

Making Everything Easier!™

EMC Isilon Special Edition

Big Data Storage

FOR
DUMMIES[®]
A Wiley Brand

Brought to you by

EMC²[®]

Will Garside
Brian Cox



Big Data empowers organisations to gain valuable insight into their markets and their customers. This enables these organisations to make better design, safety and support decisions and ultimately deliver more tailored, better products and services. However, the foundation for getting this insight is derived through a solid infrastructure of a highly scalable, easily managed and well-protected scale-out storage solution.



About EMC Isilon

Isilon, a division of EMC, is a global leader in scale-out storage. Isilon delivers powerful yet simple solutions for enterprises that want to manage their data, not their storage. Isilon products are simple to install, manage and scale — at any size. And, unlike traditional enterprise storage, Isilon stays simple no matter how much storage is added, how much performance is required or how business needs change in the future. Isilon is challenging enterprises to think differently about their storage, because when they do, they'll recognise there's a better, simpler way. Learn more at www.EMC.com/Isilon.

Big Data Storage

FOR
DUMMIES[®]

A Wiley Brand

EMC Isilon Special Edition

by Will Garside and Brian Cox

FOR
DUMMIES[®]
A Wiley Brand

Big Data Storage For Dummies®, EMC Isilon Special Edition

Published by:
John Wiley & Sons, Ltd
The Atrium
Southern Gate
Chichester
West Sussex
PO19 8SQ
England
www.wiley.com

© 2013 John Wiley & Sons, Ltd, Chichester, West Sussex.

For details on how to create a custom *For Dummies* book for your business or organisation, contact CorporateDevelopment@wiley.com. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@wiley.com.

Visit our homepage at www.customdummies.com.

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 or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher 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 THE 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.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

ISBN: 978-1-118-71392-1 (pbk)

Printed in Great Britain by Page Bros

Introduction

Welcome to *Big Data Storage For Dummies*, your guide to understanding key concepts and technologies needed to create a successful data storage architecture to support critical projects.

Data is a collection of facts, such as values or measurements. Data can be numbers, words, observations or even just descriptions of things.

Storing and retrieving vast amounts of information, as well as finding insights within the mass of data, is the heart of the Big Data concept and why the idea is important to the IT community and society as a whole.

About This Book

This book may be small, but is packed with helpful guidance on how to design, implement and manage valuable data and storage platforms.

Foolish Assumptions

In writing this book, we've made some assumptions about you. We assume that:

- ✔ You're a participant within an organisation planning to implement a big data project.
- ✔ You may be a manager or team member but not necessarily a technical expert.
- ✔ You need to be able to get involved in a Big Data project and may have a critical role which can benefit from a broad understanding of the key concepts.

How This Book Is Organised

Big Data Storage For Dummies is divided into seven concise and information-packed chapters:

- ✔ **Chapter 1: Exploring the World of Data.** This part walks you through the fundamentals of data types and structures.
- ✔ **Chapter 2: How Big Data Can Help Your Organisation.** This part helps you understand how Big Data can help organisations solve problems and provide benefits.
- ✔ **Chapter 3: Building an Effective Infrastructure for Big Data.** Find out how the individual building blocks can help create an effective foundation for critical projects.
- ✔ **Chapter 4: Improving a Big Data Project with Scale-out Storage.** Innovative new storage technology can help projects deliver real results.
- ✔ **Chapter 5: Best Practice for Scale-out Storage in a Big Data World.** These top tips can help your project stay on track.
- ✔ **Chapter 6: Extra Considerations for Big-Data Storage.** We cover extra points to bear in mind to ensure Big Data success.
- ✔ **Chapter 7: Ten Tips for a Successful Big Data Project.** Head here for the famous *For Dummies* Part of Tens – ten quick tips to bear in mind as you embark on your Big Data journey.

You can dip in and out of this book as you like, or read it from cover to cover – it shouldn't take you long!

Icons Used in This Book

To make it even easier to navigate to the most useful information, these icons highlight key text:



The target draws your attention to top-notch advice.



The knotted string highlights important information to bear in mind.



Check out these examples of Big Data projects for advice and inspiration.

Where to Go from Here

You can take the traditional route and read this book straight through. Or you can skip between sections, using the section headings as your guide to pinpoint the information you need. Whichever way you choose, you can't go wrong. Both paths lead to the same outcome – the knowledge you need to build a highly scalable, easily managed and well-protected storage solution to support critical Big Data projects.

Chapter 1

Exploring the World of Data

.....

In This Chapter

- ▶ Defining data
 - ▶ Understanding unstructured and structured data
 - ▶ Knowing how we consume data
 - ▶ Storing and retrieving data
 - ▶ Realising the benefits and knowing the risks
-

The world is alive with electronic information. Every second of the day, computers and other electronic systems are creating, processing, transmitting and receiving huge volumes of information. We create around 2,200 petabytes of data *every day*. This huge volume includes 2 million searches processed by Google each minute, 4,000 hours of video uploaded into YouTube every hour and 144 billion emails sent around the world every day. This equates to the entire contents of the US Library of Congress passing across the internet every 10 seconds!

In this chapter we explore different types of data and what we need to store and retrieve it.

Delving Deeper into Data

Data falls into many forms such as sound, pictures, video, barcodes, financial transactions and many other containers and is broken into multiple categorisations: structured or unstructured, qualitative or quantitative, and discrete or continuous.

Understanding unstructured and structured data

Irrespective of its source, data normally falls into two types, namely structured or unstructured:

- ✓ **Unstructured data** is information that typically doesn't have a pre-defined data model or doesn't fit well into ordered tables or spreadsheets. In the business world, unstructured information is often text-heavy, and may contain data such as dates, numbers and facts. Images, video and audio files are often described as unstructured although they often have some form of organisation; the lack of structure makes compilation a time and energy-consuming task for a machine intelligence.
- ✓ **Structured data** refers to information that's highly organised such as sales data within a relational database. Computers can easily search and organise it based on many criteria. The information on a barcode may look unrecognisable to the human eye but it's highly structured and easily read by computers.

Semi-structured data

If unstructured data is easily understood by humans and structured data is designed for machines, a lot of data sits in the middle!



Emails in the inbox of a sales manager might be arranged by date, time or size, but if they were truly fully structured, they'd also be arranged by sales opportunity or client project. But this is tricky because people don't generally write about precisely one subject even in a focused email. However, the same sales manager may have a spreadsheet listing current sales data that's quickly organised by client, product, time or date – or combinations of any of these reference points.

So data can be different flavours:

- ✓ **Qualitative data** is normally descriptive information and is often subjective. For example, Bob Smith is a young man, wearing brown jeans and a brown T-shirt.
- ✓ **Quantitative data** is numerical information and can be either discrete or continuous:
 - **Discrete data** about Bob Smith is that he has two arms and is the son of John Smith.
 - **Continuous data** is that John Smith weighs 200 pounds and is five feet tall.



In simple terms, discrete data is counted, continuous data is measured.

If you saw a photo of the young Bob Smith you'd see structured data in the form of an image but it's your ability to estimate age, type of material and perception of colour that enables you to generate a *qualitative* assessment. However, Bob's height and weight can only be truly *quantified* through measurement, and both these factors change over his lifetime.

Audio and video data

An audio or video file has a structure but the content also has qualitative, quantitative and discrete information.



Say the file was the popular 'Poker Face' song by Lady Gaga:

- ✓ Qualitative data is that the track is pop music sung by a female singer.
- ✓ Quantitative continuous data is that the track lasts for 3 minutes and 43 seconds and the song is sung in English.
- ✓ Quantitative discrete data is that the song has sold 13.46 million copies as of January 1st 2013. However, this data is only discovered through analyses of sales data compiled from external sources and could grow over time.

Raw data

In the case of Bob Smith or the ‘Poker Face’ song, various elements of data have been processed into a picture or audio file. However, a lot of data is raw or unprocessed and is essentially a collection of numbers or characters.



A meteorologist may take data readings for temperature, humidity, wind direction and precipitation, but only after this data is processed and placed into a context can the raw data be turned into information such as whether it will rain or snow tonight.

