabilities include

- » **Anywhere access to content:** ECS multi-site, active-active architecture, single global namespace, and universal accessibility (object, file, and HDFS) enables anywhere access to content from any application or device. Buckets of data span site, allowing reads and writes to and from any location across the globe. ECS provides strong consistency semantics which simplifies application development and facilitates anywhere access to data. ECS also provides a geo-caching capability which identifies multi-site access patterns and caches data at the location where it's being accessed most frequently.
- » **Remote office back office (ROBO) connectivity:** ECS, in conjunction with Dell EMC CloudArray, enables streamlining of remote offices. Frequently accessed SAN and NAS data is cached locally while archive data is tiered to ECS back in the data center, enabling organizations to reduce equipment sprawl at remote offices.

## Modern archive

Moving cold data off Tier 1 storage for archive and long term retention purposes in an important activity for every organization today. Public cloud storage services can offer flexible virtual

machine (VM) deployments, easy to manage platforms, and a low upfront cost to start. However, concerns around security, compliance, vendor lock-in, and unpredictable costs can quickly become a major issue. ECS provides all the benefits of public cloud in an on-premises, cloud-scale object archiving solution for inactive workloads.



TIP

Key capabilities and benefits include

- » **Multi-site, active-active architecture and access:** ECS features a truly geo-efficient architecture that stores, distributes, and protects both local and geographically remote data. This eliminates any single point of failure and provides a seamless failover from site to site with no impact to the business. ECS automatically maximizes throughput, maintains high availability and data durability, and increases capacity and the reliability of applications.
- » **Tape replacement:** ECS provides a series of dense configurations (D-Series) designed to replace tape archives for backup, long-term retention (LTR), and near-line (NL) purposes. Cold or inactive content can still deliver value and ECS allows you to monetize this data and provide an active archive with the same scalability and low cost benefits of tape-based solutions, but without the operational challenges, lack of IT agility, and reliability concerns.
- » **Cloud-scale economics:** ECS provides an easy to manage, globally accessible archive delivered at a lower cost than public cloud storage providers. Policy-based management allows you to seamlessly move content off Tier 1 storage to ECS to better optimize performance. Geo-distribution and geo-caching provide high availability on the system and lower storage overhead as you scale. ECS also provides instant access to your archived content so you don't have to wait days or weeks for data retrieval.
- » **Seamless integration:** ECS integrates seamlessly with existing Dell EMC storage platforms and investments through CloudPools, CloudBoost, and Data Domain Cloud Tier. Integrated tiering software reduces the risk and complexity of implementing data tiering of your cold data from primary storage to ECS. ECS also supports protection modes for litigation hold as well as U.S. Securities and Exchange Commission (SEC) 17 a-4(f) and Federal Information Processing Standards (FIPS) 140 compliance.



## Storage for next-generation applications

Enjoy quick and easy access to storage from your next-generation applications — mobile, cloud, big data, and IoT with ECS. Like public cloud storage, ECS is simple to manage as a singular entity from a central location with its globally distributed infrastructure. With “Exabyte scale” storage, IT organizations can manage billions of multi-size objects with ease. ECS is optimized to support ingest of and access to both small and large files with strong global consistency.

## IN THIS CHAPTER

- » Exploring Data Domain data protection solutions
- » Meeting the Data Protection Suite family
- » Introducing the Integrated Data Protection Appliance

# Chapter 7

# Getting to Know Dell EMC Data Protection Solutions

In this chapter, you discover the Dell EMC data protection solutions, including Dell EMC Data Domain, Dell EMC Data Protection Suite, and cloud backup and protection options.

## Dell EMC Data Domain

Dell EMC Data Domain systems reduce the amount of disk storage needed to retain and protect data by ratios of 10 to 30 times and greater. With throughput of up to 68 terabytes (TB) per hour, Data Domain systems make it possible to complete more backups in less time and provide faster, more reliable restores.

