Be Prepared with a Doomsday Disaster Recovery Plan

Written by Keri Farrell, product manager, Quest; Brian Hymer, senior solutions architect, Quest; and Mike Daniels, solutions architect, Quest

ABSTRACT

IT disasters are both unpredictable and inevitable. Recovery from disasters, however, should be planned, predictable and controlled. This tech brief explores the common causes of IT disasters, outlines best practices for establishing a doomsday disaster recovery plan, and offers two tools that can help you get your business running again quickly after an inadvertent mistake, malicious attack or natural disaster.

INTRODUCTION

What is a disaster?

In an everyday situation, a disaster is any calamitous event, especially one that occurs suddenly and causes great loss of life, damage or hardship. In IT, a disaster is an unexpected event that causes a substantial loss of service levels in critical business systems for an unacceptable period of time.

But the practical definition of the word "disaster" depends a great deal on your role at the time. For an IT admin, for instance, a disaster might be best characterized as any situation that causes you to stay at work late or work on a long holiday weekend.

What is a doomsday disaster?

What makes a disaster a doomsday disaster? Simply put, it's a disaster that makes it hard for the organization to continue to function and that may even threaten its continued existence. Since every company of any size today relies on technology, IT outages can quickly snowball into doomsday disasters.

The most obvious examples of doomsday disasters are complete system failures — users are not able to log in or use the applications they need to get their jobs done, and communications are down. But even a lesser failure, such as downed communications alone for an extended period, can qualify as a doomsday disaster if it means business comes to a halt and losses start to mount.

For an IT admin, all too often, a doomsday disaster can be a resume-updating scenario (RUS).

> Organizations must be prepared for disasters, no matter how experienced their IT staff might be.

The key to avoiding this consequence is having a complete disaster recovery plan. This tech brief explores the common causes of disasters, explains how to prepare for a doomsday disaster and offers tools to help.

CAUSES OF DISASTERS

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. Or, given that Active Directory (AD) is the backbone of any Windows environment, a disaster can easily begin with an AD issue. For example, a glitch in Active Directory was reported to be the cause of an IT outage that delayed treatment for over 700 patients at an NHS trust, according to Information Age magazine.

It's important to understand, however, that it's not just junior-level people who make mistakes in AD that cause disasters. In many cases, the problem is caused by senior-level IT professionals who know what they're doing. At one Quest customer site, for instance, a senior administrator made an invalid configuration change in AD that was nevertheless accepted and applied. Within eight hours, all of the domain controllers (DCs) received the change via replication, and every one fell into an endless reboot that required a full forest recovery. The OS vendor simply has not taken into account every possibility, and therefore organizations must be prepared for disasters, no matter how experienced their IT staff might be.

Malicious

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 U.S. industries from groups or countries opposed to U.S. actions or policies. For example, a December 2012 attack on an oil and gas company in Saudi Arabia affected tens of thousands of company computers, and U.S. financial institutions have been battling long-running denial-of-service attacks.

In fact, an August 2013 article by Bloomberg notes that the "probability that cyber-attacks will become a major weapon of countries whose militaries are outmatched by the U.S. has been growing for several years." In particular, it says, the "U.S. is planning for a possible wave of computer attacks against companies by hackers connected to Syria or Iran," with a particular focus on possible disruptions to power grids, financial systems or other critical infrastructure. The article also notes that 2013 attacks on both the New York Times website and the Associated Press's Twitter account by hackers known as the Syrian Electronic Army "showed higher sophistication, indicating a growing expertise." Clearly, organizations need to be prepared for disasters caused by malicious actions.

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 very 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. They can be caused by hardware failures or human error, as in the example earlier of the admin who made a change to AD that caused all the organization's DCs to fall into an endless reboot.