Creating, Consuming and Storing Data

Information generated by computer systems is typically created as the result of some task. Data creation often requires an input of some kind, a process and then an output. For example, standing at the checkout of your local grocery store, the clerk scanning barcodes on each item at the cash register collects barcode data read by the laser scanner at the register. This process communicates with a remote computer system for a price and description, which is sent back to the cash register to add to the bill. Eventually a total is created and more data such as a loyalty card might also be processed by the register to calculate any discounts. This set of tasks is common in computer systems following a methodology of data-in, process, data-out.

Gaining value from data

That one grocery store may have 10 cash registers and the company might have 10 stores in the same town and hundreds of stores across the country. All the data from each register and store ultimately flows to the head office where more computer systems process this sales data to calculate stock levels and re-order goods.

The financial information from all these stores may go into other systems to calculate profit and loss or to help the purchasing department work out which items are selling well and which aren't popular with customers. The flow of data may then continue to the marketing departments that consider

special offers on poorly performing products or even to manufacturers who may decide to change packaging.

In the example of a chain of grocery stores, data requires four key activities:

- ✓ Capture
- ✓ Transmission
- ✓ Storage
- ✓ Analysis

Storing data

Only half of planet earth's 7 billion people are online so the already huge volume of digital data will grow rapidly in the future. Traditional information stored on physical media such as celluloid films, books and X-ray photos are quickly transitioning to fully digital equivalents that are served to computing devices via communication networks.

Data is created, processed and stored all the time:

- ✓ Making a phone call, using an ATM machine, even filling up a car at a petrol station all generate a few kilobytes of information.
- ✓ Watching a movie via the internet requires 1,000 megabytes of data.
- ✓ Facebook ingests more than 500 terabytes of new data each day.

Massive amounts of data need to be stored for later retrieval. This could be television networks who want to broadcast the movie *The Wizard of Oz*, newspaper agencies who want to retrieve past stories and photos of Mahatma Gandhi or scientific research institutions who need to examine past aerial mappings of the Amazon basin to measure the rate of deforestation. Other organisations may need to keep patient files or financial records to comply with government regulations such as HiPPA or Sarbanes-Oxley. This data often doesn't require analytics or other special tools to uncover the value of the information. The value of a movie, photograph or aerial map is immediately understood.

Other records require more analysis to unlock their value. Amongst the massive flows of ‘edutainment’, petabytes of critical information such as geological surveys, satellite imagery and the results of clinical trials flow across networks. These larger data sets contain insights that can help enterprises find new deposits of natural resources, predict approaching storms and develop ground-breaking cancer cures.

This is all Big Data. The hype surrounding Big Data focuses both on storing and processing the pools of raw data needed to derive tangible benefits, and we cover this in more detail in Chapter 4.

Knowing the potential and the risks

The massive growth in data offers the potential for great scientific breakthroughs, better business models and new ways of managing healthcare, food production and the environment.



Data offers value in the right hands but it is also a target for criminals, business rivals, terrorists or competing nations. Irrespective of whether data consists of telephone calls passing across international communications networks, profile and password data in social media and eCommerce sites or more sensitive information on new scientific discoveries, data in all forms is under constant attack. People, organisations and even entire countries are defining regulations and best practices on how to keep data safe to protect privacy and confidentiality. Almost every major industry sector has several regulations in place to govern data security and privacy. These laws normally cover:

- ✓ Capture
- ✓ Processing
- ✓ Transmission
- ✓ Storage
- ✓ Sharing
- ✓ Destruction



Data security and compliance

One of the most commonly faced data security laws is around credit card data. These laws are defined by the Payment Card Industry (PCI) compliance used by the major credit card issuers to protect personal information and ensure security for transactions processed using a payment card. The majority of the world's financial institutions must comply with these standards if they want to process credit card payments. Failure to meet compliance can result in fines and the loss of

Credit Card Merchant status. The major tenets of PCI and most compliance frameworks consist of:

- ✓ Maintain an information security policy
- ✓ Protect sensitive data through encryption
- ✓ Implement strong access control measures
- ✓ Regularly monitor and test networks and systems

Chapter 2

How Big Data Can Help Your Organisation

.....

In This Chapter

- ▶ Meeting the 3Vs – volume, velocity and variety
 - ▶ Tackling a variety of Big Data problems
 - ▶ Exploring Big Data Analytics
 - ▶ Break down big projects into smaller tasks with Hadoop
-

The world is awash with digital data and, when turned into information, can help us with almost every facet of our lives. In the most basic terms, Big Data is reached when the traditional information technology hardware and software can no longer contain, manage and protect the rapid growth and scale of large amounts of data nor be able to provide insight into it in a timely manner.

In this chapter we explore Big Data Analytics, which is a method of extracting new insights and knowledge from the masses of available data. Like trying to find a needle in a haystack, Big Data Analysis projects can make a start by trying to find the right haystack!

We also dip into Hadoop, a programming framework that breaks down big projects into smaller tasks.

Identifying a Need for Big Data

The term Big Data has been around since the turn of the millennium and was initially proposed by analysts at technology

researchers Gartner around three dimensions. These Big Data parameters are:

- ✔ **Volume:** Very large or ever increasing amounts of data.
- ✔ **Velocity:** The speed of data in and out.
- ✔ **Variety:** The range of data types and sources.

These 3Vs of volume, velocity and variety are the characteristics of Big Data, but the main consideration is whether this data can be processed to deliver enhanced insight and decision making in a reasonable amount of time.



Clear Big Data problems include:

- ✔ **A movie studio** which needs to produce and store a wide variety of movie production stock and output from raw unprocessed footage to a range of post-processed formats such as standard cinemas, IMAX, 3D, High Definition Television, smart phones and airline in-flight entertainment systems. The formats need to be further localised for dozens of languages, length and censorship standards by country.
- ✔ **A healthcare organisation** which must store in a patient's record, every doctor's chart note, blood work result, X-ray, MRI, sonogram or other medical image for that patient's lifetime multiplied by the hundreds, thousands or millions of patients served by that organisation.
- ✔ **A legal firm working on a major class action lawsuit** needs to not only capture huge amounts of electronic documentation such as emails, electronic calendars and forms, but also index them in relation to elements of the case. The ability to quickly find patterns, chains of communication and relationships is vital in proving liability.
- ✔ **For an aerospace engineering company**, testing the performance, fuel efficiency and tolerances of a new jet engine is a critical Big Data project. Building prototypes is expensive, so the ability to create a computer simulation and input data across every conceivable take off, flight pattern and landing in different weather conditions is a major cost saving.
- ✔ **For a national security service**, using facial recognition software to quickly analyse images from hours of video surveillance footage to find an elusive fugitive is another

example of a real world Big Data problem. Having human operators perform the task is cost prohibitive, so automation by machine requires solving many Big Data problems.

Not really Big Data?

So, what *isn't* a Big Data problem? Is a regional sales manager trying to find out how many size 12 dresses bought from a particular store on Christmas Eve a Big Data problem? No; this information is recorded by the store's stock control systems as each item is scanned and paid for at the cash register. Although the database containing all purchases may well be large, the information is relatively easy to find from the correct database.

But it could be...

However, if the company wanted to find out which style of dress is the most popular with women over 30, or if certain dresses also promoted accessories sales, this information might require additional data from multiple stores, loyalty cards or surveys and require intense computation to determine the relevant correlation. If this information is needed urgently for the spring fashion marketing campaign, the problem could now become a Big Data one.

You don't really have Big Data if:

- ✓ The information you need is already collated in a single spreadsheet.
- ✓ You can find the answer to a query in a single database which takes minutes rather than days to process.
- ✓ The information storage and processing is readily handled by traditional IT tools dealing with a moderate amount of data.

Introducing Big Data Analytics



Big Data Analytics is the process of examining data to determine a useful piece of information or insight. The primary goal of Big Data Analytics is to help companies make better business decisions by enabling data scientists and other users

to analyse huge volumes of transaction data as well as other data sources that may be left untapped by conventional business intelligence programs.