The Data Domain Operating System (DD OS) is the intelligence that powers Data Domain. It provides the agility, security and reliability that enables the Data Domain platform to deliver scalable, high-speed, and cloud-enabled protection storage for backup and recovery, archive, and disaster recovery.

The key to DD OS performance while minimizing disk requirements is the Stream-Informed Segment Layout (SISL) scaling architecture. Specifically, SISL leverages the continued advancement of CPU performance to continuously increase Data Domain system performance by minimizing disk accesses required to de-duplicate data. SISL de-duplicates data by identifying duplicate data segments in memory, which minimizes disk usage. This enables Data Domain throughput to be CPU-centric, not “spindle bound.”

Data Domain solutions include

- » **Entry-level to midrange Data Domain Systems:** Protection storage for small to midsize environments
- » **Enterprise Data Domain Systems:** Protection storage for enterprise environments
- » **Data Domain Virtual Edition:** Software-defined protection storage for use cases such as remote office/branch offices and deployments in the cloud
- » **Data Domain Software:** Delivers powerful capabilities to get more value from your investment.
  - **Data Domain Boost:** Accelerates backups by up to 50 percent through advanced integration with leading backup and enterprise applications.
  - **Data Domain Cloud Tier:** Natively tiers deduplicated data to the public, private or hybrid cloud for long-term retention.
  - **Data Domain Replicator:** Replicates from up to 540 remote sites to a single system and reduce bandwidth requirements by up to 99 percent.
  - **Data Domain Extended Retention:** Provides cost effective, internal tiering for long term retention of backups.
  - **Data Domain Cloud Disaster Recovery:** Provides low-cost disaster recovery to the cloud.

## Dell EMC Data Protection Suite

The Dell EMC Data Protection Suite simplifies data protection choices, making it easier than ever to access best-of-breed backup, recovery, and archive solutions that fit your specific

needs. The Data Protection Suite provides comprehensive data protection to enterprise organizations of any size, for a broad range of applications from edge to core to cloud.

Protecting data using technology from replication to snapshot to backup and archive, Data Protection Suite delivers coverage across all consumption models. With purpose-built solutions built to meet the needs of all types of organizations and data protection environments, the Data Protection Suite protects data no matter where it is and against whatever happens. Data Protection Suite offerings include

- » **Data Protection Suite Enterprise Edition:** Comprehensive protection including continuous replication, snapshot-based backup, de-duplicated and traditional backup and recovery, and archive. Enterprise Edition meets the needs of organizations of all sizes via a single, comprehensive offering.
- » **Data Protection Suite for Backup:** Complete cloud-enabled backup and recovery across the broadest technology landscape. Designed for use in physical and virtual environments and supporting many different deployment models and protection technologies including deduplication backup, backup to disk, snapshot-based backup, and backup to tape, Data Protection Suite for Backup provides the flexibility to mix and match components to quickly optimize performance and data protection while also reducing costs.
- » **Data Protection Suite for Applications:** Provides unparalleled efficiency to meet stringent service-level agreements (SLAs) on mission-critical applications. Data Protection Suite for Applications provides up to 20 times faster backup by enabling direct backup to Data Domain from either primary storage or the application server, and up to 10 times faster recovery. It empowers application owners and database administrators to backup directly to Data Domain through native application utilities (like Oracle Recovery Manager, RMAN) and also allows them to discover, automate, and optimize copies for superior copy data management. Data Protection Suite for Applications reduces the hidden risk and cost associated with distributed, self-service copy creation by providing global oversight of the copy data ecosystem. Finally, it eliminates or significantly reduces the impact on application servers, because little or no data flows through the application server.

