# Always On: Business Considerations for Continuous Availability

- ■ **Trends, observations, and perspectives**

- ■ **Business incentives and inhibitors**

- ■ **Challenges and recommendations**

# Executive overview

In 2013, IBM® conducted an IBM Academy of Technology initiative in which we reviewed the end-to-end aspects of *continuous availability* (governance, technological advancements, architectural patterns, design and operational practices). We identified specific business requirements that mandate continuous availability and also found organizations that have continuous availability as part of their DNA.

However, most organizations have invested in high availability or disaster recovery but not in continuous availability. For these organizations, continuous availability might not be a requirement today, but as we increasingly depend on IT, it becomes a critical success factor.

Managers of organizations might not know the precise exposure per hour of an outage, but they have a good idea of what the business would lose for each hour of an outage. They can also quantify the cost of continuous availability as a resiliency solution. By combining these two pieces of data, they can determine when continuous availability makes sense – when the value of continuous availability is greater than the cost.

But we found that this quantification might not be enough for an organization to embark on the continuous availability journey. A different approach is needed, and this approach needs to be top-down, driven by the business leaders.

This IBM Redguides™ publication describes the findings and recommendations from that IBM Academy of Technology project. First, we define continuous availability from the business perspective and discuss new business influences, such as social networks and mobile devices, which emphasize the demand for continuous availability. Then, we describe specific continuous availability incentives and benefits, along with the inhibitors, such as cost and complexity. With that as context, we provide recommendations to address these findings and discuss communicating with C-level executives and quantifying cost.

**1**

# What continuous availability means to your business

High availability (HA) and disaster recovery (DR) are relatively mature approaches that are typically well-known even by non-technical people. We found that continuous availability is not as mature and can be confused with HA or DR. People still think in terms of "how many 9s" they can achieve (99.99% uptime, for example). But that is an HA topic. This is a challenge for continuous availability. We need common terminology to be able to have a conversation with the different stakeholders, and we need to start by defining what *continuous availability* means.

The following three principles summarize what continuous availability requires:

► The ability to withstand a single component or multiple component failures

Components will fail. These components can be anywhere in the solution stack. Examples include physical facility failures (power), networking components, hardware, operating systems, subsystems, and application components. A continuously available solution must be able to withstand failure of single component and, potentially, multiple components.

► The ability to continue to service customer requests if there is a catastrophic disaster

Unplanned catastrophic events include weather (hurricanes, tornados, floods), earthquakes, fires, terrorist attacks, and so on. These events can affect a specific facility or a metro area or region. Therefore, the infrastructure that supports this requirement is typically in different geographic regions. This requirement falls under *business continuity*.

► The ability to introduce scheduled change without disruptions

Scheduled changes to any part of the solution must not affect availability. Any change should be able to be made without disruption from the end user perspective. The ability to do this involves technology, such as clustering, and operational models for rolling changes over infrastructure components.

Although this summarizes continuous availability for a technical audience, we believe that the following definition makes more sense for non-technical people:

"Continuous availability infers that services that consumers want to use are available whenever they want to use them, independent of the time of day or day of the week."

In other words, continuous availability means "always on." From the users' perspectives, the application (such as mobile banking) works for them, always. The business owners of business applications are not interested in conversations about disaster recovery or planned outages. Instead, they are interested in discussions about availability: How can this business application be always available to our customers?

From an end user perspective, it is about perception. Outages can happen, but if they are not visible to the end user, as far as they are concerned, the business application is available. For example, it might be acceptable if the balance in your checking account happens to be data[1] that is five minutes old. Or, there can be an ATM network problem but the alternative ATM switch automatically allows transactions under a certain amount. However, if you cannot access a business application because the cell phone or WiFi connection doesn't work (which is outside of the control of the application provider), from your standpoint as a user, the application is not available.

---

[1] However. there are many other use cases where up-to-the-minute data availability is a requirement Shares trading is one example.

# Unique continuous availability requirements

Businesses in certain industries have regulations or requirements that mandate continuous availability.

## Loss of life

There are situations where lack of continuous availability results in significant injuries or loss of life. This is true in healthcare for life support equipment and where lack of access to medical records, possible allergies, or drug interactions can have a catastrophic impact. This is true in many other sectors also. Think air traffic control, aircraft navigation systems, or utilities monitoring systems.

