IBM

IBM Real-time Compression Appliance Version 4.1

- See how seamless compression integrates into storage environments
- Get familiar with the leading edge real-time compression solution
- Learn how to set up and configure the appliance

Roland Tretau
Salahaldean Alamayri
Christian Prediger Appel
João Marcelo Barros de Almeida
Andre Luis Giglioli
William Luiz de Souza

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM Real-time Compression Appliance Version 4.1

April 2013

SG24-8112-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (April 2013)

This edition applies to the IBM Real-time Compression Appliance Version 4.1 from February 2013.

Contents

Notices v
Trademarks vi

Preface vii
The team who wrote this book vii
Now you can become a published author, too! viii
Comments welcome ix
Stay connected to IBM Redbooks ix

Part 1. Introduction 1

Chapter 1. Compression technology discussed 3
1.1 Compression technology history 4
1.2 Data efficiency technologies 5
1.2.1 Space efficient technology 5
1.2.2 FlashCopy (space efficient) 5
1.2.3 Easy tiering 6
1.2.4 Archiving and space management 7
1.2.5 Data deduplication 7
1.3 Data compression technologies 8

Chapter 2. Introduction to RTCA design 11
2.1 The RTCA approach 12
2.2 IBM Real-time Compression 13
2.3 IBM Random Access Compression Engine technology 14
2.3.1 More about RACE 14
2.3.2 Link status mirroring 16
2.3.3 Unified Protocol Manager 16
2.3.4 Monitoring and Reporting Manager 17
2.4 Performance considerations 17
2.5 Data integrity 18
2.5.1 Data integrity at the network level 18
2.5.2 Data integrity at the protocol layer 19
2.5.3 Data integrity at the storage layer 20
2.6 Storage advantages 21

Chapter 3. What is new in RTCA v4.1 23
3.1 New appliance: STN7800 24
3.2 Revised terminology for compression modes 24
3.3 Improved GUI support of NFS subexports 26
3.4 Autosupport feature enhancements 26
3.4.1 Support for Autoconsult 26
3.4.2 SMTP authentication 27
3.5 New ifconfig command for IPv6 27

Chapter 4. Hardware update 29
4.1 Base model configuration 30
4.1.1 IBM Real-time Compression Appliance STN7800 (2452-780) 30
4.1.2 Hardware summary 30

4.2 Front-panel components . . . 31
4.3 Rear-panel components . . . 32
4.4 Technical specifications of STN7800 . . . 33
4.4.1 IBM System x3650 M4 . . . 33
4.4.2 Operating environment . . . 33
4.5 Interoperability matrix . . . 34
4.5.1 Additional storage system support . . . 34
4.5.2 Additional client system support . . . 34
4.5.3 SAMBA interoperability . . . 34

Part 2. Implementation . . . 35

Chapter 5. Implementation planning . . . 37
5.1 Introduction to installation, reporting, validation, testing . . . 38
5.2 Physical installation . . . 38
5.2.1 Solution requirements . . . 38
5.2.2 Architectural planning . . . 38
5.2.3 Network planning . . . 40
5.2.4 Antivirus services considerations . . . 44
5.2.5 Network considerations summary . . . 45
5.3 Logical installation . . . 46
5.3.1 Configuration planning . . . 46
5.3.2 High availability for the compression appliances . . . 47
5.4 Reporting considerations summary . . . 54
5.5 Validating and testing the installation . . . 54

Chapter 6. Software installation or upgrade . . . 55
6.1 Before you begin . . . 56
6.2 Obtaining the software file . . . 56
6.3 Installing and upgrading on a single RTCA . . . 58
6.3.1 Installing and upgrading on a high availability RTCA system . . . 60
6.3.2 Updating the hardware configuration file . . . 61
6.3.3 Configuring system options . . . 62

Chapter 7. Performance considerations . . . 67
7.1 Performance requirements . . . 68
7.2 Performance suggestions . . . 69
7.2.1 Interval between files . . . 69
7.2.2 Compression ratio . . . 71
7.2.3 Transparent versus compressed traffic . . . 72

Related publications . . . 73
IBM Redbooks . . . 73
Other publications . . . 73
Online resources . . . 74
Help from IBM . . . 74

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DIA Data Integrity Assurance®	IBM®	Redpapers™
Easy Tier®	Real-time Compression™	Redbooks (logo) ®
FlashCopy®	Real-time Compression Appliance™	System Storage®
Global Technology Services®	Redbooks®	System x®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

LTO, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Continuing its commitment to developing and delivering industry-leading storage technologies, IBM® is introducing the IBM Real-time Compression™ Appliance for NAS, an innovative new storage offering that delivers essential storage efficiency technologies, combined with exceptional ease of use and performance.

In an era when the amount of information, particularly in unstructured files, is exploding, but budgets for storing that information are stagnant, IBM Real-time Compression technology offers a powerful tool for better information management, protection and access. IBM Real-time Compression can help slow the growth of storage acquisition, reducing storage costs while simplifying both operations and management. It also enables organizations to keep more data available for use rather than storing it offsite or on tape that is more difficult to access, so they can support improved analytics and decision-making.

IBM Real-time Compression Appliance™ provides online storage optimization through real-time data compression, delivering dramatic cost reduction without performance degradation.

This IBM Redbooks® publication is for system administrators and IT architects. It describes the enhancements made in version 4.1 of the Real-time Compression Appliance as compared to previous releases.

This book is a companion to the publication *Introduction to IBM Real-time Compression Appliances*, SG24-7953.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Roland Tretau is an Information Systems professional with over 15 years of experience in the IT industry. He holds Engineering and Business master's degrees, and is the author of many storage-related IBM Redbooks publications. Roland's areas of expertise range from project management, market enablement, managing business relationships, product management, and consulting, to technical areas including operating systems, storage solutions, and cloud architectures.

Salahaldean Alamayri is an IT professional holding two master's degrees, MBA and Master of Science in Electrical Engineering (MSEE), with both management and technical backgrounds that span over a decade in the IT industry. He gained embedded systems experience, storage, and backup and recovery technologies experience, and additional background in database design, performance tuning, recovery, and security. He received multiple awards for his contributions in these fields.

Christian Prediger Appel is a Server Specialist with IBM Global Technology Services® with IBM in Brazil. He has worked with networking, storage, and server components since 2000. Christian worked for an Internet service provider (ISP) and a server management company before joining IBM in 2005, where he designs, implements, and manages projects, and also is an Instructor in the area of server technologies. His focus is on virtualization and cloud. He holds certifications from Microsoft, Citrix, VMware, and IBM, as a certified IT Specialist.

João Marcelo Barros de Almeida is an IT Specialist at IBM in Brazil. He has extensive experience in operating systems and server virtualization, designing, installing, administering and supporting several environments. He is a Microsoft Certified Systems Administrator (MCSA) and Microsoft Certified Trainer (MCT), has worked in the IT industry for 12 years.

Andre Luis Giglioli is an IT Specialist at IBM in Brazil and a NetApp Certified Professional. He supports IBM Business Partners worldwide. Andre is a focal account and subject matter expert, and has more than two years of experience with storages NetApp and N-Series. He works with projects such as deployment and implementation of storage environments.

William Luiz de Souza is a Technical Leader at IBM in Brazil Strategic Outsourcing Delivery with over 10 years of experience in the IT industry. He holds Systems Information and Business master's degrees, and has the IBM IT Specialist Level 2 certification. He is also the author of other Redbooks and IBM Redpapers™ publications. William is experienced and certified in Microsoft, VMware, and Citrix solutions with IBM System x® platforms.

Thanks to the following people for their contributions to this project:

Bertrand Dufrasne
International Technical Support Organization, San Jose Center

Eyal Traitel
Manager, Development Support, IBM Real-time Compression
IBM Israel

Fabrice Labiau, Tom Provost
IBM Belgium

Uwe Heinrich Mueller, Uwe Schweikhard
IBM Germany

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

Introduction

This part of the book introduces basic concepts for the IBM Real-time Compression Appliance. The part provides an introduction to compression technology and hardware features.

This part contains the following chapters:

- Chapter 1, “Compression technology discussed” on page 3
- Chapter 2, “Introduction to RTCA design” on page 11
- Chapter 3, “What is new in RTCA v4.1” on page 23
- Chapter 4, “Hardware update” on page 29

1

Compression technology discussed

A general overview of data compression and data efficiency technologies is provided in the following topics:

- 1.1, “Compression technology history” on page 4
- 1.2, “Data efficiency technologies” on page 5
- 1.3, “Data compression technologies” on page 8

1.1 Compression technology history

For the last 100 years, IT evolved from very large systems that were built to execute simple mathematical operations. Now, small devices are able to execute, generate, and manipulate massive amounts of data. Devices that are able to store and capture data are showing an exponential growth from year to year.

Because of requirements that data must be stored over a long period of time for long-term reference, compliance, and security purposes, new ways to optimize the capacity utilization are needed.

Historically, over the last 200 years, the technologies that are available to reduce the amount of data that is stored or transported from one place to another have greatly improved. One of the first methods for reducing the amount of data was the use of symbols and representations in mathematical format. For example, instead of writing the words “multiplied by,” the related representation used is the asterisk character (*). In the same way, the word “minus” is represented with the dash character (-).

In 1838, the invention of Morse code allowed messages to be transmitted quickly over long distances. Roman letters and Arabic numbers were replaced with symbols that were formed from lines and dots. To reduce the amount of dots or lines that are used to represent each letter, statistical analysis of the commonality of letters was performed.

The most common letters are represented with a shorter combination of dots and lines. The commonality differs for each language, as is the Morse code. For example, in the English language, the letter “c” is represented in Morse code by three dots; the letter “h” is represented by four dots. The representation will therefore consist of seven dots. However, in some languages “ch” is a common combination, so the dots were replaced by lines, and “ch” is represented by four lines, effectively saving transmission time.

Later in the twentieth century, the development of IT technologies raised the need for complex algorithms, able to reduce the amount of data. This approach is done by interpreting the information beyond the simple substitution of specific strings or letters.

One of the first techniques of mathematical data compression was proposed by Claude E. Shannon and Robert Fano in 1949. In the Shannon-Fano coding, symbols are sorted from the most probably to the least probable, and then encoded in a growing number of bits. For example, if the source data contains A B C D E, where A is the most common letter and E is the least common letter, the Shannon-Fano coding will be 00-01-10-110-111.

In 1952 a Ph.D. student at MIT named David A. Huffman proposed a more efficient algorithm for mapping source symbols to unique string of bits. In fact, Huffman has proved that his coding is the most efficient method for this task, with the smallest average output bits per source symbol.

In 1977, Abraham Lempel and Jacob Ziv proposed a method of replacing repeating words with code words. The method was applicable also to a pattern of text such as expressions. This event was the actual “dawn” of modern data compression.

In 1984, Terry Welch improved the algorithm proposed by Lempel and Ziv (also known as LZ78) and developed a method known as LZW. Today, this algorithm is the basis of modern compression techniques that, for example, are used in PKZIP for general file compression, or within GIF and TIFF file formats for images.

1.2 Data efficiency technologies

Data compression technologies can be found in various implementations over the last 15 years. These range from compression at the application level to in-band solutions that are able to compress data as it is transferred from a host to a storage device. During recent years, many new technologies and new concepts were developed to address the need for optimized storage space and more efficient capacity usage.

This chapter briefly describes some of these technologies to help avoid confusion between data compression and other space optimization methods.

Comparison: All the technologies presented in this section are used to optimize the way that the available storage capacity is used. None of those technologies change information but rather optimize how this information is stored or how much of this information is duplicated within the storage system. Data compression technologies are able to interpret data. For that, they are able to provide storage optimization, both regarding the amount of data stored and the performance needed for the storage system used in back-end systems.

1.2.1 Space efficient technology

Space efficient technology is sometimes also referred to as *thin provisioning*. It allows the storage to declare the required capacity at the host level, but to allocate on the physical storage media (hard disk drives) only the actual used capacity in terms of space. Figure 1-1 presents the difference between a traditional volume and a space efficient volume.

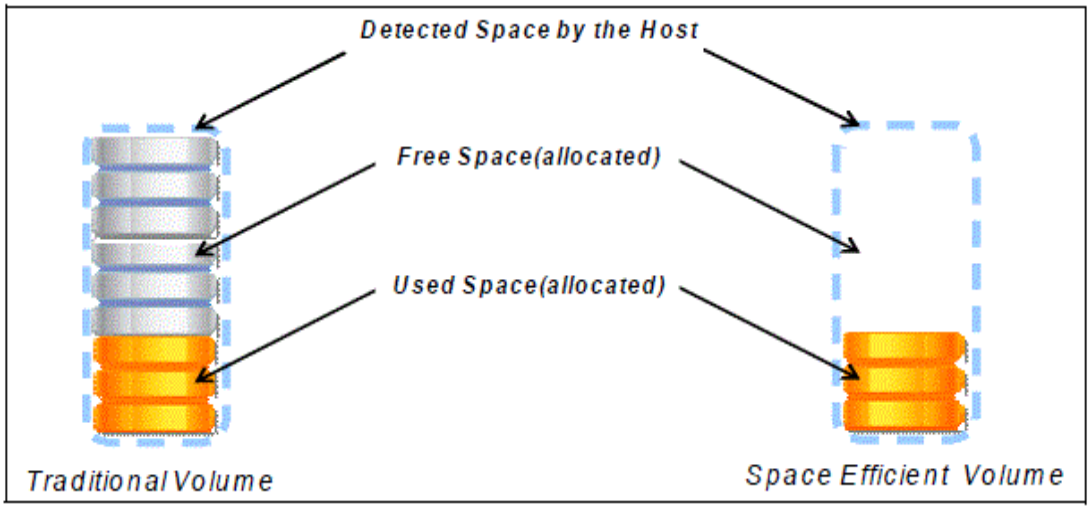

Figure 1-1 Space efficient volume

1.2.2 FlashCopy (space efficient)

Space efficient IBM FlashCopy® technique allows the multiplication of volume data and data instances without multiplying the required space with the number of copies. This technology allows the copy to depend on the source with all the data that is in common. When data is changed from the original volume, those changes are available and reflected only in the updated instance, while the rest of the data remains as a common foundation.

Figure 1-2 presents graphically a data volume and its dependency on a FlashCopy.

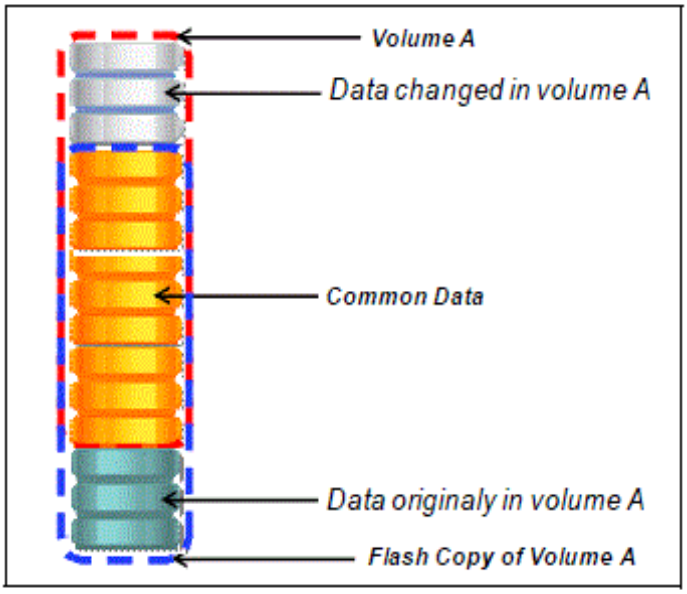

Figure 1-2 Space efficient Flash Copy

1.2.3 Easy tiering

Easy tiering is a technology that optimizes the way data is placed across types of drives (FC/SAS, SATA/SAS_NL, and SSDs). In current environments, the storage system configuration is done by using a rule set. Those rules are well known by a specialist and they usually refer to the result that is expected after configuring a storage system. The most difficult part is to address, in a price-for-performance way, the capacity needs over the performance needs.

Usually when you need capacity, you add drives; and when you need performance, you add drives. In many situations. You need to add drives to accommodate a specific performance, although those drives are not needed in terms of capacity. In these situations, the infrastructure will be oversized in terms of capacity.

IBM Easy Tier® is a feature that can provide an optimized balance between capacity and performance. It does this balancing by identifying specific “hot” zones and by automatically moving them from slower drives to faster drives. Thus, you can profit from using a combination of SSD drives with much slower but higher capacity SATA drives, instead of using many more faster drives to accommodate the needed load.

Figure 1-3 shows easy tiering data allocation.

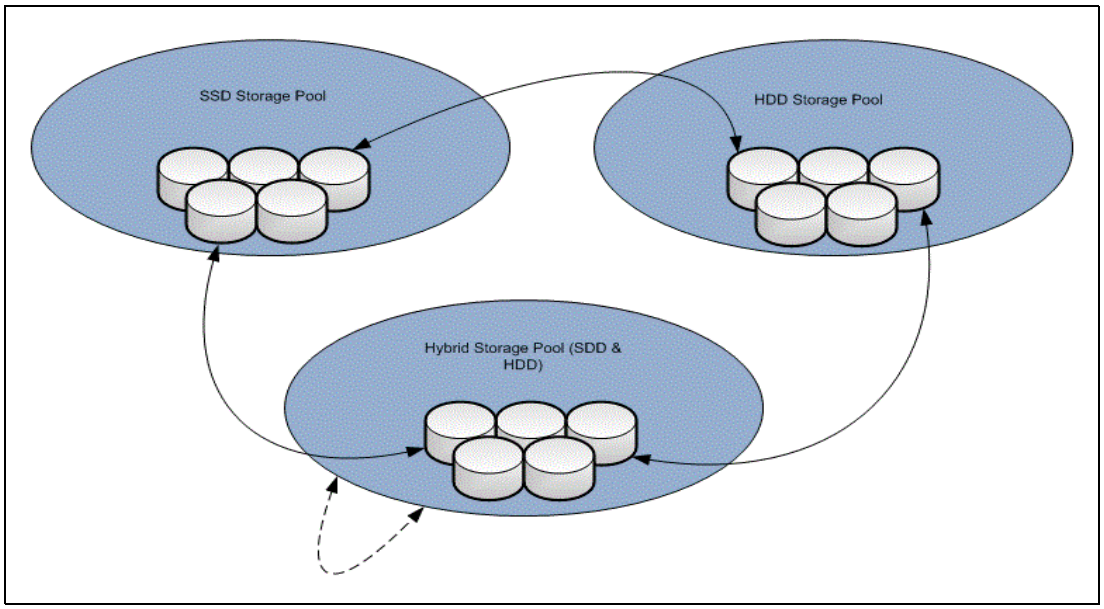

Figure 1-3 Easy tiearing data allocation

1.2.4 Archiving and space management

Archiving and space management, also known as information lifecycle management (ILM), is a concept rather than a dedicated technology solution. It can be used in environments where you need to keep data over a long period of time and you need to optimize the space occupied on various storage types (tapes, hard disk drives, and so on). The old or unused data can be moved out of the production area and placed on a much cheaper support (that is, by moving data from enterprise storage systems to low-end storage systems and then from hard drives to tapes). The historical information is usually kept in a database that holds metadata to provide accurate information about the managed data.

1.2.5 Data deduplication

The data deduplication mechanism identifies identical chunks of data within a storage container and keeps only one copy of each chunk, while all the other logically identical chunks are pointed to this chunk. There are various implementations of this method. One possibility is inline deduplication and another is post-processing. Inline solutions deduplicate information as it is stored. Post-processing solutions deduplicate data after the information is stored in the original format at certain time intervals.

1.3 Data compression technologies

The compression of data has rapidly become a focus for the IT industry. Because of the types of data and the reasons why data is compressed, two methods are used:

- Reversible data compression: This method allows the information to be rebuilt completely with no impact on the quantity or quality of the original information.
- Irreversible data compression: This method synthesizes the information and keeps only the data that is needed. The original information cannot be rebuilt completely to its original form when the data is decompressed.

Examples of irreversible data compression include audio, image, video compression, reports, and graphics that are generated to visualize large amounts of data, and statistics. One clear example can be a situation with the runners at the Olympic Games who manage to be under 10.1 seconds in the 100 meter speed running probe. We can say, for example, that 35% of the runners manage to go under 10.1 seconds and retain this value. But if we do not look at the input data where all runners have their time, we cannot say that runner “X” has that time.

In this way, reversible data compression offers the advantage of the possibility to recreate completely and accurately the input information. In comparison, the irreversible method offers only some specific information that is related to the original information.