- » **Data Protection Suite for VMware:** End-to-end data protection for VMware-based environments including backup and recovery, continuous replication, monitoring and analysis, and search capabilities. This offering is a software-only data protection solution that delivers simplified deployment and administration within virtualized and cloud environments.
- » **Data Protection Suite for Archive:** Best-in-class archive and eDiscovery solutions for email, files, and collaboration systems. As a part of the Data Protection Suite, this solution gives organizations full ownership and control over their information. Organizations have the ability to reduce costs by reclaiming valuable primary data storage space, optimizing server and operational performance, and meeting company compliance regulations, and eDiscovery and litigation needs.

## Dell EMC Integrated Data Protection Appliance

Dell EMC Integrated Data Protection Appliance (IDPA) is a pre-integrated, turnkey solution that's simple to deploy and scale, provides comprehensive protection for a diverse application ecosystem, and comes with native cloud tiering for long-term retention. IDPA combines protection storage, protection software, search, and analytics to reduce the complexity of managing multiple data silos, point solutions, and vendor relationships.

With IDPA, you reduce time-to-protection and achieve faster time-to-value with up to ten times faster deployment. You become more agile with a solution that's fast — data is protected more efficiently and can be recovered quicker — and that's reliable with industry-proven data invulnerability architecture. And, you get the high value and low TCO you expect from Dell EMC — the number one data protection provider.

## IN THIS CHAPTER

- » Determining your starting point in digital transformation
- » Planning the next steps in your digital transformation journey

# Chapter 8

# Ten Questions to Assess Your IT Transformation Maturity

In this chapter, I help you plan your organization's digital transformation journey.

## IT Transformation Assessment

To know where you're going, it's important to know first where you are on your journey to digital transformation.

Dell EMC offers a free assessment that shows where your organization stands in the IT transformation spectrum. Based on your results, you see next steps recommended to accelerate your progress.



TIP

You can access the free Dell EMC online assessment at [www.dell EMC.com/en-us/storage/it-transformation-assessment/index.htm](http://www.dell EMC.com/en-us/storage/it-transformation-assessment/index.htm).

**The IT transformation assessment asks the following questions about your organization:**

- » Does your IT organization enable end-users to provision on-premises IT resources (such as VMs, storage capacity, and so on) in a self-service fashion?
- » To what extent has your development organization adopted formal DevOps principles and best practices?
- » How frequently is the IT organization and the outcomes it delivers evaluated by C-suite business executives or the board of directors?
- » Who does the most senior IT executive at your company report to (for example, CEO or equivalent, President/COO, CFO, SVP/VP)?
- » How much progress has your IT organization made towards running applications and services on-premises with the same capabilities afforded by public cloud providers?
- » Does your organization use converged or hyper-converged infrastructure platforms to support any of its on-premises applications? If so, approximately what percentage of applications?
- » What's your company's perspective on software-defined storage (SDS)?
- » What's your company's perspective on software-defined networking (SDN)?
- » Does your organization use storage systems that utilize scale-out architectures to support any on-premises applications? If so, what approximate percentage of applications?
- » Does your organization use flash/solid-state storage to support any on-premises applications? If so, how does your organization typically deploy flash/solid-state storage?

**After completing the assessment, you can download a summary of your results with a custom action plan and an ESG benchmark study outlining the maturity model and IT industry findings.**

# Dell EMC #GetModern Assessment

With the insights you gain from the IT Transformation Assessment, your next step is to schedule a Dell EMC #GetModern assessment. A #GetModern assessment is a consultative process that produces a comprehensive proposal which demonstrates the value of a modern infrastructure. The #GetModern assessment helps you identify infrastructure throughout your data center and readily identify performance and capacity bottlenecks. Enterprises need to modernize, automate, and transform their data centers, regardless of their current level of transformation maturity. A #GetModern assessment can help you get started with an in-depth look at your storage needs. Visit [www.dell EMC.com/getmodern](http://www.dell EMC.com/getmodern) to get started.



REMEMBER

A modern infrastructure is a critical key to digital transformation for all organizations, enabling flexible and agile “cloud-like” self-service in your on-premises data centers. This capability empowers business users to leverage IT to enable digital transformation to adapt to rapidly changing market conditions and take advantage of new and evolving business opportunities.