## Exceptional client experiences

During the Academy of Technology initiative, we realized that companies with people who take continuous availability seriously are companies where people care deeply about their customers' experiences. For these companies, client experience is not an afterthought. These companies have transformed their businesses to serve their clients better. Moreover, it is easy for any employee to relate continuous availability to the company's mission statement.

We believe that a large percentage of organizations is embarking on such client experience initiatives, often driven by the chief marketing officer, and will continue to do so in the next five years. This creates a high demand for continuous availability solutions.

## Industry-specific regulations

New regulatory, scheduling, and business requirements are increasing the need for continuous availability. For example, regulatory requirements in the UK now include a sign-off from a senior bank manager in each bank to personally guarantee that disaster recovery plans exist.

Financial institutions need stringent high availability on certain days or at certain times of the day, especially for scheduled time windows where money needs to be transferred between banks. A missed window can adversely affect other businesses and result in loss of interest income or an increase in interest paid. For example, Bundesbank requires financial institutions to report at a specified time each quarter. Similarly, service level agreements (SLAs) exist with companies such as State Street and its clients to ensure posting at specific times. If that does not happen, penalties go into effect.

# Continuous availability business incentives

Continuous availability applies much more broadly than the scenarios that we've already described.

## Raised expectations in a mobile and hyper-connected world

Mobile devices are predominant across most of the globe. People constantly use their smartphones or tablets at home, at work, while commuting, or at Sunday's game. People have the device at hand, and they expect the applications that they use with that device to be available whenever they want to access them.

For example, people expect the mobile banking application to be available when they want to initiate a wire transfer at 2 AM on Saturday, even though 24x7 service might not have been part of the mobile banking application's SLA and the bank might have advertised specific planned downtimes. Also, the bank does not control the mobile network at all. Regardless, when customers cannot initiate the wire transfer, they become dissatisfied with their banks. In today's mobile and hyper-connected world, people expect *always on*.

## No time for visible planned outages in a flat world

There is an increasing demand to provide continuous service and to avoid unplanned outages. Also, the tolerance for visible planned outages is decreasing to the point where many businesses no longer have any planned maintenance windows that interrupt normal service. This is true for business-to-business, business-to-consumer, and even business-to-employee services.

For example, how do you feel if you work in Australia and cannot submit your expense report on Monday morning? As more companies establish a global presence across multiple time zones, they no longer have "slow times." The business is active somewhere in the world throughout every 24-hour period. This shrinking or elimination of the maintenance window puts pressure on people, processes, and technologies to support concurrent upgrades and concurrent maintenance.

## Reputation damage in a social network world

People expect *always on*, regardless of published SLAs. When customers try to use a business application, and the business application is not available, these clients become dissatisfied. When they are dissatisfied, they tell their friends on social networks. We have seen this happens with an airline reservation system being down or with major bank outages. We call this "reputation damage." During our work, we found this to be an important reason for companies to adopt continuous availability. This was true to some level a few years ago, but it has accelerated in the last few years and the impact will continue to increase.

## Increased availability demands from real-time analytics

Big data analytics is another huge driver in today's enterprise initiatives. We believe that the main driver for continuous availability is not so much the increased volume of data, because most of the data still does not need to be continuously available. Is the velocity of the data. The data flows so fast that even a short outage has significant effects. More importantly, initiatives such as the next best action, where the analytics become part of the business transaction. This, too, create extra demands for continuous availability.

For more information, see *Smarter Analytics: Driving Customer Interactions with the IBM Next Best Action Solution*:

http://www.redbooks.ibm.com/abstracts/redp4888.html?Open

Traditionally, data analysis has been performed in the back office and updated in batches for the online systems or end users. A large portion of analytics will continue to be in batch mode, which typically does not require continuous availability. However, there is a new wave of analytics that is becoming much closer to the front office and more in real time. For example, the Propensity to Purchase analysis model is part of a business transaction and needs to be rescored in 250 ms or less during a customer's web session.

This shift of making analytics operational or of adding analytics to operational systems creates new continuous availability requirements.

