

Optimizing Database Storage Performance

FOR
DUMMIES[®]
A Wiley Brand

Learn to:

- Identify flash storage benefits for databases
- Improve database performance and response times
- Recognize database challenges in the cloud

Lawrence C. Miller



Optimizing Database Storage Performance

FOR
DUMMIES[®]
A Wiley Brand

NetApp Edition

by Lawrence C. Miller

FOR
DUMMIES[®]
A Wiley Brand

Optimizing Database Storage Performance For Dummies,® NetApp Edition

Published by

John Wiley & Sons, Inc.

111 River St.

Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2016 by John Wiley & Sons, Inc.

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 under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <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. NetApp and the NetApp logo are registered trademarks of NetApp. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

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, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-119-25908-4 (pbk); ISBN: 978-1-119-25912-1 (ebk)

Manufactured in the United States of America

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:

Production Manager:

Carrie A. Burchfield

Editorial Manager: Rev Mengle

Acquisitions Editor:

Amy Fandrei

Business Development

Representative: Karen Hattan

Production Editor:

Siddique Shaik

Table of Contents

| | |
|--|-----------|
| Introduction | 1 |
| About This Book | 1 |
| Icons Used in This Book..... | 2 |
| Beyond the Book..... | 2 |
| Chapter 1: Understanding Market Trends | 3 |
| Forecasting a Partly to Mostly Cloud-y Future..... | 3 |
| Real-Time Analytics | 4 |
| A Few Things to Know about the Internet of Things | 5 |
| Chapter 2: Tuning for Performance..... | 7 |
| Structured and Unstructured Data | 7 |
| Hardware Considerations | 8 |
| Identifying Performance Bottlenecks | 10 |
| Chapter 3: Accelerating SQL Server Performance | 11 |
| Upgrading to SQL Server 2014 | 11 |
| Key Database Storage Capabilities | 14 |
| Deploying Flash Storage..... | 15 |
| Simplifying Database Backup, Restore, and Cloning Operations..... | 20 |

| | |
|--|-----------|
| Chapter 4: Accelerating Oracle Database Storage Performance | 23 |
| Storage Infrastructure for Oracle Database Workloads | 24 |
| Protecting Data with NetApp Data ONTAP | 29 |
| Chapter 5: Going Beyond Performance in the Enterprise..... | 35 |
| NetApp's Data Fabric Vision..... | 35 |
| Availability and Reliability..... | 37 |
| Accelerating DevOps | 37 |
| Data Protection | 38 |
| Chapter 6: Ten Key Database Storage Capabilities..... | 39 |

Introduction



Improving database performance is top of mind for many IT professionals. Their production databases in many ways are the lifeblood of the organization, supporting the transactions, applications, operations, business intelligence, and analytics that make the business work. As the volume and quality of business data grows exponentially, so too does its value to the organization. In many industries, insightful data — and how the business uses the data — has become a key competitive differentiator. Production databases are also under more pressure than ever before — they must process many more transactions at unprecedented speeds, while data is growing exponentially and IT budgets are not.

One of the most important and effective ways to address database performance and cost challenges is to not simply focus on upgrading or optimizing the database software, but to also modernize the underlying hardware infrastructure. This is where innovations such as flash storage, converged infrastructure architectures and sophisticated data management platforms can have a major impact.

About This Book

This book examines the market trends that are driving the ever demanding business requirement for better database performance, technical factors that affect

database performance tuning, different SQL Server and Oracle Database design considerations, how database storage technologies and innovations can help drive business value, and key performance optimization capabilities to look for in a storage solution.

Icons Used in This Book

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



This icon points out information that you should commit to your non-volatile memory or your noggin'!



This icon explains the jargon beneath the jargon!



This icon points out helpful suggestions and useful nuggets of information.

Beyond the Book

There's only so much I can cover in 48 short pages, so if you find yourself at the end of this book, thinking "where can I learn more?" just go to www.netapp.com.

Chapter 1

Understanding Market Trends

.....

In This Chapter

- ▶ Recognizing database challenges in the cloud
 - ▶ Analyzing the impact of big data
 - ▶ Learning things about the Internet of Things
-

Data growth is exploding. International Data Corporation (IDC) has predicted that overall digital data growth requirements will double in size every two years. Your storage infrastructure needs to handle this explosive data growth while maintaining consistent performance levels.

In this chapter, you discover several important market trends that are driving this exponential growth in data.

Forecasting a Partly to Mostly Cloud-y Future

Virtualization and cloud technologies have become mainstream as enterprise cloud strategies have quickly evolved to become competitive differentiators.

No longer relegated to limited use cases, such as application development or data backup, the cloud enables organizations to quickly adapt to rapidly changing business requirements, accelerate time to market, and seamlessly scale their infrastructure as needed.

A unified data fabric (discussed in Chapter 5) is needed to enable mission-critical applications and databases to perform seamlessly across all public, private, and hybrid cloud environments.

Real-Time Analytics

Businesses today recognize the intrinsic value of all the data they have collected as part of their normal business operations, as well as all of the new data they are now collecting. With the right storage architecture, all of this data can be mined and analyzed in interesting ways to gain competitive advantage, deliver better customer experiences, and make intelligent decisions faster.