Because of the massive amounts of data and the calculation necessary for compressing data, there are two approaches:

- Real-time compression: This method processes the data before it is written to the storage device. The key advantage of this approach is that it reduces the storage resources that are required for a given data set. If done correctly, the capacity-reduction application will preserve the inherent performance of the storage environment. Data, already optimized, will be written to storage and mitigate the capacity explosion challenge at the point of origin. It accomplishes this result by eliminating the need to allocate additional storage capacity as required with post-processing solutions. In addition, because this storage is primary storage, any compression technique must be performed in real time and must maintain the high availability features of the existing storage system.
- Post-processing optimization: These solutions eliminate the need to deal with the performance issues in real time and usually do not have any advanced high availability capabilities. The challenge with post-processing optimization is that it uses storage resources for the capacity-reduction application, which causes a significant impact on storage I/O resources. Post-processing solutions require a read operation and pulling a full copy of the original data to the storage; continued scanning the data, reading the data, and writing the optimized file; and then deleting the original file.

In addition, when data deduplication is factored into the equation, the pointer between the optimized file and its deduplicated file must be written to storage before the deletion of the original file. This process is complex and I/O-intensive. It is also the case that when you want the optimized data, there must be enough space available on disk to uncompress it, which means you do not actually improve storage use.

An important consider of post-processing optimization techniques is that all optimization operations must be scheduled carefully. Operations must be done during a time of day when free resources are available on the subsystem (non-production, non-backup windows). Also, operations must occur after a point in time when snapshots are completed and newly optimized files are written to disk. Managing a schedule such as this is next to impossible, and you risk not finding the time to optimize your capacity before you actually run out of disk space.

Over the years, IBM has introduced a series of reversible, real-time compression algorithms and solutions that are used in wide range of technologies:

- LTO-DC algorithm, which is used in IBM Linear Tape Open, formally known as LTO tape drives
- Streaming Lossless Data Compression (SLDC) algorithm
- Enterprise-class TS1130 tape drives
- Adaptive Lossless Data Compression (ALDC), used by the IBM Information Archive for its disk pool collections
- Random Access Compression Engine (RACE), used inside Real-time Compression Appliance

2

Introduction to RTCA design

Of the main concepts used in the IBM Real-time Compression Appliance (RTCA), the primary concept is the Random Access Compression Engine (RACE) algorithm. This chapter examines the concepts and the way the RTCA product manages to deliver high performance and the technology for preserving data integrity.

The following topics are included in this chapter:

- 2.1, “The RTCA approach” on page 12
- 2.2, “IBM Real-time Compression” on page 13
- 2.3, “IBM Random Access Compression Engine technology” on page 14
- 2.4, “Performance considerations” on page 17
- 2.5, “Data integrity” on page 18
- 2.6, “Storage advantages” on page 21

2.1 The RTCA approach

The industry need for data compression is clearly to be fast, reliable, and scalable. The compression algorithm must assure data consistency and a good compression rate to be implemented. In addition, the data compression solution must be easy to implement. The compression must occur without affecting the production use of the data at any time.

From the technologies described in 1.3, “Data compression technologies” on page 8, and based on industry requirements, the best model was chosen for the RTCA product. This model is a combination of both reversible data compression and a real-time compression algorithm. The RTCA solution compresses the incoming flow of the data with minimum performance impact. A generic overview of the RTCA solution is presented in Figure 2-1.

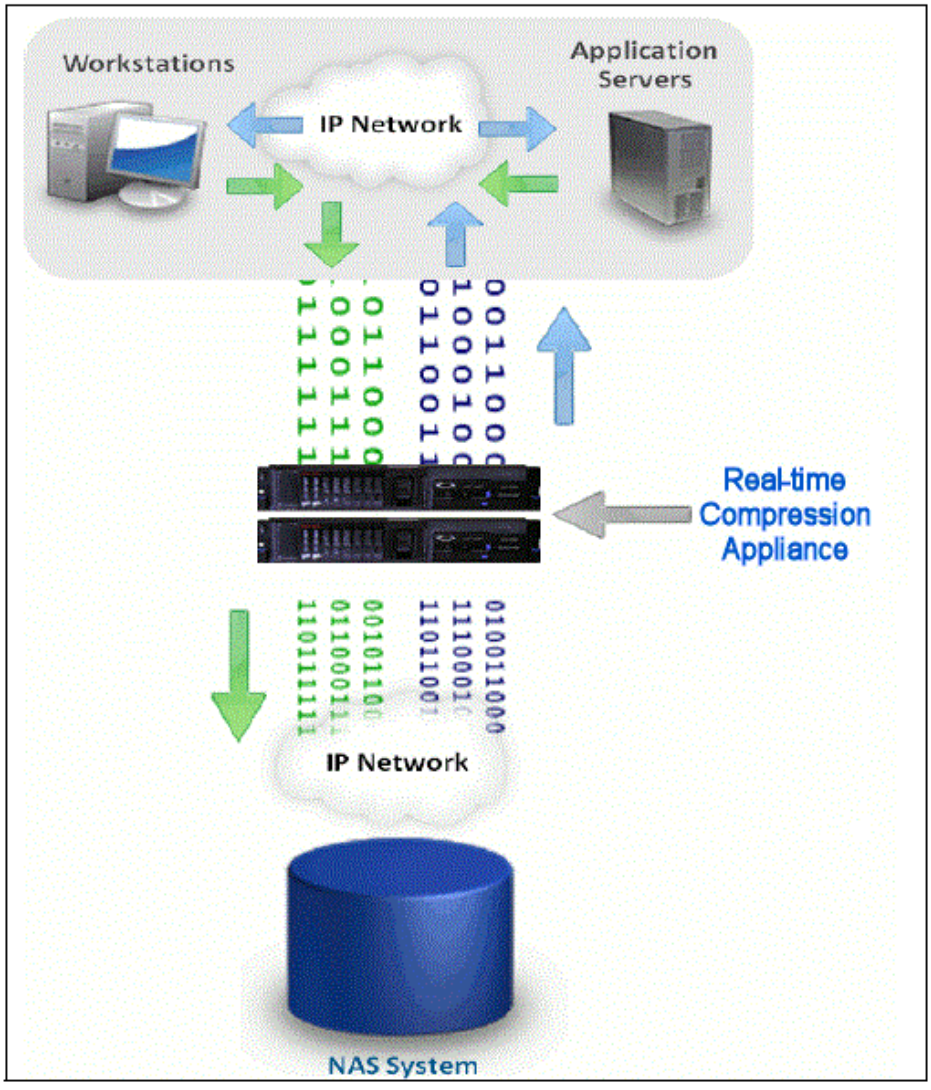

Figure 2-1 Real-time Compression Appliance overview

2.2 IBM Real-time Compression

The IBM Real-time Compression Appliance hardware and software solutions provide data compression for storage optimization. The RTCA models provide online storage optimization through real-time data compression, without performance degradation. The IBM Real-time Compression Appliance STN7800 transparently compresses primary storage by up to 80%. The compression appliances help slow the growth of storage, which reduces the amount of storage to be managed, powered, and cooled.

The RTCA solution increases the capacity of existing storage infrastructure, helping enterprises meet the demands of rapid data growth, while also enhancing storage performance and utilization. All compression appliance models apply patented IBM real-time data compression techniques to primary and existing storage, delivering optimization and savings throughout the entire storage lifecycle. The result can be exceptional cost savings, return on investment (ROI), and operational and environmental efficiencies.

The RTCA software processes incoming data streams, Network File System (NFS) and Common Internet File System (CIFS) protocol data requests, and compresses these data requests in real time to the storage system. The storage system stores the data, and then sends an acknowledgement directly back to the user or application after the write operation is committed and acknowledged by the storage system. From a data integrity perspective, the storage system must acknowledge that the write was committed.

For this reason, the RTCA software does not use a write-cache mechanism. All storage comments come from the system, which preserves the integrity of the data between the storage and the application.

To understand the basic design of the IBM Real-time Compression technology, a review of the basics of modern compression techniques is presented here. These include the Lempel-Ziv algorithm and Huffman coding.

The IBM Real-time Compression Appliance (IBM RTCA) is based on a reversible data compression algorithm that operates in a real-time method.

The IBM RTCA product compresses data on initial write operation to assure that less data is stored on primary storage. Because of this behavior, the storage system must process less data, using less CPU overhead and lower disk spindles utilization. The storage system can therefore serve more requests from its read/write cache, while some reads can be served from the RTCA product's read-ahead cache.

In addition to compressing data in real-time, with the IBM RTCA product, customers can non-disruptively compress existing data that is already saved to disk with the Compression Accelerator utility. Compression Accelerator is a high-performance and intelligent software application that is running on the IBM RTCA product, which, by policy, allows users to compress data that was already saved to disk while that data remains online and accessible by applications and users.

The policies allow users to throttle how decompressed data gets compressed so as not to affect existing storage performance. The ability to compress data that is already stored significantly enhances and accelerates the benefit to users, allowing them to see a tremendous return on their IBM RTCA investment. On initial purchase of an IBM RTCA product, users can defer their purchase of new storage. When a client determines the requirement to purchase additional storage, the IBM Real-time Compression Appliance Version 4.1 will create benefits so that less than half of the physical requirements, which were initially planned for, needs to be purchased.

2.3 IBM Random Access Compression Engine technology

The IBM Random Access Compression Engine (RACE) technology is the core of IBM RTCA products for NAS. RACE technology is based on 35 patents that are not about compression. Rather they define how to make the industry standard LZ compression of primary storage operate in real time and allow random access. The primary intellectual property behind this is the engine: RACE. This engine is located on an appliance in front of any NFS or CIFS deployment, acting as an “intelligent cable” between the IP switch and the storage. No software agents or drivers are required on clients or servers.

The RACE technology (see Figure 2-2) consists of three components:

- Random Access Compression Engine (RACE): Enables random-access data compression without compromising performance.
- Unified Protocol Manager (UPM): Enables transparent support of multiple storage and network protocols, including CIFS and NFS.
- Monitoring and Reporting Manager (MRM): Enables online storage compression trending, analysis, and reporting.

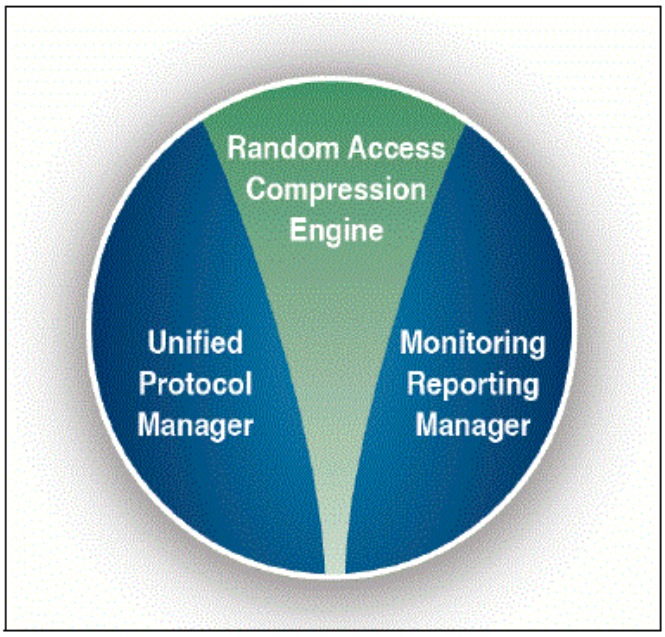

Figure 2-2 RACE technology overview

2.3.1 More about RACE

The traditional compression technologies start from a constant file size and, after compression, the result is a variable file in terms of capacity. Therefore, when using large data chunks, the performance impact is high. However, when using small data chunks, although the performance impact is small, the compression ratio is also small. Over time, there are many disadvantages that can occur. These include the need for garbage collection, poor performance while the volume of the data increases, or losing parts of metadata, such as the date of creation, date accessed, user rights, or modification dates. Another issue can be fragmentation in the target storage space. This issue occurs because, after the file is stored in its original size, the result of compression is stored in a new zone and then the input is deleted.

The Random Access Compression Engine starts from an unknown data stream and compresses data from the host. The resulting compressed file keeps all attributes from the original; metadata is not changed. Also, because of the inline approach, there is no need at the storage level to write original data, read it, write the result of compression, and finally, delete the initial file. At the end, there will not be any garbage or fragmentation on the storage system. The performance needed at the storage level will be decreased because the writes and reads will be made only in compressed format instead of a complete one.

A logical overview of the Random Access Compression Engine is presented in Figure 2-3.

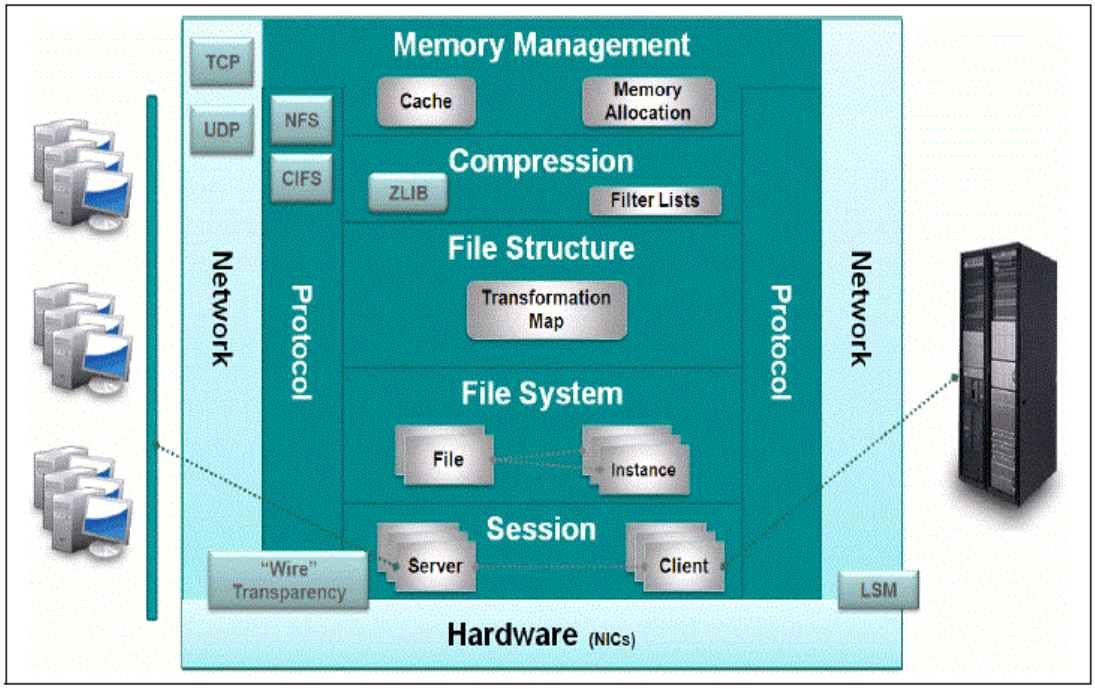

Figure 2-3 Random Access Compression Engine

RACE takes incoming data streams and compresses the data within these data requests, leaving the metadata intact, to the storage array. The data is stored in the array and the acknowledgement that the write was committed is sent directly back from the array to the user or application. This process flow is important because, from a data availability perspective, the storage array is what must acknowledge the write commitment—not the IBM RTCA product. For this reason, the RTCA product has no write cache. All storage commits come from the array, preserving the integrity of the data between the storage and the application.

As stated previously, the IBM RTCA technology uses industry standard LZ compression algorithms. What makes the compression work well by the RTCA is not necessarily only the compression algorithms, but rather the manner in which the compression is accomplished. One of the key ways that the RTCA product is able to achieve its high compression ratios and performance is by compressing data by using random access techniques.

The benefits of compressing data by using random access techniques are twofold First, the ability to read or write only the blocks of the compressed file that require read or modification means faster access performance for these operations. If you need to write only a small piece of data to update a whole file, your storage performance is maximized. Second, because the RTCA product has this capability, updates to the file are accomplished in a way that does not disrupt the other blocks in the compressed file.

The assumption is that upstream data compression significantly reduces downstream data deduplication ratios. Therefore a preferable way is to apply data deduplication technologies over decompressed data prior to performing a deduplicated backup. This way can be true for data that is compressed with traditional techniques. But the unique, random access nature of the IBM RTCA product's compression preserves data deduplication ratios. This preservation allows users to experience maximum data optimization in both primary and downstream tiers.

2.3.2 Link status mirroring

For the IBM RTCA product to appear in the network as an "intelligent cable," a technology must be available to ensure that the system provides the same functionally to act like a cable even though it is a device in the network. This solution is achieved through link status mirroring (LSM), which sits behind each cable in the data path. Suppose that any cable fails, or that any component in the data path fails. If so, the switch or storage or the I/O traffic coming into that cable, from either the sending or receiving end, must all be failed over to the other available link in the solution.

When a failure occurs and one link in the device is not available, LSM ensures that the other side of the link is also downed so that data traffic can be rerouted to the available link in the data path. This operation happens in microseconds to ensure the highest degree of availability to the data. Additionally, when the failure is fixed, and the link is restored, LSM also provides a mechanism to fail the link back to the predominant active link.

2.3.3 Unified Protocol Manager

One of the key components of RACE is the Unified Protocol Manager (UPM) that allows the IBM RTCA product to seat in line between the storage NAS (NFS and CIFS shares) and the user and to act as an "intelligent cable."

This transparent behavior means there is no change required to the user environment. Further, there are no changes required to storage, networking or applications. The IBM RTCA product does not require any changes to the existing infrastructure or to downstream processes such as snapshots, backup, replication, and archiving. The RTCA product enhances these environments because you are moving and storing less data.

The following use cases are examples:

- Snapshots: After you install the RTCA product, if your array allows you to take snapshots of your data, your snapshot data will also be compressed and be accessible, without any system changes to the server. The metadata, such as file permissions, size, ownership, and the access control list (ACL), are maintained in the compressed file.
- Conserving the file attributes: When a user performs a directory listing of the compressed data, the RTCA product reports all of the actual file metadata. This task is done as though the file was not compressed, including the file size and all of the ownership and permissions. By interpreting the file envelope and providing the native file attributes, the IBM RTCA product negates the need for applications to be changed to read the compressed data on disk. Therefore, there is no change required for IT, allowing the solution to fit seamlessly into the environment.
- No masking: When operating system commands such as df are run on the storage subsystem, the IBM RTCA product does not mask the compressed file size. This allows the server to always now exactly how much free space there is for files, and if quotas are in use, they can be enforced. These functions truly make the IBM RTCA product fully transparent in the environment.

Space information: The **df** command displays information about total space and available space on a file system.

2.3.4 Monitoring and Reporting Manager

Another component to the overall RACE technology is the Monitoring and Reporting Manager (MRM). The MRM enables online storage compression trending, analysis, and reporting. Further, MRM provides users with a host of services to efficiently manage and report on the IBM RTCA status and activities, and to quantify the benefits of the solution across the environment.

The MRM provides users with standard interfaces such as SNMP, allowing visibility into appliance health and performance through numerous traps and objects that can be polled. Downloadable SNMP MIB libraries allow users to integrate IBM RTCA monitoring into their monitoring and reporting tool of choice. MRM also allows users to integrate LDAP with the IBM RTCA product, facilitating auditable login authorizations. In an high availability (HA) configuration, MRM also automatically synchronizes the configuration setting between partner appliances to ensure the accuracy of configuration and streamline solution management.

In addition, with MRM, users can select certain file types to include or exclude from compression and also select shares to include or exclude from compression. With MRM, users can also see the amount of space saved by using the IBM RTCA product. MRM provides the ability to select the file types to include or exclude from compression. But more important, in most cases, it selects which shares to include or exclude from compression.

2.4 Performance considerations

Many factors affect storage performance, including but not limited to the following factors:

- Disk drives: Performance cannot be any faster in any system than the slowest mechanism in the chain. In storage, this is the mechanical movement of the disk drive. The farther the actuator arm has to travel, the slower the disk performance.
- I/O: The next consideration involves reads and writes, or I/O. The more reads and writes there are in a storage system, the more the storage system has to process these I/O requests. Also, the storage system takes longer to respond, and consequently, the storage system is slower. The storage system is busy processing all of these I/O requests and moving the storage arm to read and write data. So, the storage CPU has more loaded, and the response to extra storage requests becomes slower.
- Cache: At the top of the storage array performance is its cache. Storage cache is expensive; however, if configured properly for a given application, it can take a lot of the workload away from the disk, I/O, and CPU. It can provide faster response back to the user or application that is requesting data. Data that lives in cache is served significantly faster than data that must be retrieved from disk.

For any storage compression solution to fit into this environment, it cannot negatively affect any of these functions. To be a good solution, it must enhance these functions.

Also, because there is less data being written to the storage array, there is less I/O. With less I/O also come more CPU cycles to process the given read and write requests.

Finally, by compressing data in front of the storage array, a net increase in effective cache size is achieved. Whatever the compression ratio is for your data, this compression ratio transcends to your storage cache. Cache is one of the most expensive components of a storage array, and also, cache tends to have the biggest impact on storage performance. So, the more you can increase cache, the better performance users and applications will see.

