



How to Avoid Becoming Front-Page News (For All the Wrong Reasons)

A Practical Guide to Business Continuity and Disaster Recovery Planning

CONTENTS

Introduction	3
A Temperature Gauge for Response Levels	4
The Five-Step Process for Business Continuity/Disaster Recovery	5
Step 1 – Outline the Business Continuity/Disaster Recovery Technology Infrastructure	5
Step 2 – Set Service Level Expectations	6
Step 3 – Define Practical Business Continuity/Disaster Recovery Policies	7
Step 4 – Develop a Contingency Plan	8
Step 5 – Test the Contingency Plan	9
Conclusion	10

Whether organizations strive to guard themselves against terrorist attacks and Katrina-scale disasters that impact an entire region, or seek protection against more localized events like an electricity cable hit by a worker’s backhoe, enterprises never know when their operations may be threatened.

The resulting consequences of poor disaster recovery planning involve a mix of hard and soft costs. Some organizations believe that even a few hours of downtime may doom their business to extinction.

Fortunately, industry best practices have evolved to address the technical and cultural factors that create a viable disaster recovery framework. Today’s disaster recovery responses balance the duration of disruptions with what it costs to guard against them thanks to plans that create “hot,” “warm” and “cold” recovery strategies.

The key technology elements of a business continuity/ disaster recovery infrastructure consist of a main data center, a remote site that recreates the resources in that primary location, and network connections between the two that use high bandwidth broadband services. In addition, the best disaster recovery strategies follow a “redundant everything” philosophy that applies to servers, storage appliances, power supplies and other technology components.

Organizations also must create a comprehensive disaster recovery program. Help can come from industry groups like ITIL and the U.S. government’s National Institute of Standards and Technology.

Finally, disaster recovery experts say that testing and training are essential to achieving ongoing disaster recovery success.

INTRODUCTION

No one says protecting data is easy, even if an organization tries to do everything right. One company headquartered in Denver took almost all the right steps to protect its data center. It diagramed its enterprise IT resources on a large whiteboard. Then it established a comprehensive plan for safely storing real-time copies of essential data at an offsite data center in case an emergency shut down the core data center.

The business continuity/disaster recovery site was an accurate mirror of the production facility with all the necessary hardware, software and storage systems in place standing-by ready to come alive. The impressive results of emergency drills gave the company confidence. According to Mike Karp, senior analyst with technology researcher and consulting firm Enterprise Management Associates, the company believed that even if a natural or man-made emergency shut down the main processing center, the offsite facility was efficient enough that business users wouldn't notice a performance blip.

"When I asked them where this remote data center was located, one of the executives took me over to the window and said, 'That's it over there,'" Karp recalls. "I was looking diagonally across the street at the building, and I thought to myself, 'Denver is one of the centers of the oil and gas industry. All you need is for an oil truck to go through this intersection when some kid in a Camaro runs the light.' Not only do they blow up both buildings, but this business goes 'bye bye.'"

Unfortunately, the potential vulnerabilities this otherwise prudent company faces aren't unique. Guarding against every possible contingency may ultimately be impossible. Even protecting against *likely* emergencies can be complex enough that without the right plan, unexpected gaps can occur. Yet unpredictability is ever-present.

The resulting consequences of poor business continuity/disaster recovery planning involve a mix of hard and soft costs. For example, telecom companies live and die by maintaining nearly perfect uptime records. The demands are so great that one large telecom provider estimates each minute of unplanned downtime costs nearly \$1 million.

The financial costs may be even greater for financial services companies. Some organizations in the industry base their business continuity/disaster recovery strategies on the premise that three hours of downtime may doom their business as customers lose confidence and seek alternative organizations. U.S. Bureau of Labor research adds credence to assumptions like these.

In fact, U.S. Department of Labor research found that 93 percent of organizations that have experienced significant data loss are out of business within five years.

**A proper business continuity/
disaster recovery plan combines technical, management
and operational resources
to minimize risks.**

As bad as the financial setbacks may be, downtime can result in some problems so dire the costs can't even be quantified in dollars and cents. As hospitals rely more and more on digital clinical systems, including electronic patient records and digital radiology, protracted shutdowns can literally be a matter of life and death.

