ibm.com/redbooks

IBM Real-time Compression on the IBM XIV Storage System

Guenter Rebmann

Albert Verhoeff

Megan Gilge

Cloud

Storage

Redpaper

Executive overview

Today's IT companies are driven by continued data growth and economic pressure of huge amounts of data that needs to be stored on storage devices for fast access. Moving this data to lower tier devices like tape or optical devices (CD or DVD) is only a partial solution.

To comply with this requirement in today's business environment, enterprises are forced to install more storage devices, increasing purchase costs, maintenance costs, and the overall costs of data processing service centers. Implementation of more storage systems also creates these requirements:

- More space for additional racks
- Increasing costs for power and cooling
- Additional capital for licenses and software
- Larger investment to manage operational expense (OPEX)

IBM® Real-time Compression™ fully integrated in the IBM XIV® Storage System Gen3 with software 11.6 is the answer to this challenge. It provides the possibility to store 2 - 5 times more data per XIV system, without additional hardware, and effectively reduces software cost. This technology further extends the storage-replication-related bandwidth and decreases OPEX.

Using compression for replication and for volume migration with IBM Hyper-Scale Mobility is faster and requires less bandwidth for the interlink connections between the IBM XIV storage systems because the data that is transferred through these links is already compressed.

IBM Real-time Compression uses patented IBM Random Access Compression Engine (RACE) technology, achieving field-proven compression ratios and performance with compressible data. This IBM Redpaper™ publication describes the benefits of implementing IBM Real-time Compression.

This feature is introduced with XIV software version 11.6, which will be available June 4, 2015. IBM Real-time Compression with the IBM XIV Storage System requires an additional license.

© Copyright IBM Corp. 2015. All rights reserved.

Introduction

Data compression must be fast, reliable, and scalable. The compression algorithm that is used must ensure data consistency and a high compression ratio. The data compression solution must be easy to implement. The compression must occur without affecting the production use of the data at any time.

Current IT challenges

Businesses and organizations around the world are challenged with tough economic conditions. The IT environment, which historically was viewed as an expense, is now viewed as a source of innovation that can drive future revenue. However, increasing data storage requirements consume the available resources and disrupt attempts to innovate in the IT environment. The modern IT department has these challenges:

- **Support for increasing data storage requirements.** Shrinking IT budgets are pressuring IT managers to increase the lifetime of existing storage systems. Traditional methods to clean up unneeded data and archive files to auxiliary storage are time consuming. They shift one resource constraint, physical storage, to another: The human work of storage administrators.
- **Power, cooling, and floor space.** A data center provides the means to host the storage systems. However, the physical characteristics of hard disk drive-based systems limit the amount of data that can be stored per rack unit. High-power consumption and heat dissipation are major concerns for IT managers who must fit the storage systems into a limited data center. This conflicts with the increasing demand for computing power that is needed to support new types of applications.
- **High availability of data.** Digital information has become the basis for any service in use today. As a result, the underlying systems that provide access to digital information are expected to be online all the time. This requirement has made it impossible to introduce data reduction solutions that impose any type of downtime. This restriction is true whether it is an inability to access the data, or merely a major slowdown when accessing an optimized data set.

Compression of primary storage provides an innovative approach that is designed to overcome these challenges.

The solution: IBM Real-time Compression

The IBM Real-time Compression solution addresses the challenges that are listed in “Current IT challenges” because it was designed from the ground up for primary storage. Implementing IBM Real-time Compression in XIV provides the following benefits:

- Compression for active primary data. IBM Real-time Compression can be used with active primary data. Therefore, it supports workloads that are not candidates for compression in other solutions. The solution supports online compression of existing data. It allows storage administrators to reclaim disk space in an existing storage system without requiring administrators and users to clean up or archive data. This capability enhances the value of existing storage assets, and the benefits to the business are immediate. The capital expense of upgrading or expanding the storage system is delayed.
- Compression for replicated or mirrored data. Remote volume copies are always compressed if the source is compressed. This process not only reduces storage requirements, but also uses less bandwidth because the data is transferred compressed.