Performing compression in real time, before the data gets to the storage, has tremendous benefits to overall storage performance. It removes a lot of the hurdles that storage administrators face when considering compression as a technology for their primary storage.

2.5 Data integrity

As part of the solution design, IBM Real-time Compression developed the Data Integrity Assurance architecture. The multitiered data protection it provides, built into every IBM Real-time Compression data compression appliance, is a quantum leap forward in guaranteeing data integrity.

To ensure 100% data integrity, the IBM Real-time Compression product line was developed to achieve the following goals:

- Preserving existing data integrity mechanisms in the storage array
- Preserving existing data integrity mechanisms in the network
- Verifying the reliability of storing and retrieving the compressed data
- Using the industry-proven Lempel-Ziv (LZ) lossless data compression algorithm
- Using error correction code (ECC) memory modules to ensure data integrity for memory access

The IBM Real-time Compression Appliance integrates seamlessly into NAS environments, increasing the storage capacity, while assuring data integrity as follows:

- Preserving the data integrity mechanisms already in place in the network
- Providing mechanisms to guarantee the reliability of the compressed data through the IBM DIA Data Integrity Assurance® architecture.

2.5.1 Data integrity at the network level

The IBM Real-time Compression Appliance preserves data integrity mechanisms at the protocol and storage levels.

Performing verification checks at the network protocol layer ensures data integrity during transmission between components in your network. The cyclic redundancy check (CRC) checksum mechanisms verify data integrity during transmission of data between components.

A graphical representation is shown in Figure 2-4.

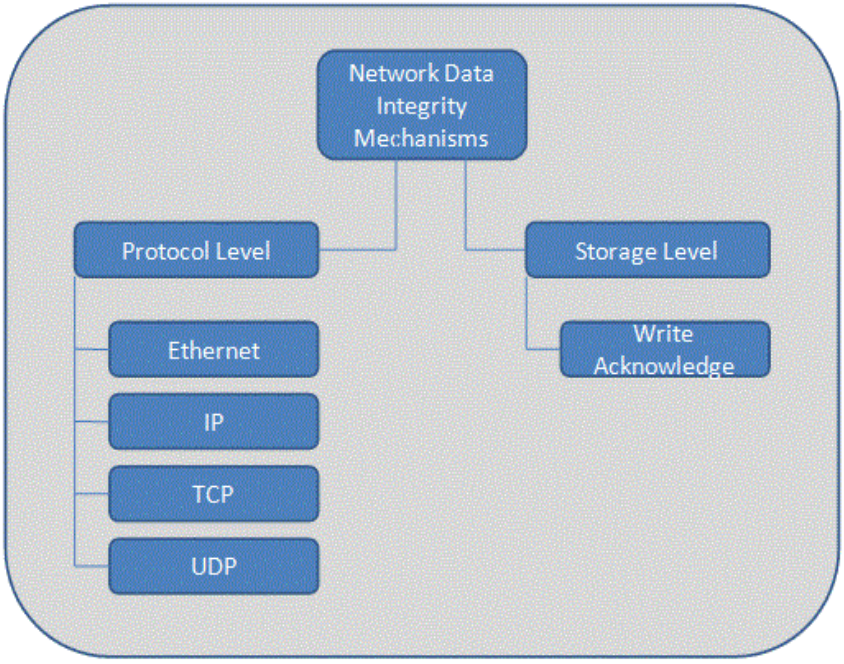

Figure 2-4 Network data integrity mechanism

2.5.2 Data integrity at the protocol layer

When you install the IBM Real-time Compression Appliance, the checksum mechanisms employed in your network continue to function as usual. The IBM Real-time Compression Appliance manages two sets of checksum mechanisms:

- Checksums on the uncompressed data in the connection between the IBM Real-time Compression Appliance and the switch
- Checksums on the compressed data in the connection between the IBM Real-time Compression Appliance and the storage array

The network protocol layer checksums include the following items:

- Ethernet-layer checksum (IEEE 802.3): A checksum that protects data against corruption during transmission on the Ethernet wire by performing a CRC checksum at the frame level
- IP header checksum (RFC 791): A checksum that ensures the integrity of the data fields in the IP header
- TCP checksum (RFC 793): A checksum that verifies the integrity of each data block during TCP/IP transmission
- UDP checksum (RFC 1122): An optional checksum that verifies the integrity of data packets associated with the UDP datagram

The following examples are of a network before and after IBM RTCA implementation:

- Figure 2-5 shows a network structure before IBM RTCA implementation.

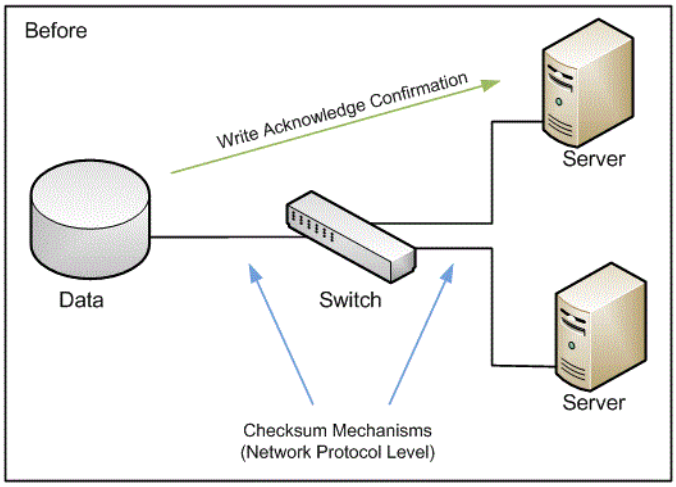

Figure 2-5 Network structure before IBM RTCA implementation

- Figure 2-6 shown a network structure after IBM RTCA implementation.

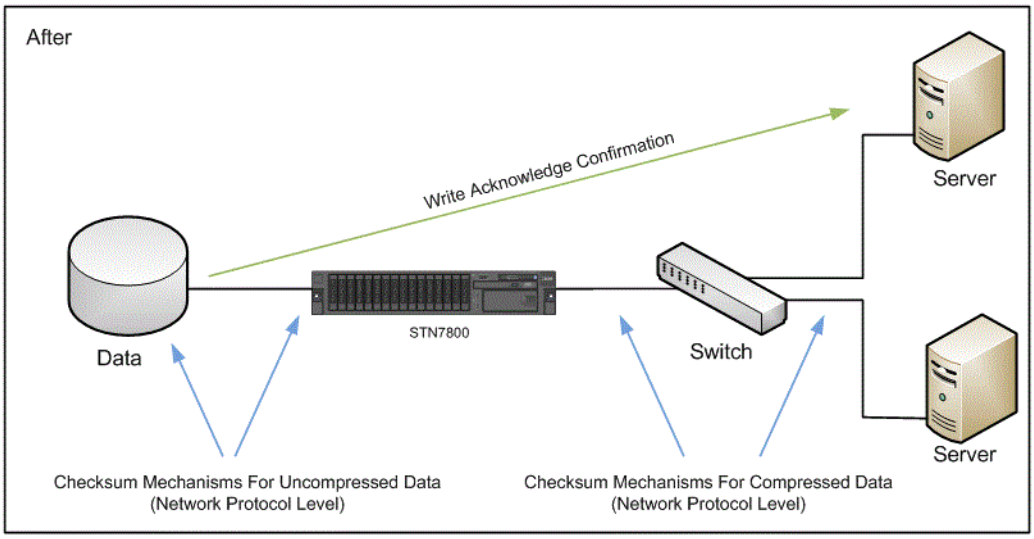

Figure 2-6 Network structure after IBM RTCA implementation

2.5.3 Data integrity at the storage layer

Write acknowledgement is a critical aspect of the functionality of client server storage applications. When the storage system completes a client write request, the storage system returns a write acknowledgement message. The write acknowledgement message provides confirmation that the data was transmitted successfully through the network and written to storage. This confirmation protects against failures in transmission caused by any component of the network, between the client server and the storage system.

When the IBM Real-time Compression Appliance is installed in your network, it transparently passes write acknowledgement messages that are received from the storage system to the client that requested the write. This way ensures that the data was stored on the storage array.

2.6 Storage advantages

Installing the IBM Real-time Compression Appliance in a storage configuration provides unmatched benefits to the storage environment:

- Primary storage compression: The IBM Real-time Compression Appliance compresses the data before it reaches the storage unit. This way provides unique advantages over other storage capacity solutions:
  - Improves response time: Reduces traffic to and from the storage system.
  - Reduces the storage system's read/write activity: Compressing data before it reaches the storage system enables the storage system to store and provide data with less read and write activity.
  - Read gains: Increases the effective size of the storage read cache, improving the performance of read operations. Storage system read performance is greatly influenced by the cache hit ratio. Better use of the read cache reduces the need for disk head movement, which is slower than cache access. Because IBM Real-time Compression Appliance compresses the data before it reaches the storage system, the effective size of the read cache increases. The read cache improvement is proportionate to the compression ratio.
  - Write gains: Increases the effective size of the write cache, improving the performance if write operations. Storage systems typically cache write operations to enable immediate acknowledgement to clients and reduce response time. The size of the write cache significantly impacts the throughput of write operations to storage. The IBM Real-time Compression Appliance compresses the data before it reaches the write cache, so the effective size of the cache increases proportionately with the compression ratio.
- Reduces infrastructure costs: Using IBM Real-time Compression Appliance reduces your storage capacity by up to 80%, and helps to eliminate the expensive infrastructure requirements of expanding your storage footprint.
- No changes or server resources are required: The IBM Real-time Compression Appliance provides storage gains with the advantages of self-contained, appliance simplicity:
  - No need to change the storage system configuration
  - No need to change your workflow
  - No server resources used

Compared to proprietary array compression or deduplication the IBM Real-time Compression Appliance has significant advantages also:

- Reduces IO workload
- Reduces CPU workload
- Cache multiplying effect
- Array vendor neutral

As a final benefit, the IBM Real-time Compression Appliance can break the IOPS per terabyte ratio. As we will be able to get larger drives, we still are limited to the physical performance capacity of the drives.

3

What is new in RTCA v4.1

IBM Real-time Compression Appliance Version 4.1 offers changes and enhancements, such as platform, NAS protocol, and compression engine.

The following topics are included in this chapter:

- 3.1, “New appliance: STN7800” on page 24
- 3.2, “Revised terminology for compression modes” on page 24
- 3.3, “Improved GUI support of NFS subexports” on page 26
- 3.4, “Autosupport feature enhancements” on page 26
- 3.5, “New ifconfig command for IPv6” on page 27

3.1 New appliance: STN7800

Version 4.1 introduces new appliance hardware, STN7800, as shown in Figure 3-1. This hardware is suitable to handle higher performance requirements and has an improved software operating system. The software optimizations in I/O processing are as follows:

- Main thread separation offers improved read/write throughput by splitting core0 work to an additional core.
- Separating the I/O work and protocols and business-logic work (which until now were done altogether in the “main” thread) is in different threads.
- Increases the throughput for read/write significantly.

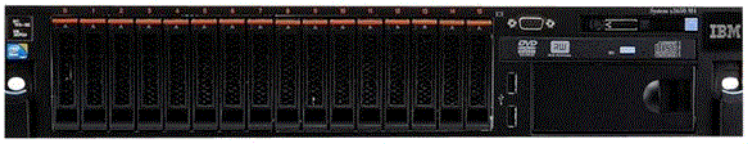

Figure 3-1 STN7800

The STN7800 hardware specification highlights are as follows:

- 2452-780 High End Appliance (STN7800)
- Based on IBM System x3650 M4 7915
- Dual 2.7 Ghz 8-cores Xeon Processors
- Eight 10Gb, or eight 1Gb plus four 10Gb Ethernet ports
- Dual power supplies
- RAM is 96 GB

More detailed hardware specifications are as follows:

- Dual Intel Xeon E5-2680 processors
- Twelve 8 GB DDR3 1333 MHz LP RDIMM or eighteen 16 GB DDR3 1333MHz LP RDIMM
- Dual 300 GB 10 K 6-Gbps SAS HDD
- UltraSlim Enhanced SATA DVD-ROM
- Dual redundant 900 W ac high-efficiency power supplies
- Integrated quad Gigabit Ethernet ports for high I/O capacity
- Two optional embedded 10 GbE ports
- Two C13-C14 PDU cables
- One 3 m Blue Cat6a cable
- One 3 m RS-232 cable

3.2 Revised terminology for compression modes

In previous releases, the Transparent/Compress mode and No-New-Compression could be attached, but was ambiguous. To create a clear meaning, the RTCA terminology is revised in the GUI.

This options set the operation mode for new files that are created on a storage interface. When the **Compress new files** option is disabled, shares that are configured to work in compressed mode have existing files read and updated based on their current state (compressed or uncompressed), but new files are created as uncompressed files.

This setting does not change the behavior of the storage system compression mode, which pertains strictly to how new files are created. Use the Compress new files mode in cases where preserving the system maintenance operations and the current state of shares is important. For example, if a storage system IP address is configured as `Compress shares/exports except selected` in its Compression Mode field, and the storage administrator wants to create a share that must work in uncompressed mode, a compression filter is created after the share is created on the storage.

Configure the system in **Compress new files** disabled mode before you add the share to avoid creating compressed files on the storage; then enable **Compress new files** after the share is added to the compression filters exclude list.

CLI unchanged: The CLI remains as before so changes are not required for scripts.

Table 3-1 provides the mapping between old to new terminology.

Table 3-1 New terminology

New terminology	Old terminology
Compress all shares/exports except selected, Compress only selected shares/exports	Compressed, Transparent
Compress New Files	No New Compression

Figure 3-2 shows a sample of the new terminology in the GUI interface.

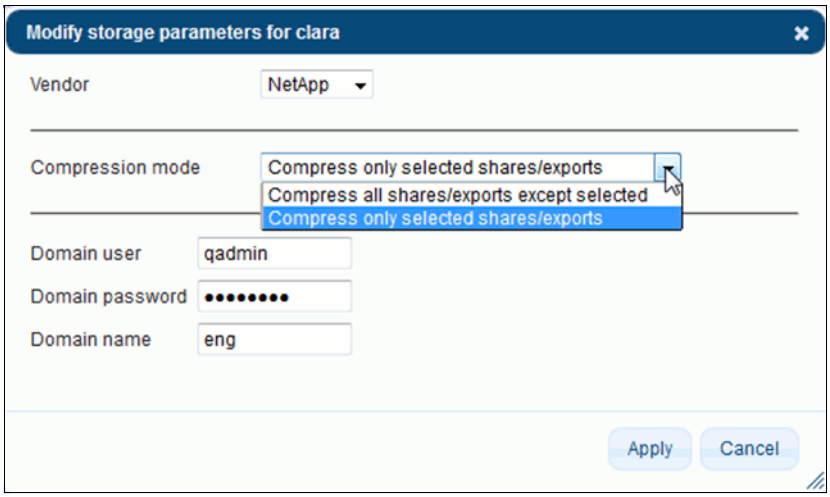

Figure 3-2 New terminology in GUI

3.3 Improved GUI support of NFS subexports

In previous releases, when users added NFS exports to the exclude list, the GUI did not display the subexport as un-compressed.

To avoid misleading information, the NFS subexport is added to the filter list to protect the data from “clear” access, not through the compression engine.

Figure 3-3 shows the NFS subexport as un-compressed.

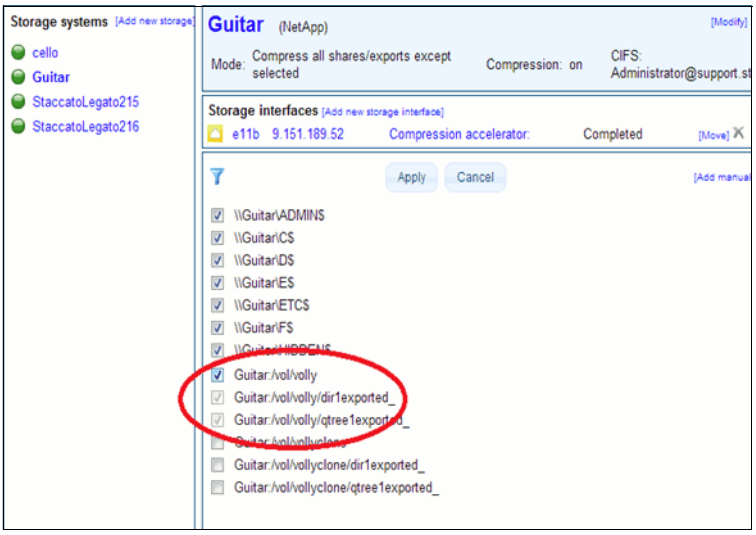

Figure 3-3 NFS export including subdirectories

3.4 Autosupport feature enhancements

Enhancements are included for autosupport.

3.4.1 Support for Autoconsult

Autoconsult is a feature that supplies error code with every autosupport trigger. It is designed to help IBM support teams to provide technical details and find the best solution.

The email body has a new structure:

- All fields are included, even if empty.
- All spacebars are changed to underscores or removed in field names.
- Two additional fields are included:
  - **swVer** = 4.1.0.01
  - **Action** = 0 will note IEPD to open a PMR, 1 will note to open a PMR (hardcoded into the matrix in the appliance) all existing alerts should say Action=1 *except* the following alert: `User-initiated reboot c. ErrorCode=1094`. Each alert is assigned a particular error code (hardcoded in the matrix). Different reboot reasons must translate into different error codes.

3.4.2 SMTP authentication

Support for SMTP authentication for the Autosupport feature is included with version 4.1. The authentication supports only cleartext passwords with no support for Transport Layer Security (TLS).

The GUI is extended by two options: SMTP User and SMTP Password. The options are shown in Figure 3-4.

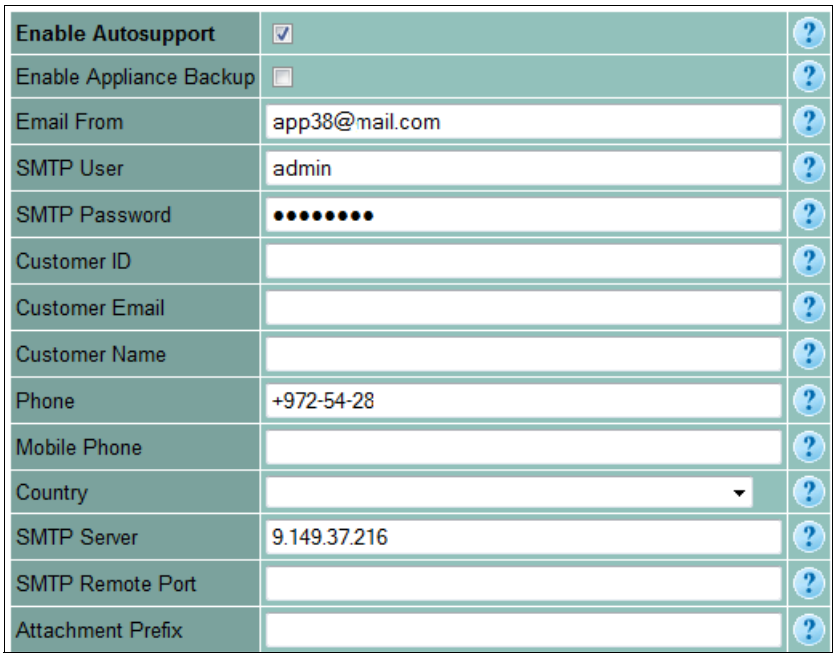

Figure 3-4 Autisupport SMTP options

In addition, the command-line interface (CLI) has an extended parameter for the support command:

```
support [-su SMTPUser] [-sp SMTPPassword]
```

3.5 New ifconfig command for IPv6

In previous versions, only one command was available to configure IPv4 and IPv6 networking environments. Previously, users had to use the **-6** options to configure IPv6 networking.

Now, two commands are available to configure the environments:

- **ifconfig mgmt**
- **ifconfig6**

Every command must be used with its own parameter set. Example 3-1 shows a sample usage of the **ifconfig6** command.

Example 3-1 The ifconfig6 command

```
ifconfig6 mgmt -type auto
user@10.10.10.5> ifconfig6 mgmt -type static
-i 2001:bf8:2000:5180:e61f:13ff:fe6d:2348/64 -gw 2001:bf8:2000:5180::4
```

4

Hardware update

This chapter has an overview of the IBM Real-time Compression Appliance (RTCA) 4.1 hardware and software components. The chapter outlines base model configurations and the technical specifications, such as number of ports, memory capacity, and the physical attributes of the products.