The Internet of Things and cloud trends discussed later in this chapter contribute further to the explosive growth of data that has to be analyzed for deep insights for businesses in all industries.



Data by itself has no value. Value comes from leveraging the data to drive business results, such as offering new services to customers and increasing revenue.

A Few Things to Know about the Internet of Things

In 2011, IDC predicted in its Digital Universe Study that the total amount of digital information created, captured, and replicated would grow to 35,000 exabytes by 2020 — and that was before the Internet of Things (IoT) really took hold and became the “next big thing!” In its 2015 Digital Universe Study, IDC updated that prediction to 44,000 exabytes by 2020!

The “things” that comprise the IoT are physical devices or objects that contain embedded electronics, software, sensors, and components that enable these things to communicate and exchange data with users, manufacturers, service providers, and/or other “things,” over the Internet.



Gartner, Inc. predicts that the number of IoT devices worldwide will grow to 25 billion by 2020, and ABI Research estimates that the data captured by IoT devices will grow to 1.6 zettabytes during that same period.

Examples of popular IoT devices include

- ✓ **Networked home appliances and components**, such as lighting, thermostats, refrigerators, and security cameras
- ✓ **Wearable technology**, such as smart watches and medical devices
- ✓ **Smart cars** that automatically adapt to changing driving conditions and can self-diagnose mechanical issues, schedule maintenance, and intelligently navigate around traffic and hazardous weather

- ✔ **Smart meters** that enable real-time, two-way communication between utilities providers and their customer buildings and homes, so that demand variations in electric power grids can be more safely and efficiently managed, and services can be delivered more cost effectively


Each of these individual “smart” devices — that will collectively number more than three times the world’s population within just a few years — will potentially collect and produce insurmountable amounts of data.

Chapter 2

Tuning for Performance



In This Chapter

- ▶ Differentiating structured and unstructured data
 - ▶ Considering hardware implications
 - ▶ Recognizing performance bottlenecks
- 

In this chapter, you discover the factors that impact the performance of your storage infrastructure so you can tune your databases for optimal performance.

Structured and Unstructured Data

Data exists in both structured and unstructured forms. Structured data includes relational databases and spreadsheets. However, much of the data growth over the past decade (see Chapter 1) has been in the form of unstructured big data. Unstructured data is not as easily ingested, stored, queried, and analyzed as structured data, and therefore has different storage performance implications.

When designing your storage infrastructure, it's important to understand the nature of your data so that you can create and manage appropriate quality of service (QoS) policies for your mission-critical databases.

Hardware Considerations

Storage hardware has a major impact on database performance. Several underlying storage/infrastructure technologies that are enabling IT to modernize their databases and improve performance include

- ✓ **Flash storage:** Because database performance is heavily dependent on throughput (measured in Input/Output Operations Per Second, or IOPS), the availability of cost-efficient enterprise-grade flash storage is changing the paradigm for database storage. Flash technology is now commonly deployed in all flash storage arrays, hybrid storage pools (combining solid state drives, or SSDs, and traditional hard disk drives, or HDDs) as flash cache in storage controllers, and as flash pools to accelerate random reads and writes.
- ✓ **Converged infrastructures:** Converged infrastructures combine servers, storage, networking and software into a purpose-built solution that can be tuned, balanced and sized specifically for the workloads handled by your databases. There are advantages in this approach as a converged infrastructure solution can simplify and accelerate deployments, while providing a pre-configured and proven platform so IT doesn't need to over-provision resources.

✓ **Server virtualization and advanced data management:** Server virtualization can further reduce total cost of ownership (TCO) and centralize management through shared storage options and application integration features. Virtualization often reduces an organization's license costs, especially in database workloads that are typically storage-bound. The server cores are frequently idle because there are pending I/Os. Virtualization allows you to securely consolidate workloads to ensure that you are using as close to 100 percent of your CPU resources as possible. The organization can achieve further savings and performance improvements by using advanced storage techniques such as compression and deduplication, saving overhead that might otherwise impact database performance. Finally, a virtualized environment is easier to run with all flash because it removes a variable from the overall performance equation.



The performance and predictability of flash storage enables organizations to confidently virtualize their database environments.

Another important hardware consideration is multi-protocol support. In order to maximize your database storage options, your storage infrastructure should support standard protocols including

- ✓ Common Internet File System (CIFS)
- ✓ Fibre Channel (FC)
- ✓ Fibre Channel over Ethernet (FCoE)

- ✓ Internet Small Computer Systems Interface (iSCSI)
- ✓ Network File System (NFS)
- ✓ Server Message Block (SMB)

Identifying Performance Bottlenecks

Latency and IOPS are key performance indicators for databases. Latency is the speed at which the database can process data, and IOPS measures the amount of data processed.

In a typical database environment with legacy hard disk drive storage, throughput is limited to how fast the hard drive system can serve up the data. At some point, the hard drive system will “top out” on throughput, where latency will increase beyond required thresholds. In order to bring latency back down to acceptable levels and increase throughput, you need more hard disk drives. Alternatively, you can use flash storage in combination with advanced storage management capabilities to drastically improve the performance of your database storage environment. Using flash storage also has the practical effect of enabling you to maximize the CPU utilization of your database servers, thereby maximizing the value of your existing server investment. Oracle Database licenses are almost always licensed on a per-CPU basis, and Microsoft SQL Server is licensed by CPU sockets or cores, thus an idle CPU might be considered an idle investment.