## Mismatched expectations

As we just saw, today's growth initiatives, mainly cloud serves, big data analytics, mobile devices, and social networks, are huge drivers for continuous availability solutions. Even though consumers and business application owners expect the applications to be continuously available, the IT infrastructure owners do not expect to provide a continuously available IT infrastructure. The application developers do not code for availability, either. We found that, most of the time, there are mismatched expectations, and continuous availability is not being discussed.

# Continuous availability inhibitors

In addition to the expectation challenge, other things also make continuous availability difficult.

## Quantifying the cost of services not being continuously available

Even as the demand for continuous availability increases, it is still difficult for our clients to justify the expense of implementing continuously available systems and procedures. General figures can be quoted, but applying this to individual business applications is difficult, still. Sources of outage costs include loss of revenue or sales through e-commerce, regulatory fines for downtime or deadlines missed, customer compensation, reputation damage that leads to lost sales, lost productivity for employees who need to use systems, or lost productivity for employees who manage the data center.

The cost of an unplanned and even planned outage is often cited only during an unplanned outage or after a recent outage. Even if the cost per minute (or other metric) of an outage can be calculated accurately, it is difficult to predict how many outages will occur in a year or how long each might last.

## Quantifying the cost of making services continuously available

As it is typically approached by organizations, disaster recovery is a regulatory or other requirement (even though not always implemented correctly). The accepted cost of DR can be compared to the cost of high availability, where the hardware and software architecture is duplicated in a distant data center and DR procedure exercises are conducted occasionally. The main disadvantage of DR solutions is that cost is allocated to idle solutions.

Continuous availability is not a regulatory requirement, currently. However, we can envision a differently stated set of regulatory requirements that would require continuous availability. In a continuous availability environment, cost is allocated to an active solution. Regardless, decision-markers in organizations typically think that continuous availability is cost-prohibitive.

## Confusing continuous availability with high availability or disaster recovery

To justify the cost of continuous availability, a measurable impact on the business must be calculated. But often, as soon as the outage is past or if a data center has been fairly stable, the business side forgets the critical need for continuous availability until the next outage occurs. Many businesses seem to understand the impact of a disaster and have disaster recovery plans in place, but they cannot justify continuous availability despite the fact that an unplanned outage is much more likely to occur than an actual disaster.

An organization sometimes even considers the disaster recovery solution to be the solution to a smaller-scale outage. This is problematic because *disaster* has a specific meaning in this context, such as a natural disaster or civil disruption. When an outage occurs on one piece of equipment or an application, affecting one line of business, organizations are reluctant to "call a disaster" (because it is not a disaster) and do the all-or-nothing failover, instead. Understandably, this incurs too much risk for the business[2].

## Heterogeneous end-to-end view

Outages are not caused only by technology (failed hardware, software bugs, and so on). End-to-end continuous availability also includes business architecture, operations, process, and support organizations. Taking an end-to-end view of availability is challenging to organizations that operate with "silos," or segmented rather than interconnected workgroups, and possibly have conflicting goals.

In addition, some clients have multiple strategic outsourcing vendors (based on data center locations, mainframe or distributed systems, for example), and they have deployed software and hardware products from several vendors across the stack: storage, network, database, middleware, and so on. These heterogeneous environments result in multiple integration points and complex interactions. Determining what portions of the environment require continuous availability and coordinating continuous availability across multiple environments is a challenge both technically and organizationally. For our clients, it is not enough that IBM tackles continuous availability. They also need their other vendors to tackle continuous availability, which creates an impetus and need for industry standards.

## Poor communication between business and IT

Even when continuous availability for a particular business is a requirement, the recognition of that need is frequently triggered by an unplanned outage. Often, the business managers seem surprised that the outage and resulting loss of service occurred.

There can be a disconnect between business and IT, where business promises continuous availability to its clients even without mentioning "continuous availability." For example, a bank recently launched a marketing campaign that promised mobile banking services anywhere and anytime. Does the marketing department recognize the cost or complexity of keeping this promise? Is the continuous availability requirement communicated to IT and backed up by IT?

---

[2] This confusion denotes an organization's low level of maturity or interest in continuous availability.

There is a need to improve communication between the business organization and the IT organization and to employ terminology and metrics that are meaningful to both and provide a way to communicate continuous availability goals. There is also a need for active and visible leadership from the chief information officer (CIO), chief technology officer (CTO), and other C-level executives and at the board level.