The following topics are included in this chapter:

- 4.1, “Base model configuration” on page 30
- 4.2, “Front-panel components” on page 31
- 4.3, “Rear-panel components” on page 32
- 4.4, “Technical specifications of STN7800” on page 33
- 4.5, “Interoperability matrix” on page 34

4.1 Base model configuration

Currently, the IBM RTCA product portfolio offers models that can be ordered as base version or optional features may be added

These details are valid at the time of writing, but for current details, see the following website:

http://www.ibm.com/systems/storage/network/rtc/

4.1.1 IBM Real-time Compression Appliance STN7800 (2452-780)

The IBM Real-time Compression Appliance STN7800 is a new model, based on standard IBM System x3650 M4, created with innovative design and flexibility and optimized for the best performance.

Featuring IBM System x3650 M4

The IBM System x3650 M4 blends outstanding uptime, performance, and I/O flexibility for cost efficiency and solid reliability. The powerful server offers an energy-smart, affordable, and easy-to-use rack solution with a pay-as-you-grow design to help lower costs and manage risks. With more computing power per watt and the latest Intel Xeon E5-2600 series processors, advanced memory support, and greater disk capacity for businesses requiring extreme storage, the x3650 M4 offers balanced performance and density. Figure 5-1 shows the IBM System x3650 M4.

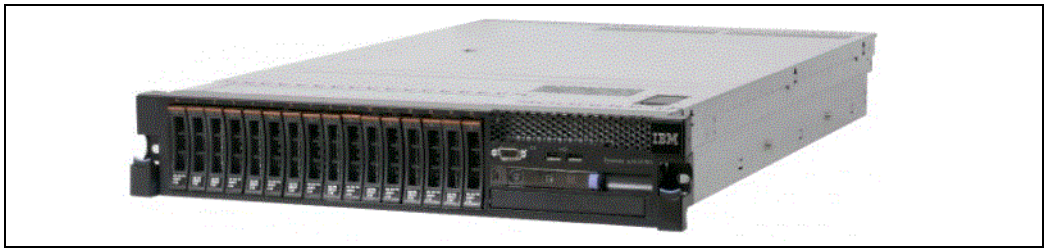

Figure 4-1 IBM System x3650 M4

Outstanding uptime

With redundant hot-swap fans, disks and power supplies, the x3650 M4 provides a resilient architecture that is ideal for business-critical applications. Predictive Failure Analysis and light path diagnostics provide advanced warning on power supplies, fans, voltage regulator modules (VRMs), disks, processors, and memory. Redundant hot-swap components can help you more easily replace failures without shutting down your system.

4.1.2 Hardware summary

The following list provides a hardware overview of STN7800 base components:

- 2U rack server provides expand ability and high performance in a dense form factor
- Up to two Intel Xeon Processor E5-2600 product family
- Up to 288 GB memory and up to two DIMMs per channel at 1333 MHz memory speed
- Integrated quad-port Gigabit Ethernet and optional embedded dual-port 10 GbE
- Up to 16 TB of 1.8-inch hot-swap SSDs or 16 TB of 2.5-inch hot-swap SAS/SATA/SSDs or 18 TB of 3.5-inch hot-swap or simple-swap SAS/SATA HDDs

- Embedded 6 Gbps hardware RAID-1
- Hot-swap disk/fan/power supply, two fan zones with N+1 fans design, light path diagnostics and Predictive Failure Analysis, better thermal design, balanced efficiency, uptime, and serviceability

The STN7800 is available in three hardware versions: mixed fiber, full fiber, or 1 Gb only. Therefore RTCA 4.1 is applicable to different types of businesses and needs of different types of infrastructure.

Note: PCIe cards are what determine the type, quantity, and speed of ports that are available on RTCA 4.1. For each, there is a need to use indispensable lift configuration requirements and ascertain the needs of the environment.

4.2 Front-panel components

Figure 4-2 shows the front view of the STN7800.

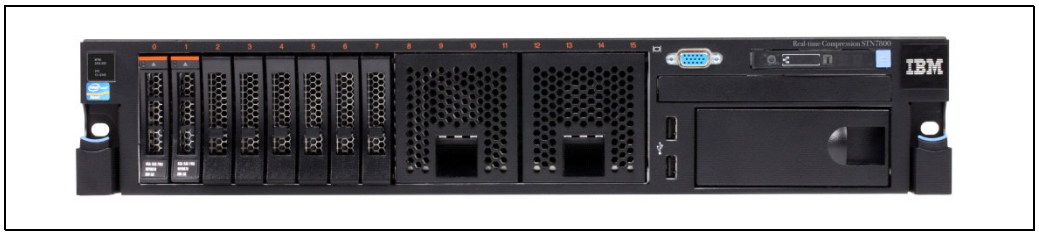

Figure 4-2 IBM STN7800 front panel

The front panel includes these components:

- LED indicators and pinholes Power pinhole: Turns the compression appliance power supply on or off.
- Power LED: Indicates if the compression appliance is receiving power.
- Light path diagnostics LED: Indicates whether there are any compression appliance alerts.
- Light path diagnostics panel: Provides monitoring and alerting for system components. If a failure occurs, a light is illuminated on the front panel of the compression appliance. Any component failure is also illuminated by LEDs on the light path panel itself, and also within the system.
- USB port: Used for manufacturing purposes (initial installation of the compression appliance).
- Hot-swappable RAID-1 mirrored hard drives: Holds the IBM Real-time Compression Appliance software and configuration.
- Optical drive bay, ultraSlim enhanced SATA DVD-ROM.
- Dual 300 GB 10 K 6-Gbps SAS HDD.

4.3 Rear-panel components

Figure 1-2 shows the rear view of the STN7800.

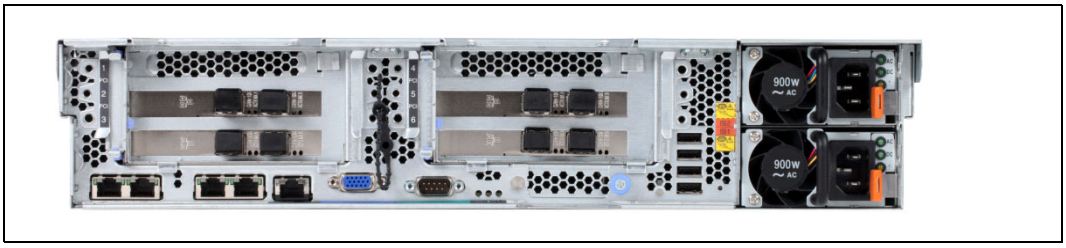

Figure 4-3 IBM STN7800 rear view

The default configuration of STN7800 ports and connections serve the following functions:

- 4x 1Gb Ethernet ports
- 2x 10Gb Ethernet ports
- Serial console port (RJ-45)
- Power supply units (PSUs)
- AC connectors
- PCI Express slots for Ethernet cards

Improved performance: The STN7800 processor performance is significantly improved compared to previous generations. The new platform and design, 10Gb Ethernet, and network bandwidth are important tasks in high performance. Also available at the same time, 1Gb, 10Gb Ethernet onboard, power efficiency, and flexible network configuration when 80 PIUs Platinum certified power supply is used.

4.4 Technical specifications of STN7800

This section discusses the technical specifications for the STN7800.

4.4.1 IBM System x3650 M4

Technical specifications for IBM System x3650 M4 are listed in Table 5-1.

Table 4-1 Technical specifications for IBM System x3650 M4

Feature	Value
Operating weight	Minimum configuration: 25 kg (55 lb.) Maximum configuration: 30 kg (65 lb.)
Form factor	2U Rack
Dimensions width	445 mm (17.5 in)
Depth	746 mm (29.4 in)
Height	86.5 mm (3.4 in)
Power	Up to two redundant hot-swap 550 W ac, 750 W ac, or 900 W ac power supplies (all 80 PLUS Platinum certification), or -48V 750 W dc power supply options.

4.4.2 Operating environment

Operating environment specifications for STN7800 are listed in Table 4-2.

Table 4-2 Environment specifications for STN7800

Feature	Value
Temperature server on	41.0 F to 104 F
Server off	41.0 F to 113 F
Humidity server on	20 - 80%
Humidity server off	8 - 80%
Acoustical noise emissions: Declared sound power, idle	6.4 bels
Acoustical noise emissions: Declared sound power, idle, operating	6.6 bels

Energy: The STN7800 offers the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to the “green” environment.

4.5 Interoperability matrix

This section summarizes the changes in the v4.1 interoperability matrix relative to the previous version (v3.9).

The full interoperability matrix for RTCA v4.1 is available at the IBM System Storage® Interoperation Center (SSIC) website:

http://www.ibm.com/systems/support/storage/ssic/interoperability.wss

4.5.1 Additional storage system support

In the RTCA v4.1 release, support is included for the following NAS storage system:

- EMC DART 7.0

4.5.2 Additional client system support

In the RTCA v4.1 release, support is included for the following operating system clients:

- Windows 8
- Windows 2012 server

For more information, see the release notes:

http://www-304.ibm.com/support/docview.wss?uid=ssg1S7003976&aid=1

4.5.3 SAMBA interoperability

Version 4.1 adds the following SAMBA version to its interoperability matrix:

- SAMBA 3.5.19

Part 2

Implementation

This part of the book provides the necessary details for a successful Real-time Compression Appliance implementation.

This part contains the following chapters:

- Chapter 5, “Implementation planning” on page 37
- Chapter 6, “Software installation or upgrade” on page 55
- Chapter 7, “Performance considerations” on page 67

It also contains a list of publications and other resource material to help you with the implementation.

5

Implementation planning

This chapter offers a planning guide to provide proper solution design. It describes the physical and logical installation of the IBM Real-time Compression Appliance (RTCA). It also describes validation and testing of the final installation and reporting considerations.

The following topics are included in this chapter:

- 5.1, “Introduction to installation, reporting, validation, testing” on page 38
- 5.2, “Physical installation” on page 38
- 5.3, “Logical installation” on page 46
- 5.4, “Reporting considerations summary” on page 54
- 5.5, “Validating and testing the installation” on page 54

5.1 Introduction to installation, reporting, validation, testing

RTCA physical installation includes and addresses the following items:

- Solution requirements
- Architectural planning
- Network planning
- Antivirus services considerations
- Network considerations

Logical installation of the RTCA includes the following items:

- Configuration planning
- High-availability for the compression appliances

Reporting considerations include information regarding the following items:

- Capacity consumption
- File types
- Capacity growth

Validation and testing of the final installation includes the following activities:

- Checking cabling, hardware, power on, and system logs
- Meeting known memory limitations for NFS and CIFS

5.2 Physical installation

This section describes the planning of the physical RTCA installation.

5.2.1 Solution requirements

The following list is an outline of important solution requirements:

- NAS storage platform: The IBM Real-time Compression Appliance (RTCA) product supports only NAS protocols. Ensure that your Storage Server and required NAS protocols are supported. Information about supported features is at the following website:

  `http://www.ibm.com/support/entry/portal`

- Data type: Be sure to identify compressible data before starting the configuration of the RTCA product.
- Capacity: Be sure to have enough capacity so that the solution is cost-effective. From experience, in most cases this is more than 20 TB of NAS data.
- Growth rate: Having a negligible annual growth rate is not what you want. The higher the growth rate, the more cost-effective the RTCA product will be.

5.2.2 Architectural planning

Architectural planning for RTCA solutions involves, physical installation and cabling, network interfaces, rack considerations, cables and ports, and understanding physical bypass.

Physical installation and cabling considerations

Evaluate physical and cabling considerations.

Network interfaces

Table 5-1 lists the three network options of the RTCA product.

Table 5-1 Network interfaces

Network interfaces	STN6500	STN7800 with feature code 1001[a]	STN7800 with feature code 1010[b]
1 Gbps copper ports	16	8	-
10 Gbps copper ports	-	4	8
1 Gbps copper management port	1	1	1

a. Feature code 1001: 4x10Gb and 8x1Gb Ethernet ports; this feature provides both 1Gb and 10Gb Ethernet configurations.
b. Feature code 1010: 8x10Gb Ethernet ports; this feature provides only 10Gb configurations.

Feature code 1010 *or* feature code 1001 must be included in the initial order of STN7800.

Rack considerations

IBM RTCA hardware must be installed in a standard 19-inch rack by using the rack mounting hardware provided with the machine. The rack cabinet must have a 2U slot available for each RTCA product.

Cables and ports

Your network cables must be unplugged from current storage, and plugged into the compression appliance. IBM supplies cables that go from the compression appliance to the storage:

- Feature code 3003: 3m Green Cat5e cable

  This feature provides one 3-meter Green Cat5e cable with RJ-45 connectors for 1Gb Ethernet connections.

- Feature code 3705: 5m LC-LC optical cable

  This feature provides one 5-meter LC-LC Optical cable for 10Gb Ethernet connections.

If you do not purchase cables from IBM, you must have an additional set of network cables to connect the RTCA product with your NAS storage server.

If the NAS storage is connected with eight 1-Gb copper cables to one Ethernet switch, you will need, in additional, eight 1-Gb copper cables and one 1-Gb copper cable for the management port for installing the RTCA product. Therefore, only one additional network port per RTCA product is required on the Ethernet switch.

Physical bypass

Under normal conditions, bypassing the RTCA product for any Ethernet traffic is unnecessary.

Tips:

- To avoid direct access to compressed data by clients, without bypassing the RTCA product, the best approach is to connect all network interfaces of the NAS storage server with the RTCA product.
- If you have available Ethernet ports on your NAS storage server and want to use these ports for iSCSI or replication purposes, at these Ethernet ports all NAS protocols must be deactivated.

5.2.3 Network planning

For proper solution design, you should understand networking concepts and terms that are associated with the RTCA product. You should understand about certain switches. Also, understand about planning the physical connection

RTCA network terminology

The RTCA product acts as an “intelligent cable” between the Ethernet switch and the NAS storage server. To provide network connectivity, having bonds and bridges is necessary.

Bond

A bond combines physical RTCA ports to a single virtual interface; a bond is the vehicle for port aggregation. The RTCA product supports the common port aggregation protocols EtherChannel and 802.3ad (LACP). Also, you can create active backup bonds.

Bridge

A bridge combines two physical interfaces or bonds to establish a network connection between the Ethernet switch and the NAS storage server. For every bridge, a new IP address must be configured. This IP address is used by the RTCA product to poll the NAS shares and exports.

The general use of bridges and bonds is illustrated in Figure 5-1.

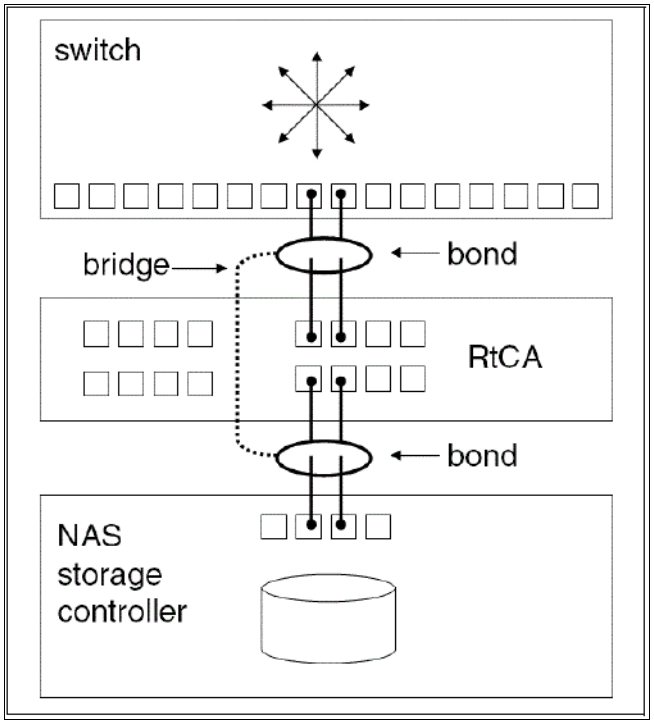

Figure 5-1 General use of bonds and bridges

Logical interfaces

If your NAS storage system is configured with port aggregation, look at the port aggregation configuration and document the allocation of virtual interfaces to physical interfaces. Preferably, adapt the aggregated physical interfaces of your NAS storage server to corresponding bonds at your RTCA product (all physical interfaces of an aggregated interface will be in the same bond).

Example 5-1 shows the configuration of types of virtual interfaces of an IBM System Storage N series controller. The physical interfaces `e0a` and `e0c` are trunked by `vifa` as a multi-mode VIF. The physical interfaces `e0b` and `e0d` are trunked by `vifb` as another multi-mode VIF. Both `vifa` and `vifb` are configured as members in a second level VIF `vif1` as a single mode VIF.

In an HA configuration, you must use one bond for `vifa` at the first RTCA product and one bond for `vifb` at the second RTCA product. Both bonds will be bridged to the corresponding bonds, which include the Ethernet connection to the switch. The VIF configuration remains unchanged at the N series system.

Example 5-1 Virtual interface configuration of N series with four physical interfaces

```
bran1> vif status
default: transmit 'IP Load balancing', VIF Type 'multi_mode', fail 'log'
vifa: 2 links, transmit 'Round-Robin Load balancing', VIF Type 'multi_mode' fail
'default'
VIF Status Up Addr_set
trunked: vif1
up:
e0a: state up, since 23Mar2011 15:24:49 (22:18:23)
mediatype: auto-1000t-fd-up
flags: enabled
```

```
input packets 19529, input bytes 2358737
output packets 30533, output bytes 1509659
up indications 10, broken indications 7
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 15:24:49
consecutive 170269, transitions 17
e0c: state up, since 23Mar2011 10:00:55 (1+03:42:17)
mediatype: auto-1000t-fd-up
flags: enabled
input packets 278152, input bytes 16971759
output packets 30757, output bytes 1522361
up indications 8, broken indications 5
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 10:00:55
consecutive 170755, transitions 13
vifb: 2 links, transmit 'Round-Robin Load balancing', VIF Type 'multi_mode' fail
'default'
VIF Status Up Addr_set
trunked: vif1
down:
e0b: state up, since 23Mar2011 09:55:33 (1+03:47:39)
mediatype: auto-1000t-fd-up
flags: enabled
input packets 346488, input bytes 101244051
output packets 86193, output bytes 41936062
up indications 4, broken indications 2
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 09:55:33
consecutive 171415, transitions 6
e0d: state up, since 23Mar2011 09:55:33 (1+03:47:39)
mediatype: auto-1000t-fd-up
flags: enabled
input packets 347356, input bytes 98485301
output packets 84014, output bytes 43590209
up indications 2, broken indications 0
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 09:55:33
consecutive 171880, transitions 2
vif1: 1 link, transmit 'none', VIF Type 'single_mode' fail 'default'
VIF Status Up Addr_set
up:
vifa: state up, since 23Mar2011 10:00:55 (1+03:42:17)
mediatype: Enabled virtual interface
flags: enabled favored
input packets 297681, input bytes 19330496
output packets 61290, output bytes 3032020
output probe packets 56815, input probe packets 167902
strike count: 0 of 10
up indications 7, broken indications 4
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 10:00:55
consecutive 170785, transitions 11
down:
vifb: state down, since 23Mar2011 10:01:27 (1+03:41:45)
mediatype: Enabled virtual interface
```

```
flags: enabled unfavored
input packets 693844, input bytes 199729352
output packets 170207, output bytes 85526271
output probe packets 56831, input probe packets 170706
strike count: 0 of 10
up indications 2, broken indications 0
drops (if) 0, drops (link) 0
indication: up at 23Mar2011 09:55:33
consecutive 171880, transitions 2
```

Important: All members of an active/active port aggregation, such as multi VIF or LACP VIF of Data ONTAP, must be connected to the same RTCA product. If the same shares or exports are accessed through separate links, an active/active setup can lead to data access issues because concurrent access to files through separate RTCA products is prohibited.

For further support information, see the RTCA Release Notes and the RTCA Administration Guide at this website:

http://www.ibm.com/support/entry/portal

Spanning-tree portfast

On Cisco switches, if you enable spanning-tree portfast, a warning is issued that this feature is not recommended for use by the bridge interface.

Disabling this feature is a requirement for all data ports in a RTCA setup, as shown in Example 5-2

Example 5-2 Disabling spanning-tree portfast on Cisco switches

```
Support#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Support(config)#interface gigabitEthernet 0/1
Support(config-if)#spanning-tree portfast
%Warning: portfast should only be enabled on ports connected to a single host.
Connecting hubs, concentrators, switches, bridges, etc... to this interface when
portfast is enabled, can cause temporary bridging loops.
Use with CAUTION
%Portfast has been configured on GigabitEthernet0/1 but will only have effect when
the interface is in a non-trunking mode.
Support(config-if)#
```

Planning the physical connection

Before you begin creating bridges and bonds, document the existing cabling diagram to determine which RTCA ports will be used to bring the switch-to-RTCA and the RTCA-to-NAS connection together.

Consider at least two possible way for planning the physical connections; that is, the paired cables (the cable from the switch to RTCA and the cable from NAS) can reside as follows:

- Reside on separate PCI cards on RTCA.
- Reside on the same PCI card on RTCA.