Data corruption is particularly challenging because no one may even realize what's happening at first, so a small problem can escalate into a disaster, and corruption can also be very difficult to troubleshoot. For example, a Quest customer was splitting their forest into two separate domains, and the FSMO role of an existing DC was hijacked with the introduction of a new DC into the forest with the same name. This stopped the domain from functioning and nearly led to a

forest failure. Accurately diagnosing and remediating the problem required level three support from Microsoft — someone who knew just where to dive within the configuration to get the computers reporting properly.

Site disasters

In its 2013 survey, the Disaster Recovery Preparedness Council reports that the most common cause of invoking a business continuity plan was natural disasters such as storms, hurricanes, tornadoes and earthquakes. Power outages and fire also make the full list in the survey:

- Extreme weather and natural disasters (storms, hurricanes, tornadoes, earthquakes) (55 percent)
- Power outages (49 percent)
- IT failures (server failure) (36 percent)
- Foods (28 percent)
- Fire (18 percent)
- Telecom failure (14 percent)

Organizations may think they have a good plan by having an off-site backup, but if the backup is physically close to the main site, both can be destroyed by the same natural disaster. For example, one company in New York had two data centers, but both were on the East Coast. So when Hurricane Sandy hit, the company lost both data centers, and it took them about three months to get back online.

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 by unintentionally damaging human actions. 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 the rack that it was on, the entire SAN went down — the entire controller was fried. 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. (This particular storage failure became an RUS for the employee who decided to put carpet in the data center.)

Situations like these 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.

PREPARING FOR A DOOMSDAY DISASTER

Although disasters are unpredictable, recovery shouldn't be. In fact, recovery should be planned, predictable and controlled. The following best practices will help you prepare for a doomsday disaster.

- **Understand your goal** — The goal of disaster recovery is not primarily to understand the cause of an outage, it is to get the company back up and running as quickly as possible.
- **Get all stakeholders involved** — Key stakeholders for all your business units need to be involved in the planning phase. They need to decide what their priorities and service-level agreements are, and sign off on the disaster recovery plan.
- **Document your recovery strategy for the people who will use it** — 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. Although some organizations say they want instructions written so that a layman could follow them, this is not feasible for a disaster recovery plan; disaster recovery situations really require IT professionals with the skills to troubleshoot issues and not merely follow a static set of instructions. Keep this in mind as you document your plan.
- **Disseminate information and make sure the recovery plan is accessible** — 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

> "The probability that cyber-attacks will become a major weapon of countries whose militaries are outmatched by the U.S. has been growing for several years."

Although disasters are unpredictable, recovery shouldn't be.

strategy where it can be accessed during a disaster, not on a public share in your Exchange folders. Ideally, it should be in multiple locations.

- **Plan for unavailable team members** — Remember that when a disaster strikes, your entire staff might not be available. In fact, key staff members could be on vacation, away for training or home sick. In the case of a natural disaster, you might not be able to locate certain employees. Make sure your disaster recovery plan does not rely on all team members being available.
- **Automate where you can** — Automate the recovery process wherever you can to minimize the chance of human error and speed recovery by eliminating manual steps.
- **Test and practice your plan** — People often say, "Practice makes perfect." A better saying might be, "Practice makes progress." No organization ever gets to perfection with its disaster recovery 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. "Once you've put the systems in place, you need to make sure they're up and running on a regular basis, or you risk the failure of the [disaster recovery] plan during an actual disaster," explains one TechRepublic article. Make sure that everyone who has a role to play attends the practice sessions, even if you hold them, for example, on Sundays.
- **Ensure a clear go/no go decision** — Ensure there is an executive-level sponsor who is authorized to decide whether to begin the disaster recovery process.
- **Update the plan periodically** — It's important to regularly review your plan. Key personnel may go on leave or terminate their employment, IT might migrate to new hardware or operating systems, or the company might acquire another company. Your plan needs to reflect the current state of the organization.
- **Establish an internal communication plan** — Effective communication is critical to recovering from a disaster. Not only does IT staff need to communicate in order to get systems back up, they need to communicate status updates to management and users in order to prevent panic and avoid a constant barrage of questions that will slow recovery efforts. Be sure to plan to use an external communication method — remember that email and IP phone systems rely on the network and internal systems, and therefore can be compromised by a disaster. Cell towers may be out or overloaded.
- **Communicate with your software vendors** — If you're not on the phone with the vendors of your mission-critical applications during a disaster recovery, you're making a big mistake. By working in conjunction with application vendors, whether they join you on site or by phone, you can help ensure a faster and more effective recovery. Plus, you'll be helping vendors learn how they can improve their reliability of their applications for the next disaster.