These other data sources may include Web server logs and Internet clickstream data, social media activity reports, mobile-phone call records and information captured by sensors. As well as unstructured data of that sort, large transaction processing systems and other highly structured data are valid forms of Big Data that benefit from Big Data Analytics.

In many cases, the key criterion is often not whether the data is structured or unstructured but if the problem can be solved in a timely and cost effective manner!

The problem normally comes with the ability to deal with the 3Vs (volume, velocity and variety) of data in a timely manner to derive a benefit. In a highly competitive world, this time delay is where fortunes can be made or lost. So let's look at a range of analytics problems in more detail.

A small Big Data problem



The manager of a school cafeteria needs to increase revenue by 10% yet still provide a healthy meal to the 1,000 students that have lunch in the cafeteria each day. Students pay a set amount for the lunchtime meal, which changes every day, or they can bring in a packed lunch. The manager could simply increase meal costs by 10% but that might prompt more students to bring in packed lunches. Instead, the manager decides to use Big Data Analytics to find a solution.

1. **First step** is the creation of a spreadsheet containing how many portions of each meal were prepared, which meals were purchased each day and the overall cost of each meal.
2. **Second step** is an analysis over the last year in which the manager discovers that the students like the lasagne, hamburgers and hotdogs but weren't keen on the curry or meatloaf. In fact, 30% of each serving of meatloaf was being thrown away!
3. **Results** suggest that simply replacing meatloaf with another lasagne may well provide a 10% revenue increase for the cafeteria.

A medium Big Data problem



An online arts and crafts supplies retailer is desperate to increase customer order value and frequency, especially with more competition in its sector. The sales director decides that data analysis is a good place to start.

- 1. First step** is to collate a database of products, customers and orders across the previous year. The firm has had 200,000 products ordered from a customer base of around 20,000 customers. The firm also sends out a direct marketing email every month with special offers and runs a loyalty scheme which gives points towards discounts.
- 2. Second step** is to gain a better understanding of the customers by collating customer profiles collected during the loyalty card sign-up process. This includes age, sex, marital status, number of children and occupation. The sales director can now analyse how certain demographics spend within the store through cross-reference.
- 3. Third step** is to use trend analysis software, which determines that 10% of customers tend to purchase paper along with paints. Also, loyalty card owners who have kids tend to purchase more bulk items at the start of the school term.
- 4. Results** gleaned by cross-referencing multiple databases and comparing these to the effectiveness of different campaigns enables the sales director to create 'suggested purchase' reminders on the website. In addition, marketing campaigns targeting parents can become more effective.

A big Big Data problem



As the manager of a fraud detection team for a large credit card company, Sarah is trying to spot potentially fraudulent transactions from hundreds of millions of financial activities that take place each day. Sarah is constricted by several factors including the need to avoid inconveniencing customers, the merchant's ability to sell goods quickly and the legal restriction on access to personal data. These factors are further complicated by regional laws, cultural differences and geographic distances.

The effectiveness to deal with credit card fraud is a Big Data problem which requires managing the 3Vs: a high volume of data, arriving with rapid velocity and a great deal of variety. Data arrives into the fraud detection system from a huge number of systems and needs to be analysed in microseconds to prevent a fraud attempt and then later analysed to discover wider trends or organised perpetrators.

Hello Hadoop: Welcoming Parallel Processing

Even the largest computers struggle with complex problems that have a lot of variables and large data sets. Imagine if one person had to sort through 26,000 boxes of large balls containing sets of 1,000 balls each with one letter of the alphabet: the task would take days. But if you separated the contents of the 1,000 unit boxes into 10 smaller equal boxes and asked 10 separate people to work on these smaller tasks, the job would be completed 10 times faster. This notion of *parallel processing* is one of the cornerstones of many Big Data projects.

Apache Hadoop (named after the creator Doug Cutting's child's toy elephant) is a free programming framework that supports the processing of large data sets in a distributed computing environment. Hadoop is part of the Apache project sponsored by the Apache Software Foundation and although it originally used Java, any programming language can be used to implement many parts of the system.



Hadoop was inspired by Google's MapReduce, a software framework in which an application is broken down into numerous small parts. Any of these parts (also called fragments or blocks) can be run on any computer connected in an organised group called a cluster. Hadoop makes it possible to run applications on thousands of individual computers involving thousands of terabytes of data. Its distributed file system facilitates rapid data transfer rates among nodes and enables the system to continue operating uninterrupted in case of a node failure. This approach lowers the risk of catastrophic system failure, even if a significant number of computers become inoperative.



First aid: Big Data helps hospital

Boston Children's Hospital hit storage limitations with its traditional storage area network (SAN) system when new technologies caused the information its researchers depend on to grow rapidly and unpredictably.

With their efforts focused on creating new treatments for seriously ill children, the researchers need data to be immediately available, anytime, anywhere.

To address the impact of rapid data growth on its overall IT backup operations, Boston Children's Hospital deployed Isilon's asynchronous data replication software SyncIQ to replicate its research information between two EMC Isilon clusters.

This created significant time and cost savings, improved overall data reliability and completely eliminated the impact of research data on overall IT backup operations. The single, shared pool of storage provides research staff with immediate, around-the-clock access to massive file-based data archives and requires significantly less full-time equivalent (FTE) support.

With EMC Isilon, Boston Children's Hospital's research staff always have the storage they need, when they need it, enabling work to cure childhood disease to progress uninterrupted.

Chapter 3

Building an Effective Infrastructure for Big Data

.....

In This Chapter

- ▶ Understanding scale-up and scale-out data storage
 - ▶ Knowing how the data lifecycles can build better storage architectures
 - ▶ Building for active and archive data
-

Irespective of whether digital data is structured, unstructured, quantitative or qualitative (head to Chapter 1 for a refresher of these terms if you need to), it all needs to be stored somewhere. This storage might be for a millisecond or a lifetime, depending on the value of the data, its usefulness or compliance or your personal requirements.

In this chapter we explore Big Data Storage. Big Data Storage is composed of modern architectures that have grown up in the era of Facebook, Smart Meters and Google Maps. These architectures were designed from their inception to provide easy, modular growth from moderate to massive amounts of data.

Data Storage Considerations

Bear the following points in mind as you consider Big Data storage:

✓ **Data is created by actions or through processes.**

Typically, data originates from a source or action. It then flows between data stores and data consuming clients. A data store could be a large database or archive of documents, while clients can include desktop productivity tools, development environments and frameworks,

enterprise resource planning (ERP), customer relationship management (CRM) and web content management system (CMS).

- ✓ **Data is stored within many formats.** Data within an enterprise is stored in various formats. One of the most common is the relational databases that come in a large number of varieties. Other types of data include numeric and text files, XML files, spreadsheets and a variety of proprietary storage, each with their own indexing and data access methods.
- ✓ **Data moves around and between organisations.** Data isn't constrained to a single organisation and needs to be shared or aggregated from sources outside the direct control of the user. For example:
 - A car insurance company calculating an insurance premium needs to consult the database of the government agency that manages driving licenses to make sure that the person seeking coverage is legally able to drive.
 - The same insurer does a credit check with a reference agency to determine if the driver can qualify for a monthly payment schedule.
 - The data from each of these queries is vital but in some instances, the insurer isn't allowed to hold this information for more than the few seconds needed to create the policy. In fact, longer term retention of any of this data may break government regulations.
- ✓ **Data flow is unique to the process.** How data flows through an organisation is unique to the environment, operating procedures, industry sector and even national laws. However, irrespective of the organisation, the structure of the underlying technology, storage systems, processing elements and the networks that bind these flows together is often very similar.

Scale-up or Scale-out? Reviewing Options for Storing Data

Storing vast amounts of digital data is a major issue for organisations of all shapes and sizes. The rate of technological

change since data storage began on the first magnetic disks developed in the early 1960s has been phenomenal. The disk drive is still the most prevalent storage technology but how it's used has changed dramatically to meet new demands. The two dominant trends are scale-up, where you buy a bigger storage system; or scale-out, where you buy multiple systems and join them together.