The second option (reside on the same PCI card on RTCA) is described 5.2.4, “Antivirus services considerations” on page 44 for the following reasons:

- From experience with PCI card problems in IBM Series x servers, we can confirm that the difference between the first and second option, for safety reasons, is minimal.
- The second option offers easier cabling interface connections and an easier logical environment perception.

5.2.4 Antivirus services considerations

In many NAS environments, the antivirus scanning activities are off-loaded for better scalability and performance. NAS vendors often advise having a dedicated physical network connection or dedicated VLAN configuration to transmit data from the NAS storage server to the antivirus server.

A highly preferred approach is to configure your physical network with NAS storage and the RTCA product, thus all data passes the RTCA product. If you configure a physical bypass of the RTCA product, compressed files will be transmitted to the antivirus server and malicious virus code will never be detected. If necessary, you can configure a dedicated bridge to separate the antivirus scanning traffic from all other network traffic.

An example of an antivirus configuration is illustrated in Figure 5-2

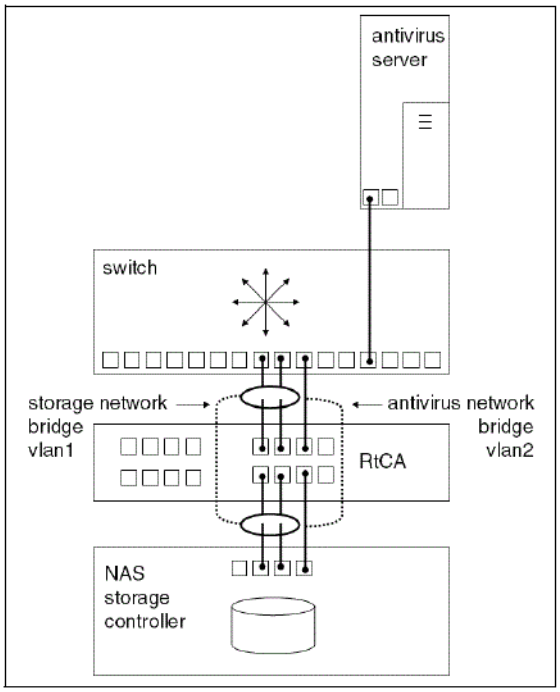

Figure 5-2 Antivirus server in an RTCA environment

5.2.5 Network considerations summary

The following network considerations are important:

- Bridges connecting between the network and storage

  Each side of a bridge must be connected to either a network switch or an NAS storage server interface. Do not connect the two sides of the bridge to the network switch; this method creates a network loop that can affect service to users.

- Direct connection to the NAS

  The Link State Mirroring (LSM) is monitoring the network link status and automatically mirrors this status between the bridge interfaces. Connecting active components such as link converters or powered patch panels in between the appliance and the NAS storage server can cause the LSM mechanism to fail. Therefore, this approach is not supported.

- Number of network interface ports

  The compression appliance acts a bridge between two “sides,” so the number of required ports is 50% higher, of the total number of links that are configured in the NAS storage server. For example, if an N series system is configured with four 1 Gb links, a total of eight 1 Gb links will be required on the RTCA product to support this configuration.

  In addition, the total number of links must support the expected aggregate throughput. For example, a particular NAS storage server is expected to provide approximately 500 MBps total throughput. Because each 1 Gb link provides approximately 80 - 100 MBps, in this environment, the best approach is to configure at least five bridges.

- Number of bridges

  Up to eight physical bridges can be configured on a single appliance, and up to 128 bridges configured with VLAN tagging (802.1q).

- Number of storage system interfaces

  Each individual storage system physical interface and VLAN-tagged interface must be represented in the RTCA network configuration. Each RTCA product provides up to 16 ports, so an NAS storage server with up to eight physical interfaces can be implemented. Connect all network interfaces through the appliance unless some interfaces are used solely for NDMP or replication traffic that is not based on NFS or CIFS.

- Number of IP addresses configured in the NAS

  Each NAS IP address must be configured in the compression appliance to support communication between the client and the NAS storage server. This step includes IP aliases and virtual server IP addresses.

- Distribution of network links across NICs

  The RTCA product is configured with either quad port 1Gb Ethernet cards, or dual-port 10Gb Ethernet cards. To provide the highest redundancy possible, the preferred way is to configure bridges from ports that are located on separate cards, as shown in Example 5-3.

Example 5-3 Bridge assignment

```
br0 – eth0 and eth4
br1 – eth1 and eth5
br2 – eth2 and eth6
br3 – eth3 and eth7
br4 – eth8 and eth12
br5 – eth9 and eth13
br6 – eth10 and eth14
br7 – eth11 and eth15
```

5.3 Logical installation

Planning your RTCA configuration involves configuration planning, understanding high availability for the compression appliances, understanding reporting considerations, validating, and testing.

5.3.1 Configuration planning

Before powering on your RTCA product for the first time, use the configuration worksheets to gather the information that the setup process requires:

1. Use Table 5-2 to gather the information for the initial setup.
2. Use Table 5-3 to enter the required information for configuring autosupport.
3. Use Table 5-4 to provide information for NTP, DNS, and logging.
4. Use Table 5-5, for every NAS storage server, to gather the required information for later configurations.
5. Use either Table 5-6 or Table 5-7, depending on the interface configuration of your NAS storage server, to gather the network information.

Table 5-2 can help you collect the necessary port management information for initial setup.

Table 5-2 Port management port setup

RTCA product host name	
IP Address	
Subnet Mask	
Gateway	
MTU (default 1500)	
Link mode	auto 10baseT/Half 10baseT/Full 100baseT/Half 100baseT/Full 1000baseT/Full

Table 5-3 can help you collect the necessary autosupport information.

Table 5-3 Autosupport configuration

Email from **(the “from” user that sends email)**	
SMTP server (IP address or host name)	
SMTP remote port (default - 25)	
Attachment prefix (The names of email attachments are prefixed with this name to avoid the possibility that the email will be rejected by some email servers)	

Table 5-4 can help you collect miscellaneous information.

Table 5-4 Miscellaneous NTP and DNS configuration information

IP address NTP server 1	
IP address NTP server 2	
IP address remote syslog server	
DNS domain	
DNS search list	
DNS name server IP address	

Table 5-5 can help you collect the necessary NAS information.

Table 5-5 NAS storage configuration: one table for each NAS storage server IP address

Storage system name	
Storage Vendor	
Storage IP address	

Table 5-6 can help you collect the necessary physical interface information.

Table 5-6 NAS storage physical interfaces to bridge directly switch - one table per interface

Name of the interface
IP address for the bridge

Table 5-7 can help you collect the virtual interface information.

Table 5-7 NAS storage virtual interface configuration - one table per virtual interface

Name of virtual interface	
Member of virtual interface (physical interfaces)	
Type	multi (EtherChannel) multi (802.3ad LACP) single (active-backup)
IP address for the bridge	

5.3.2 High availability for the compression appliances

Among many possible design solutions, decide which one is important and which will be used before implementing RTCA appliance. The solution in question is a high availability solution for the compression appliance.

Four high-availability configurations are possible. They are listed here and described in the following section, for both before and after connecting the IBM Real-time Compression Appliance products. The manner in which the compression appliance preserves the high availability features of the NAS is illustrated in the configuration diagrams.

- Active/passive path failover
- Active/active link aggregation

- Path failover using two switches
- Clustered environment (preferred configuration)

Active/passive path failover

The active/passive path failover is configured according to a single NAS storage controller, which uses link aggregation with two or more physical interfaces. Only one link can be active at any particular time. A standby link is activated only if the NAS storage controller detects a fault in the active link. That is when all traffic goes through the new link that was activated.

Figure 5-3 shows an active/passive path failover configuration without (on the left) and with (on the right) RTCA.

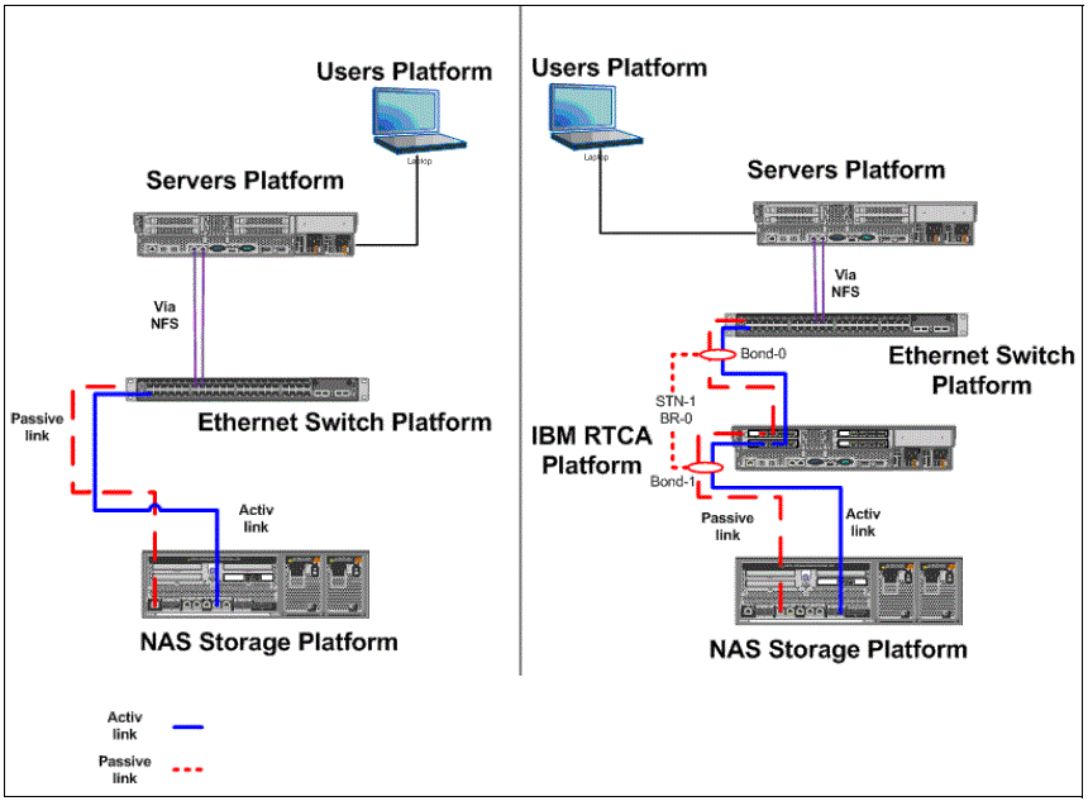

Figure 5-3 Active/passive path failover: without and with a compression appliance

Use the following steps to configure active/passive path failover:

1. Test the NAS and switch for support of active/passive (failover) in direct connection mode.
2. Connect two network cables from the compression appliance to the NAS.
3. Connect two network cables from the compression appliance to the network switch.
4. By using the compression appliance web interface, create two bonds, one for the NAS and one for the network, each holding two physical interfaces.
5. Configure the mode on both bonds as active/passive.
6. Configure a bridge STN-1 BR-0 with bond0 and bond1 as members.
7. Configure the file server.
8. Start the compression engine.
9. Continue with the standard configuration.

Active/active link aggregation

The active/passive path failover is also configured according to a single NAS storage controller, which uses link aggregation with two or more physical interfaces. All links can be, at the same time, active and connected to a switch that supports aggregation through multiple port connections.

The switch is what determines the way in which incoming packets will be forwarded to the NAS storage controller. If a link in the aggregate shuts down, the remaining links that are active share the load and traffic that is supposed to be in the NAS storage controller, which then goes through the remaining active interfaces.

Figure 5-4 shows an active/active link aggregation configuration without and with RTCA.

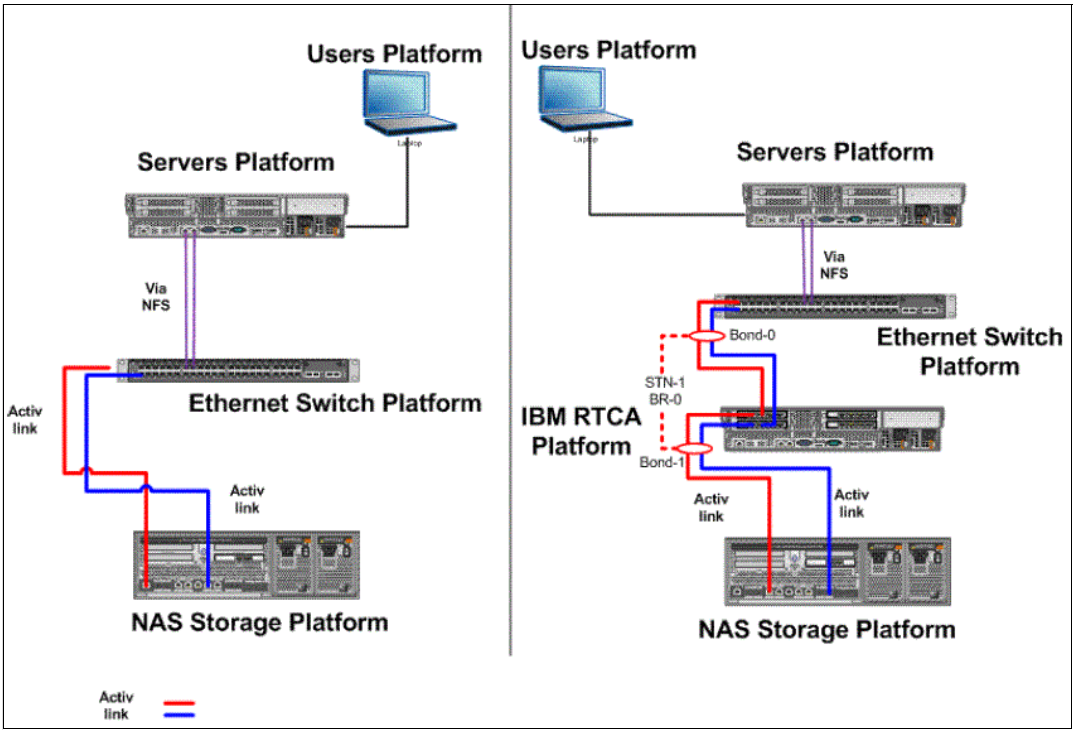

Figure 5-4 Active/active link aggregation: without and with a compression appliance

Use the following steps to configure active/active link aggregation:

1. Test the NAS storage and switch for support of active/active (link aggregation) in direct connection mode.
2. Connect two network cables from the compression appliance to the NAS.
3. Connect two network cables from the compression appliance to the network switch.
4. By using the compression appliance web interface, create two bonds, one for the NAS and one for the network, each holding two physical interfaces.
5. Configure the mode on both bonds as EtherChannel or 802.3ad, depending on your network support at the switch and NAS.
6. Configure a bridge with bond0 and bond1 as members.
7. Configure the file server.
8. Start the compression engine.
9. Continue with the standard configuration.

Path failover using two switches

Configuration for the path failover, which is using two switches, is based according to a single NAS storage controller that uses a link with two or more physical interfaces. Only one link can be active at any particular time.

If the NAS storage controller detects an error in the link that is active, the passive link becomes active with the `UP` status. This method is how passive link becomes active, allowing the traffic that is supposed to be in the NAS storage controller to go through the second switch.

The IBM Real-time Compression Appliance product can be inserted into such configuration using either one or both appliances.

Figure 5-5 shows a path failover configuration with RTCA.

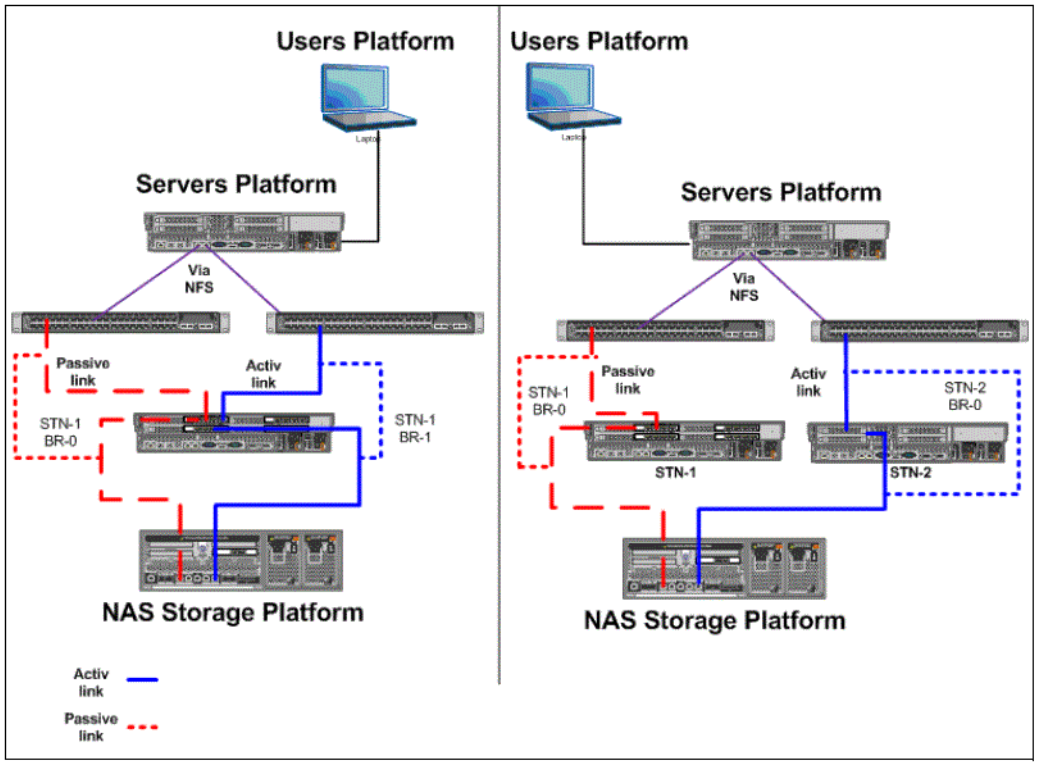

Figure 5-5 Path failover configuration using two switches: option 1 and option 2

The options are as follows:

- Option 1: When the link that is active stops working, the link status mirroring (LSM) feature changes the status of the link on the other side to the DOWN state. The change in the link is detected by the NAS storage controller, which changes the status of the passive link to active and the passive link status is now in the UP state. This option is when the LSM feature changes the link status on the LAN side to the UP state, causing traffic that is supposed to be in the NAS storage controller to travel through the interface, which just became active at the second switch.
- Option 2: When the link that is active stops working, the LSM feature on compression appliance STN-1 changes the status of the link on the other side to the DOWN state. If the compression appliance itself shuts down, the status of the links at both ends (both the switch and the NAS storage controller) changes to DOWN. When the NAS storage controller detects that the status of the link changed, it changes the passive link to active,

causing it to be in the UP state. Furthermore, the LSM mechanism on compression appliance STN-2 changes the status of the link at the LAN side to UP and makes the traffic that is supposed to be in the NAS storage controller go through compression appliance B, and activating the interface on the second switch.

A procedure for configuring path failover, using two switches, is shown in the following steps.

Option 1

Use the following steps:

1. Test the NAS and switch for support of active/passive (failover) in direct connection mode.
2. Connect two network cables from the appliance to the NAS.
3. Connect one network cable from the appliance to the first network switch, and the other network cable from the appliance to the second network switch.
4. By using the compression appliance web interface, create two bonds, one for the NAS and one for the network, each holding two physical interfaces.
5. Configure the mode on both bonds as active/backup.
6. Configure a first bridge STN-1 BR-0 and second bridge STN-1 BR-1.
7. Configure the file server.
8. Start the compression engine.
9. Continue with standard configuration.

Option 2

Use the following steps:

1. Test the NAS and switch for support of active/passive (failover) in direct connection mode.
2. Connect two network cables from the appliance to the NAS.
3. Connect two network cables from the appliance to the network switch.
4. By using the compression appliance web interface, create two bonds, one for the NAS and one for the network, each holding two physical interfaces.
5. Configure the mode on both bonds as active/backup.
6. Configure a bridge with bond0 and bond1 as members.
7. Configure the file server.
8. Start the compression engine.
9. Continue with standard configuration.

   The configuration is defined as two separate appliances, each with a single link to the NAS and a single link to the network. The NAS initiates the move from the active link to the backup.
10. Configure compression filters on both appliances.
11. Continue with the standard configuration.

Clustered environment (preferred solution)

This configuration is based on two NAS storage controllers within the configuration of a clustered failover. Only one of the links can be active at any particular time on each NAS storage controller.

If the NAS storage controller detects an error in the active link, a link that is on standby is activated (if available). If the hardware of one of the NAS storage controllers stops working or

if a request for failover is processed, then a failover procedure starts. The remaining NAS storage controller takes over the NAS storage controller disks, file systems, and IP addresses that stopped working, and starts asking for both entities.

