

NINE STEPS TO BUILDING A BUSINESS-ORIENTED DISASTER RECOVERY PLAN



Quest®



Introduction

In our e-book, “Downtime Costs How Much?,” we talked about why disaster recovery (DR) is absolutely essential to business, despite not increasing top-line revenue or reducing cost. We talked about how to calculate the value of DR and how it takes both technical know-how and business-like thinking to illustrate the importance of DR to the stakeholders who allocate the budgets. In this e-book, we start by highlighting the top business-threatening data disasters, and then offer nine steps you can implement to build a business-oriented DR plan.

Getting on the same page

To begin, everyone needs to get on the same page with regard to definitions and key assumptions. You have the opportunity to help others understand the nature and probability of different disasters and what constitutes recovery. The meaning of disaster and recovery can vary — not only across the industry but also within an organization.

Business managers, application owners and IT admins likely have different ideas of what is considered an acceptable level of downtime and data loss. They probably don't agree on what systems are critical to the survival of the company.

Most people (outside of IT) consider a “disaster” to be a calamitous event — one that occurs suddenly and causes great loss of life, damage or hardship. But within the world of IT, a disaster is any unexpected event that causes a substantial loss of service levels in critical business systems for an unacceptable period of time.

It's likely that the next data-center disaster will be caused by something more mundane, like a power failure.

A disaster could be caused by something calamitous, yes, but statistically speaking, it's more likely that the next data-center disaster will be caused by something more mundane, such as a power failure. It's important to make everyone in the organization aware of this fact.



The top threats to business continuity aren't what you think

HUMAN ERROR

Human errors often cause systems to become logically corrupt or unusable. An accident as simple as an employee tripping on a cord can bring down an entire storage system.

While accidental human actions are common problems, intentional actions are increasingly common.

MALICIOUS ATTACK

While accidental human actions are common problems, intentional actions are increasingly common. Today, of course, organizations are even more aware of the possibility of malicious acts causing disasters. For example, disgruntled or former employees can attack and bring down IT systems, and so can viruses. An even more pressing concern of late is cyberterrorism, especially threats against critical industries or government offices from groups or countries opposed to their actions or policies.

DATA CORRUPTION

A data corruption outage occurs when a corrupt hardware or software component causes corrupt data to be read or written to the database. Data corruption takes many forms. It can be widespread or it can be localized. The impact of a data corruption outage will vary accordingly.

Corruption in a single database block might affect few users, while corruption in a large portion of the database would make it essentially unusable. Most IT professionals have seen some form of data corruption in their careers, although organizations understandably tend not to publicize these problems. Such data corruption can be caused by hardware failures or human error.

STORAGE FAILURES

A storage failure outage occurs when the storage holding some or all of the database contents becomes unavailable because it has shut down or is no longer accessible. Many companies have had complete storage failures — often caused unintentionally by pesky humans. For example, at one organization, someone stacked a set of disk drives against a wall, inadvertently turning off a switch and causing system failures — an issue that was difficult to track down.

Another company that relied heavily on its storage area network (SAN) made the seemingly simple choice to lay carpet in its data center to reduce noise. When an authorized employee walked in to check the SAN and touched its racking bay, the static electricity discharge shorted the controller unit and the entire SAN went down. Without knowing that the cause of the problem was the electrostatic charge built up by walking on the carpet, the company put in a new controller. After it was up and running, someone else touched the rack again, and the new controller was also fried.



Any business-centric DR plan should also include a redundant local area network (LAN) infrastructure as well as steps on how to restore the LAN.

POWER AND NETWORK FAILURES

Power failures can seem mundane, yet they can have a crippling effect on business. In August 2016, Delta Airlines experienced an outage that affected business for three days, with 2,300 flight cancellations and \$150 million in losses. The outage started with a circuit breaker that needed to be reset in the company's Atlanta headquarters. The down circuit breaker caused a transformer to shut down, which in turn led to loss of power to Delta's data center. Systems were moved to backup power, but not all of the servers were connected to that source, which led to a number of further issues — including a huge number of customer complaints and damage to the company's overall reputation.