Imagine you start the Speedy Orange Company, which delivers pallets of oranges:

✔ **Scale-up:** You buy a big warehouse to receive and store oranges from the farmer and a large truck capable of transporting huge pallets to each customer. However your business is still growing. New and existing customers demand faster delivery times or more oranges delivered each day. The scale-up option is to buy a bigger warehouse and a larger truck that's able to handle more deliveries.

The scale-up option may be initially cost-effective when the business has only a few, local and very big customers. However, this scale-up business has a number of potential points of failure such as a warehouse fire or the big truck breaking down. In these instances, nobody gets any oranges. Also, once the warehouse and truck have reached capacity, serving just a few more customers requires a major investment.

✔ **Scale-out.** You buy four smaller regional depots to receive and store oranges from the farmer. You also buy four smaller, faster vans capable of transporting multiple smaller pallets to each customer. However your business is still growing. The scale-out option is to buy several more regional depots closer to customers, and additional small vans.

With the scale-out option, if one of the depots catches fire or a van breaks down, the rest of the operation can still deliver some oranges and may even have the capacity to absorb the loss, carry on as normal and not upset any customers. As more business opportunities arise, the company can scale out further by increasing depots and vans flexibly and with smaller capital expenditure.

For both options, the Speedy Orange Company is able to scale both the capacity and the performance of its operations. There’s no hard and fast rule when it comes to which methodology is better as it depends on the situation.



Scale-up architectures for digital data can be better suited to highly structured, large, predictable applications such as databases, while scale-out systems may better fit fast growing, less predictable, unstructured workloads, such as storing logs of internet search queries or large quantities of image files. Check out Table 3-1 to see which system is best for you.



The two methodologies aren’t exclusive; many organisations use both to solve different requirements. So, in terms of the Speedy Orange Company, this might mean that the firm still has a large central warehouse that feeds smaller depots via big trucks while the network of regional depots expands with smaller sites and smaller vans for customer deliveries.

Table 3-1 Scale-out or Scale-up Checklist

<i>Scale-out</i>	<i>Scale-up</i>
The amount of data we need to store for processing is rising at more than 20% per year	Our data isn’t growing at a significant rate
The storage system must support a large number of devices that access the system simultaneously	Most of our data is in one big database that’s highly optimised for our workload
Data can be spread across many storage machines and recombined when retrieval is needed	All data is synchronised to a central repository
We’d rather have slower access than no access at all in the event of a minor issue	Access requirements to our data stores are highly predictable
Our data is mostly unstructured, large and access rates are highly unpredictable	The data sets are all highly structured or relatively small

Understanding the Lifespan of Data to Build Better Storage

Irrespective of where data comes from, is processed, or ultimately resides, it always has a useful lifespan. A digital video of a relative's wedding needs to be kept forever. However, the three digit code from the back of a credit card used for security verification can never be stored in a merchant's sales record at any time after processing.

Real time data must be available quickly

Some data is essential for real time analysis so must be available almost instantaneously to other systems or users. For example, a police officer about to make a traffic stop needs to know quickly if the licence plate of the intercepted car is connected to an armed robbery.



The accessibility and long-term storage of data has major significance in terms of cost and accessibility. In general, data that's accessed frequently or continuously as part of a business process or other operation requires higher performance equipment and service specifications than storage of inactive data, which is accessed less frequently.

See the sidebar 'Real time data storage: Jaguar Land Rover' for an example of real time data storage.

Managing less frequently used data

Data archiving is the process of moving data that's no longer actively used to a separate data storage device for long-term retention. Data archives consist of older data that's still important and necessary for future reference, as well as data that

must be retained for regulatory compliance. Data archives are indexed and have search capabilities so that files and parts of files can be easily located and retrieved. See the sidebar ‘Archive data storage: HathiTrust’ for a great example of data archiving.



Real time data storage: Jaguar Land Rover

Jaguar Land Rover designs, engineers and manufactures some of the world’s most desirable vehicles and its success depends on innovation.

As part of its design and manufacturing processes, Jaguar Land Rover engineers rely on an innovative computer-aided engineering (CAE) process. Historically, engineers had worked with CAE applications to design and build a range of prototype models for simulation. But building physical models is expensive and time consuming. Jaguar Land Rover wanted an innovative process that would increase collaborative efficiency, flexibility and cost-effectiveness, while also reducing time to market.

To address the challenge, the company needed to refresh its IT infrastructure with a high-performance computing (HPC) environment that would drive virtual simulation for all of its engineers.

Jaguar Land Rover virtual simulations generate over 10 TB of data per day, and the company uses EMC Isilon X-Series scale-out storage capabilities to add capacity to their

original 500 TB storage configuration. Over six months with EMC Isilon, the HPC environment grew by over 250%. Storage capacity increased by over 500%, and network management architecture saw a tenfold increase.

Virtual simulation programs, driven by EMC Isilon technologies, enable teams to look at problems in much more detail, easily test new ideas, and make changes faster than ever before. Engineers can now create spatial images and resolve challenges prior to prototyping, which also significantly reduces costs.

Because the teams can quickly access the hundreds of TBs of design iterations on the EMC Isilon system, they can turn around new ideas in a matter of days, and see new designs prior to prototyping. Now, Jaguar Land Rover is doing simulations in the early phases, even before some of the design and geometric data has been created. The team can view information in real time to understand where a simulation is going and decide whether they need to take any corrective action.

EXAMPLE



Archive data storage: HathiTrust

In 2008, the University of Michigan (U-M) in conjunction with the Committee on Institutional Cooperation (CIC), embarked on a massive project to collect and preserve a shared digital repository of human knowledge called HathiTrust.

The initial focus of the partnership has been on preserving and providing access to digitised book and journal content from the partner library collections. Foremost was the challenge of being able to create a data storage infrastructure robust enough to support over 10 million

digital objects and handle the rapid scaling that the ambitious project would demand.

The EMC Isilon scale-out NAS system is the primary repository for the HathiTrust Digital Library. In partnership with Google and others, HathiTrust has successfully digitised more than 10.5 million volumes — 3.6 billion pages — from the collective libraries of the partnership to create a massive digital repository of library materials consuming over 470 terabytes.

Active and archive data are both important

Many Big Data projects use both active and archive data to deliver insight. For example, active or real-time data from the stock market can help a trader to buy or sell stocks, while archived data around a company's long-term strategy, market growth and products is useful for better overall portfolio management. The real-time information coming in from stock indexes needs to arrive as soon as it's available, while older reports and market trends can be recovered from an archive and analysed over a longer time frame.

Faster data access is normally more costly

REMEMBER



In simplistic terms, real-time, active or continuous data that enables rapid decision-making typically resides on the fastest available storage media. Normally, the faster the media, the more costly it is compared to the available capacity. This is

often called the 'Cost per Gigabyte (GB) or Terabyte (TB)'. These different types of storage performance and cost are often designated as different tiers of storage, as shown in Figure 3-1.

Fastest performance – Highest cost per GB/TB
RAM – Computer memory that requires constant power to retain information.
SSD – A solid state disk drive uses a type of RAM to store data.
SATA / SAS disks – Two types of magnetic media with SAS the faster technology.
TAPE – Using spools of magnetic media which requires no power to retain data indefinitely. However, verification of data integrity is much more difficult.
Slowest performance – Lowest cost per GB/TB

Figure 3-1: Faster performance equals the highest cost per GB or TB.

Chapter 4

Improving a Big Data Project with Scale-out Storage

.....

In This Chapter

- ▶ Understanding the storage node concept
 - ▶ Building a storage cluster
 - ▶ Avoiding Big Data storage infrastructure problems
-

Data comes in many forms and has varying requirements around volume, velocity and variety. Big Data projects may need elements of structured and unstructured, as well as real-time and archive data to achieve results, and all of this data needs to reside somewhere accessible to the applications needed for analysis.