- No changes to the existing environment are required. IBM Real-time Compression is part of the storage system. It was designed with transparency in mind so that it can be implemented without changes to applications, hosts, networks, fabrics, or external storage systems. The solution is not visible to hosts, so users and applications continue to work as-is. Compression occurs within the XIV system itself. The conversion from non-compressed to compressed is in-line and does not require downtime.
- Overall savings in operational expenses. More data can be stored in a system, so less physical capacity is required to store a data set. This reduced rack space has the following benefits:
  - Reduced power and cooling requirements. More data is stored in a system, requiring less power and cooling per gigabyte or used capacity.
  - Reduced software licensing for your backup applications. More data stored per enclosure reduces the overall spending on licensing.

Key features and functions

The IBM RACE implementation for IBM Real-time Compression is based on industry requirements. It is a combination of a lossless data compression algorithm with a real-time compression technology.

The RACE technology is the core of the IBM Real-time Compression solution, and it is available in these IBM storage systems:

- IBM XIV Gen3
- IBM Spectrum Virtualize (SAN Volume Controller)
- IBM Storwize® V7000 and V7000 Gen2
- IBM FlashSystem™ V840 and V9000

RACE technology is based on over 70 patents that are not primarily about compression. Rather, they define how to make industry-standard compression of primary storage operate in real time and allow random access. The primary intellectual property behind this is the RACE. At a high level, the RACE component compresses data that is written into the storage system dynamically. This compression occurs transparently, so Fibre Channel and iSCSI connected hosts are not aware of the compression. RACE is an in-line compression technology, meaning that each host write is compressed as it passes through the XIV system to the disks. This has a clear benefit over other compression technologies that use post processing. These technologies do not provide immediate capacity savings, and therefore are not a good fit for primary storage workloads, such as databases and active data set applications.

The following sections describe where RACE technology can be applied and the expected results in addition to tools to analyze workloads.

Comprestimator

It is important to determine which volumes are potentially worth compressing and which are not. The comprestimator utility can be used to estimate an expected compression ratio by using the IBM Real-time Compression technology. The utility uses advanced mathematical and statistical algorithms to perform the sampling and analysis process in a quick and efficient way.

This utility is embedded into the XIV software to estimate savings even before you activate compression. This way you can, without knowing the data types, analyze whether compression will have benefits on a specific volume. The comprestimator sends only read commands, so it has no effect on the data that is stored on XIV.

Tip: Comprestimator is also available as a stand-alone tool that can be used to estimate compression savings for data that is not yet on XIV storage.

The utility has an accuracy range that deviates plus or minus 5 percent based on the formulas that are used by the RACE implementation.

For more information, see the following website:

http://www-304.ibm.com/webapp/set2/sas/f/comprestimator/home.html

Use cases

IBM Real-time Compression is appropriate for data that has the following characteristics:

- Any data for which the comprestimator tool estimates 25% or higher savings
- Volumes that contain data that is not already compressed (for example, uncompressed image and video files)
- Data for which application-based encryption is not used or data that is not sent encrypted to the XIV. Self-encrypting drives (SEDs) and other back-end encryption methods are supported.

This section describes the most common use cases for implementing compression:

- General-purpose volumes
- Databases
- Virtualized infrastructures

Explanation: The expected compression ratios in this section are based on customer data that is gathered from the field. Real-life ratios can vary from the stated values, and depend on individual customer environments and data structures.

General-purpose volumes

Most general-purpose volumes are used for highly compressible data types, such as home directories, CAD/CAM, oil and gas data, and log data. Storing such types of data in compressed volumes provides immediate capacity reduction to the overall used space. More space can be provided to users without any change to the environment.

Many file types can be stored in general-purpose volumes. The estimated compression ratios are based on field experience. Expected compression ratios are 50% - 80%.

Databases

Database information is stored in table space files. It is common to observe high compression ratios in database volumes. Examples of databases that can greatly benefit from IBM Real-time Compression are IBM DB2®, Oracle, and Microsoft SQL Server. Expected compression ratios are 50% - 80%.

Tip: Some databases offer optional built-in compression. In some cases, IBM Real-time Compression can provide additional compression to save even more space.

Virtualized infrastructures

The proliferation of open systems virtualization in the market has increased the use of storage space, with more virtual servers, virtual desktop infrastructure (VDI), and backups kept online. The use of compression reduces the storage requirements at the source.

Examples of virtualization solutions that can greatly benefit from IBM Real-time Compression are VMware, Microsoft Hyper-V, and KVM. Expected compression ratios are 40% - 75%.

> **Tip:** Virtual machines with file systems that contain compressed files are not good candidates for compression.