Chapter 3

Accelerating SQL Server Performance

In This Chapter

- ▶ Taking advantage of SQL Server 2014
 - ▶ Recognizing key database storage capabilities
 - ▶ Exploring flash storage options
 - ▶ Simplifying SQL storage management tasks
-

Database performance is a top challenge in today's enterprise. Database administrators (DBAs) must ensure their mission-critical database infrastructures deliver the reliable performance and 24/7 availability to meet service-level agreement (SLA) requirements for their business.

In this chapter, you learn how to optimize your SQL Server database infrastructure with NetApp storage solutions.

Upgrading to SQL Server 2014

Many organizations today are running mixed environments and older versions of Microsoft Windows and SQL Server that support legacy applications and databases.

These older operating system (OS) and database versions are quickly becoming performance bottlenecks that stifle innovation and erode competitive advantage for the enterprise, while introducing growing security and compliance risks as they near their end of life cycle support.

Upgrading to Microsoft SQL Server 2014 empowers organizations to leverage advanced new features and capabilities that can accelerate your database applications and make your IT infrastructure more scalable, available, and efficient, including

- ✔ **In-Memory Online Transaction Processing (OLTP) Engine:** The In-Memory OLTP Engine provides significant performance improvements by moving selected tables into memory. The built-in Analyze, Report, and Migrate (ARM) tool helps you select the tables to be moved into memory and the stored procedures to compile into machine code for high-performance execution.
- ✔ **Buffer Pool Extension (BPE):** The BPE feature uses high-speed flash drives to extend SQL Server 2014's buffer pool and thereby improves performance in systems that have maxed out their memory capability. BPE provides the best performance gains for read-heavy OLTP workloads.
- ✔ **Updateable Columnstore Index:** Microsoft introduced the columnstore index in SQL Server 2012, providing significantly better performance for data warehousing types of queries. The original columnstore indexes required that the underlying tables be read-only. SQL Server 2014 eliminates that requirement and enables updates to be performed on the underlying table without first needing to drop the columnstore index.

- ✔ **Improved scalability:** SQL Server 2014 scales up to 640 logical processors and 4 terabytes (TB) of memory on a physical server. It scales to 64 virtual processors and 1TB of memory when running in a virtual machine (VM).
- ✔ **High availability and disaster recovery:** AlwaysOn Availability Groups can have up to eight replicas (increased from four in SQL Server 2012). You can mix synchronous and asynchronous replicas to provide both high availability and disaster recovery. AlwaysOn Availability Groups can also be integrated with Windows Azure for cloud-based disaster recovery.
- ✔ **Resource Governor for I/O:** First introduced in SQL Server 2008, the Resource Governor lets you limit the amount of CPU and memory that a given workload can consume. SQL Server 2014 extends that capability to manage storage input/output (I/O) usage. The SQL Server 2014 Resource Governor can now limit the physical I/O issued for user threads in a given resource pool, providing more predictable application performance.
- ✔ **New backup capabilities:** SQL Server 2014 provides a new ability to back up to Windows Azure that's completely integrated with SQL Server Management Studio and Transact SQL (T-SQL). You can use SQL Server Azure backups to restore on-premises databases or to quickly restore database backups to a Windows Azure VM.



NetApp supports all of the new SQL Server 2014 features, including in-memory databases, BPE on SSDs, and eight-replica-node AlwaysOn Availability Groups.

Key Database Storage Capabilities

Storage is one of the most important factors to consider when optimizing your SQL Server 2014 database environment. Advanced storage capabilities such as consolidation and flexible deployment options enable a robust database architecture that leverages the latest technological innovations.

Consolidating storage improves performance by taking advantage of flash hybrid and all flash solutions, and provides storage efficiencies with resource pooling, data deduplication, and thin provisioning.

Multiprotocol support for data access across multiple hypervisors — such as Microsoft Hyper-V, Oracle VM, and VMware vSphere — provides maximum flexibility for deploying solutions on-premises and in the cloud.

The NetApp clustered Data ONTAP operating environment is a key component of the NetApp storage solution for SQL Server and is the software foundation that underlies the NetApp FAS storage system. The system optimizes SQL Server storage according to your business requirements with capabilities such as

- ✔ **Quality of Service (QoS):** Built-in QoS prioritizes database throughput for more predictable performance in business-critical applications. This ensures that less critical resource-intensive applications don't consume too much throughput capacity, thereby causing performance degradation across the entire SQL Server infrastructure.

- ✔ **Scalability:** NetApp storage provides seamless scalability for OLTP workloads and data warehouses that require petabytes of storage.
- ✔ **High availability:** NetApp reduces or eliminates planned and unplanned downtime through non-disruptive operations. You can also move production databases to higher performance storage subsystems on-the-fly.
- ✔ **Data protection:** NetApp data protection tools are built on top of its innovative Snapshot technology. Snapshot captures a point-in-time image of a data volume, and only consumes additional storage as the volume changes, for up to a 10:1 storage space savings.