Servers and storage often get all the attention in DR planning, but any business-centric DR plan should also include a redundant local area network (LAN) infrastructure as well as steps on how to restore the LAN. Network failures are the third most common cause for unplanned downtime. The loss of a single network switch could quickly turn into a major, time-consuming outage for the organization.

NATURAL DISASTERS

Natural disasters such as Hurricane Harvey or the 8.2 magnitude earthquake in southern Mexico, both of which occurred in 2017, are the types of disasters that will jump to the mind of your business colleagues when you mention disaster recovery. These types of disasters, however, happen less frequently than IT-related failures.

IT disasters happen every day, with even the most experienced IT professionals on staff. Organizations must accept that disasters will happen, and build a sound recovery strategy to minimize their impact.

Nine steps to building a business-oriented DR plan

Although IT disasters are unpredictable, recovery shouldn't be. In fact, recovery should be planned, predictable and controlled. The following steps will help you organize your thoughts, ask the right questions and develop a strategy for your DR plan that is closely aligned with your business.

1. CONDUCT AN ASSET INVENTORY

Disaster recovery planning should always start with an inventory of all your IT assets. This step is necessary to untangle the complexity of your environment. Start by listing all the assets under IT management, including all servers, storage devices, applications, data, network switches, access points and network appliances. Then map where each asset is physically located, which network it is on, and identify any dependencies. See example on Page 7, Table 1.

2. PERFORM A RISK ASSESSMENT

Once you have mapped out all your IT assets, networks and their dependencies, list the potential internal and external threats to each of those assets. Imagine the worst-case scenario — and be thorough. Threats could include natural disasters or mundane IT failures.

Next, include the probability of each happening and the impact it could have. How would it affect the business if each scenario were to occur? Enlist the help of your business colleagues for this exercise — but be sure to emphasize that mundane events happen much more frequently than natural disasters. Move the conversation away from earthquakes and

hurricanes and toward higher probability events such as power outages or IT hardware failures. See example on Page 7, Table 2.

3. DEFINE CRITICALITY OF APPLICATIONS AND DATA

Before you begin to build out your business-oriented DR plan, you'll need to classify your data and applications according to how critical they are to the business. Start by speaking to your business colleagues and support staff.

Look for commonalities and group them according to how important each is to the business, frequency of change and retention policy. You do not want to apply a different DR technique to every application or dataset you have. Grouping assets with similar characteristics will allow you to implement a less complex strategy.

You'll need to classify your data and applications according to how critical they are to the business.

Classifying data in a vacuum based on assumptions could come back to haunt you. Be sure to involve other business managers and support staff in this exercise. You will undoubtedly have to make some trade-offs to limit the number of data classes you have. For medium-sized businesses, the number of classes should likely be between three and five. See example on Page 7, Table 3.

Location	Server/VM	OS and hypervisor	Application	IP address	Disk allocated	Disk used	Dependencies
SFO-1	Orcl-001	RHEL 5.x	Oracle 11g	10.10.10.1	5 TB	1 TB	
	Exch-001	Win 2012 r2	Exchange 2013 (DAG1)	10.10.10.2	20 TB	7 TB	Exch-002
	Exch-002	Win 2012 r2	Exchange 2013 (DAG2)	10.10.10.3	20 TB	7 TB	Exch-001
	MOSS-001	Win 2012 r2	SharePoint 2010	10.10.10.4	10 TB	8 TB	SQL-01, SQL-02
	SQL-001	Win 2012 r2	SQL Server 2008	10.10.10.5	5 TB	3 TB	
	SQL-002	Win 2012 r2	SQL Server 2008	10.10.10.6	5 TB	2 TB	
	SQL-003	Win 2012 r2	SQL Server 2008	10.10.10.7	5 TB	2 TB	
	SQL-004	Win 2012 r2	SQL Server 2008	10.10.10.8	5 TB	2 TB	
	AD-001	Win 2012 r2	AD Domain Controller	10.10.10.9	3 TB	1 TB	

Table 1: An example of an inventory of one company's IT assets

Location	Assets	Threat (internal and external)	Probability	Impact
SFO-01	Orcl-001, Exch-001, SQL-001, SQL-002, SQL-003, SQL-004, SQL-005, FLS-001	Natural disaster — earthquake	Low	High
		Network failure	Medium	Medium
		Power failure	High	High

Table 2: Risk assessment

Class	Description
Low impact	All data and systems that are needed to achieve the business' strategic objectives but that do not need to be immediately restored for the business to continue to operate.
Moderate impact	All data and systems that are important to achieving business objectives. The business can operate but in a diminished state.
High impact	All data and systems that are critical to business operations. Business comes to a halt without the associated services.