Architecture

Unlike other IBM Real-time Compression implementations, XIV uses an *above cache* architecture, where the data is compressed or decompressed between the I/O interface and the cache. Reads are stored compressed in cache, and are decompressed after they are read from cache. Data in cache and SSD (as an extension of the read cache) is stored compressed, which means that effectively more data is cached.

Figure 1 shows the flow of data with and without IBM Real-time Compression.

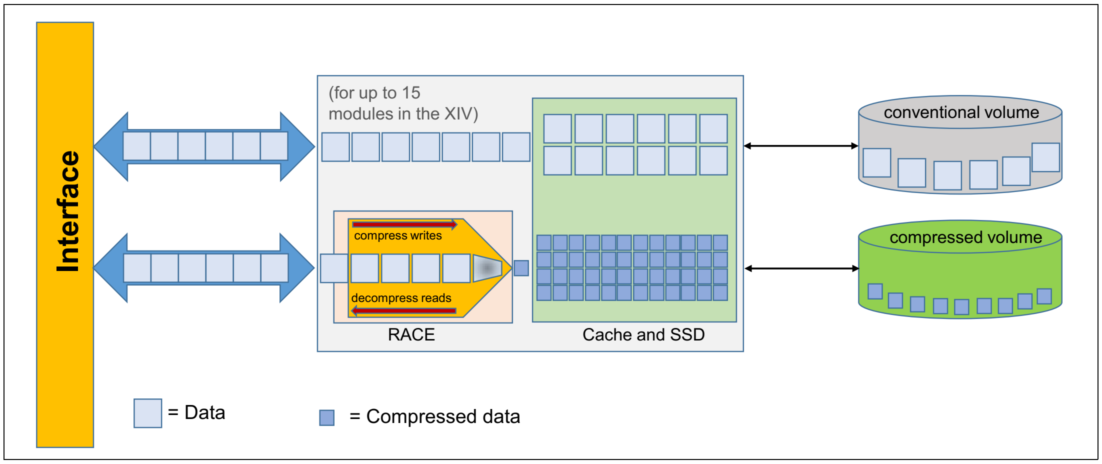

Figure 1 The flow of data with and without IBM Real-time Compression

A new software node, using RACE, is added to each module. It compresses only the parts of the volume that belong to the module. This implementation means that IBM Real-time Compression takes advantage of the massively parallel XIV design. The additional workload that is generated by IBM Real-time Compression is distributed to all modules. Because of this, IBM Real-time Compression in XIV has nearly no impact on the overall performance of the XIV. When there is more data in cache and SSD, compressed volumes can have a better cache hit ratio and less latency.

Volumes that are not compressed bypass this software node and follow the usual data path.

Writes

Writes are compressed by RACE inflight before they enter the cache. After the compressed data is written into cache, the write is acknowledged back to the host. The added latency is proportional to the block size. RACE gets a maximum I/O size of 64 KB from the interface. If the I/O size is larger than 64 KB, the interface splits the I/O.

Reads

Reads are decompressed inflight when they are read from cache before they are passed back to the interface. IBM Real-time Compression uses the XIV prefetch mechanism for compressed application data.

Random read requests from the host also become random reads from the storage because RACE needs to decompress each read request. Because the cache is compressed, the effective disk cache is significantly larger. This increase improves the cache hit ratio and the random read performance.

Performance information

IBM Real-time Compression compresses a data stream as it is written and uses an "above cache" architecture as described in "Architecture" on page 5. When RACE is used, an incoming write stream turns into a sequential stream of compressed data. This process occurs even if the incoming write stream is random. Thus, any random small block I/O write stream is converged into a single chunk of data to be compressed. This single chunk contains multiple independent writes that are each compressed independently. The compressed data is then written.

In real-life applications, when this data is read back, the compression engine reads the sequential stream of compressed data. Because this stream contains more data than is requested in this I/O request, the cache hit ratio might be effectively higher.

This implementation typically improves system performance because you have more available cache because the data is compressed above cache. There is also less traffic if you destage down to the physical disk, meaning that XIV is required to handle less I/O.

For real-world applications, I/O is almost never random. The reality is that an application reads and writes objects or groups of data. These groups of I/O requests form a repeatable pattern, with the same group of I/Os occurring one after another, even if they are written to random locations on disk. IBM Real-time Compression analyzes these patterns to provide better compression ratios and return the best performance.