Deploying Flash Storage

When you're evaluating database storage options, consider your performance needs along with your capacity and cost requirements. Figure 3-1 presents a summary of the range of enterprise storage options. Server flash — typically implemented directly into the server via Peripheral Component Interconnect Express (PCIe) cards — is the highest-performing option, but it also tends to be the most expensive and offers the lowest storage capacity. All flash arrays are the next highest performing option. They provide slightly lower performance than server flash, but they also provide higher capacities. Hybrid arrays combine flash storage and traditional hard disk drives (HDDs) for a combination of high performance and larger storage capacity. Traditional HDD storage is the slowest of the enterprise storage options, but it also provides the highest capacities and is the least expensive option.

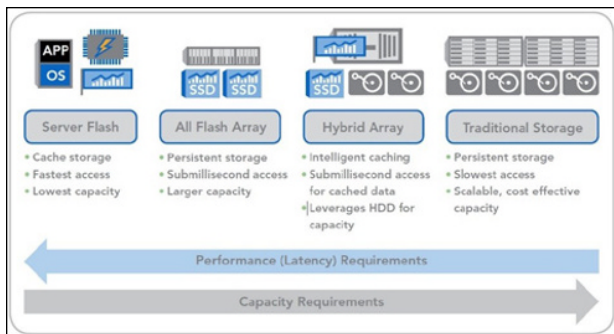


Figure 3-1: Enterprise storage options.



TIP NetApp offers a complete portfolio of both hybrid and all flash storage solutions. Hybrid flash systems such as the FAS and E-Series Storage Systems combine flash with traditional HDDs. NetApp's Virtual Storage Tier (VST) — a self-managing, data-driven service layer for NetApp storage — achieves maximum flash performance benefits with less flash storage. VST uses intelligent caching to leverage flash-based technology with minimal I/O and CPU overhead. Data is dynamically and automatically moved between flash and HDD storage based on actual data usage. VST is compatible with both SAN and NAS environments.



TECHNICAL STUFF NetApp AFF All Flash storage system and EF-Series solutions are all flash systems that deliver very high I/O per second (IOPS) and extremely low sub-millisecond latency for business-critical database applications that

demand the highest levels of performance. AFF features a scalable architecture that delivers up to 2 million IOPS and as low as 800 microseconds of latency.

RP Data turns real estate property data into actionable information with NetApp EF-Series Flash Arrays

RP Data is the largest provider of real estate property valuation information, analytics and risk management services in Australia and New Zealand. Owned by CoreLogic, one of the world's largest data and analytics companies, RP Data takes in large volumes of data every day. Real estate and property data, including images and financial information, is the foundation of its business. Success hinges on how quickly the company can acquire, cleanse, link and store the information to provide timely reports to its customers, which include real estate and financial services companies as well as homeowners, investors and government organizations.

Challenges

With more than 100TB of new data coming in every year, RP Data stores an average of five images of nearly every piece of real estate in Australia and New Zealand. The company's property management database, which runs on Microsoft SQL Server 2012, is its most critical workload. The database contains more than 500 million property decision points spanning more than 9 million properties. Fast database transfer times are absolutely crucial to providing accurate, relevant, and actionable data to

(continued)

(continued)

customers faster. As the database grew, storage I/O became a concern.

To overcome performance issues and improve the customer experience, RP Data wanted to implement storage technology that would allow it to

- ✔ Reduce database transfer times
- ✔ Produce more frequent customer reports
- ✔ Provide high availability and rock-solid reliability

Solution

Initially, RP Data considered overprovisioning its storage with more partially filled spinning drives to solve the immediate storage I/O problem. However, with space at a premium at the company's two active-active Brisbane data centers, it needed a solution that could provide greater density and long-term scalability. RP Data turned to NetApp.

NetApp flash storage technologies were instrumental in helping RP Data achieve its business goals for faster data processing. "We were confident that the NetApp EF-Series was the best flash solution for our business," says Adrian Jansz, head of Information Communication and Technology (ICT) at RP Data. "It's a mature, enterprise-ready flash offering that doesn't sacrifice reliability for speed."

RP Data's team leveraged the EF540 reference configuration for Microsoft SQL Server 2012. "It was reassuring to base our deployment on a well-thought-out reference

configuration,” says Jansz. Incoming data from three financial institutions is stored on solid-state drives (SSDs) on a NetApp EF540 flash array connected to servers by Fibre Channel. The extract, transform, and load (ETL) process runs and property valuation reports are generated.

Results

- ✔ **Driving down latency:** RP Data was able to reduce database transfer times by over 70 percent (50 minutes versus 3 hours), decrease the time needed to process reports by approximately 70 percent (3.5 hours versus 12.5 hours) and achieve in excess of 100,000 IOPS running their ETL process in real-world use cases.
- ✔ **Simplifying storage management:** RP Data engineers use NetApp SANtricity Storage Manager software to manage the EF-Series array. SANtricity allows users to achieve maximum performance and utilization through configuration flexibility and custom performance tuning. All management tasks can be performed while the storage remains online with complete read/write data access, enabling users to make configuration changes and conduct maintenance without disrupting I/O.
- ✔ **Maintaining high availability:** The NetApp EF-Series flash array captures and monitors extensive diagnostic data for comprehensive fault isolation and analysis. SSD wear life is proactively tracked, and alerts are issued when the threshold is reached. To allow recovery in the event of an environment issue or component failure, the array includes fully redundant I/O paths with automated failover.