Figure 5-6 shows a clustered environment configuration without and with RTCA.

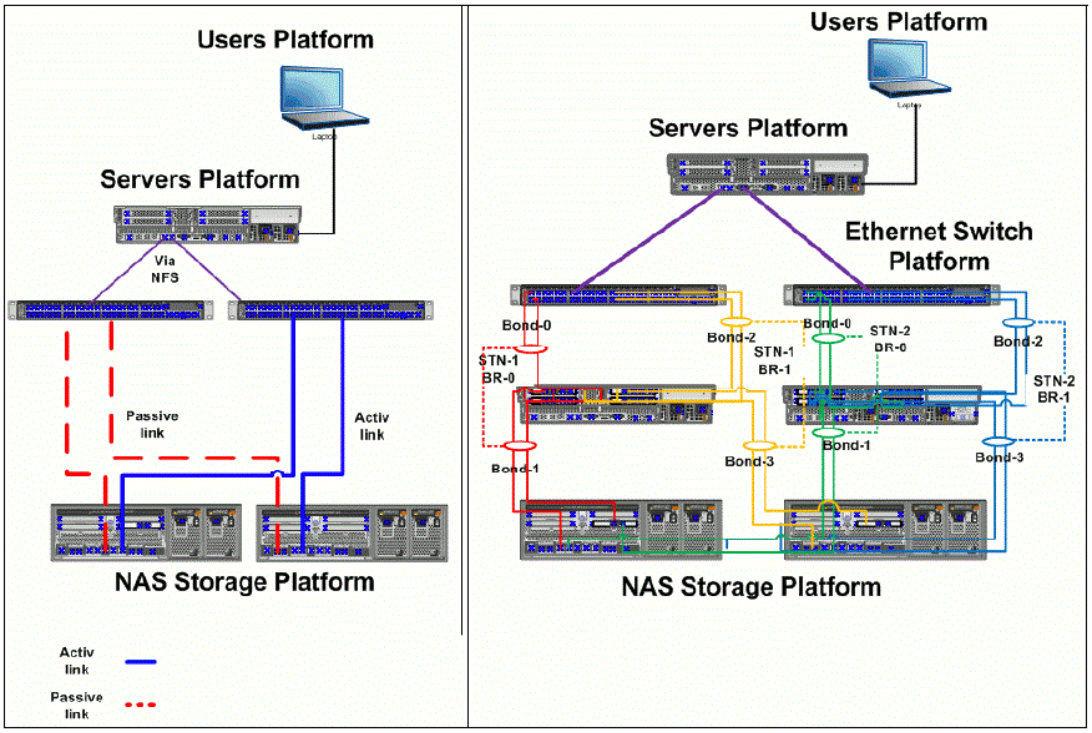

Figure 5-6 Clustered environment without and with compression appliances

Use the following steps to configure a clustered environment:

1. Configure each appliance for active/passive support.
2. Test both NAS storage controllers and switches for support of active/passive (failover) in direct connection mode.
3. Connect two network cables from the compression appliance to the NAS.
4. Connect two network cables from the compression appliance to the network switch.
5. By using the compression appliance web interface, create two bridges, one for each NAS, each holding two physical interfaces.
6. Configure two bridges STN-1 BR-0, STN-2 BR-1 on each compression appliance.
7. Configure the NAS Storage controllers accordingly to the defined bridges.
8. Start the compression engine.
9. Continue with the standard configuration.

Synchronizing high availability

When you configure a compression appliance in the high availability (HA) configuration, both appliances must have the same compression filters, both appliances must be synchronized.

Two modes are available in the process of synchronization:

- Manually synchronizing high-availability mode
- Automatically synchronizing high-availability mode

Manually synchronizing high-availability mode

This process synchronizes the configuration of shares in the high availability environment:

1. Exporting a file that lists the shares for one compression appliance
2. Importing that file to other appliances in the high availability configuration

The following procedure creates manual synchronization in the HA configuration:

1. Click **High Availability** to open the High Availability page.
2. Click **MANUAL**.
3. Click **Exports For High Availability**. The File Download window opens:
4. Click **Save** to save the file on your local computer.
5. Log in to the other compression appliance in the high availability configuration.
6. Browse for the file that was saved from the first appliance.
7. Click **Upload**.

Automatically synchronizing mode high availability

The automatic option for HA synchronization requires enabling a compression appliance to automatically update other appliances (remote partners) in the HA environment.

The following procedure creates automatic synchronization in the HA configuration:

1. Click **High Availability** to add remote partners for automatic synchronization.
2. On the High Availability Auto Synchronization page, verify that **Enable Auto Sync** is selected.
3. In the Remote Partners List table, enter the following information for the partner compression appliance that is to be added:
   - IP address: XYZ
   - User name: XYZ
   - Password: XYZ
4. Click **Apply Changes to Remote Partners and Synchronize**.
5. Click **High Availability** to remove a compression appliance from the list of remote partners.
6. On the Availability Auto Synchronization page, in the Remote Partners List, select **Delete** for the partner compression appliance to remove.
7. Click **Apply Changes to Remote Partners and Synchronize**.

5.4 Reporting considerations summary

Use the IBM Real-time Compression Appliance reporting feature to get reports about the compression savings in a storage environment. The system provides a top-down view about interfaces and shares that are contained across the storage system. The system also supports customizing reports. This feature enables viewing all reports that are offered by the system for all scanned objects or for only a subset of them.

IBM Real-time Compression Appliance has three types of reporting mechanisms:

- Capacity Report: This report shows the capacity consumption and the achieved savings by using the RTCA, per storage systems, interfaces, and in the systems shares. It is available in a a pie chart and textual tabular report that can be exported to a CSV file.
- File Types Report: This report presents the scanned files as they are categorized into file name extensions. It also presents the scanned files as they are discovered in the scan or the file type categories configured in the system. It is available in a bar chart or a pie chart, and in a textual tabular report that can be exported to a CSV file.
- Capacity Growth Report: This report presents charts of the previous scans that were performed with the system. It provides historical graphical view in the last 36 months of the compressed capacity and the uncompressed or effective capacity in the system.

5.5 Validating and testing the installation

The proper installation requires checking and validating all steps during the installation, from the hardware side:

- Proper power cabling installation is completed.
- Proper connectivity cabling installation is completed.
- All hardware properly works after power on.
- All system logs are checked and verified (RTCA logs, Switch logs, NAS Storage logs).
- All host systems are checked after power on.
- All host system application logs are checked and verified.
- All storage shares are properly mapped and verified from the host side.
- All expectations must be consistent with system limitations

6

Software installation or upgrade

This chapter describes how to install or upgrade the RTCA software.

The following topics are included in this chapter:

- 6.1, “Before you begin” on page 56
- 6.2, “Obtaining the software file” on page 56
- 6.3, “Installing and upgrading on a single RTCA” on page 58

6.1 Before you begin

Note: Download the installation file from the following website:

`http://www.ibm.com/support/entry/portal`

Currently, minor software release upgrades can be implemented without stopping the RTCA. For example, software updates for the web interface can be performed without stopping the compression engine. Upgrades to other parts of the software might require restarting the compression engine or rebooting the system. Before you begin, be sure you have a valid configuration backup file.

For additional information about RTCA updates, also see the following online resource:

`http://www.ibm.com/support/docview.wss?uid=ssg1S7003976`

6.2 Obtaining the software file

To obtain the software file, go to the IBM Support website:

`http://www.ibm.com/support`

Ignore any pop-ups that might open.

To more quickly locate available downloads for your product, specify the full type and model or specify part of the type and model in the text box, as Figure 6-1 shows.

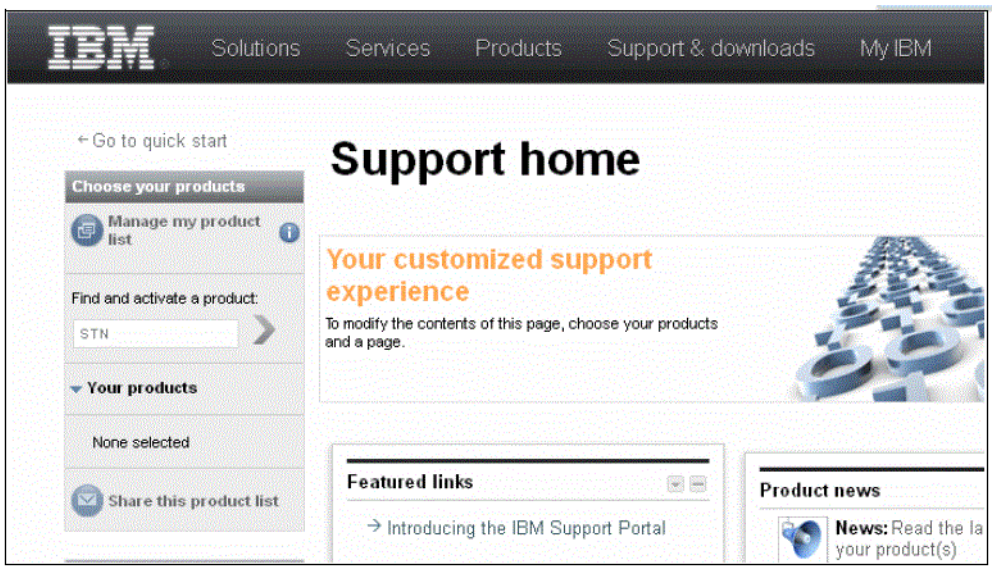

Figure 6-1 IBM Support website

Click the right angle bracket (>) next to the text box. You can then select your product or the website selects it for you if you gave it enough details. For example, we typed `STN`, which was enough information to select the IBM Real-time Compression Appliance. See Figure 6-2 on page 57.

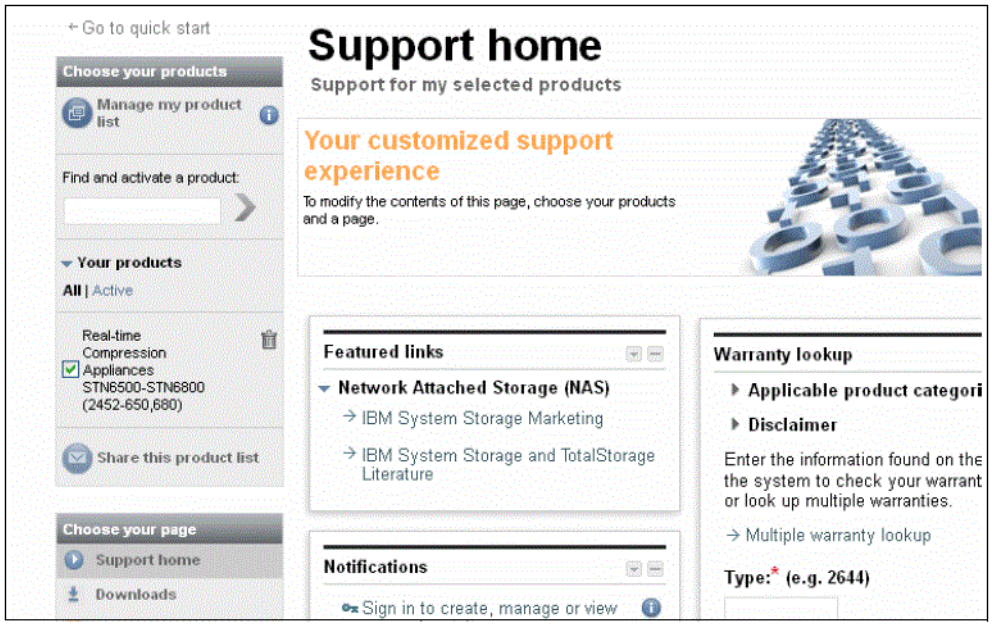

Figure 6-2 IBM Support website

After the correct product is selected, specify that you are looking for available downloads by clicking the **Downloads** section of the Choose your product menu (Figure 6-3).

The Downloads section of the selected product opens, as shown in Figure 6-3. Click **View STN7500 RTCA STN6500-STN6800 downloads**.

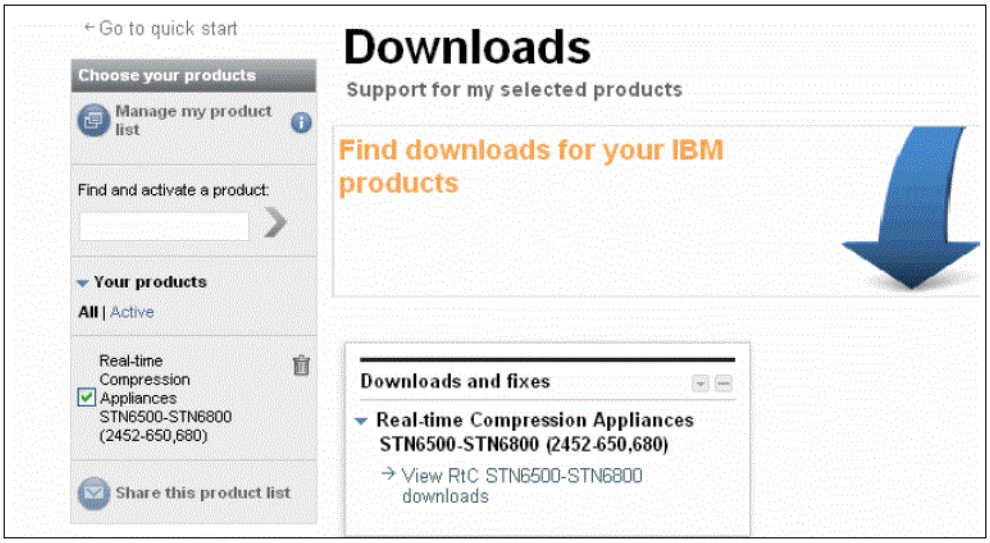

Figure 6-3 Downloads section

In the Available downloads, you can select the fix pack to upgrade the software. Select the check box next to the fix pack and click the **Continue** link.

There is always a license agreement to review and accept. Read it and agree with it so you can download the new software.

Note: When you use HTTP to download, right-click the link to the file and select **Save As**.

6.3 Installing and upgrading on a single RTCA

Important: If a reboot of a single RTCA is required, the NAS storage server is not available during the time of the reboot. You must reboot during a planned downtime.

To install the software, complete the following steps:

1. Log in to the web interface.
2. Click **Upgrade** under the Administration section. See Figure 6-4.
3. Click **Browse** and go to where the software installation file is located (the installation file is named `IBM_Real-time_Compression_4.x.x.xx.stn`) and click **Upgrade**.

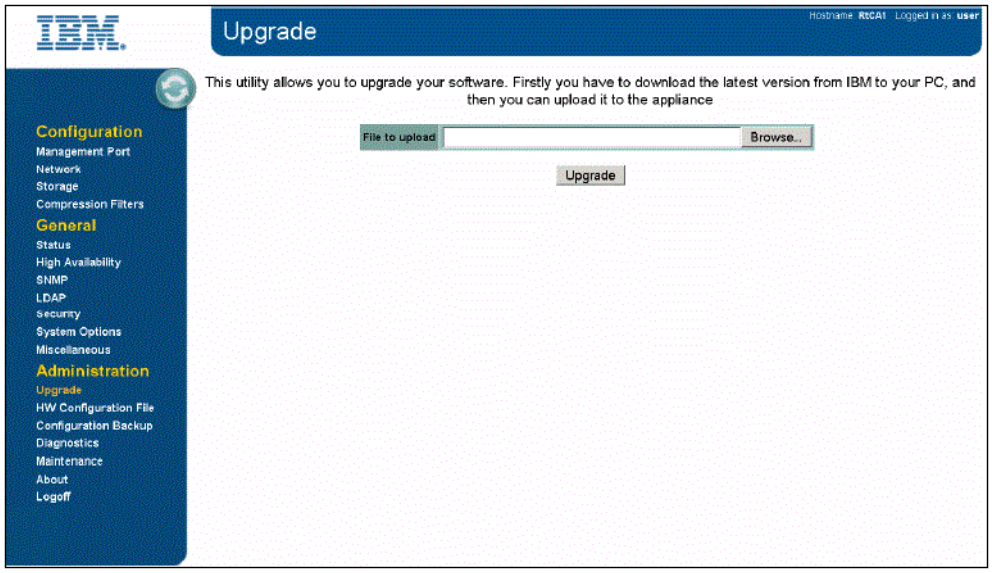

Figure 6-4 Software upgrade

Messages displayed: Uploading the file can take up to five minutes to complete. Several warning messages might be displayed during this time but you can safely ignore them.

After the file is uploaded successfully, read the license agreement (Figure 6-5 on page 59) and accept it. The Install button is displayed but not available for selection until you accept the new license agreement. After you accept the agreement, you can click **Install** to install the new software to the RTCA.

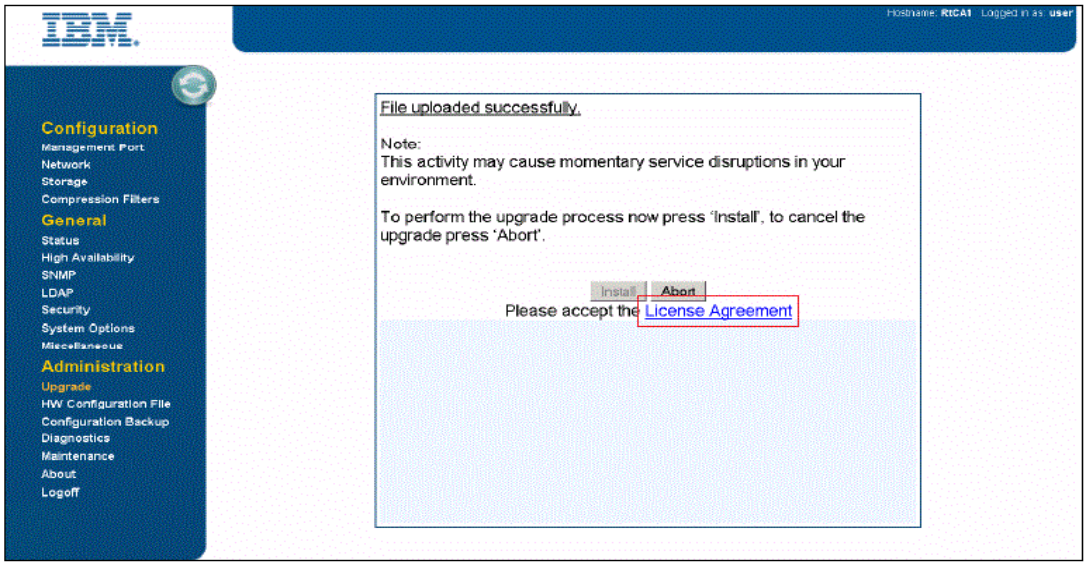

Figure 6-5 License agreement for upgrade

If the software upgrade is successful, an `Upgrade Successful` message is displayed, as Figure 6-6 shows.

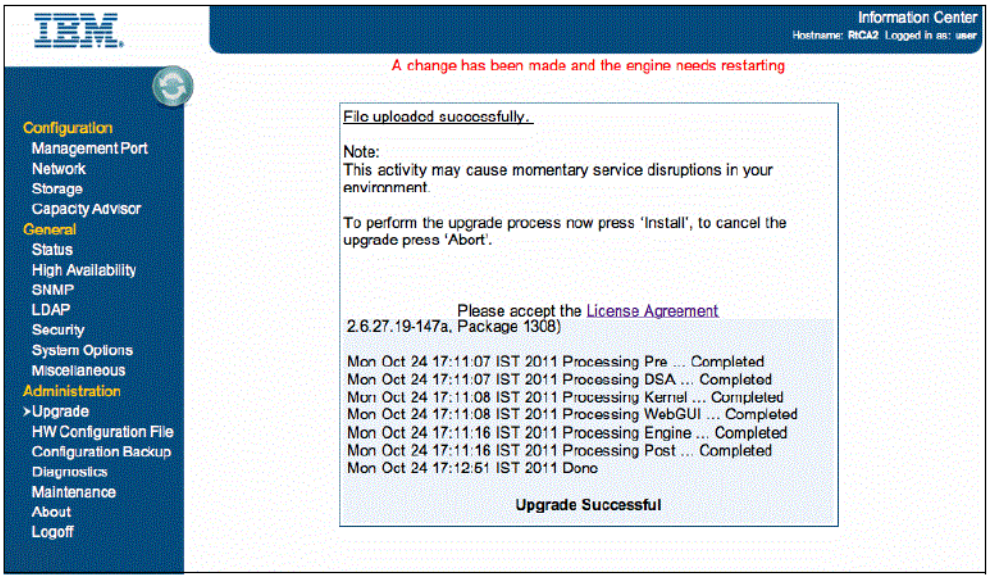

Figure 6-6 Successful upgrade

As previously mentioned, the IBM Real-time Compression Appliance compression engine must be restarted. The following message is displayed on all pages until the engine is restarted:

`A change has been made and the engine needs restarting`

To restart, go to the **Configuration** → **Network** and click **Restart Engine**. See Figure 6-7.