## Ownership and funding

Along with this communication, there is a need to determine who "owns" continuous availability and who pays for it. Which organizations specify the requirements and which organizations are responsible for providing resources and doing the necessary work to meet those requirements?

Crucial, ongoing conversations must occur at the highest levels of organizations to understand and appreciate the necessity for continuous availability, the ability to fund its implementation, and the reality of availability "guarantees." This is not purely a business or IT issue; it is an enterprise issue. The continued emphasis on cutting costs, especially across areas of information technology, will not fulfill the objective of continuous availability, nor will the business areas that rely on IT ever realize increased availability requirements, because they are being driven by their end users.

# Addressing the cost challenge

In this paper, we described what continuous availability means for a non-technical person. We then described unique continuous availability requirements and strategic initiatives that mandate continuous availability without explicitly calling it continuous availability. We then described the things that make continuous availability difficult to implement. In this section, we provide recommendations and considerations for how to address the challenge of cost.

Overwhelmingly, cost has been one of the main reasons that customers have not pursued continuous availability for mission-critical applications. To aid in addressing the disparities that exist between the business areas and the IT organization that develops the solutions to support them, we offer the following recommendations.

Our top recommendation is to guide the conversation away from ROI and toward competitive advantage.

## Being realistic about the cost of continuous availability

The cost of implementing continuous availability initiatives can create a high hurdle (especially the geographically dispersed Active-Active approach, sometimes called Active-Active-Active or 3-Active). At the same time, we observed organizations with unrealistic expectations about implementing continuous availability "on the cheap," given the low price of commodity hardware. Simply throwing hardware at the problem and building a custom solution is a recipe for high maintenance and unattainable SLAs.

The reality is that continuous availability is much more than just hardware and software. The facility costs and resource requirements greatly exceed the initial implementation costs. As a result, continuous availability needs to be viewed in terms of total cost of ownership (TCO), and it must be cost-justified based on the inherent risks of *not* providing continuous availability.

Although many believe that cost of the 3-Active approach is exorbitant, that is an invalid analysis for any organization that manages an unused DR site. Many organizations that implement DR without using this capacity can optimize their approach in terms of both cost *and* resiliency by adopting a 3-Active architecture to optimize continuous availability, with a focus on balancing absolute data consistency requirements with eventual data consistency patterns.

## Case study: ibm.com

As an example of the cost aspects of eventual data consistency, with our flagship website, ibm.com, we provide a continuous availability solution at less than two times the cost of an equivalent HA or DR solution. A requirement for complete Atomicity, Consistency, Isolation, Durability (ACID) data consistency can increase the cost by up to eight times more, depending on business risk identification and the required HA patterns within each of the three data centers.

However, it is important to consider that it is not an all-or-nothing effort. Middleware clusters are independent of the ACID requirements and can be less expensive with a 3-Active environment than with a 2-Active environment. The ibm.com case study shows that 3-Active requires only 50% of the compute, memory, and network capacity per cluster, but 2-Active requires 100% per cluster.

The data layer is where the cost difference stands out. Eventual data consistency supports the 3-Active concept of 50% compute, memory, and network capacity for each cloud, yet the storage requirement remains at 100% per cloud. ACID data consistency supports only a 2-Active within a metro area. Out of Region fast failover requires 100 - 200% capacity for each cloud at the data layer and 50 - 100% capacity for each application cluster. Therefore, although cost differs by each organization's requirements, this underscores the need to truly understand the client environment and consistency requirements, given the increasing need to provide continuous availability.

## Assessing the business impact of unavailability

Various methods are used to quantify the cost of an outage (lost sales, fines from a regulator, or others). However, quantifying the cost of not providing continuous availability has not been a strong enough incentive for organizations to adopt continuous availability.

We recommend that organizations compare their business strategies and business plans for the next 12 - 18 months with their IT strategies and plans. We believe that the business strategy and business plans will mandate specific continuous availability capabilities, especially for organizations that are compelled by their marketing, sales, and support organizations to emphasize the quality of their customers' experiences as a primary benefit.

The following white paper provides insight about the financial consequences of IT failure. This paper is a good starting point to build a fact-based business case for continuous availability:

*Understanding the economics of IT risk and reputation*

http://www.ibm.com/services/us/gbs/bus/html/risk_study.html

## A continuous availability conversation with the CFO

Arguably, chief financial officers have become one of the most critical contacts within a business. Their roles and responsibilities have changed greatly since the economic downturn of 2008. CFOs now have a broader scope of responsibilities: cash flow, company liabilities, overall business performance, financial relations and obligations, and much more.