Simplifying Database Backup, Restore, and Cloning Operations

Organizations typically require multiple backup copies of production databases for various purposes, such as

- ✔ **Disaster recovery and business continuity:** Traditional approaches to database backup and recovery can be complex and time-consuming, and may impact server performance.
- ✔ **Development and testing:** Application development, patch testing, upgrades, and integration often require production database clones.
- ✔ **Training:** Real data is often required for various training scenarios, such as training for new users and on new or upgraded applications.
- ✔ **Troubleshooting:** When help desk incidents are escalated, various IT teams including applications, database, networking, and storage may need to get involved. These teams must often re-create the problem with a replica of the production database in order to troubleshoot and resolve the problem.

NetApp SnapManager for SQL Server (SMSQL) automates and simplifies SQL database backup, restore, and cloning operations, and provides data protection for SQL applications by consolidating the database backup of multiple SQL instances and databases.

SMSQL capabilities and features include

- ✔ **Fast backup:** Creates fast, efficient backups without negatively affecting the storage system's performance with NetApp SnapShot technology.

- ✔ **Rapid recovery:** Utilizes NetApp SnapRestore for point-in-time data recovery, reducing SQL data recovery from hours to minutes.
- ✔ **Efficient cloning:** Leverages NetApp FlexClone technology to quickly create virtual SQL database clones and reduce storage requirements up to 10X.
- ✔ **High-availability replication:** Tight integration with SQL Server 2014 AlwaysOn Availability Groups accelerates the setup of AlwaysOn Availability Groups and enables rapid backup and restore of SQL databases in these groups. You can also mirror AlwaysOn Availability Groups to remote locations using NetApp SnapMirror.
- ✔ **Flexible cloud deployment options:** NetApp supports SQL backups and restores to Microsoft Azure and Amazon AWS, as well as SQL database failover between different cloud environments. NetApp Private Storage for Cloud (NPS), provides complete control of SQL databases on dedicated, private NetApp storage located 'next to' multiple clouds in Equinix data centers to meet compliance and data sovereignty requirements.
- ✔ **Self-service tools:** Simple but powerful, self-service GUI tools, as well as a complete set of PowerShell 3.0 cmdlets ("command-lets"), enable DBAs to streamline and automate database backup, restore and cloning operations.

Chapter 4

Accelerating Oracle Database Storage Performance

In This Chapter

- ▶ Deploying Oracle Database on NetApp storage
 - ▶ Tapping into Data ONTAP's data protection features
-

Designing storage infrastructure to optimize Oracle Database workloads is complicated. Delivering the right level of performance for each workload without sacrificing reliability and availability is of paramount importance. But because there is so much variability between Oracle Database implementations, there are many design considerations and issues that must be addressed. A database project could consist of a single, mission-critical database or thousands of legacy databases, ranging in size from a few gigabytes to hundreds of terabytes. Clustering and virtualization options introduce still more variables. The right solution depends on both the technical details of the implementation and the business requirements driving the project.

In this chapter, you learn how to optimize your Oracle Database infrastructure with NetApp all flash storage array solutions.

Storage Infrastructure for Oracle Database Workloads

Not all Oracle workloads are created equal. A solution that works for online transaction processing (OLTP) may be completely wrong for your data warehouse. You could deploy different storage for each workload, but why create more complexity?

NetApp AFF All Flash storage system delivers superior performance, availability and scalability, transforming storage into a strategic operational resource for Oracle Database workloads (see Figure 4-1). All flash arrays provide the raw performance to meet the needs of the most demanding Oracle Database workloads and absorb unexpected spikes in demand.

AFF addresses enterprise storage requirements with high performance, combined with superior flexibility and best-in-class data management. Built on NetApp's clustered Data ONTAP storage operating system, AFF speeds up your Oracle databases without compromising efficiency, reliability, or the flexibility of your database operations. As a true enterprise-class all-flash array, it accelerates, manages, and protects your business-critical data. With AFF systems, you can

- ✓ Accelerate databases with four to twelve times higher Input/Output Operations Per Second (IOPS) and twenty times faster response

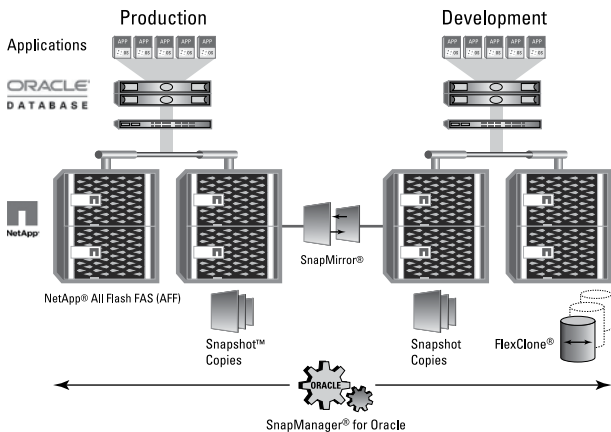


Figure 4-1: All flash arrays flexibly adapt to meet the performance and management needs of your entire Oracle environment.