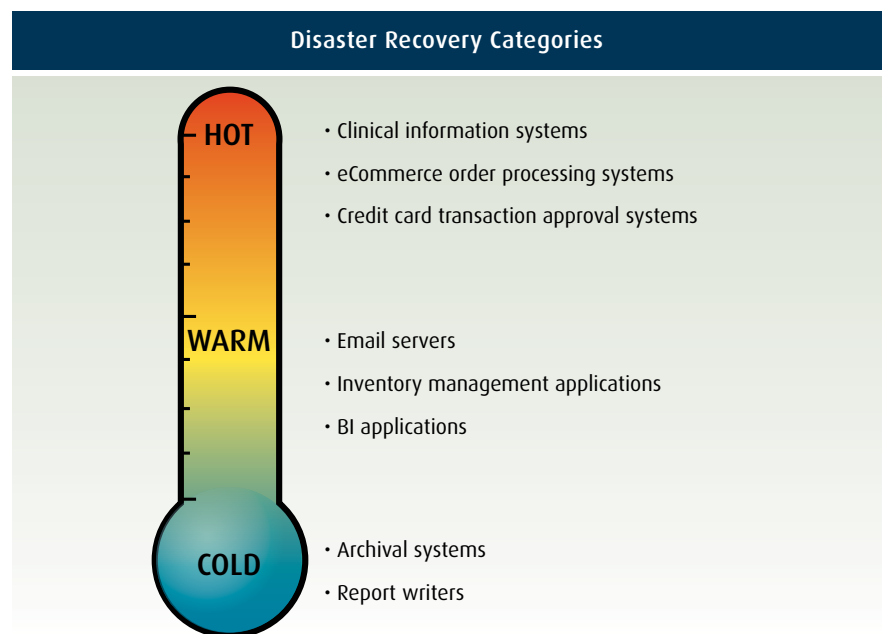
Realities like these are changing attitudes about disaster recovery. Concerns about the consequences of unplanned downtime are fueling the growth internationally of data-protection and recovery-management products and services, according to technology researcher IDC. The company forecasts the market will grow from nearly \$58 million in 2006 to \$200 million by 2011.

"People used to say that if their IT systems took a hit and came back up in two days, that was acceptable," recalls Becky Albin, chief IT architect at Software AG. "Today, more and more organizations are saying, 'we have to be up in either a matter of minutes or in a matter of hours.'"

Fortunately, the news isn't all gloom and doom. Organizations can minimize or eliminate many IT-system vulnerabilities with the proper business continuity/disaster recovery plans that combine technical, management and operational resources to minimize risks.

A Temperature Gauge for Response Levels

Not all IT disasters are caused by natural disasters. According to technology consultant Forrester Research, preventable incidents account for the majority of outages. IT systems are vulnerable to a variety of disruptions, ranging from mild (power brownouts or a single disk-drive failure) to severe (building fires and floods). Just as the degrees of "disasters" differ, so too do the levels of responses, which should balance the duration of disruptions with what it costs to guard against them. While the costs are high when business-critical systems go down even for seconds, some business applications don't require instantaneous emergency responses. Thus, industry best practices separate disaster recoveries into "hot," "warm" and "cold" categories (see graphic below).



Hot means if a primary system crashes, remote backup resources come up almost instantaneously thanks to exact data, applications and hardware “mirrors” that receive updates in real-time.

A “warm” recovery allows for a window of perhaps a couple of hours for activities to rebound. To accomplish this, organizations use duplicate configurations of hardware and system software, but the clone applications and databases aren’t actively running during normal periods. Instead, if a shutdown occurs, IT managers bring the standby system up to date. After anywhere from two to six hours, the system is ready for business people to use.

A “cold” backup allows for a couple days before becoming fully operational. In this scenario, organizations commit hardware to the recovery task, but the needed operating systems and applications aren’t running until they’re required to be made available to receive new data via FTP connections or tape cartridges.

THE FIVE-STEP PROCESS FOR BUSINESS CONTINUITY/DISASTER RECOVERY PLANNING

Creating a reliable approach for keeping the business running during an incident (large or small) is not a matter of hocus-pocus or guess-work. But neither is the task simple or easy. Since failure to put a comprehensive and reliable business continuity/disaster recovery process in place invites business catastrophe, here are five steps that can provide some solid direction.

Step 1 – Outline the Business Continuity/Disaster Recovery Technology Infrastructure

The key technology elements of a business continuity/disaster recovery infrastructure consist of a main data center, a remote site that duplicates the resources in that primary location and network connections using the highest bandwidth—typically “T3” business broadband services.