KEY COMPONENTS OF A DISASTER RECOVERY PLAN

With those best practices in mind, let's turn to several key components of a disaster recovery plan. As noted earlier, the goal of disaster recovery is to get the business up and running as quickly as possible. But a good plan will quantify that with two key figures:

- **Recovery time objective (RTO)** — The amount of time between the beginning of a critical data loss incident to the time the data and all systems are restored.
- **Recovery point objective (RPO)** — The point in time when data must be recovered after an outage.

These values will differ for different kinds of data, so a third key component of a disaster recovery plan is a classification of the data by recovery priority. Let's look more closely at each of the components.

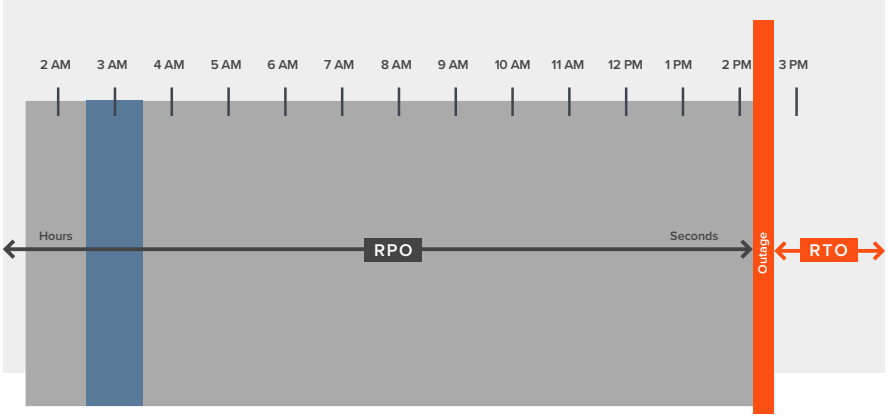

Figure 1: Your RTO is the amount of time between the beginning of a critical data loss incident to the time the data and all systems are restored.

Determining your RTO

How long can any of your systems be down and your business still survive? This is your recovery time objective (see Figure 1).

To calculate the RTO for an application, consider how much money 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 for determining 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 a disk-based backup with continuous data protection features.

Determining your RPO

How much data can your organization afford to lose? That is your recovery point objective (see Figure 2).

How long can any of your systems be down and your business still survive? This is your recovery time objective.

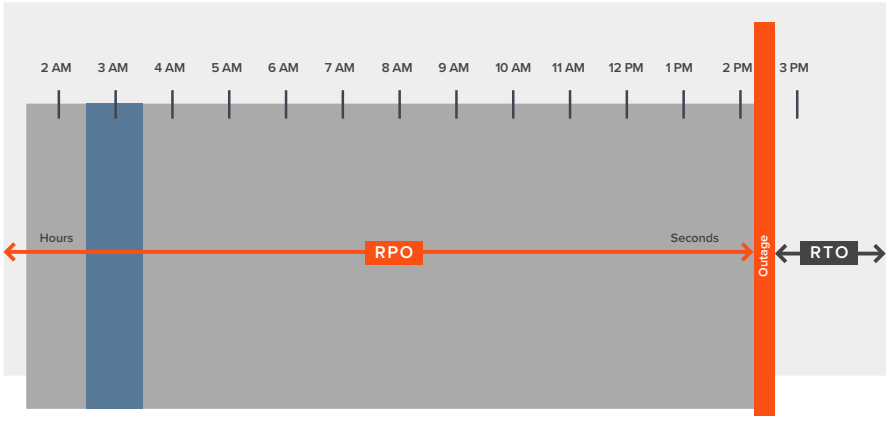

Figure 2: Your RPO is the point in time to which data must be recovered to after an outage.