- ✓ Reduce solid-state drive (SSD) storage by five to ten times on average with data reduction technologies
- ✓ Move data between the performance tier and the capacity tier, on premises or in the cloud
- ✓ Scale out to 24 nodes in a cluster and move data between flash and hard disk drive (HDD) tiers non-disruptively

With flexible scale-out capabilities, support for both storage area network (SAN) and network-attached storage (NAS) protocols, plus non-disruptive operations, quality of service (QoS), and Integrated Data Protection,

AFF can support all of your Oracle environment needs, including

- ✓ **OLTP:** OLTP applications process hundreds or even thousands of transactions per minute, generating a mix of reads and writes on the storage system. In many cases write activity is quite high. High storage throughput and low latency are critical, as is high availability. All flash arrays can address OLTP throughput and latency demands and integrated write cache technology drives write-intensive OLTP environments.
- ✓ **Data warehouse:** All flash arrays have the flexibility to address diverse data warehouse requirements. A typical data warehouse environment contains large quantities of data, with read-intensive reporting operations and extremely write-intensive data ingest.
- ✓ **Oracle applications:** Your Oracle application environment will benefit from the high availability and flexible performance options of an all flash array. Consolidate multiple applications on a single storage system while simplifying and accelerating provisioning for new and existing applications. QoS and the ability to fine-tune flash usage for each application allow you to address unique application demands from a single storage system.
- ✓ **Virtualization:** With the increasing focus on virtualization and the cloud, more organizations are running hypervisors such as VMware and Oracle VM (virtual machine) — and running Oracle Database instances within these hypervisors. Virtual environments create unique data duplication and performance challenges for storage systems. AFF lets you consolidate your server virtualization environment on a single storage

system. NetApp deduplication technology eliminates redundancy in virtual environments and achieves maximum acceleration with minimum flash using cache amplification, which allows a single cached block of deduplicated data to be accessed by multiple virtual machines.



Tools such as the NetApp Virtual Storage Console for VMware vSphere and the NetApp Plug-in for Oracle VM enable provisioning and cloning of virtual machines quickly and space efficiently from within native tools. This significantly reduces the cost and complexity of managing virtual machines that rely on NetApp storage.

Drilling for opportunities at Apache Corporation with fast access to seismic data

Data is constantly growing, and seismic models are becoming more complex. For years Apache IT has used NetApp flash solutions to keep pace with scientists' needs, which are always evolving.

Challenges

A popular trend in the industry and on the horizon for Apache is a new computing model for remote visualization of the Earth's subsurface that will greatly improve decision making in the field. Instead of deploying onsite infrastructures, which can lead workers into locales with unstable power grids or challenging political climates, Apache will deliver Schlumberger Petrel seismic-to-simulation

(continued)

(continued)

software and other critical applications to scientists via NVIDIA GPU-accelerated virtual desktops. However, for the solution to work storage latency had to be extremely low.

Solution

Seeking to build upon its past success and remove disk bottlenecks for remote, high-end graphics, Apache turned again to NetApp. With critical operations on the line, Apache deployed an AFF system in its new Aberdeen, Scotland, data center to support remote visualization throughout the region, encompassing Europe, the Middle East, and Africa (EMEA). Additionally, Apache upgraded the NetApp storage in its U.S. data centers to the clustered Data ONTAP operating system. Hybrid flash clusters in Houston and Dallas, Texas now support exploration and corporate IT, so data is always available.

With the flash solution in place, Apache's scientists will now be able to access 3D models in near real time, allowing them to achieve better, faster results. Scientists can generate more iterations of seismic data models in less time, with access to data on the depths and paths of thousands of existing wells and the ability to visualize seismic data remotely. At the same time, Apache can shift from onsite infrastructures in challenging environments to a centralized, global data center model, reducing business risk and costs.

NetApp flash solutions have reduced storage latency to microseconds for Apache's core applications and Oracle databases, giving back hours of productivity every day for both scientists and business users. In addition, NetApp solutions are enabling Apache to adopt remote visualization sooner than some of its competitors, helping the company "do more with less" and compete in a volatile energy market.

Protecting Data with NetApp Data ONTAP

The NetApp clustered Data ONTAP storage operating system is a powerful data management platform with native capabilities that include inline compression, non-disruptive hardware upgrades, and the ability to import a logical unit number (LUN) from a foreign storage array. Up to 24 nodes can be clustered together, simultaneously serving data through numerous file sharing and communication protocols including Common Internet File System (CIFS), Network File System (NFS), and Oracle Direct NFS (dNFS), among others.



Oracle dNFS is an integrated NFS client that's optimized for Oracle Database file I/O access patterns. The benefits of dNFS include

- ✓ **Performance:** Concurrent direct I/O and asynchronous I/O allows higher throughput and greater overall performance than traditional NFS while consuming fewer system resources.
- ✓ **Optimization:** Automatically load balances requests across all specified paths.
- ✓ **Scalability and fault tolerance:** Supports up to four parallel network paths.