Our example of the Speedy Orange Company in Chapter 3 explained the two basic philosophies for building a storage architecture – scale-up and scale-out – and you can use each exclusively or together to achieve the desired result. However, the requirements of Big Data projects that deal with both structured and unstructured data sets combined with high velocity, volume and variety of data tend to lend themselves to scale-out storage technologies.

Exploring a Common Scale-out Architecture

Many scale-out storage technologies follow a similar architectural design which is recognised by experts as best practice. Most systems follow a basic structure that makes them

recognisable as either leaning towards fully scale-out or a hybrid scale-out/up design.

The storage node as the first building block



One of the fundamental aspects of a scale-out architecture is, as the name says, the ability to scale. To achieve this aim, designers use the principle of a *node*. A node is a self-contained storage device that works together with other nodes to store and deliver data between producers and consumers of data.

Here's a very simple analogy. A node is like a gallon bucket and the water within is the data. If you want to store more than a single gallon of water, you add more buckets. Now you have more water and more people can simultaneously draw water from the buckets. However, if one bucket runs dry, people need to wait in line until a bucket with water is free to use.

So to solve this sharing problem, each bucket has pipes connecting them so that multiple people can access the common water source simultaneously from different buckets. If you add more water into one bucket, it simply flows down the connecting pipes and distributes evenly to all the connected buckets, as shown in Figure 4-1. If one bucket develops a leak, you can simply disconnect it from the other buckets and move its pipe to another bucket while you repair it.




Figure 4-1: The buckets (nodes) distribute the water (data).

Many scale-out architectures carve up files and spread the pieces across nodes, thus allowing them to flow like water.

Be warned: in legacy scale-up architectures that attempt to act like true scale-out systems, the files on each node aren't carved up into pieces and spread around and, unlike water, data often stays in the node with merely a copy sent to another client.

Nosing inside a node

A scale-out storage node is essentially a highly optimised server with a dedicated software application that's designed to manage the storage and flow of data between itself and the other nodes within the group. Each node can also communicate with external clients via the network to either store or send data, as shown in Figure 4-2.

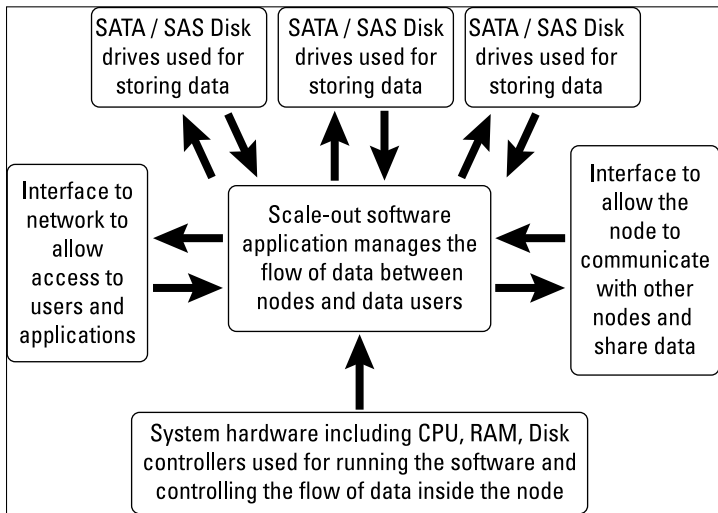


Figure 4-2: True scale-out architecture.

In a true scale-out architecture, each node contains one or more central processing units (CPUs), an amount of Random Access Memory (RAM) and a number of hard disk drives. The unit also has a connection to the network and often a method of interconnecting the nodes.

Connecting the nodes to form a cluster

Each scale-out storage node is normally connected in two ways. The first is a connection that links all the nodes together via a switch to form a storage cluster, as shown in Figure 4-3.

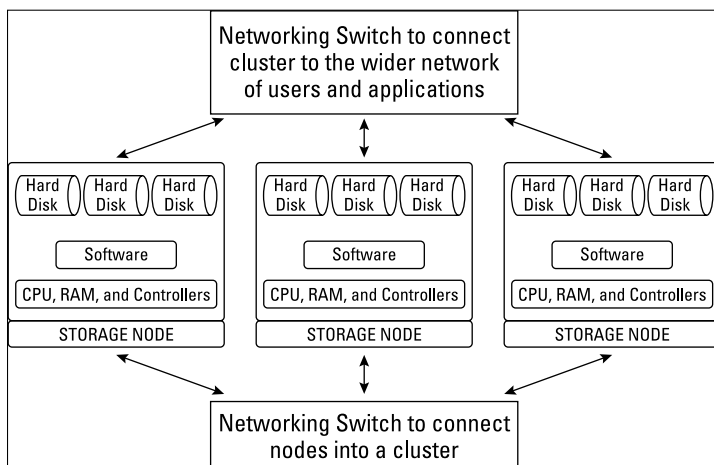


Figure 4-3: The nodes are linked together to form a storage cluster.

This interconnect enables the cluster to share data between nodes, which offers data resilience and provides a performance boost because now any node can serve up the data that actually resides on a different node within the cluster. These dedicated interconnects normally use a fast network interface such as 1Gbit (gigabit) or 10Gbit Ethernet or even faster 40Gbit Infiniband.

Connecting the cluster to the wider network

The scale-out cluster is also connected to the wider network to allow access from applications and users both locally and across the wider area network. Most scale-out clusters (and

especially when working with Big Data projects) tend to have the fastest possible links to the analysis applications running on application servers. These are typically 1Gbit and 10Gbit Ethernet connections.

Communicating with confidence

The scale-out cluster can communicate with applications in a number of ways. These *communication protocols* allow access to different communications standards that are typically adopted by different operating system vendors. For example:

- ✓ Common Internet File System (CIFS) also called Server Message Block (SMB) is commonly used by Microsoft Windows applications.
- ✓ Network File System (NFS) is a distributed file system that's common with UNIX and Linux open source applications.
- ✓ File Transfer Protocol (FTP) is an older network protocol often used for basic file transfer.
- ✓ Hypertext Transfer Protocol (HTTP) is an application protocol used extensively for web applications.



The scale-out storage cluster is often able to use multiple file systems and communication protocols but some offer better compatibility between applications, while others may have a raw performance advantage. There's a balancing act between computability and performance, so a key stage in designing the storage architecture for a Big Data project is to gain an understanding of how applications need to communicate between clients, storage and users.

Understanding the Benefits and Limitations

In general, Big Data projects have a combination of high volume, high velocity and high variety of data. A fraud detection system may have all three, while a simulation of a jet engine in a wind tunnel might have just two elements. Both

projects pose a challenge to the storage system collating or feeding data into the analysis systems.

By using scale-out storage architectures, it's relatively easy to increase capacity and performance due to the basic node architecture. Each node contains storage capacity, network connectivity and powerful processing units (CPUs) to receive, store and transmit data. Each node offers an increase to the overall capability of the storage cluster.

✔ **Benefits:** More nodes means more capability. For example, a single node has 100 Terabytes of storage capacity, 2Gbit of network bandwidth and 256GByte of RAM for caching data to improve transfer rates. Adding another node effectively doubles this available pool. Using four nodes delivers four times the available storage, bandwidth and caching pool.

✔ **Limitations:** This additive effect of more nodes isn't a constantly increasing vector. Just like fitting a faster engine in a car, at certain points other factors limit the top speed, such as wind resistance, tyres or the length of the road. In the case of a scale-out cluster, the network connectivity often proves a limitation on increasing throughput performance. Also, the hard disks within each node have a theoretical maximum transfer rate. Even using faster solid state disks eventually causes a bottle neck.

Another limitation is physical size. Although each node is relatively small, to build a scale-out cluster that could store all the pictures on Facebook would require a space the size of a football pitch and enough power to keep a small town running.

Preventing Problems

It's a sad fact, but everything mechanical eventually fails. Even the simplest machines eventually jam, break or wear out through use. Digital technology has mostly replaced mechanical elements like valves or relays with integrated circuits and silicon chips that drastically improve reliability. However, even this change doesn't make a computer or scale-out node last forever.