“A business continuity/disaster recovery site needs to be 100 to 150 miles away from your primary site, at a minimum. Beyond that, it could be literally anywhere in the world, depending on the networking capabilities between the two sites.”

– Becky Albin, Chief IT Architect, Software AG

Creating a reliable approach for keeping the business running during an incident... is not a matter of hocus-pocus or guess-work.

**The best business continuity/
disaster recovery strategies
follow a “redundant
everything” philosophy
throughout the data center.**

Just as a redundant data center doubles an organization’s chances of staying in operation, the best business continuity/disaster recovery strategies follow a “redundant everything” philosophy throughout the data center. Multiple mainframes and servers should run in the production and backup data facilities. That’s why if one piece of equipment encounters problems, the disaster recovery system fails over to the other resources using utilities in mainframe operating systems or clustering software in RISC and Intel/AMD blade servers. Similarly, multiple network-attached storage applications use failover techniques to guard against any one hard drive becoming a single point of failure.

Power supplies may be unglamorous compared to today’s elegantly engineered hardware and software, but they’re one of the most critical components in a business continuity/disaster recovery strategy. Power outages rank among the leaders in most common and preventable disruptions, according to industry analyses.

And no matter how fat the high-bandwidth networking pipeline may be, it’s of little use if a careless construction crew accidentally severs a fiber cable. Therefore, network connections must not only be redundant, they also need to follow different paths within a wider WAN topology to keep a single threat from bringing businesses to a standstill.

Step 2 – Set Service Level Expectations

Once organizations have sketched out the topology of their business continuity/disaster recovery infrastructure, the next activity is to develop an accurate inventory of their IT assets. This enables the organization to understand the resources and business processes that need to be protected. “Organizations need to look at their critical applications and analyze them to determine what is required to support them,” says Bruce Beaman, senior director of Adabas and Natural product marketing for Software AG.

A wide range of tools and resources is available to help organizations develop and maintain accurate inventories of IT resources. Vendors of enterprise-management tools offer modules that use software agents to scour the IT infrastructure. These investigative agents gather information for developing detailed summaries of organization-wide assets – ranging from mainframes, servers and storage devices to applications with their latest version numbers and peripherals, such as printers and multifunction products (MFPs). Similarly, storage resource management (SRM) software focuses on shared-storage devices to provide details about storage-area network (SAN) appliances.

Configuration management databases (CMDBs) store details about hardware and software assets running within organizations. CMDBs were originally a best practice introduced by the IT Infrastructure Library (ITIL), a comprehensive, vendor-neutral framework for managing IT services (<http://www.itil-itsm-world.com/index.htm>).

Organizations may use ITIL guidelines to develop their own CMDBs, or they can consider a growing number of commercial products that sell pre-built versions. Either way, a successful CMDB can give organizations an accurate view of hardware and software inventory, as well as service-level agreements created to define their uptime and recovery parameters.

Developing service-level agreements can be daunting for organizations. “IT managers need to talk to business users to see exactly what their expectations are if an outage occurs,” Beaman says. For example, the IT department may put a particular system in a recovery category of four hours or less, while the affected business managers assume a faster response will occur.

“When I’ve been involved at this stage at customer sites, I emphasize having an IT inventory when IT talks to business users, then putting in writing what the expectations are. Because if it’s not in writing, there could be finger pointing if a disruption does happen.”

– Becky Albin, Chief IT Architect, Software AG

Organizations that need help with developing service-level agreements for business continuity/disaster recovery plans should consider using software programs and best-practices documentation from ITIL.

Step 3 – Define Practical Business Continuity/Disaster Recovery Policies

Of course, data protection isn’t just a matter of plugging in the latest and greatest technologies and expecting everything to work perfectly during a crisis. Success requires a business continuity/disaster recovery framework. Fortunately, industry best practices offer help with sorting out the cultural factors that can either lead to the successful creation of a framework or torpedo the negotiation process.

Armed with 1) clear data about the configuration of the production data center, 2) a summary of critical IT and business processes and 3) service-level agreements, organizations next need to define the key elements of their business continuity/disaster recovery program. In addition to advice from industry groups like ITIL, the U.S. government’s National Institute of Standards and Technology (NIST) has compiled a comprehensive examination of the subject titled “Contingency Planning Guide for Information Technology Systems,” available at <http://csrc.nist.gov/publications/nistpubs/800-34/sp800-34.pdf>. NIST. This guide outlines common characteristics in approaches for developing and maintaining contingency plans.