Application benchmark results

This section contains the results of performance benchmark testing by using typical XIV Gen3 model 114 and 214 system configurations.

Preliminary: The performance results in this publication are preliminary and are likely to change.

Benchmark testing configurations

The main user benefit of compression is data reduction. The testing configurations that are used for the benchmarks in this paper had the configurations that are listed in this section.

The test used six IBM x3550 M4 hosts with the configuration that is shown in Table 1.

Table 1 Host configuration

Hardware	Description
Memory	128 GB
CPU	Xeon E5-2650 at 2.00 GHz

Hardware	Description
Fibre Channel HBA	QLogic QLE2562 8Gb FC dual-port

Table 2 shows the XIV Gen3 configurations that were used.

Table 2 XIV configurations that were used during benchmarks

Model	Configuration
XIV Gen3 model 114	► Number of modules: 15 ► RAM: 24 GB x 15 ► 180 x 3 TB 7.2K RPM disks ► Total FC Ports: 12 per node ► SSD: 15 x 512 GB
XIV Gen3 model 214	► Number of modules: 15 ► RAM: 48 GB x 15 ► 180 x 4 TB 7.2K RPM disks ► Total FC Ports: 12 per node ► SSD: 15 x 800 GB

Table 3 lists the tests that were run.

Table 3 Test descriptions

Test	Description
CSOP Benchmark	► Data Set: 2.4 TB (24 volumes x 100 GB) ► Compression Rate: 65%
TPC-C Benchmark	► Quest benchmark factory (v 6.6.1) ► Oracle 11.2.03 ► Scale 100 ► 1.2 TB Oracle database ► Oracle memory (SGA) 10 GB

Oracle online transaction benchmark

This benchmark uses Oracle to simulate an order-entry application by running a mixture of read-only and update-intensive transactions typical of online transaction processing (OLTP) environments. The benchmark incorporated five transaction types (for example, New Order, Delivery, and Payment). Throughput is measured in transactions per minute. The benchmark also reports the response time per transaction, which is broken out by transaction type.

The benchmark ran these transactions against the database:

- STOCK LEVEL: Checking the stock level
- DELIVERY: Processing a batch of 10 orders
- ORDER STATUS: Monitoring the status of orders
- PAYMENT: Processing a payment
- NEW ORDER: Entering a complete order

The benchmark measures transactions per minute (tpmC), which indicates new order transactions run per minute and provides a measure of business throughput. The benchmark also measures response time, which is the average time that a user got a response for each transaction.

Table 4 shows the results of the Oracle benchmarks.

Table 4 XIV using Oracle

Transaction	XIV Gen3 model 114, 30,000 users	XIV Gen3 model 214, 47,000 users
Stock Level	9.188	6.601
Delivery	6.022	3.575
Order status	1.086	0.79
Payment	1.135	0.83
New order	3.293	1.69
Average response time	2.62174	1.55604
tpmC (Throughput)	35594.295	45875.058
IOPS	84,413	135,152

Note: *Average Response Time in Seconds* is not the statistical/mathematical average of the response time, but rather an average time that the benchmark delivers based on the weight of transactions and their rate per minute. It is also worth mentioning that the metric "Seconds" is an application benchmark response time that is composed of hundreds of I/Os.

Synthetic workloads

To demonstrate the variability between workloads that demonstrate zero and 100% temporal locality (that is, data that is written at approximately the same time and might be part of the same compressed stream), you must modify traditional block benchmark tools to enable repeatable random workloads. This process results in best case and worst case raw block performance workloads when using compressed volumes. These tests were run with a known compressibility of data blocks by using a 1 MB pattern file that has a 65% and 80% compression ratio. These ratios are average for typical compression rates as seen in the field across IBM storage systems using IBM Real-time Compression.

Table 5 shows the performance results that can be used as a generic framework for realistic upper-limit expected performance for several models of XIV Gen3 using compression. The workload is ideal DB type with 65% compression, random I/O, 8 K I/Os, and 80% read.

Table 5 Synthetic workloads on the XIV Gen3

	XIV Gen3 model 114		XIV Gen3 model 214	
Operation	**Worst case**	**Best case**	**Worst case**	**Best case**
Read 4 KB Random IOPS	47K	305K	58K	368K
70/30 4 KB Random IOPS	64K	306K	93K	326K
Write 4 KB Random IOPS	64K	211K	85K	263K

User interface

Figure 2 shows the integration of compression from the IBM XIV GUI. The “HDM-CADCAM” volume is online and converting from uncompressed to compressed. The “RDM-IOTMeter” volume is not yet compressed. The comprestimator tool estimates the amount of space that can be saved.