In addition, NetApp Snapshot technology is the basis for creating tens of thousands of online backups and fully operational database clones. Native data protection capabilities and features in Data ONTAP include

- ✓ **Write Anywhere File Layout (WAFL) integration:** NetApp WAFL integrates with RAID (Redundant Array of Independent Disks) to provide data protection without negatively impacting write

operations. Write operations are coalesced in RAM and prepared as a complete RAID stripe, including parity generation. There is no need to perform a read to service a write, which means that Data ONTAP and WAFL avoid the RAID penalty. Performance for latency-critical operations such as redo logging is unimpeded, and random data-file writes do not incur any RAID penalty resulting from a need to regenerate parity.

- ✔ **Snapshots for application development:** NetApp SnapManager for Oracle enables efficient administration of critical Oracle-based business processes. SnapManager for Oracle creates database clones in seconds, either on primary storage or straight to development and test environments. Use clones to engage in parallel QA (quality assurance), development, testing, and other processes to deploy applications faster.
- ✔ **Snapshot-based backup and recovery:** The most important consideration for a file-system layout is the plan for leveraging NetApp Snapshot copies. A *crash-consistent backup* of a database requires the capture of the entire database structure, including datafiles, redo logs, and control files, at a single point in time. If a database spans volumes, a consistency group (CG) snapshot must be created. When more granular point-in-time recovery is required, a *Snapshot-protected hot backup* is preferable. This procedure yields a set of snapshots containing datafiles in hot-backup mode and the critical archive logs generated while in hot-backup mode.

NetApp volume-based SnapRestore (VBSR) technology allows you to almost instantly revert the state of a volume to an earlier point in time. In addition, many databases can be restored by using file-based single-file SnapRestore (SFSR) or by simply copying files from the snapshot back into the active filesystem.

CERN unlocks the secrets of the universe with Oracle Database and NetApp storage

The European Organization for Nuclear Research (CERN) conducts fundamental research about how the universe works. CERN's Information Technology department manages a broad array of IT services and data for a demanding scientific community that comprises nearly half of the world's particle physicists. "They will turn the knob until it breaks," remarks Frédéric Hemmer, head of CERN's IT department.

Challenges

The big science being done at CERN introduces equivalently big data management challenges. The Large Hadron Collider (LHC), for example, creates 600 million collisions per second, producing raw data at the rate of 1 million gigabytes per second. Current experiments produce more than 20 petabytes (PB) of new data annually.

CERN IT delivers this functionality faces the universal challenge of providing more services with limited funding

(continued)

(continued)

and the same or decreasing resources. CERN continually balances technical demands for performance, reliability, and scalability with the constancy of financial constraints.

Oracle technology is used throughout the organization and plays a critical role in accelerator control systems, engineering and administrative applications, and LHC experiments.

On the data storage side, essential requirements include manageability, availability, and scalability to respond to fast-changing or unexpected requirements. For example, heavy lead ions cause especially complicated collisions that can make estimating data rates an inexact science. In one case, incoming data rates were five times higher than predicted. Hemmer further quantifies: "Data can come into our computer center at rates up to 6GB per second — that's equivalent to the contents of two DVDs every three seconds."

Solution

In 2007, CERN selected NetApp technology for the LHC logging database built on an Oracle database with Real Application Clusters (RAC) technology. CERN has since unified its entire Oracle infrastructure on NetApp and today stores 99 percent of all Oracle data on NetApp solutions.

Eric Grancher, database services architect at CERN, says deploying Oracle databases on NetApp enables his team to balance requirements for efficiency with necessary stability, performance, and scalability. He cites vital functionality:

- ✔ **Oracle Direct NFS (dNFS)** enables multiple paths to storage. This technology contributes to scalability and, because it bypasses the server operating system, typically doubles the performance of traditional NFS.
- ✔ **NetApp FlexClone software** enables efficient creation of temporary, writable copies. CERN required space-efficient Snapshot copies and writable copies of large databases, but also needed to make sure that replication processes did not impact performance.
- ✔ **NetApp Data ONTAP 8 operating in Cluster-Mode** makes it possible to maintain peak application performance and storage efficiency by adding storage and moving data without disrupting ongoing operations.

Chapter 5

Going Beyond Performance in the Enterprise

In This Chapter

- ▶ Creating a cloud data fabric
 - ▶ Recognizing the importance of reliability
 - ▶ Making DevOps a reality
 - ▶ Looking to the cloud for backup and recovery
-

In this chapter, you explore the benefits of a cloud data fabric, enterprise requirements for reliability and availability, how your database storage infrastructure can accelerate DevOps in your organization, and innovative ways to address backup and recovery challenges.

NetApp's Data Fabric Vision

NetApp's approach to hybrid cloud facilitates the seamless connection of cloud resources, with a highly efficient transport between systems and clouds and a single purview of data management over a cloud data fabric.

NetApp's storage operating system, Data ONTAP, is the foundation underpinning this cloud data fabric. Data ONTAP has evolved from humble beginnings to become a data framework that simultaneously manages multiple hybrid cloud endpoints through its four variations:

- ✔ **Data ONTAP** offers an adaptable, always-on storage infrastructure for enterprise applications and private clouds.
- ✔ **NetApp Private Storage** extends the reach of Data ONTAP by placing application data next to, but not inside of, public clouds — allowing enterprises to build hybrid clouds that offer the speed and control of private clouds.
- ✔ **Cloud ONTAP** is a software-only version of Data ONTAP that can be run on top of cloud services, such as Amazon AWS, bringing unified data management to the public cloud.
- ✔ **Data ONTAP Edge** is a virtual, VMware-ready version of Data ONTAP designed for remote offices that don't require dedicated storage.