First, organizations need to create a policy statement for contingency planning. The goal behind building a formal policy statement, according to NIST, is to ensure that everyone in the organization understands why a contingency plan is important in the first place and how each person contributes to its success. To do this, NIST recommends that the document clearly defines the overall IT contingency objectives and describes responsibilities throughout the organization.

Support from CIOs and other senior managers is key, say disaster recovery experts, not only to elicit buy-in from staff members, but also to validate the importance of making redundant technology a spending priority.

Support from CIOs and other senior managers is key, say disaster recovery experts, not only to elicit buy-in from staff members, but also to validate the importance of making redundant technology a spending priority.

NIST emphasizes that while disaster recovery is important, the best scenario is disaster avoidance.

Think about “who is going to pay for the disaster recovery capabilities,” advises Karp.

“Disaster recovery/business continuity projects must have buy-in from across IT and the multiple business units in order to work. That’s because this approach not only represents an outlay of money for hardware and software; it’s also about allocating people.”

– Mike Karp, Senior Analyst, Enterprise Management Associates

At this point the IT staff and business managers should discuss whether each individual system requires a hot, warm or cold disaster-recovery plan, he adds.

After creating a policy statement, organizations should perform a business impact analysis (BIA) to flesh out details about system requirements, processes and systems inter-relationships. These profiles will define the ultimate requirements and priorities of the overall plan. NIST says the BIA should also perform one other critical service—provide a description of the potential consequences of system disruptions.

From this baseline, organizations can then determine allowable lengths of time for shut-downs and what systems receive priority for recovery efforts, NIST adds. For background, the NIST contingency planning guide includes a sample BIA process.

NIST emphasizes that while disaster recovery is important, the best scenario is disaster avoidance. To accomplish preventative measures where possible, IT managers should shore up infrastructure components that support computers and software. This includes installing uninterruptible power supplies (UPSs) to keep systems running during brown-outs and power spikes (or to allow critical systems to switchover to backup resources should longer-term power outages occur). Also on the list of preventative and problem-mitigating tools are:

- Gasoline- or diesel-powered generators
- Air-conditioning systems with spare critical components, such as compressors
- Fire suppression systems
- Computer-room water sensors
- Tarps readily available to protect IT equipment from water damage

Step 4 – Develop a Contingency Plan

The IT contingency plan details the roles and responsibilities of departments and individuals in keeping technology systems available. In addition, it outlines the procedures required to restore IT systems during an emergency, according to NIST.

Senior IT officials should play direct roles in developing the contingency policy, its structure and objectives, as well as defining roles and responsibilities. Other areas to focus on during the policy-creation stage are resource requirements, training needs, the frequency of training exercises and testing, plan maintenance schedules and data-backup schedules.

The NIST contingency-plan approach consists of two components intended to encourage comprehensiveness and three additional modules with detailed recommendations for emergency responses.

To ensure completeness, organizations should provide any necessary support information to help people understand why contingency planning is important and what requirements are necessary for implementing and maintaining plans. Second, disaster recovery managers should make clear where people can find any additional information related to but not spelled out in the core contingency plan. This may include technical, operational and management specifications for running each system. Organizations should also compile a directory of contact information for key emergency-response staff members, as well as for vendors and service providers, NIST advises. Other important ancillary information ranges from checklists for system recovery or processes, inventories of hardware and software, notes on connecting to or physically traveling to an offsite backup facility, and a copy of the business impact analysis.

For effective responses to emergencies, IT managers need to be ready to launch a three-prong attack.

First, the notification/activation phase defines the initial steps to be taken when an emergency strikes. This phase spans notification of emergency personnel, system assessments and the roll out of the disaster recovery plan, NIST explains. The goal of these activities is to launch temporary contingency measures as quickly as possible.

The next stage begins the more complete recovery phase based on damage assessments and having emergency crews activated. Second-stage activities include making any necessary repairs to IT systems and getting business processes running at main production sites, if possible, or at dedicated disaster recovery facilities.

Finally, in NIST's third emergency-response phase, organizations either revert to normal operations or, when day-to-day resources have been damaged beyond repair, activities continue using the backup resources as the organization works to launch a new production facility.