Table 3: This example of a classification scheme includes three categories.



4. DEFINE RECOVERY OBJECTIVES

Different classes of assets and data will have different recovery objectives. For instance, a critical e-commerce database may have very aggressive recovery objectives because the business simply can't afford to lose any transactions or be down for long. On the other hand, a legacy internal system may have less stringent recovery objectives because the data involved doesn't change very often and it's less critical to get back online.

Many IT pros fall short when it comes to this step. Setting recovery objectives without consulting the business line managers is the No. 1 cause for misalignment. It's imperative that you involve them in this process.

Here is a sample list of questions to ask your business colleagues when devising your DR plan:

- What applications and data does your department use?

- What is your tolerance for downtime for each?
- What is your tolerance for data loss for each?
- Are there times when these applications are not being used by employees, partners or customers?
- Would you ever need to restore data that is older than 90 days old? How about six months old? How about one year old?
- Are there any requirements (internal or external [i.e., industry or regulatory]) for the organization to retain the data for a designated period of time?
- Are there any requirements (internal or external [i.e., industry or regulatory]) that prevent us from moving the data from one geographical region to another?
- Are there any requirements (internal or external [i.e., industry or regulatory]) with regard to security and encryption?

The key here is to understand business needs and provide a differentiated level of service availability based on business priority. Once you have that information in hand, it needs to be translated into recovery objectives to be included in your DR plan. See Page 12, Table 4 for an example asset class breakdown.

Recovery time objective (RTO) — What is the acceptable time any of your data and production systems can be unavailable? This is your recovery time objective. To calculate the RTO for an application, consider how much revenue your organization would lose if the application went down for a given length of time.

For example, how much would you lose if your customer portal went down for an hour or a day? How much cost would be incurred if none of your employees can work because email is down?

Calculating your RTO is necessary to determine the features you need in your backup systems and products.

Calculating your RTO is necessary to determine the features you need in your backup systems and products. For example, if you have a very high RTO (say, more than four hours), you will probably have time to back up from tape, but if you have a very low RTO (such as just a few minutes), you need to use host-based replication or disk-based backup with continuous data protection features.

Recovery point objective (RPO) — What is the acceptable amount of data your organization can afford to lose? That is your recovery point objective. If your organization has a high tolerance for data loss, your RPO can be high, from hours to days. If your business can't afford to lose any data, or very little, your RPO will be seconds.

The RPO you set will determine the minimum frequency for backing up your data. If you can only afford to lose an hour's worth of data, you should back up the data at least every hour. That way, if an outage begins, for example, at 2:30 p.m., you can retrieve the 2 p.m. backup and meet the RPO requirement.

5. DETERMINE THE RIGHT TOOLS AND TECHNIQUES

Once you have identified all your IT assets, mapped their dependencies, grouped them together based on how critical they are to your business and set their recovery objectives, it's now time to choose what tools and techniques to use.

The good news is that numerous disaster recovery solutions are on the market today. Just make sure that what you choose offers the appropriate level of protection. Over-protection can cost the company needless money and introduce unnecessary complexity. (Complexity is the enemy of productivity and will likely increase the possibility for human error.) Under-protection is obviously bad because it puts important business functions at risk.

For instance, nightly backups using traditional (file-based) methods are more than sufficient for low-impact data, but this method would be inappropriate for high-impact data and applications. A continuous data protection (CDP) solution is great for high-impact data and systems, but it can add overhead to production servers and storage costs.



Perhaps the most critical component of your DR plan is offsite protection — use it regardless of the type of backup method you choose. Offsite protection (be it a tape vaulting service or replication to the cloud) should be commensurate to your recovery objectives. Make sure your data is sent to a location that is far enough away that it is not in the same geographic risk zone. Typically, this is at least 25 miles away from the primary location.