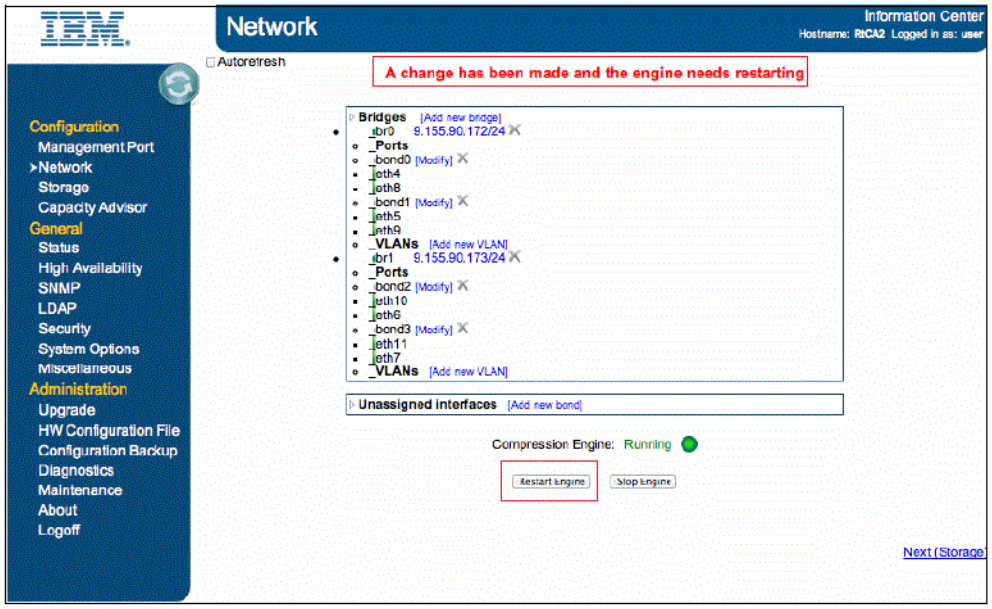

Figure 6-7 Network page of the IBM RTCA: A restart of the engine is necessary

The software upgrade completes.

Important: Make sure that there is no activity through the IBM Real-time Compression Appliance, or connectivity problems can occur.

After the engine is restarted, check the RTCA status panel. It lists the actual firmware level next to Version, under Device Information.

6.3.1 Installing and upgrading on a high availability RTCA system

When installing the software in a high availability environment, ensure that the same version is running on a partner RTCA system. To avoid any downtime on a high availability environment, you must install the software on one RTCA system after another:

1. Perform the same steps as described for installing on a single appliance. On the Status page, verify that the expected version is running after the engine restarted.
2. On a second appliance, perform the steps described for installing on a single appliance. On the Status page, verify that the expected version is running after the compression engine of the second appliance restarted.

Note: The bonds probe is displayed only when the system has configured bonds.

6.3.2 Updating the hardware configuration file

The RTCA requires a valid hardware configuration file. Typically the appliance is preinstalled with a proper hardware configuration file as part of the manufacturing process. However, you might have to update the hardware configuration file. The Hardware Configuration File page enables updating of the compression appliance with details of the hardware configuration file, which specifies the hardware identity of the components in the appliance. Among other information, the hardware configuration file includes the amount of memory in the appliance, its serial number, and the hardware (MAC) addresses that is used in the network interfaces. To update the hardware configuration file, complete the following steps:

1. Select the **HW Configuration File** option under the Administration section. See Figure 6-8.
2. If a hardware configuration file is not installed or provided, click **Generate Hardware Configuration File Request** first, save the file locally, and then send the file to IBM Support.
3. After IBM provides the new hardware configuration file, you can upload it by clicking **Browse** and searching for your hardware configuration file.
4. Select the hardware configuration file, and then click **Upload**.

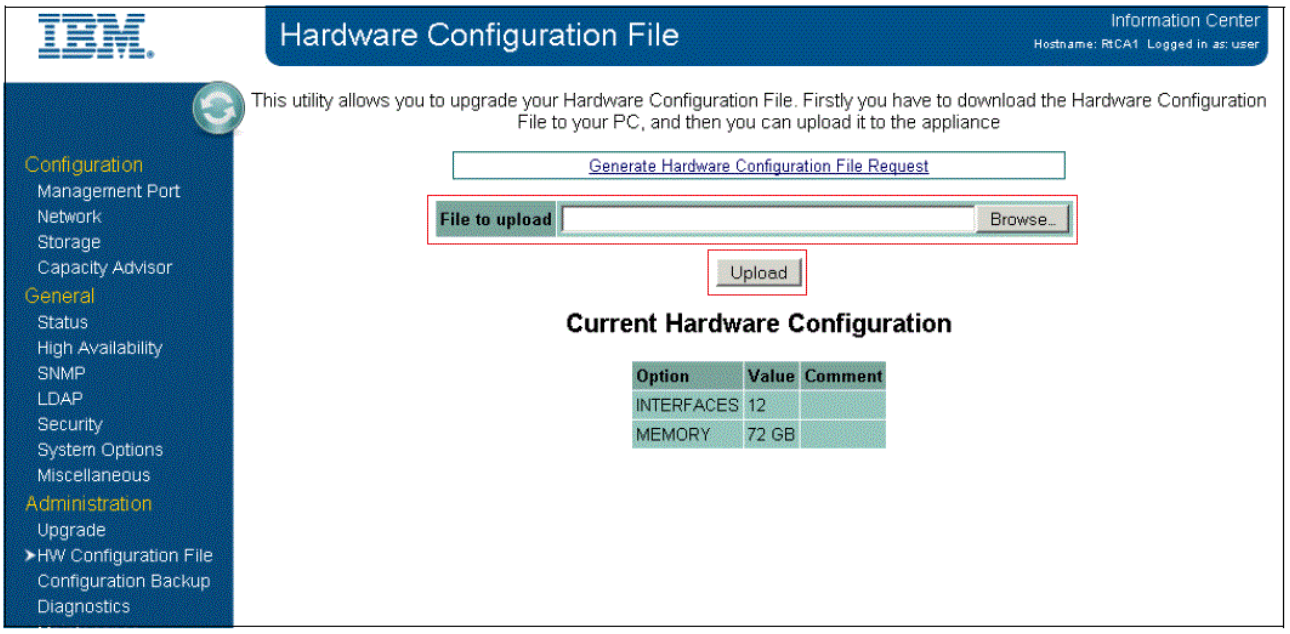

Figure 6-8 Hardware configuration file

6.3.3 Configuring system options

To modify general options for the IBM RTCA product, select **System Options**, under General, as shown in Figure 6-9.

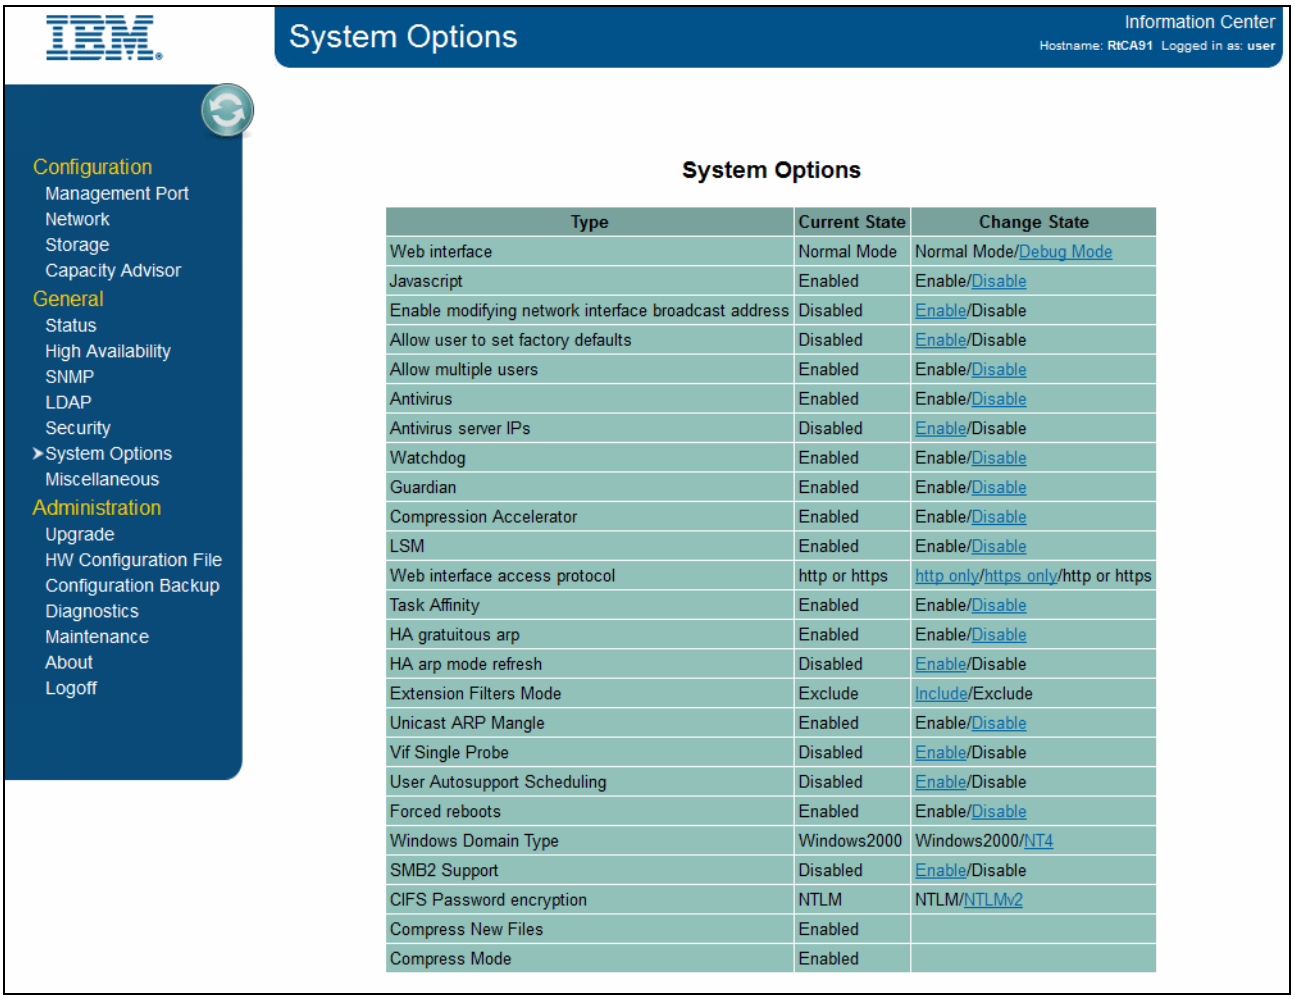

Figure 6-9 System options

The following options are accessible:

- WebGui: The default setting is normal mode. If you change to debug mode, additional debug information is displayed at the top of each page.
- JavaScript: Choose Enable or Disable to activate or stop the use of JavaScript in the web interface pages. JavaScript is mainly used for pop-ups when you use the web interface. The default setting is enabled.
- Network Traffic Capture line: This option shows the status of the tcpdump network capture. This line only displays the status of the *tcpdump* option; to activate tcpdump, select **Diagnostics panel** → **Run tcpdump**.

- Enable modifying network interface broadcast address: Use this option to enable or disable the user from modifying the broadcast address for the management and bridge interfaces. By default, this option is disabled. If you enable this option, the Broadcast field on the Management Port and Modify Bridge pages is enabled. See Figure 6-10.

  Attention: The broadcast address is assigned automatically by the management interface; modifying the broadcast address is typically not required. Consult with your network administrator before you change the broadcast address.

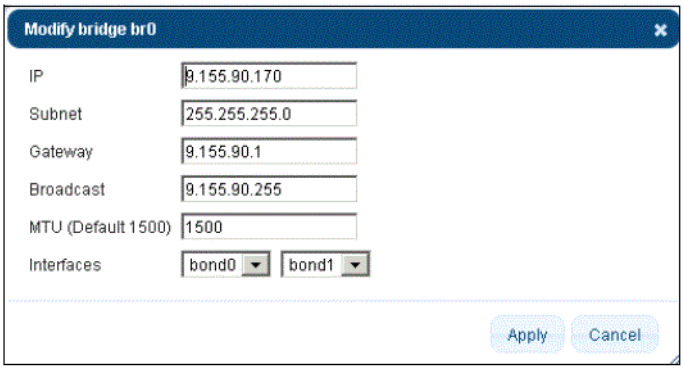

Figure 6-10 Broadcast field

- Allow user to set factory defaults: This option prevents the user from restoring the IBM Real-time Compression Appliance configuration to factory settings. By default, this option is disabled. If you enable this option, a `Restore Factory Defaults` link becomes available in the Configuration Backup panel.

  Attention: Restoring the factory defaults erases all your settings and completely stops the traffic through the appliance. Compressed data becomes available after you configure the RTCA back to your own settings.

- Allow multiple users: Enable or disable access to the web interface from multiple client IP addresses. If this option is enabled (default), multiple users can log in and use the web interface concurrently. If you disable this feature, when an additional client accesses the appliance, the following message is displayed on the login page:

  `Someone else on IP_address is already logged in. Logging in will terminate the other session.`

- Antivirus: Enable or disable support for storage-server-based antivirus products. By default, this option is enabled.

- Antivirus server IPs: Enable or disable support for specific McAfee virus scanning products. By default, support is disabled. If this option is enabled, a new Antivirus Server IP column is displayed in the Storage panel. See Figure 6-11.

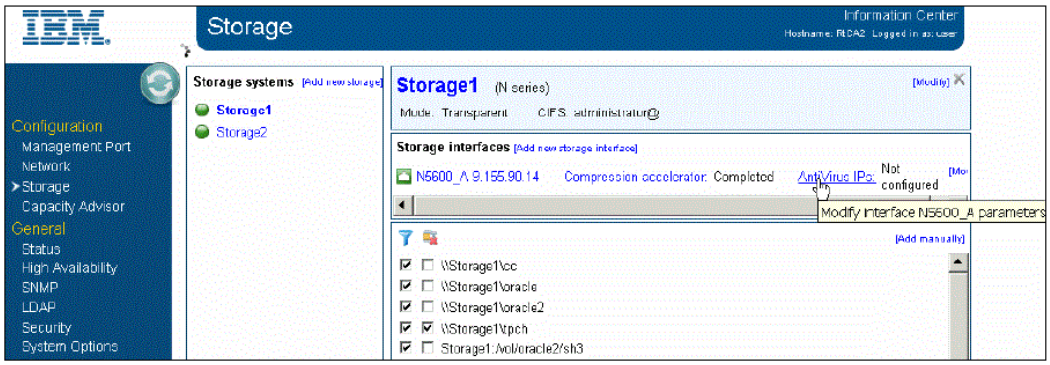

Figure 6-11 System Options - Change Antivirus Server IPs

- Watchdog: Enable or disable the hardware watchdog feature. By default, this option is enabled.
- Guardian: Enable or disable the guardian process that automatically recovers from engine failures. By default, this option is enabled.

> **Guardian and Watchdog difference:** Although Guardian and Watchdog features can both operate with traffic interruption, the difference between them is as follows:
>
> - Guardian can restart the process based on ID only when and if that process hangs.
> - Watchdog issues a complete restart and initialization of the hardware machine.
>
> Because both processes are disruptive, do not disable them unless instructed by IBM Support.

- CA Full Scan: Enable or disable the Compression Accelerator feature from the Storage panel that enables the user to compress existing data. By default, this option is enabled. If you disable this feature, the result might be that a volume will have both compressed and uncompressed data. Do not disable this feature unless instructed by IBM Support.
- LSM: Enable or disable the link status mirroring (LSM) feature described in 2.3.1, “More about RACE” on page 14. By default, this feature is enabled. Do not disable this feature unless instructed by IBM Support.
- Web GUI Access: Enable or disable the protocol used for GUI management. The HTTP-only mode allows only standard HTTP (port 80) access to the web interface; HTTPS-only allows only secured HTTP (port 443) access to the web interface. HTTP or HTTPs (the default) accepts either.

> **Tip:** The web interface supports HTTP and HTTPS. To configure HTTPS certificates, use the same procedure described in “Setting Up SSH authentication” in *IBM Real-time Compression Appliances CLI Command Reference Version 4 Release 1*, GA32-2219-01.

- Task Affinity: Enable or disable the automatic tuning of CPU affinity. By default, this feature is enabled for highest performance as internal modules of the compression appliance are balanced across the CPU cores using optimal CPU affinity. Do not disable this feature unless instructed by IBM Support.

- HA gratuitous ARP: Enable or disable support for passing gratuitous ARP messages. By default, this feature is enabled. Do not disable this feature unless instructed by IBM Support.
- HA ARP mode refresh: Enable or disable support for high availability environments. By default, this feature is enabled. Do not disable this feature unless instructed by IBM Support.
- Extension Filters Mode: Choose Include or Exclude to configure the mode in which the extension filters are applied. By default, the extension filters that are configured and listed at the bottom of the Compression Filters panel define which extensions are excluded from compression (this is the Exclude mode). If Include is configured, the appliance compresses only files with extensions that are defined in the list.
- Unicast ARP Mangle: Enable or disable support for passing unicast ARP messages. By default, this feature is enabled. Do not disable this feature unless instructed by IBM Support.
- Vif Single Probe: Enable or disable support for Data ONTAP single VIF configurations. By default, this feature is disabled. Do not enable this feature unless instructed by IBM Support.
- User Autosupport Scheduling: Allow or prevent the user from modifying the default schedule for Autosupport reports. By default, this feature is enabled.
- Forced reboots: Enable or disable the forced reboots after frequent engine restarts. This option is applicable only for high availability configurations (in single appliance configurations, frequent engine warm restarts eventually become cold restarts; in high availability configurations, the warm restarts become forced reboots if the partner appliance and also its engine are running). By default, this option is enabled.
- SMB2 Pass Thru Support: Enable or disable the support for SMBv2-based CIFS sessions. By default, this option is disabled; CIFS sessions from SMBv2-enabled clients (such as Windows Vista, Windows 7, or Windows Server 2008) are downgraded automatically to SMBv1 protocol.
- Compress New Files: This setting does not change the behavior of the storage system compression mode, which pertains strictly to how new files are created. Use the Compress new files mode in cases where it is important to preserve the system maintenance operations and the current state of shares. For example, a storage system IP address is configured as Compress shares/exports except selected in its Compression Mode field, and the storage administrator wants to create a share that must work in uncompressed mode, so a compression filter is created after the share is created on the storage.
- Compressed mode: This feature changes the default behavior of RTCA when new storage is added.

From the CLI, use the **`options`** command followed by the appropriate option and either the **`on`** or **`off`** parameter to enable or disable that feature as presented Example 6-1 on page 66.

Values 0 and 1: A value of 0 (zero) is associated with the **`off`** option; a value of 1 is assimilated with the **`on`** option.

Example 6-1 Check “System Options” status

```
user@9.155.66.158> option
CA_FullScan 1
ForcedReboots 1
LSM 1
VifSingleProbe 0
allowSelectCompressedMode 1
allowSetFactoryDefaults 1
allowTriggersScheduling 0
antivirus 1
avip 0
extensionFiltersMode exclude
guardian 1
haArpRefreshMode 1
haGratuitousArp 1
javascript 1
multipleUsers 1
noNewCompression 1
probeBonds 1
probeDuplicateIPs 1
probeGateways 1
probeVendors 1
setbroadcastmode 0
smb2passthrusupport 0
taskAffinity 1
unicastArpMangle 1
watchdog 1
webGuiAccess http_https
webgui (debug mode) 0
```

To list all options, type the word `option` and then press the Tab key. This step can help you identify the correct parameter. See Example 6-2.

Example 6-2 List options

```
user@9.155.66.158> option
CA_FullScan allowSetFactoryDefaults guardian
noNewCompression setbroadcastmode webGuiAccess
ForcedReboots allowTriggersScheduling haArpRefreshMode
probeBonds smb2passthrusupport webgui
LSM antivirus haGratuitousArp
probeDuplicateIPs taskAffinity
VifSingleProbe avip javascript
probeGateways unicastArpMangle
allowSelectCompressedMode extensionFiltersMode multipleUsers
```

Example 6-3 shows changing a parameter.

Example 6-3 Change a parameter

```
user@9.155.66.158> option allowSetFactoryDefaults off
```

7

Performance considerations

Several factors must be considered when you design an optimal configuration with the IBM Real-time Compression Appliance (RTCA).

A key benefit of the RTCA solution is being able to perform compression while maintaining high performance. The compression appliance is part of the overall storage solution, and as such, it must be reviewed together with the storage system itself. Various performance requirements must be part of this review.