Step 5 – Test the Contingency Plan

Disaster recovery experts say one of the most important yet frequently overlooked aspects of disaster recovery planning comes after the formal policies and procedures are delineated. Plans must be tested initially for their completeness and effectiveness, and then on an ongoing basis to make sure that any subsequent changes to the IT infrastructure and business processes haven't created a need for policy modifications.

"The advice here is don't just put a plan down on paper and then never test it again," Karp says. Training is closely related to testing. Unless everyone in the organization understands the plan and his or her role in implementing it, gaps can occur that under the stress of an emergency result in data losses and unnecessary downtime.

Plans must be tested initially for their completeness and effectiveness, and then on an ongoing basis.

...for disaster recovery strategies to be tested appropriately, the drills must simulate real-world conditions as closely as possible.

Regularly working through these two areas is “something that most companies still don’t do, and the issue of having untested plans is just a gaping wound,” Karp says.

Each IT contingency plan element should be tested to confirm the accuracy of individual recovery procedures and the overall effectiveness of the plan, NIST says. This covers both technical and personnel components, including system recovery on an alternate platform from backup media; coordination among recovery teams; internal and external connectivity; system performance using alternate equipment; restoration of normal operations; and notification procedures.

Karp also warns that for disaster recovery strategies to be tested appropriately, the drills must simulate real-world conditions as closely as possible. “Some companies test their plans under rarified atmospheres,” he says. “They ask everybody to please stay off the production system from 10 o’clock Friday evening until 3 a.m. Monday morning because ‘that’s when we are going to do disaster testing.’”

Instead, organizations should create test beds that accurately reflect day-to-day business conditions. So, for example, a company might replicate a database, with the terabytes of new information that’s been added in the last 12 hours, for example, and run it on the same hardware that’s used in the production system. “Then bang the hell out of the test system,” Karp says. “That’s how you get the peace of mind that comes from having tested it in a real-world environment.”

CONCLUSION

Life may be too complex for organizations to protect themselves against every disaster contingency. But with the right technologies, clear service-level expectations, practical disaster recovery policies, thorough contingency plans and rigorous testing methodologies, organizations can minimize the consequences of the most common disruptions. Diligence will also protect organizations against obvious mistakes – such as a back-up data center that’s vulnerable to a wayward Camaro – and keep them a safe distance from the news headlines as well.

MARKET TRENDS ON DISASTER RECOVERY

These figures from Forrester Research show the importance of disaster recovery. Figure 1 shows it is “very critical” to the majority of enterprises surveyed. Figure 2 illustrates the driving need that enterprises have to improve their disaster recovery plans.

Figure 1

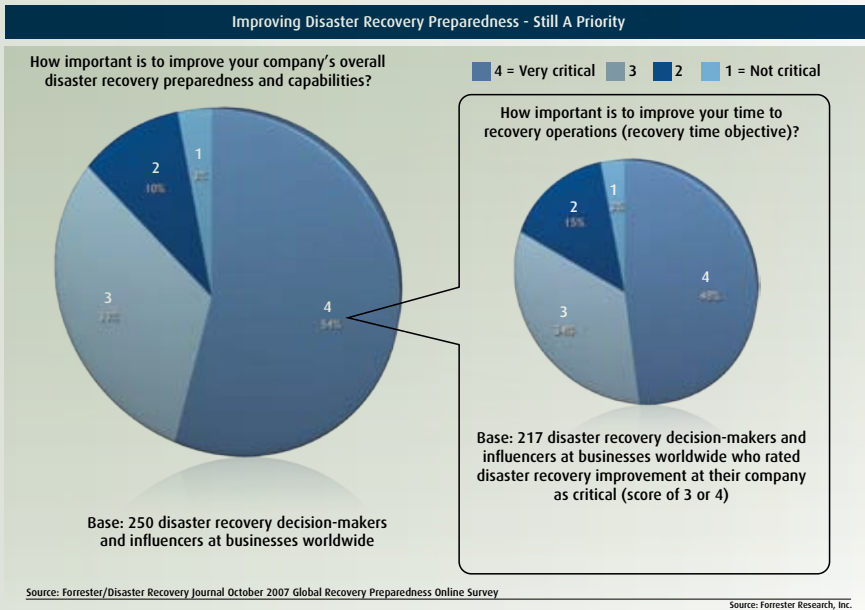
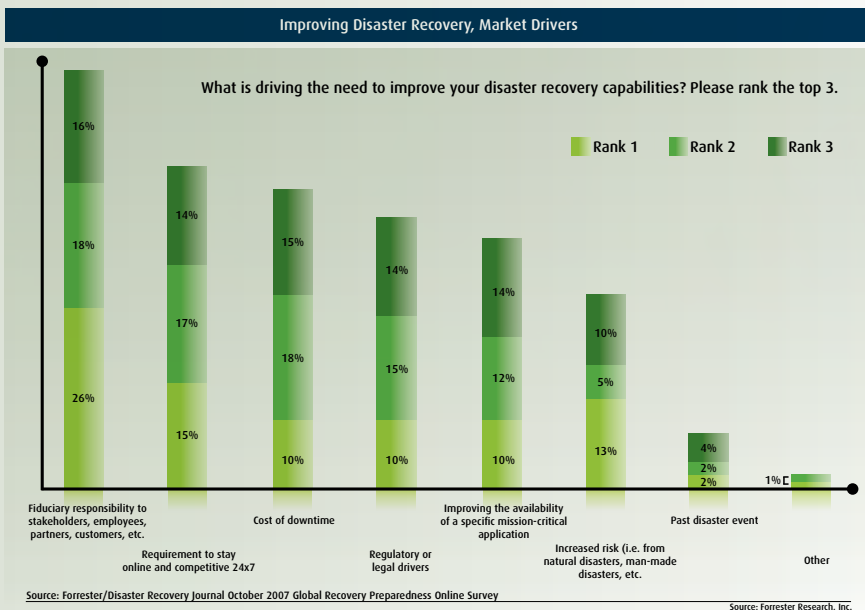