Crucial conversations must be initiated at the CFO level to appeal to the business aspects of supporting or not supporting continuous availability for those applications that warrant the highest level of availability. With the responsibilities held by today's CFOs, it is advisable to make an appeal that is based on the business consequences of not investing in continuous availability. Develop the business case that explains business impacts if the most critical solutions are not available within the enterprise. What are the cash flow implications if the finance and accounting applications are down for two days due to data corruption or hardware failure? What liabilities exist if the business is unable to fulfill contractual commitments because applications are offline for hours or days? What financial obligations exist or what penalties can be imposed if applications that run day-to-day operations are affected for several days? This is the first part of the approach: Documenting the benefits of continuous availability.

The second part of the approach is to measure the cost of continuous availability and compare it to the expected benefits. For example, regulations such as Basel II in the finance sector are causing companies to set aside contingency funds, based on their risk rating, which includes the risk of an IT outage. Business cases are now being approved for improving availability, based on not having to keep as much money in a contingency fund.

## A continuous availability conversation with the CMO

In our experience with clients, we are seeing many initiatives being driven from the office of the chief marketing officer. The CMO has money.

We believe that the nature of the initiatives that are driven by the CMO in the next few years will require continuous availability. For example, CMOs will emphasize data analytics, client experience, and mobile services. Also, with today's marketing campaigns, prospects can respond to an advertisement anytime (YouTube advertisements or QR codes in ads on trains, for example). In that sense, CMOs already see the need for continuous availability, because the application needs to be available whenever the prospect clicks or scans. We recommend that the CMO become the business champion for continuous availability.

To increase administrators' appreciation and understanding of the IT implications of achieving goals that make continuous availability necessary, we recommend a tighter connection between the chief marketing officer and the chief information officer. Through tighter integration between these C-level roles, IT will be more involved in new marketing initiatives. That will provide opportunities for more dialog and lead to a better understanding of the need for continuous availability across the business applications that are designed, developed, and delivered.

To summarize: Cost is not the only way to talk about continuous availability. You need to talk about cost, but you also need to talk about new opportunities for revenue. It is also important to explain that the cost is lower than what people typically expect. We propose a positive approach, where organizations put their clients' experiences first. That is an approach where continuous availability becomes a competitive differentiator.

# Summary

In this IBM Redguides publication, we explored the requirements, constraints, and challenges that are associated with continuous availability and provided recommendations for how to meet these challenges.

We provided a business-level definition for continuous availability to aid in communication within your enterprise and with customers. We explained that continuous availability means "always on," and that is sometimes more about end user perception than actual availability.

We then described how social media aggravates reputation damage in the event of outages, mobile services raise expectations of always on, and real-time data analytics create availability demands that did not exist with traditional back-office analytics. After looking at the new requirements for continuous availability, we reviewed the challenges that are associated with it, such as quantifying the cost of outages and quantifying the cost of achieving continuous availability. Another challenge is related to the complexity and heterogeneity of the end-to-end environment, with multiple vendors and multiple technology platforms involved.

In the second part of this guide, we provided recommendations for how to address the challenges. It is possible to make services continuously available with less investment that some analysts advertise. We need to quantify cost as part of the recommendation, but a more positive approach is to shift the conversation from the cost of services not being continuously available (for example, fines and reputation damage) to a conversation about the your customers' or clients' experiences. In this conversation, the CMO is the main business sponsor.

Finally, remember that, for an organization, achieving continuous availability means that continuous availability is a priority for that organization. But it also needs to be a priority for their primary vendors, including strategic outsourcing vendors, software and hardware vendors, and systems Integrators.

## Author

**Bertrand Portier** is an IBM Executive IT Architect. He has written extensively on distributed computing, modeling, service orientation, and business process management. He holds a Master of Science degree in Computer Engineering. He lives in Austin, Texas.

## 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 by 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

**13**

This document, REDP-5090-00, was created or updated on July 14, 2014.

# 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:

Redbooks (logo) ®          Redbooks®
IBM®                       Redguides™

The following terms are trademarks of other companies:

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