Finally, automate and streamline the recovery process as much as you can. In the event of a disaster, key IT staff may be unavailable. Automation also lessens the risk of human error.

Later in this e-book, we'll dive deeper into the specific technologies on the market today.

Tip: David Shulman of Salomon Brothers once applied the Goldilocks principle to economics when he wrote a strategy piece titled “The Goldilocks Economy: Keeping the Bears at Bay.” In this report, he was referring to an economy that was hot enough for profit growth but cool enough to keep the fed from hiking interest rates. This principle can also easily be applied to disaster recovery planning. The method you use should be just right for the classification of data you are protecting. An obscure reference to this report is likely to elicit nods of approval in the board room.

Another high-priority best practice that's often overlooked during a disaster recovery is to establish a dial-tone email system that enables all users to send and receive new emails during a power outage. The term “dial-tone” refers to the fact that even during power outages, phones often continue to work. Your IT users really should be able to expect a similar level of dial-tone service for mission-critical communications such as email. It may take a while to restore all email history for all users, but that is largely irrelevant to the pressing need to communicate in real time

following a major outage or disaster. A dial-tone email service can also help relieve some of the pressure on the IT staff to get everything up and running as soon as possible.

6. GET STAKEHOLDER BUY-IN

Go beyond the walls of the data center and involve key stakeholders for all your business units (i.e., application owners and business managers). They need to be involved in the planning phase. And they should agree with you on the company's priorities as well as the service-level agreements (SLAs) your team will provide.

Also, consult your strategic partners and vendors to make sure you're getting the most out of your DR solution or services. A few years back at the Orleans Parish in New Orleans, two servers failed, causing the loss of critical conveyance and mortgage records dating back to the 1980s. IT staff hadn't been keeping in close contact with the parish's cloud backup and disaster-recovery-as-a-service (DRaaS) provider. Be sure not to make that mistake, and stay in close contact with any vendor you employ.

Once you have consulted all of the key stakeholders, enlist an executive-level sponsor who will get behind you and the project. The importance of collaboration, consensus and executive support to your DR plan's success cannot be emphasized enough.

7. DOCUMENT AND COMMUNICATE YOUR PLAN

In a disaster scenario, you need a documented strategy for how to get back to a working state. This document should be written for the people who will use it.

Communicate your plan. All too often, only one person in the organization really knows the whole picture, leaving the organization vulnerable if that one person is unavailable during a disaster. In addition, be sure to

store your recovery strategy where it can be accessed during a disaster — not on a public share in your Exchange folders. Ideally, it should be printed and posted in multiple locations.

Go beyond the walls of the data center and involve all key stakeholders.

8. TEST AND PRACTICE YOUR DR PLAN

People often say, "Practice makes perfect." A better saying might be, "Practice makes progress." No organization ever gets to perfection with its DR plan, but practice will help you find and rectify problems in your plan, as well as enable you to execute it faster and more accurately. Make sure that everyone who has a role to play attends the practice sessions, even if you hold them, for example, on Saturdays.

You do not need to practice executing the full disaster recovery plan every time. It's perfectly acceptable to carve out pieces of your plan to test. See example on Page 12, Table 5.

9. EVALUATE AND UPDATE YOUR PLAN

A DR plan should be a living document. It's especially important to regularly review your plan given the shifting sands of an ever-changing business environment. Tolerance for downtime and data loss may decline. Key personnel may go on leave or terminate their employment. IT might migrate to new hardware or operating systems. The company might acquire another company. Your plan needs to reflect the current state of the organization.

Classification	Application	Server/VM	RTO	RPO
Low impact	Filesystem	FLS-001	24 hrs	24 hrs
Moderate impact	SharePoint, Active Directory	MOSS-001, AD-001, SQL-001, SQL-002	12 hrs	12 hrs
High impact	Exchange, Oracle	Exch-001, Orcl-001	1 hr	10 min

Table 4: Document the RTO and RPO for each asset class.