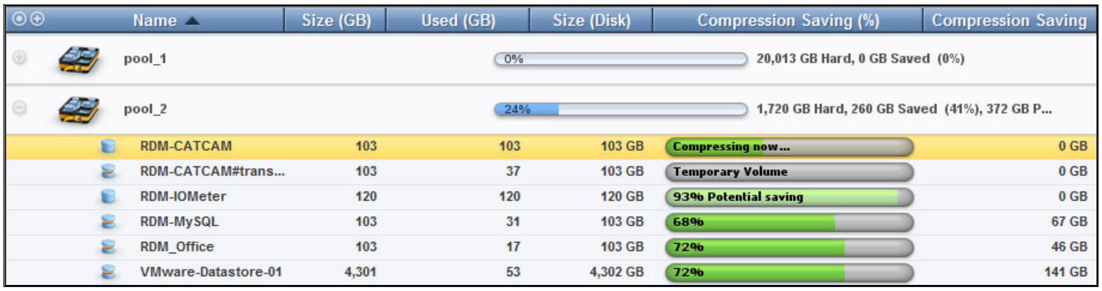

Figure 2 Compression integration with the IBM XIV GUI

Figure 3 shows where to enable the compression capabilities in the IBM XIV.

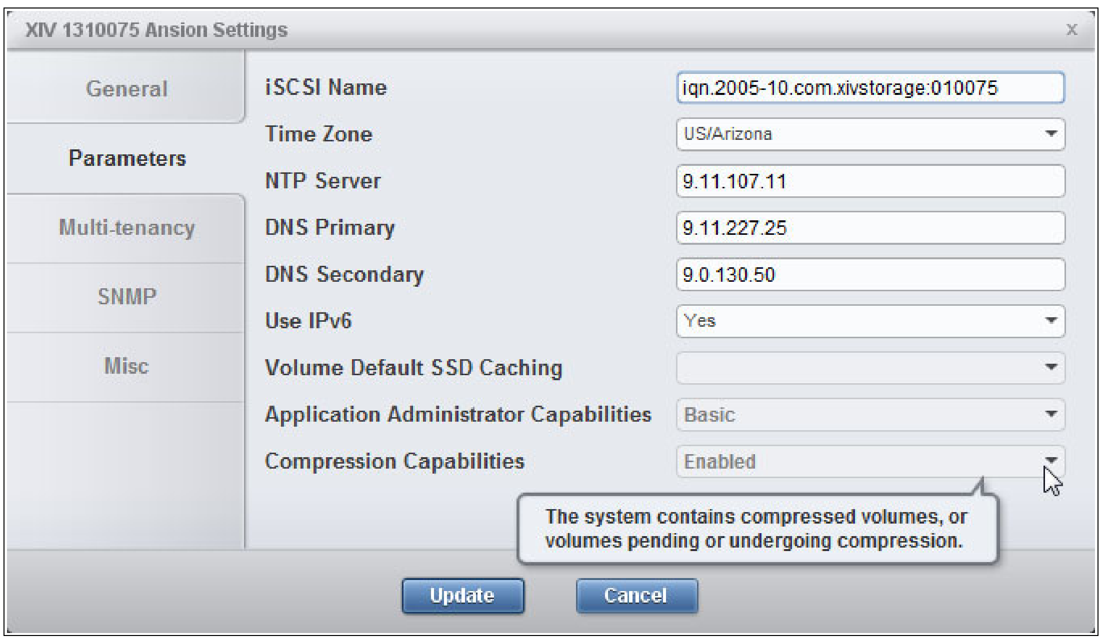

Figure 3 Enabling compression

Figure 4 shows the comprestimator output for volumes that are not yet compressed. You can also sort with various options. This example shows a sort by domain.

```
XIV 1310075 Ansion>>vol_comprestimate_list domain=itso_domain
Name          Compression Ratio (%)  Last estimation time  Status  Position
vol_145_1     insufficient data      N/A                   Idle    2
RDM-IOMeter   93                     2015-04-29 12:13:35   Idle    4
RDM-CATCAM    11                     2015-04-29 07:20:49   Idle    1
test_b        insufficient data      N/A                   Idle    3
XIV 1310075 Ansion>>
```

Figure 4 XCLI output from the comprestimator tool

Figure 5 shows the details for a compressed volume.