Failure is guaranteed

Some items are guaranteed to fail and you can predict the likelihood based on usage. For example, the hard disk technology used to store the bulk of digital information uses spinning plates that have magnetised encoded areas. The plates themselves experience microscopic wear and tear as the reading device floats nanometres above the spinning platter. Eventually, this constant magnetising of the surface causes wear and tear that results in data being either incorrectly written or read from. In a worst case scenario, the constantly spinning motor that rotates these plates simply fails.

To circumvent such wear and tear, engineers have come up with various schemes to protect digital data. The most common method is simply to copy a complete or partial duplicate of data from one drive to another. This is often known as Redundant Array of Independent Disks (or RAID). However, RAID has three major problems:

- ✓ RAID is wasteful. Considering how much data the world produces, keeping a complete or even partial backup isn't always practical.
- ✓ Rebuilding data from the copy into a new, coherent source is time-consuming and impacts the overall efficiency of the storage platform.
- ✓ As the volume of original and duplicate data increases, both sets could simultaneously experience critical data loss.

For these reasons, RAID is largely unsuitable for critical Big Data projects.



Instead, many scale-out architectures break up the data into smaller chunks and spread these chunks across multiple disks and on multiple nodes. Many of these chunks are also duplicated and the systems create mathematical checksums to enable the rebuilding of data if it's lost. A bit like the popular Sudoku puzzle, having enough pieces and clever maths can enable you to rebuild the whole data set.

This distribution of data is also able to scale in line with the number of nodes and disks to counteract the probability that a simultaneous failure will delete a critical piece of data. Another advantage is that even if both a disk and an entire node is lost, the data is still available. This type of protection, where data is spread across multiple locations or sent over multiple channels, is known as erasure coding, data striping or forward error correction.

Chapter 5

Best Practice for Scale-out Storage in a Big Data World

In This Chapter

- ▶ Understanding tiers, quotas and thin provisioning
 - ▶ Using Solid State Drive
 - ▶ Considering security and legal issues
-

We've looked at the fundamentals of data and the elements you need to create storage architecture suited to a Big Data project. But a number of technologies, management processes and best practices can help save time, money and make your system more secure. That's what we explore in this chapter.

Tiers Before Bedtime: Looking Closer at Scale-out Architecture

Although every organisation has a slightly different requirement and technical approach, a set of common elements interact with scale-out storage architectures within a typical enterprise organisation, as shown in Figure 5-1.

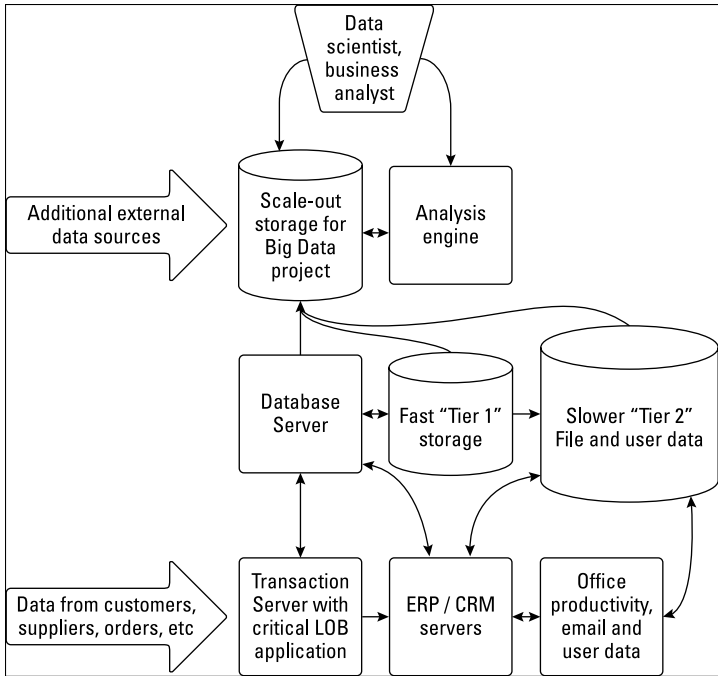


Figure 5-1: Common elements interact with scale-out storage architecture.

Aligning data value to storage costs with tiers

Big data projects are essentially about extracting valuable insight from raw data. Both the raw data and the gleaned information have an intrinsic value. This value dictates how the data is stored, who has access to it and where it lives on the physical storage devices. Data also needs to be arranged to make it accessible to the analysis engines or a particular performance or capacity requirement.

To give an analogy: The important elements of a car – the steering wheel, foot pedals, gear shift and seat – are arranged with accessibility and safety in mind. However, the less-important aspects such as the switch to open the sun-roof or to use the CD player are placed away from critical elements like the accelerator, as they’re accessed less frequently.



Data tiers describe in a logical fashion the characteristics of lots of data. These tiers can describe both the type of information and its access requirements. This is especially true in a Big Data project where the analysis may be looking over older or seemingly less valuable data that may well reside on a slower tier of storage. A simple way to look at tiers is in a hierarchical fashion:

Tier 1	Mission critical and recently accessed data	Requires the highest degree of performance, reliability and accessibility
Tier 2	Seldom used data such as previous years' sales information	Requires lesser performance
Tier 3	Archive data that may require retention for compliance	Requires slower performance, but very large capacity

Improving efficiency with automatic tiering technology

Many scale-out storage platforms have software that helps manage this tiered approach. The software uses a user definable set of rules to automate the process of moving data around both the cluster and to available external storage. For example, the rules can be set up to move data that hasn't been accessed in a year onto lower performance disk, or conversely, move a frequently requested item such as a popular video clip into a higher performance area.

Managing Growing User Demand with Quotas

IT projects tend to fill up all the available space if left unchecked, like an attic groaning at the seams. Starting from the humble email in-box, through to staging areas for data needed for projects, and even archived files that have no deletion date – data keeps on growing. For administrators with finite resources, a *quota* is a method of defining how much

space every user or project team has, and placing the responsibility for managing that capacity into a policy or into the hands of the agency responsible for that data.



Quotas are becoming increasing flexible tools in the armoury of storage admin and can be linked to tiers to create areas that are used for a specific function. These quotas can also have a monetary cost associated to them to allow internal billing or chargeback.

When technical teams are given a finite resource pool to manage, they become much more diligent about cleaning up unneeded data after themselves!

Reducing Wasted Space with Thin Provisioning

Traditional storage management technologies produce a lot of waste. Vendors talk about raw capacity (the size of each drive times the number of drives) versus useable capacity (raw capacity minus storage overheads) and these numbers may vary greatly. This discrepancy is caused by many factors which include:

- ✓ Overhead for the file systems
- ✓ Data protection method
- ✓ Metadata used to keep track of files
- ✓ Spare capacity across many storage containers

In addition to these capacity consuming overheads, a lot of data is duplicated which also wastes precious storage capacity. However, the biggest culprit is inefficient provisioning technologies where the capacity of the storage platform may be effectively assigned but actually is unused.



Scale-out storage combined with thin provisioning technology is one approach that helps to get over this problem by not actually assigning the capacity to any groups up front. Instead, it dynamically assigns the pool of storage only on request. This method placates any applications that check for available storage space before writing data by informing the system that

there's plenty of capacity for the data operation to take place. In some instances, the intelligence within the thin provisioning engine may work in tandem with automated storage tiering to move data that's never used off expensive primary storage to a different tier of storage (refer to Figure 5-1) that's cheaper or better suited to longer term archive.

Turbo-charging Your Big Data Project with Solid State

So your Big Data analysis project is underway. Data is flowing, applications are delivering new insights but demands are coming in for more performance. So what do you do?



Well, one common quick fix is to speed up the performance of the data as it moves through the storage cluster. Physical spinning disks have a maximum throughput of data that's limited by just how fast a disk can spin and how quickly data can be read from it as a magnetic signal. A faster method is to use a diskless media such as Random Access Memory chips. A Solid State Drive is, as the name suggests, a disk drive that uses memory chips instead of spinning disks. The technology comes in two main flavours:

- ✓ Flash SSDs are suited to read-only applications and mobility applications.
- ✓ DRAM SSDs have much higher read and write performance with a better cost per unit of performance than Flash but have a higher upfront cost per GB of storage.