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, your RPO will be low, perhaps even seconds.

The RPO you set will determine the minimum frequency for taking a backup of your data. If you can lose only an hour of data, you should take a backup every hour. That way, if an outage begins, for example, at 2:30 p.m., you can retrieve the 2:00 p.m. backup and meet the RPO requirement.

Classifying your data and applications

Different types of data and applications will have different RTOs and RPOs. For instance, a critical customer application may have a very low RPO and RTO because you can't afford to lose any transactions or be down for long, but a legacy internal system may have a more generous RPO (since the data doesn't change very often) and RTO (since it's less critical to get back online).

Different types of data and applications will have different RTOs and RPOs.

As you build your disaster recovery plan, you'll need to classify your data and applications so you can assign appropriate RTOs and RPOs, and choose appropriate data protection approaches. One useful classification strategy is illustrated in Figure 3.

This classification scheme includes three categories:

- Static — Static data doesn't change very often, such as files or Word documents. Static data is often subject to strict regulations, but often does not impact the bottom line of the business. Therefore, RTOs and RPOs can be less stringent, which allows for data protection approaches such as backing up to tape.
- Business vital — Business-vital data and applications are not required for the business to run but are more important than most static data, such as your customer relationship management or enterprise resource planning system.
- Mission critical — Mission-critical applications are those that must be restored as quickly as possible, such as email and Active Directory.

A COMPLETE DISASTER RECOVERY STRATEGY

Business-vital and mission-critical data and applications require more advanced data protection approaches than simple tape backups. Specifically, a complete disaster recovery strategy requires both a bare-metal recovery (BMR) solution and an Active Directory forest recovery solution. The rest of this tech brief explores those two approaches in detail.

What is bare-metal recovery?

BMR is a data recovery technique that enables you to quickly get a complete system running again after a disaster, even if the environment has no functioning OS. A BMR solution backs up not only the data, but also the OS, the application and configuration settings. Therefore, you can quickly rebuild a server, including its OS,

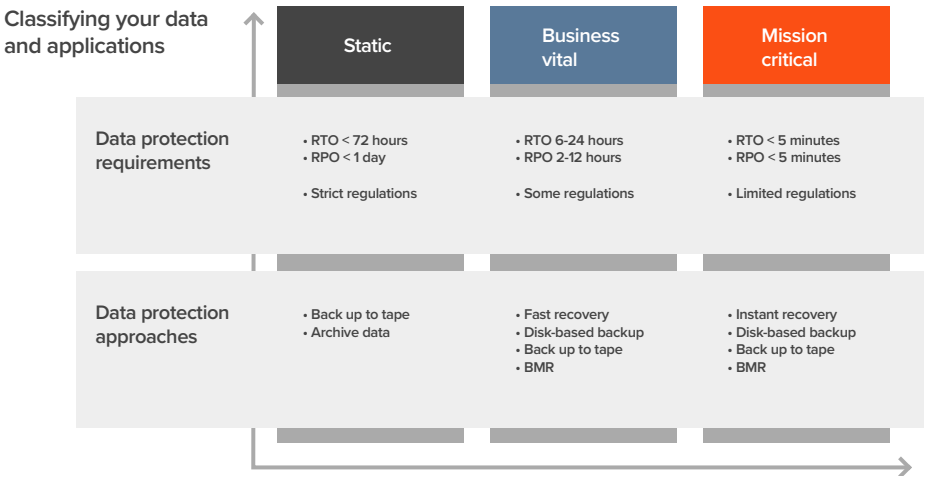

Figure 3: Classifying data and applications

network and system settings, application binaries, disk partitions and data.

A BMR solution is ideal for mission-critical applications with RTOs measured in minutes or hours, since automation eliminates much of the manual intervention and guesswork that slows other recovery techniques. Be sure to look for a solution that offers recovery to your choice of similar hardware, dissimilar hardware or virtual machines.