Each variant of Data ONTAP can operate independently or in conjunction with the other variants, creating a data fabric with common storage commands and the combined intelligence to universally apply data management policies across an entire cloud ecosystem.

Having a fabric means that data is free to move dynamically across all cloud resources. Businesses realize greater efficiencies by pairing workload requirements with cloud economic models in real time, without disruption.

NetApp's universal data platform and data portability, helps organizations remove barriers between private and public clouds, creating a reliable data fabric that enables data stewardship to be maintained across all resources.

Availability and Reliability

The importance of data availability — guaranteeing that database queries can be completed and information delivered to critical applications — continues to rise as organizations struggle to meet business demands and address competitive challenges.

NetApp AFF All Flash storage system delivers “five nines” or greater availability through a comprehensive approach that combines highly reliable hardware, innovative software, and sophisticated service analytics.

Advanced hardware features, including alternate control path and an integrated service processor, provide the utmost reliability to protect your data.

Leading software capabilities include non-disruptive operations, quality of service (QoS), and Integrated Data Protection, as well as service analytics that constantly monitor risk signatures and create alerts to proactively address issues that might affect operations.

Accelerating DevOps

DevOps describes a cultural trend of promoting greater collaboration and communications between application developers and operational IT teams to support a rapidly evolving business environment. DevOps extends Agile development principles throughout the

IT organization, so that new software can be developed, tested and deployed quickly, frequently and reliably.

A key to accelerating DevOps is the ability to rapidly provision infrastructure and deploy and clone databases for application development and testing.

NetApp's centralized storage management tools enable automated, self-service for many of the storage and database operations necessary to support a DevOps environment.



NetApp FlexClone technology can be used to create instant virtual copies of production databases, storage logical unit numbers (LUNs), and data volumes for use in development and staging environments.

Data Protection

Data protection has become increasingly important in recent years as organizations recognize data loss — whether due to natural or man-made events, such as fires, flooding, equipment failures, malware, or cyberattacks — are among the greatest risks to their business operations. Unfortunately, the explosive growth of data and complexity of application and database architectures haven't made data protection any easier.

Backup to the cloud is a rapidly emerging data protection strategy that leverages on-premises backup infrastructure to provide the first tier of data protection using snapshot-based backups, then replicates snapshot-based backup data to an integrated cloud-based backup target in order to provide the second tier of data protection offsite.

Chapter 6

Ten Key Database Storage Capabilities

In This Chapter

- ▶ Identifying key performance capabilities
- ▶ Considering other important storage features

This chapter gives you ten, okay eight, important database performance and storage optimization capabilities you need:

- ✓ **Flash technology:** Flexible all flash and hybrid storage arrays with intelligent data management software maximize database performance. Flash storage technology is commonly implemented as server and storage cache, as well as a replacement for traditional hard disk drives (HDDs).
- ✓ **Intelligent data management:** Advanced storage techniques such as compression and deduplication are important capabilities that improve database performance and reduce storage requirements.

- ✔ **Advanced data protection:** Cloning and backup/restore are critical database operations that can negatively impact production database performance. Look for advanced data protection capabilities including snapshots and mirroring.
- ✔ **Storage Quality of Service (QoS):** QoS enables database and storage administrators to prioritize resources for mission-critical databases and workloads.
- ✔ **Flexible deployment options:** Public, private, and hybrid cloud infrastructures have become integral parts of enterprise cloud computing strategies. A common data fabric ensures your database infrastructure can migrate to different cloud environments as your business needs change.
- ✔ **Multiprotocol support:** Support for different protocols, including Common Internet File System (CIFS), Network File System (NFS), Oracle Direct NFS (dNFS) and Server Message Block (SMB), among others, provides the flexibility to deploy network-attached storage (NAS) and storage area networks (SAN) solutions.
- ✔ **Nondisruptive operations:** Mission-critical databases require 24/7/365 uptime. Planned outages are no longer acceptable. Routine maintenance and other storage operations must now be performed in production environments *without* downtime.
- ✔ **Seamless scalability:** Data growth is accelerating with trends such as the Internet of Things (IoT), cloud and real-time analytics producing ever-increasing amounts of data. Your storage solution must be able to seamlessly scale up and out.

Examine the trends for better database performance

Understand the technical factors that affect database performance tuning, different SQL Server and Oracle Database design considerations, how database storage technologies and innovations can help drive business value, and key performance optimization capabilities to look for in a storage solution.

- *Modernize underlying hardware infrastructure — discover flash storage, converged infrastructure architectures, and sophisticated data management platforms*
- *Upgrade to Microsoft SQL Server 2014 — empowers organizations to leverage advanced new features and capabilities*
- *Use NetApp's approach to hybrid cloud — facilitate the seamless connection of cloud resources*



Open the book and find:

- **How to accelerate SQL server performance**
- **Information on Oracle Database Storage Performance**
- **Ten key database storage capabilities**

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

FOR
DUMMIES[®]
A Wiley Brand



Also available
as an e-book

ISBN: 978-1-119-25908-4
Not for resale

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.