However, simply changing all the disks in a scale-out storage platform from spinning platters to SSD is extremely expensive. Also, SSD don't last forever, just like disks. Plus, as flash drives become larger, there are question marks over reliability in comparison to old-fashioned spinning disks.



Scale-out architectures often use SSD in different ways. One intelligent use is to use SSD to speed up searching for items requested by the client. So, in a cluster with 20 nodes and several billion items of data, the process of actually finding a specific data item within the cluster may take a second. Moving this map (sometimes called *metadata*) of where each

item is physically located onto the SSD instead of slower spinning disks can reduce this delay. This intelligent use of SSD boosts overall system performance without having to resort to changing out every physical disk for a SSD equivalent.

Ensuring Security

Digital data is valuable. A Big Data project that aims to generate a new insight or scientific breakthrough is like a precious jewel to a thief intent on stealing the results – or even the source material. Information security is a constant concern so don't overlook it when working on any project. A Big Data project may need more protection due to the potential damage that having so much sensitive information in one place could cause.

The many important considerations around storage security include:

- ✔ Ensure the network is easily accessible to authorised people, corporations and agencies.
- ✔ Compromising the system must be extremely difficult for a potential hacker.
- ✔ The network needs to be reliable and stable under a wide variety of environmental conditions and volumes of usage.
- ✔ Provide protection against online threats such as viruses.
- ✔ Only provide access to the data directly relevant to each department.
- ✔ Assign certain actions or privileges to an individual as they match their job responsibilities.
- ✔ Encrypt sensitive data.
- ✔ Disable unnecessary services to minimise potential security holes.
- ✔ Regularly install updates to the operating system and hardware devices.
- ✔ Inform all users of the principles and policies that have been put in place governing the use of the network.

It May be Big, But Is It Legal?

Sometimes, a Big Data project can prompt an organisation to gather and store types of information that it previously hadn't retained. In some instances, the company may need to bring in data from an external source for comparison against its own data sets and doing so moves the organisation into a new legal area. For example, if a German insurance company wanted to analyse clinical outcomes of different surgical procedures against policy types and pay-outs, the project could require huge volumes of data from around the globe. If the source of the data was from the US, its storage would need to comply with the Healthcare Insurance Portability and Accountability Act (HIPAA).

As the data needed to power Big Data projects crosses international borders, there can be additional requirements to meet local regulations. For example, the European Union's Data Protection Directive means that organisations that fail to secure data or suffer a breach can expect fines or, in serious cases, imprisonment for senior executives.



Key compliance frameworks to be aware of include:

- ✓ **Healthcare Insurance Portability and Accountability Act (HIPAA)**, which keeps health information private.
- ✓ **Sarbanes Oxley Act**, aimed at the accounting sector.
- ✓ **Gramm-Leach-Bliley Act (GLB)**, which requires financial institutions to ensure the security and confidentiality of customers' information.
- ✓ **Bank Secrecy Act**, used by US government to pursue tax-related crimes.

Chapter 6

Extra Considerations for Big Data Storage

.....

In This Chapter

- ▶ Improving the data centre
 - ▶ Longer term planning to save money
 - ▶ Considering virtualisation and cloud computing
-

In this chapter we look at the other considerations or parts of the business that can often be impacted by a Big Data project. We also consider some longer term goals and strategies that may well provide an alternative to doing a Big Data project in-house.

Don't Forget the Data Centre!

Various estimates suggest that storage accounts for around 35 per cent of the power used in data centres. The drain on power stations is likely to grow as more people go online to generate and consume digital content. With energy costs rising and the potential for energy surcharges, energy consumption is rapidly becoming a major concern. As Big Data projects arrive, and with them new storage and server centres, follow these tips:

✔ **Reduce data centre hot spots to reduce cooling costs.**

As data centres grow without sufficient thought to power and cooling requirements, a hot spot can start to cause problems for the smooth operation of computing equipment. Storage racks are large units and, once placed on

the data centre floor, are difficult to move without causing disruption to applications. Instead, distribute workloads more strategically across the site.

✔ **Configure equipment racks with cold and hot rows.**

Most computer devices expel hot air from the back of the unit. If the back row is breathing in the hot exhaust from the adjacent front row, proper cool air flow is disrupted which forces data centre air conditioning units to generate more expensive cold air. Instead, ensure racks are placed with exhaust designed to expel hot air into unused areas or vented away.

✔ **Move workloads to save energy.** Virtualisation and storage management software can help data centres to reorganise where computer and storage tasks physically take place within the data centre. This can help to evenly spread or (in theory at least) move workloads to underutilised servers and turn off 'empty' storage nodes or unused servers without having to physically move racks around.

✔ **Higher density can expand valuable floor space.**

Consider increasing density of the hard drives used for data storage. Although a 4TB drive has four times the capacity of a 1TB drive, it doesn't use four times the amount of power. With some scale-out storage architectures, it's relatively easy to swap out drives without downtime. If these density upgrades are carried out a single node at a time, a 100TB cluster can expand to a 400TB cluster and consume the same physical footprint for only a few percentage points more power consumption.

Longer Term Planning for Major Cost Benefits

Irrespective of whether your Big Data project is small, medium or large, your IT infrastructure is probably growing. Even with the arrival of virtualisation, that allows computers to operate more efficiently, the criticality of IT has forced more dependence on larger, more complex systems.

Storage has become more powerful yet physically smaller. The cost per gigabyte of storage capacity has fallen faster, yet storage density, speed and performance has improved massively.



Disk based technologies for data storage is the most likely path for upgrade. As standard drives expand past 4TB and up to possible 16TB per unit over the next 5 years, the ability for organisations to upgrade capacity in situ within the same storage pool is a major advantage.

Another longer term strategy is to move data automatically off high performance Serial Attached SCSI (SAS) hard disk drives and SSDs to slower, less expensive storage such as Serial AT Attachment (SATA) drives. Duplicated entries of data are deleted and statistically unimportant information is retired. These Information Lifecycle Management (ILM) projects can help extend the viability of storage architecture.

Getting to Grips with Virtualisation

Virtualisation has been the most significant technology trend of the last decade. However, it's an umbrella term for many different types of computing:

- ✓ **Server Virtualisation:** This enables one server to run multiple operating systems (OS) at the same time, decreasing the number of physical servers needed to run multiple server applications. A virtualised server may not actually offer a visual element to the user and can simply be running a non-interactive process such as a network proxy or data processing task.
- ✓ **Desktop Virtualisation:** Often known as Virtual Desktop Infrastructure (VDI), the concept of desktop virtualisation allows each computer's preferences, OS, applications and files to be hosted on a remote server. Users can then use an access client such as a PC or thin client to view and interact with this remote desktop over a network. Desktop virtualisation has a number of benefits, both for end users and for IT departments, as a low-powered device such as a tablet can now run complex applications and the data management is simplified as it never leaves the central server.
- ✓ **Storage Virtualisation:** This is the consolidation of physical storage from multiple storage devices into what appears to be a single storage device managed from a

central location. Storage Virtualisation is the fundamental concept behind scale-out storage: a collection of storage nodes can be added on demand to increase capacity and performance in a single pool of storage with no disruption to users or applications. This has many benefits in terms of reduced management overhead, less physical space and ability to reduce duplicated data. Storage virtualisation simplifies and often reduces the number of physical storage devices needed for any given volume of data due to efficiencies gained.

Using Cloud Technology for Big Data Projects

Given the rigorous demands that Big Data places on networks, storage and servers, it's not surprising that some customers would outsource the hassle and expense to somebody else. This is an area where cloud computing can potentially help.



Public or private cloud computing is the use of hardware and software resources that are delivered as a service over a network, including the internet.