NetVault Bare Metal Recovery

NetVault Bare Metal Recovery is an integrated solution capable of recovering an entire system, including the OS, applications, system settings, partition information and data for any supported client. NetVault Bare Metal Recovery provides BMR with either offline/cold backups or online/hot backups.

In the event of a system failure, the administrator can boot the system using the minimal OS or LiveCD to initiate the recovery process. You can recover a Windows or Linux system to similar or dissimilar hardware (physical to physical), or even to a virtual machine (physical to virtual).

To help you meet your RTOs, NetVault Bare Metal Recovery automates many of the manual BMR steps, as illustrated in Figure 4.

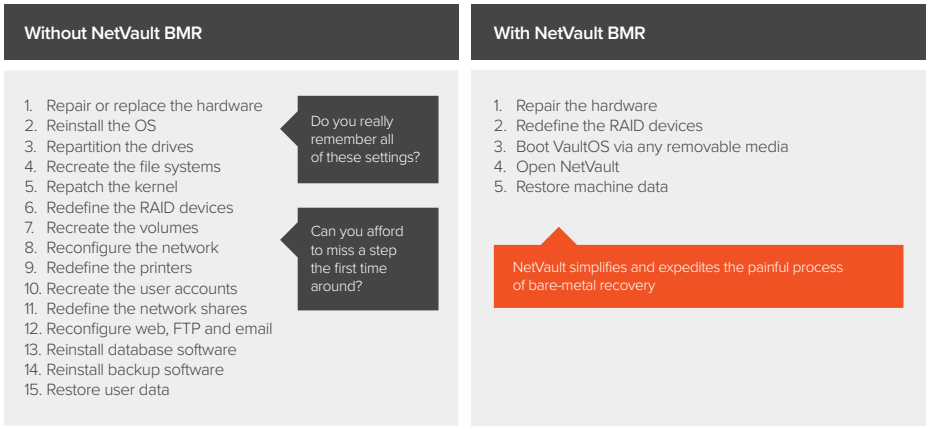

Figure 4: NetVault Bare Metal Recovery speeds the BMR process.

BMR is a data recovery technique that enables you to quickly get a complete system running again after a disaster.

After a bare-metal recovery, you need to perform an AD forest recovery.

With NetVault Bare Metal Recovery, you can:

- Quickly rebuild the server, including its OS, settings, application, disk partitions and data
- Ensure your servers are back online as soon as possible
- Establish online full-partition backups, offline block-level backups and Windows VSS-based backups
- Recover servers from bare metal to similar hardware, dissimilar hardware or a virtual machine

Active Directory recovery

After the BMR is complete and the OS is back in a functional state from before the disaster, you can now focus on getting Active Directory back in sync and up to date.

Let's go through the process. Each DC uses an update serial number (USN) to know where it is in its process. When you perform a bare-metal restore, you use a backup to restore the main DC from an earlier time. In the time since the backup was taken, objects may have been created on that DC and replicated to other DCs; after the BMR, those objects no longer exist on the DC they were created on, and the DC's USN is set back to what it was at the time of the backup.

Therefore, after the BMR, if the DC creates new objects and attempts to replicate them out, the other DCs will reject those updates based on the main DC's reset USN. Since the USN is lower than the most recent update already processed, the DCs will assume there is no new information to process.

Having DCs out of sync in this way can cause a variety of problems. The first is lingering objects. These are objects that exist on one DC but, because of the USN problem, don't replicate to other DCs, which can cause attributes to fall out of sync. This will eventually cause passwords to fall out of sync and, ultimately, people won't be able to log on. Microsoft addressed this issue in 2003, but the fix was detecting a USN rollback, which causes the DC to quarantine itself and put out an error message.