The following topics are included in this chapter:

- 7.1, “Performance requirements” on page 68
- 7.2, “Performance suggestions” on page 69

7.1 Performance requirements

The focus areas for performance are throughput, response time, and scalability:

- Throughput: The storage solution is typically measured by the number of megabytes per second of data being read or written to it. We can measure this throughput during work hours and also peak time, to make sure the system is designed to provide expected service levels, even in high load time windows.
- Response time: Users are sensitive to how responsive the system is to the actions and jobs that they perform in the computing systems. The user response time is affected by many components: user actions themselves, client-side application processing time, operating systems, networking layers, physical links, file system processing, and the disk sub-system.
- Scalability: Any computer system is typically designed to handle current and future load levels. Planning ahead requires a detailed review of the performance metrics of an overall storage solution.

One of the major objectives of RTCA v4.1 is to improve the throughput. While significant performance boost is gained by the new IBM System x3650 M4 hardware platform, which supports the Xeon E-2600 CPU, there are also software changes that are targeted to improve the performance. For more information about IBM System x3650 M4, read Chapter 4, “Hardware update” on page 29.

Release v4.1 also delivers higher performance compared to previous releases. The newly introduced hardware is suitable to handle higher performance loads and has an improved and optimized software operating system.

Next we review processes for configuring the highest performance Real-time Compression Appliances solutions.

7.2 Performance suggestions

Various configuration items and parameters affect the performance of a Network Attached Storage (NAS) setup and that includes the IBM Real-time Compression Appliance solutions.

7.2.1 Interval between files

This feature can be set from RTCA GUI. It defines how many milliseconds to wait before a new file is processed by the RTCA engine, and describes how to improve the Compression Accelerator performance.

Use the following procedure to change the interval between files:

1. Click **Storage**, as shown in Figure 7-1, and select the correct storage name. We selected the N3300A option.

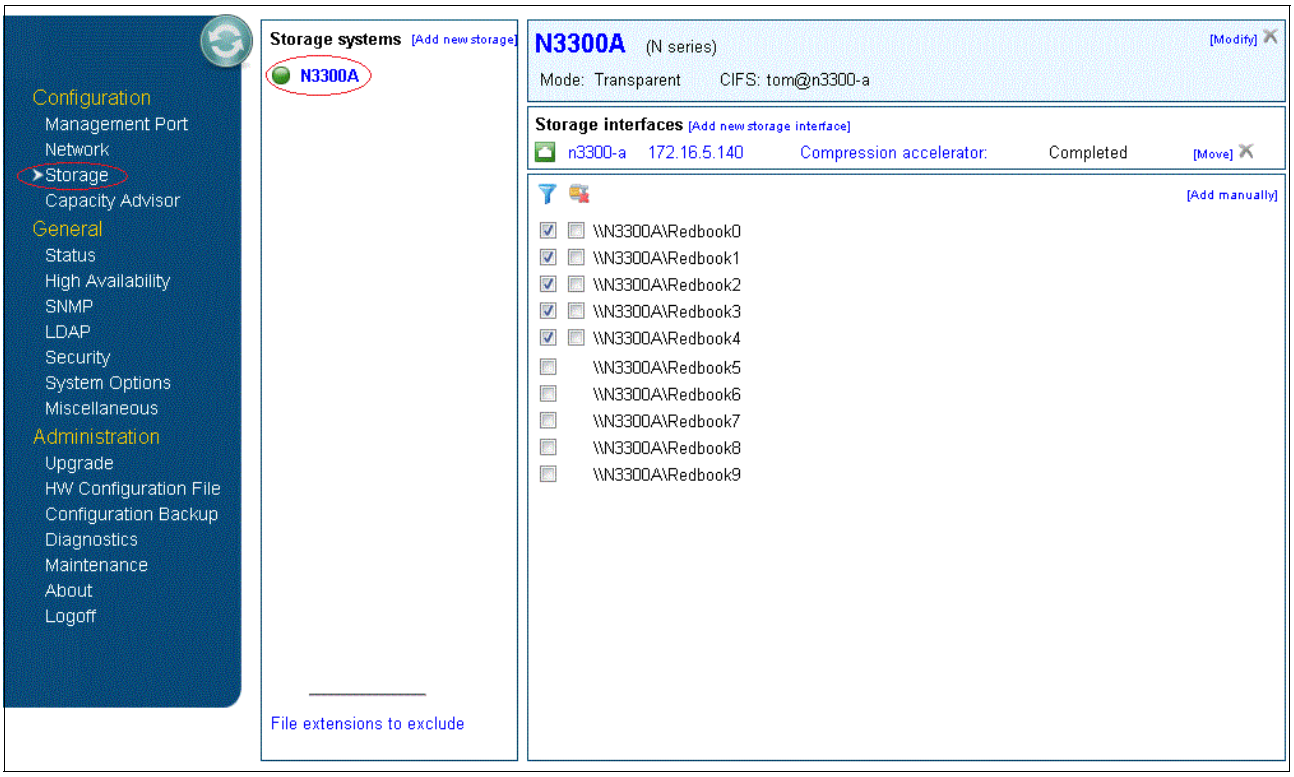

Figure 7-1 "Storage" configuration panel

2. For the N3300A selection, which shows what kind of files are compressed and which mode is selected (transparent or compressed), click **Compression accelerator**, as shown in Figure 7-2 on page 70.

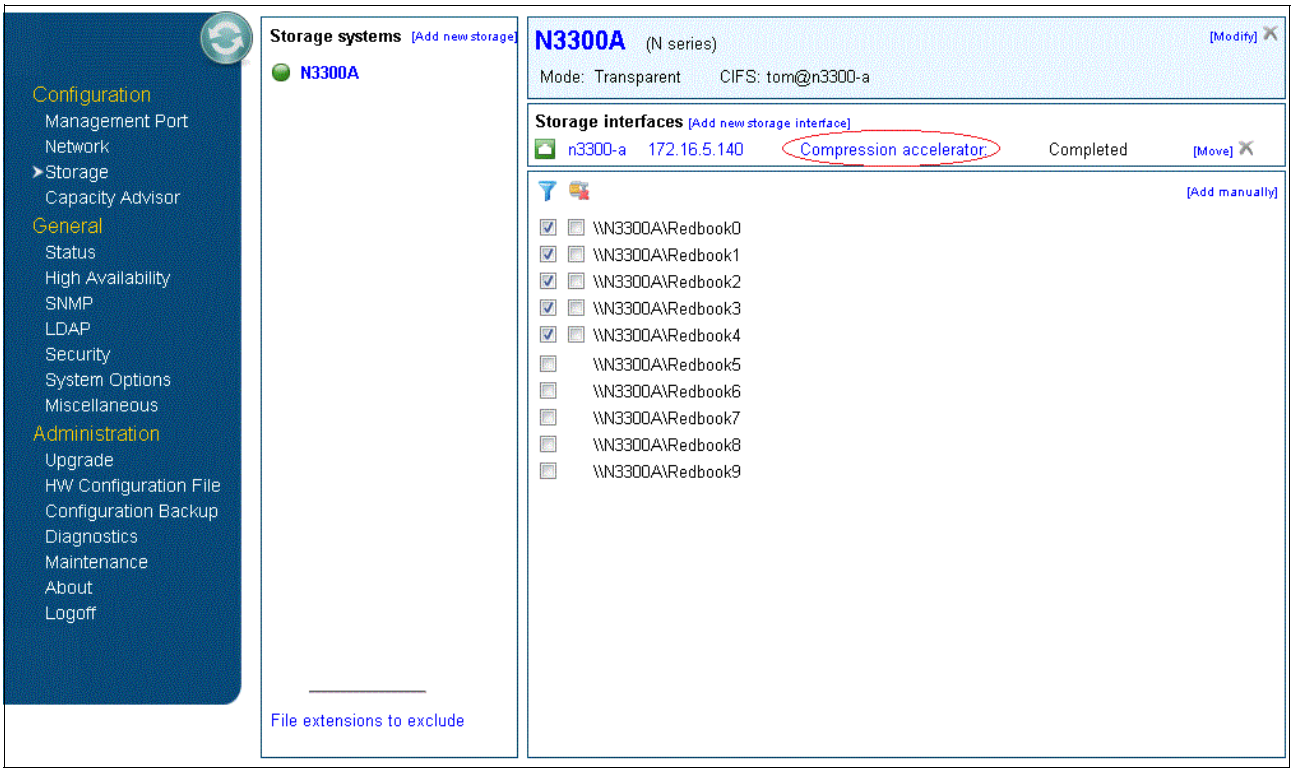

Figure 7-2 “Compression Accelerator” option

3. Change the **Interval between files (millisecond)** value to 1 (Figure 7-3) to get the highest performance.

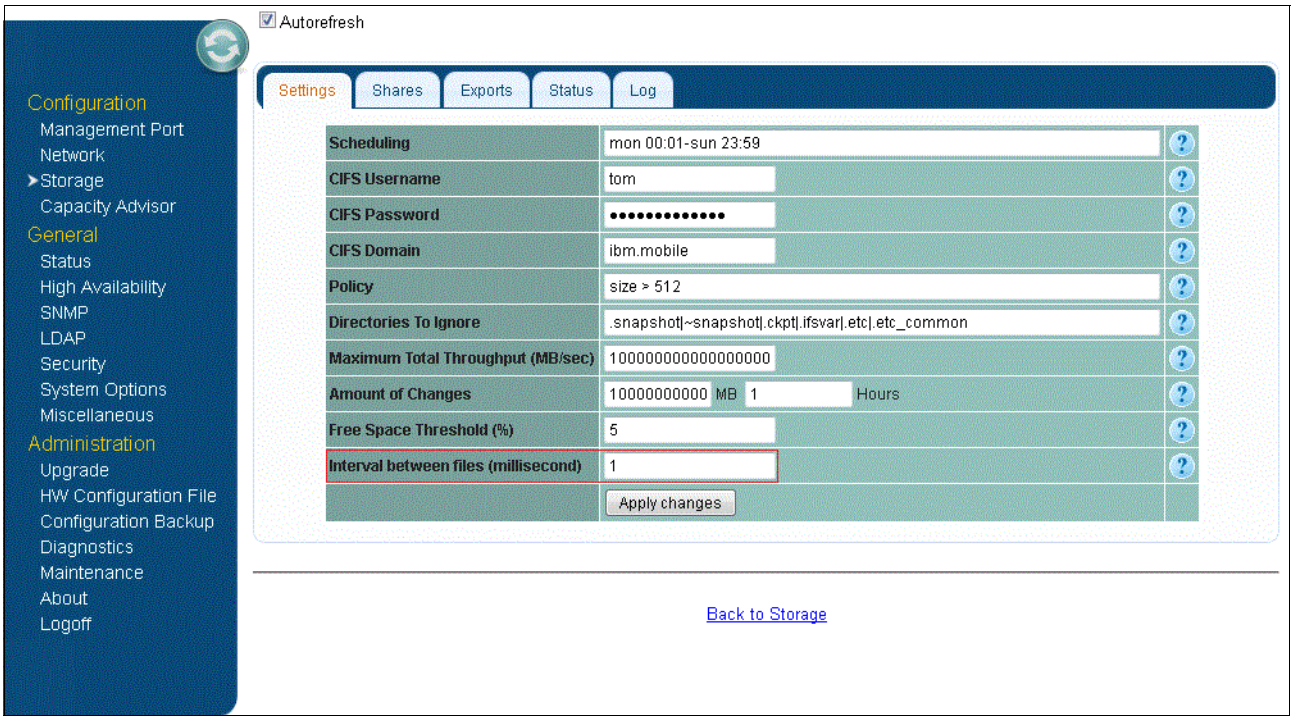

Figure 7-3 “Interval between files (millisecond)” option

7.2.2 Compression ratio

A useful principle is that the higher the compression ratio, the lower the CPU time that is spent on compressing the data. The same rule applies for decompression also. Therefore, a preferable approach is to configure file types with low expected compression ratio in the File Extensions to Exclude field on the Storage page. This approach reduces CPU time that is spent attempting to compress files with low expected file size reduction.

Use the following procedure to create a list of file extensions to exclude:

1. Click **Storage**, and select **File extensions to exclude**, as shown in Figure 7-4.

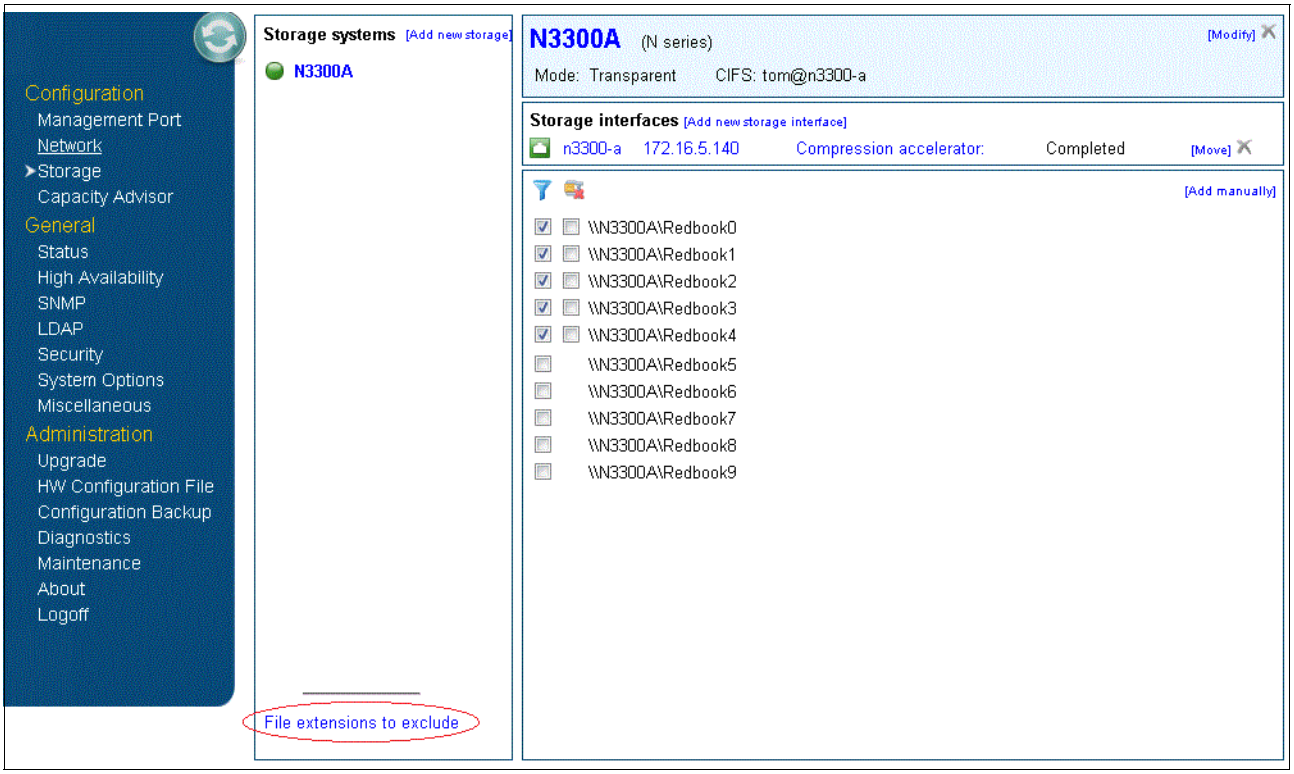

Figure 7-4 "File extensions to exclude" option

2. Specify the list of file extensions to exclude, as shown in Figure 7-5. For example, specify the following extensions:

```
zip gz tgz arj xlm
```

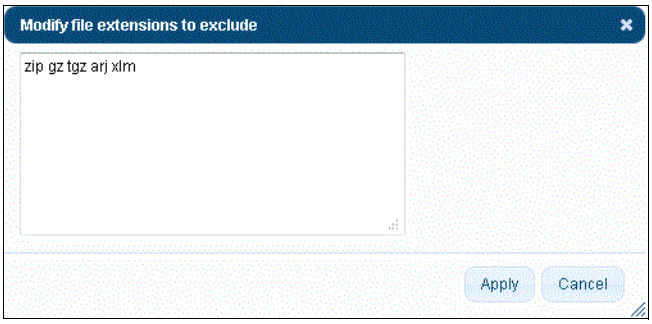

Figure 7-5 List of file extensions to exclude

Important: Changing file extension filters might cause client session disconnections on this appliance and affect partner appliances (in a high-availability configuration).

3. Click **Apply** to make the changes to file extensions.

7.2.3 Transparent versus compressed traffic

A basic concept of the compression appliance is that it adds minimal latency when traffic arrives and gets processed by its network interfaces. But, it reduces the latency of I/O requests that are sent to the storage, because less data is sent to the storage compared with the original data that is sent by the client.

However, suppose that the appliance is implemented without compressing data, or when only a minimal amount of data is compressed. In this case, the latency that is added by the appliance, being an additional network hop, will not be compensated, and the overall throughput will be lowered. It is best to implement the appliance and start compressing files right away.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topics in this document. Note that some publications referenced in this list might be available in softcopy only:

- *IBM Real Time Compression Appliance Application Integration Guide,* SG24-8029
- *IBM Real-Time Compression Appliance: New Features in Version 3.9*, REDP-4879
- *IBM Real Time Compression Appliance: What is new in RTCA Version 3.8*, REDP-4861
- *IBM System Storage N series Hardware Guide*, SG24-7840
- *IBM System Storage N series MetroCluster,* REDP-4259
- *IBM System Storage N series MetroCluster Planning Guide,* REDP-4243
- *IBM System Storage N series Software Guide,* SG24-7129
- *IBM System Storage Solutions Handbook,* SG24-5250
- *Implementing IBM Storage Data Deduplication Solutions*, SG24-7888
- *Introduction to IBM Real-time Compression Appliances*, SG24-7953
- *NAS / N series Solution Design for Real-time Compression Appliances*, REDP-4770

You can search for, view, or download Redbooks publications, Redpaper publications, Technotes, draft publications, and Additional materials, and also order hardcopy Redbooks publications, at this website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *IBM Real-time Compression Appliances Installation and Planning Guide Version 4 Release 1*, GI13-1221-04
- *IBM Real-time Compression Appliances Configuration and Administration Guide Version 4 Release 1*, GA32-2217-01
- *IBM Real-time Compression Appliances CLI Command Reference Version 4 Release 1*, GA32-2219-01
- The following location lists other publications:

  http://pic.dhe.ibm.com/infocenter/rtca/rtcaic/index.jsp?topic=%2Fcom.ibm.storage.rtca.help.doc%2Fcmn_r_publications.html

Online resources

These websites are also relevant as further information sources:

- IBM Real-time Compression:
  http://www.ibm.com/systems/storage/solutions/rtc/
- IBM Real-time Compression Appliance STN6500:
  http://www.ibm.com/systems/storage/network/rtc/stn6500/
- IBM Real-time Compression Appliance STN6800:
  http://www.ibm.com/systems/storage/network/rtc/stn6800/index.html
- IBM RTCA Support overview:
  http://www.ibm.com/storage/support/rtc
- Reducing NAS Costs with Real-time Data Compression:
  http://wikibon.org/wiki/v/Reducing_NAS_Costs_with_Real-time_Data_Compression
- Network attached storage:
  http://www.ibm.com/systems/storage/network/
- IBM support; documentation:
  http://www.ibm.com/support/entry/portal/Documentation

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Redbooks

IBM Real-time Compression Appliance Version 4.1

(0.2”spine)
0.17”<->0.473”
90<->249 pages

IBM Real-time Compression Appliance Version 4.1

See how seamless compression integrates into storage environments

Get familiar with the leading edge real-time compression solution

Learn how to set up and configure the appliance

Continuing its commitment to developing and delivering industry-leading storage technologies, IBM is introducing the IBM Real-time Compression Appliance for NAS, an innovative new storage offering that delivers essential storage efficiency technologies, combined with exceptional ease of use and performance.

In an era when the amount of information, particularly in unstructured files, is exploding, but budgets for storing that information are stagnant, IBM Real-time Compression technology offers a powerful tool for better information management, protection and access. IBM Real-time Compression can help slow the growth of storage acquisition, reducing storage costs while simplifying both operations and management. It also enables organizations to keep more data available for use rather than storing it offsite or on tape that is more difficult to access, so they can support improved analytics and decision-making.

IBM Real-time Compression Appliance provides online storage optimization through real-time data compression, delivering dramatic cost reduction without performance degradation.

This IBM Redbooks publication is for system administrators and IT architects. It describes the enhancements made in version 4.1 of the Real-time Compression Appliance as compared to previous releases.

This book is a companion to the publication Introduction to IBM Real-time Compression Appliances, SG24-7953.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-8112-00 ISBN 073843793X