# Modern Storage Deserves a Modern Network

**BROCADE**

Gartner reports that modern storage networks are becoming a critical priority thanks to the rapid adoption of high-throughput, solid-state, and flash storage.

Gartner also states, "Fibre Channel will remain as the data center storage protocol of choice for the next decade."


See why Gen 6 Fibre Channel SAN solutions, now available from Dell EMC as Connectrix B-Series products, are the right choice to maximize your storage ROI.

\*Gartner "The Future of Storage Protocols" Valdis Filks, Stanley Zaos, 29 June 2016



Flash storage demands maximum network speed and throughput

## Workload Completion Time: Less Is More



## ~4x Workload Completion Time Improvement with 32 GFC\*\*

\*\*As compared to 8 GFC. Emulex/Broadcom TPC-H benchmark Testing with Dell EMC PowerEdge Servers and XtremIO all-flash storage. <https://www.emc.com/collateral/software/white-papers/h15688-xtremio-poweredge-broadcom-brocade-tech-brief.pdf>



Learn more and start modernizing with a free SAN Health assessment at <https://dell.emc.brocade.com>

With 20+ years of partnership, Dell EMC and Brocade deliver industry leading Fibre Channel storage networks that help modernize, automate and transform your IT infrastructure.

Brocade  
**GEN6**  
FIBRE CHANNEL



Greater Availability



Predictable Performance



More Automation



Deep Visibility



Mass Scalability



Increased Security

© 2017 Brocade Communications Systems, Inc. All Rights Reserved.

These materials are © 2017 John Wiley & Sons, Ltd. Any dissemination, distribution, or unauthorized use is strictly prohibited.



## Accelerate Your Technology Adoption and Optimization...With Expert Help

To remain relevant, IT groups need to be agile when introducing new technologies. Once new technologies are in place, it's critical that staff are able to take advantage of new features with little or no ramp time.

### **Let Dell EMC Residency Services help.**

Imagine technology experts to optimize configurations, processes and procedures, and share knowledge with your IT staff. Transition to your new platform capabilities quickly, while keeping your data center running at its peak. With total control of these experts' priorities and time, you'll have the flexibility to adjust resources and budget to your evolving business needs.

And every resident is backed by Dell EMC expert engineers, solution architects and implementation specialists—supplementing their own skills with an arsenal of proven tools, ITIL-aligned processes, and industry best practices.

Contact your local Dell EMC representative or authorized reseller...Or, watch this: [http://bit.ly/dellemc\\_residency](http://bit.ly/dellemc_residency)



# Modernize your technology infrastructure

Modern infrastructure is key to keeping pace with and exceeding customer demands in today's digital world. Many modern organizations are taking a hybrid approach to IT, consisting of on-premises, private, public, and hybrid cloud environments. To effectively integrate these diverse operating environments, IT infrastructure (and architecture) must be modern and agile and deliver higher performance, lower TCO, and better business outcomes for you and your company.

## Inside...

- Modern infrastructure requirements
- All-Flash storage solutions
- Converged and hyper-converged solutions
- Software-defined storage solutions
- Data protection solutions
- Assess your IT transformation maturity

**DALLEMC**

**Lawrence C. Miller** has worked in information technology for more than 25 years. He is the co-author of *CISSP For Dummies* and has written more than 100 other *For Dummies* books.

Go to **Dummies.com**<sup>®</sup>  
for videos, step-by-step photos,  
how-to articles, or to shop!

**for  
dummies**<sup>®</sup>  
A Wiley Brand

ISBN: 978-0-730-34866-5  
Not for resale

# **WILEY END USER LICENSE AGREEMENT**

Go to [www.wiley.com/go/eula](http://www.wiley.com/go/eula) to access Wiley's ebook EULA.