Class	Description	Frequency
Walk-thru exercise	Review the contents of your DR plan.	As often as necessary to familiarize response teams and individuals with a documented plan or changes to a plan
Tabletop exercise	Using a scenario, discuss the response and recovery activities of a documented plan.	At least 4 times per year or any time a change is made to the business or IT operating environment
Component exercise	Physically exercise a component of a DR plan (e.g., testing automated communications services or work-from-home capabilities together with IT or partner capabilities).	At least twice per year or when a change is made to the business or IT operating environment
Full-scale simulation	Using a scenario, carry out the response and recovery activities of a DR plan throughout the entire organization.	At least once or twice per year or when a change is made to the business or IT operating environment

Table 5: Here is an example of different test types with their frequency.

DR tools and techniques

Choosing the right approach (or approaches, as may be the case) to protect business-critical data and applications can be daunting. Savvy IT organizations choose flexible and powerful strategies and technologies that not only protect the IT environment and accelerate system, application and data recovery, but that also help them proactively avoid unplanned outages and data loss in the first place.

IT optimization ensures that your systems are right-sized for your current and future workloads. Set yourself up to be proactive rather than reactive when it comes to disaster recovery planning.

The best way to choose a solution to successfully recover after a disaster is, again, to think like a business leader. Choose a solution that offers the best value — one that doesn't cost the company needless money, introduce unnecessary complexity or put important business functions at risk.



FILE-LEVEL BACKUP

Some vendors may try to convince you that traditional backup (aka file-level) methods are outdated. However, file-level backups often form the foundation of a comprehensive data protection and disaster recovery plan. With these solutions, you select the files, folders and databases you'd like to back up. The backup solution scans the file system and makes a copy of each dataset to a different destination. This backup is performed on a regular schedule, typically once a day.

If traditional backup meets your recovery objectives, go with it. Don't succumb to marketing hype. That approach is absolutely the right business decision. Just be sure to choose a solution that offers a wide range of support for various applications and that doesn't complicate — and therefore jeopardize — the recovery process.

IMAGE-LEVEL BACKUP

Image-level backup (aka infinite incremental backup) offers the ability to take multiple incremental image backups in a short time span, such as every five minutes, to achieve great RPOs and get near-continuous data protection. With this data recovery method, you can quickly get a complete system running again after a disaster — even if the environment has no functioning operating system.

A bare metal restore (BMR) is a type of image-level backup that not only backs up the data but also the operating system and the application and configuration settings. With a BMR solution, you can quickly rebuild a server, including its operating system, network and system settings, application binaries, disk partitions, and data.

CONTINUOUS DATA PROTECTION

CDP solutions have become more popular because they offer impressive backup and recovery speeds as well as very granular recovery points (from seconds to minutes). A CDP solution works at the sub-file level, watching all of the new or changed blocks of data and only capturing those blocks.

Some modern CDP solutions periodically create restore points using a snapshot and volume filter device driver to track disk changes. This method enables the solution to perform a restore of a failed server or virtual machine (VM) in minutes.

These backups are also application-consistent, meaning the buffers are flushed and files are closed and so on, to protect the integrity of the data and to avoid issues during a restore. To ensure recoverability, one solution even validates its latest backup by completing an integrity check for any sort of corruption and mounting the backup copy of the database.

ARRAY-BASED SNAPSHOTS

Many companies are complementing their backups with array-based snapshots because they are low-impact, near-instantaneous and space-efficient. They also enable a business to quickly recover entire volumes of data at the granular level. Once a base snapshot is taken of any data written to a volume, only incremental changes are captured in subsequent snapshots. This not only saves disk space but also speeds local recovery.

Users can create many snapshots without setting aside extra disk space, and those snapshots can all be scheduled at very short intervals (depending on your RPO). In rare cases, these snapshots are also application-consistent — meaning that all open transactions have been committed, buffers have been flushed, files have been closed and the application is ready for the snapshot to occur.

Tip: If you are creating array-based snapshots, make sure to conduct a test failover from that snapshot so you can validate application-consistency.

Snapshots can cause complexity because they operate outside your normal backup operations. However, some backup applications on the market today complement array-based snapshots, allowing users to generate, schedule and recover snapshots through the same user interface. This allows you to simplify and centralize management.

Array-based snapshots are low-impact, near-instantaneous and space-efficient.

SINGLE-FILE RESTORE

Some backup applications also support single-file restores. They do this by recording checkpoints throughout the backup and saving the file history information on a per-file basis. To restore a single file, the backup application only has to read a small part of the backup data to find and restore the requested file(s).