Figure 5 Volume properties

Requirements

Some requirements need to be considered when IBM Real-time Compression is used with the XIV Storage System.

Software levels

IBM Real-time Compression requires software version 11.6 or later on the XIV Storage System.

Volumes and pools

Compressed volumes must be thin provisioned volumes that are part of a thin provisioned pool. The process to compress and decompress a volume requires free space in the pool.

A compressed volume must be at least 103 GB. XIV supports up to a maximum of 1024 compressed volumes (including snapshots).

Mirroring

Compressed volumes are eligible for synchronous and asynchronous mirroring, and also for three-way mirroring. All included XIV storage systems must run software level 11.6 or later and have a compression license.

IBM XIV Real-time Compression trial

IBM is offering current IBM Storage XIV Gen3 customers the opportunity to evaluate the new Real-time Compression capability at no cost for 45 days. Customers who want to take advantage of this opportunity to evaluate Real-time Compression in their own environment with actual applications must enable the function, and run the evaluation guide. The 45-day evaluation period begins when the XIV Real-time Compression function is enabled. At the end of the evaluation period, customers must either purchase the required licenses for Real-time Compression or disable the function. Customers taking advantage of this opportunity must work with their local IBM Business Partner or their IBM sales representative through the evaluation process to apply the correct IBM XIV Storage software version, resolve any issues and questions related to the evaluation, and conclude the evaluation period. If the customers decide to use IBM Storage XIV Gen3 Real-time Compression after the evaluation period is over, it is their responsibility to ensure that they are properly licensed.

Summary

Using IBM Real-time Compression not only saves an average of 50% in storage space, potentially reducing capital expense, but can also decrease effective RPO because less data must be transferred across the replication mechanism. You might also see increased performance because compression is done in an above cache architecture, making more cache available for serving I/O. You might also save money for operational expenses because you can store much more data on an XIV system. External products can also benefit. For example, fewer backup licenses might be required and backups can run faster. With the integrated comprestimator tool, you can estimate your data reduction without knowing what type of data is stored on the XIV system.

Authors

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

Guenter Rebmann is a certified XIV Product Field Engineer in Germany. He joined IBM in 1983 as a Customer Engineer for large-system clients. After 10 years of experience with all large system products, he joined the DASD-EPSG (EMEA Product Support Group) in Mainz. In 2009, he became a member of the XIV PFE EMEA-Team. He has more than 30 years of experience providing technical support for past and present high-end DASD products.

Albert Verhoeff is a Client Technical Specialist working in pre-sales for IBM storage solutions. He has a deep technical knowledge of storage solutions, and supports Belgium, Luxembourg, and the Netherlands. He has 12 years of experience in IT and has held roles in first-level support for IBM Tivoli® Storage Manager, Tivoli Productivity Center, and IBM storage software. He has delivered enhanced technical support for enterprise accounts and worked as a product field engineer for EMEA and CEMEA on IBM N series (NetApp). His areas of expertise include not only storage but also hypervisors (VMware and Hyper-V), cloud orchestrators (OpenStack), Microsoft solutions, and SAN technology. As part of the IBM NAS Tiger program, he has deep knowledge of NAS storage solutions.

Megan Gilge is a Project Leader at the IBM International Technical Support Organization. Before joining the ITSO, she was an Information Developer in the IBM Semiconductor Solutions and User Technologies areas.

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

Bert Dufrasne
International Technical Support Organization

Paul Braren
Nir Rigai
Tzahi Shahak
Steve Solewin
Bosmat Tuv-El
IBM Systems

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

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 IBM 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

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.

© Copyright International Business Machines Corporation 2015. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by
GSA ADP Schedule Contract with IBM Corp.

This document REDP-5215-00 was created or updated on May 8, 2015.

Send us your comments in one of the following ways:
- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email 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 U.S.A.

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:

DB2®	Real-time Compression™	Storwize®
FlashSystem™	Redbooks®	Tivoli®
IBM FlashSystem™	Redpaper™	XIV®
IBM®	Redbooks (logo) ®	

The following terms are trademarks of other companies:

Microsoft, 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.