Clouds can serve different purposes (as shown in Figure 6-1) and include:

- ✓ **Infrastructure as a service (IaaS):** One or many computers with storage and network connectivity that you can access via a network connection.
- ✓ **Software as a service (SaaS):** Access to a specific software application complete with your own data via a network connection.
- ✓ **Platform as a service (PaaS):** Provides the core elements such as software development tools needed to build your own remote IT environment that users can access, possibly via virtual desktops across a network.
- ✓ **Storage as a service (STaaS):** A remote storage platform that has a specific cost per GB for data storage and transfer.

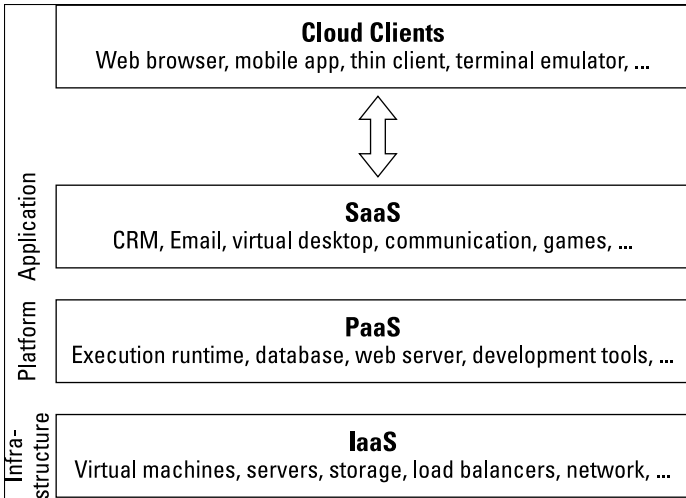


Figure 6-1: Different flavours of cloud computing.

Some Big Data projects may be well suited to running in a public cloud as its elasticity enables it to scale quickly. Also, many public clouds allow a great deal of resources to be rented on a short-term basis without the upfront and often extremely expensive capital expenditure costs.

However, there are still concerns regarding the security, reliability, performance and transfer of data using public cloud technologies:

- ✓ For projects that need to move large quantities of data around the internet, the limitations and cost of network bandwidth may actually make a public cloud solution for a Big Data project more expensive than an on-premise, private cloud equivalent.
- ✓ For organisations who have high-value intellectual property or highly sensitive personal information such as healthcare files or student records, having data stored in an unknown location managed by unknown people is a major cause of concern. In fact, many government entities have data residency or sovereignty laws that require that data created in a certain set of boundaries stay within that jurisdiction, such as not being stored across national borders. Also, the data protection policies and procedures of information in the public cloud can't be easily audited.

- ✔ Performance of data written to and read from the public cloud can be slow and expensive depending on the distance and network type used and the billing rates that the public cloud vendor charges for writing and retrieving that data.
- ✔ Once large amounts of data are stored in the public cloud it can be difficult and expensive to move that data to a different public cloud vendor. The divorce of leaving one public cloud vendor and marrying another can be very painful.



Many organisations are pursuing a private cloud strategy where freedom of cloud access is enabled by the internet, but access control is still maintained by the organisation. Furthermore, the physical security, back-up, disaster recovery and performance of the data is controlled by the organisation.

Chapter 7

Ten Tips for a Successful Big Data Project

.....

In This Chapter:

- ▶ Identifying data types and flows
 - ▶ Preparing for data growth
 - ▶ Avoiding costly mistakes with sensible data management
 - ▶ Planning for worse case scenarios
-

If you're reading this chapter first, we're guessing it's because you're keen to avoid any of the mistakes that can derail a Big Data project. Here are a few issues to consider.

- ✔ **Start any Big Data project with a data review and classification process.** Defining whether data falls into the category of structured, unstructured, qualitative or quantitative is a useful precursor to designing storage architectures (head to Chapter 1 for a refresher). It's also useful to estimate data growth based on past trends and future strategy.
- ✔ **Create a simple map of how the data flows around the organisation.** Having a simple diagram showing where data is created, stored and flows to is helpful when working within a multi-discipline group. Having everybody reading from the same page can avoid costly misunderstandings.
- ✔ **Consider your future data storage requirement based on the success of the Big Data project.** Big Data projects may well uncover new insights or force changes to

operational processes. The resulting information delivered by a project may in turn have an additional data storage requirement causing an exponential growth in capacity requirements. Always consider the longer term view.

- ✔ **Be flexible.** Many projects use both scale-up and scale-out storage technologies in harmony (explained in Chapter 3). Every organisation and project is unique. The selection of a storage technology needs to be goal orientated instead of fixed around a particular technical architecture. Multiple vendors have both scale-up and scale-out products that can work well together.
- ✔ **Data storage requirements may grow but consider automatically moving less frequently accessed data to less costly, slower storage. Deletion is also a viable longer term option.** Irrespective of where data comes from, is processed by, or ultimately resides, it always has a useful lifespan. Deciding when to delete data is a complex task but it can provide a massive cost saving over the longer term. Automatically demoting data to slower storage is an easier task that still reaps massive benefits.
- ✔ **Ask technology vendors about what happens when you reach a theoretical capacity or performance limit.** Although your Big Data project might start out small, it will probably grow over time. Understanding the upgrade path for your chosen technology direction helps you avoid unpleasant surprises a few years down the line.
- ✔ **Plan for the worst case scenario.** Even the simplest machine eventually wears out through use, jams or breaks. When working with a technology supplier, ask what happens if different elements within the storage platform fail. A well-designed system should never have a single point of failure.
- ✔ **Create a quota system early in a project to ease future management issues.** IT projects tend to fill up all the available space if left unchecked. A quota is a method of defining how much space every user or project team has. Place the responsibility for managing that capacity either into a policy or into the hands of the agency responsible for that data.

- ✔ **Always include IT security experts within any Big Data project.** Digital data is valuable. Although a Big Data project might sit within a research group, the overall IT security team needs to be involved from the earliest stages so security is built into the heart of the project.
- ✔ **Remember to include management time when calculating storage costs.** The overall cost of storage needs to include how much time is required in the provisioning and management of the platform. A self-healing and highly automated system that removes the need for a full-time administrator offers considerable longer term cost savings over cheaper hardware that requires lots of manually intensive tasks.

CAPACITY

PERFORMANCE

SCALABILITY

SIMPLICITY



BIG DATA IT'S ON ISILON

THE SOLUTION FOR YOUR BIG DATA STORAGE NEEDS



Your Industry Has Storage Issues. EMC Isilon Has Solutions.

EMC Isilon scale-out NAS solutions enable customers across a variety of industries to dramatically increase the productivity of mission critical applications, workflows, and processes—accelerating business through faster time-to-money and new revenue opportunities.

EMC.COM/ISILON



EMC²

EMC2, EMC, and the EMC Logo, are registered trademarks of EMC in the United States and other countries. © Copyright 2012 EMC Corporation. All rights reserved.

A handy guide for designing, implementing and managing your storage architecture to support innovative Big Data Projects.

The massive growth in electronic data holds the potential for great scientific breakthroughs, better business models and new ways of managing healthcare, food production and the environment. Gaining valuable insight and use of Big Data requires innovative storage architectures, management processes and practices to help save time, money and make you more secure. Read all about it in this book.

- **Understand data fundamentals** – learn about the structure and use of data
- **Build a scalable architecture** – maximise your knowledge of the building blocks for Big Data storage
- **Learn best practice** – discover the key skills to deliver successful Big Data projects

Will Garside is a freelance editor, a published author and IT journalist with more than 17 years' experience across web, print and broadcast media. He wears his techie badge with pride, and loves talking technology with anyone who'll listen. **Brian Cox** is Senior Director of Product and Solutions Marketing at the EMC Isilon Storage Division. He was around in the early days of personal computing, having owned a TI-99/4A and the original model of the two floppy drive IBM PC. In his spare time he plays Taxi Dad for his four kids, taking them to all their school and sporting activities.



Open the book and find:

- What you need to know to build a complete storage architecture
- Techniques to ensure a successful Big Data project
- Best practice tips around data security and compliance
- Practical examples of Big Data in the real world

Go to Dummies.com
for videos, step-by-step examples,
how-to articles, or to shop!