Manual AD recovery

It is an understatement to say that resolving a USN rollback problem can be a difficult undertaking. Microsoft provides a recovery guide that includes more than 40 manual steps — and the complexity can multiply fivefold for a multi-domain forest. The details of a manual AD recovery are beyond the scope of this document, but here are just some of the complexities:

- You must manually select your backups.
- An administrator must physically access each recovered DC in each location, since you have to do things you cannot do remotely.
- You have to manually quarantine your DCs so that a DC that's being recovered can't talk to a DC that may still be corrupt or in an invalid state.
- The process requires expertise in command-line tools, increasing the probability of human error.
- You can reinstall AD on only one DC at a time.

Recovery Manager for Active Directory Forest Edition

There is a better way to bring Active Directory into sync after a BMR. Recovery Manager for Active Directory Forest Edition simplifies and automates all of the tasks necessary to perform the recovery. Specifically, Recovery Manager:

- Automates backup selection because it makes system state backups of all your DCs
- Remotely and simultaneously restores all DCs in the forest, eliminating the need for administrators to be physically present at each DC
- Automatically quarantines all affected copies of AD so the corruption can't propagate to the new environment
- Eliminates the need to remember seldom-used command-line tools
- Provides a single console so IT staff can decide which DCs to restore and track the progress of the restoration
- Creates and maintains necessary recovery process documentation

Moreover, Recovery Manager uses multiple agents, so many steps can

happen simultaneously, drastically reducing the time for recovery. In fact, the solution has more agents working on the recovery than you would have staff to do the work manually.

Testing your disaster recovery plan

As noted earlier, it's crucial to test your disaster recovery plan at least once a year. However, adequate testing requires a test lab that closely matches your production environment, and most organizations simply don't have the resources to create one. Fortunately, there is no longer any need to settle for suboptimal testing in a lab with only a few machines.

Recovery Manager for Active Directory Forest Edition includes the Active Directory Virtual Lab, which creates a virtual lab from your production data, enabling you to have a replica that meets your needs. The virtual lab can include all of your DCs and even member servers, or you can choose to exclude certain domains or certain servers. The lab creation options include physical to virtual or virtual to virtual. Whichever you choose, the replica has your production data, so when you test your disaster recovery plan, your results will more accurately match what you may run into when you need to perform a real recovery.

CONCLUSION

Organizations cannot know exactly what disasters they will face, but they should expect disasters to happen. Inadvertent mistakes, malicious hackers, cyberterrorism and natural disasters are all facts of life for businesses today. By following the best practices detailed in this tech brief, you can build — and test — a sound disaster recovery strategy.

Together, NetVault Bare Metal Recovery and Recovery Manager for Active Directory Forest Edition provide a complete recovery solution. With these tools, you can recover quickly and restore accurate and up-to-date AD data, ensuring that a disaster does not turn into a resume-updating scenario.

FOR MORE INFORMATION

You can read more about these Quest solutions online, and even download a free trial:

- For NetVault Bare Metal Recovery, please visit quest.com/products/netvault-backup/bare-metal-recovery.aspx.
- For Recovery Manager for Active Directory Forest Edition, please visit quest.com/products/recovery-manager-for-active-directory-forest-edition.

Quest offers a webcast that presents actual forest disasters and explains how Recovery Manager for Active Directory Forest Edition helped those organizations recover; please see "Active Directory Forest Recovery: What You Don't Know WILL Hurt You."

If you want to learn more about advanced persistent threats, read "Advanced Persistent Threads: The New Reality."

Together, NetVault Bare Metal Recovery and Recovery Manager for Active Directory Forest Edition provide a complete recovery solution.

ABOUT QUEST

Quest helps our customers reduce tedious administration tasks so they can focus on the innovation necessary for their businesses to grow. Quest® solutions are scalable, affordable and simple-to-use, and they deliver unmatched efficiency and productivity. Combined with Quest's invitation to the global community to be a part of its innovation, as well as our firm commitment to ensuring customer satisfaction, Quest will continue to accelerate the delivery of the most comprehensive solutions for Azure cloud management, SaaS, security, workforce mobility and data-driven insight.

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 Web site (www.quest.com) for regional and international office information.