Figure 2



INTRODUCING A DISASTER RECOVERY SOLUTION FROM SOFTWARE AG

To make it possible for organizations to achieve real-time disaster recovery, Software AG offers Event Replicator for Adabas, which allows organizations to create an offsite backup data center almost anywhere in the world and use it as a "hot standby" facility. Event Replicator for Adabas is the only technology that provides for immediate Adabas-to-Adabas replication, according to Software AG.

"The difference between real-time replication and what other technologies offer can be between five to 10 minutes – which for a hot site just isn't sufficient," says Becky Albin, chief IT architect at Software AG.

Instead, Event Replicator for Adabas can be part of a hot, warm or cold business continuity/disaster recovery effort for organizations that use Adabas. The real-time performance of Event Replicator for Adabas is possible because the technology removes a number of steps that otherwise require manual intervention in a disaster recovery process.

For example, other replication products for Adabas must first create a copy of the database's protection log. "This requires making a copy of the log, then going to a disaster recovery tool to pull updates from the protection log copy and apply them to the disaster recovery database before actually bringing it up to take over in a recovery situation," Albin says.

By contrast, because Event Replicator for Adabas is tightly integrated with Adabas, it can pull the updates directly from the production database's nucleus and apply them to the disaster recovery database. "By completely eliminating this manual step, time, labor and the chance for errors are all reduced," says Bruce Beaman, senior director of Adabas and Natural product marketing for Software AG.

In addition, Event Replicator for Adabas can ease most day-to-day database maintenance tasks. So when an organization adds tables, files, fields or columns to the production Adabas database, the updates replicate in real-time to remote sites. "You don't have to do maintenance in both locations," Beaman says.

Business continuity requires a comprehensive approach to protecting data centers that encompasses database replication as well as redundant hardware, applications, security technologies and other technologies essential to running the IT environment. "Event Replicator for Adabas is one important piece in that 'big picture,'" Albin says.

**TO FIND THE SOFTWARE AG OFFICE
NEAREST YOU, PLEASE VISIT
WWW.SOFTWAREAG.COM**

Take the next step to get there – faster.

ABOUT SOFTWARE AG

Software AG is the world's largest independent provider of Business Infrastructure Software. Our 4,000 global customers achieve measurable business results by modernizing and automating their IT systems and rapidly building new systems and processes to meet growing business demands.

Our industry-leading product portfolio includes best-in-class solutions for managing data, enabling service oriented architecture, and improving business processes. By combining proven technology with industry expertise and best practices, our customers improve and differentiate their businesses – faster.

Software AG – Get There Faster

© 2008 Software AG. All rights reserved. Software AG and all Software AG products are either trademarks or registered trademarks of Software AG. Other product and company names mentioned herein may be the trademarks of their respective owners.