A close-up photograph of a person's hands working at a desk. One hand holds a white pen over a laptop keyboard, while the other hand touches a tablet displaying a data visualization with a bar chart and a line graph. The background is softly blurred, showing a window with natural light.

Replicate

Regardless of the backup method you choose, it is essential that the backup data be replicated to a secondary location or the cloud. Some storage solutions will only replicate the incremental changes following initial site synchronization. This approach not only reduces hardware costs but also minimizes bandwidth requirements and accelerates recovery in the event of a disaster. Virtualization is a great option, allowing you to maintain standby VMs at your secondary location ready for deployment when disaster strikes.

Cloud-based data replication can be a great addition to your disaster recovery plan. Some organizations choose to maintain a hybrid environment in which some data is replicated on premises and some replicates to the cloud. Either way, the cloud option is cost-effective, efficient and highly scalable. Best of all? Recovery times are fast, which means near-zero downtime for critical apps.

In hybrid environments, some data is replicated on premises and some replicates to the cloud.

Virtual standby is a feature within many backup and recovery solutions that allows you to continually send updates to a virtual machine — either on premises or in the cloud — that can be activated in the event that you have an issue with your primary machine. This option allows you to maintain business continuity under nearly any circumstance.



Vendor considerations

Now that you've got a plan for how to build your business-centric DR plan, you may find you need to either completely replace or complement your existing solutions. You'll find dozens of products on the market from many different vendors, so it may be difficult to distinguish between them and ultimately make that decision on who you want as your partner.

Look for a vendor with solutions that are easy to acquire and deploy. Vendors have increasingly become more flexible when it comes to licensing. Choose a licensing model that will accommodate growth. Buying backup by component is a good idea if you aren't planning to add a lot of servers and applications. Alternatively, buying backup by capacity may not make sense when you have a few NAS filers with petabytes of data to protect.

Finally, be on the lookout for any "gotchas." We mentioned earlier that some vendors charge for the amount of data you recover!

Stand-alone deduplication solutions can help when you can't afford to completely rip and replace your existing backup solution. Some offer "accelerator" technology, which can greatly improve backup and restore performance as well as reduce the backup traffic over your network. They may also include replication capabilities, so you can safely and efficiently send your backup data to your DR site.

Regardless of the vendor you choose, ensure that the solution strikes the right balance between capabilities and manageability. Some solutions on the market today include a lot of bells and whistles but have an overwhelming amount of options for scheduling, tracking, streaming data and other options, and require admins to define every minute detail of the backup process. This can introduce the possibility for human error. The right solution will offer enterprise-grade capabilities and an intuitive user experience.

Support and maintenance is also a key consideration. Look for a vendor that has a proven track record of supporting the latest applications and operating system releases.

Look for a vendor with solutions that are easy to acquire and deploy.

Last, when selecting your solution, take the total cost of ownership into consideration. That includes not only license fees but also maintenance renewal fees and professional service fees required for upgrades and tuneups, as well as all of the hardware costs required to run the backup system. You'd be surprised at how quickly these add up.



Conclusion

Sometimes there is misalignment between what the business expects and what IT can actually deliver — especially when it comes to disaster recovery. As an IT professional, your ultimate goal is undoubtedly to keep the business running smoothly, even in worst-case situations. Following the guidelines we've outlined can help you organize your thoughts, ask the right questions and develop the right strategy to begin taking a more proactive, business-centric approach to disaster recovery planning. These steps can also help you become an IT hero — or possibly even get you a seat in the C-suite.

Your ultimate goal is undoubtedly to keep the business running smoothly.

ABOUT QUEST

At Quest, our purpose is to solve complex problems with simple solutions. We accomplish this with a philosophy focused on great products, great service and an overall goal of being simple to do business with. Our vision is to deliver technology that eliminates the need to choose between efficiency and effectiveness, which means you and your organization can spend less time on IT administration and more time on business innovation.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (www.quest.com) for regional and international office information.

© 2018 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at www.quest.com/legal

Trademarks

Quest and the Quest logo are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit www.quest.com/legal/trademark-information.aspx. All other trademarks are property of their respective owners.